Essays on Monetary Policy and Macroeconomics

Doctoral Thesis
submitted to Justus Liebig University Giessen,
Department of Economics and Business Studies

1*st* Supervisor: Prof. Dr. Peter Tillmann
2*nd* Supervisor: Prof. Dr. Peter Winker

David Finck

November 24, 2022

Preface

This dissertation consists of nine independent papers that I wrote between 2017 and 2022. To be honest, I can't remember exactly how I (or my coauthors) came up with the idea of writing the individual papers at the time. I do know, however, that I never wanted to commit to a particular topic when I started my dissertation. Fortunately, I got excited about several topics, which resulted in this colorful bouquet of topics. The papers contributed to this dissertation were written at a time when the impact of the Great Recession on the dynamics of macroeconomic variables was still in the forefront. Nevertheless, since 2020, people as well as institutions in almost all countries around the world have been struggling with new challenges that are still being felt since the outbreak of the pandemic. While the situation in most countries has improved significantly and most economies are recovering from the slump in economic activity in 2020, the consequences of the pandemic, for example in the form of supply shortages as a result of pandemic measures, are still clearly felt today. One bright spot is that research has also made tremendous progress in this area, and publicly available innovative data sets are helping to conduct structural analyses that help us better understand, for example, the far-reaching effects of the pandemic.

Many of the following papers are only loosely related. For instance, two papers are purely model-theoretic, some papers combine both empirical models and insights from structural macro models, and some papers use only empirical models. On the following few pages, I will briefly attempt to discuss each paper in order to prepare and motivate the reader for the individual papers in the main body of this dissertation.

The first paper, *Has Monetary Policy Really Become Less Effective in the Euro Area? A Note*, tries to answer the question whether the overall effectiveness of monetary policy in the euro area became less effective over time. In order to do so, I apply a time-varying VAR (TVP-VAR) model with both drifting coefficients and a drifting variance-covariance matrix. This model is well suited to uncover sudden structural breaks and gradual changes alike in the transmission mechanism of monetary policy shocks. The results show that, although there are weak tendencies towards less effective monetary policy, the transmission of monetary policy shocks did not change significantly over time. However, analyzing unanticipated policy changes begs for answers of how systematic monetary policy changes potentially affect the economy. To answer this question, I plant the most recent average policy rule, i.e. the coefficients in the interest rate equation, into 2005 and re-simulate both the inflation rate and the unemployment rate. The results indicate that the loose monetary policy rule pursued by the ECB would not have resulted in a less severe recession.

The second paper, submitted as *Mortgage Debt and Time-Varying Monetary Policy Transmission*, studies the role of monetary policy for the dynamics of US mortgage debt. Household debt in general and mortgage debt in particular are important for monetary policy, as the build-up of household debt in the US (and also other countries) is often interpreted as a potential risk to financial stability. An unexpected monetary tightening should curb the build-up of mortgage debt. Hence, monetary policy shocks can have important effects on the overall dynamics of mortgage debt. In this paper, we use a time-varying VAR model and study the sensitivity of US mortgage debt to monetary policy shocks. We find that an identically sized policy shock became less effective over time. Using a dynamic stochastic general equilibrium model (DSGE), we find that a fall in the share of adjustable rate mortgages (ARMs) replicates this finding. We invoke the concept of the sacrifice ratio to shed light on when a debt reduction as a result of a monetary tightening is relatively expensive or cheap, respectively, and find that toward the end of the sample, it is particularly costly in terms of employment to use a monetary tightening in order to initiate a reduction in household debt. These results have an important policy implication, as a weaker transmission of policy impulses to the mortgage market, but not to the real economy, implies that monetary policy is not the right instrument to reduce household sector debt. A policy tightening at the end of the sample period to curb the build-up of mortgage debt is expensive in terms of forgone employment.

The third paper, *Do Credit Supply Shocks Have Asymmetric Effects?*, also looks at the role of the credit market, but from a different angle. We examine whether credit supply shocks in the US have asymmetric effects on macroeconomic variables as well as on credit variables. A large body of literature points at amplifications and asymmetries in the propagation of sudden (financial) distortions as, ranging from (i) occasionally binding borrowing constraints, (ii) market imperfections (in terms of asymmetric information) to (iii) the role of asymmetric central bank behavior. In order to uncover asymmetries in the transmission of credit supply shocks, we first identify a credit supply shock within a linear VAR model and separate this shock into positive and negative. We then rely on local projections which allow us to directly evaluate the effect of positive (negative) credit supply shocks for adjacent horizons. Although our framework is not suitable to trace the exact transmission mechanisms, the contribution of our paper is rather to see whether credit supply shocks do have asymmetric effects in general. And they do. Our results clearly point to asymmetric effects concerning overall private debt, mortgage debt and house prices. For instance, we find that positive credit supply shocks tend to have stronger and more prolonged effects. Also, we find key macro headline variables (prices, production, short-term interest rate) as well as some key variables determining the demand and supply side of the economy to respond asymmetrically. Overall, following an expansionary shock of credit supply, our results clearly point to the well established boom-bust cycle. In the case of a tightening shock, we do not observe such a pattern. Importantly, we get relatively mild responses for debt and prices in a linear model, such that the true effects tend to be underestimated when not taking possible asymmetries in transmission mechanism into account.

The fourth paper, *Optimal Monetary Policy Under Heterogeneous Beliefs*, and the fifth paper, *Forward Guidance Under the Cost Channel*, are both model-theoretic and focus in particular on the way households form expectations, albeit in different ways. In the first of the two papers, I relax the gold standard of rational expectations formation and examine how optimal monetary policy changes when a fixed share of households within an economy does not form rational expectations, but instead naively extrapolates the most recent observations into the future. The twist in this paper is that the central bank, which conducts optimal monetary policy in a discretionary fashion, knows that a fixed share of households forms expectations naively, but it does neither know this share, nor does it know how strongly the non-rational households extrapolate current observations into the future. In order to incorporate this kind of model uncertainty, I derive a robust optimal monetary policy plan in a min-max framework where the central bank plays a zero-sum game versus a fictitious, malevolent "evil agent". Calibrating the model to the US economy, I find that the central bank can improve welfare if it accounts for uncertainty while the model is distorted. Interestingly, implementing the robust policy plan favorably affects the inflation-output stabilization tradeoff which results in a reduction of welfare losses even if the model is not being distorted.

As a result of the economic collapse in the wake of the pandemic, many central banks, including the Federal Reserve, have lowered the federal funds rate back to zero percent after previously raising it steadily. When the zero lower bound (ZLB) on interest rates is binding, central banks rely on unconventional measures such as Quantitative Easing and Forward Guidance to fight deflationary pressures. The fifth essay of this dissertation, *Forward Guidance Under the Cost Channel*, analyzes how the cost channel of monetary policy affects the credibility of forward guidance. A cost channel is present when firms' marginal costs depend on the nominal rate of interest. In a standard New Keynesian model, this assumption introduces a direct supply-side effect of monetary policy. In a discretionary framework, I allow for recurring episodes of binding constraints on interest rates to account for recent observations. Calibrating the model to the US economy, the results show that, compared to a standard New Keynesian model, (i) the cost channel makes forward guidance more powerful by reducing welfare losses at the zero lower bound on interest rates, (ii) the central bank's best strategy requires a shorter forward guidance horizon, and (iii) that these effects depend on the strength of the cost channel. I show that these findings are a result of the breakdown of the *divine coincidence*, i.e. the central bank's ability to perfectly stabilize demand shocks by adjusting the short-term nominal interest rate appropriately. Although the cost channel is empirically well documented in the literature, it is typically ignored when investigating optimal monetary policy. In a poor man's approach, I therefore investigate how my findings change if a cost channel is present but mistakenly ignored by the central bank. I show that ignoring the cost channel is costly in terms of steady-state consumption.

In early 2020, the global spread of the COVID-19 pandemic led to a sharp contraction of economic activity in almost all economies around the world. The sixth essay, *Pandemic Shocks and Household Spending*, investigates how pandemic shocks, therein referred to as the surprise number of fatalities, affects household spending in the US. Specifically, we derive the pandemic shock as the difference between daily real-time and unrevised projections on deaths based on a (machine learning) data-driven approach and the actual reported number of deaths. We allow for state-dependent effects based on the growth rate of the daily number of new infections as our state variable. Our results show that a pandemic shock has a negative effect on spending: a surprise increase in the number of deaths leads to a sharp reduction in expenditures. We also show that this effect is depending on the position of the US economy with respect to the infection curve. With the number of new infections increasing, the effect of a shock is much stronger. If the growth rate of the number of infections is small, in contrast, the pandemic shock has almost no effect.

The ongoing global integration of financial and goods markets changes are often assumed to affect the inflation process in general, and the relationship between inflation and economic activity in particular. For instance, the sharp drop in GDP during the Great Recession did not lead to a further drop in inflation. This phenomenon is often referred to as the "missing disinflation" in the literature, and such flattening of the Phillips curve has important consequences for monetary policy: for given inflation expectations, the argument goes, a flatter Phillips curve would require a deeper recession in order to bring high inflation back to the target. However, quantifying the extent to which global forces explain inflation is not straightforward. The seventh paper of this dissertation, *The Role of Global and Domestic Shocks for Inflation Dynamics: Evidence from Asia*, studies inflation dynamics and the output-inflation tradeoff in small open economies. We estimate a series of VAR models for a set of six Asian emerging market economies, in which we identify a battery of domestic and global shocks using sign restrictions. Focusing on the co-movement between domestic and global variables, our identification strategy also allows us to distinguish between global and domestic shocks. We find that global shocks explain large parts of inflation and output dynamics. In order to investigate how global shocks affect the Phillips curve, we first derive counterfactuals in which global shocks are simulated away. We then estimate reduced-form Phillips curve regressions based on alternative counterfactual scenarios (i.e. alternative decompositions of output into global and domestic components). Our results show a positive and significant correlation between inflation and the fraction of GDP driven by domestic shocks. While including the output component driven by oil prices seems to *flatten* the Phillips curve, though the effect is not significant, the component driven by global demand shocks *steepens* the inflation-output nexus. These results highlight the difficulties facing inflation targeting central banks in the respective region: while monetary policy primarily affects domestic demand, global demand, which drives the bulk of inflation, is not under the control of monetary policy.

In the spring of this year 2022, images went viral on Twitter showing hundreds of container ships waiting to be unloaded outside the port in Shanghai. This backlog of container ships was a result of the ultra-restrictive COVID-19 measures taken by Chinese authorities to contain the spread of the pandemic. There has been a lot of discussion in the media about the impact of this backlog on inflation, especially against the backdrop of the lack of inputs in other parts of the world for the production process due to the containers not being loaded. The eighth paper of this dissertation, *The Macroeconomic Effects of Global Supply Chain Disruptions*, picks up such scenarios and investigates how unexpected disruptions in global supply chains affect economic activity in the euro area. Specifically, we estimate a structural VAR model using a novel identification scheme based on a combination of traditional sign restrictions and narrative sign restrictions to identify a global supply chain shock. The idea of narrative restrictions is to constrain the admissible set of structural parameters by ensuring that both model-consistent structural shocks and historical decompositions align with established narratives around key historical events, in our case namely the following three: the Tōhoku earthquake in 2011, the Suez Canal obstruction in 2021, and the Shanghai backlog in 2022 mentioned above. We show that only a handful of narrative sign restrictions dramatically sharpens and even changes the inference of SVARs originally identified via traditional sign restrictions. Among other things, our results show that a global supply chain shock causes a drop in euro area real economic activity and a strong increase in consumer prices. Importantly, we show that our shock is uncorrelated with a whole range of other supply-side shocks from the literature, i.e., that our shock, and hence its consequences, have been previously ignored in the literature.

The last paper, *On the Empirical Relevance of the Exchange Rate as a Shock Absorber at the Zero Lower Bound*, also examines the consequences of the ZLB on interest rates, but from a different standpoint and with a different set of questions in mind. The motivation for this paper is based on the common finding in the literature that within an open economy New Keynesian model with flexible exchange rates, the real exchange rate appreciates in response to an asymmetric negative demand shock in a ZLB scenario and exacerbates the adverse macroeconomic effects. This is because the real interest rate differential rises from the perspective of the country located at the zero lower bound on interest rates, as the real interest rate is driven only by inflation expectations, which in turn rise very sharply at the ZLB. Consequently, flexible exchange rate regimes cannot stabilize cyclical developments in response to adverse demand shocks. We show that, when monetary policy is able to accommodate the adverse effects of the negative demand shock via unconventional measures, the model can generate a real depreciation when interest rates are at the ZLB. We further examine these counteracting exchange rate channels empirically. We estimate the effect of a negative asymmetric demand shock on the real exchange rate and inflation expectations as well as output and prices by employing state-dependent and sign-restricted local projection methods for the euro area vis-à-vis the US, Canada, and Japan. We find that the real exchange rate depreciates when interest rates are not at the ZLB but also when they are. Furthermore, our empirical results show that the real exchange rate can absorb considerable variations in output, confirming its shock-absorbing capacity before but also during the ZLB episode.

I want to close the preface of this dissertation with some acknowledgements. I am well aware that during my time as a doctoral student, I have always been very lucky with the people around me. Nevertheless, I would like to single out a few people. First and foremost, I would like to thank my first supervisor, Prof. Dr. Peter Tillmann. Not only for his tireless commitment and help whenever it was needed, but also for the very pleasant working atmosphere. His constant undying support and insightful feedback had a huge positive impact on each of the nine papers. Moreover, I have always learned a lot in our numerous joint research projects. I am already very much looking forward to continuing my research with him in the future. I would also like to thank my second supervisor, Prof. Dr. Peter Winker. His outstanding knowledge in the field of econometrics and statistics helped me to solve many problems I encountered during my time here. A very special thanks also goes to my former supervisors at the Deutsche Bundesbank, Dr. Mathias Hoffmann and Dr. Patrick Hürtgen, for the great time I had during my internship in 2021. During this time, I learned an incredible amount from both of them and I am very happy that this resulted in a very exciting research project and that further projects will follow in the future. Furthermore, I thank my former colleagues, Dr. Annette Meinusch, Dr. Immaculate Machasio, Dr. Lucas Hafemann, as well as my former office buddies Dr. Jörg Schmidt and Paul Rudel not only for their support and many insightful suggestions, but also for the wonderful time we spent together during the beginning of my journey. I would also like to thank my current colleagues Anisa Tiza Mimun, Sinem Kandemir and Moritz Grebe. Besides their professional capabilities, they shaped a very pleasant and fruitful working environment. I would also like to thank our former and current student assistances Omar Omari, Salah Hassanin, Niklas Benner, Justin Berndt, Svea Lembcke, Svenja Küchler, Marlene Alt, and our secretary Cornelia Strack for their tireless effort.

Finally, I would like to thank my family, especially my mother Judith, my two brothers Dennis and Daniel, my cousin Jil and my friends Ruben Zimmermann, Dennis Kimm, Michael Knorr, Kevin Blasiak, Rahii Katyal, David Worms, Daniel Grabowski, Karsten Kucharczyk, Benjamin Fiorelli, Gabriel Mansour and Pasqual Siegel for their mental support, especially during struggling times.

I dedicate this thesis to my parents Judith and Wolfgang, my two brothers Dennis

and Daniel and my better half Vanessa. You mean everything to me.

Table of Contents

Essay I:

Has Monetary Policy Really Become Less Effective in the Euro Area? A Note

This paper has been published as:

Finck, David, "Has Monetary Policy Really Become Less Effective in the Euro Area? A Note", *Applied Economics Letters*, 2019, *26*(13), 1087–1091.

Acknowledgement

I thank Jouchi Nakajima, Giorgio Primiceri and Peter Tillmann for insightful comments.

Has Monetary Policy Really Become Less Effective in the Euro Area? A Note

David Finck*

Abstract

This paper applies a time-varying VAR model with stochastic volatility to the euro area. In contrast to the literature, we find that (i) monetary policy has not become less effective and that (ii) the expansionary policy that is currently pursued would not have resulted in a less severe recession in 2009.

Keywords: Bayesian TVP-VAR, Non-Systematic Monetary Policy

JEL classification: C11, E58, E52

*University of Giessen, email: david.finck@wirtschaft.uni-giessen.de

I Introduction

Since the recent Global Financial Crisis that started in 2007, the monetary transmission mechanism has become one of the hottest and most studied fields in monetary economics for many currency areas. Focusing on the euro area, there is a heating debate about whether monetary policy became less effective. In fact, there is a broad range of literature that discusses potential sources that are able to affect the transmission of monetary policy, ranging from structural, demographic, and institutional changes to the above mentioned financial crisis itself and globalization (see, for instance, Kara and von Thadden, 2016; Oinonen and Paloviita, 2014; Georgiadis and Mehl, 2016; Weber et al., 2009). However, most authors focus on the development of a specific transmission channel. Therefore, the question whether the overall effectiveness of monetary policy has weakened, cannot clearly be answered. One reason is that changes in a specific transmission channel can partly offset effects that possibly influence the effectiveness of other transmission channels as well.

This paper takes a more general view and tries to answer the question whether the overall effectiveness of monetary policy in the euro area has changed. Therefore, we estimate a time-varying VAR model in the style of Primiceri (2005) including three variables, namely inflation, unemployment and a short-term (shadow) interest rate. This class of model is well suited to uncover structural breaks in the economy, as continuous drifting coefficients are able to capture learning dynamics of both, private agents and policy makers. There are two research questions that we try to answer: (i) Has the effectiveness of non-systematic monetary policy (i.e. the way the economy responds to monetary policy shocks) changed over time? (ii) Does a shifting policy rule explain substantial parts of the fluctuations of our endogenous variables, i.e. would loose policy have resulted in a less severe recession?

This paper is organized as follows. Section II provides a short model overview. Section III presents pairwise differences of impulse responses to identified monetary policy shocks as well as a counterfactual simulation. Section IV concludes.

II The Time-Varying VAR in a Nutshell

We estimate a time-varying VAR model for the euro area. The model contains the following $n = 3$ variables at quarterly frequency, mainly taken from the AWM database provided by Fagan et al. (2005): the inflation rate and the unemployment rate represent the non-monetary policy block.[1] Furthermore, a short-term (shadow) interest rate captures the central bank's instrument. More precisely, for periods when the EONIA rate was at the zero lower bound, it is replaced by the shadow rate provided

[1] Choosing unemployment rather than real GDP aims to minimize the problem of data revisions. While output is typically revised several years after the first release, revisions of unemployment figures are very small. In particular, data revisions should matter in a time-varying model, in which drifting parameters should not be confused with revised data.

by Wu and Xia (2016).[2] To capture structural breaks in the parameters and non-linearities in the lag structure, the model allows for both time-varying coefficients and a time-varying variance-covariance matrix. For two lags, the structural VAR takes the form

$$A_t y_t = d_t + F_{1,t} y_{t-1} + F_{2,t} y_{t-2} + u_t \qquad u_t \sim \mathcal{N}(0, \Sigma_t \Sigma_t'), \tag{1}$$

where y_t is an (3×1) vector of observed endogenous variables. d_t is an (3×1) vector of time-varying intercepts. A_t, $F_{1,t}$ and $F_{2,t}$ are (3×3) matrices of time-varying coefficients, u_t is an (3×1) vector of structural shocks. We assume A_t as the lower triangular matrix of the form

$$A_t = \begin{bmatrix} 1 & 0 & 0 \\ \alpha_{21,t} & 1 & 0 \\ \alpha_{31,t} & \alpha_{32,t} & 1 \end{bmatrix} \tag{2}$$

and Σ_t is the diagonal matrix of the form

$$\Sigma_t = \begin{bmatrix} \sigma_{1,t} & 0 & 0 \\ 0 & \sigma_{2,t} & 0 \\ 0 & 0 & \sigma_{3,t} \end{bmatrix}. \tag{3}$$

Multiplying the SVAR in (1) by A_t^{-1}, the model can hence be rewritten into the reduced form as

$$y_t = c_t + B_{1,t} y_{t-1} + B_{2,t} y_{t-2} + A_t^{-1} \Sigma_t \varepsilon_t, \qquad \varepsilon_t \sim \mathcal{N}(0, I_3), \tag{4}$$

where $c_t = A_t^{-1} d_t$ and $B_{i,t} = A_t^{-1} F_{i,t}$ for $i = 1, 2$. The variance-covariance matrix Ω_t of the reduced form VAR can be reduced to

$$A_t \Omega_t A_t^T = \Sigma_t \Sigma_t. \tag{5}$$

We now define the vectors $x_t' = I_n \otimes [1, y_{t-1}', y_{t-2}']$ and $B_t = \text{vec}[c_t, B_{1,t}, B_{2,t}]$, where B_t is a vector of all stacked right-hand side coefficients of (4) and $\otimes$ denotes the Kronecker product. Thus (4) can be rewritten as

$$y_t = x_t' B_t + A_t^{-1} \Sigma_t \varepsilon_t. \tag{6}$$

[2]The shadow rate refers to the short-term policy rate which reflects unconventional measures such as Quantitative Easing and Forward Guidance. It is derived from a term structure model as the implicit short rate consistent with the evolution of long-term rates. In other words it provides an estimate of how the policy rate would be if the zero lower bound was not binding. We also tested the shadow rate provided by Krippner (2013) for robustness, the results were very much the same.

Define α_t as the vector of free elements in A_t, such that $\alpha_t = [\alpha_{21,t}, \alpha_{31,t}, \alpha_{32,t}]'$. Also define σ_t as the vector of diagonal elements of Σ_t, so $\sigma_t = [\sigma_{1,t}, \sigma_{2,t}, \sigma_{3,t},]'$. Summing up, the following system of equations for our parameters of interest specifies the dynamics of the time-varying parameters

$$\begin{aligned} B_t &= B_{t-1} + \nu_t, && \nu_t \sim \mathcal{N}(0, Q) \\ \alpha_t &= \alpha_{t-1} + \kappa_t, && \kappa_t \sim \mathcal{N}(0, S) \\ \log \sigma_t &= \log \sigma_{t-1} + \eta_t, && \eta_t \sim \mathcal{N}(0, W). \end{aligned}$$

The elements of B_t as well as those of A_t are modeled as random walks. The deviations σ_t are assumed to follow a geometric random walk and belong to the so-called class of stochastic volatility.

All innovations in the model are assumed to be jointly normally distributed. Taken together, this can be modeled with the following variance-covariance matrix:

$$V = Var\left(\begin{bmatrix} \varepsilon_t \\ \nu_t \\ \kappa_t \\ \eta_t \end{bmatrix}\right) = \begin{bmatrix} I_3 & 0 & 0 & 0 \\ 0 & Q & 0 & 0 \\ 0 & 0 & S & 0 \\ 0 & 0 & 0 & W \end{bmatrix}, \tag{7}$$

where Q, S and W are positive definite matrices.[3] We construct a Markov Chain using a two-block Gibbs sampler, as in Del Negro and Primiceri (2015).

More precisely, we draw Σ^T using a Metropolis-Hastings step in the first block and $[B^T, A^T, V, s^T]$ as well as s^T in the second block, where B^T, A^T and s^T correspond to the histories of the respective parameters.[4] Summing up, the Gibbs sampling procedure takes the form

1. Draw a candidate $\tilde{\Sigma}^T$ from $\tilde{p}(\Sigma^T | y^T, \theta, s^T)$

$$\Sigma^{(m)^T} = \begin{cases} \tilde{\Sigma}^T & \text{with probability} \quad \alpha \\ \Sigma^{(m-1)^T} & \text{with probability} \quad 1-\alpha \end{cases}$$

2. Draw (θ, s^T) from $p(\theta, s^T | y^T, \Sigma^T)$ by:

   (i) drawing θ from $p(\theta | y^T, \Sigma^T)$

   (ii) drawing s^T from $\tilde{p}(s^T | y^T, \Sigma^T, \theta)$.

[3]We further simplify inference by specifying S as $S = Var(\kappa_t) = \begin{bmatrix} S_1 & 0_{1\times 2} \\ 0_{2\times 1} & S_2 \end{bmatrix}$, where $S_1 = Var(\kappa_{21,t})$ and $S_2 = Var([\kappa_{31,t}, \kappa_{32,t}]')$.

[4]Notice that s^T corresponds to the history of mixture of normals. This is necessary since drawing volatility states requires a state space form with non-Gaussian but $\log_{\chi^2}(1)$ distributed errors in the measurement equation.

The data ranges from 1971Q1 to 2015Q4. The first ten years are used as a training sample.
We apply a Cholesky-decomposition of the variance-covariance matrix Ω_t to identify monetary policy shocks. In particular, we assume that unemployment and inflation respond to an unanticipated interest rate hike with at least one period of delay. The ordering of variables is $y_t = [\pi_t, \psi_t, r_t]'$, where π_t is the inflation rate, ψ_t the unemployment rate and r_t the short-term (shadow) interest rate. The priors are mainly

	Median	Mean	Min	Max	10^{th}	90^{th}
V	14.96	18.91	12.70	143.72	13.48	18.04
Σ^T	13.55	16.43	4.99	79.64	6.36	27.86
A^T	4.62	8.35	1.30	84.50	1.45	19.00
B^T	2.63	3.36	1.30	13.05	1.71	6.26

Table 1: Distribution of inefficiency factors for the entire parameter space.

calibrated as in Primiceri (2005), with only slight modifications for the scaling factors of $k_Q(=0.05)$, $k_W(=0.04)$ and $k_S(=0.2)$. We get 50000 draws and discard the first 30000 draws. Further, we thin our Gibbs sequence by picking every 20^{th} draw, such that 1000 draws remain. Relying on inefficiency factors, our Markov Chain seems to mix fast as reported in Table (1).[5]

III Results

A. *Non-Systematic Monetary Policy*

The first question in the introduction we try to answer is whether non-systematic monetary policy has changed. This is done in two ways. First, we analyze the stochastic volatility of our identified monetary policy shocks. Second, we compare impulse

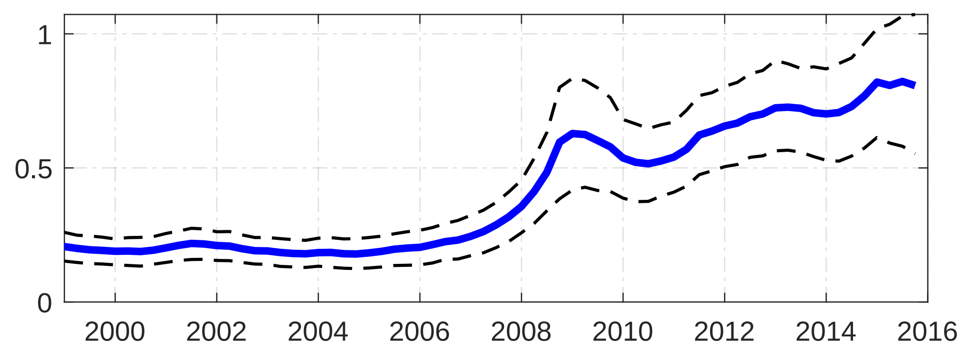

Figure 1: Posterior mean (solid), 16^{th} and 84^{th} (dotted) percentiles of the standard deviation of monetary policy shocks.

response functions of both, inflation and unemployment at three different points in

[5] Further convergence checks include the Raftery-Lewis statistics as well as the decay of autocorrelation between draws which are available on request.

time. Figure (1) plots the stochastic volatility of our identified monetary policy shocks. Two things stand out. First, the period 2008 up to now exhibits a substantially higher variance, which marginally declines after 2009 but remains at a remarkably persistent higher level than in the pre-crisis period. This is not surprising and reflects the effects

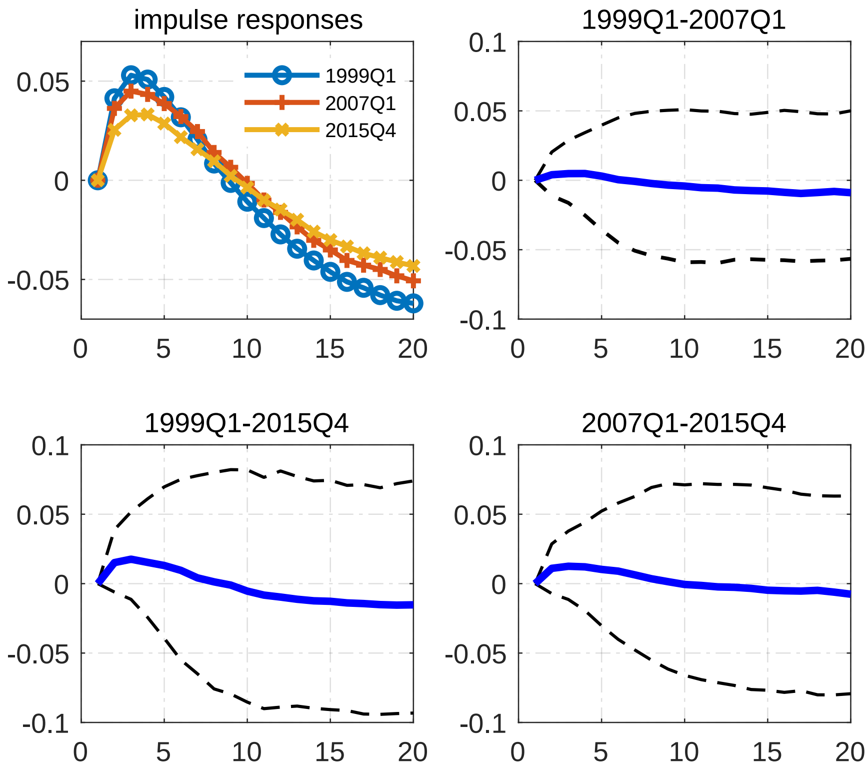

Figure 2: (a) Impulse response of inflation to monetary policy shocks, (b)-(d) pairwise differences with 16^{th} and 84^{th} percentiles (dotted) between (b) 1999Q1 and 2007Q1, (c) 1999Q1 and 2015Q4 and (d) 2007Q1 and 2015Q4.

of the Global Financial Crisis that started in 2007 as well as the subsequent European debt crisis. The reason for this is, given we model a simplified policy rule only taking unemployment and inflation into account, any other variable which has an impact on the interest rate is captured as a residual. Since the disturbances that arose as a result of the Lehman collapse in 2008 for example clearly forced the central bank to also respond to other movements (such as financial sector disturbances) than mainly to unemployment and inflation, this in turn is captured as stochastic volatility movements of structural shocks.

Figures (2) and (3) show the median and pairwise differences of the time-varying impulse responses for inflation and unemployment, respectively, at three different points in time, namely 1999Q1, 2007Q1 and 2015Q4, such that Quantitative Easing that started in 2015 is also included. For the inflation case, it stands out that the transmission of an unexpected monetary policy shock gradually became weaker on the course of time. Even though the differences seem to be small, this finding does also hold for other dates. However, the pairwise differences show that this finding

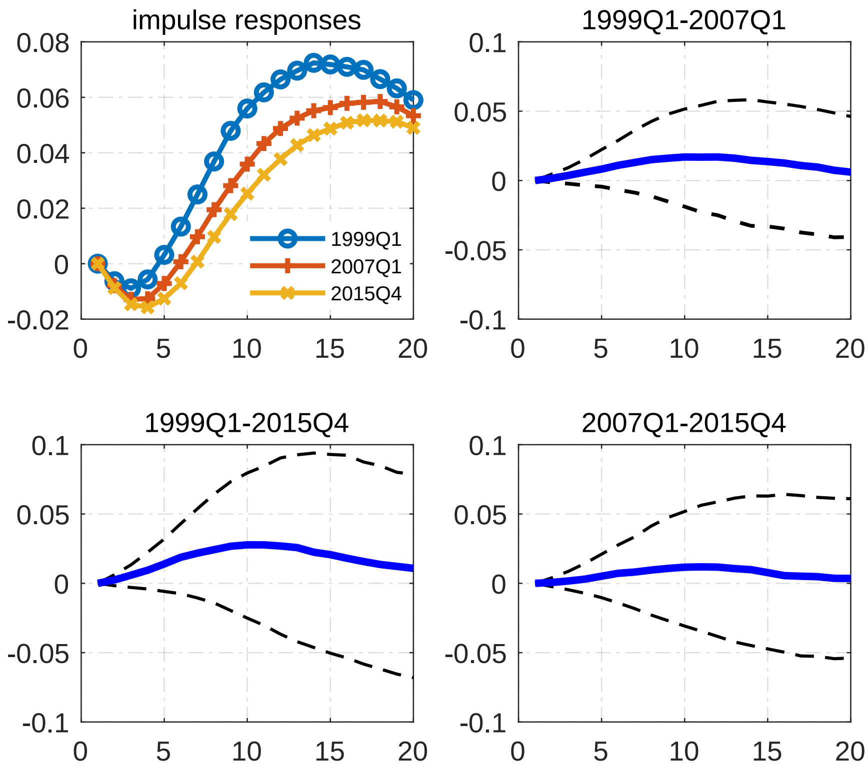

Figure 3: (a) impulse response of unemployment to monetary policy shocks, pairwise differences with 16^{th} and 84^{th} percentiles (dotted) between (b) 1999Q1 and 2007Q1, (c) 1999Q1 and 2015Q4 and (d) 2007Q1 and 2015Q4.

might be spurious as in any case the impulse responses are not significantly different compared to their counterparts. The impulse responses for unemployment indicate a similar picture insofar as the transmission became less effective. The perceived time-variation is higher than in the former case. But also for unemployment we cannot conclude that the reaction became significantly weaker. Summing up the findings, we conclude that there is an infirm pronounced trend towards a weaker transmission, which is however not significant.

B. *Shifting Policy Rule*

Potentially important implications for the way the economy responds to monetary shocks are based on how the ECB reacts to economic activity, thus referring to the second question mentioned in the introduction. To put it differently, analyzing unanticipated policy changes begs for answers of how systematic monetary policy changes potentially affect the economy. We therefore plant the most recent average policy rule of 2014Q1-2015Q4 into 2005Q1. This counterfactual simulation helps to analyze whether different policy rules would have resulted in different outcomes for the non-policy block, i.e. inflation and unemployment. Figure (4) shows the simulated coun-

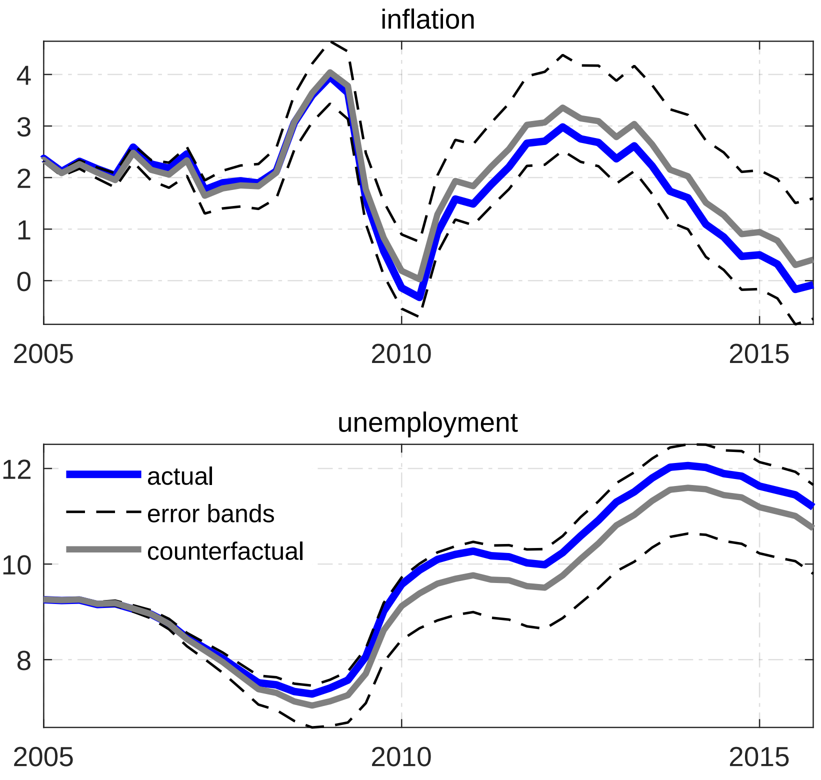

Figure 4: Counterfactual historical simulation of (a) inflation and (b) unemployment with 16^{th} and 84^{th} percentiles.

terfactual paths of both, inflation and unemployment.[6] Notice that the simulated path for inflation lies above, and the counterpart for unemployment below the actual path, reflecting a more "aggressive" policy rule. However, the loose monetary policy that is pursued by the ECB, would not have resulted in a significantly less severe recession, as in both cases the counterfactual paths do not lie outside the 16^{th} and 84^{th} percentiles. We, therefore, conclude that a shifting policy rule does not explain substantial parts of the fluctuation of both inflation and unemployment.

IV Conclusion

There are weak tendencies towards less effective monetary policy. However, our results show that neither the impulse responses differ significantly, nor would a shifting policy rule explain substantial parts of the fluctuation in the real economy.

[6] Note that potential drawbacks of counterfactual simulations like the Lucas critique are mitigated in Bayesian frameworks, see Primiceri (2005) for discussions.

References

Del Negro, Marco and Giorgio E Primiceri, "Time Varying Structural Vector Autoregressions and Monetary Policy: A Corrigendum," *The Review of Economic Studies*, 2015, *82* (4), 1342–1345.

Fagan, Gabriel, Jerome Henry, and Ricardo Mestre, "An Area-Wide Model for the Euro Area," *Economic Modelling*, 2005, *22* (1), 39–59.

Georgiadis, Georgios and Arnaud Mehl, "Financial Globalisation and Monetary Policy Effectiveness," *Journal of International Economics*, 2016, *103*, 200–212.

Kara, Engin and Leopold von Thadden, "Interest Rate Effects of Demographic Changes in a New Keynesian Life-Cycle Framework," *Macroeconomic Dynamics*, 2016, *20* (1), 120–164.

Krippner, Leo, "Measuring the Stance of Monetary Policy in Zero Lower Bound Environments," *Economics Letters*, 2013, *118* (1), 135–138.

Oinonen, Sami and Maritta Paloviita, "Updating the Euro Area Phillips Curve: The Slope Has Increased," *Bank of Finland Research Discussion Papers 31*, 2014, *Bank of Finland.*

Primiceri, Giorgio E, "Time Varying Structural Vector Autoregressions and Monetary Policy," *The Review of Economic Studies*, 2005, *72* (3), 821–852.

Weber, Axel A, Rafael Gerke, and Andreas Worms, "Has the Monetary Transmission Process in the Euro Area Changed? Evidence Based on VAR Estimates," *BIS Working Paper No. 276*, 2009, *Bank for International Settlements.*

Wu, Jing Cynthia and Fan Dora Xia, "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound," *Journal of Money, Credit and Banking*, 2016, *48* (2-3), 253–291.

Essay II:

Mortgage Debt and Time-Varying Monetary Policy Transmission

This paper has been published as:

> **Finck, David, Jörg Schmidt, and Peter Tillmann**, "Mortgage Debt and Time-Varying Monetary Policy Transmission," *Macroeconomic Dynamics, forthcoming*, 2021.

This paper was presented at the following workshops and international conferences:

I: 2nd Workshop on Financial Econometrics and Empirical Modeling of Financial Markets (Kiel, Germany)
Date: May 2018
Presenter: David Finck

II: 93rd WEAI Annual Conference (Vancouver, Canada)
Date: June 2018
Presenter: David Finck

III: 24th CEF Conference (Milan, Italy)
Date: June 2018
Presenter: Jörg Schmidt

IV: 33rd Annual Congress of the EEA (Cologne, Germany)
Date: August 2018
Presenter: David Finck

Acknowledgement

We thank the anonymous referees, Mathias Klein, Jouchi Nakajima and Sarah Zubairy for important input. Participants of the 2nd Workshop on Financial Econometrics and Empirical Modeling of Financial Markets, the 93rd WEAI Annual Conference, the 33rd Annual Congress of the EEA, the 24th CEF Conference, and of seminars at University Giessen provided important feedback. We thank the Fritz Thyssen Stiftung for financial support (grant no. 10.16.2.004WW).

Mortgage Debt and Time-Varying Monetary Policy Transmission

David Finck* Jörg Schmidt† Peter Tillmann‡

Abstract

This paper studies the role of monetary policy for the dynamics of US mortgage debt, which accounts for the largest part of household debt. A time-varying parameter vector autoregressive (VAR) model allows us to study the variation in the sensitivity of mortgage debt to monetary policy. We find that an identically sized policy shock became less effective over time. We use a dynamic stochastic general equilibrium model to show that a fall in the share of adjustable rate mortgages (ARMs) can replicate this finding. Calibrating the model to the drop in the ARM share since the 1980s yields a decline in the sensitivity of housing debt to monetary policy which is quantitatively similar to the VAR results. A sacrifice ratio for mortgage debt reveals that a policy tightening directed toward reducing household debt became more expensive in terms of a loss in employment. Counterfactuals show that this result cannot be attributed to changes in monetary policy itself.

Keywords: Mortgage Debt, Monetary Policy, Debt Reduction, Time-Varying VAR, DSGE

JEL classification: E3, E5, G2

*University of Giessen, email: david.finck@wirtschaft.uni-giessen.de
†University of Giessen, email: joerg.h.schmidt@wirtschaft.uni-giessen.de
‡University of Giessen, email: peter.tillmann@wirtschaft.uni-giessen.de

I Introduction

The build-up of household debt in the USA and other countries is often interpreted as a potential risk to financial stability (see Jordà et al., 2016) and a determinant of the overall credit cycle (see Mian et al., 2017). Since the recent financial crisis originated in the US housing market, the mortgage market receives a lot of attention. Mortgage debt is by far the largest component of household debt, as it reflects the single, most important financial decision of most households. Monetary policy should affect not only the value of houses, but also the dynamics of mortgage debt. A monetary tightening should curb the build-up of mortgage debt. This is one important channel for monetary policy transmission to households.[1]

In this paper, we study the time-varying sensitivity of US mortgage debt to monetary policy. To do so, a model that allows for time variation in the link between the Fed and the mortgage market is needed. We therefore use a time-varying parameter vector autoregressive (TVP-VAR) model along the lines of Primiceri (2005).

The choice of this model is based on two observations: first, there is strong evidence that the US economy underwent both structural and institutional changes and also faces structural changes in the world economy, as emphasized by Canova and Gambetti (2009), Boivin (2006) and Mishkin (2009). Second, financial liberalization and deregulation changed the process of financial intermediation in the US economy. These structural changes cannot be picked up by a linear model with constant VAR parameters and a constant variance-covariance matrix. As a result, there is a danger of incorrect policy implications for central banks if the econometric model is not flexible enough to detect possible time-varying effects. We therefore rely on a model which allows for drifting coefficients and a time-varying variance-covariance matrix.

A model with discrete breaks would be an alternative estimation framework to describe the sensitivity of household mortgages to monetary policy shocks. However, a TVP-VAR is far more flexible than a model with discrete breaks and thus captures any form of structural shifts.[2]

Our model includes four variables: civilian unemployment, GDP deflator inflation and the deviation of real debt from its trend, intended to represent the non-policy block, and a short-term interest rate representing the monetary policy instrument. We assume that the Fed only responds to inflation and employment and restrict all time-varying VAR coefficients in the policy rule other than those related to inflation and employment to zero. Our estimation procedure relies on Markov chain Monte Carlo (MCMC) methods as in Nakajima (2011).

Our key result is that the reaction of mortgage debt to an identically sized monetary

[1] See Jordà et al. (2015) for historical evidence of the link between loose monetary policy and real estate lending booms.

[2] Huber and Punzi (2020) also use a TVP-VAR model to study the transmission of monetary to the housing market at the zero lower bound.

policy shock became much smaller over time.[3] Hence, a monetary policy tightening today reduces household debt much less than the same shock in the 1970s. For instance, a 25bp tightening in 1960 led to a drop in the cyclical component of mortgage debt by about 0.09 percentage points after eight quarters, while the same shock today would result in an insignificant reduction of mortgage debt. Importantly, we show that non-mortgage debt of households does not exhibit such a decline in its sensitivity to monetary policy.

We invoke the concept of the sacrifice ratio to shed light on when a debt reduction as a result of a monetary tightening is relatively expensive or cheap, respectively. We find that toward the end of the sample, it is particularly costly in terms of employment to use a monetary tightening in order to initiate a reduction in household debt.

To rule out that the change in the transmission to mortgages stems from shifts in the systematic behavior of monetary policy itself, we simulate counterfactuals in the spirit of Sims and Zha (2006) and keep the reaction function of the central bank fixed over time. The results are indistinguishable from our baseline finding. Hence, our results fit into the "mortgage rate conundrum" put forward by Justiniano et al. (2022). They argue that the link between Treasury yields, which party reflect monetary policy, and rates on mortgages weakened over time. Hence, both papers stress the role of some underlying structural changes in the mortgage market that impair the monetary policy transmission mechanism.

In order to understand our finding, we study the interaction of the interest rate elasticity of mortgages with the share of adjustable mortgages (ARMs). The literature points to the ARM share as a key determinant of policy transmission to the housing market (see, among others, Calza et al., 2013; Ehrmann and Ziegelmeyer, 2017; Garriga et al., 2017). The idea is that as long as mortgages are closely tied to short-term interest rates, unexpected increases in the policy rate will quickly lead to a shift in cash flows and mortgage payments, which is particularly important for existing borrowers. In this sense, changes in the policy rate also affect the initial cost of new home loans, which in turn affects the demand for housing. The ARM share in the USA exhibits a decline since the early 1980s, where the ARM share data starts. This encourages us to employ the dynamic stochastic general equilibrium (DSGE) model of Alpanda and Zubairy (2017) and examine exogenous variation in the ARM share. The model features a housing market and an occasionally binding credit constraint along the lines of Iacoviello (2005). The impact of the ARM share is nonlinear: a drop in the ARM share causes a less than proportional decline in the interest rate elasticity of mortgage debt. We calibrate the model to the ARM share in 1985, leaving all other parameters at their sample average, and to the ARM share at the end of the data sample. The resulting responses of mortgage debt to a simulated 25bp tightening

[3]Eickmeier and Hofmann (2013) use a factor-augmented VAR model to study the monetary policy transmission to the housing market. They show that monetary policy contributed to the housing pre-crisis housing boom.

shock quantitatively match the estimated decline in the mortgage response sensitivity to policy shocks.

Our results have important policy implications. First, a weaker transmission of policy impulses to the mortgage market, but not to the real economy, implies that monetary policy is not the right instrument to reduce household sector debt. A policy tightening at the end of the sample period to curb the build-up of mortgage debt is expensive in terms of forgone employment. Hence, we would advise against using monetary policy to counteract financial risks related to household borrowing. This finding is in line with Alpanda and Zubairy (2017) who find that, for example, maximum loan-to-value (LTV) ratios are better macroprudential tools to counteract high household indebtedness.

Second, the results can be interpreted as a case against the "too low for too low" argument. That is, keeping interest rate low for an extended period of time does not contribute to inflating a housing credit boom. Our counterfactuals show that the systematic component of monetary policy contributes little to the time-varying sensitivity of mortgage debt to monetary policy.

Our results are related to several branches of the literature. In the following, we highlight only those papers which we consider most relevant for us. The first field studies structural features of the US mortgage market and relates them to the strength of the transmission process. Calza et al. (2013) present VAR results consistent with the notion that the monetary transmission to housing investment is stronger for a high share of ARM. We extend this line of research by looking at the US economy over time, not at the cross section of countries which differ in the average ARM share. Zeev (2016) presents a partial equilibrium model of the housing market, in which a high share of ARM mortgage contracts amplifies the effect of an interest rate shock. He also presents empirical evidence consistent with this finding. The ARM share is used to interact the economy's response to a credit supply shock. Garriga et al. (2017) build a general equilibrium model with incomplete asset markets in which monetary policy affects housing investment by changing the cost of new mortgages. A high ARM share again intensifies the response of the economy to the monetary policy impulse.[4] Secondly, the result of this paper can be interpreted as an equivalent to the "mortgage rate conundrum" of Justiniano et al. (2022). They establish the finding that the connection between mortgage interest rates and Treasury rates broke down in 2003. Hence, the policy tightening of the Fed in 2004 did not lead to higher mortgage rates. Paul (2020) also uses a TVP-VAR model to study the policy transmission to various asset prices. He finds that the transmission to house prices was particularly weak before the 2008-2009 financial crisis.

The remainder of this paper is structured as follows: Section II sketches the evolution

[4]Harding and Klein (2019) also study the nonlinear impact of monetary policy on the US housing market. They show that monetary policy is effective in periods of deleveraging but not during the build-up of household leverage.

of mortgage debt. Section III introduces the TVP-VAR model and details about the Bayesian estimation. Section IV discusses the main results. Section V focuses on the role of adjustable rate mortgages for our results and presents results from a DSGE model. Sections VI and VII present counterfactual analyses and robustness checks, respectively. Section VIII concludes. An appendix contains additional information.

II Mortgage Debt

This paper focuses on mortgage debt held by US private households, which is the main component of overall household debt. The left panel in Figure (1) depicts the development of postwar real household debt. There has been a steady increase in household debt and mortgage debt, respectively, until the eve of the Great Recession.

For the purpose of this paper, we look at the cyclical component of real mortgage debt, which we derive from Baxter and King (1999) filtering the original series. The filter has a band length of eight and lets frequencies pass that are between 4 and 64 quarters. With this calibration, we account for the average length of financial and debt cycles in the USA and take into account that these cycles are about twice as long as the business cycle (Alpanda and Zubairy, 2019). The resulting cyclical series, the mortgage gap, is depicted in the right panel of Figure (1).

Figure 1: Household debt in the USA

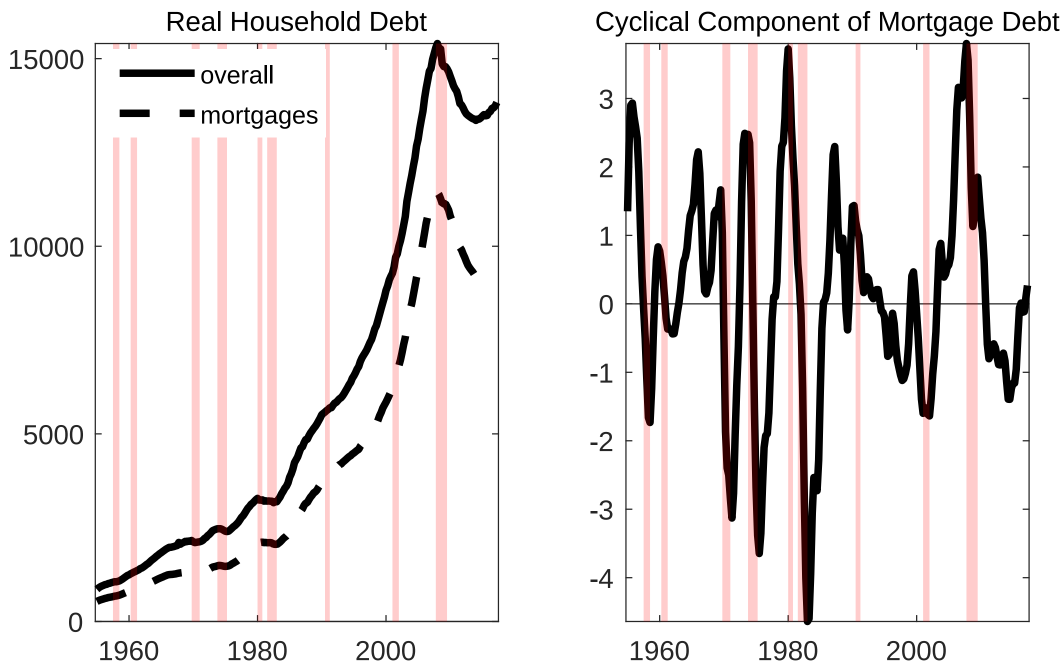

Notes: The solid line in the left panel is overall real household debt. The dashed line in the left panel shows real mortgage debt. The right panel shows the cyclical component of Baxter-King-filtered mortgage debt (in %). The shaded areas reflect NBER-dated recessions.

We find that the mortgage gap fluctuates between 4% and -5% relative to its trend and peaks before each recession. During a recession, the gap turns negative.

III The Empirical Model

We rely on a VAR model with drifting coefficients as well as a drifting variance-covariance matrix. To introduce a few key elements and fix notation, we start with a time-invariant VAR model.

A. *Structural VAR Model*

A standard time-invariant structural VAR is defined as:

$$\mathbf{A}\mathbf{y}_t = \mathbf{d} + \mathbf{F}_1\mathbf{y}_{t-1} + ... + \mathbf{F}_s\mathbf{y}_{t-s} + \varepsilon_t, \qquad t = s+1, ..., T \tag{1}$$

where $\mathbf{y}_t$ is a $k \times 1$ vector of observed endogenous variables, $\mathbf{d}$ is a $k \times 1$ vector which includes deterministic components, and $\mathbf{A}, \mathbf{F}_1, ..., \mathbf{F}_s$ are $k \times k$ matrices of VAR coefficients. Finally, ε_t is a $k \times 1$ vector of structural shocks, where $\varepsilon_t \sim \mathcal{N}(0, \Sigma\Sigma')$.

The vector $\mathbf{y}_t$ contains four variables. The first variable is the unemployment rate, which is our measure of real economic activity. The second variable is inflation, measured as the year-on-year growth rate of the GDP deflator. The third variable is the cyclical component of real mortgage debt derived before. Our fourth variable is supposed to reflect the Fed's policy instrument. We therefore use the effective federal funds rate, augmented with the Wu and Xia (2016) shadow rate during periods in which the federal funds rate was characterized by the zero lower bound (ZLB).[5] Note that our choice of variables other than the cyclical component of mortgage debt is almost the same as in Primiceri (2005).[6] We refer to the first three variables as the non-policy block of the VAR system, for reasons to become clear below.

The simultaneous relations of structural shocks are specified by recursive identification, assuming that $\mathbf{A}$ is lower triangular as:[7]

$$\mathbf{A} = \begin{bmatrix} 1 & 0 & \cdots & 0 \\ a_{21} & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ a_{k1} & \cdots & a_{kk-1} & 1 \end{bmatrix}. \tag{2}$$

Premultiplying both sides by $\mathbf{A}^{-1}$, the model in (1) can hence be rewritten as

$$\mathbf{y}_t = \mathbf{c} + \mathbf{B}_1\mathbf{y}_{t-1} + \cdots + \mathbf{B}_s\mathbf{y}_{t-s} + \mathbf{A}^{-1}\Sigma\varepsilon_t, \qquad \varepsilon_t \sim \mathcal{N}(0, \mathbf{I}_k) \tag{3}$$

[5] The appendix provides additional details on the sources of the variables.

[6] Primiceri (2005) uses the yield on 3-month Treasury bills as opposed to the federal fund as in this paper.

[7] Note that the lower-triangular specification (2) of $\mathbf{A}$ (and thus $\mathbf{A}^{-1}$) is widely used and enables us to easily identify structural shocks (e.g., monetary policy shocks) by recursive ordering, although the examination of implications for the economic structure may require more complicated identification schemes.

where $\mathbf{c} = \mathbf{A}^{-1}\mathbf{d}$ and $\mathbf{B}_j = \mathbf{A}^{-1}\mathbf{F}_j$ for $j = 1, ..., s$. The standard deviations of the structural shocks are captured in the $k \times k$ matrix $\mathbf{\Sigma}$ which is specified as

$$\mathbf{\Sigma} = \begin{bmatrix} \sigma_1 & 0 & \cdots & 0 \\ 0 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & \sigma_k \end{bmatrix}. \tag{4}$$

Stacking all elements of $\mathbf{c}$ and $\mathbf{B}_j$, we get the $k + (k^2 s) \times 1$ vector $\mathbf{B}$. Defining $\mathbf{X}_t' = \mathbf{I}_k \otimes [1, \mathbf{y}_{t-1}', ..., \mathbf{y}_{t-s}']$, the model can be rewritten in reduced-form as

$$\mathbf{y}_t = \mathbf{X}_t'\mathbf{B} + \mathbf{A}^{-1}\mathbf{\Sigma}\varepsilon_t. \tag{5}$$

We order the unemployment rate in $\mathbf{y}_t$ first, the inflation rate second, mortgage debt third and the effective federal funds rate last. This ordering of variables implies that all variables of the non-policy block belong to the information set of the central bank, such that we allow the central bank to react contemporaneously to these variables. At the same time, our timing restriction implies that the monetary policy shock, which we are primarily interested in, does not contemporaneously affect the non-policy block. That is, we assume that unemployment, inflation and mortgage debt react to monetary policy shocks with at least one period delay. While this assumption is standard in the literature (see Christiano et al., 1999) for unemployment and inflation, it is also reasonable to assume that mortgage debt does not contemporaneously respond to unexpected changes in the federal funds rate.

Note that including mortgages into a VAR model also implies that, in principle, the policy reaction function incorporated in the VAR model allows for a feedback from mortgages to monetary policy. Hence, the Fed could respond to changes in household indebtedness. Since we want to keep the model as close as possible to the standard VAR model, even when including mortgages, we restrict the response of monetary policy to mortgages to zero across all lags.[8]

The model is estimated on quarterly data from 1955Q3 to 2017Q2. We use $s = 2$ two lags. This choice comprises a reasonable trade-off between the risk of over-parameterization of our model on the one hand, and autocorrelated residuals on the other.[9]

[8] It turns out that even if we do not impose these restrictions, our results remain indistinguishable.

[9] However, while the marginal likelihood might be one natural candidate to find the optimal lag length within a TVP-VAR, several articles criticized that the harmonic mean method (as was typically used in the literature) could be biased. Therefore, a rolling window estimation can provide an alternative indication of the optimal lag length, because the calculation of the marginal (or log) likelihood within time-invariant VAR models is tractable and well understood. Based on an estimation window of size 160, for each estimation we calculate and store standard information criteria over time. It turns out that this procedure suggests a lag length of four. Although we estimate our baseline model with two lags in order not to overload the model, we also estimate our TVP-VAR with a lag length of $s = 4$. Besides a significantly increased estimation time, our main results remain qualitatively mostly similar, as can be seen in the appendix.

B. *The TVP-VAR Model*

Our TVP-VAR is specified as:

$$\mathbf{y}_t = \mathbf{X}_t'\mathbf{B}_t + \mathbf{A}_t^{-1}\Sigma_t\varepsilon_t, \tag{6}$$

where all parameters, that is, the VAR coefficients captured in $\mathbf{B}_t$, the simultaneous relationships among the endogenous variables captured in $\mathbf{A}_t$ as well as the stochastic volatility of our structural shocks captured in Σ_t, are time-varying. Let $\mathbf{a}_t = [a_{21,t}, \cdots, a_{kk-1,t}]'$ be the vector of non-zero and non-one elements in $\mathbf{A}_t$ (i.e. the lower-triangular elements in $\mathbf{A}_t$) and let $\mathbf{h}_t = [h_{1,t}, \cdots, h_{k,t}]'$ with $h_{i,t} = \log \sigma_{i,t}^2$ for $i = 1, \cdots, k$. Following Primiceri (2005) and others, we assume the following dynamics of the model's parameters

$$\mathbf{B}_{t+1} = \mathbf{B}_t + \eta_{B,t} \tag{7}$$

$$\mathbf{a}_{t+1} = \mathbf{a}_t + \eta_{a,t} \tag{8}$$

$$\mathbf{h}_{t+1} = \mathbf{h}_t + \eta_{h,t}. \tag{9}$$

As regards the initial states of the time-varying parameters, we follow Nakajima (2011) and assume that $\mathbf{B}_{s+1} \sim \mathcal{N}(\mu_{B_0}, \Sigma_{B_0}), \mathbf{a}_{s+1} \sim \mathcal{N}(\mu_{a_0}, \Sigma_{a_0})$ and $\mathbf{h}_{s+1} \sim \mathcal{N}(\mu_{h_0}, \Sigma_{h_0})$. Finally, we follow Primiceri (2005) and assume that all innovations in the model are jointly normally distributed as

$$\mathbf{V} = \mathrm{Var}\left(\begin{bmatrix} \varepsilon_t \\ \eta_{B,t} \\ \eta_{a,t} \\ \eta_{h,t} \end{bmatrix}\right) = \begin{bmatrix} \mathbf{I}_k & 0 & 0 & 0 \\ 0 & \Sigma_B & 0 & 0 \\ 0 & 0 & \Sigma_a & 0 \\ 0 & 0 & 0 & \Sigma_h \end{bmatrix}. \tag{10}$$

Note that the parameters follow a random walk process. Although the random walk process is nonstationary, its assumption can capture both gradual and sudden structural changes.[10] We follow the literature and choose the random walk assumption because we do not want to impose further restrictions on the transition equation a priori when estimating our state-space model. Moreover, the advantage of the random walk assumption is that (1) it allows for permanent shifts and (2) drastically reduces the number of parameters in the estimation procedure.[11]

[10] See Nakajima (2011) and Primiceri (2005) for detailed discussions.

[11] Note that the problem of potential spurious movements is negligible as we choose rather tight priors for the covariance matrix of the disturbance in the random walk process.

C. *Estimation Algorithm and Priors*

In order to estimate the TVP-VAR, we rely on MCMC methods. Our estimation procedure mainly follows Nakajima (2011) and can be summarized as follows: given the data $\mathbf{y} = \{\mathbf{y}_t\}_{t=1}^{T}$, $\boldsymbol{\omega} = (\boldsymbol{\Sigma}_B, \boldsymbol{\Sigma}_a, \boldsymbol{\Sigma}_h)$ and our prior density $\pi(\boldsymbol{\omega})$, we use the following MCMC algorithm to draw samples from the posterior $\pi(\mathbf{B}, \mathbf{a}, \mathbf{h}, \boldsymbol{\omega}|\mathbf{y})$[12]:

1. Initialize $\mathbf{B}, \mathbf{a}, \mathbf{h}$ and $\boldsymbol{\omega}$.
2. Sample $\mathbf{B} \mid \mathbf{a}, \mathbf{h}, \Sigma_B, \mathbf{y}$.
3. Sample $\boldsymbol{\Sigma}_B \mid \mathbf{B}$.
4. Sample $\mathbf{a} \mid \mathbf{B}, \mathbf{h}, \Sigma_a, \mathbf{y}$.
5. Sample $\boldsymbol{\Sigma}_a \mid \mathbf{a}$.
6. Sample $\mathbf{h} \mid \mathbf{B}, \mathbf{a}, \Sigma_h, \mathbf{y}$.
7. Sample $\boldsymbol{\Sigma}_h \mid \mathbf{h}$.
8. Go back to 2.[13]

Priors need to be specified for the starting values of our MCMC algorithm (i.e. for the initial state of the time-varying parameters $(\mu_{B_0}, \Sigma_{B_0})$, $(\mu_{a_0}, \Sigma_{a_0})$ and $(\mu_{h_0}, \Sigma_{h_0})$) and for the i^{th} diagonals of the covariance matrices (i.e. the i^{th} diagonals of $\boldsymbol{\Sigma}_B$, $\boldsymbol{\Sigma}_a$ and $\boldsymbol{\Sigma}_h$). As regards the initial states for the VAR parameters, there are mainly two common practices for specifying the initial states. The first approach follows Primiceri (2005), who uses the point estimates of a time-invariant VAR model estimated from, say, the first 10 years, in order to calibrate the prior. The potential drawback of this approach is that we lose these observations for the estimation of our TVP-VAR, as in a standard Bayesian setting the prior should not contain any information based on the sample. Second, from the standpoint that we do not have any information about the initial state a priori, setting diffuse priors for the initial states is another option (see Nakajima, 2011 for a discussion). In particular, the initial state of our parameters has flat priors, set as

$$\mathbf{B}_{s+1} \sim \mathcal{N}(0, 10 \cdot \mathbf{I}), \qquad \mathbf{a}_{s+1} \sim \mathcal{N}(0, 10 \cdot \mathbf{I}), \qquad \mathbf{h}_{s+1} \sim \mathcal{N}(0, 10 \cdot \mathbf{I}). \tag{11}$$

Of course, the prior choice for the hyperparameters can affect posterior inference, although $\boldsymbol{\Sigma}_B, \boldsymbol{\Sigma}_a$ and $\boldsymbol{\Sigma}_h$ do not parameterize time variation in the first line, but only prior beliefs about the time variation. However, as emphasized by Primiceri (2005), Nakajima (2011) and Koop and Korobilis (2010), the priors should be carefully chosen because our model has many parameters to estimate. This is even more important from the standpoint that all of our model's parameters are modeled as nonstationary random walks. Thus, tight priors for the variance-covariance matrices of the disturbances in the random walk process should avoid ill-determined (implausible) behaviors of

[12]Accordingly, $\mathbf{B} = \{\mathbf{B}_{s+1}, ..., \mathbf{B}_T\}$, $\mathbf{a} = \{\mathbf{a}_{s+1}, ..., \mathbf{a}_T\}$ and $\mathbf{h} = \{\mathbf{h}_{s+1}, ..., \mathbf{h}_T\}$.
[13]For a useful exposition of the MCMC algorithm, see Nakajima (2011).

the parameters. Moreover, it should be noted that in general the time-varying VAR coefficients require a tighter prior than the time-varying variance-covariance matrix. To address this problem, we rely on rather tight priors.

In particular, we set the following priors for the i^{th} diagonals of $\mathbf{\Sigma}_B, \mathbf{\Sigma}_a$ and $\mathbf{\Sigma}_h$:[14]

$$(\mathbf{\Sigma}_B)_i^{-2} \sim \Gamma(25, 6 \cdot 10^{-5}), \qquad (\mathbf{\Sigma}_a)_i^{-2} \sim \Gamma(4, 0.05), \qquad (\mathbf{\Sigma}_h)_i^{-2} \sim \Gamma(4, 0.05).$$

To compute the posterior estimates, we draw $n = 20,000$ draws and discard the first $10,000$ draws, as samples that have been generated in early iteration steps are likely to be not representative for the true posterior distribution. The appendix reports several convergence checks that are common in the literature, including inefficiency factors as well as the convergence checks proposed by Geweke (1992) and Raftery and Lewis (1992). It stands out that all convergence checks justify our prior choice and that our Markov chain mixes quickly.

We further investigate whether the time variation found in our coefficients is a feature of the data or driven by our priors. For that purpose, we apply the test proposed by Cogley and Sargent (2005) and compare the trace of the prior with the trace of the posterior distribution of $\mathbf{\Sigma}_\beta$. We would point to time variation in the parameters stemming from the data rather than the priors if the trace of the prior is below the 16^{th} percentile of the trace of posterior distribution. As can be seen in the appendix, the trace of the prior is indeed well below the 16^{th} percentile of the trace of the posterior of $\mathbf{\Sigma}_\beta$.

IV Results

The advantage of our TVP-VAR model is that we can show time-varying effects of monetary policy shocks. This time variation is not only driven by the estimated VAR coefficients in $\mathbf{B}_t$, but also by the shock impact matrix, $\mathbf{A}_t^{-1}$, as well as the stochastic volatility of the covariance matrix, $\mathbf{\Sigma}_t$.[15]

A. *Responses to a Monetary Policy Shock*

In this section we discuss the time-varying effects of monetary policy shocks on the endogenous variables. To get an impression of the changing nature of monetary policy transmission that accounts for estimation uncertainty, Figure (2) plots the impulse responses following monetary policy shocks originating in 1985Q1 and 2011Q3,

[14]We therewith follow Nakajima (2011) by separate priors for the i^{th} diagonals of the covariance matrices in the transition equations. Primiceri (2005), on the other hand, specifies priors for $\mathbf{\Sigma}_B, \mathbf{\Sigma}_a$ and $\mathbf{\Sigma}_h$ based on an inverse-Wishart distribution. Note that the Wishart distribution is nothing but a multivariate generalization of the gamma distribution. In this case, drawn samples are positive-definite matrices rather than positive real numbers.

[15]In the appendix, we also show that the TVP-VAR model delivers a superior fit to the data compared to a constant parameter VAR model.

respectively. These dates are chosen among episodes of a high and low ARM share, which will become important in the next section. The left column shows the impulse responses following a 25bp monetary policy shock in 1985Q1, when the ARM share was about 60%, whereas the middle column shows the corresponding impulse responses in 2011Q3, when the ARM share reached a minimum of about 10%.

An unexpected policy tightening in 1985 raises unemployment by 0.05 percentage points and leads to a significant fall in mortgage debt by about 0.03%. The response of the inflation rate exhibits a price puzzle, that is, a positive response in the first few quarters, before it starts to fall. A shock in 2011 triggers a similarly sized increase in unemployment and a drop in inflation. However, mortgage debt is no longer responsive to monetary policy.

The right panel shows the difference between the impulse responses in 1985Q1 and 2011Q3, respectively. We would refer to a significantly different transmission of monetary policy shocks when the error bands do not enclose the zero line. Contrary to this, the transmission would not be significantly different across the two points in time when the the zero line lies inside the error bands.

For mortgages the impulse response is significantly stronger in the early 1980s than toward the end of our sample. The responses of unemployment and inflation, however, are not significantly different across the two points in time. Hence, the sensitivity of mortgage debt to monetary policy shocks is lower at the end of our sample period, while the transmission of policy to inflation and unemployment remains unaffected. This is the most important finding of this paper.

Notice that this finding does not only hold for the two selected dates, that is, for 1985Q1 and 2011Q3. Rather, we find that the transmission of monetary policy shocks to mortgage debt gradually decreases. We best see that when we study the cross sections of the time-varying impulse response functions with their respective 16^{th} and 84^{th} percentiles of the posterior distribution.

Figures (3) to (5) show the time-varying reaction of the endogenous variables 4, 8 and 16 quarters after the shock. The full profiles of impulse responses, that is, the response as a function of the timing of the occurrence of the shock and the time after the shock are shown as three-dimensional plots in the appendix.

After four quarters, see Figure (3), the policy tightening has led to higher unemployment. The response of unemployment is significant between the 1970s and 1990s. For inflation, we observe a price puzzle in the 1970s. However, the initial significance of our observed price puzzle begins to vanish over time. Mortgage debt falls significantly following a policy shock originating in the 1960s and 1970s, but remains unchanged for shocks occurring more recently.

Eight quarters after the shock, see Figure (4), unemployment is still higher as a result of the tightening, though the impact on unemployment becomes slightly smaller over time. The response of inflation 2 years after the shock remains insignificant for most

Figure 2: Response to a monetary policy shock for selected dates

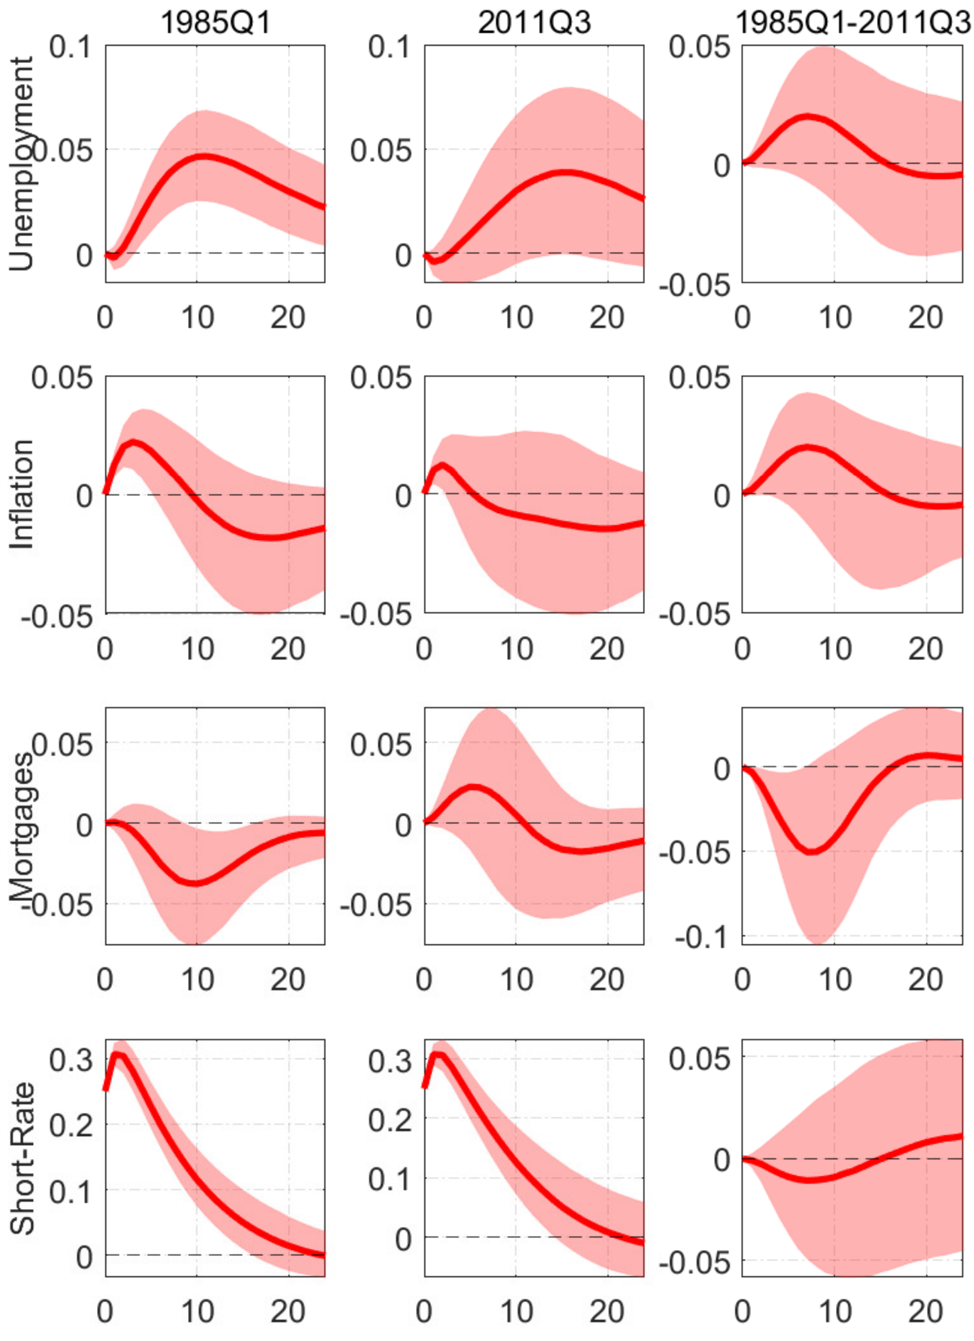

Notes: The left panel shows the mean responses (red-solid) with 16^{th} and 84^{th} percentiles to a 25bp monetary policy shock in 1985Q1, whereas the middle panel shows the mean responses to a 25bp monetary policy shock in 2011Q3. The right panel shows the difference between the responses in 1985Q1 and 2011Q3.

of the sample period. Importantly, we see a very pronounced tendency in the response of mortgage debt: the sensitivity falls steadily over time and disappears in the 2000s.

The impact of the policy shock after 4 years is shown in Figure (5). Again, the response of mortgage debt is negative between the 1970s and the 1990s. After that, mortgages remain insignificantly different from their trend. From a first glance at Figure (5), one might mistakenly infer that the mortgage response will regain strength over time. However, this is due to the fact that the response to a monetary policy shock becomes somewhat more sluggish over time with the peak response occurring later. Thus, especially for the last 20 years of our sample, Figure (5) shows the response of mortgages at the time when the response to monetary policy is strongest, before mortgages return to the trend.

The persistence of the interest rate response to the monetary policy shock does not fluctuate much over time.

Finally, we extract the maximum impact of the monetary policy shock on mortgage debt as well as the time at which the maximum response occurs and show the evolution of both in Figure (6).[16] This allows us to better describe the shift in the sensitivity of mortgage debt to monetary policy over the sample period. We can see the general tendency of diminishing effects of monetary policy shocks with responses becoming insignificant at the end of our sample. We also find that the maximum impact of the policy shock occurs later over time.

The decline in the sensitivity of mortgage debt has strong implications for monetary policy, which we will now study in detail. Keeping in mind that the reduced-form model is not suitable for normative policy analysis, our results suggest that using monetary policy to curb household debt is difficult. This is because the effectiveness of monetary policy shocks strongly varies over time. This also implies that the costs of a debt reduction in terms of unemployment vary over time. We will pick up this thought in the following section.

B. *The Sacrifice Ratio of Debt Reduction*

Let us consider a central bank which aims at using an increase in the policy rate in order to foster a debt reduction of households. How much additional unemployment does the central bank need to accept? We invoke the concept of the sacrifice ratio which in its initial form answers the question of how costly a disinflation is in terms of unemployment. We construct a measure that shows how expensive a debt reduction is in terms of unemployment. This measure, $\varphi_{t,t+h}$, is the ratio of the cumulative response of mortgage debt relative to the cumulative response of unemployment, h periods after the shock occurs. That is, the nominator and the denominator are the

[16]The calculation of the peak timing is based on the posterior mean. In a previous version we showed the median of the peak timing.

Figure 3: Response to a monetary policy shock after four quarters

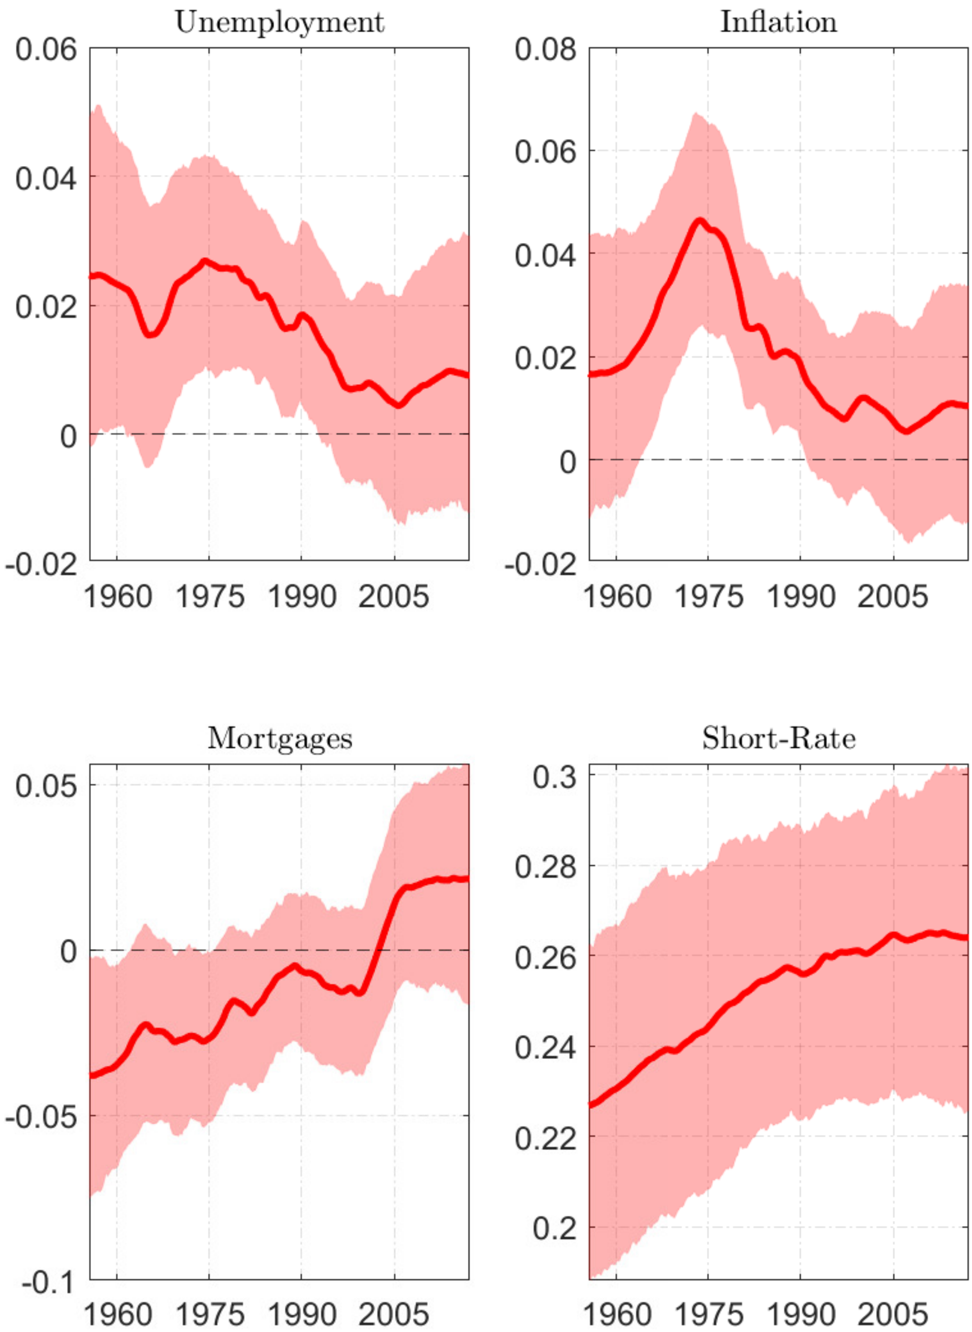

Notes: Mean responses (red-solid) after four quarters with 16^{th} and 84^{th} percentiles to a 25bp monetary policy shock.

Figure 4: Response to a monetary policy shock after eight quarters

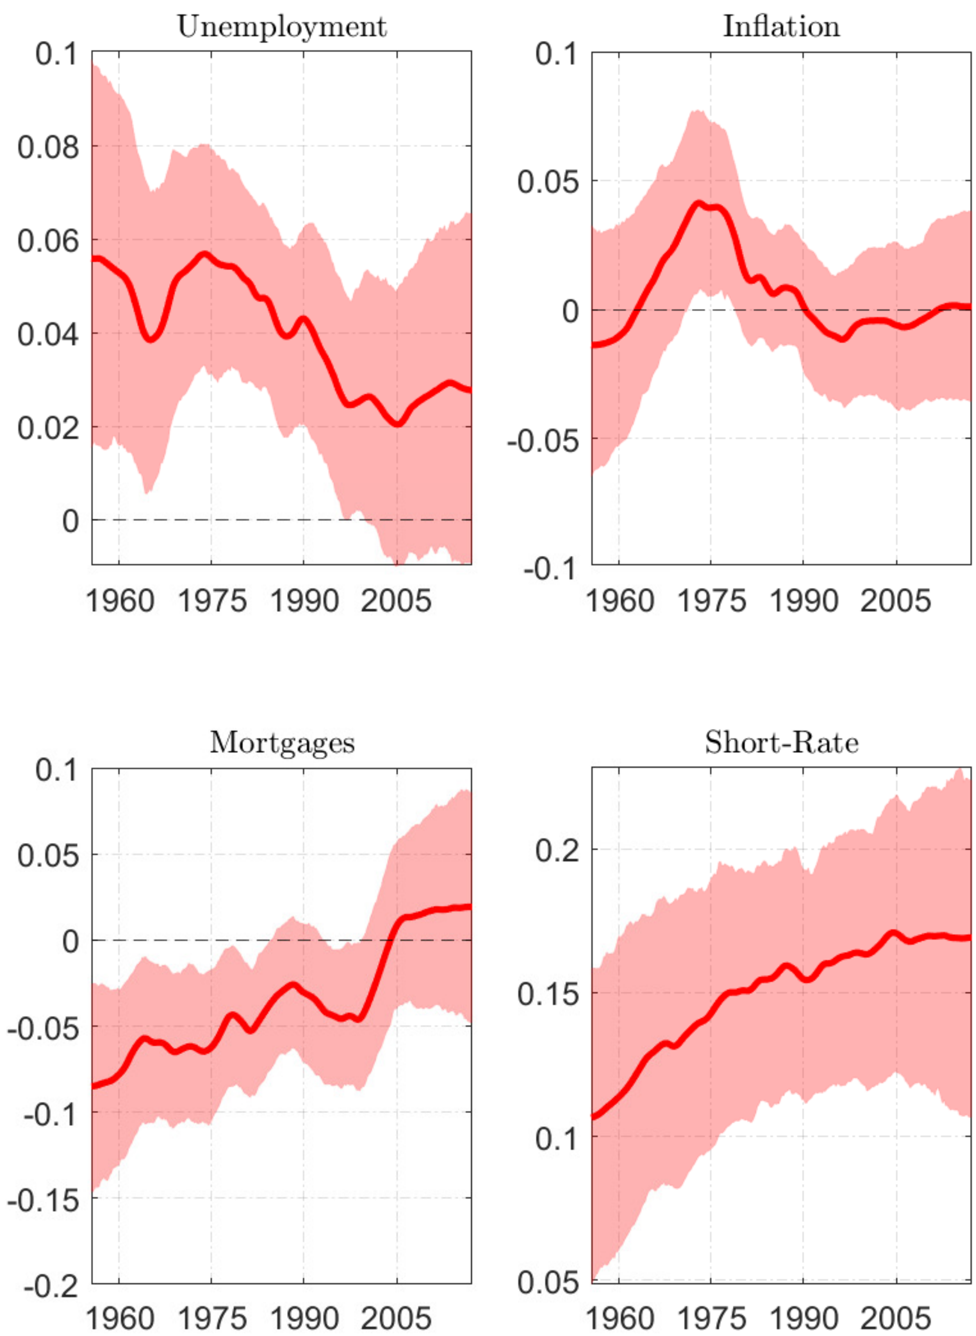

Notes: Mean responses (red-solid) after eight quarters with 16^{th} and 84^{th} percentiles to a 25bp monetary policy shock.

Figure 5: Response to a monetary policy shock after 16 quarters

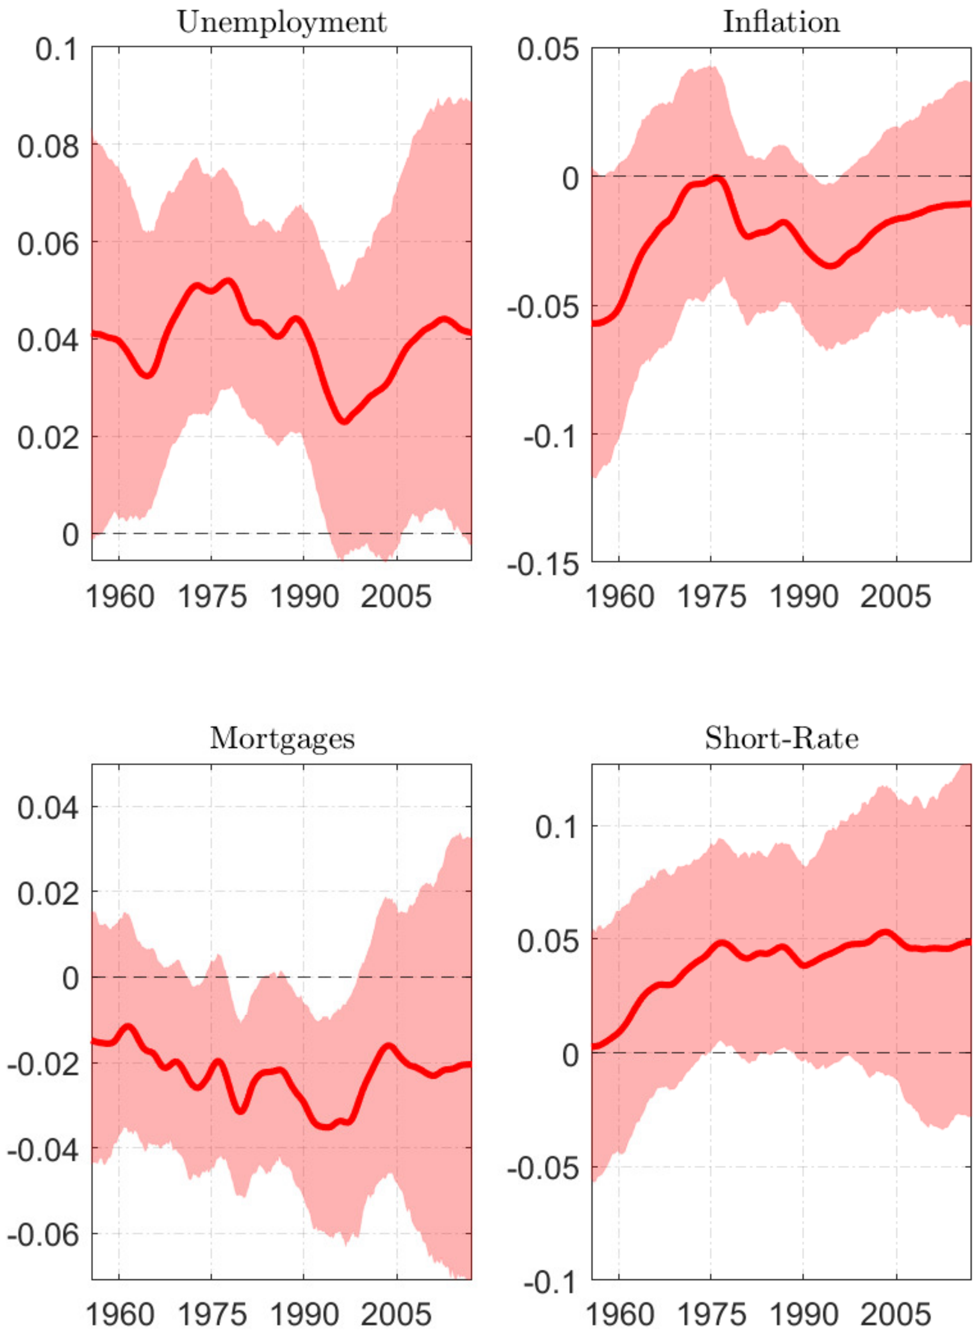

Notes: Mean responses (red-solid) after 16 quarters with 16^{th} and 84^{th} percentiles to a 25bp monetary policy shock.

Figure 6: Peak response of mortgage debt

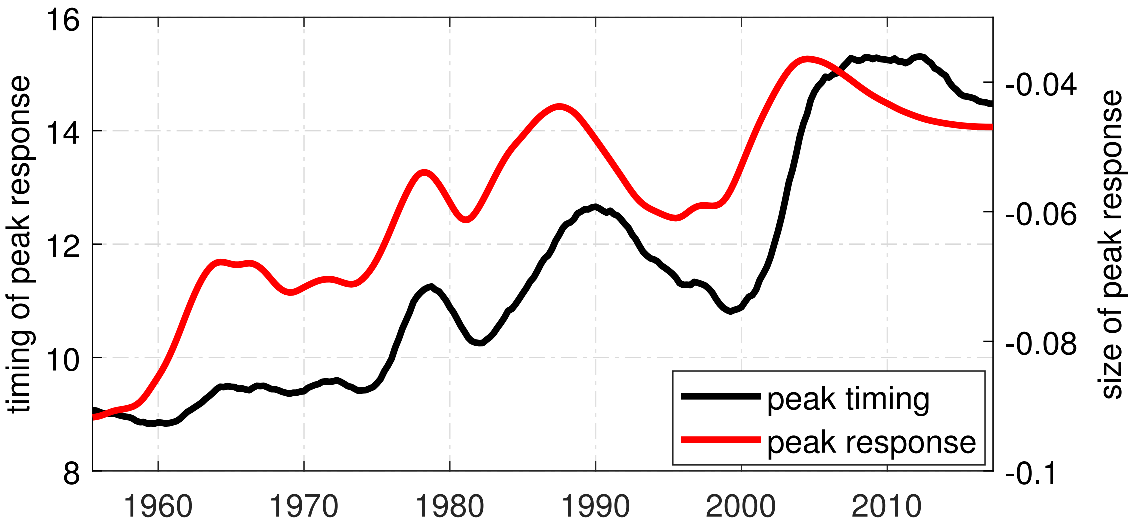

Notes: Peak response and timing of the peak response following a 25bp monetary policy shock.

cumulative mean impulse responses:

$$\varphi_{t,t+h} = \frac{\sum_{i=0}^{h} \text{IRF}_{t+i}^{\text{debt}}}{\sum_{i=0}^{h} \text{IRF}_{t+i}^{\text{unemp}}}, \tag{12}$$

where the horizon h can be interpreted as the relevant horizon for policy-makers. Consider a policy shock that leads to a reduction in debt and an increase in unemployment. Hence, the ratio is negative. If for the same loss in employment debt falls more, this reduction becomes less costly. In this case, the index falls. By the same token, we would observe a drop in the sacrifice ratio if unemployment increases less, given that the reduction of debt remains unchanged. That is, policy-makers were able to achieve a high reduction in mortgage debt per unit increase in unemployment. An increase in $\varphi_{t,t+h}$ would thus be consistent with a more costly reduction, that is, for a one unit increase in unemployment we only get a relatively small reduction in debt. We consider two alternative horizons: $h = 8$ and $h = 12$ quarters. Since we base this thought experiment on a reduced-form model, we should stress that we do not intend to derive normative implications. Rather, we want to find out at what time a debt reduction as a result of a monetary tightening was relatively expensive or cheap, respectively.

The blue-solid lines in Figure (7) display the ratio for the two horizons h, whereas the black-dashed lines plot $\varphi_{t,t+h}$ with h corresponding to the timing of the peak response. We can see an upward trend in both series. Hence, a debt reduction becomes more expensive over time. The trend is broken only during the 1990s.

Technically, the break in the sacrifice ratio results from the changes in the impulse responses of unemployment and mortgages in the first half of the 1990s. The response of unemployment eight quarters after the shock remains horizontal at that time, while the response of mortgages temporarily become somewhat more pronounced. One possible explanation for the behavior of the unemployment response is the sudden weakening of Okun's Law (Owyang and Sekhposyan, 2012). Hence, when the Fed tightened, both mortgage debt and real economic activity fell. Due to the shift in Okun's Law, however, the loss in terms of employment was smaller than before and the debt reduction appeared more favorably when compared to the loss in employment. The sudden shift in the responsiveness of mortgages can best be explained in terms of changes in the ARM share. As shown in below, the ARM share jumped from 20% to more than 50% in the early- to mid-1990s, before it continued its downward trend. This jump is directly reflected in a stronger response of mortgages to monetary policy shocks.

Figure 7: The costs of debt reduction

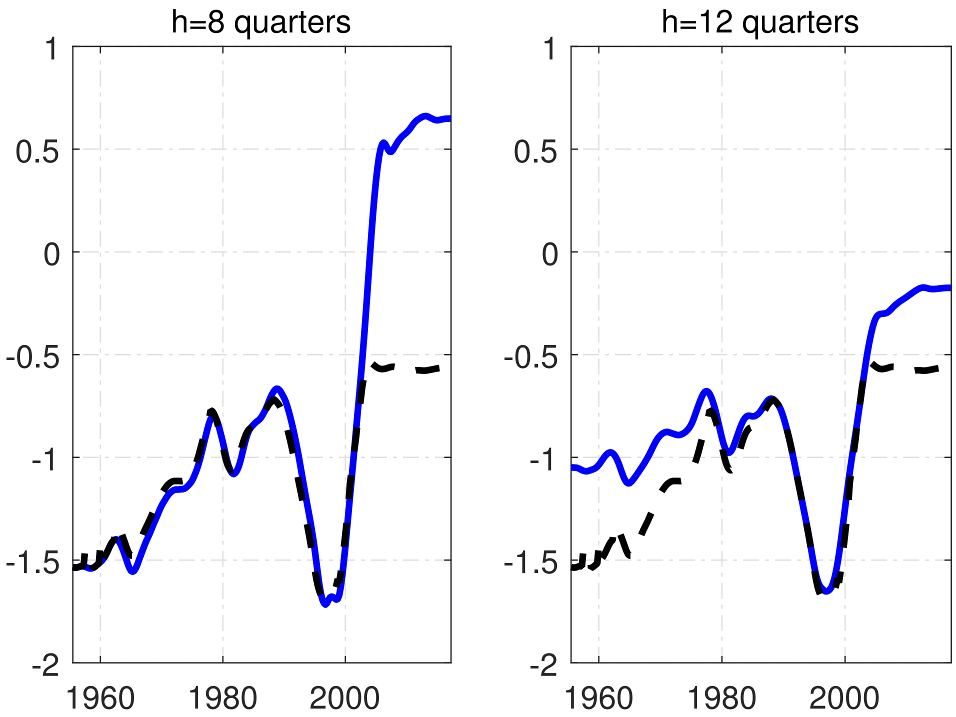

Notes: The blue-solid lines reflect the ratio $\varphi_{t,t+h}$ for horizon $h = 8$ (left panel) and $h = 12$ (right panel) over time. In both panels, the black-dashed line reflect $\varphi_{t,t+h}$ where h corresponds to the timing of the peak response.

Overall, our results mirror-image the time-varying sensitivity of unemployment and debt, that is, that monetary policy-induced debt reductions are (relatively) cheap if either the sensitivity of unemployment to monetary policy shocks is relatively low while the sensitivity of mortgage debt is relatively high or vice versa.

V The Role of Adjustable Rate Mortgages

Many structural forces could potentially be responsible for the decline in the sensitivity of mortgages to Fed policy. Calza et al. (2013) highlight key parameters that have an effect on how sensitive mortgages are to monetary policy. They find that the transmission of monetary policy shocks to consumption is stronger in countries where mortgages with adjustable rates dominate. In the US economy, this share declined over time. Figure (8) plots the ARM share since the early 1980s. While in 1985 the share was about 60%, it fell to only 10% by the end of 2011, implying that

Figure 8: Share of adjustable rate mortgage contracts

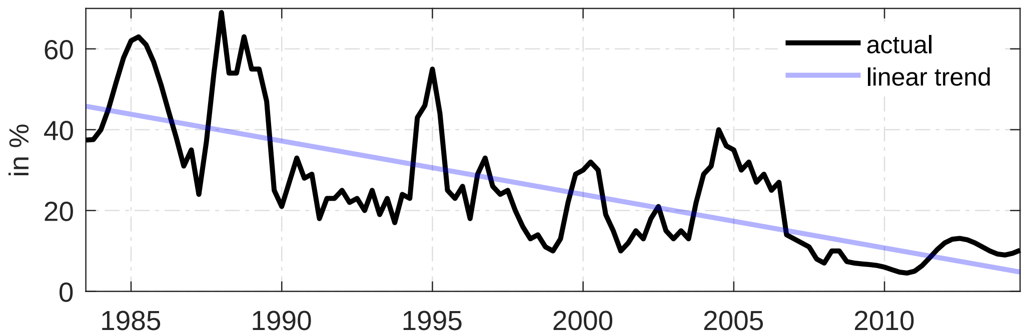

Notes: The data is taken from Federal Housing Agency, Monthly Interest Rate Survey. We use chained data based on data availability. Thus, we use interpolated data extracted from annual basis from 1982Q2 until 1984Q4 and end-of-quarter monthly data from 2008Q4 until 2014Q3.

home-buyers mostly prefer fixed-rate mortgages over adjustable rate mortgages.[17] As emphasized by Alpanda and Zubairy (2017), under fixed-rate contracts the interest rate channel of monetary policy can be impaired, as agents are unable to lower the interest burden on their existing debt by refinancing, especially when their existing debt levels are high and they do not possess adequate equity. We therefore expect that the transmission of monetary policy shocks to the economy is linked to the ARM share. In particular, we expect that the transmission is more powerful when the ARM share is high.

With this in mind, let us briefly return to Figure (2). We show the impulse responses for shocks originating at a time when the ARM share is high (1985Q1) and low (2011Q3), respectively. The two impulse responses are indeed significantly different.

[17] Moench et al. (2010) discuss the reduction in the ARM share in the USA and explain it in terms of financial innovations such as an increase in securitization and a shifting term structure of interest rates.

A. *Illustrative Evidence*

We want to obtain some illustrative evidence on the determinants of time variation in the policy impact on mortgages. For that purpose, we regress the dynamic impulse response functions on the ARM share. We also control for other characteristics of the prevailing contracts in this market.[18] This is done for both the time-varying impulse responses and the cumulative effects of mortgage debt for different periods h after the shock. Formally, we regress

$$\mathrm{IRF}_{t,t+h} = c + \gamma \mathbf{x}_t + \delta t + \varepsilon_t,$$

as well as for impulse responses accumulated up to horizon h,

$$\sum_{i=0}^{h} \mathrm{IRF}_{t,t+i} = c + \gamma \mathbf{x}_t + \delta t + \varepsilon_t,$$

where t is the period the shock occurs and i denotes the number of periods after the shock.

The vector γ collects the coefficients for the control variables in $\mathbf{x}_t$, c is a constant and ε_t is an error term. The matrix $\mathbf{x}_t$ contains the ARM share, the annual growth rate of GDP, the S&P/Case-Shiller US National Home Price Index (in logs), the LTV ratio, and the average effective interest rate paid for mortgage contracts (effective rate), which is also provided by the Federal Housing Agency's Interest Rate Survey. To account for possible nonlinearities between the impulse responses and the ARM share, $\mathbf{x}_t$ also includes the squared ARM share. This is primarily motivated by the implications of our DSGE model, which will be described in the following section. Lastly, we add a linear trend, which is meant to account for the fact that both the dependent variable and some of the explanatory variables are very persistent. HAC standard errors are used with a lag length of eight, as recommended by Wooldridge (2016).

Table (1) shows the outcome of the regressions for $h = 8, 12, 16$ and h being equal to the timing of the peak response. A few things stand out. First, for all horizons under consideration, the joint explanatory power of our control variables is high and in all but one case well above 85%. Most control variables also have a significant impact and can explain some of the variation in the time-varying effects of mortgages. Second, the coefficient for the ARM share shows the expected negative sign in all cases and is significant in almost all cases except for $h = 8$ in the upper panel as well as for the peak response in the bottom panel. This implies that mortgage debt falls more strongly after a policy tightening when the ARM share increases, which is in line with the results found in Alpanda and Zubairy (2017), Calza et al. (2013) and Garriga et al. (2017).

[18] In the next subsection, we simulate a structural model to corroborate our findings.

Table 1: Explaining the time-varying policy impact on mortgages

	Impulse Response			
	Peak	$h = 16$	$h = 12$	$h = 8$
Constant	0.025 (0.039)	0.010 (0.021)	0.052 (0.071)	0.054 (0.123)
ARM share	−0.349** (0.158)	−0.207* (0.119)	−0.452** (0.231)	−0.580 (0.370)
ARM share squared	0.706*** (0.175)	0.366** (0.147)	0.948*** (0.261)	0.125*** (0.039)
Real GDP	0.530* (0.317)	0.995** (0.396)	0.594 (0.544)	−0.192 (0.645)
House prices	0.030*** (0.006)	0.012*** (0.001)	0.022** (0.009)	0.026 (0.021)
LTV	−2.706*** (0.617)	−1.466*** (0.255)	−2.99*** (0.661)	−3.157*** (0.751)
Effective rate	1.065* (0.542)	1.818*** (0.298)	2.873*** (0.746)	2.741* (1.344)
Trend	yes	yes***	yes**	yes
R^2	0.877	0.683	0.878	0.875
adj. R^2	0.817	0.664	0.870	0.867
	Cumulative impulse response			
	Peak	$h = 16$	$h = 12$	$h = 8$
Constant	0.425 (0.975)	0.640 (1.204)	0.550 (1.133)	0.285 (0.683)
ARM share	−3.410 (2.568)	−6.263* (3.558)	−5.194* (3.261)	−3.015* (1.785)
ARM share squared	0.696** (0.297)	0.129** (0.376)	1.080*** (0.342)	0.622*** (0.204)
Real GDP	2.463 (3.462)	2.273 (6.355)	−0.674 (5.234)	−1.993 (3.145)
House prices	0.127 (0.151)	0.285 (0.185)	0.216 (0.194)	0.113 (0.123)
LTV	−19.955*** (0.610)	−35.496*** (8.362)	−27.583*** (6.676)	−13.967*** (3.553)
Effective rate	26.151** (8.424)	26.151*** (8.424)	26.395** (11.427)	12.966* (6.995)
Trend	yes	yes	yes	yes
R^2	0.872	0.888	0.878	0.872
adj. R^2	0.865	0.881	0.870	0.865

Notes: The dependent variable is the response of mortgage debt following a 25bp monetary policy shock when the response peaks (first column), after 16 periods (second column), after 12 periods (third column), after 8 periods (fourth column). The coefficients and standard errors of the ARM share, real GDP, the loan-to-value ratio, the short-rate and trend are multiplied by 1000. The coefficients and standard errors of the squared ARM share are multiplied by 10,000. HAC standard errors are used with a lag length of eight and a Newey-West fixed bandwidth of 5. A significance level of 1%, 5% and 10% is denoted by ***, ** and *.

Second, the coefficient on the squared share of adjustable rate mortgages is positive and significant in all cases. Note that a positive coefficient implies that the effect of an increase in the ARM share is decreasing, the higher the ARM share is. This means that the response of mortgages after an increase in the ARM share is stronger in absolute terms when the ARM share is relatively low. If, on the other hand, the ARM share is relatively high, a further increase in the ARM share will still lead to a stronger response, this effect is however less pronounced than before.

Third, the coefficient on the LTV ratio is negative and significant in all cases. This result is in line with Alpanda and Zubairy (2017) and implies that the effects of monetary policy shocks are amplified when the LTV ratio increases.

Overall, our results support our assumption that the ARM share affects the transmission of monetary policy shocks.

B. *A DSGE Model with Mortgage Contracts*

While the previous regressions are illustrative, eventually a structural model is needed to shed light on the impact of a shift in the ARM share on the strength of policy transmission. Therefore, we resort to the DSGE model by Alpanda and Zubairy (2017).

The model builds on a closed-economy DSGE with housing and household debt as well as an occasionally binding credit constraint. The model features two types of households: patient households (savers) and impatient households (borrowers). In the model, excessive household debt arises as a result of exuberance shocks on expectations on house prices, which drives a wedge between actual and fundamental values.

Importantly, the model allows the average duration of the fixed interest rate for loans to be shorter than the full amortization duration of the underlying loan itself. That is, the interest rate on new mortgage loans is decomposed into a fraction carrying a fixed mortgage interest rate and a fraction of existing loans that is refinanced each period.

We use the model to simulate impulse responses to monetary policy shocks for different calibrations of the ARM share. We simulate impulse responses for mortgage debt to a 25bp monetary policy shock, which is consistent with the definition of the policy shock in the TVP-VAR. In the first case, the "low share" case, we use the same overall calibration as Alpanda and Zubairy (2017), including the interest rate adjustability of mortgages based on a 10% ARM share, which we can observe at the end of our sample. In the second case, the "high share" case, we recalibrate the interest rate adjustability parameter based on an ARM share of 60% as observed in 1985, keeping anything else similar to the benchmark case.[19]

In Figure (9), we compare the impulse responses to 25bp monetary policy shock for both the DSGE model and the TVP-VAR. First, similar to our time-varying VAR, there is a weaker reaction of mortgage debt to monetary policy shocks when the ARM share is low. Second, the amplitude of both the DSGE and the TVP-VAR are nearly identical in the high share case with a 0.05 percentage point drop in mortgages. For the low share case, the TVP-VAR shows a weaker response. However, this can most likely be attributed to the fact that our recalibration was solely based on the different fraction of ARM shares keeping all other parameters unchanged. In the model, the interest rate duration is based on factors other than the ARM share, including home equity loans and repayments. Additionally, we can see that the relationship between the ARM share and the size of the mortgage reaction to a restrictive 25bp monetary policy shock shows a nonlinear pattern.

To conclude this section, we calibrate the DSGE model not just for a high and a low share, but for the full range of observable ARM shares. Figure (10) plots the resulting

[19] In the appendix we explain in detail how we account for different values for the ARM share.

Figure 9: Responses of mortgage debt to monetary policy in the DSGE model and the TVP-VAR

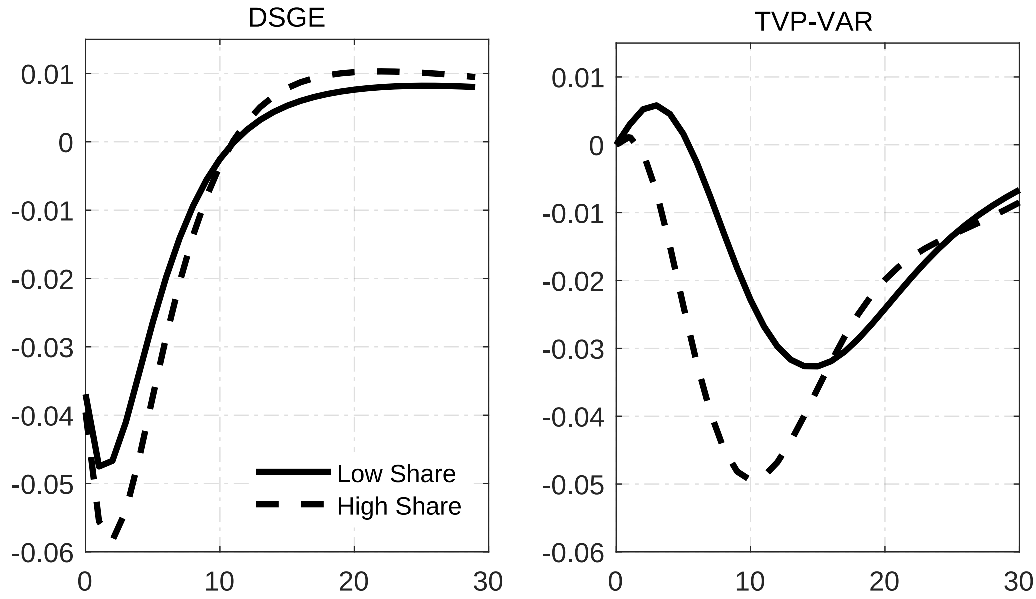

Notes: In both cases, the shock is a surprise increase in the interest rate by 25bp. The DSGE model of Alpanda and Zubairy (2017) is calibrated to a "high share" state with an ARM share of 60% and a "low share" state with a share of 10%.

peak responses to the monetary policy shock against the ARM share. The sensitivity to the policy shock increases nonlinearily in the ARM share, which is why we allow for a nonlinear effect in our regressions reported in Table (1). Nevertheless, although we also find a nonlinear relationship here, the sign in this case does not match the sign from the illustrative regression results.

Summing up, the DSGE model provides further evidence for the role of ARM shares in the transmission of monetary policy shocks. The empirically observed drop in the ARM share since the early 1980s leads to impulse responses which are quantitatively very similar to the responses derived from the TVP-VAR.

VI Counterfactual Impulse Responses

Thus far, our results provide evidence that the transmission of monetary policy shocks may have become weaker over time. Moreover, from the standpoint that our model is supposed to uncover the time-varying structure of the economy, relative cumulative responses provide evidence that periods exist in which reducing debt might be less costly than in others.

One possible driver of a weaker transmission of monetary policy is the structural change in the Fed's policy rule itself. That is, if in the 2000s the Fed responds systematically different to inflation and unemployment than in the 1970s, the responsiveness of the rest of the economy could change. In this section, we want to rule out that

Figure 10: Model-implied peak response of real mortgage debt to monetary policy shock for alternative ARM shares

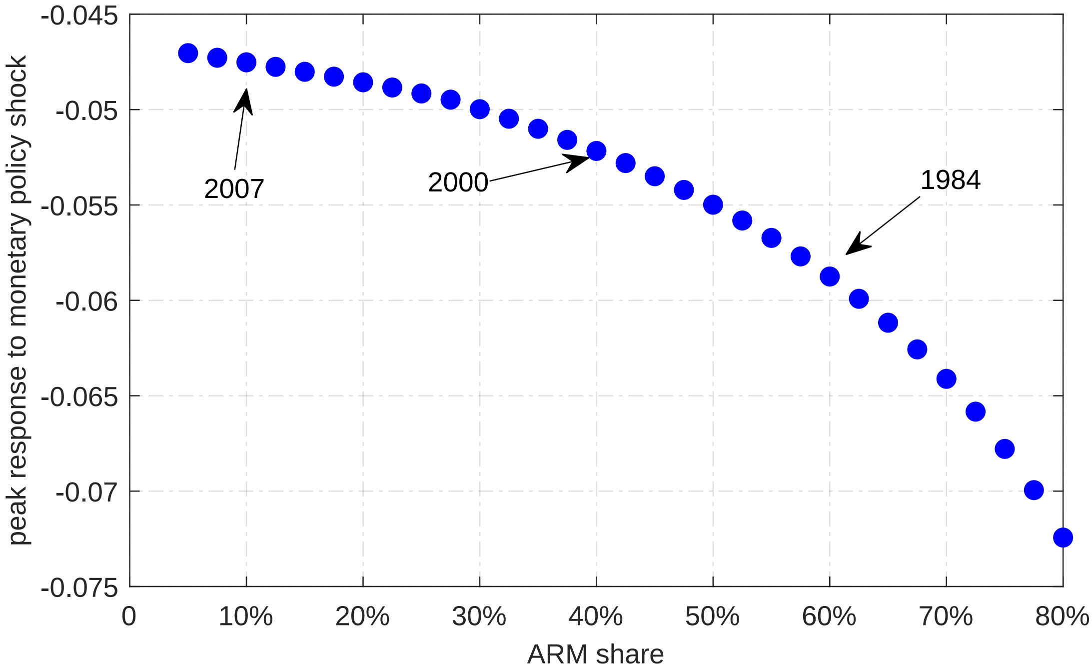

Notes: Response of real mortgage debt in the DSGE model of Alpanda and Zubairy (2017) for different calibrations of the ARM share. To provide some orientation, we highlight three historical ARM shares.

monetary policy itself is driving the change in the transmission mechanism.

We report results for counterfactual experiments that have been widely used (see, for instance, Primiceri, 2005; and Sims and Zha, 2006) and are an informative possibility to establish the role of the Fed in the weaker transmission of monetary policy shocks observed in Section IV. Specifically, we provide counterfactual impulse response functions and counterfactual historical simulations, assuming that the monetary policy rule under Fed chairman Bernanke (February 2006 to January 2014) prevails over the full estimation sample. We therefore fix the parameters in the interest rate equation of the VAR model to the mean from their posterior during the Bernanke tenure, while all other parameters are allowed to fluctuate.[20]

The counterfactual impulse responses of mortgages are shown in Figure (11), which reports also the baseline results from Figure (2) as a benchmark. For both static impulse responses as well as the responses after four and eight quarters, the impulse responses almost completely overlap with the baseline results. This corroborates the notion that

[20]Figure (18) in the appendix shows the corresponding counterfactual historical simulation for the exercise described above. Here, too, we see that a possible shift in the policy rule has hardly contributed to a change in the historical development of our endogenous variables. Only the unemployment rate would have been somewhat lower with the "Bernanke Rule" in the 1980s. The difference is not significant, however, as the observed data lie within the 16^{th} and 84^{th} percentiles of the counterfactual simulation.

the change in the transmission process does not stem from shifts in the Fed's policy rule.

Figure 11: Counterfactual impulse responses under fixed policy rule

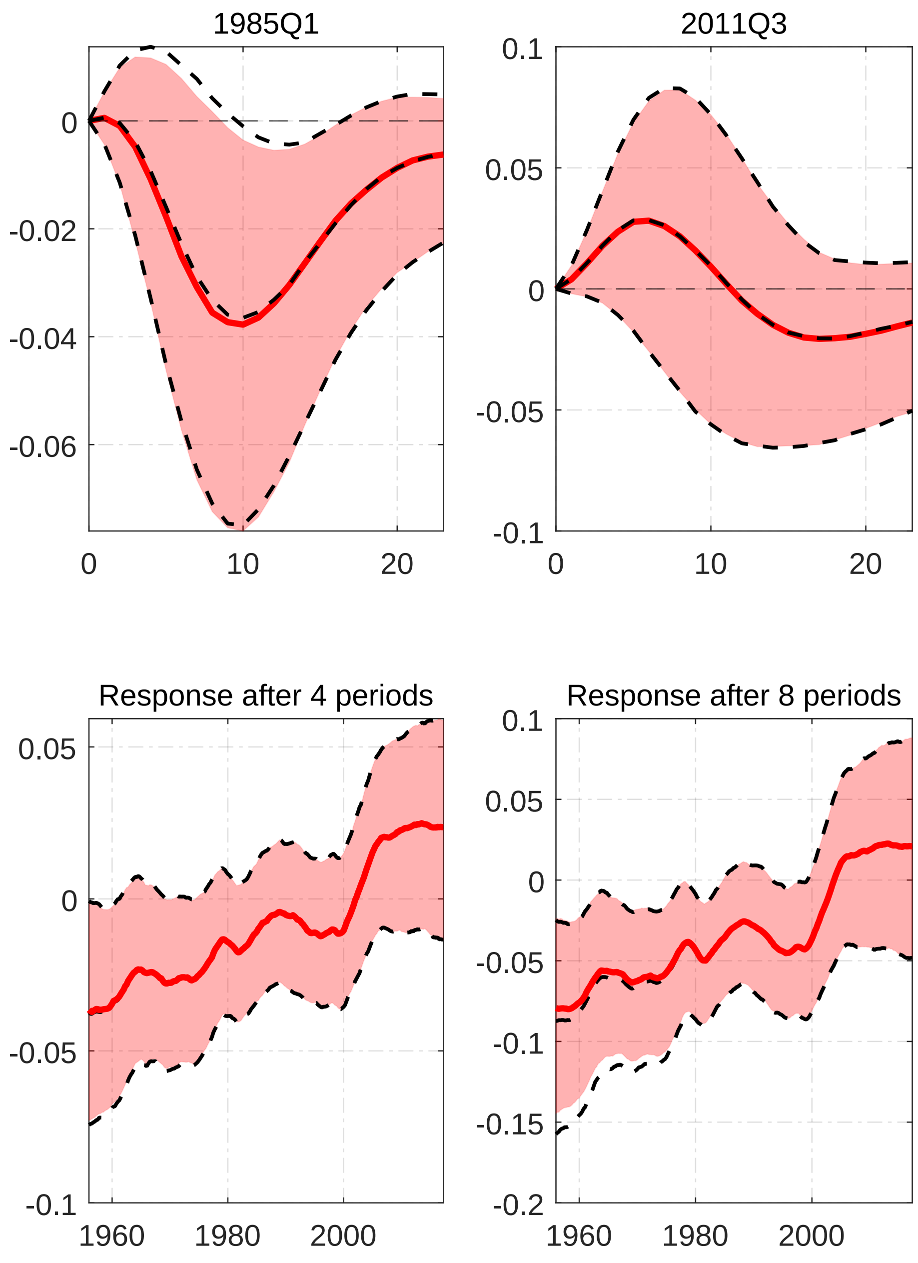

Notes: Mean responses from the baseline model (red-solid) and counterfactual responses (black-dashed) with 16^{th} and 84^{th} percentiles to a 25bp monetary policy shock.

VII Robustness

The following section examines the sensitivity of our results to a battery of robustness checks. These checks include the estimation of a rolling window VAR model, an alternative identification scheme, an analysis of the non-mortgage component of household debt, the role of the ZLB and a model with the ratio of mortgage debt to GDP.

A. *Results of a Rolling Window Estimation*

The results in the main section are based on a TVP-VAR model in which both the coefficients and the variance-covariance matrix are time-varying. This provides us with a set of estimated parameters for each individual period of our estimation sample. In this subsection, we report the results obtained from the estimation of a conventional fixed-parameter VAR over a rolling window.

We estimate the fixed-parameter VAR with the same variables and lag order of two over a fixed estimation window of 40 years, which we move in steps of one quarter over the sample period. After adjusting for lags as well as for the size of the estimation window, we thus obtain 86 sets of estimated parameters. As in the main part of the paper, we identify our shock assuming that the matrix of simultaneous relationships is lower triangular. In order to compare the results based on the same size of the shock, we again assume that an unexpected 25bp monetary policy shock hits the economy. The error bands for the impulse responses are obtained using a standard Monte Carlo simulation with 500 replications.

In order to save space, we only present the most important results here. The lower left panel in Figure (12) shows the response of mortgage debt four periods after the shock over time, and the lower right panel shows the same response eight periods after the shock over time. The left panel of the top row shows the impulse response of mortgage debt based on the sample spanning from 1957Q1 to 1996Q4, whereas the right panel shows the response of mortgage debt based on the sample running from 1977Q3 to 2017Q2.

Note that the results for individual periods are not exactly comparable to those of our TVP-VAR as the estimated parameters of the rolling window procedure are average effects over the selected estimation window. As in the main part of the paper, the sensitivity of mortgages to monetary policy shocks decreases significantly over time. That is, our results point to a drop in the responsiveness of mortgage debt to monetary policy. Furthermore, the results show the same pattern as regards the timing of the peak response, that is, the number of periods it takes for the monetary policy shock to reach its maximum impact.[21]

As an additional backup, we repeat the rolling window estimation based on an optimal lag length which we derive from standard information criteria. Based on the Bayesian information criterion, we choose seven lags as an alternative lag length.[22] Our qualitative results remain untouched, as can be seen in the appendix. Although

[21] The impulse responses for the other variables are available on request. However, it stands out that the price puzzle disappears as the impulse response of the inflation rate now shows the expected negative sign much earlier than before.

[22] We allow for a maximum lag length of eight lags.

Figure 12: Response of mortgage debt from a rolling window VAR

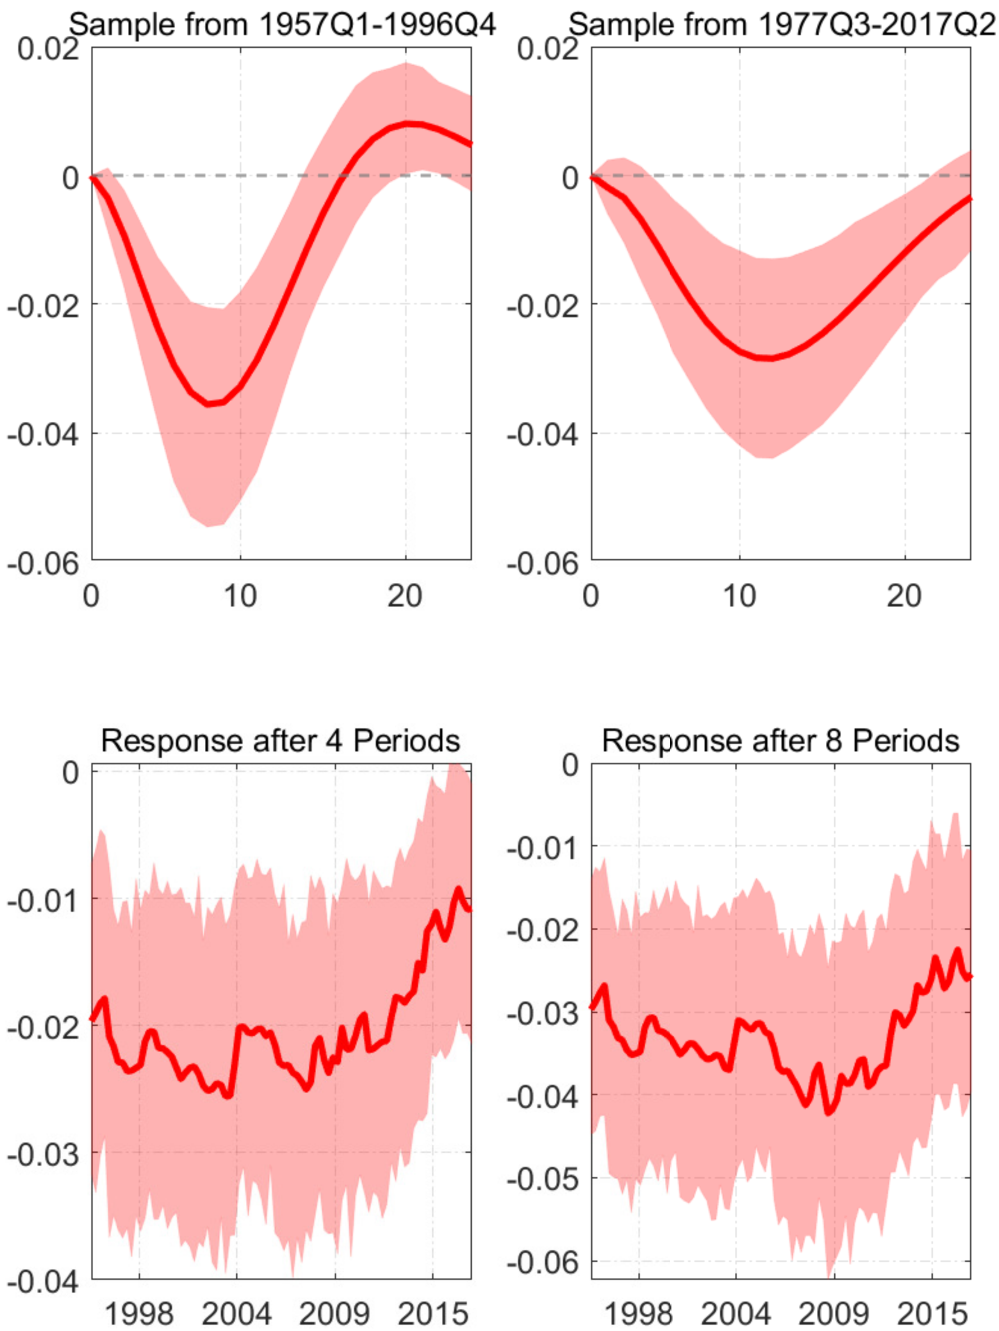

Notes: Mean responses (red-solid) with 16^{th} and 84^{th} percentiles to a 25bp monetary policy shock obtained from a rolling window VAR model with two lags.

the shape of the impulse responses is not as smooth as before, the trend of decreasing sensitivity of mortgages is more visible under four lags.

As our TVP-VAR runs into difficulties when we estimate the model in trend stationary time series, we further investigate whether our results change when our variables enter the VAR in (log) levels. We use the following variables: (1) real GDP (in logs), (2) the implicit price deflator (in logs), (3) real mortgage debt (in logs) and (4) the effective federal funds rate. We use two lags as suggested by the Bayesian information criterion and the Hannan-Quinn information criterion.

Figure (24) in the appendix summarizes our results for mortgage debt and shows impulse responses for the same samples as above as well as the time-varying effect four and eight quarters after the shock, respectively.[23] Here, too, we see that the sensitivity decreases over time.

Overall, the results from a rolling estimation with different specifications support our findings from the main part of the paper.

B. *Alternative Identifications*

Choosing an alternative Choleski ordering can, in principle, affect our results, as we alter the linear combinations of structural shocks which lie behind the reduced-form error terms. The question is whether mortgage debt still shows a negative reaction when we order the short-term interest before mortgage debt. In general, this ordering implies that mortgage debt does not belong to the information set of the central bank. At the same time, however, the contemporaneous reaction of mortgage debt is not restricted as it can react in the same period in which the monetary policy shock occurs.

Figure (13) reports the impulse responses for this ordering for the same selected dates as in Figure (2) as well as the time-varying response four and eight quarters, respectively. Overall, the alternative ordering of variables does not change our results too much. Even when the response of mortgage debt is unrestricted, in almost all periods we see that the response of mortgage debt shows the same negative sign as in the main part of the paper. In some periods, the contemporaneous response of mortgage debt is not different from zero, suggesting that we get the same qualitative response on impact. We see again that especially toward the end of our sample, the responsiveness of mortgage debt gradually decreases. This trend, however, is less pronounced than in the main part of the paper.

As an alternative identification scheme, we rely on restrictions on the signs of the impulse response functions (see, for instance, Uhlig, 2005; Arias et al., 2018). We estimate a constant parameter VAR with two lags for the full sample period, using the same variables as in the baseline model. The nature of the restrictions is consistent with a wide consensus in the literature and is backed by standard models of the business cycle. According to our partial identification scheme, a monetary policy shock raises

[23] The results are similar when we use the unemployment rate instead of (log) real GDP.

Figure 13: Response of mortgage debt for an alternative ordering

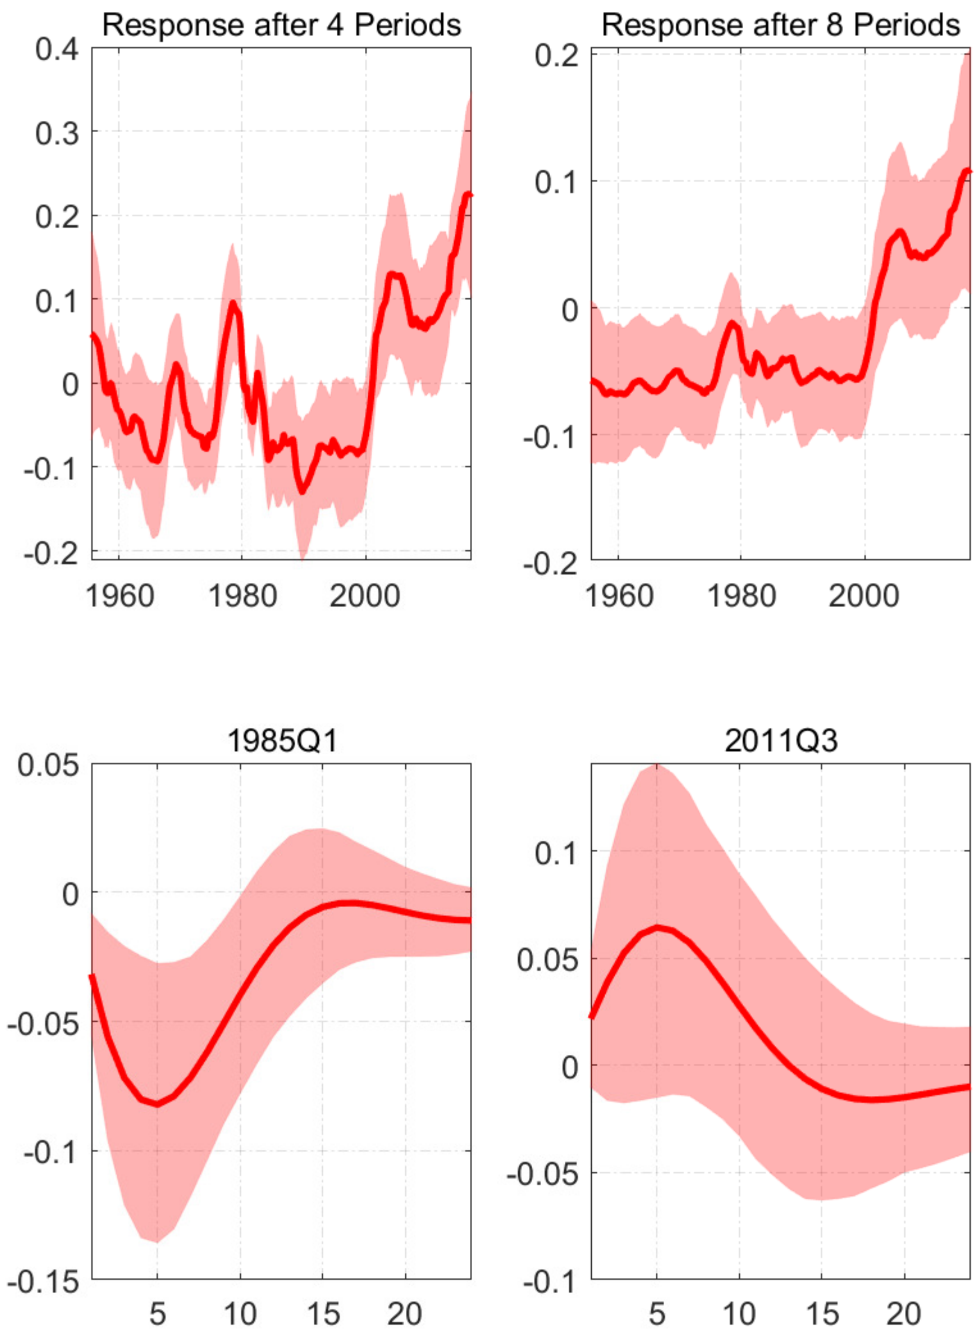

Notes: Mean responses (red-solid) with 16^{th} and 84^{th} percentiles to a 25bp monetary policy shock.

the short-rate, leads to a fall in inflation and an increase in unemployment. Finally, the response of mortgages is constrained to zero. All responses are constrained on impact only. These restrictions imply that we cannot observe a price puzzle as was found in the results from our baseline model.

Figure 14: Responses for an identification through sign-restrictions

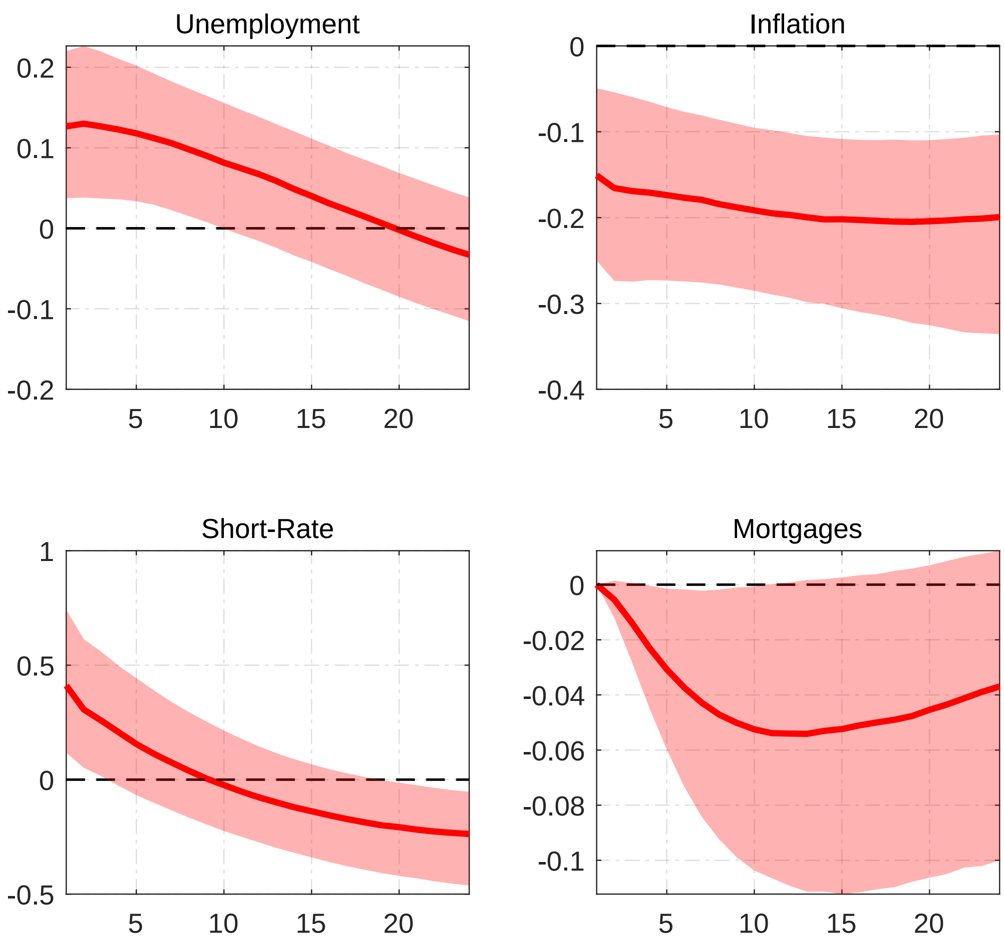

Notes: Mean responses (red-solid) with 16^{th} and 84^{th} percentiles to a 25bp monetary policy shock.

Figure (14) contains the impulse responses. While we cannot learn from this graph about potential time variation, it serves as a consistency-check for our baseline results. In particular, we find that mortgages react negatively already one period after the shock. Thus, the direction of the impulse response qualitatively corresponds to the results from the main section. Moreover, the magnitude of the responses roughly equals the sample means of the impulse responses shown earlier.

C. *Results for Non-Mortgage Debt*

We now ask whether the drop in the sensitivity to monetary policy does also hold for non-mortgage household debt. For that purpose, we estimate our baseline model and replace mortgage debt by Baxter and King (1999) filtered non-mortgage household debt.

Figure 15: Response of non-mortgage debt

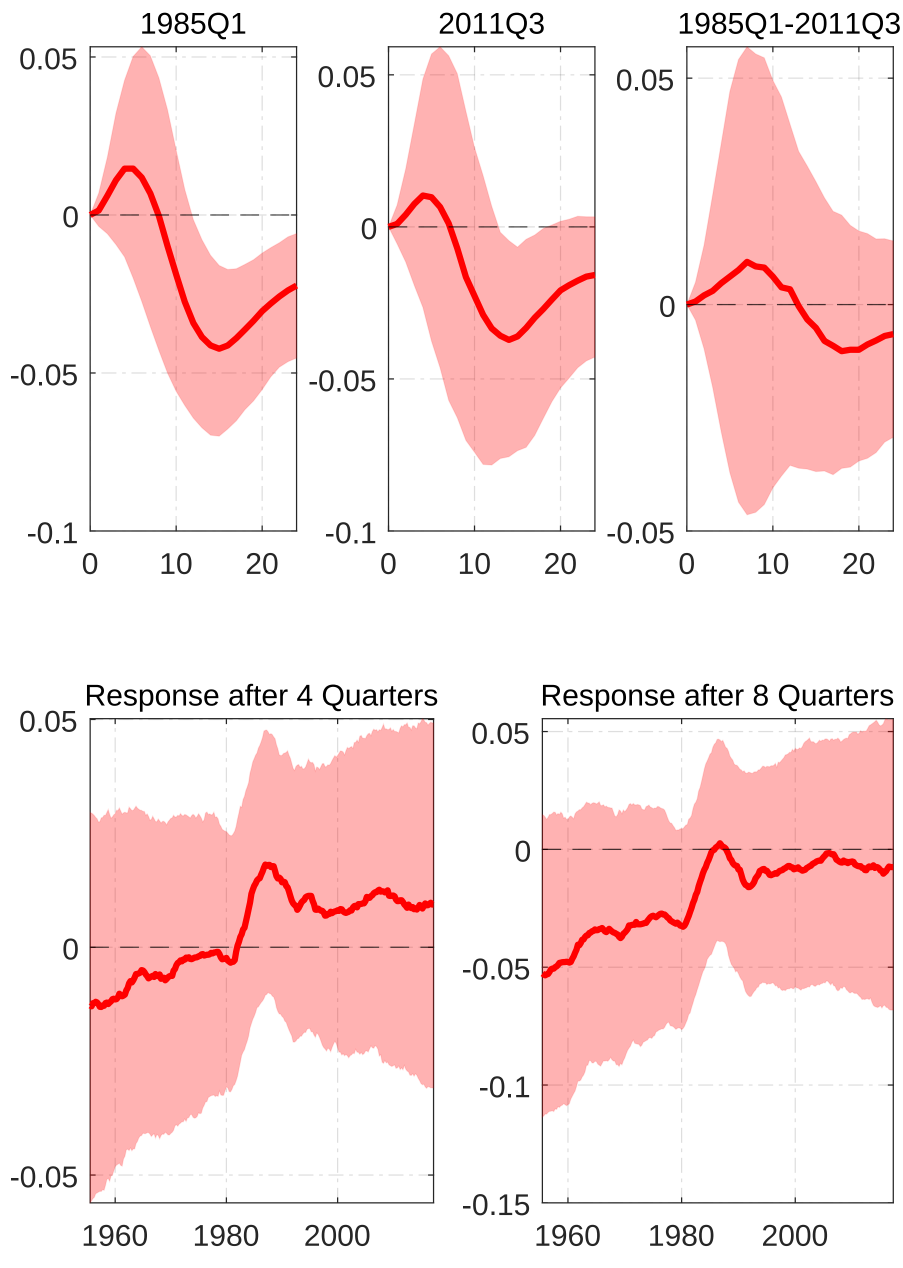

Notes: Mean responses (red-solid) with 16^{th} and 84^{th} percentiles to a 25bp monetary policy shock.

Figure (15) reports the resulting impulse responses. A policy tightening in 1985Q1 or 2011Q3 leads to a significant contraction of non-mortgage debt. Importantly, both impulse responses are statistically indistinguishable. Hence, we are unable to reject a stable transmission of monetary policy impulses to non-mortgage debt. In addition, the cross-sectional responses shown at the bottom of the figure suggest that the responses after four and eight quarters do not fluctuate over the sample period and remain insignificant. From that, we conclude that the reduced sensitivity of mortgages with respect to monetary policy is indeed a feature of the mortgage market that is not shared by other types of household debt.

D. *The Role of the ZLB*

Our sample also covers the period in which the US economy was at the ZLB on nominal interest rates. In the main parts of the paper, we show results for a VAR model that includes the federal funds rate for most of the sample period but uses a shadow short rate for the ZLB episode. We now drop the shadow short-rate and use the federal funds rate for the full estimation sample. Hence, our model includes a period when the short rate was constrained by zero.

The resulting impulse responses of mortgage debt are shown in Figure (16). We find a pattern that is very close to our baseline results. Therefore, we can conclude that the ZLB or the use of the shadow rate, respectively, does not affect our findings.

E. *A Model with Mortgage Debt to GDP*

The analysis thus far rests on the cyclical component of real mortgage debt, which we obtain from band-pass filtering the original time series. Alternatively, one could express real mortgage debt relative to real GDP. As our estimation framework relies on stationary time series, however, we need to include this series in first differences.

Figure (17) depicts the key results, that is, the responses of mortgage debt to GDP at two points in time as well as the evolution of the cross-sectional responses over time. The similarity with our baseline findings is striking: again, we find that mortgage debt fell stronger after a policy shock in 1985 compared to the same shock in 2011. Furthermore, the cross section of the impulse responses document a decline in the effectiveness of monetary policy which is qualitatively similar to the results from our baseline model.

VIII Conclusions

In this paper, we studied the role of monetary policy for the dynamics of US mortgage debt, the largest and most important component of overall household debt. In the aftermath of the recent financial crisis, which originated in the US housing market, the mortgage market received much attention.

The main tool of our analysis, a time-varying VAR model with stochastic volatility, allowed us to study the sensitivity of mortgage debt to monetary policy over time. We find that since the 1960s, the impact of monetary policy on mortgage debt has declined. A policy shock in 2011 has a much smaller effect on mortgage debt than an identically sized shock occurring in 1985. This finding, which is new to the literature,

Figure 16: Response of mortgage debt: the role of the ZLB

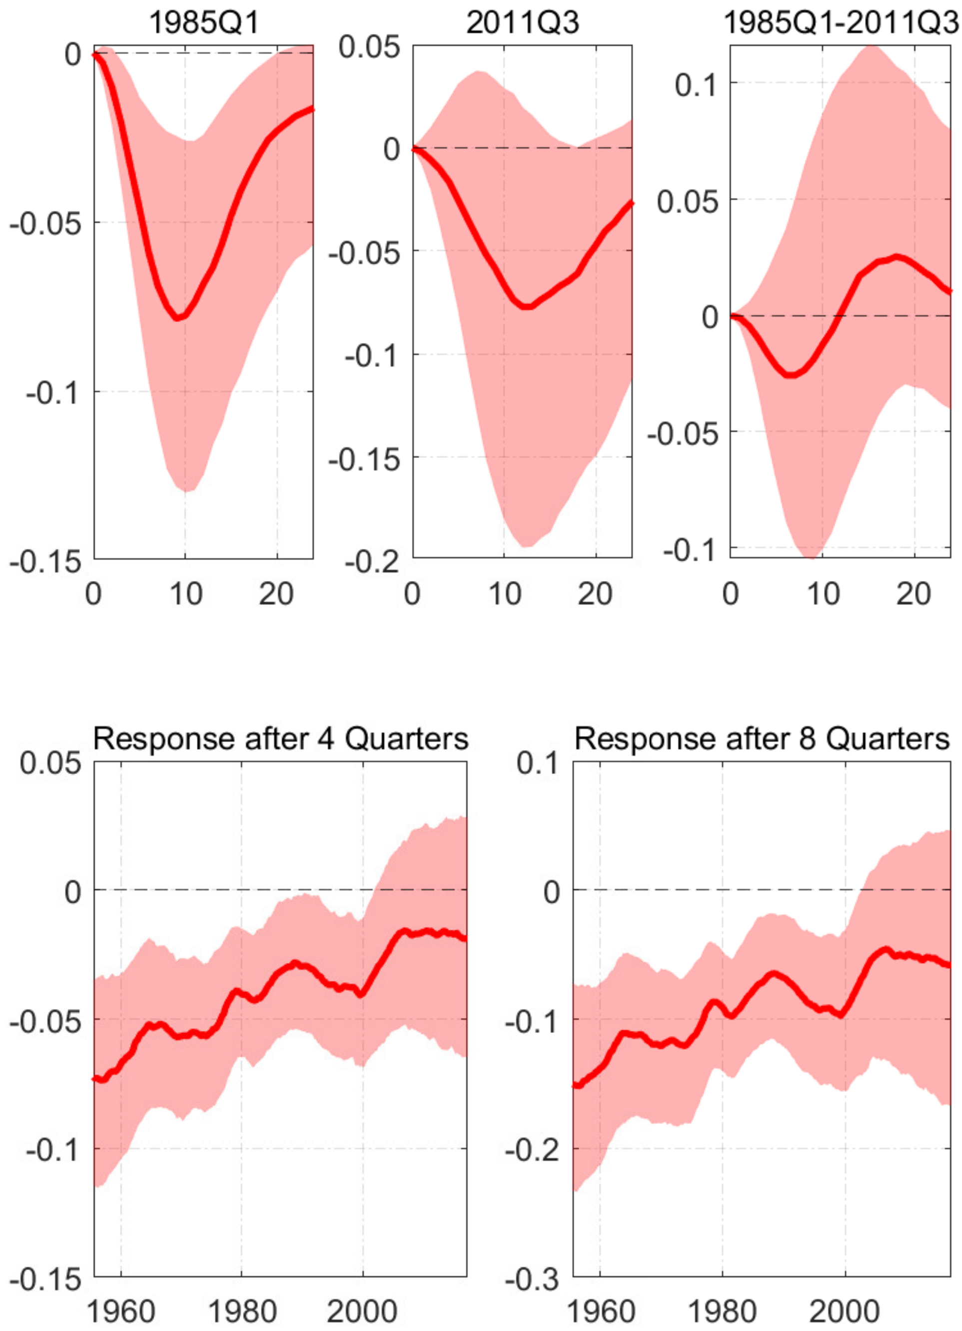

Notes: The left panel shows the mean responses (red-solid) with 16^{th} and 84^{th} percentiles to a 25bp monetary policy shock in 1985Q1, whereas the middle panel shows the mean responses to a 25bp monetary policy shock in 2011Q3. The right panel shows the difference between the responses in 1985Q1 and 2011Q3.

Figure 17: Response of mortgage debt relative to GDP

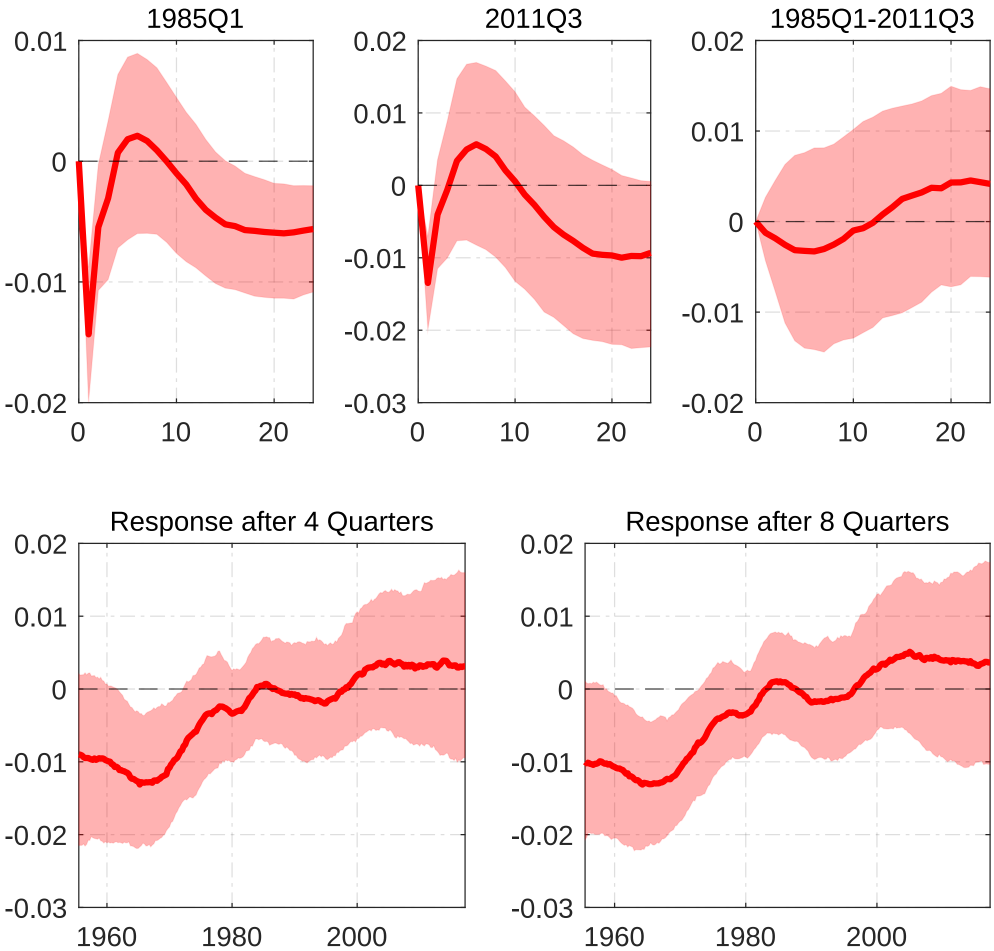

Notes: Mean responses (red-solid) with 16^{th} and 84^{th} percentiles to a 25bp monetary policy shock.

is robust to variations of the model and the estimation approach and is not observed for non-mortgage debt.

We also estimate a DSGE model for the U.S. economy in order to replicate our empirical findings. The ARM share, a key parameter in the determination of the model-based impulse responses, is shown to have declined strongly since the early 1980s. Once we calibrate the model to alternative realizations of the ARM share, we are able to replicate the decline in the response of debt to monetary policy quantitatively. To the extent the ARM share could be taken as given, this offers a consistent explanation for our results.

These findings have several implications for monetary policy and the mortgage market. First, our results suggest that, nowadays, monetary policy is a blunt and inefficient tool to engineer a reduction of household debt. The decline in the sensitivity to monetary policy implies that a large policy adjustment is needed in order to have a sizable effect on mortgage debt. This, however, would cause a deep recession. Hence, our results speak against using monetary policy as an instrument to prevent the build-up of household debt.

A second interpretation of our results addresses the role of the Fed in the run-up to the recent financial crisis. It is often claimed that the Fed contributed to inflating house prices by keeping the federal funds target rate too low for too long. Our results put this claim into perspective. If the sensitivity of mortgage debt to monetary policy in the mid-2000s is low, which is our main result, even persistently low levels of the federal funds rate should contribute little to the rise in mortgage debt before the crisis. Likewise, tightening monetary conditions, as the Fed did after June 2004, should translate into a small decrease in mortgage debt.

Our results fit the "mortgage rate conundrum" diagnosed by Justiniano et al. (2022). These authors argue that the empirical link between mortgage rates and longer-term interest rates broke. Hence, there seem to be strong structural changes in the mortgage market and its link to monetary policy.

References

Alpanda, Sami and Sarah Zubairy, "Addressing Household Indebtedness: Monetary, Fiscal or Macroprudential Policy?," *European Economic Review*, 2017, *92*, 47–73.

— **and** — , "Household Debt Overhang and Transmission of Monetary Policy," *Journal of Money, Credit and Banking*, 2019, *51* (5), 1265–1307.

Arias, Jonas E, Juan F Rubio-Ramírez, and Daniel F Waggoner, "Inference Based on Structural Vector Autoregressions Identified With Sign and Zero Restrictions: Theory and Applications," *Econometrica*, 2018, *86* (2), 685–720.

Baxter, Marianne and Robert G King, "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series," *Review of Economics and Statistics*, 1999, *81* (4), 575–593.

Boivin, Jean, "Has US Monetary Policy Changed? Evidence from Drifting Coefficients and Real-Time Data," *Journal of Money, Credit and Banking*, 2006, *38* (5), 1149–1174.

Calza, Alessandro, Tommaso Monacelli, and Livio Stracca, "Housing Finance and Monetary Policy," *Journal of the European Economic Association*, 2013, *11* (suppl_1), 101–122.

Canova, Fabio and Luca Gambetti, "Structural Changes in the US Economy: Is There a Role for Monetary Policy?," *Journal of Economic Dynamics and Control*, 2009, *33* (2), 477–490.

Christiano, Lawrence J, Martin Eichenbaum, and Charles L Evans, "Monetary Policy Shocks: What Have We Learned and to What End?," *Handbook of Macroeconomics*, 1999, *1*, 65–148.

Cogley, Timothy and Thomas J Sargent, "Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII US," *Review of Economic Dynamics*, 2005, *8* (2), 262–302.

Ehrmann, Michael and Michael Ziegelmeyer, "Mortgage Choice in the Euro Area: Macroeconomic Determinants and the Effect of Monetary Policy on Debt Burdens," *Journal of Money, Credit and Banking*, 2017, *49* (2-3), 469–494.

Eickmeier, Sandra and Boris Hofmann, "Monetary Policy, Housing Booms, and Financial (Im)Balances," *Macroeconomic Dynamics*, 2013, *17* (4), 830–860.

Garriga, Carlos, Finn E Kydland, and Roman Šustek, "Mortgages and Monetary Policy," *The Review of Financial Studies*, 2017, *30* (10), 3337–3375.

Geweke, John, "Evaluating the Accuracy of Sampling-Based Approaches to the Calculations of Posterior Moments," *Bayesian Statistics*, 1992, *4*, 641–649.

Greenspan, Alan and James E Kennedy, "Estimates of Home Mortgage Originations, Repayments, and Debt on One-to-Four-Family Residences," *FEDS Paper*, 2005, (2005-41).

— **and James Kennedy**, "Sources and Uses of Equity Extracted from Homes," *Oxford Review of Economic Policy*, 2008, *24* (1), 120–144.

Harding, Martin and Mathias Klein, "Monetary Policy and Household Deleveraging," *DIW Berlin Discussion Paper No. 1806*, 2019.

Huber, Florian and Maria Teresa Punzi, "International Housing Markets, Unconventional Monetary Policy, and the Zero Lower Bound," *Macroeconomic Dynamics*, 2020, *24* (4), 774–806.

Iacoviello, Matteo, "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle," *American Economic Review*, 2005, *95* (3), 739–764.

Jordà, Òscar, Moritz Schularick, and Alan M Taylor, "Betting the House," *Journal of International Economics*, 2015, *96*, S2–S18.

— , — , **and** — , "The Great Mortgaging: Housing Finance, Crises and Business Cycles," *Economic Policy*, 2016, *31* (85), 107–152.

Justiniano, Alejandro, Giorgio E Primiceri, and Andrea Tambalotti, "The Mortgage Rate Conundrum," *Journal of Political Economy*, 2022, *130* (1), 121–156.

Koop, Gary and Dimitris Korobilis, *Bayesian Multivariate Time Series Methods for Empirical Macroeconomics*, Vol. 3, Foundations and Trends in Econometrics, 2010.

Mian, Atif, Amir Sufi, and Emil Verner, "Household Debt and Business Cycles Worldwide," *The Quarterly Journal of Economics*, 2017, *132* (4), 1755–1817.

Mishkin, Frederic S, "Globalization, Macroeconomic Performance, and Monetary Policy," *Journal of Money, Credit and Banking*, 2009, *41*, 187–196.

Moench, Emanuel, James I Vickery, and David Aragon, "Why Is the Market Share of Adjustable-

Rate Mortgages So Low?," *Current Issues in Economics and Finance*, 2010, *16* (8).

Nakajima, Jouchi, "Time-Varying Parameter VAR Model with Stochastic Volatility: An Overview of Methodology and Empirical Applications," *Monetary and Economic Studies*, 2011, *29*, 107–142.

Owyang, Michael T and Tatevik Sekhposyan, "Okun's Law over the Business Cycle: Was the Great Recession All That Different?," *Federal Reserve Bank of St. Louis Review*, 2012, *94* (September/October 2012).

Paul, Pascal, "The Time-Varying Effect of Monetary Policy on Asset Prices," *Review of Economics and Statistics*, 2020, *102* (4), 690–704.

Primiceri, Giorgio E, "Time Varying Structural Vector Autoregressions and Monetary Policy," *The Review of Economic Studies*, 2005, *72* (3), 821–852.

Raftery, Adrian and Steven Lewis, "How Many Iterations in the Gibbs Sampler?," *Bayesian Statistics*, 1992, *4*, 765–776.

Sims, Christopher A and Tao Zha, "Were There Regime Switches in U.S. Monetary Policy?," *American Economic Review*, 2006, *96* (1), 54–81.

Uhlig, Harald, "What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure," *Journal of Monetary Economics*, 2005, *52* (2), 381–419.

Wooldridge, Jeffrey M, *Introductory Econometrics: A Modern Approach*, Nelson Education, 2016.

Wu, Jing Cynthia and Fan Dora Xia, "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound," *Journal of Money, Credit and Banking*, 2016, *48* (2-3), 253–291.

Zeev, Nadav Ben, "Household Debt, Adjustable-Rate Mortgages, and the Shock-Absorbing Capacity of Monetary Policy," *unpublished, Ben-Gurion University of the Negev*, 2016.

Appendix

A Data and Sources

Table (2) reports the four variables which are used in our baseline model. All data series are taken from the FRED data base of the Federal Reserve Bank of St. Louis.

Table 2: Data & description

unemployment	civilian unemployment rate
prices	implicit price deflator
short-rate	effective federal funds rate
mortgages	households and NPO; home mortgages

As the unemployment rate we take the civilian unemployment rate (UNRATE), i.e. the number of people 16 and older actively searching for a job as a percentage of the total labor force. We calculate the inflation rate as the year-on-year growth rate of the implicit price deflator (GDPDEF). As the policy rate, we choose the effective Federal funds rate (EFFR). For the time when the effective federal funds rate is at the Zero Lower Bound, we follow common practice and amend this series by the Wu and Xia (2016) shadow short rate. Lastly, mortgage debt is taken as the cyclical component of home mortgages to households and nonprofit organizations (HHMSDODNS). This is obtained by Baxter and King (1999) filtering the original series. Following the observation of Alpanda and Zubairy (2019) that debt cycles are about twice as long as the business cycle, we choose a band-length of eight and let frequencies between four and 64 quarters pass.

B Convergence Diagnostics

This appendix assesses the convergence of our MCMC algorithm in the baseline case presented in Section III in the main paper. We apply different metrics in order to judge how well our chain mixes. Remember that we use 20,000 iterations and discard the first 10,000. It stands out that choosing different burn-in periods delivers exactly the same results.

It is common practice to observe the inefficiency factors for convergence analysis. The inefficiency factor is the inverse of the relative numerical efficiency measure of Geweke (1992) and defined by $1 + 2\sum_{j=1}^{\infty} \rho_j$, where ρ_j is the autocorrelation of j^{th} order for the underlying parameter. Inefficiency factors of around 20 are regarded as satisfactory. Table (3) reports the distribution of inefficiency factors of our entire parameter space. Except for the hyperparameters, the inefficiency factors are on average far below 20. Not taking single outliers too serious as our parameter space is large, we conclude that our chain mixes quite fast.

Table 3: Distribution of inefficiency factors

	Mean	Median	Min	Max	10^{th} Percentile	90^{th} Percentile
$\mathbf{V}$	36.98	33.32	22.07	96.15	23.57	62.36
$\mathbf{B}$	4.05	3.04	1.29	15.88	2.06	7.14
$\mathbf{A}$	3.13	2.09	1.00	14.96	1.42	6.79
$\boldsymbol{\Sigma}$	4.46	3.05	1.05	34.41	1.61	8.20

Notes: The elements of **V** are the elements of the covariance matrix of the model's innovations. The elements of **B**, **A** and $\boldsymbol{\Sigma}$ are the time-varying coefficients, the time-varying volatilities of structural shocks and the time-varying simultaneous relations among our endogenous variables.

As a second check, we look at the Raftery and Lewis (1992) diagnostic. They use a two-state Markov chain assumption to construct a univariate diagnostic which is aimed to report the recommended number of iterations needed for a given level of precision in posterior samples.[24] Table (4) reports the corresponding distribution of the recommended minimum number of draws.

Overall, for both the hyperparameters in **V** as well as the parameters in **B**, **A** and $\boldsymbol{\Sigma}$, we have far more than the required number of draws, conditional on our desired accuracy goal.

Table 4: Distribution of Raftery-Lewis statistics

	Mean	Median	Min	Max	10^{th} Percentile	90^{th} Percentile
$\mathbf{V}$	1330	1115	813	4340	882	2353
$\mathbf{B}$	382	182	157	2340	161	1141
$\mathbf{A}$	245	164	153	1611	157	342
$\boldsymbol{\Sigma}$	457	182	157	2640	161	1141

Notes: The elements of **V** are the elements of the covariance matrix of the model's innovations. The elements of **B**, **A** and $\boldsymbol{\Sigma}$ are the time-varying coefficients, the time-varying volatilities of structural shocks and the time-varying simultaneous relations among our endogenous variables.

As a last check, we also apply the Geweke (1992) convergence diagnostic test, which can be sketched as follows: for each single parameter, the idea is to compare the mean of the first n_0 draws of the chain to the mean of the n_1 draws by dropping the corresponding draws in between.[25] It turns out that for 97.67% of the hyperparameters

[24] We follow the most common values (see, for instance, Raftery and Lewis, 1992; Primiceri, 2005) and set our quantiles of the marginal posteriors to 2.5% and 97.5%, the minimum probability needed to achieve our stationary posterior distribution of 95% and the desires accuracy of 2.5%, such that our reporting based on a 95% interval result in the true posterior values with a probability lying between 92.5%-97.5%.

[25] The statistics are calculated as $G = (\bar{x}_0 - \bar{x}_1)/\sqrt{\widehat{\sigma}_0^2/n_0 + \widehat{\sigma}_1^2/n_1}$, where $x_j = (1/n_j)\sum_{i=m_j}^{m_j+n_j} x^i$ with x^i being the ith draw. $\widehat{\sigma}_j^2/n_j$ is the standard error of $\bar{x}_j$ for $j = 0, 1$, where $\widehat{\sigma}_j^2$ is computed using a Parzen

in **V**, for 99.99% of the parameters in **B**, for 99.92% of the parameters in **A** and for 100% of the parameters **Σ**, the chains converges to the target distribution within the first 2000 draws, suggesting that the influence of the priors decays fast.

To sum up, the convergence diagnostics seem satisfactory and justify our prior choice, considering the large parameter space of our model.

C Details on the DSGE Model

We use the DSGE model from Alpanda and Zubairy (2017). For our purposes, the key element of the model is the description of the mortgage market, which we present here in order to facilitate the interpretation of our results. The rest of the model can best be found in Alpanda and Zubairy (2017).

The law of motion of the stock of mortgage debt is

$$\frac{D_t}{P_t} = (1-\kappa)\frac{D_{t-1}}{P_t} + \frac{L_t}{P_t}. \tag{13}$$

Here, D_t is the stock of nominal mortgage debt, P_t is the price level, L_t are new mortgages and κ is the amortization rate, which is assumed to be constant.

The fixed interest rate on new mortgages is R_t^F. Each period, a fraction $\mathbf{\Phi}$ of the stock of mortgage debt is refinanced at the rate of new mortgages, R_t^F. The effective interest rate on the stock of mortgages R_t^M is the weighted average of the previous effective rate, R_{t-1}^M, and the rate on new mortgages. The weights are given by the share of old loans which are not refinanced and the sum of new loans and refinanced existing loans, respectively. Hence, the expression for the effective rate is

$$R_t^M = \underbrace{(1-\mathbf{\Phi})\left(1-\frac{L_t}{D_t}\right)}_{\text{share of old loans not refinanced}} R_{t-1}^M + \left[\underbrace{\frac{L_t}{D_t}}_{\text{new loans}} + \underbrace{\mathbf{\Phi}\left(1-\frac{L_t}{D_t}\right)}_{\text{share of old loans refinanced}}\right] R_t^F$$

The calibration of $\mathbf{\Phi}$ proposed by Alpanda and Zubairy (2017) contains information on the ARM share, among other things. In short, the ratio of repayments resulting from refinance originations to the stock of mortgages in Alpanda and Zubairy (2017) relies on data in Greenspan and Kennedy (2005) and Greenspan and Kennedy (2008) and averages around 4.35% quarterly, which corresponds to a duration of 31.16 quarters (7.79 years) using the half life formula.[26] Alpanda and Zubairy (2017) then

window. We follow common practice and choose n_0 as the first 10% and n_1 as the last 50% of our Markov chain. Note, however, that G is of course below 0.05 if the whole chain is stationary, saying that the means of the first n_0 and the last n_1 values are quite similar.

[26] That is, the duration is calculated as $31.16 = 2 \times \frac{\ln 0.5}{\ln(1-0.0435)}$.

calculate an average duration assuming a 10% share for adjustable rate mortgages:

$$\Phi = 1 - \exp\left(\frac{\ln 0.25}{4 \times (0.9 \times 7.79 + 0.1 \times 1)}\right).$$

This results in a duration of 28.43 quarters (7.11 years) which is converted back to a $\Phi = 0.0475$ used in the paper:[27]

In our simulation exercise, we assume different ARM (FRM) shares ranging from 5% to 80% and calibrate the ratio of repayments resulting from refinance originations to the stock of mortgages Φ according to

$$\Phi = 1 - \exp\left(\frac{\ln 0.25}{4 \times (\text{FRM share} \times 7.79 + \text{ARM share} \times 1)}\right)$$

For each Φ under consideration, we then re-simulate the model, keeping the entire calibration of all remaining parameters unchanged, i.e. identical to the calibration as in Alpanda and Zubairy (2017).

Note that for a larger ARM share, Φ is closer to one, such that there is a closer connection between the effective interest rate and the rate on new mortgages.

D Explanatory Power

Table (5) compares the empirical fit of the TVP-VAR model and the constant parameter VAR model, respectively. We use the R^2 as a metric of explanatory power. In both cases, i.e. for the restricted and the unrestricted short-rate equation, the TVP-VAR has a superior empirical fit for unemployment, inflation and mortgages. Only for the short-rate, the fit of the constant parameter VAR is superior.

Table 5: Comparison of R^2

	Restricted Policy Rule			
	Unemployment	Inflation	Mortgages	Short-Rate
TVP-VAR	0.962	0.984	0.978	0.900
Constant Parameter VAR	0.953	0.977	0.950	0.911
	Unrestricted Policy Rule			
	Unemployment	Inflation	Mortgages	Short-Rate
TVP-VAR	0.962	0.984	0.978	0.906
Constant Parameter VAR	0.953	0.977	0.950	0.913

Notes: The upper part of the table refers to the model where the coefficients of the short-rate equation to changes in mortgage debt are restricted, whereas the bottom part of the table refers to a specification where all parameters in the short-rate equation remain unrestricted.

[27] We are thankful to Sarah Zubairy for insightful comments and for providing us data and details on the calibration of Φ.

E Additional Results on Counterfactual Simulations

Figure 18: Counterfactual historical simulation

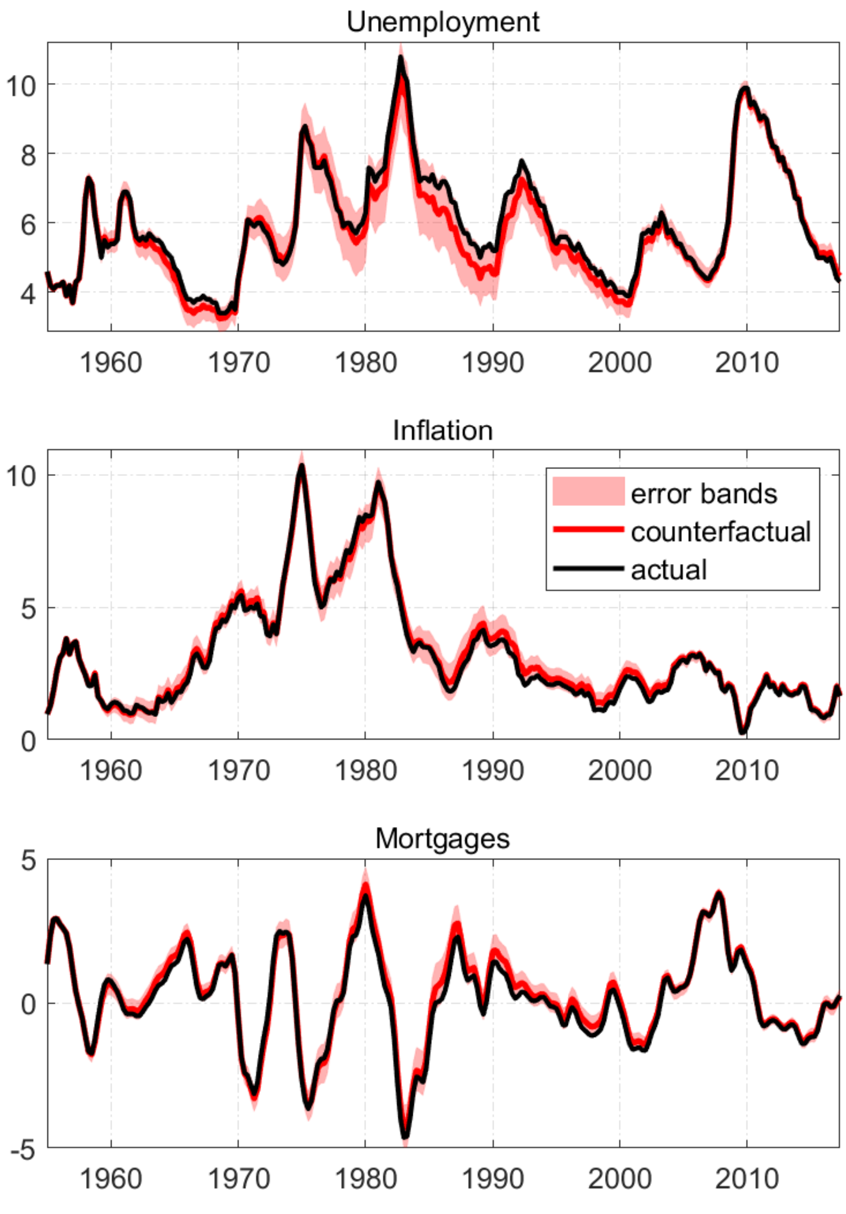

Notes: Counterfactual historical simulation for unemployment, inflation and mortgage debt. The parameters for the short-rate equation (monetary policy rule) are drawn from the posterior mean distribution of the Bernanke tenure.

F Additional Results for TVP-VAR Model

Figures (19) and (20) report the full three-dimensional profiles of the time-varying impulse responses.

Figure (21) shows the posterior mean as well as the 16^{th} and 84^{th} percentiles of the standard deviation of the structural shocks.[28] Remember that stochastic volatility is meant to capture possible heteroscedasticity in shocks. In this sense, the conditional variance of the shocks is a possible driver of time variation in the linear structure of our model. Overall, our results indicate time variation in the volatility of shocks. Thus, some of the variation in the model's linear structure comes from the variance-covariance matrix in addition to the VAR coefficients. Besides the high volatility of structural shocks from the short rate equation around 1980, it is particularly noticeable that the volatility of structural shocks from the inflation equation and from the mortgages equation is much more persistent.

To investigate whether the time variation in our paper is a feature of the data or driven by our priors, we follow the formal test proposed by Cogley and Sargent (2005) and compare the trace of the prior for the variance-covariance matrix of $\mathbf{\Sigma}_\beta$ with the trace of the posterior distribution of $\mathbf{\Sigma}_\beta$. Following Cogley and Sargent (2005), we would point to time variation in the parameters stemming from the data rather than the priors if the trace of the prior is below the 16^{th} percentile of the trace of posterior distribution.

Table 6: Trace test

16^{th}	50^{th}	84^{th}	Prior
0.053	0.054	0.056	0.034

Notes: The first three columns report the 16^{th}, the 50^{th} and the 84^{th} percentiles of the trace of the posterior of $\mathbf{\Sigma}_\beta$, respectively. The fourth column reports the trace of the prior for $\mathbf{\Sigma}_\beta$.

These results suggest that the data rather than the priors are the source of the time variation in the coefficients.

[28]Note that our MCMC algorithm obtains draws from $\mathbf{h}_t = [h_{1,t}, \cdots, h_{k,t}]'$ with $h_{i,t} = \log \sigma^2_{i,t}$ for $i = 1, \cdots, k$. Therefore, the volatility of structural shocks is obtained by calculating $\sigma_{i,t} = \sqrt{\exp(h_{i,t})}$.

Figure 19: Mean responses to a monetary policy shock

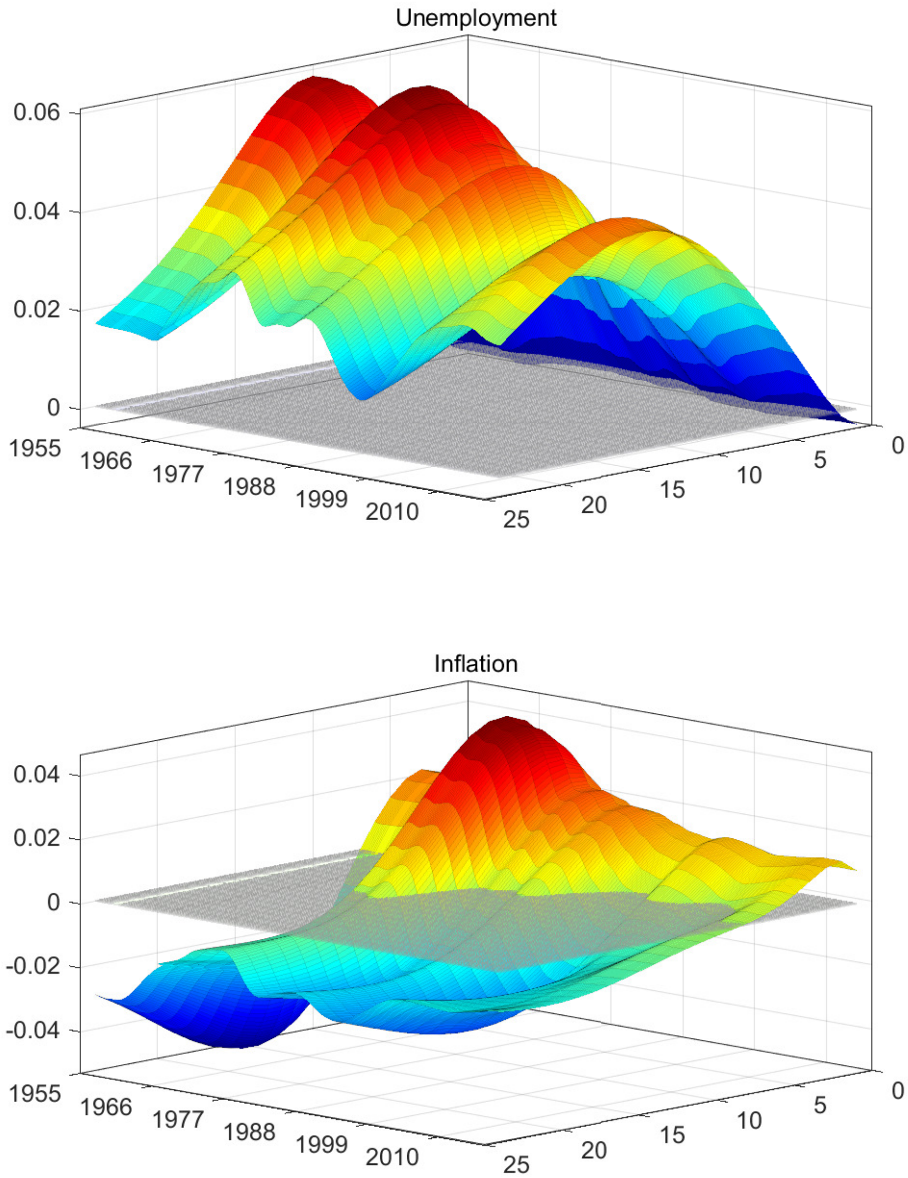

Notes: Mean responses to a 25bp monetary policy shock.

Figure 20: Mean responses to a monetary policy shock

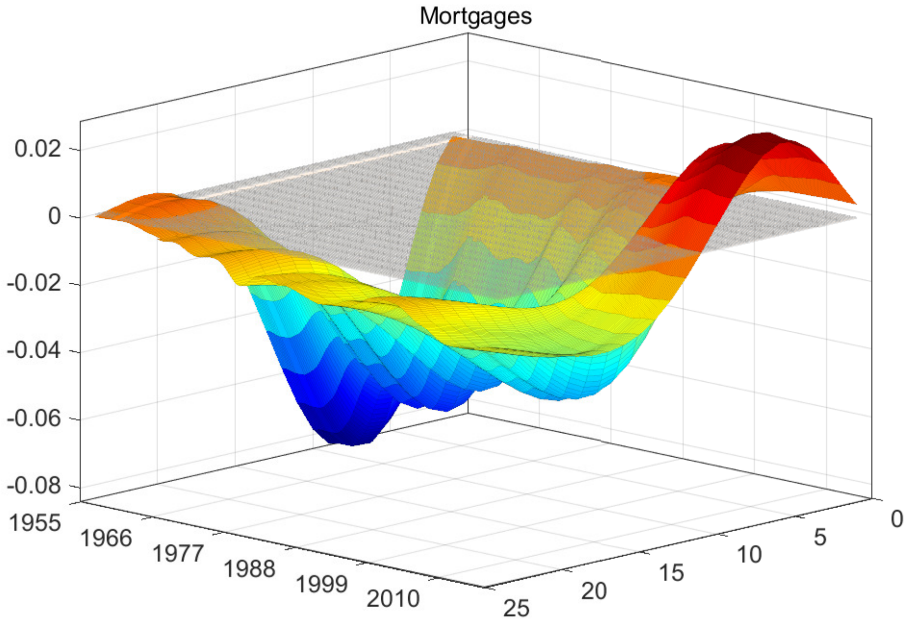

Notes: Mean responses to a 25bp monetary policy shock.

Figure 21: Stochastic volatility

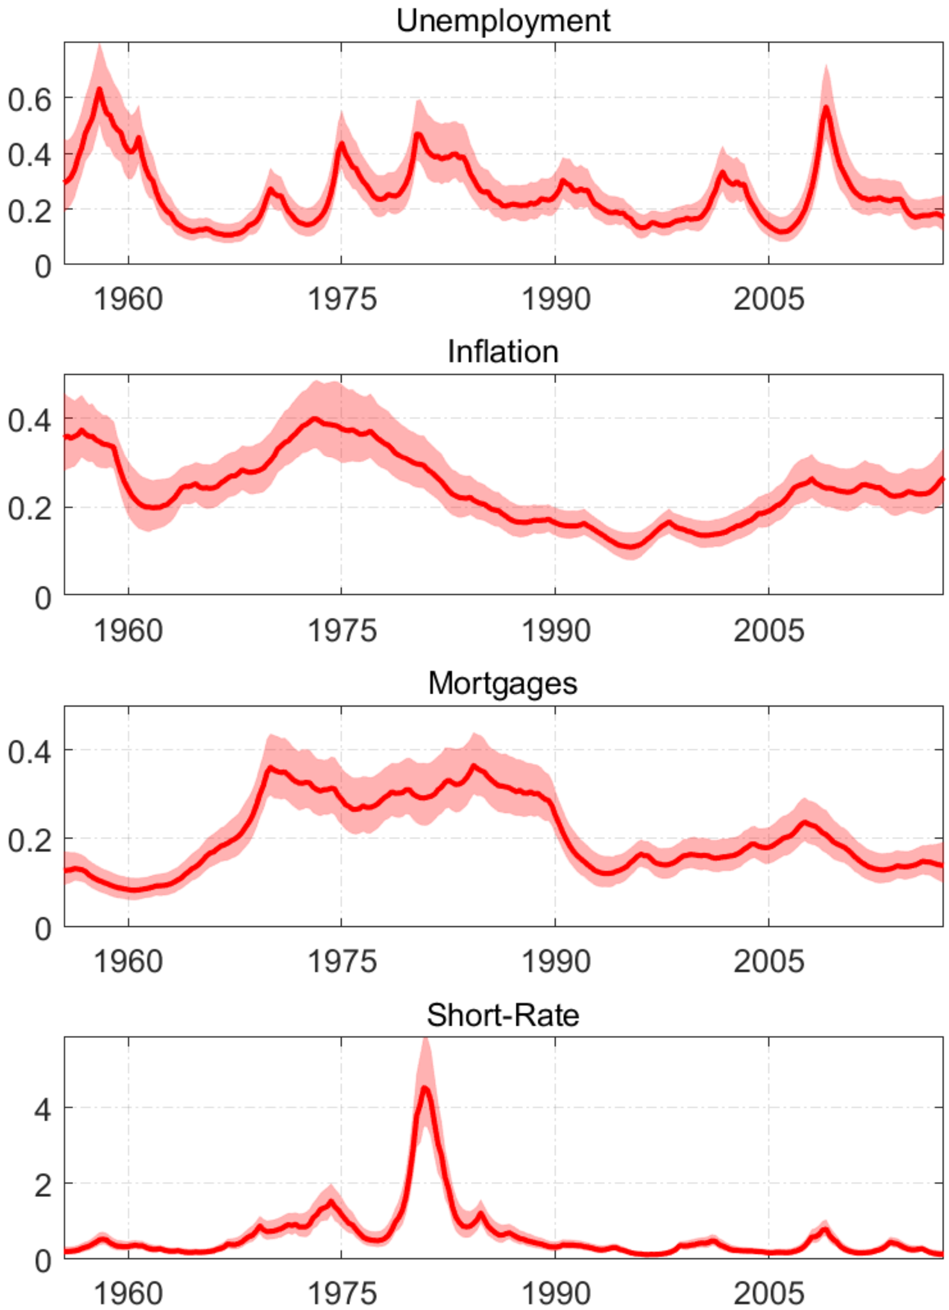

Notes: Posterior mean and 16^{th} and 84^{th} percentiles of the standard deviation of structural shocks.

G Additional Results for Rolling Window VAR Model

The main text reports the results from a rolling window VAR with two lags. Here, we show the results from a rolling window VAR with the optimally chosen lag order.

For each estimation within our rolling window procedure, we calculate the log likelihood as

$$\text{loglik} = -\frac{kT}{2}(1 + \ln 2\pi) - \frac{T}{2} \ln \det \widehat{\Sigma},$$

where k is the number of endogenous variables, T is the sample size and $\det \widehat{\Sigma}$ is the determinant of the estimated variance-covariance matrix. Note that also for the TVP-VAR, the marginal (or log) likelihood would in general be a natural candidate to determine the optimal lag length. However, several articles criticized that the harmonic mean method (as was typically used in the literature for TVP-VARs) could be biased.[29]

Figure 22: Optimal lag order for the rolling window VAR Model

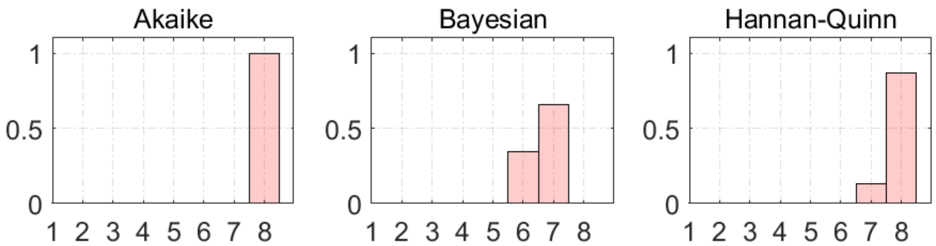

Notes: Distribution of preferred lag length by different information criteria.

We use the log likelihood to calculate the values of standard information criteria. Figure (22) shows the share of each lag length as preferred by the different information criteria. As the Bayesian Information Criterion penalizes the complexity of the model the most among all information criteria, we choose seven lags as an alternative lag order.

For the alternative lag order, our results remain qualitatively unchanged, see Figure (23). If anything, the shape of the impulse responses is not as smooth as with a lag length of two.

[29]We are grateful to Jouchi Nakajima and thank him for several helpful and insightful comments on this topic.

Figure 23: Response of mortgage debt from a rolling window VAR

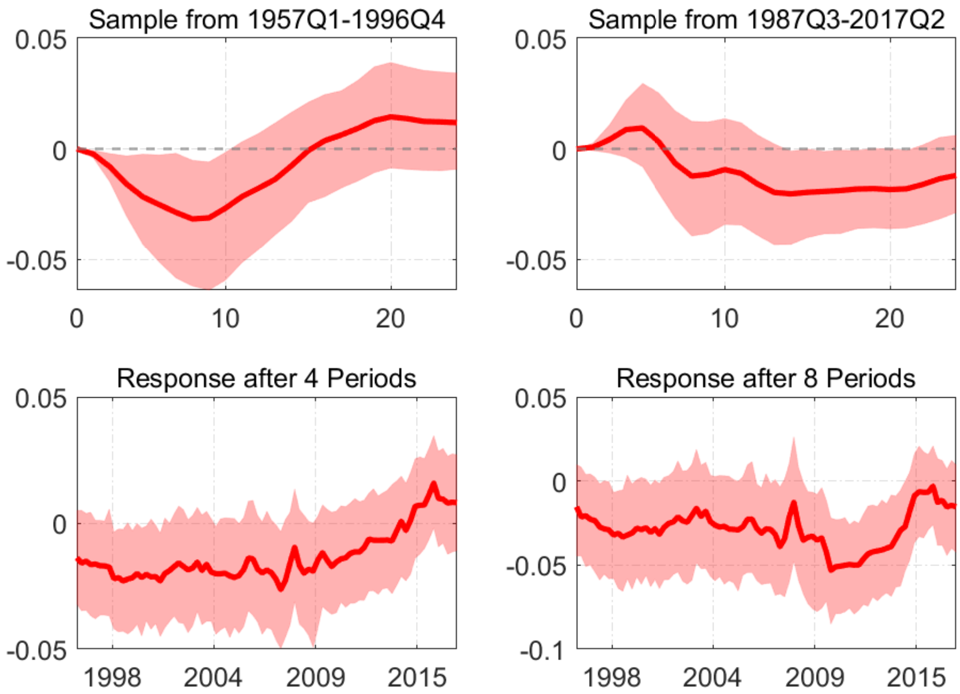

Notes: Mean responses (red-solid) with 16^{th} and 84^{th} percentiles to a 25bp monetary policy shock obtained from a rolling window VAR model with four lags.

Figure 24: Response of real mortgage debt from a rolling window VAR in (log) Levels

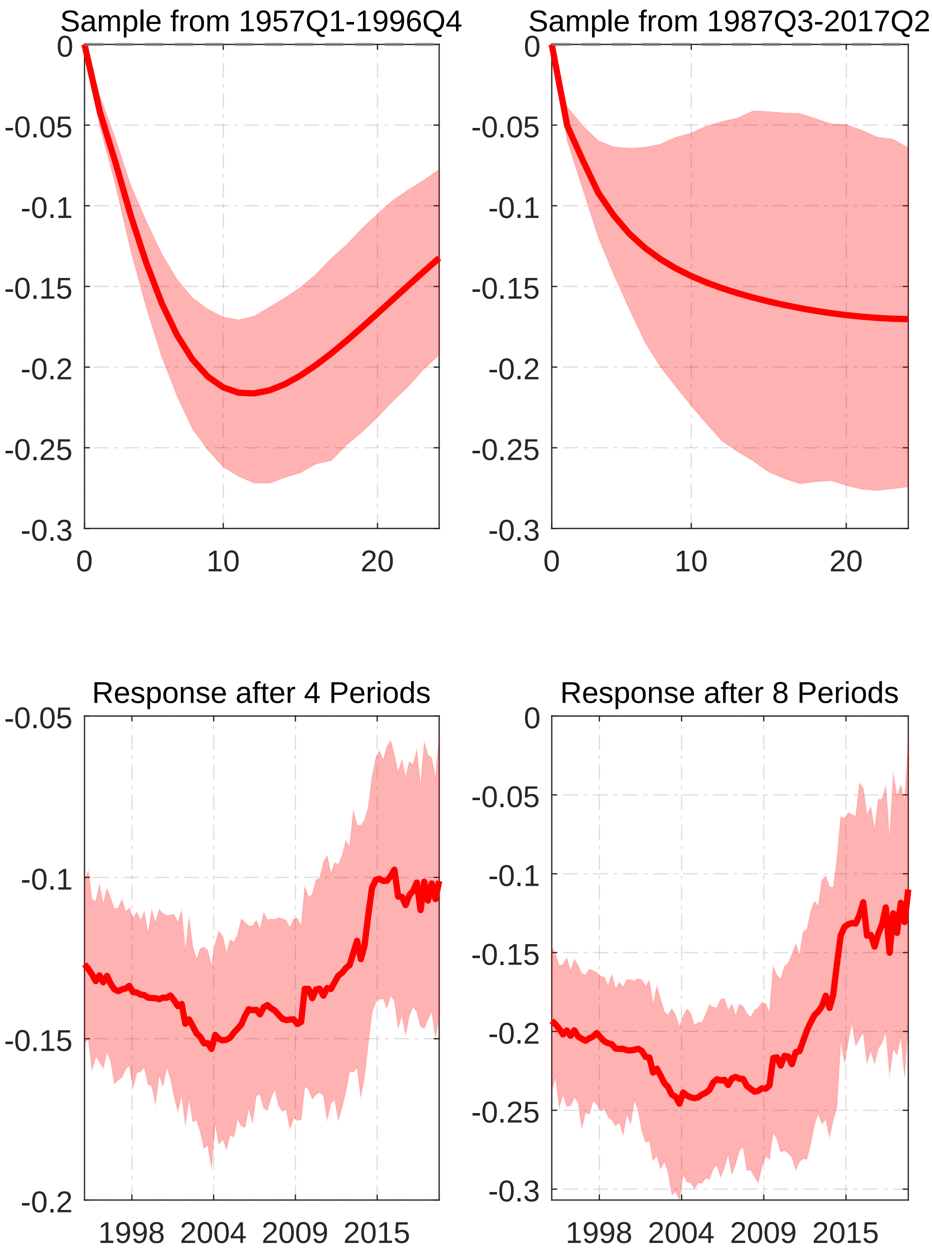

Notes: Mean responses (red-solid) with 16^{th} and 84^{th} percentiles to a 25bp monetary policy shock obtained from a rolling window VAR model in (log) levels.

Essay III:

Do Credit Supply Shocks Have Asymmetric Effects?

This paper is published as:

Finck, David and Paul Rudel, "Do Credit Supply Shocks Have Asymmetric Effects?," *Empirical Economics*, *forthcoming*, 2022.

Acknowledgement

We thank an anonymous referee and Peter Tillmann for very helpful comments and suggestions. Vanessa Thelen, Katrin Luise Gärtner and Karsten Kucharczyk provided insightful comments. We also thank Egon Zakrajšek and Luca Gambetti for providing their data. We thank the Fritz Thyssen Stiftung for financial support (grant no. 10.16.2.004WW).

Do Credit Supply Shocks Have Asymmetric Effects?

David Finck[*] Paul Rudel[†]

Abstract

They do. Partly. We identify credit supply shocks via sign restrictions in a Bayesian VAR and separate them into positive and negative. Using local projections, we find that positive credit supply shocks leave notably different prints in private debt, mortgage debt, and debt-to-GDP, as opposed to negative credit supply shocks. This pattern is caused by the response of household mortgage debt. Furthermore, we find evidence that positive credit supply shocks are the driving force behind boom-bust cycles. Yet, developments behind the boom-bust cycle cannot explain the strong and persistent response in debt; but house prices tend to. However, if we abstract from potential asymmetries, we get rather mild results, which underestimate the true effects of credit supply shocks.

Keywords: Credit Supply Shocks, Household Debt, Asymmetry, Local Projections

JEL classification: C11,E21, E22, E32

[*]University of Giessen, email: david.finck@wirtschaft.uni-giessen.de
[†]University of Giessen, email: paul.rudel@wirtschaft.uni-giessen.de

I Introduction

The financial crisis and the subsequent Great Recession made sternly clear that credit markets can take a leading role for economic activity. As such, there is a renewed interest in the nexus between credit supply shocks and economic activity.

One common finding in the literature is that an unexpected contraction in credit conditions has adverse effects. For example, Gambetti and Musso (2017) estimate a time-varying vector autoregressive model (VAR) model with drifting parameters and stochastic volatility for the US, the UK, and the euro area for 1980 to 2011. They find significant effects of credit supply shocks which increase over time. Their results also imply that the effects of credit market distortions are stronger during recessions. In another seminal paper, Gilchrist and Zakrajšek (2012) show that lower credit spreads improve the costs of debt finance, which increases spending and production and, in turn, increases asset prices and thus stimulates economic activity through the financial accelerator mechanism (for further recent contributions, see Jordà et al., 2013; Mumtaz et al., 2018; Mian et al., 2017; and Gertler and Gilchrist, 2018, among others). However, these analyses have in common that they examine symmetric responses to distortions in credit supply, i.e. the impact of a positive shock is identical to the impact of a negative shock in absolute terms.

The contribution of this paper is to examine asymmetries (and potentially nonlinearities) in the propagation of credit supply shocks. The objective here is to provide new evidence on the transmission of credit supply shocks by relying on a simple, yet flexible framework. We follow Tenreyro and Thwaites (2016), who investigate asymmetric effects of monetary policy shocks on key macroeconomic variables via local projections. The advantages of local projections over vector autoregression (VAR) models are well documented (Jordà, 2005). Above all, it is not necessary to impose dynamic restrictions as is done in VAR models. Moreover, local projections allow us to parsimoniously test for asymmetric effects. We do so by splitting an identified credit supply shock à la Gambetti and Musso (2017) into its positive and negative parts which are then planted into a set of seemingly unrelated equations. This allows us to directly evaluate the effect of a credit supply shock for different adjacent horizons.[1] Our set of variables includes key variables for the US on the overall debt cycle of households, but also various variables describing real activity as well as the demand and supply side of the economy ranging from the early 1970s until late 2018.

Overall, our results clearly point to asymmetric effects concerning the different debt measures considered in this paper. More specifically, we find that the response of overall household indebtedness to an unexpected increase in credit supply is substan-

[1]Such a credit supply shock can be associated with various events, such as unexpected changes in bank capital availability for loans due to changes in regulatory capital ratio requirements or unanticipated changes in the degree of competition in the banking sector. More precisely, an exogenous drop in credit supply is assumed to lead to an increase in the lending rate, which ultimately leads to a drop in economic activity and deflationary pressure.

tially stronger in absolute terms than in the opposite case of an unexpected decrease. That is, we find that an expansionary shock of one standard deviation leads to a significant increase of household indebtedness by about 1.5 percent after roughly three years. However, a negative shock of the same size leads to a decrease of household indebtedness which is significant only for the first year and also significantly different from the response following a credit supply easing. Importantly, this result is shown to clearly be driven by the responsiveness of households' mortgages. Also, we find key macro headline variables (prices, production, short-term interest rate) as well as some key variables determining the demand and supply side of the economy to respond asymmetrically. Overall, following an expansionary shock of credit supply, our results clearly point to the well established boom-bust cycle. In the opposite case we do not observe such a pattern.

While our framework allows us to flexibly uncover asymmetric responses following shocks, tracking the exact mechanisms that drive our results is beyond the scope of this model and, thus, of this paper.

Nevertheless, a large body of literature points at amplifications and asymmetries in the propagation of sudden (financial) distortions as, for example, (i) occasionally binding borrowing constraints, (ii) market imperfections (in terms of asymmetric information), or (iii) the role of asymmetric central bank behavior.

Regarding occasionally binding constraints, suppose an unexpected easing of credit conditions. In this case, both, firms and households that previously were excluded from the credit market are now able to borrow. The resulting increase in demand will in turn stimulate the economy. Following an unexpected deterioration in credit conditions, in contrast, the credit constraint could eventually become binding for both, households and firms. In this case, firms and households alike could either (i) no longer being able to rely on financial intermediaries to borrow externally or (ii) not being able to borrow the amount demanded.[2]

Another important mechanism relies on the role of asymmetric information. In general, the external finance premium is the premium that banks charge due to asymmetric information regarding a project to be financed. As such, the balance sheet or net worth is of particular importance for the financing decision. Fluctuations in the net worth could thus increase the effect of shocks hitting the economy. Consider again a negative credit supply shock which reduces the availability of credit in the economy. In this case, the role of the balance sheet with respect to the credit conditions of firms will become more important than before. A positive credit supply shock, in contrast, increases the availability of credit. A firm with a given balance sheet may therefore

[2]In this respect, as pointed out by Sedláček and Sterk (2017) and emphasized by Barnichon et al. (2022), episodes of high financial stress can prevent high-potential firms to emerge. More precisely, Sedláček and Sterk (2017) show that cohorts of large firms tend to be born during periods of booming consumer demand, i.e., when it is easy for firms to acquire new customers. Phases of booming consumer demand, in turn, may depend on credit conditions.

find it easier to borrow money, but disproportionately more difficult to borrow money in the case of a negative shock.

Finally, the central bank may play an important role in the transmission of credit supply shocks. Central banks monitor the lending behavior to firms and households very closely. Eventually, the central bank has asymmetric preferences when it comes to stabilizing shocks stemming from the credit market. For example, it is possible that the central bank reacts stronger during boom phases than during bust phases. Also, the central bank eventually reacts different when overall credit conditions are loose already. Finally, the reaction of the central bank may also depend on the zero lower bound. For example, if the key interest rate is close to zero, the central bank has less scope to counter inflationary pressures and may have to resort to unconventional monetary policy measures. These effects on the systematic component of monetary policy can all contribute to making the transmission of credit supply shocks non-linear or asymmetric.[3]

Overall, our paper fits well into a relatively new strand of literature which investigates non-linearities and asymmetries in the transmission of credit supply shocks. The first paper to mention is of Colombo and Paccagnini (2020), who estimate a smooth transition VAR (STVAR) and investigate the role played by credit supply shocks across the business cycle. They find that contractionary credit supply shocks trigger asymmetric and negative effects. However, the paper that fits closest is Barnichon et al. (2022), who estimate a vector moving-average (VMA) with functional approximations of impulse responses. They find that the effects of financial shocks on the economy depend on their sign and size. That is, their results imply that the mild and short-lived effects of financial market disruptions typically found in SVAR models can be explained by asymmetric effects. However, we do not solely focus on key macro variables, but also investigate whether asymmetric effects can be found in credit volumes as well as different measures of both, aggregate supply and aggregate demand. Finally, investigating asymmetries helps not only to get a better understanding of the consequences of credit supply shocks but rather is crucial for policymaking. This is because the events leading to and following the latest financial crisis raised the question whether monetary policy should take credit developments into account when making policy decisions. For this reason, it is not surprising that many economists propose that the central bank should eventually lean against the credit cycle. If this is the case, it is crucial for policy design to know whether credit shocks propagate symmetric or asymmetric before adjusting the short-term nominal interest rate.

The remainder of this paper is organized as follows. Section II explains the methodology. Section III contains our main results. In Section IV, we conduct a battery of robustness checks. Section V concludes.

[3] As will be shown in the robustness section, we find that the central bank's systematic response to credit volumes and lending conditions has no important role in the transmission of credit supply shocks and is thus not a driving force for the asymmetric effects we find in our paper.

II Methodology

Before we show our results, we introduce our methodological approach in this section. First, we describe the derivation of our credit supply shock. We then explain our econometric methodology to uncover asymmetries in the transmission of these shocks. Finally, we set out our approach to statistical inference.

A. *Deriving the Credit Supply Shock*

In order to derive a credit supply shock, we use an auxiliary structural VAR model identified by means of sign restrictions for the US. Let $\mathbf{y}_t$ be an $n \times 1$ vector including real GDP, consumer prices, loan volumes, a composite lending rate, and a reference short-term (shadow) interest rate. The SVAR reads

$$\mathbf{y}_t'\mathbf{A}_0 = \sum_{j=1}^{J} \mathbf{y}_{t-j}'\mathbf{A}_j + \mathbf{c} + \varepsilon_t' \,, \tag{1}$$

where ε_t is an $n \times 1$ vector of structural shocks, $\mathbf{A}_j$ is an $n \times n$ matrix of structural parameters for $0 < j \leq J$ lags with $\mathbf{A}_0$ invertible and $\mathbf{c}$ is a $1 \times n$ vector of parameters.

The SVAR model in (1) can be rewritten as

$$\mathbf{y}_t'\mathbf{A}_0 = \mathbf{x}_t'\mathbf{A}_+ + \varepsilon_t' \,, \tag{2}$$

where $\mathbf{A}_+' = [\mathbf{A}_1', ..., \mathbf{A}_J', \mathbf{c}']$ and $\mathbf{x}_t' = [\mathbf{y}_{t-1}', ..., \mathbf{y}_{t-J}', 1]$ for $1 \leqslant t \leqslant T$. Hence, we estimate the reduced-form VAR

$$\mathbf{y}_t = \mathbf{x}_t'\mathbf{B} + \mathbf{u}_t' \,, \tag{3}$$

where $\mathbf{B} = \mathbf{A}_+\mathbf{A}_0^{-1}$, $\mathbf{u}_t' = \varepsilon_t\mathbf{A}_0^{-1}$, and $\mathbf{E}[\mathbf{u}_t\mathbf{u}_t'] = (\mathbf{A}_0\mathbf{A}_0')^{-1}$.

Since the structural parameters are not identified, we need to impose some restrictions. We therefore rely on the identification strategy of, among others, Gambetti and Musso (2017) who draw inference concerning the response of (log) real GDP, (log) consumer prices, (log) loan volumes, a composite lending rate, and a reference short-term interest rate to four structural shocks: (i) aggregate supply, (ii) aggregate demand, (iii) monetary policy, and (iv) credit supply.[4] The latter is of main interest for our further analysis. Table (1) summarizes the identification restrictions. The identification of an aggregate supply shock, an aggregate demand shock and a monetary policy shock is quite standard, we therefore will not discuss it here.[5]

The identification of the expansionary credit supply shock, on the other hand, deserves further discussion. As pointed out by, inter alia, Christiano et al. (2010) and

[4] In order to estimate (3), we rely on Bayesian techniques using a Minnesota prior. Inference is based on 20000 draws, where the first 10000 draws are discarded, as samples that have been generated in early iteration steps are likely to be not representative for the true posterior distribution. Data is compiled as in Gambetti and Musso (2017) (see supplementary material therein) and extended until 2018Q4.

[5] For a comprehensive description, see Gambetti and Musso (2017).

Table 1: Identification restrictions

Shock	Real GDP	Prices	Short-term interest	Lending rate	Loan Volume
Aggregate Supply	+	–	No restriction	No restriction	No restriction
Aggregate Demand	+	+	+	+	No restriction
Monetary Policy	+	+	–	No restriction	No restriction
Credit Supply	+	+	+	–	+

Notes: The identifying assumptions are imposed on impact, where '+' means an increase and '–' a decrease in the underlying variable. All shocks are normalized as leading to an increase in real GDP.

Gambetti and Musso (2017), a credit supply shock can be associated with various events, such as unexpected changes in bank capital availability for loans due to changes in regulatory capital ratio requirements or unanticipated changes in the degree of competition in the banking sector. The identification we use can be thought of as (i) shocks to the bank funding technology or bank reserve demand, as in Cúrdia and Woodford (2010), as well as (ii) shocks to bank's capital quality and bank's net worth, as in Gertler and Karadi (2011). An expansionary credit supply shock consequently leads to an increase in real GDP and prices. This notion is in line with the identification scheme of all model-specific credit supply shocks in the models of Cúrdia and Woodford (2010) and Gertler and Karadi (2011) as well as some specific credit supply shocks in Gerali et al. (2010) and Christiano et al. (2010).

An exogenous increase in loan supply (credit supply) is also assumed to lead to a contemporaneous drop in the lending rate as well as an increase in the short-term interest rate, which is under the control of the central bank. An exogenous expansion of the supply of loans to the private sector via a decrease in the lending rate has expansionary effects as the lower costs of external funds enable to expand consumption while firms can expand their investments.[6] The central bank counteracts the resulting price pressure by increasing the short-term interest rate.

Finally, note that while it is common practice to derive shocks from an auxiliary SVAR model, our hybrid VAR-LP approach is internally inconsistent. This is because in the SVAR model, structural shocks are identified under the assumption that the data generating process is linear and symmetric, as pointed out by Barnichon et al. (2022).[7] Of course, one might consider deriving the structural shock in a state-dependent local projections model as well. However, this approach has two limitations: first, also within an LP approach along the lines of Plagborg-Møller and Wolf (2021), we need to proceed in two steps, where, in the first step, a linear local projections model is fitted to the observed data to extract a credit supply shock. In a second step, this

[6] We take a closer look at this transmission channel in Section III.

[7] They directly estimate a vector moving-average model to derive financial shocks and rely on functional approximations of impulse responses (FAIR). However, their robustness checks imply that a hybrid VAR-LP approach (as ours) and the internally-consistent approach using FAIR yield very similar results.

shock is then used in another local projections model to uncover possible asymmetries in the transmission of credit supply shocks. The reason is that within both, an LP framework as well as a VAR framework, it is not possible to separately identify positive and negative credit supply shocks a priori without knowing the shock in advance. The second point is more of a technical nature. Structural identification in a local projection approach works very similarly to a VAR approach, whereby the variance-covariance matrix is derived from the $h = 1$ projection residuals for the identification. The (only) crucial difference, as pointed out by Plagborg-Møller and Wolf (2021), is that within a local projections framework, the reduced-form impulse responses come directly from the direct projections rather than from iterative forecasts as in a VAR model. Importantly, as pointed out by Plagborg-Møller and Wolf (2021), the projection residuals from a local projections approach equal these Wold innovations. As a result, the variance-covariance matrix $\mathbf{\Sigma}$ as obtained within a local projections framework contains the same information as obtained from a VAR framework. Hence, if and only if we restrict the responses of the endogenous variables on impact only, as is done in this paper, then we get the same set of structural shocks in the local projections framework as in the VAR framework. Since we are not interested in the impulse responses from the first stage but only in the identified credit supply shock for the second stage, it makes no difference to us whether we identify our shock in a VAR framework or a local projections framework.

Nevertheless, we use an alternative measure of a credit supply shock obtained from a proxy SVAR in the robustness section.

B. *Econometric Setup*

In order to uncover the asymmetric effects of structural shocks, we follow Tenreyro and Thwaites (2016) and rely on local projections, as proposed by Jordà (2005). The local projection method provides a flexible framework and is easy to implement. Moreover, it is well documented that local projections have several advantages over VAR models. Above all, local projections are more robust to possible misspecifications, at least under a finite lag structure. Moreoever, they allow us to parsimoniously model asymmetric effects and, in effect, saves degrees of freedom relative to a multivariate approach. That is, even though we lose observations from adjusting for leads and lags, we ultimately save degrees of freedom because our set of control variables on the right-hand side is relatively sparse as we do not need to describe the dynamics of the endogenous variables conditional on the shock.

Local projections base on the idea to directly regress the dependent variable at different horizons $t + h$ for $h = 0, 1, ..., H$, conditional on an information set Ω_t that consists of a set of control variables. In the linear case, the regression equation reads

$$y_{t+h} = \alpha_h + \gamma_h \mathbf{x}_t + \beta_h shock_t + e_{t+h} \ , \tag{4}$$

where y_{t+h} is the variable of interest at horizon $t+h$, $\mathbf{x}_t$ is a vector of control variables, and $shock_t$ is the identified structural shock.[8] The coefficient β_h measures the average response of the dependent variable to the shock that hits the economy at time t. Thus, one constructs the impulse responses as a sequence of the β_h's estimated in a series of separate regressions for each horizon.

Note that (4) is easily adapted to estimate a model that allows for non-linear effects. More precisely, we want to test whether positive shocks have the same impact as negative shocks. This can be done by regressing

$$y_{t+h} = \alpha_h + \gamma \mathbf{x}_t + \beta_h^+ \max\{shock_t, 0\} + \beta_h^- \min\{shock_t, 0\} + e_{t+h} \,. \tag{5}$$

While the information sets in (4) and (5) do not differ, the coefficients β_h^+ and β_h^- do now allow us to test for sign-dependent impulse responses. In particular, the response of y_{t+h} on a shock in t is now given by

$$\frac{\partial y_{t+h}}{\partial shock_t} = \begin{cases} \beta_h^+, & \text{if} \quad shock_t \geqslant 0, \\ \beta_h^-, & \text{if} \quad shock_t < 0. \end{cases} \tag{6}$$

It is important to stress that a perfectly symmetric transmission of credit supply shocks would imply that $\beta_h^+ = \beta_h^-$. That is, exogenous expansions and contractions of credit supply would have the same effects in absolute terms. Contrary to this, we would point to asymmetric effects when the divergence between β_h^+ and β_h^- is significantly different from zero.

Our vector $\mathbf{x}_t$ contains control variables that are supposed to have an effect on the endogenous variable y_t. We therefore include $p = 2$ lagged values of the short-term federal funds rate, the consumer price index, and real GDP, where consumer prices and real GDP are in logs and multiplied by 100. Finally, $\mathbf{x}_t$ also includes $q = 3$ lags of the dependent variable.[9]

C. *Inference*

Regressing the dependent variable at different horizons on the same set of control variables will likely result in autocorrelated residuals. In order to calculate standard errors that account for the possibility of serially correlated residuals within and across

[8] More precisely, $\mathbf{x}_t$ summarizes p lagged values of a vector of control variables, $controls_t$, and q lagged values of the dependent variable. Hence, we get $\gamma_h \mathbf{x}_t \equiv \gamma_h^{controls} \sum_{j=1}^{p} controls_{t-j} + \gamma_h^y \sum_{k=1}^{q} y_{t-k}$.

[9] The choice of p and q is based on the Schwartz Bayesian Information Criterion given by $-2\ln(\widehat{L}) + k\ln(n)$, where k is the number of parameters, $(\widehat{L})$ is the maximized value of the likelihood function, and n is the effective sample size after adjusting for leads and lags. We tried different combinations of p and q ranging from 1 to 4 lags each. For each combination of p and q, we follow the strategy of Tenreyro and Thwaites (2016) and sum up the resulting information criteria over both, the horizon h and over all different dependent variables. Finally, the minimum value of this operation results in the optimal lag length of $p = 2$ and $q = 3$.

equations, we follow Ramey and Zubairy (2018) and Tenreyro and Thwaites (2016) and estimate seemingly unrelated equations as proposed by Driscoll and Kraay (1998). To be more precise, we estimate the parameters of interest of each equation separately and, in a second step, average the moment conditions across horizons $h = 0, ..., H$ when deriving Newey-West standard errors. In effect, the Driscoll-Kraay standard errors account for autocorrelation across both, time t and horizons h. We follow standard practice (see Jordà, 2005) and set the maximum autocorrelation lag for the Newey-West procedure $L = h + 1$.

III Results

In this section, we present the baseline results of our paper. In the baseline setting, the idea is to uncover possible asymmetries in the responses of the economy to a credit supply shock. Hence, the baseline regression focuses on the asymmetric effects of positive and negative shocks on the economy. The sample size covers data from 1975Q3 to 2018Q4, consisting of 174 observations. After adjusting for leads and lags, the effective sample size starts in 1976Q2 and ends in 2013Q4 and, hence, consists of 151 observations.

A. *Baseline Results*

Do credit supply shocks have asymmetric effects? Before we answer this question, it is important to bear in mind that throughout the paper, we present impulse response coefficients rather than impulse responses.[10] So, for example, a *positive* value of β_h^- indicates that after a *negative* shock (i.e., an unexpected credit curtailment), the effect, i.e. impulse response, is *negative*. In other words, there is a uniflow response to the shock.

This being said, Figures (1) - (4) read as follows: The rows correspond to the dependent variables that are affected by the credit supply shocks. The first column depicts the impulse response coefficients β_h^+ (red-solid lines) following a positive credit supply shock. In contrast, the second column depicts the impulse response coefficients β_h^- (red-solid lines) describing the response following a negative credit supply shock to the particular dependent variable, both accompanied by their respective 90 percent confidence bands. For further comparison, dashed lines represent the impulse response coefficients from the linear model (equation 4), i.e. β_h.

The third column shows $t-$statistics testing the null hypothesis $H_0 : (\beta_h^+ - \beta_h^-) = 0$ for adjacent horizons $h = 0, ..., H$, where the shaded area covers the $t-$critical values for a 90 percent confidence interval, i.e. ± 1.645. If the $t-$ statistics (red-solid) lie outside the shaded-area, we reject the null that the difference in the response of the endogenous variable is not distinguishable from zero in favor of the alternative hypothesis that

[10] In order to get impulse responses, one simply needs to flip the impulse response coefficients.

they are significantly different from zero at the 90 percent confidence level and thus, indicate non-negligible asymmetric effects.

Figure 1: Response of credit volumes to credit supply shocks

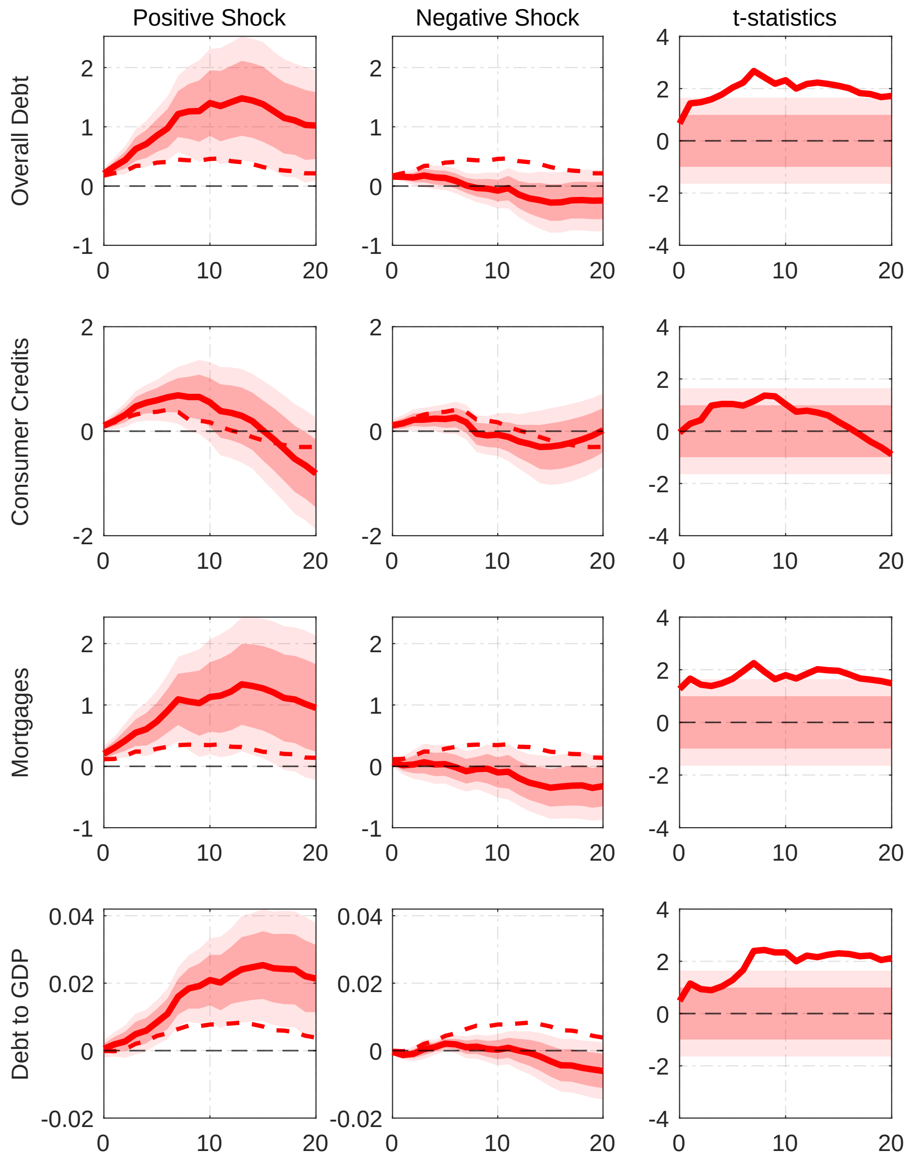

Notes: The first column shows the impulse response coefficients (red-solid) β_h^+ for $h = 0, ..., H$ for a positive (one standard deviation) credit supply shock, the second column shows the impulse response coefficients (red-solid) β_h^- for a negative (one standard deviation) credit supply shock. In both cases, the dark (pale) red-shaded area corresponds to the 68 (90) percent confidence interval, relying on Driscoll-Kraay standard errors. The red-dotted lines in the first two columns show the impulse response coefficients β_h from a linear model without testing for asymmetric effects. The third column shows the t–statistics testing the null that $H_0 : (\beta_h^+ - \beta_h^-) = 0$ for each horizon h using the Driscoll-Kraay method. The dark (pale) red-shaded area covers the t–critical values for a 68 (90) percent confidence interval, i.e. ±0.995 (±1.645). The rows show, from top to bottom, the responses of overall debt volume (in percent), the volume of consumer credit (in percent), mortgage credit volume (in percent), and the share of debt-to-GDP (in percentage points).

Figure (1) shows the impulse response coefficients corresponding to different debt volumes, namely overall household debt, consumer credit, mortgages, as well as the response of debt-to-GDP for each horizon h (x-axis) after a shock hits the economy in t. Given a positive credit supply shock, overall debt steadily and significantly increases by 1.5 percent. This effect is highly persistent as it holds for more than 20 quarters. The volume of consumer credit increases significantly, reaching a peak of 0.9 percent after ten periods before it wears out. The effect of a credit supply shock on consumer credit is more short-lived than overall debt. The response of mortgages shows where the sluggish behavior of overall debt stems from: mortgages respond very similar to a positive credit supply shock as does overall debt in both, magnitude and duration.[11] The positive effect proceeds over more than 20 quarters with a peak response of 1.5 percent. Debt-to-GDP shows a picture similar to both, the response of overall debt and mortgage debt. The positive response, which peaks at 0.025 percent, becomes significant after five periods and remains thereafter, shows evidence that debt is more sensitive to a credit supply shock than real GDP.

Turning to the responses to a negative shock, recall that a positive response coefficient indicates a concurrent reaction such that, e.g., concerning overall debt, a detrimental credit supply shock leads to a decrease in overall debt for two quarters. In the medium-term, i.e., after roughly ten quarters, the negative shock goes into reverse and leads to an increase of debt, though not significant, as the confidence bands of the response coefficients comprise zero. This is true for all debt categories, as well as the debt-to-GDP ratio. Above all, it stands out that the responses are in all cases other than the response of consumer credit, less pronounced than in the case of a positive credit supply shock. By the same token, when not accounted for asymmetries, one gets misleadingly rather mild responses to credit supply shocks.[12]

Furthermore, the responses in all cases are more sticky than in the case given a positive shock. Even though the response of the overall indebtedness, mortgage debt, and the debt-to-GDP ratio is slightly significantly different from zero, the concurrent effect is short-lived, at most, and tends to increase debt in the medium-term, with response coefficients significantly different from zero at the 90 percent confidence level.

Most importantly, in all cases except consumer credit, the responses to a positive (negative) shock reside well above (below) their symmetric counterparts, indicating that there are non-negligible asymmetric effects.

The third column of Figure (1) underpins this visual impression. The $t-$statistics for overall debt, mortgage debt, and the debt-to-GDP ratio show that we reject the null whereupon private debt equally responds to positive and negative shocks. Solely for consumer credit, we do not reject the null hypothesis based on our $t-$statistics

[11] This comes at no surprise as the share of mortgage credit to the overall indebtedness amounts to over 90 percent in the mid-2000's which further explains the results.

[12] A finding in line with Barnichon et al. (2022) statement regarding financial shocks.

and conclude that, following an exogenous shock to credit supply, consumer credit responds similarly in absolute terms.

Three insights from this first exercise need to be highlighted: first, in the linear model, the responses are driven by the positive component of the shock, thus, suppressing a debt-increasing effect of adverse shocks in the medium-run, as seen in the second column of Figure (1).[13] Second, the differences in the responses between the linear and non-linear models are most pronounced for overall debt, mortgage debt, and debt-to-GDP; measures that are central to the debate whether central banks should actively lean against the credit cycle in their policymaking.[14] Third, not accounting for asymmetries leads to misleadingly attenuated responses of the aforementioned debt measures.

Turning to the effects on key macro variables, the first column of Figure (2) reports the results for real GDP and consumer prices (both in logs) as well as for the short-term (shadow) interest rate. Following a positive credit supply shock, real GDP increases immediately by about 0.25 percent, peaking at 0.5 percent after four quarters and steadily reverts afterward. This effect is significant for about seven periods. What stands out is that the effect of the shock does not wear out but instead leads to a subsequent decline in real GDP. After 16 quarters, there is a significant drop in output of 0.5 percent. This is in line with the well-established notion that credit supply expansions lead to subsequent episodes of economic downturn, as in, e.g., Schularick and Taylor (2012) or Mian et al. (2017).

The economic upturn caused by additional availability of credit leads to an increase (though not significant) in consumer prices by up to 0.2 percent which starts to decline after 10 quarters. As real GDP subsequently decreases, so do prices, leading to deflationary pressure after 18 quarters which is significantly different from zero. In light of the responses of debt-to-GDP both, the responses of real GDP and prices add up to Fisher's *debt-deflation hypothesis* whereupon an economic slowdown increases the real burden of debt, which in turn puts downward pressure on aggregate demand, slowing down economic activity even further.

As the monetary authority responds to output and prices in the conduct of its policies, it follows the boom-bust pattern. First, it increases its short-term interest rate in response to the increase in those indicators. The peak response is an increase of one percentage point after four quarters. As the economic downturn comes into action, interest rates respond with a substantial decrease.

The effects following an unexpected credit crunch are not very different to those from the linear model. A negative credit supply shock leads to a significant decrease

[13]For example, the response coefficient for mortgage debt given an expansionary credit supply shock is three times higher than the equivalent response coefficient from the linear model.

[14]See, for example, Lambertini et al. (2013) and Svensson (2017). In the robustness section, we evaluate the role of monetary policy in the transmission of credit supply shocks in a counterfactual experiment and derive policy implications.

Figure 2: Credit supply shocks and headline variables

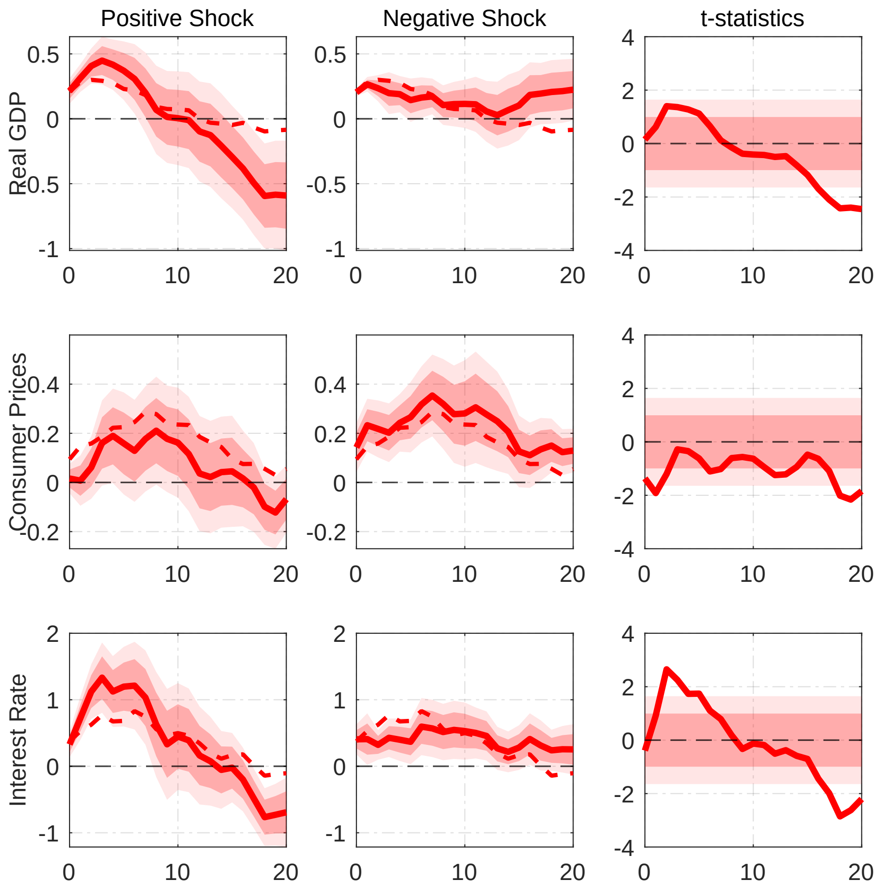

Notes: The first column shows the impulse response coefficients β_h^+ (red-solid) for $h = 0, ..., H$ for a positive (one standard deviation) credit supply shock, the second column shows the impulse response coefficients β_h^- (red-solid) for a negative (one standard deviation) credit supply shock. In both cases, the dark (pale) red-shaded area corresponds to the 68 (90) percent confidence interval, relying on Driscoll-Kraay standard errors. The red-dotted lines in the first two columns show the impulse response coefficients β_h from a linear model without testing for asymmetric effects. The third column shows the t–statistics testing the null that $H_0 : (\beta_h^+ - \beta_h^-) = 0$ for each horizon h using the Driscoll-Kraay method. The dark (pale) red-shaded area covers the t–critical values for a 68 (90) percent confidence interval, i.e. ± 0.995 (± 1.645). The first row shows the response of real GDP (in percent), the second row the response of consumer prices (in percent), and the third row shows the response of the effective federal funds rate (amended by the Wu-Xia shadow rate) in percentage points.

in output and prices, and as such, interest rates decrease to counteract the economic downturn. Here, after four years, prices hint to patterns of asymmetry, mounting to a stronger deflationary pressure in the presence of negative shocks. This also explains the asymmetric effects of real debt variables in the mid-run, because the strong response of prices following a negative shock seem to dampen the overall response of real debt following an unexpected negative credit supply shock.

Summing up, our results point to symmetric effects in the short-run, as the t-statistics fluctuate within, rather than outside of its critical values for the first four years or so. Thereafter, however, we find that all three key variables respond asymmetrically. More precisely, in all three cases, the response following a positive credit supply shock is stronger in absolute values than in the case of a negative credit supply shock. Furthermore, it is worth noting that our results suggest that in the presence of a negative shock, prices exhibit stronger deflationary pressure, which are underestimated in a linear model.

B. *Digging Deeper*

There are different channels through which a credit supply deterioration can transmit into the (real) economy. An expansion in credit supply, therewith decrease in lending rates, could, for example, boost the supply side of the economy through additional (funds for and realizations of) investments due to decreasing credit costs, as argued by inter alia Gilchrist and Zakrajšek (2012), which in turn would lead to an increase in employment. Aggregate demand, on the other hand, can be stimulated by means of credit expansion by enabling households to increase consumption as their balance sheets improve, as noted by i.a. Mian and Sufi (2014) and Gertler and Gilchrist (2018). For this reason, we search for traces of asymmetry in the responses of some surrogates for both, aggregate supply and aggregate demand to credit supply deteriorations.

Beginning with the supply side, Figure (3) shows how external financing, measured via the credit spread provided by Gilchrist and Zakrajšek (2012), the total non-farm payroll employment, investments, and real credits to non-financial corporations respond.

As the top panel shows, we find that, following a negative credit supply shock, the response of the external finance premium is different from zero only on impact and indistinguishable from zero afterwards. This is contrary to the response following a positive credit supply shock. Here, we see no significant response on impact and the subsequent three years. Afterwards, however, the finance premium increases, which coincides with the responses to the bust pattern we observe for our macroeconomic variables.[15]

In the case of employment, total non-farm payrolls clearly exhibits the boom-bust pattern. If a positive credit supply shock hits the economy, employment significantly increases on impact by 0.05 percent and increases up to 0.2 percent before it reverts and becomes significantly negative after almost four years. In the case of a negative credit supply shock, total non-farm payrolls co-move by 0.1 percent for 12 periods, which translates into an decrease in payrolls. Both results are mostly indistinguishable,

[15] As financing conditions can play a crucial rule for economic fluctuations (see, for example, Adrian et al., 2010 or Bruno and Shin, 2015), we have also looked at the subcomponents of the Chicago Fed National Financial Conditions Index (NFCI), which comprise a wide range of indicators concerning volatility and funding risk, credit conditions, as well as debt and equity measures. In short, the impulse responses of the subindexes are remarkably similar to the impulse responses of the GZ spread, which is why we refrain from presenting them in the paper. Of course the results are available upon request.

Figure 3: Credit supply shocks and the supply side

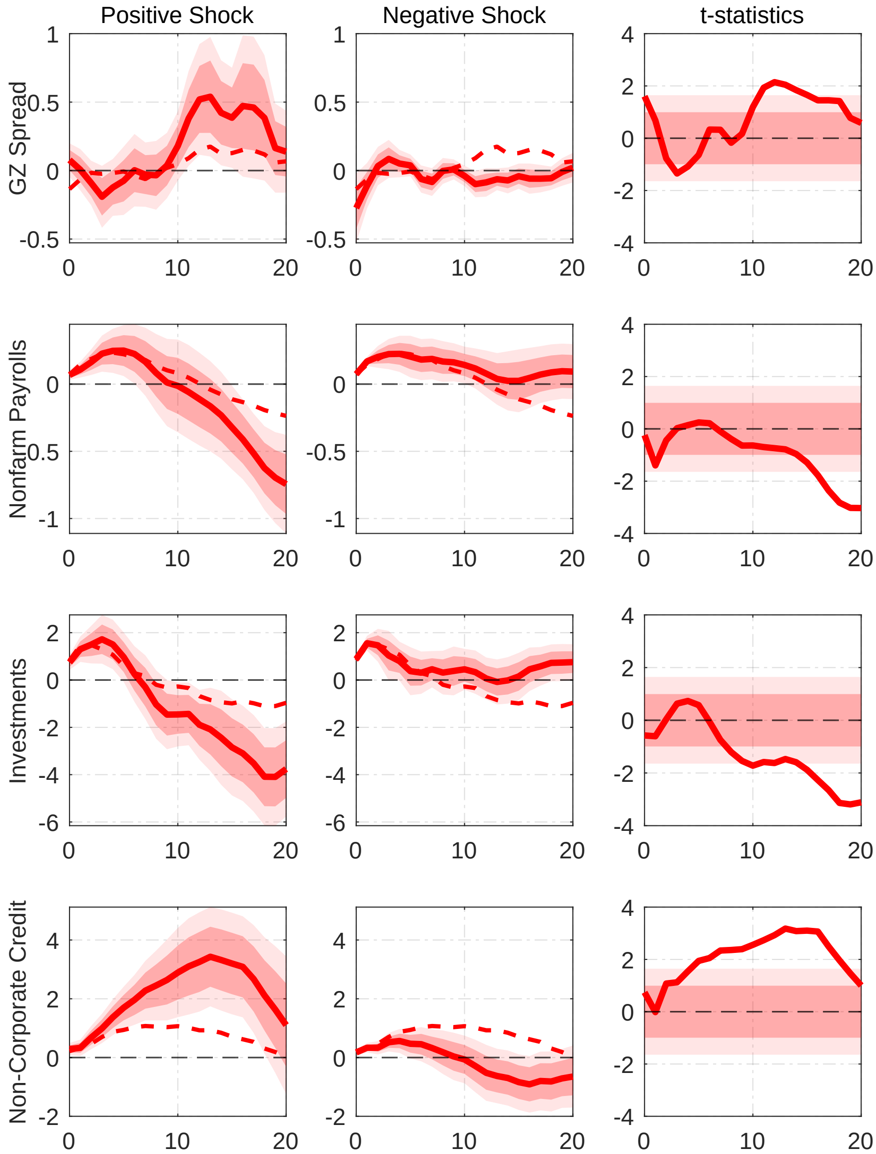

Notes: The first column shows the impulse response coefficients β_h^+ (red-solid) for $h = 0, ..., H$ for a positive (one standard deviation) credit supply shock, the second column shows the impulse response coefficients β_h^- (red-solid) for a negative (one standard deviation) credit supply shock. In both cases, the dark (pale) red-shaded area corresponds to the 68 (90) percent confidence interval, relying on Driscoll-Kraay standard errors. The red-dotted lines in the first two columns show the impulse response coefficients β_h from a linear model without testing for asymmetric effects. The third column shows the t−statistics testing the null that $H_0 : (\beta_h^+ - \beta_h^-) = 0$ for each horizon h using the Driscoll-Kraay method. The dark (pale) red-shaded area covers the t−critical values for a 68 (90) percent confidence interval, i.e. ±0.995 (±1.645). The first row shows the response of the Gilchrist and Zakrajšek (2012) spread (in percentage points), the second depicts the response of total non-farm payrolls (in percent), the third shows the response of real investments (in percent), and the last row shows the response of real credits to non-financial corporations (in percent).

except, as before, at the long-end of our analysis, and a result of the changes in economic activity given the respective shocks.

Investment responds very similarly to both, positive and negative credit supply shocks, as does real GDP. After a significant and somewhat persistent increase on impact, peaking at 1.5 percent, the response mean-reverts and becomes significantly negative by 1.5 percent after 11 quarters. As before, we observe asymmetric effects after 16 quarters, i.e. at the farther end of our projection. Also here, we find that the response of a positive credit supply shock is stronger in absolute value than in the case of a negative credit supply shock, as indicated by the significantly negative t-statistic in the respective panel.

Taken together, the responses of the external finance premium, non-farm payrolls, as well as real investments represent the well established boom-bust cycle caused by a credit supply deterioration.

Finally, real credits to non-financial corporations exhibit a significant and persistent increase in response to a positive credit supply shock. The peak median response, which amounts to 3.4 percent, is observed after 13 quarters, before the response reverses. The timing also matches the responses of overall debt and mortgage debt, which comes at no surprise, as these variables follow the same cycle. The same is true in the case of a negative credit supply shock: the initial decrease in real credits wears out after roughly two years and tends to increase thereafter. As in the case of the other credit variables, the responses show significant asymmetries.

Mian et al. (2017) stress that the boom-bust cycles of the past four decades were primarily driven by household debt operating through the household demand channel. For example, the combination of rising house prices (and thus improvement in the household's balance sheet) and declining lending rates (i) led to an increase in residential investments and (ii) enabled additional consumption of both, domestic and foreign goods and services.[16] Hence, we take a look at the responses of house prices, real residential investments, and real household expenditures, which are depicted in Figure (4).

The first row shows that the response of house prices (real Shiller Index) resembles remarkably well the response of mortgage debt. Positive credit supply shocks lead to a persistent and also significant increase in house prices. While negative shocks lead to quite similar opposite effects, we find clear patterns of asymmetry as the t-statistics are very close at or above the critical values from quarter 4 onwards. This is because, at the back-end, house prices (i) remain on a relatively high level after a positive shock and (ii) tend to exhibit mean-reverting behavior after a negative shock, which in turn increases the difference between the β coefficients. Furthermore, the linear response is mitigated by the response to a negative shock, underestimating the stark effect of

[16]For further impressions, see the evolution of real mortgage rates depicted in Figure (10) in the appendix.

Figure 4: Credit supply shocks and the demand side

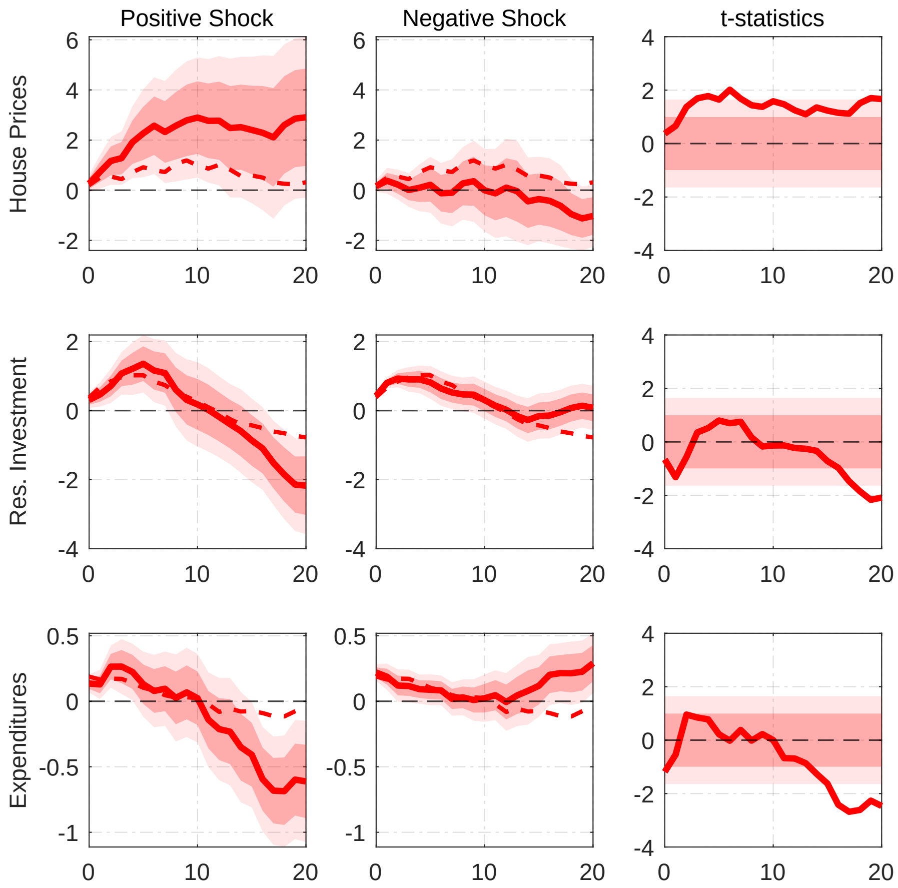

Notes: The first column shows the impulse response coefficients β_h^+ (red-solid) for $h = 0, ..., H$ for a positive (one standard deviation) credit supply shock, the second column shows the impulse response coefficients β_h^- (red-solid) for a negative (one standard deviation) credit supply shock. In both cases, the dark (pale) red-shaded area corresponds to the 68 (90) percent confidence interval, relying on Driscoll-Kraay standard errors. The red-dotted lines in the first two columns show the impulse response coefficients β_h from a linear model without testing for asymmetric effects. The third column shows the t–statistics testing the null that $H_0 : (\beta_h^+ - \beta_h^-) = 0$ for each horizon h using the Driscoll-Kraay method. The dark (pale) red-shaded area covers the t–critical values for a 68 (90) percent confidence interval, i.e. ± 0.995 (± 1.645). The first row shows the response of house prices (real Shiller index, in percent).The second row depicts real residential investments (in percent). The last row shows the response of real personal expenditures (in percent).

positive credit supply shocks on house price inflation.

The response of residential investment is similar to the response of total investments. This can be explained by higher demand for housing due to an improvement in the household's balance sheets. The response of residential investment peaks at one percent after one year before it mean-reverts, exhibiting the boom-bust pattern. After 20 quarters, residential investments significantly decrease by approximately two percent. Concerning a negative credit supply shock, the same story applies as in the case of total investments.

It stands out that for both shocks, consumption expenditures follow the response of consumer credit and real GDP very closely. Given a positive credit supply shock, expenditures increase initially and show mean-reverting behavior that eventually results in a significant decrease. That is, beneficial credit conditions improve the household's balance sheets (e.g. through higher asset prices such as housing) such that they increase borrowing (especially consumer credit) with which they expand consumption. In the case of a contractionary credit supply shock, expenditures decrease significantly by 0.2 percent and tend to remain declining in the long-run.

To summarize, three things stand out: First, the responses of consumer credit, real GDP, the interest rate, non-farm payrolls, total and residential investment, as well as personal expenditures to a positive credit supply shock all contribute to the notion that credit supply expansions lead to a vicious boom-bust cycle.

Second, the response of credits to non-financial corporations, house prices, and mortgage debt is highly persistent and does not follow the boom-bust pattern. One explanation could be the following: Positive credit supply shocks increase funding possibilities for corporations as well as house prices and, thus, stimulate economic activity, which in turn leads to an increase in demand for (new or the expansion of existing) mortgage loans. As the economy, however, transits into the bust-phase, monetary policy steps into place, decreasing interest rates. This decrease, in turn, leads to an increase in asset prices, such as house prices. Thus, while the economy shows the boom-bust pattern, house prices do not. For example, a closer look at the responses of house prices and real GDP reveals that the pace of change in house prices decreases after 6 to 10 quarters; the period where the boom goes bust.

Last, asymmetries appear in the back-end of the responses and are primarily driven by the positive portion of the shock.[17] In turn, this implies that not accounting for asymmetries leads to underestimated effects of credit supply shocks in the medium to long-run.[18]

IV Sensitivity Analysis

So far, our results indicate that the transmission of credit supply shocks is asymmetric. This section seeks to underpin our results via a battery of robustness checks, includ-

[17] At the back-end, the t-statistics are either (i) overwhelmingly negative, as in the case of consumer prices, non-farm payrolls, as well as total and residential investments or (ii) significantly negative, as consumer prices, non-farm payrolls, investment-to-GDP, and expenditures, which indicates stronger responses to positive than negative shocks.

[18] Another observation, that an anonymous referee thankfully brought to our attention, is that we observe a boom-bust pattern in the impulse responses of economic variables, on the one hand, but no such pattern in the responses of the debt variables, on the other hand. One potential reason is the difference in the underlying cycles. As far as credit and debt variables are concerned, these variables are rather sluggish compared to classical economic variables and are merely linked to the credit cycle. There is much evidence that the letter is twice as long as the business cycle (see, for instance, Alpanda and Zubairy, 2019). This could explain why, for some variables, we do not yet observe a change in the sign of the response over our 20-quarter horizon.

ing an alternative choice of credit supply shock as well as several alternative model specifications.

A. *Credit Supply Shocks from a Proxy SVAR*

Our results in Section III rely on the identification strategy following Gambetti and Musso (2017) through sign restrictions. Since the identified credit supply shock comes from a different model and enters our local projection framework exogenously, we crosscheck our results using an alternative credit supply shock. Using a Monte Carlo experiment, Mumtaz et al. (2018) provide detailed evidence that the performance of various structural vector autoregression models seeking to identify credit supply shocks varies substantially.[19] Using data for the US, they find that the Gambetti and Musso (2017) identification performs well in replicating DSGE model-implied impulse response functions. Several other authors use alternative proxies for credit supply, including Gilchrist and Zakrajšek (2012), Bassett et al. (2014), and Lown and Morgan (2006), and use it as an endogenous variable in an otherwise standard VAR, where a shock to the proxy is interpreted as a credit supply shock. However, Mumtaz et al. (2018) show that proxy variables in a recursive SVAR do not perform well due to a large attenuation bias increasing in the variance of the measurement error. They also show that the measurement error has little effect on the performance of proxy VARs because the proxy of credit supply does not enter the model directly. We, therefore, extend the sample of Mumtaz et al. (2018) and estimate a proxy SVAR model where the excess bond premium of Gilchrist and Zakrajšek (2012) is considered as an instrument rather than an additional endogenous variable.[20]

The correlation of the estimated credit supply shock from our proxy SVAR and our credit supply shock identified as in Section II is 0.59. Both series show the same pattern most of the time and the dynamics during periods of financial turmoil as in 2008 overlap almost perfectly.[21]

Figures (11) to (14) in the appendix show that in most cases, the qualitative directions of the impulse responses do not change. Also, the difference between the coefficients

---

[19] The authors consider three different DSGE models featuring credit supply in order to construct artificial data, namely the DSGE model by Gertler and Karadi (2011), the estimated DSGE model from Christiano et al. (2014) and, finally, the model by Cúrdia and Woodford (2010). In a nutshell, Mumtaz et al. (2018) use these DSGE models as the true data generating processes and consider various competing SVAR models with different identification strategies in order to shed light on their replication performance of the true impulse responses following the model-implied credit supply shocks.

[20] We use the same time span as in Section III, namely 1975Q3 to 2018Q4. Our estimation is based on one lag, as suggested by the Bayesian Schwartz Information Criterion. For details on the estimation, see Mumtaz et al. (2018).

[21] The reliability statistic proposed by Mertens and Ravn (2014), as the squared correlation between the proxy and the credit supply shock, is 0.16 in our case and notably higher as in Mumtaz et al. (2018), thus indicating higher reliability as a strong instrument. Most strikingly, however, is that all impulse responses within our proxy SVAR show the expected sign. This is interesting from the standpoint that in our proxy SVAR, the identification is far less restrictive than in our baseline model from Section II.

β_h^+ and β_h^- looks very much like in our benchmark case. However, while the t–statistics for the macro variables still point to asymmetric effects, uncertainty in the impulse response coefficients of debt volumes after a positive shock is remarkably higher than after negative shocks, resulting in t-statistics that are not statistically different from zero. Overall, we conclude that our results still point to asymmetric effects in most cases when we use a credit supply shock from a proxy SVAR instead of a shock identified by means of sign restrictions.

B. *A Note on the Role of Estimation Uncertainty Surrounding the Shock*

So far, our impulse response coefficients are the mean responses to the median of a series of identified shocks, that satisfy the imposed restrictions. Consequently, estimation uncertainty surrounding the shock is not taken into account. Our set of $B = 10000$ draws for the credit supply shock from the VAR nevertheless allows us to test how representative our shock is. For each series $shock_t^{(1)}, ..., shock_t^{(10000)}$, we therefore estimate our baseline model and report the uncertainty across all draws. Note that the potential disadvantage of the median is that the median in t and $t + 1$ can come from different draws (and thus different models, see Fry and Pagan, 2011 on this). By using the entire time series for each draw, we avoid this inconsistency.
Figure (5) reports the baseline point estimates (red-solid) with the 68 percent and 90 percent confidence bands (red-shaded areas) as well as the median over all point estimates (black-solid) and 90 percent (black-dashed) confidence intervals across all 10000 draws. It stands out that our median shock seems to be very representative of the population, as the median across all 10000 estimates is very close to our baseline impulse responses.

C. *Is Monetary Policy a Potential Cause for Asymmetric Effects?*

As mentioned in the introduction, it is quite conceivable that monetary policy decisions can be a cause of asymmetric responses, if, for example, monetary policy systematically responds differently to expansionary shocks than to contractionary shocks. To examine the role of monetary policy in the transmission of credit supply shocks in our setup, we re-estimate the VAR model based on the assumption, that the policy rate does not react to contemporaneous and past movements of the composite lending rate and credit volumes by means of block exogeneity. In this scenario, the policy rate only reacts to movements in those variables via feedback from the remaining variables of the model, but not directly to movements of those variables.

In a second step, we repeat our baseline estimation

$$y_{t+h} = \alpha_h + \gamma_h \mathbf{x}_t + \beta_h^+ \max\left\{shock_t^{cf}, 0\right\} + \beta_h^- \min\left\{shock_t^{cf}, 0\right\} + e_{t+h},$$

where y_{t+h} is the variable of interest at horizon $t+h$, $\mathbf{x}_t$ is a vector of control variables,

Figure 5: Impulse responses across all draws of the credit supply shock

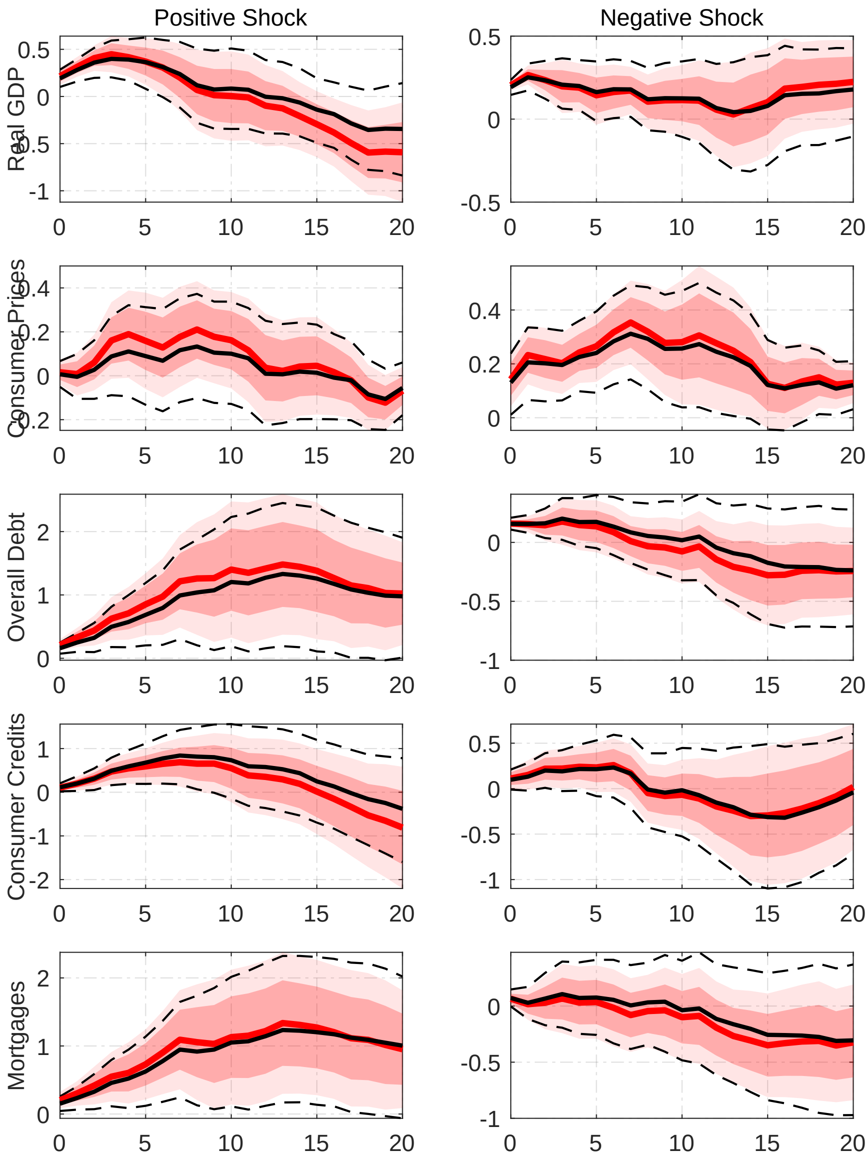

Notes: Point estimates (red-sold) with the 68 and 90 percent confidence bands (red-shaded areas) from the baseline model as well as the median over all point estimates (black-solid) and 90 percent (black-dashed) confidence bands across 10000 draws.

and $shock_t^{cf}$ is the credit supply shock of this counterfactual exercise.

Figure (6) compares the credit supply shock from our baseline setup with the shock we obtain in this exercise. Here, the red-solid line reports the mean of the baseline specification with the corresponding 90 percent confidence bands (shaded area), respectively. In contrast, the black-solid line corresponds to the median shock of the counterfactual scenario. As can be seen, the two shocks mostly overlap.

Figure (7) reports the corresponding impulse response coefficients for this exercise together with the original impulse response coefficients for a subset our endogenous

Figure 6: Baseline shock versus counterfactual shock

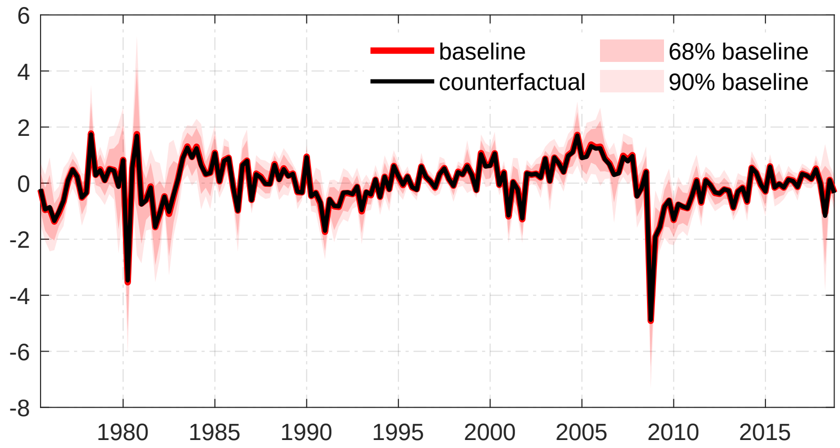

variables, in order to conserve space.[22] The red-solid line corresponds to the baseline impulse response coefficients, whereas the black-solid line corresponds to the impulse response coefficients obtained from the counterfactual experiment credit supply shock. Note that we only report the baseline specification's 68 percent and 90 percent confidence bands.

In all cases, there is virtually no difference in the responses to the shocks, as the impulse response coefficients are basically alike. In other words, our findings suggests that the central bank's systematic response to credit volumes and lending conditions has no important role for the transmission of credit supply shocks and is thus not a driving force for the asymmetric effects we find in our results.

What are the policy implications for monetary policy? The results of the counterfactual analysis by no means imply that monetary policy does not influence the transmission of the credit supply shock at all. This is because the restriction within the counterfactual experiment merely restricts the central bank's systematic response to credit volumes and the composite lending rate. However, a credit supply shock also affects other variables, including real GDP and prices, i.e., variables to which the central bank reacts. Consequently, monetary policy acts in the background, but is not in the position to directly influence the course of events. In other words, leaning against the credit cycle would not meet the desired objectives. Hence, from a policy point of view, it is necessary to asymmetrically respond to unexpected disruptions in the credit market by means of macroprudential activity. One conceivable measure would be to implement buffers whose amount is aligned with the sign or size of the credit shock – similar to the counter-cyclical capital buffers, which are determined by the business cycle. However, we would like to leave a more thorough analysis of this to further research.

[22] The impulse responses for the remaining variables are available on request.

Figure 7: Baseline IRFs vs counterfactual IRFs

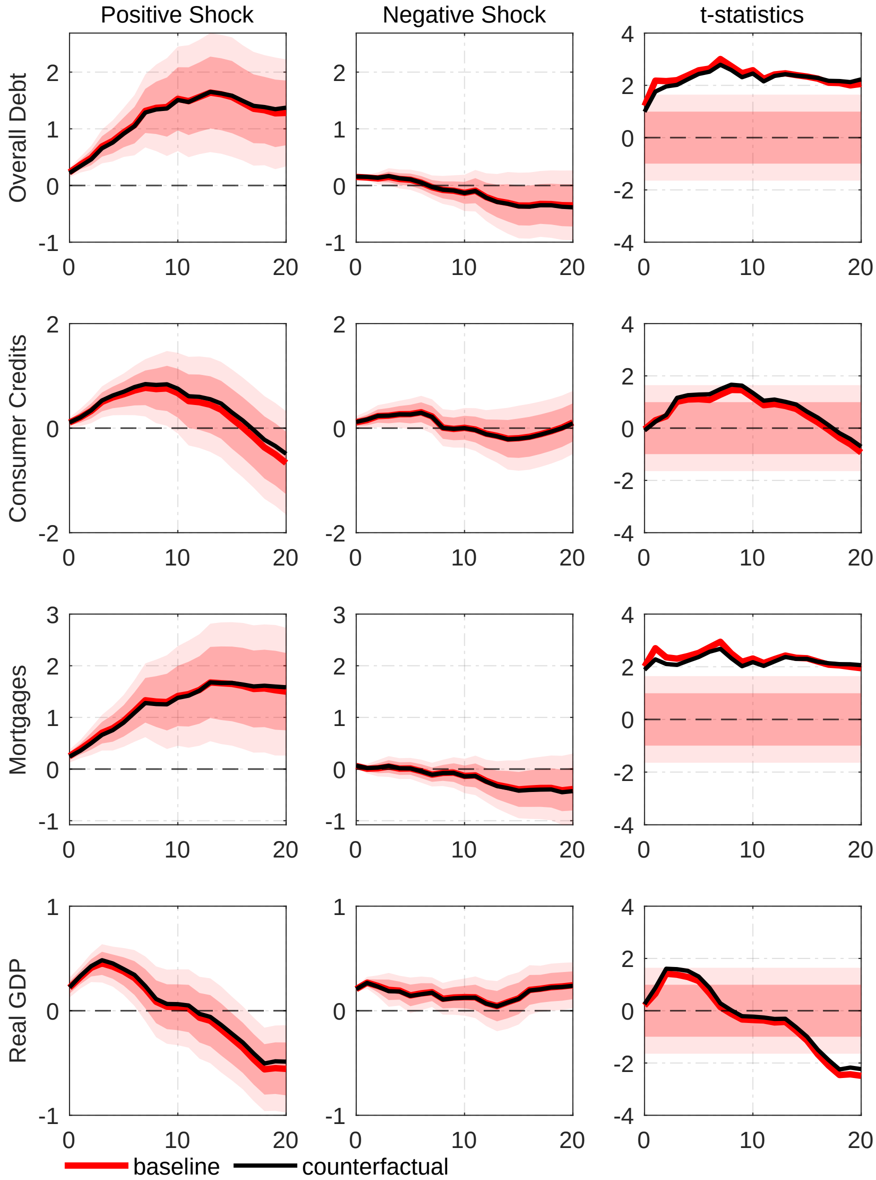

Notes: Impulse response coefficients for the baseline (red-solid) and counterfactual (black-solid) credit supply shocks. Dark (pale) red-shaded areas are the 68 (90) percent confidence intervals of the responses to the baseline shock, relying on Driscoll-Kraay standard errors. The last column shows the t−statistics testing the null that $H_0 : (\beta_h^+ - \beta_h^-) = 0$ for each horizon h using the Driscoll-Kraay method. Here, the dark (pale) red-shaded areas cover the t−critical values for a 68 (90) percent confidence interval, i.e. ±0.995 (±1.645). All responses are expressed in percent.

D. *The Effects of the Overall Business Cycle and Credit Conditions*

One explanation for our observed asymmetries could be that the distribution of our identified credit supply shocks itself is asymmetric, depending on either the overall business cycle or credit conditions. First, we devote ourselves to the effects of the overall business cycle.

One could argue that positive credit supply shocks occur mainly during periods of expansion, while adverse credit supply shocks could predominantly occur during recessions. If this was the case, we probably would have to account for this state dependence in our regression analysis. In order to test for this possibility, we compare the distribution of shocks during different phases of the business cycle. To do so, we follow the procedure of Tenreyro and Thwaites (2016), which amounts to estimating state-depending probability distributions via a smoothly increasing logistic function $F(z_t)$ as a weighting function of the kernel, where z_t is an indicator of the state of the economy. To do so, we take the annualized quarterly growth rate of real GDP. Since this series is very volatile, we smooth out short-term volatility by filtering the series using a seven-quarter moving average. Subsequently, we calculate the logistic function over time. We follow Teräsvirta and Granger (1993) and employ a function of the form

$$F(z_t) = \frac{\exp\left(\kappa \frac{z_t - \mu}{\sigma_z}\right)}{1 + \exp\left(\kappa \frac{z_t - \mu}{\sigma_z}\right)}, \tag{7}$$

where μ is used to control the proportion of the sample the economy spends in either state (boom or recession), and σ_z is the sample standard deviation of the state variable z_t. The parameter κ controls how abruptly the economy switches from one state to the other following movements of the state variable. That is, higher values of κ mean that rather small movements of the state variable are needed to induce a switch from one regime to the other. We follow Auerbach and Gorodnichenko (2012) and Tenreyro and Thwaites (2016) and choose a value of $\kappa = 3$, indicating an intermediate degree of intensity of the regime-switching. We calibrate μ such that we spend 11.3% of the time in recessions. According to the NBER recession indicators, this corresponds to the share of periods we spend in recessions within our sample.

With the logistic function at hand, we estimate both, a probability density function (PDF) and cumulative density function (CDF) using our logistic function $F(z_t)$ and $1 - F(z_t)$, respectively, for weighting our (Normal Gaussian) kernel function in order to get state-dependent distributions of our credit supply shock. The resulting state-depending PDFs and CDFs are shown in Figure (8). The bottom panel shows the weighting function $F(z_t)$ (black-solid line), where the red-shaded areas highlight NBER recessions. As expected, we observe drops in our function $F(z_t)$ during recessions. This happens with a short delay since we use a moving average.

Figure 8: Estimated state-dependent PDFs and CDFs

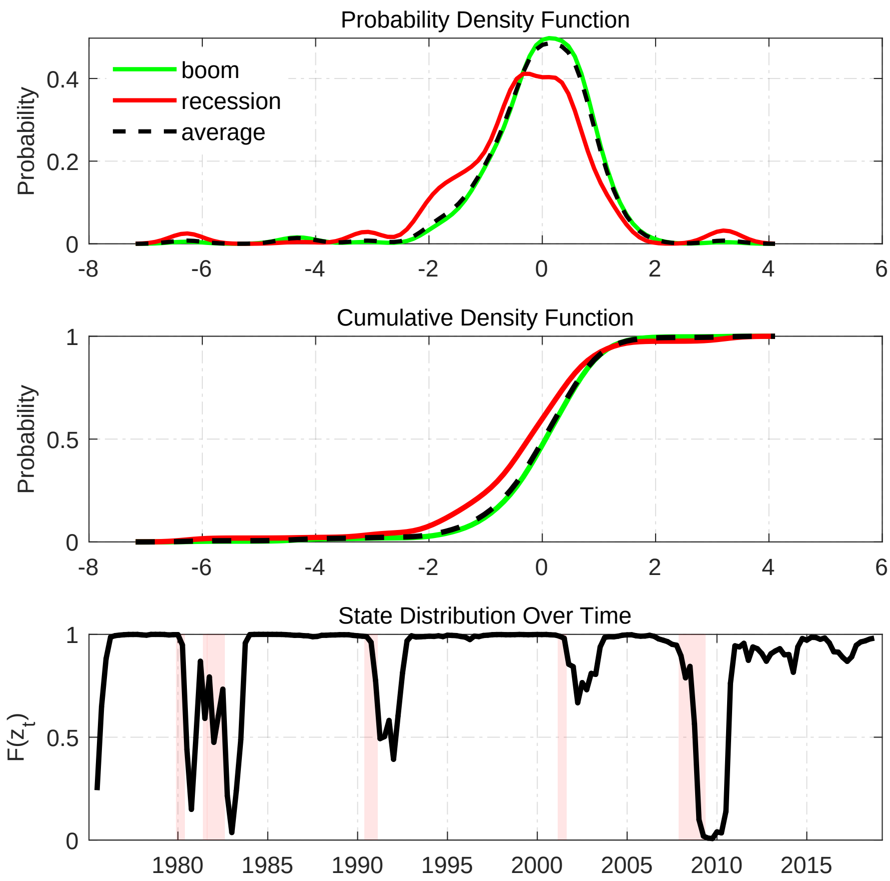

Notes: Estimated probability density functions (top panel), cumulative density functions (middle panel) and transition function $F(z_t)$ (bottom panel) over time. The green lines show the estimated distributions of credit supply shocks during booms, and the red lines show the estimated distributions during recessions using $F(z_t)$ and $1 - F(z_t)$ for weighting our kernel function. The black-dotted lines correspond to the average distributions of credit supply shocks using a normally distributed kernel. In the bottom panel, the red-shaded areas highlight NBER recession dates.

Turning to the estimated probability distribution functions, the average PDF of our credit supply shock appears to be normally distributed, which is not surprising because our shock is distributed as $shock_t \sim \mathcal{N}(0, I)$ by construction within our VAR. The distribution during booms almost perfectly overlaps the average distribution. The estimated PDF during recessions is also clearly centered around zero. Not surprisingly, we observe that large (negative) shocks occur mainly during recession. This is also visible in the estimated CDF, where the probability of getting a shock of minus one standard deviation or less (in absolute terms) in size is slightly higher during recessions. Nevertheless, the overall picture clearly points to common central tendencies of the distributions across the business cycle. We therefore conclude that the distribution of shocks does not depend on the business cycle.

However, it is possible that positive credit supply shocks primarily occur when financial conditions are looser than average, whereas negative credit supply shock may predominantly occur when financial conditions are tighter than average. If we find this to be true, this would imply that we possibly would have to take different states of financial conditions into account and therefore have to adjust our regression. We, therefore, take the NFCI Credit subindex and repeat our exercise in order to investigate whether the distribution of our credit supply shock systematically depends on financial credit conditions.[23]

Figure 9: Estimated state-dependent PDFs and CDFs | z_t : Financial Conditions Credit Subindex

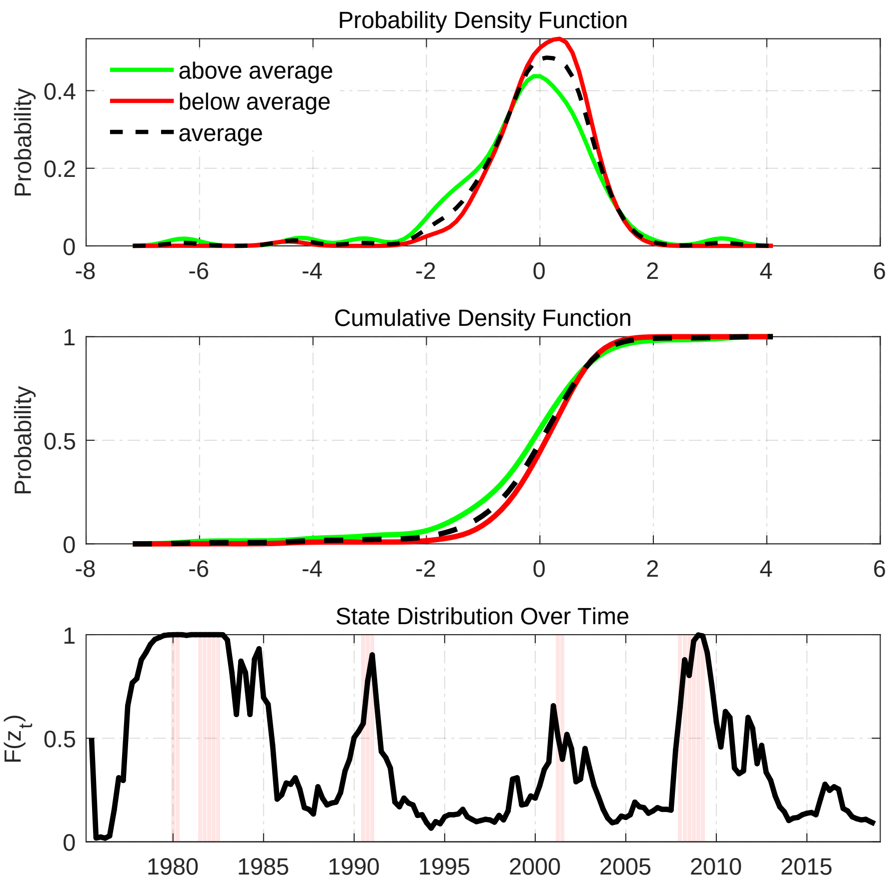

Notes: Estimated probability density functions (top panel), cumulative density functions (middle panel) and transition function $F(z_t)$ (bottom panel) over time. The green lines show the estimated distributions of credit supply shocks during booms. The red lines show the estimated distributions during recessions using $F(z_t)$ and $1 - F(z_t)$ for weighting our kernel function. The black-dotted lines correspond to the average distributions of credit supply shocks using a normally distributed kernel. In the bottom panel, the shaded areas highlight the NBER recession dates.

Starting with the dynamics of $F(z_t)$ at the bottom panel of Figure (9), it turns out that

[23] We repeat this exercise for different state variables, namely, the output gap and the adjusted Chicago Fed National Financial Conditions Index. The resulting estimated probability density functions as well as the estimated cumulative density functions can be found in appendix (B), Figures (19) and (20), respectively.

whenever we are in a recession as indicated by the red-shaded area, credit conditions are above average, i.e. tighter on average. The estimated probability density functions at the top panel show that small shocks in absolute values are slightly more likely when credit conditions are below average, i.e., when credit conditions are relatively loose (red line), whereas larger shocks in absolute terms seem to be somehow more common during periods of credit tightening (green line). However, both estimated state-dependent distributions are clearly centered around zero and not much different from the estimated average probability distribution. This finding is mirror-imaged in the middle panel, where the cumulative probability distributions do not differ notably. We conclude that, similar to the previous exercise, the distributions of credit supply shocks both, when credit conditions are above and below average, follow common central tendencies.

E. *Altering Specifications*

In order to further stress test our results, we make changes to the sample size as well as take a closer look at the role of lags and trends in order to ensure that these factors do not distort our results.

First, we take a look at the role of the Great Recession. Throughout the paper, we estimate our model over an effective period from 1975Q3 to 2013Q3. The asymmetric effects we found so far could possibly be a relic of the Great Recession. We test for this possibility and estimate our model, excluding NBER recession dates that mark the Great Recession. The blue-dashed line in Figures (15)-(18) in the appendix show the results of this exercise. In short, the results are generally similar to our baseline such that we conclude that our findings are not driven by the Great Recession.

The next exercise is to investigate the effects of lags and trends. In our baseline setting, we do not include a log-linear trend into our regression equation in order to keep the regression equation consistent across our endogenous variables.[24] This subsection examines whether adding a trend changes our main results. We do so by estimating

$$y_{t+h} = \delta t + \alpha_h + \gamma \mathbf{x}_t + \beta_h^+ \max\{0, shock_t\} + \beta_h^- \min\{0, shock_t\} + e_{t+h} \, , \tag{8}$$

where δ denotes the effect of a log-linear time trend t, while the remainder is similar to our baseline regression (5). The results do not change, as can be seen form the course of the black-dotted line in Figures (15)-(18).

As explained in Section II, the choice of lags for both, the endogenous variable as well as the set of control variables other than the endogenous variables, relies on the Bayesian Schwartz Information Criterion. We, therefore, test whether the re-

[24]Note that our effective sample size runs from 1975Q3 to 2013Q4. For some variables under consideration during this time span, it can be seen with the naked eye that adding a linear time trend would probably deliver misleading results.

sults change if our lag lengths are chosen according to the more restrictive Akaike Information Criterion (AIC) given by

$$2k - 2\ln(\widehat{L}) \,. \tag{9}$$

The AIC suggests $p = 4$ lags of the endogenous variable and $q = 3$ lags for the remaining variables, as opposed to the benchmark case where two lags of the endogenous variable have been used. Overall, we find that the choice of lags does not change our results remarkably, as the respective impulse responses (black-dashed line) are very similar to our baseline results (red-solid), as Figures (15)-(18) show. Furthermore, the $t-$statistics as well show the same pattern as in Section III.

V Conclusion

The Great Recession has renewed the interest in the nexus between credit cycles and economic activity. One important finding is that credit developments can lead to vicious boom-bust cycles. That is, what starts with a stimulation of economic activity eventually results in economic downturn (e.g. Schularick and Taylor, 2012; Jordà et al., 2013). Another finding is that the boom-bust phases observed in the past four decades goes back to credit supply expansions (e.g. Mian et al., 2020; Justiniano et al., 2019). However, the research so far is limited to symmetric effects of credit supply developments. As pointed out by Barnichon et al. (2022), this may lead to somewhat counterintuitively lenient results. One reason is that originally differently operating shocks will eventually level out in a symmetric setup. This, as well as the theoretical literature that shows asymmetries (and non-linearities) in the response to financial shocks, raises the question, whether credit supply shocks cause asymmetric effects.

They do. Especially overall private debt, mortgage debt, debt-to-GDP, and house prices exhibit asymmetries in the response to credit supply shocks. In this respect, positive credit supply shocks tend to have stronger and more prolonged effects.

Furthermore, our results underpin the narrative of the boom-bust cycle in the presence of financial distortions which, again, is more pronounced in the presence of positive credit supply shocks. After an initial increase in economic activity (output, prices, investments, expenditures, and consumer credit) that lasts for five to ten quarters, the economy transits into a bust phase with a notable slowdown in economic activity. In contrast, negative credit supply shocks cause notably stronger deflationary pressure.

Looking at some surrogates of aggregate demand and aggregate supply, the boom-bust narrative is further confirmed, yet it does not explain the highly asymmetric and persistent response of overall debt, mortgage debt, and debt-to-GDP. However, even though this paper does not aim at tracing down the specific causes of asymmetries, we find that house prices, and thus the household-driven demand channel (e.g. Mian and Sufi, 2018) are key for the persistence in the response of mortgage debt and debt-to-

GDP to credit supply shocks.

Lastly, if we abstract from asymmetries, we get relatively mild responses for debt and prices in the presence of credit supply shocks, such that the true effects tend to be underestimated.

However, we do not trace down the causes of asymmetries, which may be rooted in a variate of reasons: occasionally binding borrowing constraints, asymmetric information due to market imperfections, or behavioral biases, to name a few. We leave further going down the rabbit hole for future research.

References

Adrian, Tobias, Emanuel Moench, and Hyun Song Shin, "Macro Risk Premium and Intermediary Balance Sheet Quantities," *IMF Economic Review*, 2010, *58* (1), 179–207.

Alpanda, Sami and Sarah Zubairy, "Household Debt Overhang and Transmission of Monetary Policy," *Journal of Money, Credit and Banking*, 2019, *51* (5), 1265–1307.

Auerbach, Alan J and Yuriy Gorodnichenko, "Measuring the Output Responses to Fiscal Policy," *American Economic Journal: Economic Policy*, 2012, *4* (2), 1–27.

Barnichon, Regis, Christian Matthes, and Alexander Ziegenbein, "Are the Effects of Financial Market Disruptions Big or Small?," *Review of Economics and Statistics*, 2022, *104* (3), 557–570.

Bassett, William F, Mary Beth Chosak, John C Driscoll, and Egon Zakrajšek, "Changes in Bank Lending Standards and the Macroeconomy," *Journal of Monetary Economics*, 2014, *62*, 23–40.

Bruno, Valentina and Hyun Song Shin, "Capital Flows and the Risk-Taking Channel of Monetary Policy," *Journal of Monetary Economics*, 2015, *71*, 119–132.

Christiano, Lawrence J, Roberto Motto, and Massimo Rostagno, "Financial Factors in Economic Fluctuations," *ECB Working Paper*, 2010, *No. 1192.*

_ , _ , and _ , "Risk Shocks," *American Economic Review*, 2014, *104* (1), 27–65.

Colombo, Valentina and Alessia Paccagnini, "Does the Credit Supply Shock Have Asymmetric Effects on Macroeconomic Variables?," *Economics Letters*, 2020, p. 108958.

Cúrdia, Vasco and Michael Woodford, "Credit Spreads and Monetary Policy," *Journal of Money, Credit and Banking*, 2010, *42*, 3–35.

Driscoll, John C and Aart C Kraay, "Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data," *The Review of Economics and Statistics*, 1998, *80* (4), 549–560.

Fry, Renee and Adrian Pagan, "Sign Restrictions in Structural Vector Autoregressions: A Critical Review," *Journal of Economic Literature*, 2011, *49* (4), 938–960.

Gambetti, Luca and Alberto Musso, "Loan Supply Shocks and the Business Cycle," *Journal of Applied Econometrics*, 2017, *32* (4), 764–782.

Gerali, Andrea, Stefano Neri, Luca Sessa, and Federico Signoretti, "Credit and Banking in a DSGE Model of the Euro Area," *Journal of Money, Credit and Banking*, 2010, *42*, 107–141.

Gertler, Mark and Peter Karadi, "A Model of Unconventional Monetary Policy," *Journal of Monetary Economics*, 2011, *58* (1), 17–34.

_ and Simon Gilchrist, "What Happened: Financial Factors in the Great Recession," *Journal of Economic Perspectives*, 2018, *32* (3), 3–30.

Gilchrist, Simon and Egon Zakrajšek, "Credit Spreads and Business Cycle Fluctuations," *American Economic Review*, 2012, *102* (4), 1692–1720.

Jordà, Òscar, "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 2005, *95* (1), 161–182.

_ , Moritz Schularick, and Alan Taylor, "When Credit Bites Back," *Journal of Money, Credit and Banking*, 2013, *45* (s2), 3–28.

Justiniano, Alejandro, Giorgio E Primiceri, and Andrea Tambalotti, "Credit Supply and the Housing Boom," *Journal of Political Economy*, 2019, *127* (3), 1317–1350.

Lambertini, Luisa, Caterina Mendicino, and Maria Teresa Punzi, "Leaning Against Boom–Bust Cycles in Credit and Housing Prices," *Journal of Economic Dynamics and Control*, 2013, *37* (8), 1500–1522.

Lown, Cara and Donald Morgan, "The Credit Cycle and the Business Cycle: New Findings Using the Loan Officer Opinion Survey," *Journal of Money, Credit, and Banking*, 2006, *38* (6), 1575–1597.

Mertens, Karel and Morten Ravn, "A Reconciliation of SVAR and Narrative Estimates of Tax Multipliers," *Journal of Monetary Economics*, 2014, *68*, S1–S19.

Mian, Atif, Amir Sufi, and Emil Verner, "Household Debt and Business Cycles Worldwide," *The Quarterly Journal of Economics*, 2017, *132* (4), 1755–1817.

—, —, and —, "How Does Credit Supply Expansion Affect the Real Economy? The Productive Capacity and Household Demand Channels," *The Journal of Finance*, 2020, *75* (2), 949–994.

— and —, "What Explains the 2007–2009 Drop in Employment?," *Econometrica*, 2014, *82* (6), 2197–2223.

— and —, "Finance and Business Cycles: The Credit-Driven Household Demand Channel," *Journal of Economic Perspectives*, 2018, *32* (3), 31–58.

Mumtaz, Haroon, Gabor Pinter, and Konstantinos Theodoridis, "What Do VARs Tell Us About the Impact of a Credit Supply Shock?," *International Economic Review*, 2018, *59* (2), 625–646.

Plagborg-Møller, Mikkel and Christian K Wolf, "Local Projections and VARs Estimate the Same Impulse Responses," *Econometrica*, 2021, *89* (2), 955–980.

Ramey, Valerie A and Sarah Zubairy, "Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data," *Journal of Political Economy*, 2018, *126* (2), 850–901.

Schularick, Moritz and Alan Taylor, "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870-2008," *American Economic Review*, 2012, *102* (2), 1029–1061.

Sedláček, Petr and Vincent Sterk, "The Growth Potential of Startups over the Business Cycle," *American Economic Review*, 2017, *107* (10), 3182–3210.

Svensson, Lars E O, "Cost-Benefit Analysis of Leaning Against the Wind," *Journal of Monetary Economics*, 2017, *90*, 193–213.

Tenreyro, Silvana and Gregory Thwaites, "Pushing on a String: US Monetary Policy Is Less Powerful in Recessions," *American Economic Journal: Macroeconomics*, 2016, *8* (4), 43–74.

Teräsvirta, Timo and Clive Granger, "Modelling Nonlinear Economic Relationships," *OUP Catalogue, Oxford University Press,* 1993.

Appendix

A Other Figures

Figure 10: Real mortgage rates over time

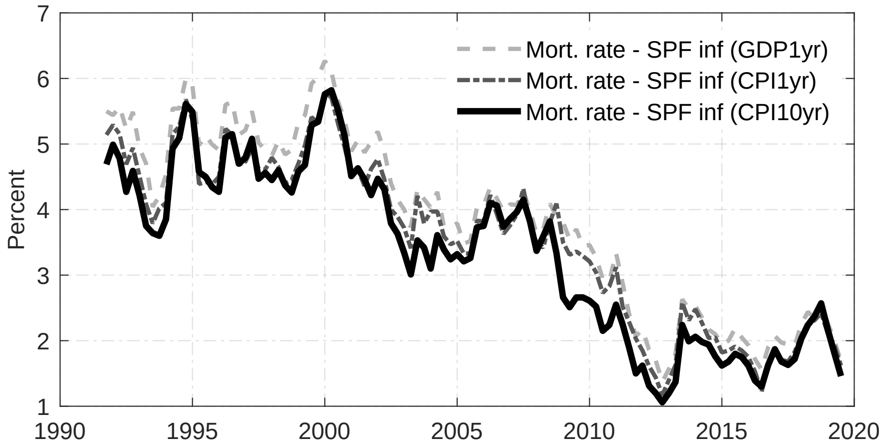

Notes: Real mortgage rate from 1990Q1 to 2019Q3 as the difference between the 30-year fixed mortgage rate and different measures of inflation expectations: GDP deflator-based one year-ahead inflation rate forecast (light grey-dotted), CPI-based one year-ahead inflation rate forecast (dark grey-dotted) and CPI-based 10 years-ahead inflation rate forecast (black-solid). The 30-year fixed mortgage rate is taken from the FRED, and the SPF forecasts are taken from the Philadelphia Fed.

Figure 11: Impulse response of credit volumes to alternative credit supply shock

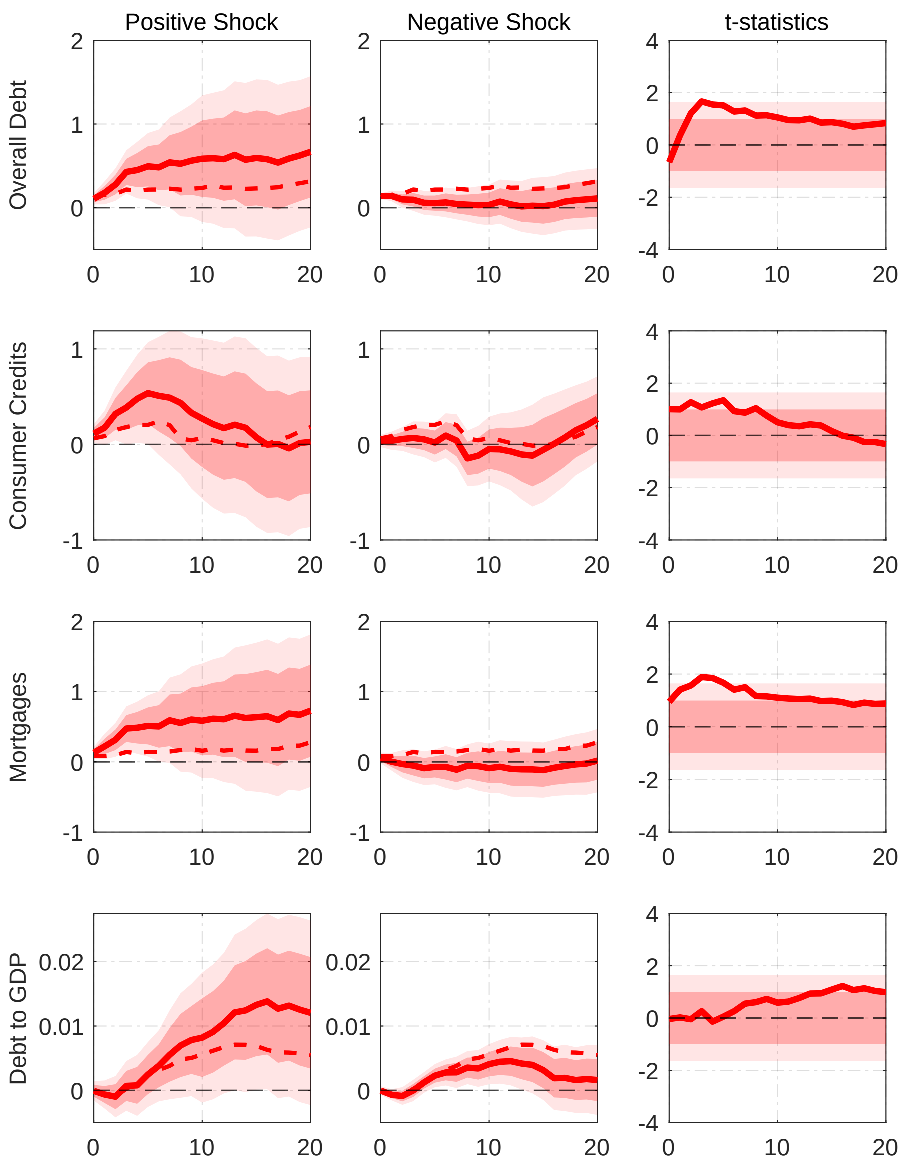

Notes: The first column shows the impulse response coefficients (red-solid) β_h^+ for $h = 0, ..., H$ for a positive (one standard deviation) credit supply shock from the proxy SVAR, the second column shows the impulse response coefficients (red-solid) β_h^- for a negative (one standard deviation) credit supply shock from the proxy SVAR. In both cases, the dark (pale) red-shaded area corresponds to the 68 (90) percent confidence interval, relying on Driscoll-Kraay standard errors. The red-dotted lines in the first two columns show the impulse response coefficients β_h from a linear model without testing for asymmetric effects. The third column shows the t–statistics testing the null that $H_0 : (\beta_h^+ - \beta_h^-) = 0$ for each horizon h using the Driscoll-Kraay method. The dark (pale) red-shaded area covers the t–critical values for a 68 (90) percent confidence interval, i.e. ±0.995 (±1.645). The first row shows the response of overall debt volume (in percent), the second row the response of consumer credit volume (in percent), the third row the response of mortgage credit volume (in percent) and the fourth row shows the response of the share of debt-to-GDP (in percentage points).

Figure 12: Impulse responses of headline variables to alternative credit supply shock

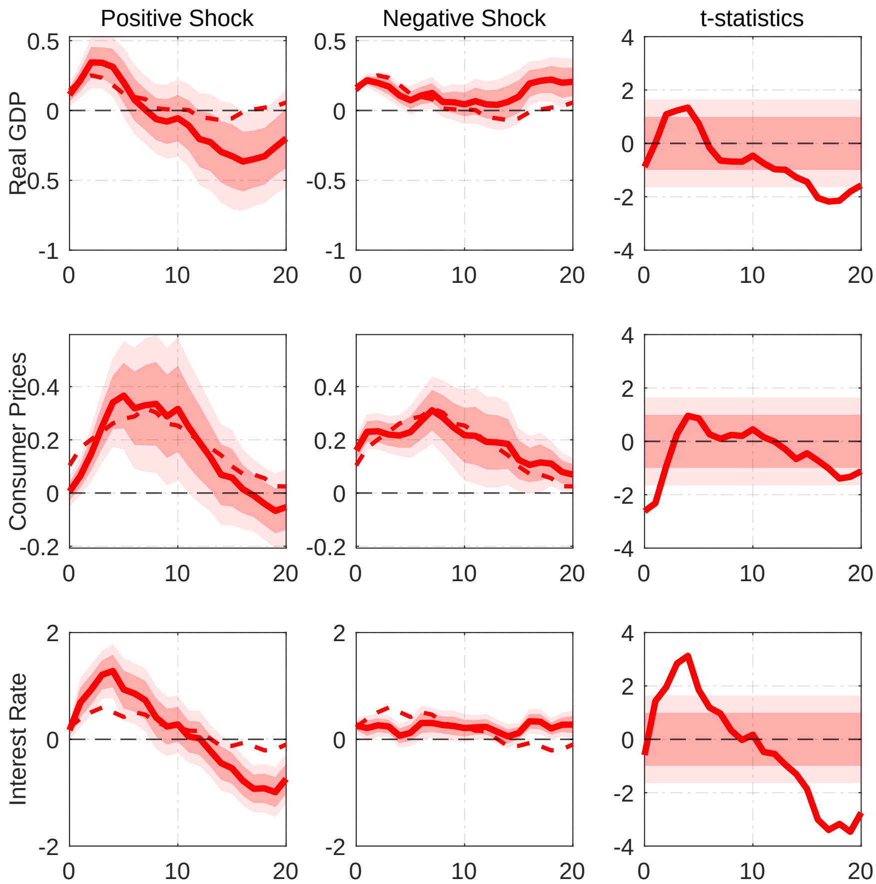

Notes: The first column shows the impulse response coefficients (red-solid) β_h^+ for $h = 0, ..., H$ for a positive (one standard deviation) credit supply shock from the proxy SVAR, the second column shows the impulse response coefficients (red-solid) β_h^- for a negative (one standard deviation) credit supply shock from the proxy SVAR. In both cases, the dark (pale) red-shaded area corresponds to the 68 (90) percent confidence interval, relying on Driscoll-Kraay standard errors. The red-dotted lines in the first two columns show the impulse response coefficients β_h from a linear model without testing for asymmetric effects. The third column shows the t−statistics testing the null that $H_0 : (\beta_h^+ - \beta_h^-) = 0$ for each horizon h using the Driscoll-Kraay method. The dark (pale) red-shaded area covers the t−critical values for a 68 (90) percent confidence interval, i.e. ±0.995 (±1.645). The first row shows the response of real GDP (in percent), the second row the response of consumer prices (in percent) and the third row shows the response of the effective federal funds rate (amended by the Wu-Xia shadow rate) in percentage points.

Figure 13: Impulse responses of supply side variables to alternative credit supply shock

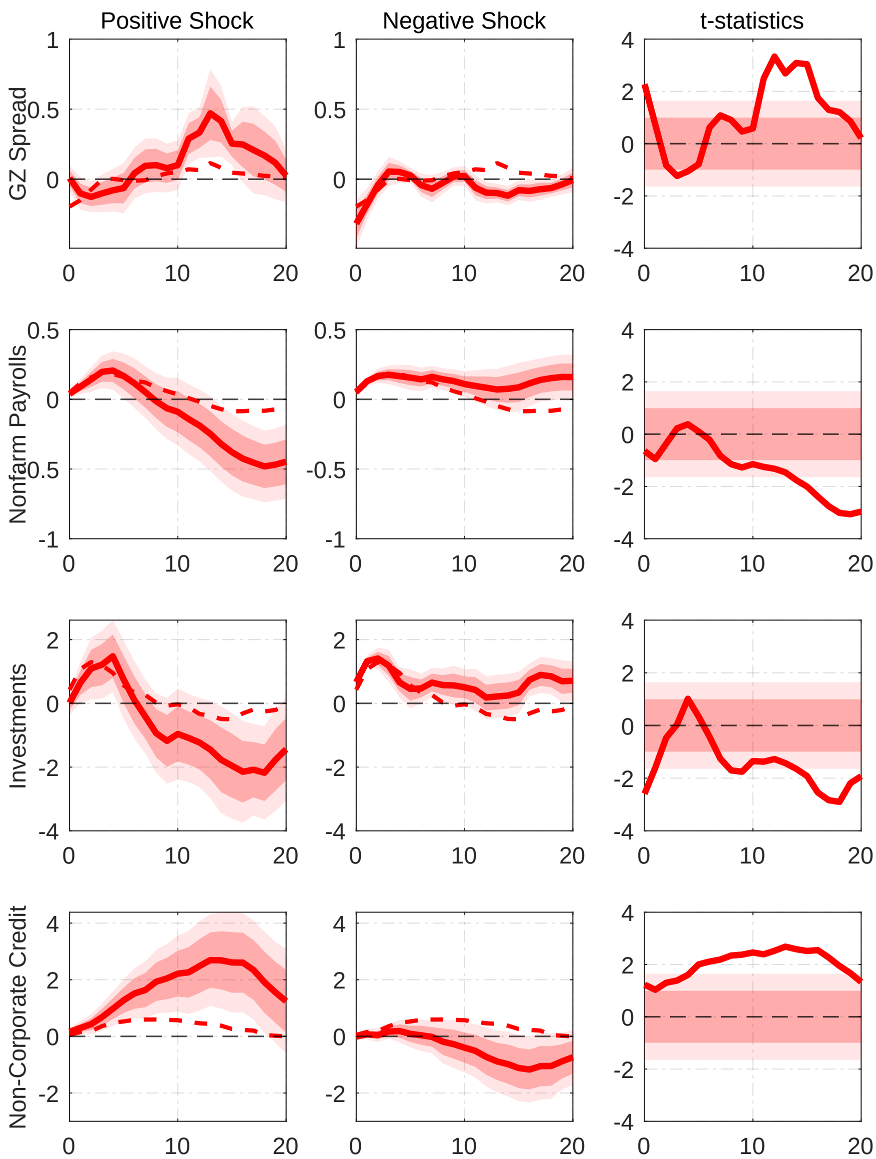

Notes: The first column shows the impulse response coefficients (red-solid) β_h^+ for $h = 0, ..., H$ for a positive (one standard deviation) credit supply shock from the proxy SVAR, the second column shows the impulse response coefficients (red-solid) β_h^- for a negative (one standard deviation) credit supply shock from the proxy SVAR. In both cases, the dark (pale) red-shaded area corresponds to the 68 (90) percent confidence interval, relying on Driscoll-Kraay standard errors. The red-dotted lines in the first two columns show the impulse response coefficients β_h from a linear model without testing for asymmetric effects. The third column shows the t−statistics testing the null that $H_0 : (\beta_h^+ - \beta_h^-) = 0$ for each horizon h using the Driscoll-Kraay method. The dark (pale) red-shaded area covers the t−critical values for a 68 (90) percent confidence interval, i.e. ±0.995 (±1.645). The first row shows the response of the Gilchrist and Zakrajšek (2012) spread (in percentage points), the second depicts the response of total non-farm payrolls (in percent), and the third shows the response of real investments (in percent).

Figure 14: Impulse responses of demand side variables to alternative credit supply shock

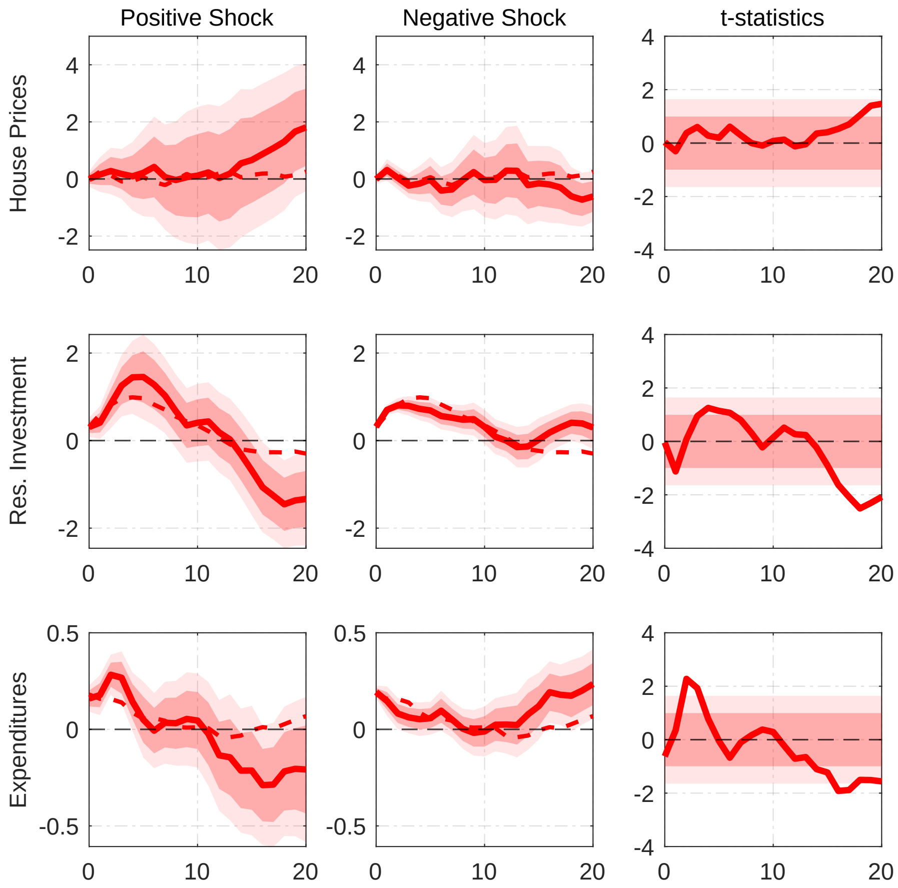

Notes: The first column shows the impulse response coefficients (red-solid) β_h^+ for $h = 0, ..., H$ for a positive (one standard deviation) credit supply shock from the proxy SVAR, the second column shows the impulse response coefficients (red-solid) β_h^- for a negative (one standard deviation) credit supply shock from the proxy SVAR. In both cases, the dark (pale) red-shaded area corresponds to the 68 (90) percent confidence interval, relying on Driscoll-Kraay standard errors. The red-dotted lines in the first two columns show the impulse response coefficients β_h from a linear model without testing for asymmetric effects. The third column shows the $t-$statistics testing the null that $H_0 : (\beta_h^+ - \beta_h^-) = 0$ for each horizon h using the Driscoll-Kraay method. The dark (pale) red-shaded area covers the $t-$critical values for a 68 (90) percent confidence interval, i.e. ± 0.995 (± 1.645). The first row shows the response of house prices (real Shiller index, in percent).The second row depicts real residential investments (in percent). The last row shows the response of real personal expenditures (in percent).

Figure 15: Robustness checks: response of credit volumes

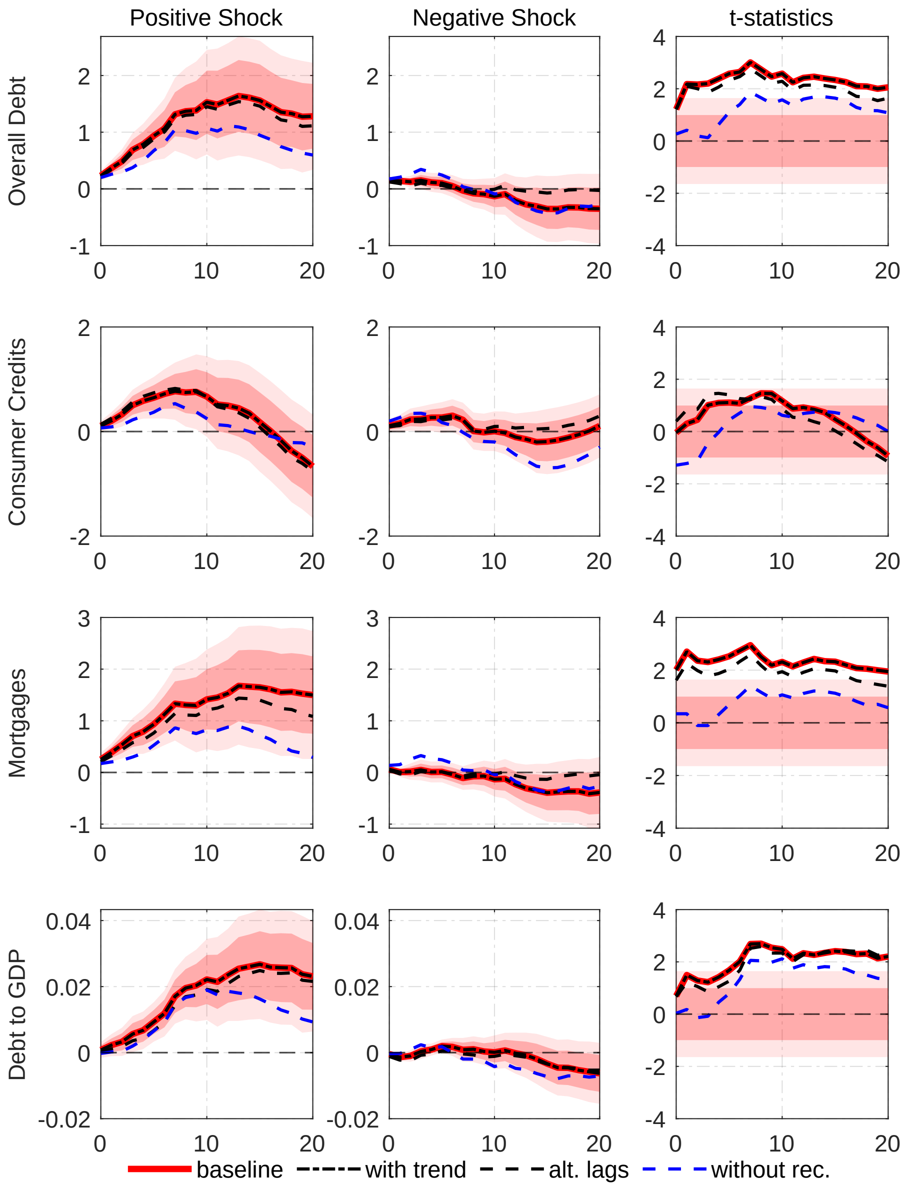

Notes: The first column shows the impulse response coefficients β_h^+ for $h = 0, ..., H$ for a positive (one standard deviation) credit supply shock from the baseline model (red-solid) alongside the impulse responses given alternative model specifications such as a model including a linear trend (black-dotted), a model with an alternative lag specification (black-dashed) as well as the results from an estimation of the baseline model, where the great recession has been excluded (blue-dashed). The second column shows the impulse response coefficients β_h^- for a negative (one standard deviation) credit supply shock. In both cases, the dark (pale) red-shaded area corresponds to the 68 (90) percent confidence interval from the baseline model, relying on Driscoll-Kraay standard errors. The third column shows the corresponding t−statistics testing the null that $H_0 : (\beta_h^+ - \beta_h^-) = 0$ for each horizon h using the Driscoll-Kraay method. The dark (pale) red-shaded area covers the t−critical values for a 68 (90) percent confidence interval, i.e. ± 0.995 (± 1.645), from the baseline model. The rows show, from top to bottom, the responses of of overall debt volume (in percent), the volume of consumer credit (in percent), the responses of mortgage credit volume (in percent), and the share of debt-to-GDP (in percent).

Figure 16: Robustness checks: response of headline macro variables

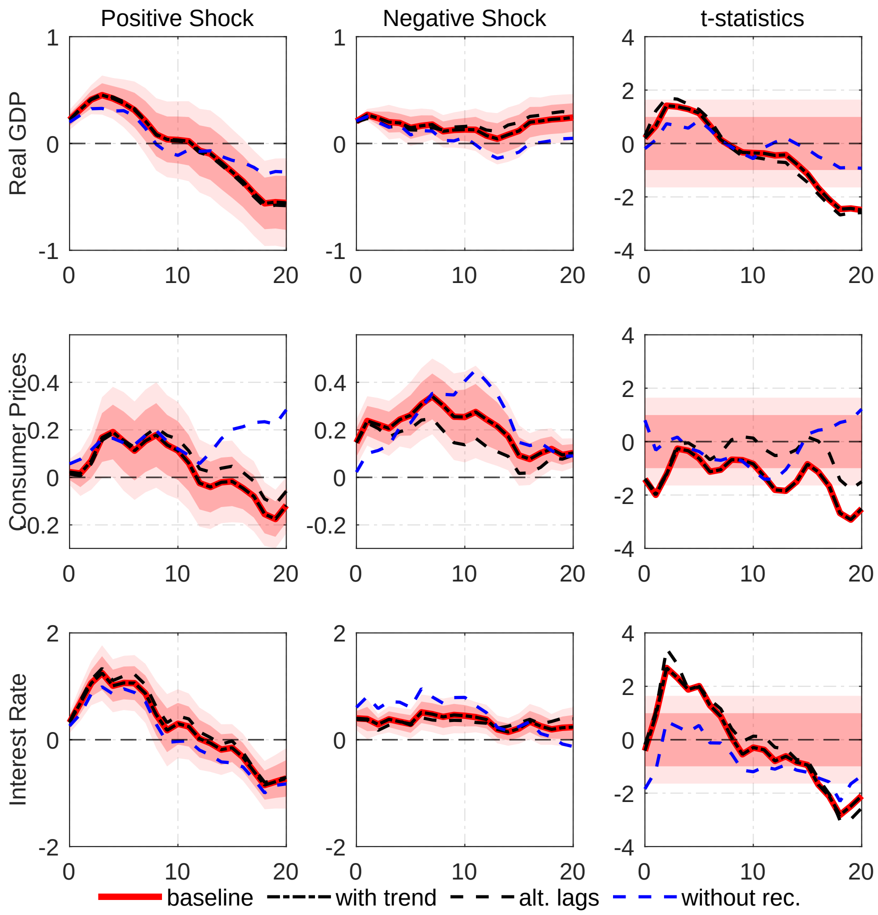

Notes: The first column shows the impulse response coefficients β_h^+ for $h = 0, ..., H$ for a positive (one standard deviation) credit supply shock from the baseline model (red-solid) alongside the impulse responses given alternative model specifications such as a model including a linear trend (black-dotted), a model with an alternative lag specification (black-dashed) as well as the results from an estimation of the baseline model, where the great recession has been excluded (blue-dashed). The second column shows the impulse response coefficients β_h^- for a negative (one standard deviation) credit supply shock. In both cases, the dark (pale) red-shaded area corresponds to the 68 (90) percent confidence interval from the baseline model, relying on Driscoll-Kraay standard errors. The third column shows the corresponding t–statistics testing the null that $H_0 : (\beta_h^+ - \beta_h^-) = 0$ for each horizon h using the Driscoll-Kraay method. The dark (pale) red-shaded area covers the t–critical values for a 68 (90) percent confidence interval, i.e. ± 0.995 (± 1.645), from the baseline model. The first row shows the responses of real GDP (in percent).The second row depicts the responses of consumer prices (in percent). The last row shows the responses of the effective federal funds rate (amended by the Wu-Xia shadow rate (in percentage points).

Figure 17: Robustness checks: supply side variables

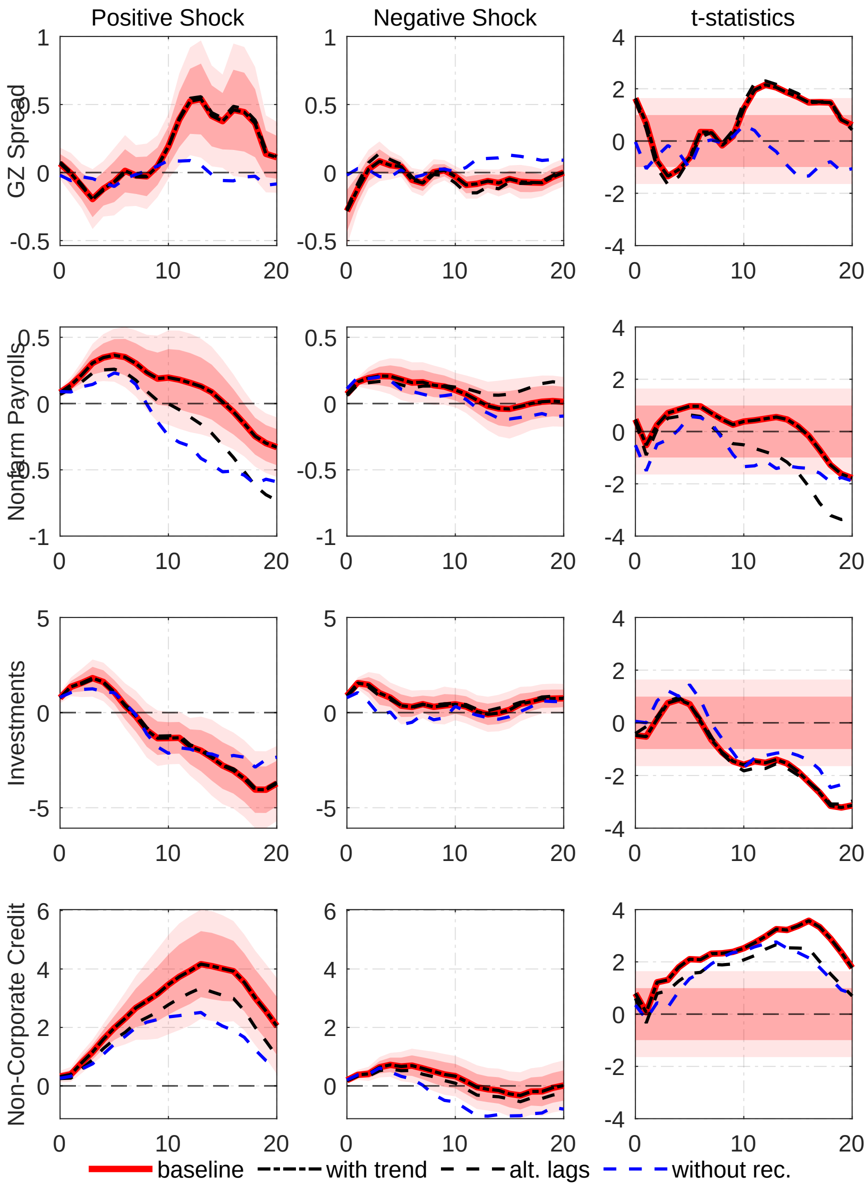

Notes: The first column shows the impulse response coefficients β_h^+ for $h = 0, ..., H$ for a positive (one standard deviation) credit supply shock from the baseline model (red-solid) alongside the impulse responses given alternative model specifications such as a model including a linear trend (black-dotted), a model with an alternative lag specification (black-dashed) as well as the results from an estimation of the baseline model, where the great recession has been excluded (blue-dashed). The second column shows the impulse response coefficients β_h^- for a negative (one standard deviation) credit supply shock. In both cases, the dark (pale) red-shaded area corresponds to the 68 (90) percent confidence interval from the baseline model, relying on Driscoll-Kraay standard errors. The third column shows the corresponding t−statistics testing the null that $H_0 : (\beta_h^+ - \beta_h^-) = 0$ for each horizon h using the Driscoll-Kraay method. The dark (pale) red-shaded area covers the t−critical values for a 68 (90) percent confidence interval, i.e. ±0.995 (±1.645), from the baseline model. The first row shows the responses of the Gilchrist and Zakrajšek (2012) spread (in percentage points).The second row depicts the responses of total non-farm payrolls (in percent). The last row shows the responses of real investments (in percent).

Figure 18: Robustness checks: demand side variables

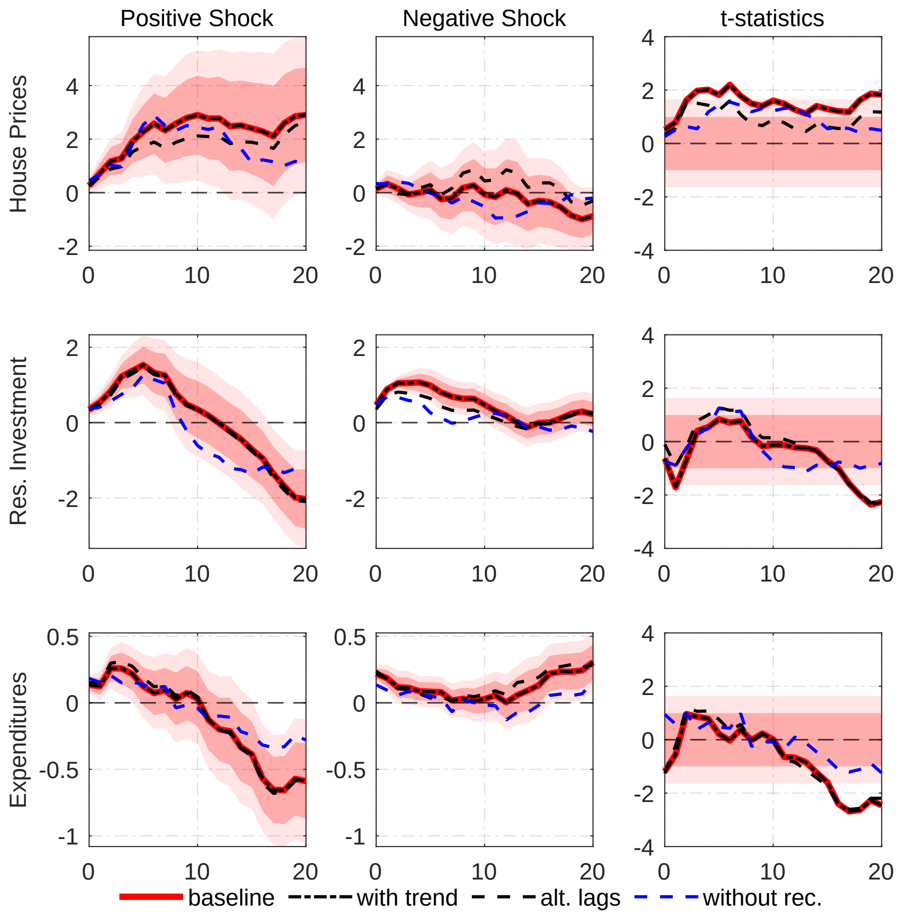

Notes: The first column shows the impulse response coefficients β_h^+ for $h = 0, ..., H$ for a positive (one standard deviation) credit supply shock from the baseline model (red-solid) alongside the impulse responses given alternative model specifications such as a model including a linear trend (black-dotted), a model with an alternative lag specification (black-dashed) as well as the results from an estimation of the baseline model, where the great recession has been excluded (blue-dashed). The second column shows the impulse response coefficients β_h^- for a negative (one standard deviation) credit supply shock. In both cases, the dark (pale) red-shaded area corresponds to the 68 (90) percent confidence interval from the baseline model, relying on Driscoll-Kraay standard errors. The third column shows the corresponding t−statistics testing the null that $H_0 : (\beta_h^+ - \beta_h^-) = 0$ for each horizon h using the Driscoll-Kraay method. The dark (pale) red-shaded area covers the t−critical values for a 68 (90) percent confidence interval, i.e. ± 0.995 (± 1.645), from the baseline model. The first row shows the responses of house prices (real Shiller index, in percent).The second row depicts real residential investments (in percent). The last row shows the responses of real personal expenditures (in percent).

B State-Dependent Distributions of Shocks

Figure 19: Estimated state-dependent PDFs and CDFs | z_t : output gap

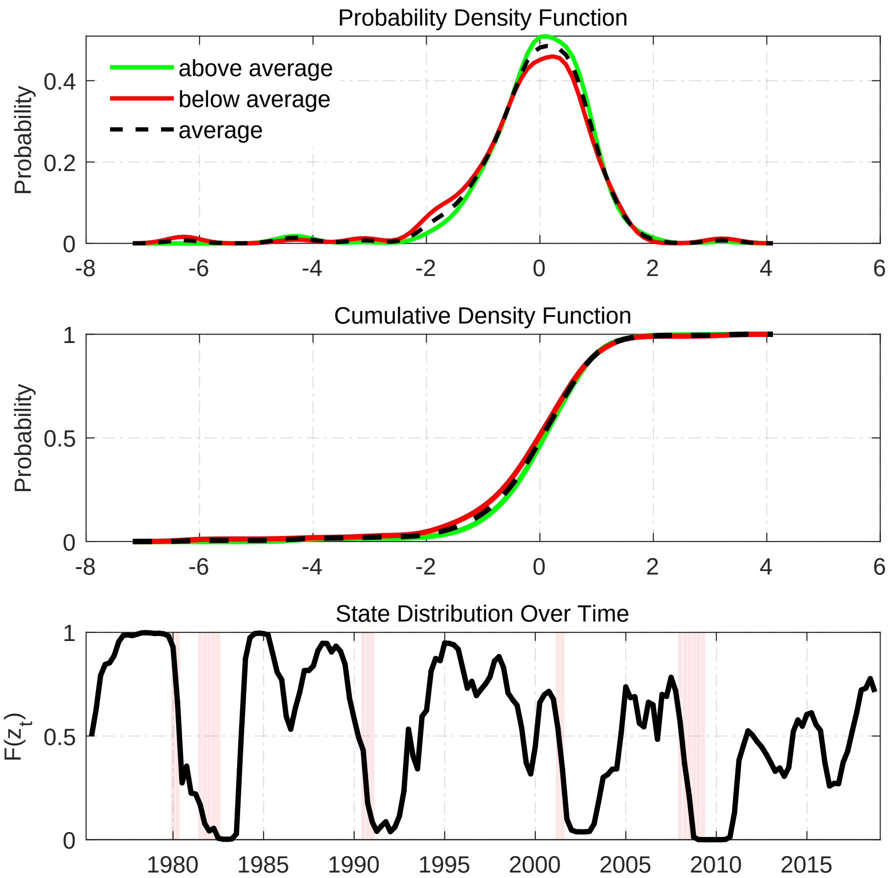

Notes: The state variable z_t is the output gap. The smoothed series relies on seven lags. The calibration is chosen to $\kappa = 3$ and $\mu = 0$. Red-shaded areas in the bottom panel show NBER recession dates.

Figure 20: Estimated state-dependent PDFs and CDFs | z_t : adjusted CFNFCI

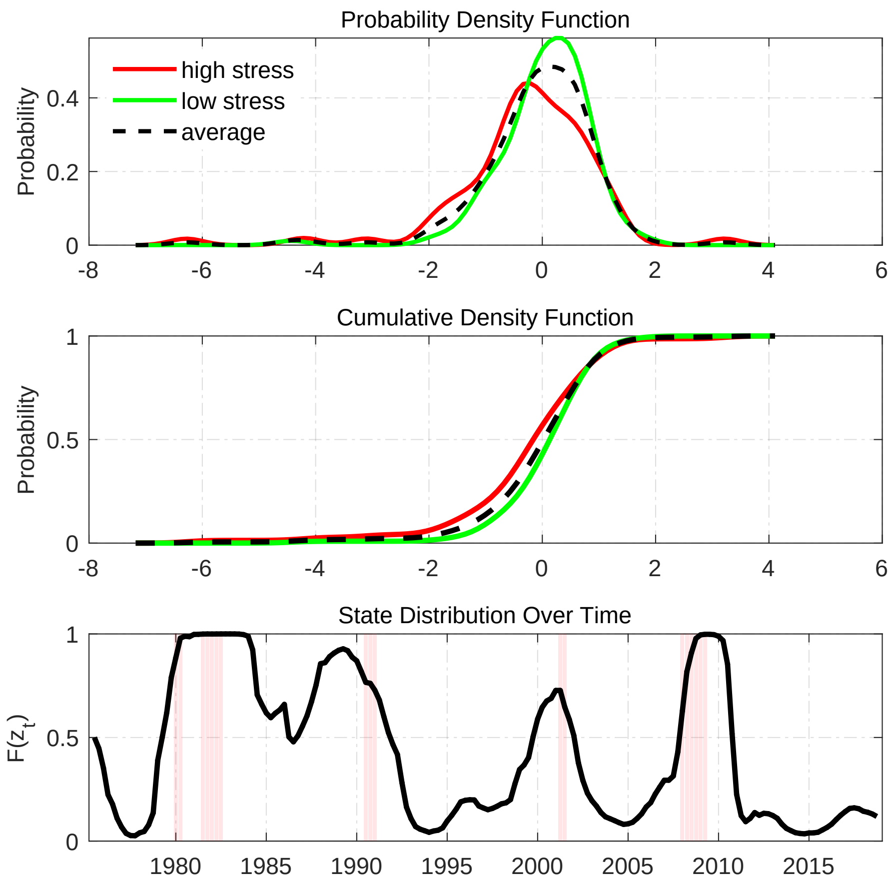

Notes: The state variable z_t is the adjusted Chicago Fed National Financial Conditions Index. The smoothed series relies on seven lags. The calibration is chosen to $\kappa = 3$ and $\mu = 0$. Red-shaded areas in the bottom panel show NBER recession dates.

Essay IV:

Optimal Monetary Policy Under Heterogeneous Beliefs

This is a revised version from 2022. The revised manuscript is currently under review at the *Scottish Journal of Political Economy*.

This paper was presented at the following workshops and international conferences:

I: 27th Annual Symposium of the Society for Nonlinear Dynamics and Econometrics (Dallas, USA)
Date: March 2019
Presenter: David Finck

II: 2nd Behavioral Macroeconomics Workshop (Bamberg, Germany)
Date: June 2019
Presenter: David Finck

III: 25th International Conference, Computing in Economics and Finance (Ottawa, Canada)
Date: June 2019
Presenter: David Finck

Acknowledgement

I thank Peter Tillmann for fruitful and helpful suggestions. I also thank Domenico Massaro, Emanuel Gasteiger, Karsten Kucharczyk, Dennis Finck, Joerg Schmidt and seminar participants from the University of Giessen for insightful comments. Conference participants at the 27th Annual Symposium of the Society for Nonlinear Dynamics and Econometrics, participants at the 2nd Behavioral Macroeconomics Workshop, and participants at the 25th International Conference, Computing in Economics and Finance provided valuable feedback.

Optimal Monetary Policy Under Heterogeneous Beliefs

David Finck*

Abstract

We use a New Keynesian model that features rational and non-rational households. Assuming that both the fraction of rational households and the expectations formation process are uncertain from the perspective of the central bank, we derive robust optimal discretionary monetary policy in a simple min-max framework where the central bank plays a zero-sum game versus a fictitious, malevolent evil agent. We show that the central bank is able to improve welfare if it accounts for uncertainty while the model is distorted. Even if the central bank accounts for the worst possible outcomes while the model is undistorted, the central bank can still reduce the welfare loss by implementing a more aggressive targeting rule that favorably affects the inflation-output stabilization trade-off.

Keywords: Heterogeneous Expectations, Robust Monetary Policy, Policy Implementation, Uncertainty

JEL classification: E52, D84

*University of Giessen, email: david.finck@wirtschaft.uni-giessen.de

I Introduction

Over recent decades, New Keynesian models have become the workhorse in both monetary policy and theory, acknowledged by researchers and policy makers alike. In its basic version (see, for instance, Woodford, 2003; Walsh, 2017; Galí, 2015), the assumptions for price flexibility as well as perfect competition among producers are relaxed which - in a nutshell- implies nonneutral effects for monetary policy. In fact, there is a broad consensus that the New Keynesian framework has been both an effective and intuitive tool as it provides guidance on expectations management which is an important aspect for stabilization policy.

However, the broad strand of literature on New Keynesian models reflects that some important aspects are missing in the standard framework. Among others, abandoning the strong assumption of rational expectations has proved to provide important evidence on business cycle theory. In fact, central banks increasingly care about the importance of the development of households' expectations. This is because on the one hand, households' inflation expectations signal future inflationary risks, while on the other hand, heterogeneous expectations are regarded as potential drivers of business cycle amplification (see, for instance, King, 2012; Mokhtarzadeh and Petersen, 2021).

Understanding the nature of expectations is particularly important at the Zero Lower Bound. To the extent central banks rely on the management of expectation through policies such as forward guidance, the effectiveness of these policies rests on the assumption that the central bank knows how agents receive and process information and generate forecasts about economic variables in the future (see, for instance, Andrade et al., 2019; Hommes and Lustenhouwer, 2019; Beqiraj et al., 2019).

The failure of the rational expectations framework is well documented in survey data and laboratory experiments alike, which is ultimately based on the criticism that agents (1) are assumed to have too much information available in general and (2) can process an unlimited amount of information in particular. In particular, rational agents must know the true structure and probability distribution of the economy to form rational expectations (see, for instance, Evans et al., 2001; and Carroll, 2003). To overcome this issue, Branch and McGough (2009), Di Bartolomeo et al. (2016), Gasteiger (2014) and Massaro (2013) incorporate a fraction of non-rational households into an otherwise standard New Keynesian framework and show that both determinacy and optimal monetary policy properties are strongly affected when non-rational households are incorporated.

However, it is important to stress that incorporating heterogeneous expectations into stochastic equilibrium models ultimately stems from empirical and experimental evidence in laboratories that provide strong evidence that a significant share of agents use simple heuristics and adaptive mechanisms for forecasting (see, for instance, Branch and McGough, 2016; and Hommes, 2011). Branch (2004) uses data from the Michigan surveys of consumers and finds that belief formation across participants is both heterogeneous on the one hand and dynamic on the other. Distinguishing between vector autoregression (VAR) based, adaptive and naive expectations when forming beliefs, he finds that respondents switch their predictor depending on the size of relative mean-squared errors. Pfajfar and Žakelj (2014) study the process of inflation expectation formation by focusing on both rational and adaptive agents and find that on the one hand, they cannot reject the hypothesis of rational expectations for at least 40% of the laboratory subjects, while on the other hand 20% of the subjects are best described by adaptive learning models. In another paper, Pfajfar and Žakelj (2018) explore the formation of inflation expectations within a New Keynesian framework. Similar to their previous paper, they find that about 40% of households form rational expectations, while 35% extrapolate expectations more than one-to-one into the future. 20% of households seem to employ adaptive learning models. Chavas (2000) investigates beef price equations under different expectations regimes and finds that nearly half of the market participants seem to form naive expectations.

Even though the recent literature well documents how optimal monetary policy is affected in models with heterogeneous expectations, all above mentioned authors mostly assume that central bankers do precisely know the fraction of non-rational households and also how expectations of these non-rational households are formed. It is however likely that these parameters (the fraction of non-rational households and the expectation formation process) are unknown, as reflected in different shares of adaptive agents in the empirical (experimental) literature. Therefore, a natural question is how optimal monetary policy is affected if the central bank is uncertain about the true distribution of rational and non-rational households.

This paper tries to fill this gap by modeling uncertainty in a New Keynesian framework with heterogeneous beliefs. We use a simple min-max mechanism where the central bank plays a zero-sum game against a fictitious, malevolent 'evil agent' that draws the potentially distorted parameters from a feasible set of models such that welfare loss is maximized. The application of a min-max framework is based on two reasons. First, if the central bank conducts discretionary monetary policy (as assumed in this paper), the specific targeting rule depends on both the fraction of non-rational households as well as the adaption parameter (as a part of the expectations formation process). In simple words, the trade-off of output gap and inflation stabilization depends on the possibly mispecified parameters. A central bank that is uncertain about the distribution is therefore able to shield the economy against perturbations by accounting for uncertainty. Second, a min-max approach has proven to be a simple and intuitive tool to model uncertainty (see, for instance, Giannoni, 2002; Giannoni, 2007; Hansen and Sargent, 1993; and Tillmann, 2009). Typical outcomes in min-max equilibria imply that the central bank is better-off accounting for uncertainty by implementing worst-case beliefs into its targeting rule, i.e. it can reduce welfare losses if it accounts for uncertainty while the worst-case beliefs turn out to be warranted (see, for instance, Giannoni, 2007; and Tillmann, 2009).

In this paper, we derive a robust optimal policy plan that accounts for uncertainty. Hereby, we refer to the same definition as in Giannoni (2007), i.e. those policy plans that perform best in a worst-case scenario, i.e. a scenario that delivers the worst possible welfare outcomes within a pre-specified set of parameter configurations.[1] We assume that the central bank does not know the true values of the unknowns, but it knows ranges (i.e. intervals) where the true values lie within. Calibrating our model to the US economy, we find that a central bank that accounts for uncertainty can substantially reduce welfare losses if it implements a robust policy plan. From the standpoint that the fraction of non-rational households is indeed unknown, our results are important insofar as they reveal that welfare losses can become large if the worst possible parameter perturbations turn out to be true. Interestingly, a central bank that implements the robust policy is able to further reduce welfare losses, relative to the case of certainty, i.e. when the true values are known. In other words, we find that the central bank can reduce the well known stabilization bias by incorporating a more 'aggressive' policy plan relative to the case of certainty. The reason for this finding is that discretionary monetary policy suffers from the stabilization bias as long as the supply shock is not white noise. That is, in the absence of commitment, inflation is too volatile. We find that adopting a robust policy approach can reduce the stabilization bias. A desire for robustness thus makes monetary policy more aggressive. As a result, inflation volatility is reduced and welfare improves compared to the rational expectations solution.

The rest of this paper is organized as follows: Section II sketches the underlying baseline model. In Section III, we introduce the min-max framework and derive the robust optimal policy plan and equilibrium outcomes. Section IV reports simulation exercises and discusses our most important results in detail. Section V argues that the implementation of robust optimal policy rules is preferable over a standard Taylor rule. Section VI concludes.

[1] Gasteiger (2021) also investigates expectations mismeasurement by comparing long-run losses for three different reference models. Therein, robust optimal monetary policy is defined as being a policy plan that yields determinacy across reference models for *given* determinants that characterize expectations heterogeneity. Di Bartolomeo et al. (2016) also consider a case where the central bank is being unable to recognize how households form their beliefs and explore a robust optimal policy consisting of minimizing the maximum regret of choosing a wrong share rational households. They find that implementing a wrong share of rational households always results in higher welfare losses, even though this effect is stronger if the central bank implements a belief that is greater than the true value.

II The Model

Our underlying model substantially relies on recent work of Branch and McGough (2009) and Gasteiger (2021) who incorporate heterogeneous expectations into a simple New Keynesian model. The economy is populated by a continuum of households, firms and the central bank. Time is discrete and indexed by $t \in (0, ..., \infty)$. The IS-curve (1) and the New Keynesian Phillips curve (2) can be summarized as

$$x_t = \widehat{E}_t x_{t+1} - \sigma^{-1}\left(i_t - \widehat{E}_t \pi_{t+1}\right) \tag{1}$$

$$\pi_t = \beta \widehat{E}_t \pi_{t+1} + \phi x_t + e_t, \tag{2}$$

where x_t is output gap, i_t the short-term interest rate, π_t inflation and e_t a cost push shock which follows an AR(1) process specified as

$$e_t = \rho e_{t-1} + \varepsilon_t^e, \sim \mathcal{N}(0, \sigma_e^2).$$

Similar to the recent literature, we assume that aggregate expectations are a linear combination of two types different types of agents, namely rational households (indexed $\mathcal{R}$) and non-rational households (indexed $\mathcal{B}$), such that aggregate expectations for a generic variable z_t read

$$\widehat{E}_t z_{t+1} = \alpha_{\mathcal{R}} E_t^{\mathcal{R}} z_{t+1} + \alpha_{\mathcal{B}} E_t^{\mathcal{B}} z_{t+1}.$$

As is common in the specific literature, we assume that rational agents have a one-step ahead perfect foresight on economic variables, i.e. for a generic variable z_t it holds that $E_t^{\mathcal{R}} z_{t+1} = E_t z_{t+1}$, where E_t corresponds to the rational expectations operator in t. However, rational households are supposed to be not "hyperrational" in a sense that they do not know the expectation formation process of non-rational households.[2] As regards the non-rational households, we assume that expectations of non-rational households are specified as $E_t^{\mathcal{B}} z_{t+1} = \gamma z_t$.[3] That is, we depart from the assumption as in Branch and McGough (2009) and assume that non-rational households, besides the variables they seek to maximize, can observe current outcomes. In simple words, tomorrow's expectations equal today's observation, adjusted for an adaption parameter γ which is equal for all outcomes.[4] For any variable z_t this implies that, when forecasting in period t what will prevail in $t + 1$, these households simply look at z_t and extrapolate it forward. The parameter γ therefore is an adaption parameter.

[2] This is an important axiomatic assumption. It implies that $E_t^{\mathcal{R}} E_t^{\mathcal{B}} z_{t+1} = E_t^{\mathcal{R}} z_{t+1}$, see Branch and McGough (2009) for a detailed discussion.

[3] Albeit this assumption is admittedly crude, it does in no way affect our assumptions on aggregate expectations. However, it is likely that many people are more sophisticated in their expectation formation process.

[4] We could also assume that different values for γ prevail for different aggregate variables, i.e. (in our case) that $\gamma_x \neq \gamma_\pi$.

$\gamma < 1$ implies that non-rational households place less weight on current outcomes while $\gamma = 1$ means that today's expectation of tomorrow's outcome is equal to the current observation. $\gamma > 1$ states that non-rational households are 'extrapolative' in a sense that they overweight today's observation by extrapolating it more than one-to-one into the future. It should be stressed that in our assumption, both households, i.e. rational and non-rational households, are able to solve the underlying model, but only a fraction α is able to form rational expectations while the remainder does not.

Since we focus on two types of households, it holds that $\alpha_{\mathcal{B}} = 1 - \alpha$, such that aggregate expectations read

$$\widehat{E}_t z_{t+1} = \alpha E_t z_{t+1} + (1-\alpha)\gamma z_t.$$

It follows that our IS-curve and New Keynesian Phillips curve read

$$\begin{aligned}
x_t &= \underbrace{[\alpha E_t x_{t+1} + (1-\alpha)\gamma x_t]}_{\widehat{E}_t x_{t+1}} - \sigma^{-1}\Big(i_t - \underbrace{[\alpha E_t \pi_{t+1} + (1-\alpha)\gamma \pi_t]}_{\widehat{E}_t \pi_{t+1}}\Big) \\
\pi_t &= \beta\,[\alpha E_t \pi_{t+1} + (1-\alpha)\gamma \pi_t] + \phi x_t + e_t.
\end{aligned}$$

Importantly, both the IS-curve as well as the New Keynesian Phillips curve are microfounded and nest the standard model in the absence of rational expectations. More precisely, our model collapses to the textbook model when $\alpha = 1$.

III Uncertainty about Heterogeneous Expectations

It is assumed that the central bank does neither know the share of rational households, nor the exact value of the adaption process, i.e. there is uncertainty about the true values of α and γ. However, the central bank knows that α and γ lie within an interval of a lower bound and an upper bound, respectively, i.e.

$$\alpha \in [\alpha_l,\ \alpha_h], \qquad \gamma \in [\gamma_l,\ \gamma_h],$$

where $\alpha_h > \alpha_l > 0$ and $\gamma_h > \gamma_l > 0$.[5] This assumption enables us to model uncertainty about α and γ in a simple min-max approach. More precisely, we assume the central bank plays a zero-sum game versus an evil, malevolent agent who chooses the vector $\psi = (\alpha,\ \gamma) \in \Psi$, where $\Psi = [\psi_1, ..., \psi_m]$ is a feasible compact set $\Psi \in \mathbb{R}^m$ satisfying

$$\Psi \equiv \{\psi = (\alpha, \gamma) \mid \alpha_l \leqslant \alpha \leqslant \alpha_h,\ \gamma_l \leqslant \gamma \leqslant \gamma_h\},$$

such that welfare loss is maximized.

The game consists of four stages. In the first stage, the evil agent and the cen-

[5] Even if the central bank does not know the true lower and upper bounds for α and γ, we can interpret these as the desire for robustness, i.e. a measure of caution.

tral bank as well as the entire private sector observe the cost push shock e_t. After observing the shock, in the second stage the central bank implements a policy plan $f(\psi) \in F$, where F denotes the compact set of all feasible non-inertial policy plans, i.e. $F = \{f(\psi) \mid f(\psi) \geqslant 0\}$ satisfying $F \in \mathbb{R}^n$, and designs optimal monetary policy in a discretionary fashion while still being uncertain about the true values of α and γ. Then, in the third stage, uncertainty is resolved, i.e. it turns out whether the model is distorted or not. Finally, the equilibrium outcomes are realized in the last (fourth) stage.

A. *The Robust Optimal Monetary Policy Plan*

As mentioned above, the central bank adopts a policy plan $f(\psi) \in F$ after observing the shock, but before $\psi \in \Psi$ is known. For simplicity, we assume that no commitment technology is available, such that the central bank conducts optimal monetary policy in a discretionary fashion by minimizing output and inflation volatility subject to the Phillips Curve. In particular, the central bank is assumed to be concerned about household's well-being by minimizing both output gap and inflation volatility due to monopolistic competition and sticky prices.[6] The objective function reads

$$L_t = E_t \sum_{k=0}^{\infty} \beta^k \frac{1}{2} \left[\pi_{t+k}^2 + \lambda x_{t+k}^2 \right],$$

where λ is the relative weight the monetary authority places on the output gap. Recall that the central bank is assumed to be concerned about potential parameter perturbations, that is in a min-max equilibrium it implements a policy plan $f^\star(\psi^\star)$ that takes the worst possible outcomes into account. Formally, the resulting robust policy plan $f^\star(\psi^\star)$ reads

$$f^\star(\psi^\star) = \arg\min_{f \in F} \left\{ \max_{\psi \in \Psi} E \left[\mathcal{L}_t(f(\psi)), \psi \right] \right\}.$$

[6]Notice that the loss function as used here is only model consistent as long as we focus on a standard utility function in the homogeneous expectations framework, i.e. in the absence of non-rational households. That is, the loss function we assume here is actually *ad hoc* in a sense that we ignore additional terms due to the presence of heterogeneous expectations. As shown by Di Bartolomeo et al. (2016), a model-consistent loss function in the presence of heterogeneous expectations and backward-looking agents requires two additional components, implying eight additions terms in the case of two types of expectation formation, one of them being due to higher price dispersion caused by adaptive agents. The second term is based on a dispersion in consumption across households types. See Gasteiger (2021) for a detailed discussion that justifies the application of an *ad hoc* loss function as opposed to the model-consistent loss function.

It can be shown that the specific targeting rule for the discretionary case is given as

$$\pi_t = \underbrace{-\frac{\lambda}{\phi}\left[1-\beta(1-\alpha)\gamma\right]}_{f(\psi)} x_t.$$

Assume that the central bank is faced with a positive cost-push shock which creates a trade-off between inflation and output stabilization objectives. The policymaker balances them by creating a negative output gap.

The malevolent fictitious evil agent will try to harm the central bank by maximizing the loss function with respect to the unknown parameters, i.e. α and γ. The maximization problem reads

$$\max_{\alpha,\,\gamma}\left\{\frac{f(\psi)^2+\lambda}{\left[f(\psi)\left[1-\beta\left(\alpha\rho+(1-\alpha)\gamma\right)\right]+\phi\right]^2}\right\}.$$

Proposition 1. If the central banker is uncertain about the true value of α and γ, a robust policy rule will incorporate $\alpha^\star = \alpha_l$ and $\gamma^\star = \gamma_h$.

Proof. (i) The first order condition with respect to α is given by

$$\frac{\partial \mathcal{L}_t}{\partial \alpha} = -\frac{2\beta f(\psi)\,(\gamma-\rho)}{\left[\phi+f(\psi)\left[1-\beta(\alpha\rho+(1-\alpha)\gamma)\right]\right]^3}\left(f(\psi)^2+\lambda\right)$$

The derivative is negative for the entire parameter space under consideration, i.e. $\frac{\partial \mathcal{L}_t}{\partial \alpha} < 0$. It follows that $\alpha^\star = \alpha_l$ maximizes $\mathcal{L}_t$.
(ii) The first order condition with respect to γ is given by

$$\frac{\partial \mathcal{L}_t}{\partial \gamma} = \frac{2\beta f(\psi)\,(1-\alpha)}{\left[\phi+f(\psi)\left[1-\beta(\alpha\rho+(1-\alpha)\gamma)\right]\right]^3}\left(f(\psi)^2+\lambda\right)$$

The derivative is positive for the entire parameter space under consideration, i.e. $\frac{\partial \mathcal{L}_t}{\partial \gamma} > 0$. It follows that $\gamma^\star = \gamma_h$ maximizes $\mathcal{L}_t$. The worst-case belief of the central bank is thus $\psi^\star = (\alpha^\star, \gamma^\star) = (\alpha_l, \gamma_h)$. ∎

In what follows, we mainly focus on three possible scenarios. In the first scenario, the worst-case scneario, the central banker implements the robust policy plan and the model is distorted, such that the realized outcomes for output, inflation and the short-term interest rate are given as $x_t(f^\star(\psi^\star), \psi^\star)$, $\pi_t(f^\star(\psi^\star), \psi^\star)$ and $i_t(f^\star(\psi^\star), \psi^\star)$. In the second equilibrium, the approximating equilibrium, the central bank implements $f^\star(\psi^\star)$ but the model turns out to be undistorted. The equilibrium outcomes are therefore given as $x_t(f^\star(\psi^\star),\psi)$, $\pi_t(f^\star(\psi^\star),\psi)$ and $i_t(f^\star(\psi^\star),\psi)$.[7]

Now that we have derived the robust optimal policy plan $f^\star(\psi^\star)$, we next derive the equilibrium outcomes for both the worst-case and the approximating outcomes. In equilibrium, inflation, output gap and the interest rate will be linear functions of the supply shock e_t. The worst-case equilibrium outcomes are thus given as $\pi_t(f^\star(\psi^\star),\psi^\star) = b_\pi^{\text{min-max}} e_t$, $x_t(f^\star(\psi^\star),\psi^\star) = b_x^{\text{min-max}} e_t$ and $i_t(f^\star(\psi^\star),\psi^\star) = b_i^{\text{min-max}} e_t$, where

$$\begin{aligned}
b_\pi^{\text{min-max}} &= \frac{f^\star(\psi^\star)}{\phi + f^\star(\psi^\star)\left[1-\beta(\alpha_l\rho + (1-\alpha_l)\gamma_h\right]} \\
b_x^{\text{min-max}} &= -\frac{1}{\phi + f^\star(\psi^\star)\left[1-\beta(\alpha_l\rho + (1-\alpha_l)\gamma_h\right]} \\
b_i^{\text{min-max}} &= \frac{\sigma - (f^\star(\psi^\star)+\sigma)\left[\alpha_l\rho + (1-\alpha_l)\gamma_h\right]}{\phi + f^\star(\psi^\star)\left[1-\beta(\alpha_l\rho + (1-\alpha_l)\gamma_h\right]}.
\end{aligned}$$

In the approximating equilibrium, the central bank still implements the robust optimal plan $f^\star(\psi^\star)$, where, however, the model turns out to be undistorted. The outcomes for the three endogenous variables are $\pi_t(f^\star(\psi^\star),\psi) = b_\pi^{\text{approx}} e_t$, $x_t(f^\star(\psi^\star),\psi) = b_x^{\text{approx}} e_t$ and $i_t(f^\star(\psi^\star),\psi) = b_i^{\text{approx}} e_t$, where the respective coefficients are

$$\begin{aligned}
b_\pi^{\text{approx}} &= \frac{f^\star(\psi^\star)}{\phi + f^\star(\psi^\star)\left[1-\beta(\alpha\rho + (1-\alpha)\gamma\right]} \\
b_x^{\text{approx}} &= -\frac{1}{\phi + f^\star(\psi^\star)\left[1-\beta(\alpha\rho + (1-\alpha)\gamma\right]} \\
b_i^{\text{approx}} &= \frac{\sigma - (f^\star(\psi^\star)+\sigma)\left[\alpha\rho + (1-\alpha)\gamma\right]}{\phi + f^\star(\psi^\star)\left[1-\beta(\alpha\rho + (1-\alpha)\gamma\right]}.
\end{aligned}$$

Notice that the only difference between the min-max and the approximating equilibrium are the *realized* outcomes for ψ. That is, in the worst-case equilibrium, $\psi^\star = (\alpha^\star,\gamma^\star) = (\alpha_l,\gamma_h)$ while in the approximating equilibrium it holds that $\psi = (\alpha,\gamma)$. In both cases however, the robust optimal policy plan $f^\star(\psi^\star)$ is implemented which is given as

$$f^\star(\psi^\star) = f^\star(\alpha_l,\gamma_h) = \frac{\lambda}{\phi}\left[1-\beta(1-\alpha_l)\gamma_h\right] \geqslant 0$$

and is non-negative by assumption.

At this point, we can see that the stabilization between the output gap and inflation clearly depends on the set of models Ψ. To better understand how the central bank shields the economy against uncertainty, notice that the derivative of the policy plan

[7] We will later investigate a scenario in which the central bank does not account for uncertainty while the model turns out to be distorted, i.e. the central bank implements $f(\psi)$ while $\psi^\star$ holds. The equilibrium outcomes therefore read $x_t(f(\psi),\psi^\star)$, $\pi_t(f(\psi),\psi^\star)$ and $i_t(f(\psi),\psi^\star)$.

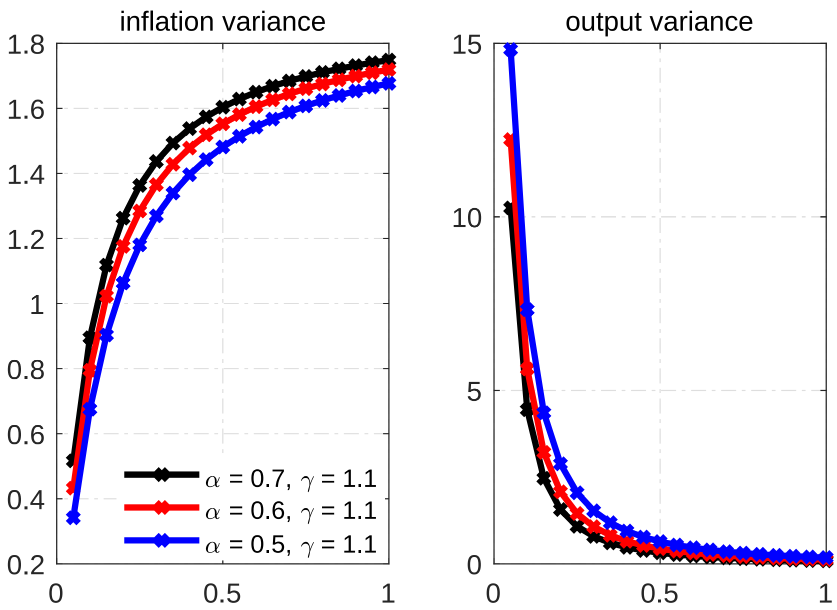

Figure 1: Inflation and output gap variance for $\lambda \in [0, 1]$ and different worst-case scenarios, in which the model is undistorted.

$f(\alpha, \gamma) = \frac{\lambda}{\phi}[1 - \beta(1-\alpha)\gamma]$ with respect to α and γ are given as

$$-\frac{\partial f(\psi(\alpha))}{\partial \alpha} = -\frac{\lambda}{\phi}\gamma\beta < 0, \qquad \frac{\partial f(\psi(\gamma))}{\partial \gamma} = -\frac{\lambda}{\phi}\beta(1-\alpha) < 0.$$

The derivates imply an important result for later purposes. If the lower bound of the interval $\alpha \in [\alpha_l,\ \alpha_h]$ increases, i.e. if the central bank considers a larger fraction of non-rational households and implements the corresponding robust policy plan, this implies that inflation can be stabilized with a smaller fall in the output gap. A similar argumentation holds for γ_h.

Confronting these results with equilibrium outcomes, Figure (1) plots inflation variances and output variances for different beliefs where the worst-case, however, does not occur, i.e. the worst-case beliefs are unwarranted.[8] More precisely, the black and the red line correspond to scenarios where the central bank implements $\alpha = 0.7$ (black line) and $\alpha = 0.6$ (red line), given $\gamma = 1.1$, whereas the blue line corresponds to the case where the central bank implements worst-case beliefs $\alpha_l = 0.5$ and $\gamma_h = 1.2$ for different values for λ between 0 and 1. Notably, the central bank can achieve a smaller inflation variance when values for α are implemented that are below the actual outcome. At the same time, however, the output gap variance increases if the desire for robustness increases, i.e. if the central bank widens its interval for α. This is an interesting result for the next section where it is shown in a simulation exercise that implementing the robust policy plan reduces welfare relative to the case of certainty, i.e. without parameter perturbations. Trivially, there will be an improvement in terms

[8]The plots in (1) rely on the calibration in Table (1), i.e. it is assumed that the 'true' values for α and γ are 0.8 and 1, respectively.

Parameter	β	ϕ	λ	σ	ρ	α_l	α_h	γ_l	γ_h
Value	0.99	0.15	0.25	1	0.35	[0.5; 0.7]	0.8	0.8	[1.1; 1.2]

Table 1: Baseline calibration

of lower welfare loss if the effect of lower inflation variance dominates the effect of higher output variance, which depends on λ, i.e. the weight in the loss function the central bank places on the output gap.

IV A Simulation Exercise

To investigate the consequences of uncertainty, we calibrate our model to the US economy and compare different scenarios for pairwise different sets of worst-case parameters $\psi^\star = (\alpha_l, \gamma_h)$.[9] More precisely, we compare inflation and output volatility as well as welfare losses for the min-max equilibrium, the approximating equilibrium and the case where the model is distorted but the central bank does not account for uncertainty, i.e. it does not implement the robust policy plan.

A. *Calibration*

We calibrate out model to the US economy with the parameters as in Table (1). The values for $\beta = 0.99$ and $\lambda = 0.25$ and $\sigma = 1$ (log utility) are standard in the literature. We follow Surico (2008) and set the slope of the Phillips curve to $\phi = 0.15$.[10] Setting $\rho = 0.35$ implies a medium persistent cost-push shock, as in Woodford (2003). The bounds for α_l are set to 0.5 and 0.7 which imply shares for non-rational households of 50% and 30%, respectively. These are even larger values the share of non-rational households than found in most of the literature and thus intended to represent a grave misspecification. $\gamma_h = 1.2$ and $\gamma_h = 1.1$ are common values for the investigation of heterogeneous beliefs in New Keynesian models (see, for example, Beqiraj et al., 2019; and Gasteiger, 2021) and intended to represent the upper bounds of the interval that is known to the central bank.

B. *Results and Discussion*

Table (2) reports the results for our simulation exercises for different pairs of α_l, γ_h and ρ. We assume that the share of non-rational households in the 'true model' is 20% which implies $\alpha = 0.8$. We further assume that non-rational households in the baseline model are naive in a sense that they forecast via a simple random walk rule, i.e. we assume that the true γ takes a value of one. Comparing the first row in a case

[9]Section A in the appendix investigates the determinacy properties of the model. We find that the Blanchard and Kahn (1980) condition is not satisfied, saying that our system of equations does not have a unique stationary solution. Hence, we cannot rule out that stationary sunspot equilibria exist.

[10]We also tried other values for ϕ, as for example in Christiano et al. (2005) who find $\phi = 0.2$. The results are not presented but available on request.

$\psi^\star\ (\alpha_l, \gamma_h)$	$\rho = 0.35$			$\rho = 0.0$		
	π^2	x^2	$\mathcal{L}$	π^2	x^2	$\mathcal{L}$
			no parameter uncertainty (case A)			
$\alpha = 0.8,\ \gamma = 1.0$	2.464	1.379	1.405	1.196	0.670	0.682
			robust policy, undistorted model (case B)			
$\alpha_l = 0.7,\ \gamma_h = 1.1$	2.306	1.832	1.382	1.420	0.907	0.684
$\alpha_l = 0.6,\ \gamma_h = 1.1$	2.136	2.414	1.370	1.082	1.223	0.694
$\alpha_l = 0.5,\ \gamma_h = 1.1$	1.916	3.325	1.374	1.001	1.737	0.717
$\alpha_l = 0.7,\ \gamma_h = 1.2$	2.264	1.967	1.378	1.127	0.980	0.686
$\alpha_l = 0.6,\ \gamma_h = 1.2$	2.063	2.696	1.368	1.055	1.379	0.700
$\alpha_l = 0.5,\ \gamma_h = 1.2$	1.795	3.919	1.387	0.954	2.084	0.738
			non-robust policy, distorted model (case C)			
$\alpha_l = 0.7,\ \gamma_h = 1.1$	3.392	1.898	1.933	1.621	0.907	0.924
$\alpha_l = 0.6,\ \gamma_h = 1.1$	4.552	2.548	2.594	2.184	1.223	1.245
$\alpha_l = 0.5,\ \gamma_h = 1.1$	6.426	3.597	3.663	3.103	1.737	1.768
$\alpha_l = 0.7,\ \gamma_h = 1.2$	3.796	2.125	2.163	1.751	0.980	0.998
$\alpha_l = 0.6,\ \gamma_h = 1.2$	5.431	3.039	3.095	2.464	1.379	1.405
$\alpha_l = 0.5,\ \gamma_h = 1.2$	8.403	4.703	4.789	3.724	2.084	2.122
			robust policy, distorted model (case D)			
$\alpha_l = 0.7,\ \gamma_h = 1.1$	3.139	2.493	1.881	1.536	1.219	0.920
$\alpha_l = 0.6,\ \gamma_h = 1.1$	3.756	4.245	2.409	1.908	2.157	1.224
$\alpha_l = 0.5,\ \gamma_h = 1.1$	4.343	7.536	3.114	2.345	4.068	1.681
$\alpha_l = 0.7,\ \gamma_h = 1.2$	3.418	2.971	2.080	1.629	1.416	0.992
$\alpha_l = 0.6,\ \gamma_h = 1.2$	4.192	5.480	2.781	2.063	2.696	1.368
$\alpha_l = 0.5,\ \gamma_h = 1.2$	4.843	10.576	3.743	2.538	5.543	1.962
			only rational households			
$\alpha = 1.0, \gamma = n.a.$	1.809	0.651	0.986	0.842	0.303	0.459

Table 2: Simulation results for different policy scenarios, see Table (1) for the calibration.

of no parameter uncertainty to the standard New Keynesian model in the absence of heterogeneous expectations, we can see that incorporating non-rational households always results in both higher inflation and output volatility as well as higher welfare losses.

In the first scenario (case B), the approximating equilibrium, we assume that the central bank implements the robust policy plan $f^\star(\psi^\star)$ while the model is undistorted, i.e. we report inflation variance $\pi^2(f^\star(\psi^\star), \psi)$, output variance $x^2(f^\star(\psi^\star), \psi)$ and the loss function $\mathcal{L}(f^\star(\psi^\star), \psi)$. In the second scenario (case C), we assume that the central bank does not implement the robust policy plan while the model is distorted, i.e. the equilbrium outcomes are $\pi^2(f(\psi), \psi^\star)$, $x^2(f(\psi), \psi^\star)$ and the loss function $\mathcal{L}(f(\psi), \psi^\star)$. In the last scenario (case D), the min-max equilibrium, the central bank implements the robust policy plan $f^\star(\psi^\star)$ while the model is distorted, i.e. the evil agent draws $\psi^\star$. The equilibrium outcomes in this case are $\pi^2(f^\star(\psi^\star), \psi^\star)$, $x^2(f^\star(\psi^\star), \psi^\star)$ and the loss function is $\mathcal{L}(f^\star(\psi^\star), \psi^\star)$.

A few things stand out. First, when accounting for uncertainty, the central bank

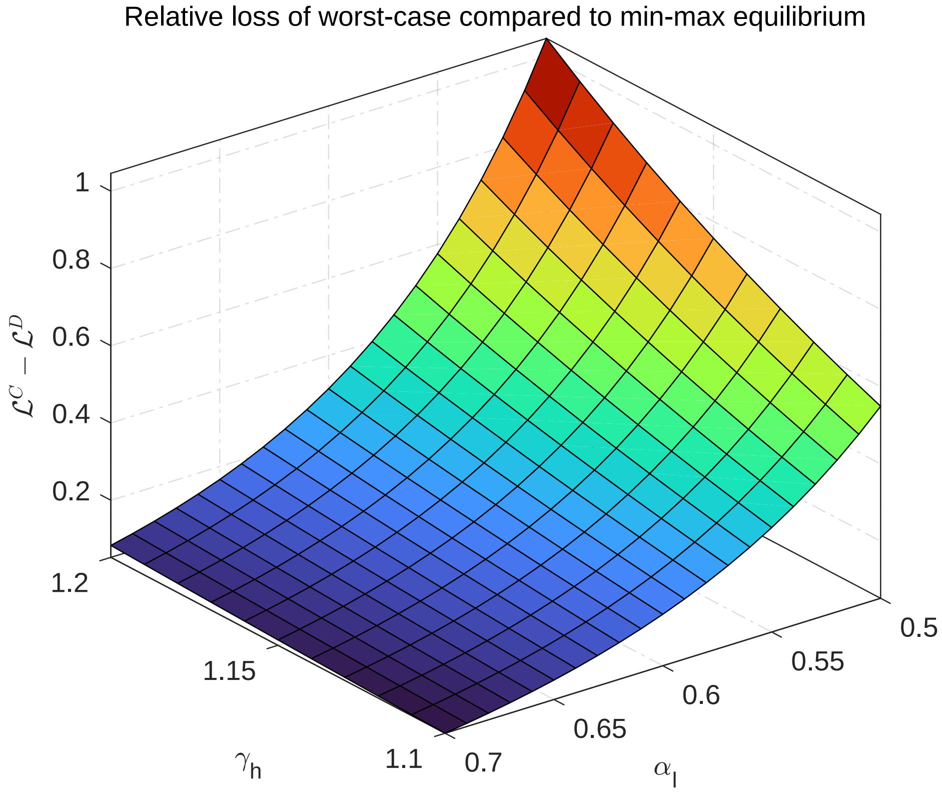

Figure 2: Difference between welfare when the central bank implements the robust policy plan and the model is and the case where there is no parameter uncertainty.

can substantially reduce welfare losses when the model is distorted relative to the case when the central bank ignores possible parameter perturbations. For instance, when the true share for non-rational households turns out to be 50% (i.e. $\alpha = 0.5$) and households are extrapolative (with $\gamma = 1.2$), inflation volatility is 8.403 and output gap volatility 4.703. If the central bank however incorporates the worst-case beliefs into the targeting rule, i.e. implements α_l and γ_h, inflation volatility is substantially reduced to 4.843 while output volatility is increased to 10.576. Importantly, the central bank does not place a weight of one on the output gap in the loss function, which results in an overall reduction of losses from 4.789 to 3.743. This is one key result in this paper and holds for all worst-case configurations under consideration, i.e. that lie within the set $\psi^\star \in \Psi$.[11] Figure (2) plots this difference for different values for α_l between 0.5 and 0.7 and for γ_h between 1.1 and 1.2. Not surprisingly, welfare always improves, meaning that the central bank can always achieve welfare gains when accounting for uncertainty, given the worst-case beliefs turn out to be warranted. Moreover, the improvement monotonically increases with respect to γ_h and decreases with respect to α_l. That is, the largest welfare improvement is achieved when the central bank is highly uncertain about the true values and these values are actually the true ones.[12]

[11] Note that this qualitative result, i.e. a lower welfare loss, is robust in the case where shocks are uncorrelated.

[12] Notice that Figure (2) also confirms that the min-max equilibrium is a global Nash-equilibrium. This

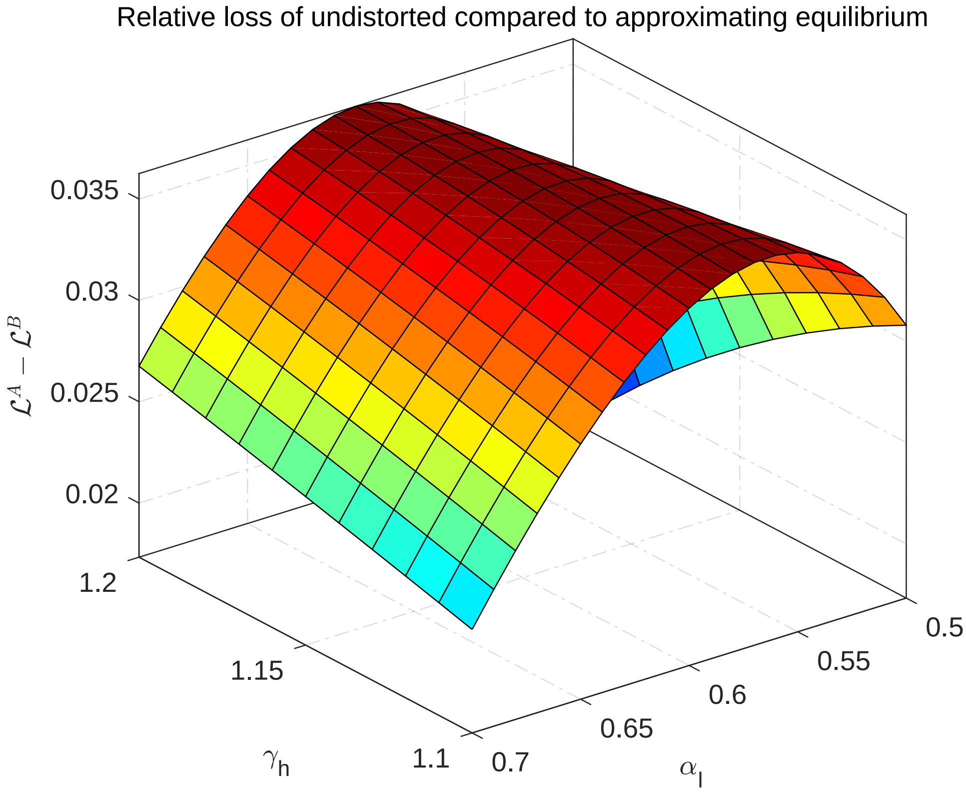

Figure 3: Difference welfare when the central bank implements the robust policy plan and the model is undistorted and the case where there is no parameter uncertainty.

Second, comparing the results for the case where the central bank incorporates the worst-case beliefs α_l and γ_h into the specific targeting-rule but the model is undistorted with the case of no parameter uncertainty at all, we can now confirm that there is an improvement in the stabilization between inflation and the output gap.

Interestingly, we observe the largest improvement in the case where we imposed the largest interval (that is known to the central bank) for the uncertain parameters. Notice that, identical to the former cases, this improvement is grounded on lower inflation volatility while output gap volatility is increased. If cost-push shocks are however uncorrelated, i.e. if $\rho = 0$, this improvement vanishes and welfare loss is higher than in the case of no parameter uncertainty. This result confirms our observation from Figure (1). Naturally, the question arises whether the central bank could not just implement $\alpha = 0$ and a γ that is greater than the worst-case belief to further exploit the improvement described above. Therefore, Figure (3) plots the loss difference between the approximating model and the model without uncertainty, i.e. $\mathcal{L}(f^\star(\psi^\star), \psi) - \mathcal{L}(f(\psi), \psi)$. As can be seen, there seems to exist an α for every given γ that maximizes the welfare improvement which (e.g. for $\gamma = 1.2$) is about $\alpha = 0.6$.[13]

Discretionary monetary policy suffers from the stabilization bias as long as the supply shock is not white noise. The reason is that in the absence of commitment, inflation is

means that there is no pair of $\widetilde{\alpha}$ and $\widetilde{\gamma}$ that violates $\mathcal{L}(f^\star(\psi^\star), \psi^\star) > \mathcal{L}(f^\star(\widetilde{\psi}), \widetilde{\psi})$.

[13] Widening the intervals to $\alpha_l = 0$, i.e. no rational agents exist, results in a negative value (not reported).

too volatile. However, we find that adopting a robust policy approach can reduce the stabilization bias. A desire for robustness makes monetary policy more aggressive. As a result, inflation volatility is reduced and welfare improves compared to the rational expectations solution.

Overall, we conclude that, given there exists a non-negative share of non-rational households, the central bank has an incentive to incorporate the worst-possible beliefs into its targeting rule, even if the model is undistorted. Figures (4) and (6) show the immediate response of the short-term interest rate when, in both cases, the central bank implements the worst-case beliefs, i.e. $f^\star(\psi^\star)$. Not surprisingly, the responses monotonically increase in the share of worst-case beliefs for the share of non-rational households (i.e. decrease in α_l) and the adaption parameter. However, we can see that in the case where the central bank implements the worst-case belief, the immediate interest rate response is systematically lower than in the case where parameter uncertainty is ignored. This is based on the degree of dispersion about the true values for α and γ and the worst-case beliefs α_l and γ_h, respectively. As a result, it turns out that in the worst-case equilibrium, the response of inflation is lower than in the case where the central bank does not take perturbations into account, as can be seen in Figures (7) and (8). At the same time, the opposite is true for the output gap, as the robust policy plan results in a larger contraction of output in order to stabilize the economy, see Figures (10) and (11).

V The Gain of Discretion over a Taylor Rule

Taylor (1993) argues that US monetary policy is well described as a simple interest-rate feedback rule that responds to inflation, output, or other economic conditions. It is well known that Taylor rules also mimic optimal discretionary policies, at least under certain circumstances, i.e. they may lead to similar dynamics as in the discretionary case. In this section, we argue that a central bank that accounts for expectations heterogeneity in general and uncertainty about expectations heterogeneity in particular, as implemented in the previous section, is able to achieve gains in terms of lower welfare costs relative to the case of an implemented Taylor rule. We follow Di Bartolomeo et al. (2016) and focus on a feedback rule that responds to current inflation π_t and current output gap x_t, i.e. our reference Taylor rule takes the form

$$i_t = \vartheta_\pi \pi_t + \vartheta_x x_t,$$

where $\vartheta_\pi = 1.5$ and $\vartheta_x = 0.125$. That is, the Taylor principle, stating that the interest-rate should respond by more than one-to-one, is satisfied. Since we have no backward-looking variables in our model, it is easy to derive the equilibrium outcomes for x_t and

$\psi^\star\,(\alpha_l,\gamma_h)$	$\rho = 0.35$			$\rho = 0.0$		
	π^2	x^2	$\mathcal{L}$	π^2	x^2	$\mathcal{L}$
		no parameter uncertainty				
$\alpha = 0.8,\ \gamma = 1$	1.722	4.307	2.799	0.975	1.925	1.456
		Taylor rule, distorted model				
$\alpha_l = 0.7,\ \gamma_h = 1.1$	2.144	6.606	3.659	1.251	2.710	1.928
$\alpha_l = 0.6,\ \gamma_h = 1.1$	2.561	8.200	4.611	1.576	3.774	2.520
$\alpha_l = 0.5,\ \gamma_h = 1.1$	3.047	11.438	5.907	2.022	5.518	3.401
$\alpha_l = 0.7,\ \gamma_h = 1.2$	2.301	6.817	4.005	1.330	2.953	2.068
$\alpha_l = 0.6,\ \gamma_h = 1.2$	2.813	9.754	5.252	1.722	4.307	2.799
$\alpha_l = 0.5,\ \gamma_h = 1.2$	3.386	14.528	7.018	2.274	6.683	3.945
		only rational households				
$\alpha = 1, \gamma = n.a.$	1.303	2.869	2.020	0.694	1.235	1.003

Table 3: Simulation results for scenarios if a Taylor rule with $\vartheta_\pi = 1.5$ and $\vartheta_x = 0.125$ is implemented.

π_t in this case and thus derive output and inflation variance

$$\mathbf{A}\begin{bmatrix} x_t \\ \pi_t \end{bmatrix} = \mathbf{B}\begin{bmatrix} E_t x_{t+1} \\ E_t \pi_{t+1} \end{bmatrix} + \mathbf{D}e_t,$$

where **A** and **B** contain appropriate parameters and **D** is a vector of zeros and ones.[14]

Table (3) reports simulation results with the same calibration as in Table (2). Comparing the results with those in the previous section, it stands out that welfare losses are higher in all scenarios under consideration. That is, if there is no parameter uncertainty at all, in the absence of non-rational households as well as in the case where the evil agent draws the worst-case parameter perturbations while the central bank implements the Taylor rule. Even though inflation volatility is higher than under discretion, the overall effect of lower welfare losses ultimately stems from higher output gap volatility. Furthermore, if the cost-push shock is serially uncorrelated, the central bank can achieve substantially gains with an optimal policy plan, independent of whether the policy maker implements the robust policy plan or not. For example, if the true value for α and γ turn out to be 0.5 and 1.1 respectively, while shocks are correlated with $\rho = 0.35$, inflation variance is 3.047 and output variance 11.438 which results in a loss of 5.907. However, if the robust optimal policy plan $f^\star(\psi^\star)$ is implemented and the central bank shields the economy against uncertainty, inflation variance is 4.343, output variance is 7.536, which results in a loss of 3.114. Even if the central bank ignores possible parameter perturbations and implements the discretionary policy plan, our simulation exercise implies that welfare losses are always lower relative to the case where the Taylor rule is implemented. Summing up,

[14]Here, $\mathbf{A} = \begin{bmatrix} 1-(1-\alpha)\gamma + \sigma^{-1}\vartheta_x & \sigma^{-1}\varphi_\pi - \sigma^{-1}(1-\alpha)\gamma \\ -\phi & 1-\beta(1-\alpha)\gamma \end{bmatrix}$, $\mathbf{B} = \begin{bmatrix} \alpha & \sigma^{-1}\alpha \\ 0 & \beta\alpha \end{bmatrix}$ and $\mathbf{D} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$.

we conclude that if the central bank has the opportunity to conduct a discretionary policy plan rather than following a Taylor rule, it should do so because of two reasons. First, a Taylor rule ignores expectations heterogeneity at all, that is, it only responds to changes in inflation and output while not accounting for different expectation formation processes. The policy plan (i.e. the specific targeting rule) under discretion however was shown to account for non-rational households as it implements both α and γ into the plan f. Second, and more importantly, this argument is even stronger if the central bank has a desire for robustness, i.e. if the central banker wants to account for uncertainty.

VI Conclusion

This paper uses a simple New Keynesian model that features fixed shares of rational and non-rational households, mainly as in Branch and McGough (2009), Gasteiger (2014), Gasteiger (2021) and Di Bartolomeo et al. (2016). Since laboratory experiments and expectations surveys alike imply that (i) a non-negative share of households does not form rational expectations and (ii) the share of non-rational households substantially varies across experiments/surveys, we implement model uncertainty about the determinants of heterogeneity across expectations. Given that both the fractions of rational and non-rational households as well as the expectation formation process is uncertain to the central bank, we derive robust optimal plans in a simple min-max equilibrium as pioneered in Giannoni (2002) and Giannoni (2007). Even though we implemented a very simple way to model adaptive expectations, our results have important implications. First, a central bank should account for uncertainty if it does not know the distributions of different households that are characterized by different expectation formation processes. In particular, in a simulation exercise calibrated to the US economy, we show that a central bank that accounts for uncertainty can *substantially* reduce welfare losses relative to the case where both heterogeneous expectations in general and uncertainty about the distribution of both household types in particular are uncertain.

Second, a central bank that seeks to implement a notable desire for robustness is better-off when it implements the worst-case beliefs, even if these beliefs turn out to be unwarranted. This is based on the fact that a policy rule (i.e. the targeting rule in the discretionary fashion) that incorporates the worst-case beliefs favorably affects the trade-off between output gap and inflation stabilization.

Finally, it has to be mentioned that a model that implements a more sophisticated learning process would probably better map the outcomes that are observable in laboratory experiments (see, for instance, Pfajfar and Žakelj, 2018). However, even though we abstract from backward-looking adaptive expectations formation processes, our model allows the application of a simple min-max approach that is shown to give fruitful structural insights into the mechanism behind robust optimal monetary policy when the strong paradigm of purely rational expectations is abandoned.

References

Andrade, Philippe, Gaetano Gaballo, Eric Mengus, and Benoit Mojon, "Forward Guidance and Heterogeneous Beliefs," *American Economic Journal: Macroeconomics*, 2019, *11* (3), 1–29.

Beqiraj, Elton, Giovanni Di Bartolomeo, and Marco Di Pietro, "Beliefs Formation and the Puzzle of Forward Guidance Power," *Journal of Macroeconomics*, 2019, *60*, 20–32.

Blanchard, Olivier Jean and Charles M Kahn, "The Solution of Linear Difference Models under Rational Expectations," *Econometrica*, 1980, *48* (5), 1305–1311.

Branch, William, "The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations," *The Economic Journal*, 2004, *114* (497), 592–621.

__ **and Bruce McGough**, "A New Keynesian Model with Heterogeneous Expectations," *Journal of Economic Dynamics and Control*, 2009, *33* (5), 1036–1051.

__ **and** __ , "Heterogeneous Beliefs and Trading Inefficiencies," *Journal of Economic Theory*, 2016, *163*, 786–818.

Carroll, Christopher D, "Macroeconomic Expectations of Households and Professional Forecasters," *The Quarterly Journal of Economics*, 2003, *118* (1), 269–298.

Chavas, Jean-Paul, "On Information and Market Dynamics: The Case of the US Beef Market," *Journal of Economic Dynamics and Control*, 2000, *24* (5-7), 833–853.

Christiano, Lawrence J, Martin Eichenbaum, and Charles L Evans, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 2005, *113* (1), 1–45.

Di Bartolomeo, Giovanni, Marco Di Pietro, and Bianca Giannini, "Optimal Monetary Policy in a New Keynesian Model With Heterogeneous Expectations," *Journal of Economic Dynamics and Control*, 2016, *73*, 373–387.

Evans, George W, Seppo Honkapohja, and Ramon Marimon, "Convergence in Monetary Inflation Models with Heterogeneous Learning Rules," *Macroeconomic Dynamics*, 2001, *5* (1), 1–31.

Galí, Jordi, *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and its Applications*, Princeton University Press, 2015.

Gasteiger, Emanuel, "Heterogeneous Expectations, Optimal Monetary Policy, and the Merit of Policy Inertia," *Journal of Money, Credit and Banking*, 2014, *46* (7), 1535–1554.

__ , "Optimal Constrained Interest-Rate Rules Under Heterogeneous Expectations," *Journal of Economic Behavior & Organization*, 2021, *190*, 287–325.

Giannoni, Marc P, "Does Model Uncertainty Justify Caution? Robust Optimal Monetary Policy in a Forward-Looking Model," *Macroeconomic Dynamics*, 2002, *6* (1), 111–144.

__ , "Robust Optimal Monetary Policy in a Forward-Looking Model With Parameter and Shock Uncertainty," *Journal of Applied Econometrics*, 2007, *22* (1), 179–213.

Hansen, Lars Peter and Thomas J Sargent, "Seasonality and Approximation Errors in Rational Expectations Models," *Journal of Econometrics*, 1993, *55* (1-2), 21–55.

Hommes, Cars, "The Heterogeneous Expectations Hypothesis: Some Evidence from the Lab," *Journal of Economic Dynamics and Control*, 2011, *35* (1), 1–24.

__ **and Joep Lustenhouwer**, "Managing Unanchored, Heterogeneous Expectations and Liquidity Traps," *Journal of Economic Dynamics and Control*, 2019, *101*, 1–16.

King, Mervyn, "Twenty Years of Inflation Targeting," *Stamp Memorial Lecture, London School of Economics, London, October 9*, 2012.

Massaro, Domenico, "Heterogeneous Expectations in Monetary DSGE Models," *Journal of Economic Dynamics and Control*, 2013, *37* (3), 680–692.

Mokhtarzadeh, Fatemeh and Luba Petersen, "Coordinating Expectations Through Central Bank Projections," *Experimental Economics*, 2021, *24* (3), 883–918.

Pfajfar, Damjan and Blaž Žakelj, "Experimental Evidence on Inflation Expectation Formation," *Journal of Economic Dynamics and Control*, 2014, *44*, 147–168.

__ **and** __ , "Inflation Expectations and Monetary Policy Design: Evidence From the Laboratory," *Macroeconomic Dynamics*, 2018, *22* (4), 1035–1075.

Surico, Paolo, "The Cost Channel of Monetary Policy and Indeterminacy," *Macroeconomic Dynamics*, 2008, *12* (5), 724–735.

Taylor, John B, "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy*, 1993, *39*, 195–214.

Tillmann, Peter, "Optimal Monetary Policy With an Uncertain Cost Channel," *Journal of Money, Credit and Banking*, 2009, *41* (5), 885–906.

Walsh, Carl E, *Monetary Theory and Policy*, MIT Press, 2017.

Woodford, Michael, *Interest and Prices*, Princeton University Press, 2003.

Appendix

A Determinacy Properties

In order to establish determinacy properties of our model in the case where the central bank conducts optimal (robust) monetary policy, first rewrite the model in matrix form as functions of our forward-looking variables. To eliminate the short-term interest rate, recall that, in equilibrium, the interest rate is a function of the deep structural parameters and exogenous driving forces, i.e. $i_t = b_i e_t$. Rewrite the model as

$$\begin{bmatrix} \alpha & \sigma^{-1}\alpha \\ 0 & \beta\alpha \end{bmatrix} \begin{bmatrix} E_t x_{t+1} \\ E_t \pi_{t+1} \end{bmatrix} = \begin{bmatrix} 1-(1-\alpha)\gamma & -\sigma^{-1}(1-\alpha)\gamma \\ -\phi & 1-\beta(1-\alpha)\gamma \end{bmatrix} \begin{bmatrix} x_t \\ \pi_t \end{bmatrix} + \begin{bmatrix} \sigma^{-1} b_i \\ -1 \end{bmatrix} e_t$$

Define $\mathbf{\Omega} = \begin{bmatrix} \alpha & \sigma^{-1}\alpha \\ 0 & \beta\alpha \end{bmatrix}^{-1}$, the model can, without loss of generality, be rewritten as

$$\begin{aligned}
\begin{bmatrix} E_t x_{t+1} \\ E_t \pi_{t+1} \end{bmatrix} &= \mathbf{\Omega} \begin{bmatrix} 1-(1-\alpha)\gamma & -\sigma^{-1}(1-\alpha)\gamma \\ -\phi & 1-\beta(1-\alpha)\gamma \end{bmatrix} \begin{bmatrix} x_t \\ \pi_t \end{bmatrix} + \mathbf{\Omega} \begin{bmatrix} \sigma^{-1} b_i \\ -1 \end{bmatrix} e_t \\
\begin{bmatrix} E_t x_{t+1} \\ E_t \pi_{t+1} \end{bmatrix} &= \mathbf{A} \begin{bmatrix} x_t \\ \pi_t \end{bmatrix} + \mathbf{\Omega} \begin{bmatrix} \sigma^{-1} b_i \\ -1 \end{bmatrix} e_t,
\end{aligned}$$

where

$$\mathbf{A} = \mathbf{\Omega} \begin{bmatrix} 1-(1-\alpha)\gamma & -\sigma^{-1}(1-\alpha)\gamma \\ -\phi & 1-\beta(1-\alpha)\gamma \end{bmatrix}.$$

Blanchard and Kahn (1980) show that a system like the one above has a unique stationary solution if and only if the number of eigenvalues of $\mathbf{A}$ outside the unit circle is equal to the number of forward-looking variables, which is two in our case (x_{t+1} and π_{t+1}). It can be shown that the vector $\mathbf{z}$ of eigenvalues of $\mathbf{A}$ is given by

$$\mathbf{z} = \begin{bmatrix} \dfrac{\left(\phi + \sigma(1+\beta) - \sqrt{(\beta\sigma)^2 + 2\beta\sigma(\phi-\sigma) + \phi(\phi+2\sigma) + \sigma^2} - 2\beta\sigma\gamma(1-\alpha)\right)}{2\alpha\beta\sigma} \\ \dfrac{\left(\phi + \sigma(1+\beta) + \sqrt{(\beta\sigma)^2 + 2\beta\sigma(\phi-\sigma) + \phi(\phi+2\sigma) + \sigma^2} - 2\beta\sigma\gamma(1-\alpha)\right)}{2\alpha\beta\sigma} \end{bmatrix}.$$

It is easy to see that the only difference between both eigenvalues is the opposite sign before the square root expression. It turns out that for our entire parameter space, the first eigenvalue is always below one while the second is greater than one. Since the Blanchard and Kahn (1980) condition is not satisfied, our system of equations does not have a unique stationary solution, i.e. it is possible that stationary sunspot equilibria exist.

Notice that the same solution prevails across all different scenarios that we investigate in the main part of the paper. This is because the solutions for all three scenarios, when written as above, only affect b_i, i.e. the vector of eigenvalues $\mathbf{z}$ is unaffected by any policy rule $f(\psi(\alpha,\gamma)) \in F$.

B Additional Figures

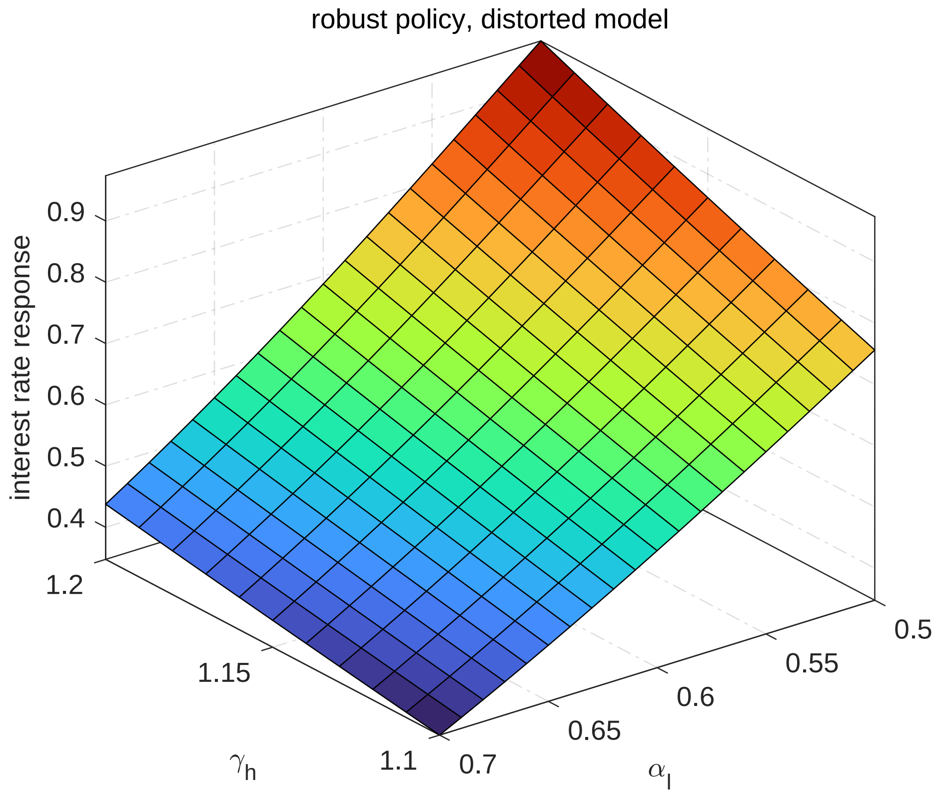

Figure 4: Interest-rate response for different worst-case combinations of α_l and γ_l.

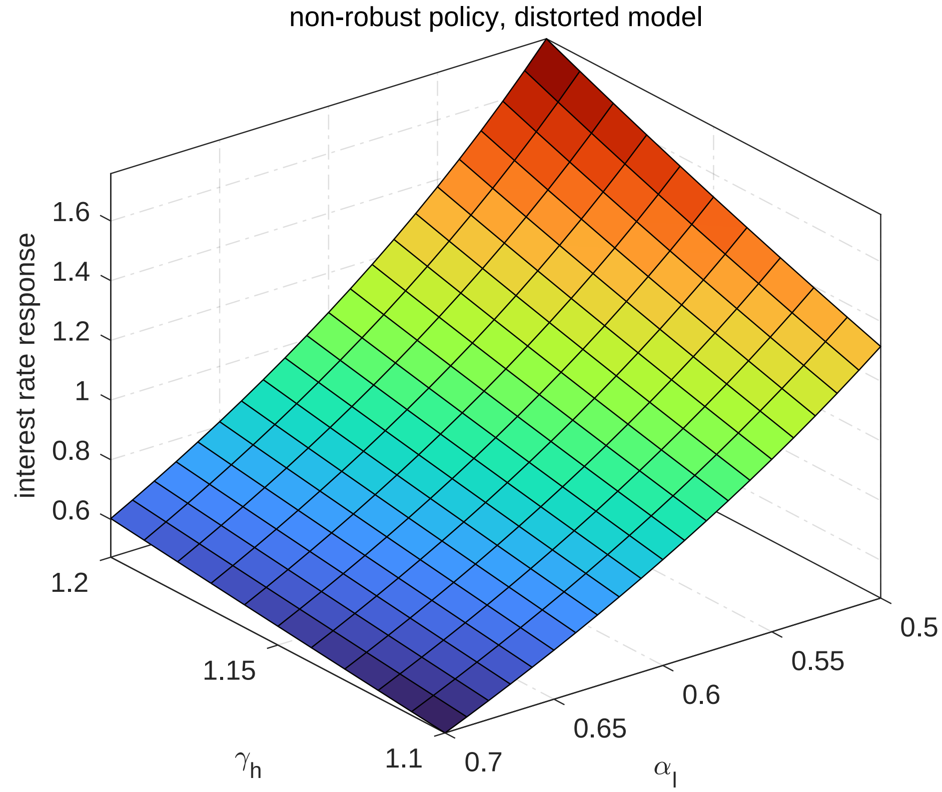

Figure 5: Interest-rate response for different worst-case combinations of α_l and γ_l.

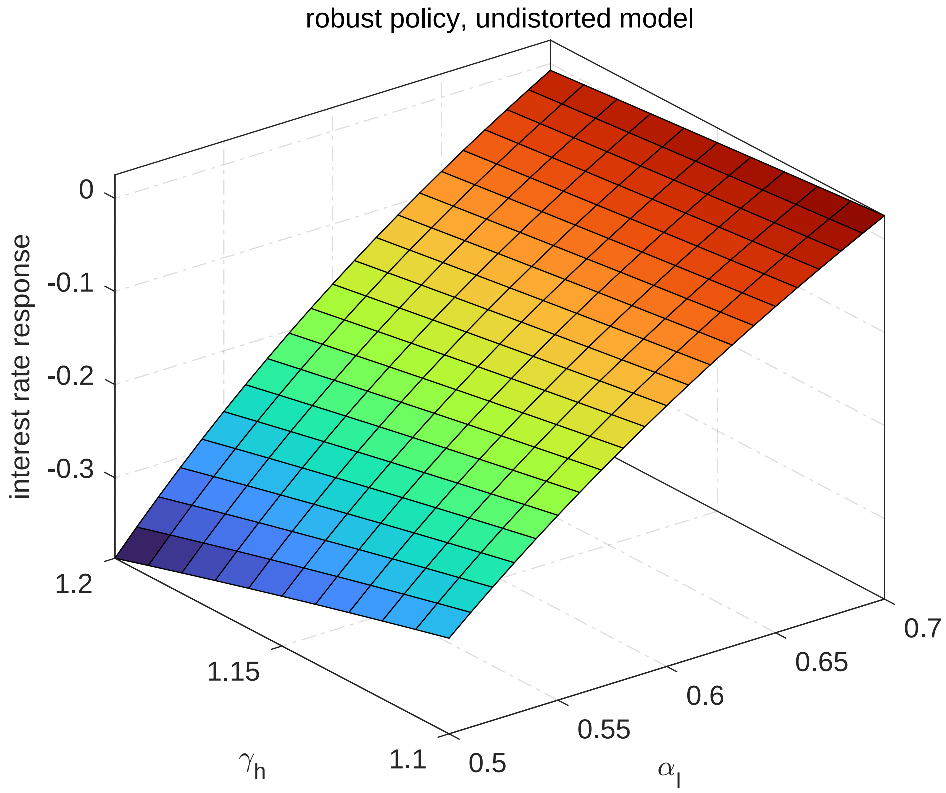

Figure 6: Interest-rate response for different worst-case combinations of α_l and γ_l.

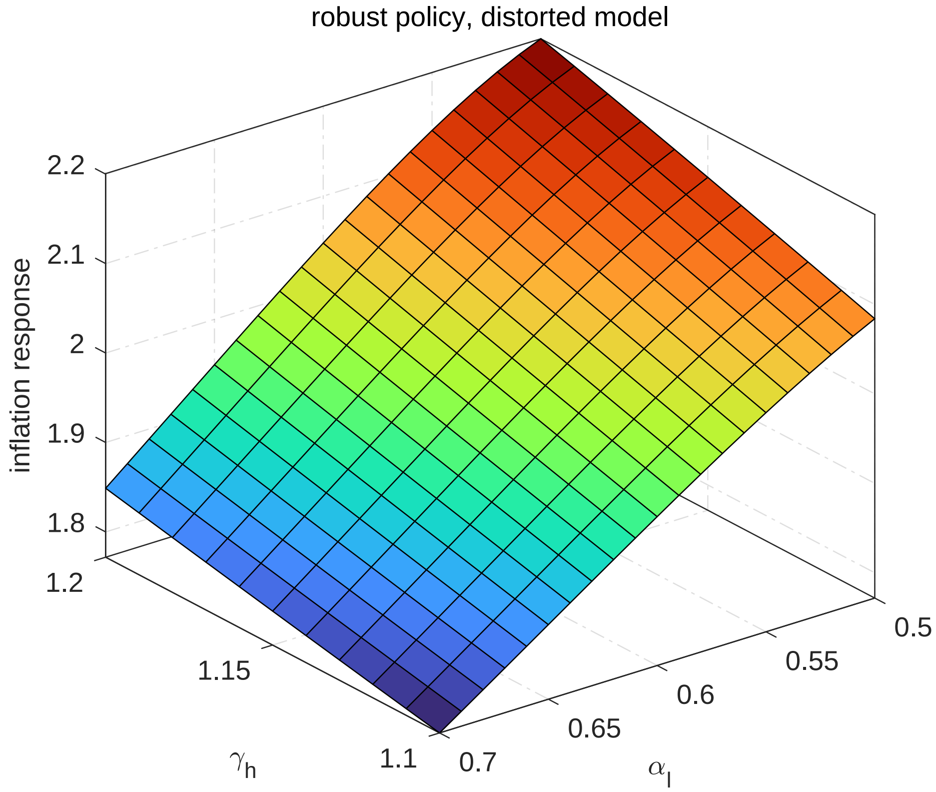

Figure 7: Inflation response for different worst-case combinations of α_l and γ_l.

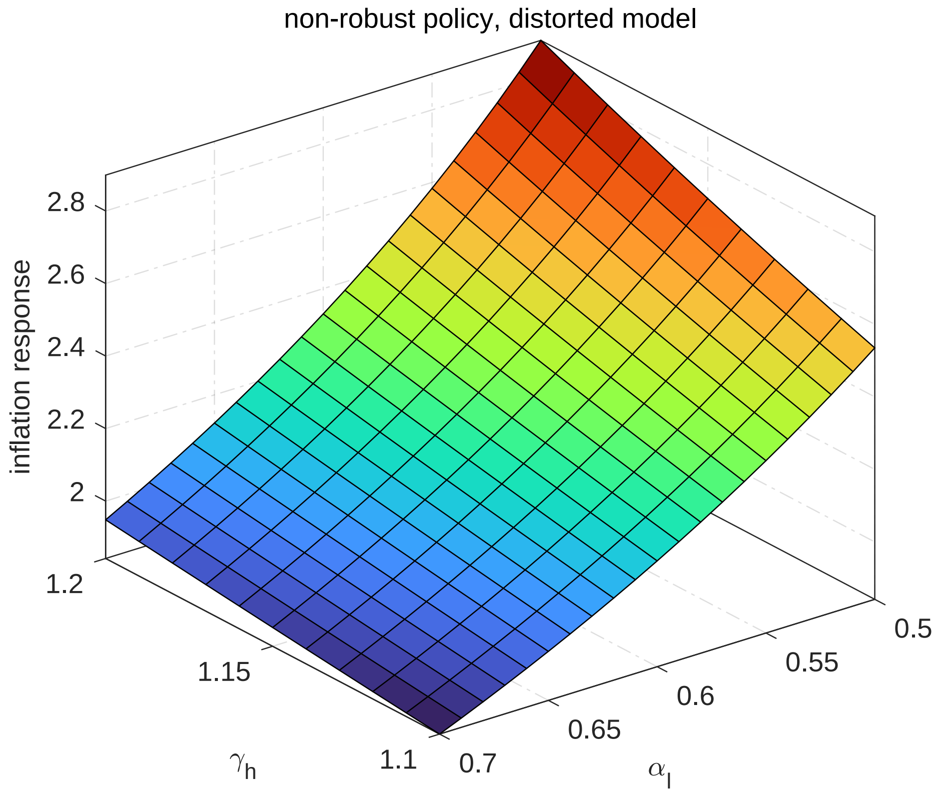

Figure 8: Inflation response for different worst-case combinations of α_l and γ_l.

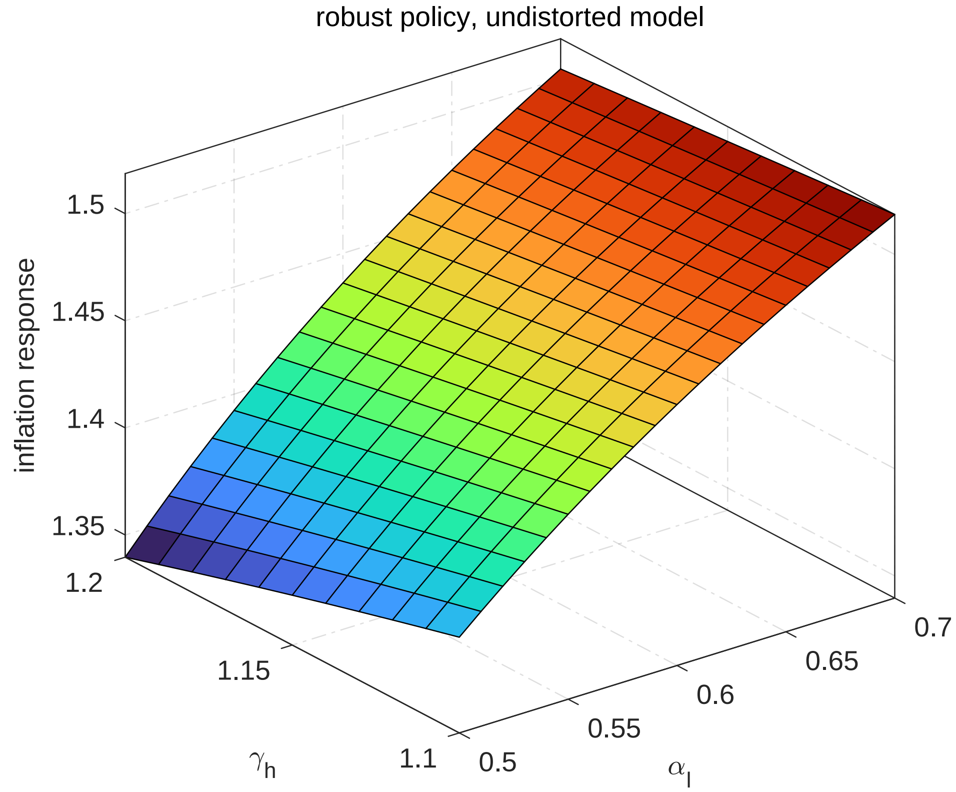

Figure 9: Inflation response for different worst-case combinations of α_l and γ_l.

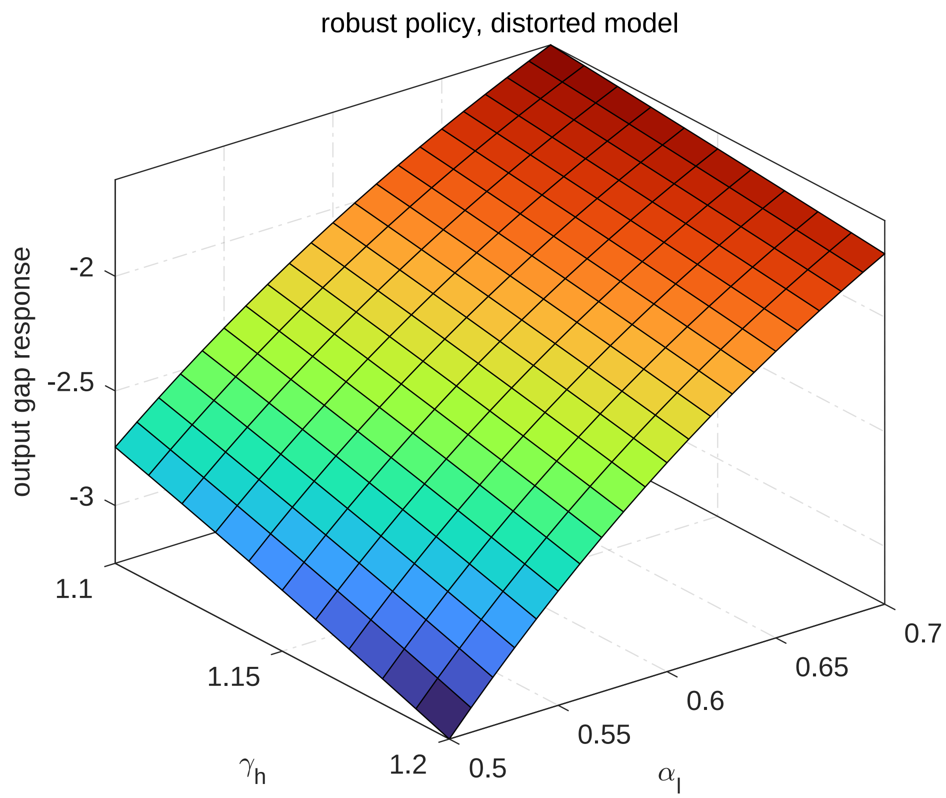

Figure 10: Output gap response for different worst-case combinations of α_l and γ_l.

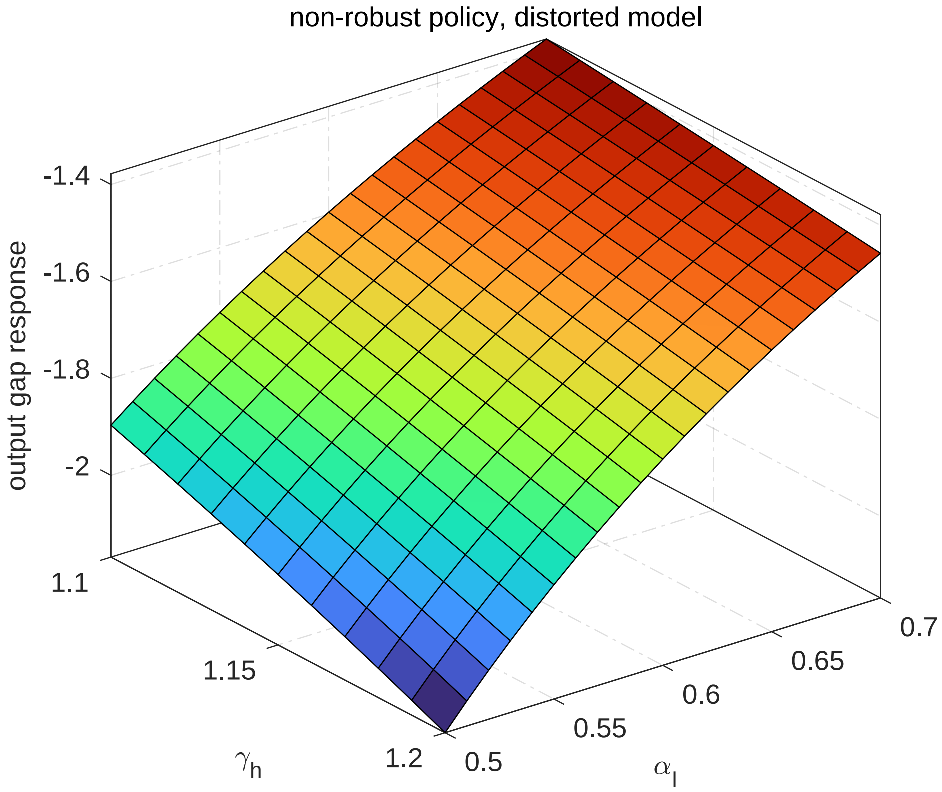

Figure 11: Output gap response for different worst-case combinations of α_l and γ_l.

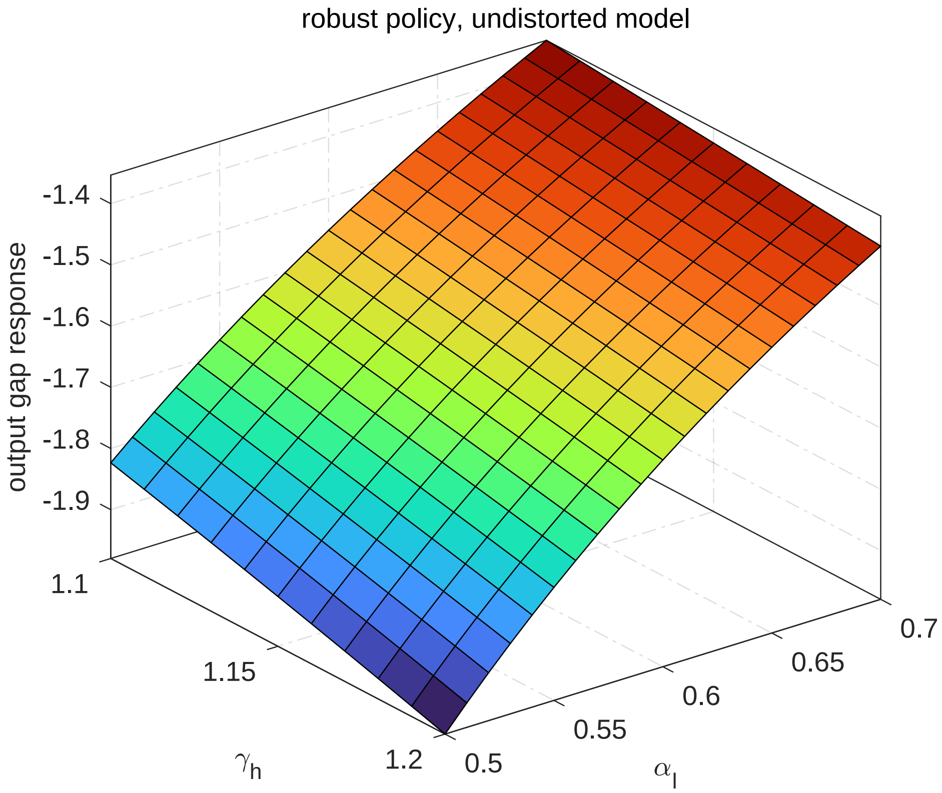

Figure 12: Output gap response for different worst-case combinations of α_l and γ_l.

Essay V:

Forward Guidance under the Cost Channel

This is a revised draft of the paper and currently under revision at *Macroeconomic Dynamics*. An earlier version is available under:

> **Finck, David**, "Forward Guidance under the Cost Channel," *MAGKS Joint Discussion Paper Series in Economics*, 2020, No. 04-2020.

This paper was presented at the following workshops and international conferences:

I: 30*th* Annual Symposium of the Society for Nonlinear Dynamics and Econometrics (Virtual)
Date: March 2022
Presenter: David Finck

Acknowledgement

I thank Peter Tillmann and Carl Walsh for reading multiple drafts and for very helpful comments and suggestions. I also thank Daniel Grabowski, Jörg Schmidt, Paul Rudel, Lucas Hafemann, Dennis Finck, Karsten Kucharczyk for useful comments and suggestions. Seminar participants at the University of Giessen and conference participants at the 30*th* Annual Symposium of the Society for Nonlinear Dynamics and Econometrics provided insightful comments.

Forward Guidance under the Cost Channel

David Finck*

Abstract

This paper analyzes how the cost channel of monetary policy affects the credibility of forward guidance. A cost channel is present when firms' marginal costs depend on the nominal rate of interest. As a result, there is a supply-side effect of monetary policy. We analyze the power of forward guidance using a calibrated New Keynesian model and compare our results to the standard New Keynesian model. We find that, compared to a standard New Keynesian model, the cost channel (i) makes forward guidance more powerful by reducing welfare losses at the zero lower bound, (ii) the central bank's best strategy requires a shorter forward guidance horizon, (iii) the strength of the cost channel plays an important role, and (iv) ignoring the cost channel is costly in terms of steady-state consumption.

Keywords: Forward Guidance, Credibility, Cost Channel, Discretion

JEL classification: E12, E43, E52, E58, E61

*University of Giessen, email: david.finck@wirtschaft.uni-giessen.de

I Introduction

In many advanced economies, short-term interest rates have reached the effective lower bound. In an environment of low interest rates, central banks have not much ammunition left to fight deflationary pressures. In response to the corona pandemic, the Federal Reserve cut interest rates to zero following two unscheduled meetings in March 2020.[1] In addition, in his statement on September 16, Fed chairman Powell showed projections by the FOMC that consider an increase in interest rates before the end of 2023 very unlikely.[2] It is therefore time to get used to the idea that the ZLB remains a constraint on monetary policy even after the recovery from the financial crisis. Central banks have established forward guidance, i.e. the central bank's commitment to future policy actions based on its assessment, as very effective tool to provide stimulus to the economy (for a recent overview, see Bernanke, 2020).

An early branch of the literature discusses the credibility of forward guidance in simple analytical frameworks and finds that the ZLB is not a serious constraint on the implementation of optimal monetary policy as long as the central bank can commit to future actions (Eggertsson and Woodford, 2003; Nakov, 2008, and Jung et al., 2005). However, models in these papers rely on the assumption that, once the ZLB episode is over, it never reoccurs in the future. Recent models have been extended in a way to allow for frequent ZLB episodes in order to analyze the power of forward guidance, even when there is no commitment technology available. According to Nakata (2018), an optimal Ramsey plan can credibly sustain as long as there is only a marginal probability that a binding zero lower bound will reoccur in the future.Walsh (2018) shows that also under discretionary policy, forward guidance can be credible as long as there are recurring episodes of binding constraints on interest rates.

This paper analyzes how the cost channel of monetary policy affects the credibility of forward guidance. A cost channel is present when firms' marginal costs depend on the nominal rate of interest. As a result, there is a supply-side effect of monetary policy. Investigating whether and, if so, how the cost channel changes the credibility of forward guidance is not least important because the cost channel is empirically relevant (see, for instance, Barth and Ramey, 2001; Ravenna and Walsh, 2006; Chowdhury et al., 2006 and Tillmann, 2008). We build on the approach of Walsh (2018) who proposes two conditions that must be satisfied to render forward guidance credible.

First, the gain of forward guidance must be positive. That is, the present discounted value of losses at the ZLB must be smaller under forward guidance compared to a reference policy (i.e., discretionary policy in our case). Second, the central bank must

[1]The Fed also began to purchase treasuries and mortgage-backed securities in order to stimulate economic activity.

[2]Even before this statement, the New York Fed's survey of primary dealers as of July 2020 reveals that 95% of respondents expect the federal funds rate to be equal to or below the range of 0.00–0.25% at the end of 2020. Also for the end of 2021 (2022), 77% (61%) of respondents expect interest rates to be within this range.

have no temptation to defect from promises made in the past. In other words, in every period until the promised lift-off date, fulfilling promises made in the past must be the central bank's best strategy.

We show that, compared to a standard New Keynesian model, the supply-side effect of the cost channel significantly changes the nature of forward guidance and the conditions under which promises about future policy actions are credible. Overall, our results are threefold.

First, we find that across both models, past promises might be honored, which is, however, only the case when fulfilling past promises is the central banks' best strategy at any point until the promised lift-off date. Across both models, we observe that the reason for forward guidance to be credible is that it improves outcomes at the ZLB by raising expectations of inflation and the output gap after exiting the ZLB. That is, we find that the promised path of zero interest rates leads to an anticipated overheating of the economy which, in turn, leads forward-looking households to increase consumption and firms to increase prices, accordingly. Nevertheless, in our calibrated model, we find that the optimal communicated horizon is shorter under the cost channel. Second, however, following the optimal strategy leads to larger improvements at the ZLB, which is found to be based on the breakdown of the divine coincidence. Finally, we find that these effects depend on the strength of the cost channel. More precisely, we find that both the improvement at the ZLB as well as the optimal horizon are strongly affected by the strength of the supply-side effect in the Phillips curve. We take these results as motivation and investigate what happens if the cost channel is simply ignored in the central bank's optimization problem. We find evidence that if the cost channel is present but ignored by the central bank and if the central bank behaves optimally under the perceived model it believes to be the true one, this can result in foregone improvements at the ZLB in terms of steady-state consumption.

This paper connects two strands of literature. The first strand assumes that each period, there is a fixed probability of leaving the ZLB. This Markovian structure is also used in, to name a few, Eggertsson and Woodford (2003), Eggertsson (2011), Bilbiie (2019).[3] However, in all of these papers, the problem is that ZLB episodes are one-off events. Even though the central bank can improve outcomes at the ZLB, following the insinuated path is of no further use if ZLB episodes cannot reoccur, because in such a situation, the central bank's best strategy is to renege on past promises. As is shown by Nakata (2018), things change if there is only a slight probability of recurring ZLB episodes. However, while Nakata (2018) considers the optimal Ramsey policy in his paper, we follow Walsh (2018) and assume that optimal monetary policy is discretionary. This is because, as discussed in Bernanke (2020), the limits of forward

[3]Other papers (see, for instance, Kiley, 2016 and Cochrane, 2017) assume that the ZLB is binding for a fixed number of periods.

guidance are rendered by how the public understands the central bank's communicated objective. In this respect, as pointed out by Walsh (2018), the simplicity of discretionary policy makes it easier to communicate to the public, even if discretionary policy might be suboptimal relative to the Ramsey policy. We find that similar to the results in Walsh (2018), also under the cost channel forward guidance can only be credible when recurring ZLB episodes are possible and if the length of horizon is not too long. Morevoer, we find a non-linear effect of the combination of high ZLB persistence and simultaneously high probability of recurrent ZLB episodes on the gain of forward guidance. More precisely, we find that even with one-period forward guidance, the gain can be negative if these two probabilities are not too high at the same time. Importantly, our results imply that this non-linear effect is much stronger under the cost channel.

The second strand of literature examines the role of the cost channel in the implementation of optimal monetary policy (including forward guidance) at zero lower bound. As we do in this paper, most of these papers rely on the Ravenna and Walsh (2006) framework which nests the standard New Keynesian model. For example, Pathberiya (2016) finds that within a discretionary framework, the optimal forward guidance horizon is longer when a cost channel is present. In his framework, he assumes that agents have perfect foresight in order to simplify the analysis. Chattopadhyay and Ghosh (2020) show that the cost channel plays an important role under forward guidance. They show that the implementation of forward guidance reduces welfare losses relative to discretion.

However, the problem in these papers is that they include simplifying assumptions (e.g., ZLB episodes as one-off events, perfect foresight) and abstract from various sources of uncertainty. For example, while perfect foresight might simplify the analysis, it also implies that both the central bank as well as households know exactly when and how long the ZLB will bind. Contrary to this, in our framework uncertainty about future ZLB episodes plays a key role in the expectation formation process. We therefore model the evolution of the economy in different states, namely at the ZLB, away from the ZLB and during every period within the forward guidance horizon.

The remainder of this paper is structured as follows. In Section II, we introduce the model and show how to solve it with numerical techniques. In Section III, we investigate how the cost channel affects equilibrium outcomes away from the ZLB in the absence of forward guidance. The latter is introduced in Section IV, where we first present the results for one-period forward guidance and subsequently extend the model to the multiperiod case. A battery of robustness checks and further results are presented in Section V. Section VI concludes.

II Methodology

A. *The Model*

Consider a simple New Keynesian model that features the cost channel as proposed in Ravenna and Walsh (2006). The private sector can be summarized by

$$x_t = E_t x_{t+1} - \sigma^{-1}\left(i_t - E_t \pi_{t+1} - r_t\right) \tag{1}$$

$$\pi_t = \beta E_t \pi_{t+1} + \kappa\left[\left(\sigma + \eta\right) x_t + \delta i_t\right], \tag{2}$$

where x_t is the output gap in period t, i_t is the short-term nominal interest rate, π_t is the current inflation rate, and r_t is an exogenous stochastic process. Contrary to the standard New Keynesian model, the Phillips curve now also includes a supply-side effect, δi_t, such that the short-term nominal interest rate appears both on the demand and the supply side. This effect relies on the assumption that firms' marginal costs depend on the nominal interest rate because firms need to raise nominal debt to finance production. Importantly, the Phillips curve as in (2) nests the conventional Phillips curve for $\delta = 0$.

We assume that the nominal interest rate must not be smaller than zero. That is, it must hold that[4]

$$i_t \geqslant 0. \tag{3}$$

We assume a two-state (states n and z) Markov chain process, so r_t can have two different values in either state, namely $r_z < 0$ or $\rho > 0$. Define $\mathbf{P}$ as the matrix of transition probabilities which is equal to

$$\mathbf{P} \equiv \begin{bmatrix} Pr(n \mid n) & Pr(n \mid z) \\ Pr(z \mid n) & Pr(z \mid z) \end{bmatrix} = \begin{bmatrix} s & 1-q \\ 1-s & q \end{bmatrix}.$$

In particular, we assume that in state z (ZLB is binding) when $r_t = r_z$, then $r_{t+1} = r_z$ with probability q and $r_{t+1} = \rho$ with $1-q$. We can therefore think of q as the probability of staying at the ZLB. Accordingly for state n , when $r_t = \rho$ (ZLB not binding), then $r_{t+1} = \rho$ with probability s and $r_{t+1} = r_z$ with $1 - s$. In other words, once we exit the ZLB, we assume that a non-negative probability $1 - s$ exists that periods can occur where the ZLB becomes binding (i.e. a reversion to the ZLB). This is the crucial difference to the Eggertsson and Woodford (2003) framework, where state z is a one-off event, i.e. where the probability of staying away from the ZLB is equal to $s = 1$.

An appropriate specification of monetary policy closes the model. It is assumed that

[4] For simplicity, we do not differentiate between the ZLB and the ELB, i.e., the rate below which it becomes profitable to exchange reserves for cash.

monetary policy minimizes the microfounded loss function

$$L_t = \frac{1}{2} E_t \sum_{k=0}^{\infty} \beta^k \left(\pi_{t+k}^2 + \lambda x_{t+k}^2 \right), \tag{4}$$

by means of discretionary policy, where λ is the relative weight the central bank places on the output gap.[5] Hence, we assume that the central bank reoptimizes the trade-off between the of output and inflation and there is no mechanism that allows the central bank to commit to future actions. The central bank minimizes (4) subject to (1)-(3) and takes households' expectations as given. The resulting targeting rule reads

$$\lambda x_t + \kappa \left[\sigma(1-\delta) + \eta \right] \pi_t = 0. \tag{5}$$

That is, once the shock r_t occurs in state n, the central bank sets the nominal interest rate in a way that (5) always holds. Keep in mind that in the absence of the cost channel, the resulting targeting rule reads $\lambda x_t + \kappa \left[\sigma + \eta \right] \pi_t = 0$. In the next section, we will discuss the implications of this change in the targeting rule.

B. *Solving the Model*

In the absence of forward guidance, we need to solve the model for two different states, namely for state n when the zero lower bound is not binding and for state z when the zero lower bound is binding. Recall that in state n, there is a non-negative probability $1-s$ that the zero lower bound will bind in the subsequent period. Accordingly, we remain away from the zero lower bound with probability s. Rational agents are able to solve the model and form expectations accordingly. Hence, under discretion, expected outcomes for the output gap and inflation in state n are weighted averages of the equilibrium values in state n and z, respectively, and read

$$E_t(x_{t+1}|n) = s x_n^d + (1-s) x_z^d$$
$$E_t(\pi_{t+1}|n) = s \pi_n^d + (1-s) \pi_z^d,$$

where the superscript d denotes outcomes under discretionary policy. The central bank implements the targeting rule (5) such that the resulting equilibrium condition in state n is given by

$$\lambda x_n^d + \kappa \left[\sigma(1-\delta) + \eta \right] \pi_n^d = 0. \tag{6}$$

[5]Note that in this paper, we assume that the appropriate reference for welfare consideration is the efficient economy, i.e., the output gap is measured as the difference between the actual output that would prevail in the absence of nominal rigidities and (in the cost channel economy) financial frictions. The advantage is that the corresponding loss functions of the central bank perfectly coincide across models. Importantly, in the standard New Keynesian model, price stickiness and market power are the only sources of inefficiency. Common practice is to introduce a fiscal subsidy of employment that is financed via a lump-sum tax. In this case, the output under flexible prices is also efficient.

The resulting equilibrium conditions for state n are obtained by using the expectations for output and inflation given we are in state n and read

$$x_n^d = \left[sx_n^d + (1-s)x_z^d\right] - \sigma^{-1}\left(i_n^d - \left[s\pi_n^d + (1-s)\pi_z^d\right] - \rho\right) \tag{7}$$
$$\pi_n^d = \beta\left[s\pi_n + (1-s)\pi_z\right] + \kappa\left[(\sigma+\eta)\,x_n^d + \delta i_n^d\right]. \tag{8}$$

We can apply the same procedure when the economy is stuck at the zero lower bound. In this case, the nominal interest rate is already determined and equal to zero. The model-consistent expected values of the output gap and inflation, given we are in state z, are

$$E_t(x_{t+1}|z) = qx_z^d + (1-q)x_n^d$$
$$E_t(\pi_{t+1}|z) = q\pi_z^d + (1-q)\pi_n^d.$$

As a result, the equilibrium conditions as implied by the supply and demand side are given by

$$x_z^d = \left[qx_z^d + (1-q)x_n^d\right] + \sigma^{-1}\left(\left[q\pi_z^d + (1-q)\pi_n^d\right] + r_z\right) \tag{9}$$
$$\pi_z^d = \beta\left[q\pi_z^d + (1-q)\pi_n^d\right] + \kappa\,(\sigma+\eta)\,x_z^d. \tag{10}$$

Importantly, the equilibrium conditions in state z coincide across the models, i.e., equations (9) and (10) do not directly depend on the presence of the cost channel. This is because in state n, the supply side effect disappears when the short-term interest rate is equal to zero. Nevertheless, outcomes at the zero lower bound, i.e., in state z, may differ across models because households attach a positive probability to equilibrium outcomes in state n.

We summarize the five unknowns in a vector $\mathbf{y} = [x_n^d\ \pi_n^d\ i_n^d\ x_z^d\ \pi_z^d]'$ and solve the model

$$\mathbf{A}\mathbf{y} = \mathbf{c},$$

where the matrix $\mathbf{A}$ contains the reduced-form coefficients and $\mathbf{c}$ is an appropriate vector that contains the corresponding entries of the stochastic process r_t, i.e. the corresponding entries of either ρ and r_z or zero otherwise.[6]

[6] In the appendix it is shown how to numerically solve the model. Because the system is purely forward-looking, we use this procedure throughout the paper. Note that many papers (see, for instance, Adam and Billi, 2006; Adam and Billi, 2007 and Nakata, 2018) model both a zero lower bound on nominal interest rates as well as a normal distribution for stochastic disturbances, which in combination renders the model highly non-linear. This is in part circumvented throughout this paper because we assume that the natural rate of real interest follows a two-state Markov process and we assign different states to equilibrium outcomes at and away from the ZLB. This simplifies our solution method.

C. *Calibration*

The baseline calibration for the structural parameters and the shock is summarized in Table (1).

Parameter	β	σ	η	κ	λ	r_z	ρ
Value	0.99	2	2	$0.02(\eta+\sigma)^{-1}$	0.003	-0.005	0.01

Table 1: Baseline Calibration

The values of the deep structural parameters for β, σ, η, and the value for λ are commonly used in the literature (see Eggertsson and Woodford, 2003; McKay et al., 2016; and more recently, Walsh, 2018).[7] While r_z is calibrated to model a shock with an annualized size of -2%, the value of β implies a long run real interest rate of $\rho = \beta^{-1} - 1 = 0.01$. It is worth noting that we need to multiply the value of κ in our model with $(\sigma + \eta)^{-1}$ to get a value of $\kappa = 0.02$ which corresponds to the slope of the Phillips curve as in the papers cited above. Throughout the paper, we compare our results from the textbook model, that is the New Keynesian model as in Clarida et al. (1999), with the model that features the cost channel as in Ravenna and Walsh (2006). For the former case, this amounts to setting $\delta = 0$, such that the model collapses to the textbook model. However, when the cost channel is considered, we use $\delta = 1$.[8]

Finally, we need to calibrate benchmark values for s and q. We follow Walsh (2018) and Nakata (2018) and discipline our choice by confronting the transition probabilities for states n and z with historical data for the effective federal funds rate running from 1960Q1-2019Q4. For this sample, the federal funds rate has been below 25bp for 11.81% of the time, while for 88.19% it was above or equal to 25bp. Thanks to the nature of Markov Chains, we can derive the steady-state behavior of our chain. For the limiting case $\lim_{k\to\infty} \mathbf{P}^k$, it must hold that

$$\Psi = \Psi \mathbf{P},$$

where Ψ contains the long-run fractions of being in either state n or state z, respectively. After some algebra, one can show that

$$\Psi(n) = \frac{1-q}{2-s-q} \overset{!}{=} 0.8819, \qquad \Psi(z) = \frac{1-s}{2-s-q} \overset{!}{=} 0.1181.$$

Following Walsh (2018) and Eggertsson and Woodford (2003) we choose values for q that are close to 0.9 and satisfy both conditions above, in particular $q = 0.875$ and

[7] The model-consistent value of λ is equal to $\lambda = \kappa(\sigma + \eta)\theta^{-1}$ where θ is the price elasticity of demand faced by firms. We follow Walsh (2018) and apply the same value for θ as in Woodford (2003), i.e., θ is equal to $\theta = 7.88$, which implies a relative weight of the output gap of $\lambda = 0.003$.

[8] This is the empirically estimated value for δ for the US. In order to test the sensitivity of our results with respect to this parameter, we will also try different values for δ.

calibration	q	s
A	0.875	0.98326
B	0.85	0.97991

Table 2: Baseline values for q and s.

$q = 0.85$. That is, once the economy is stuck at the ZLB, there is a 87.5% (85%) chance of staying at the ZLB in the next period. This in turn translates into a value of $s = 0.98326$ ($s = 0.97991$).

Recall that in state n, when the zero lower bound is not binding, the nominal interest rate needs to satisfy the non-negativity constraint, i.e. it must hold that $i_n \geqslant 0$. Hence, we need to ensure that our benchmark calibration results in an interest rate which satisfies (3).[9] Table (2) reports our baseline calibrations which we label as calibration A for $s = 0.98326$ and $q = 0.875$ and calibration B for $s = 0.97991$ and $q = 0.85$, respectively.

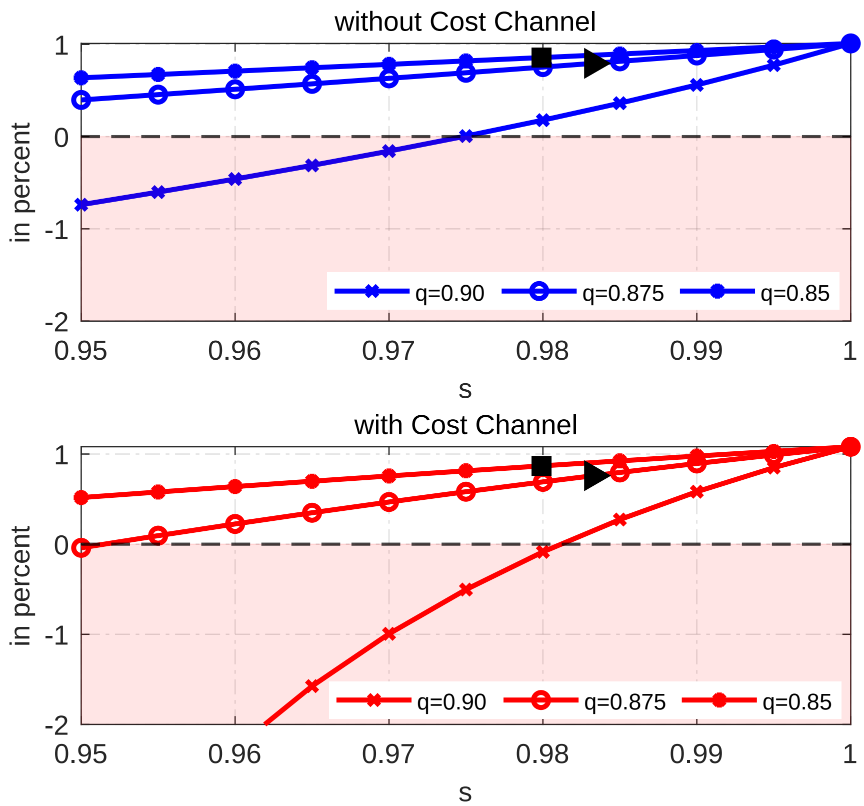

Figure 1: Short-term interest rate i_n in equilibrium for different combinations of s and q. The black squares (triangles) mark the corresponding equilibrium values under calibration A (B). Values are multiplied by 100.

Figure (1) plots the equilibrium values of i_n (in percent) for different combinations of s and q. The upper panel plots i_n in the absence of the cost channel, i.e., when $\delta = 0$, whereas in the bottom panel, the cost channel is included with $\delta = 1$. The

[9] This is why we cannot use a value of $q = 0.9$, as a 90% probability of staying at the ZLB violates the non-negativity constraint, see Figure (1).

shaded areas highlight the region in which the model violates the ZLB. Although we will discuss the effects of the cost channel later on in detail, it stands out that under both calibrations A and B, the interest rate in both models is positive so condition (3) is not violated.

Figure (1) also presents the schedule for i_n when $q = 0.90$. This is only presented because this value for the probability of staying at the ZLB is the same as in Eggertsson and Woodford (2003). Nevertheless, throughout the paper, we mainly present our results for $q = 0.875$ and $q = 0.85$.

To sum up, the baseline calibration ensures that the non-negativity constraint is satisfied for both models. However, we should be cautious when trying different combinations of s and q, as the equilibrium outcomes for i_n are very sensitive to s when the probability of staying at the ZLB is high.

III Equilibrium in the Absence of Forward Guidance

In this section, we first examine how the cost channel affects the equilibrium values in both states n and z in the absence of forward guidance. We do so by comparing the equilibrium outcomes both without and with the cost channel. Abstracting from the cost channel amounts to setting $\delta = 0$, such that the model collapses to the textbook model. On the contrary, we set $\delta = 1$ as the baseline value for the interest rate pass-through.

A. *How the Cost Channel Works in this Framework*

Before discussing the effect of the cost channel on the equilibrium outcomes of the output gap, inflation and the short-term interest rate, this subsection aims to clarify where and, if so, how the cost channel can make a difference at the ZLB in general and in this framework in particular.

We say that a cost channel is present when firms' marginal costs depend on the nominal rate of interest. This can be motivated under the assumption that firms need to hold working capital in order to paying factors of production before receiving their revenues. Technically, the cost channel translates into an additional term, δi_t, on the supply side of the model. That is, contrary to the textbook model, where the short-term interest rate solely appears on the demand side, under the cost channel it also has a direct effect on the supply side. The cost channel also introduces a direct effect on the specific targeting rule, that is the optimal relationship between output and inflation that must hold when the central bank adjusts its interest rate. Note that the targeting rule reads

$$\pi_t = -\left(\frac{\lambda}{\kappa}\right)\frac{1}{[\sigma(1-\delta)+\eta]}x_t.$$

For $\delta = 0$, the targeting rule collapses to the textbook case $\pi_t = -\lambda (\tilde{\kappa})^{-1} x_t$, where $\tilde{\kappa} \equiv \kappa (\sigma + \eta)$. However, under the cost channel, i.e. when $\delta > 0$, the optimal relationship between inflation and the output gap changes. Comparing the targeting rules across models, we see that in the presence of the cost channel, a given change in x_t goes in hand with a larger change in inflation compared to the textbook model. In other words, stabilizing inflation becomes more costly in terms of the output gap. Hence, the central bank needs to move the nominal interest rate by more compared to the case in the textbook model.

Another important aspect is that the cost channel represents a simple friction that leads to a breakdown of the divine coincidence. That is, given a shock on the demand side of the economy, the central bank can no longer perfectly stabilize the shock by counteracting appropriately. Hence, a shock on the demand side leads to inflation *and* output gap fluctuations under optimal discretionary policy. This is contrary to the standard New Keynesian model of Clarida et al. (1999), where an optimal response of the central bank to demand shocks guarantees that neither inflation nor the output gap deviate from their steady-state equilibrium values.

One advantage throughout this paper is that it is easy to trace back the source of diverging results from the models with and without the cost channel, respectively. Keep in mind that we have different states in our framework, namely one state away from the zero lower bound (state n), one state at the zero lower bound (state n) and, in the case of forward guidance, one additional state for each period of forward guidance until the lift-off date. To keep things simple, let us for a moment concentrate on the case without forward guidance. Importantly, in state z the short-term interest rate at the zero lower bound is already determined and equal to zero. This is important because in this case, the equilibrium conditions in state z across the textbook model and the model featuring the cost channel effectively coincide. Keeping this in mind helps us a lot when analyzing the sources of different results across models, because we know that any difference in equilibrium values in state z must effectively originate (via households' expectations) in state n. This argument is also true under forward guidance. That is, because the short-term interest rate is held at zero although the ZLB is no longer binding, the equilibrium conditions across models coincide in any forward guidance period until the lift-off date, i.e., until the central bank implements the optimal discretionary policy. When the economy is stuck at the zero lower bound, households form model-consistent expectations and take into account the equilibrium values that prevail when the economy has left the zero lower bound. Hence, via households' expectations, the equilibrium outcomes in state n also have an effect on equilibrium outcomes in state z. Also under forward guidance, the equilibrium values will differ across models because households form expectations accordingly and thus take outcomes in state n into account.

B. *Equilibrium Outcomes away from the ZLB*

As we discussed in the previous subsection, the breakdown of the divine coincidence in combination with the more aggressive targeting rule should have an effect on equilibrium outcomes. In this subsection, we will compare equilibrium outcomes away from the zero lower bound across both models, i.e. in the standard New Keynesian model as well as under the cost channel. State n serves as a benchmark, since all equilibrium conditions other than in state n coincide across models.[10]

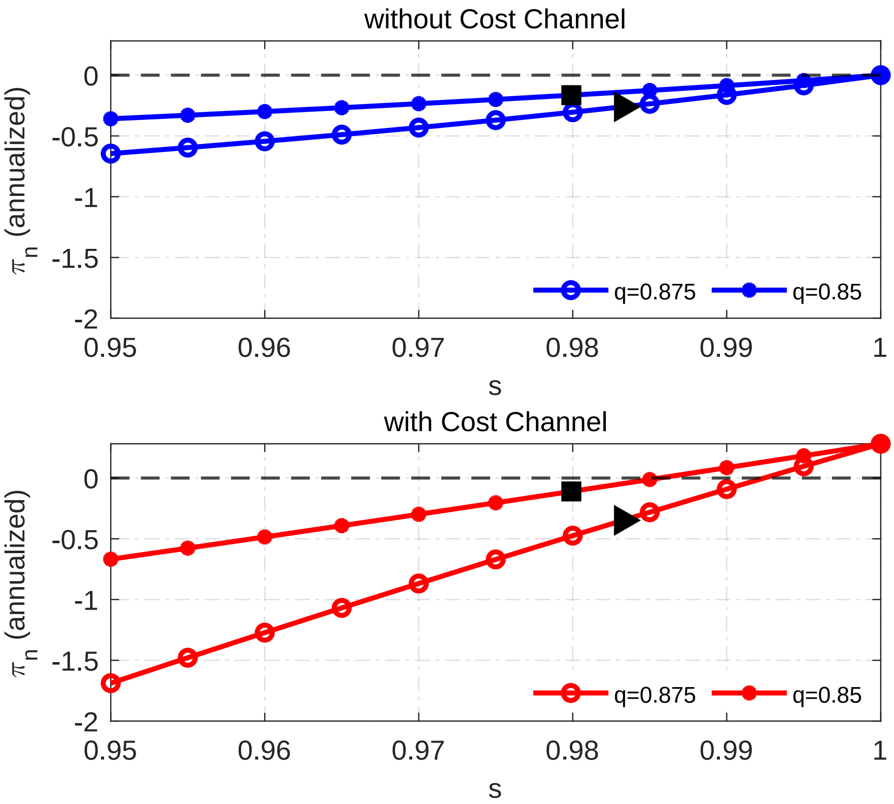

Figure 2: Outcomes for π_n in equilibrium for different combinations of s and q. The black squares (triangles) mark the corresponding equilibrium values under calibration A (B).

Figures (2) and (3) show the annualized equilibrium values of inflation and the output gap in state n for both models. In both figures, the upper panel shows equilibrium outcomes in the textbook model, while the lower panel shows the corresponding equilibrium values under the cost channel. Three key results stand out.

First, for $s = 1$, both inflation and the output gap in the textbook model are equal to zero, which is because the central bank can perfectly stabilize the demand shock. However, this is not the case when a cost channel is present. The breakdown of the divine coincidence worsens the trade-off between the stabilization of inflation and the output gap and, as a result, we have a non-zero inflation even when no recurring episodes at the ZLB are possible. Maintaining the targeting rule, which constitutes

[10]That is, equilibrium outcomes in state z will also be different because households attach a positive probability to a switch from state z to state n.

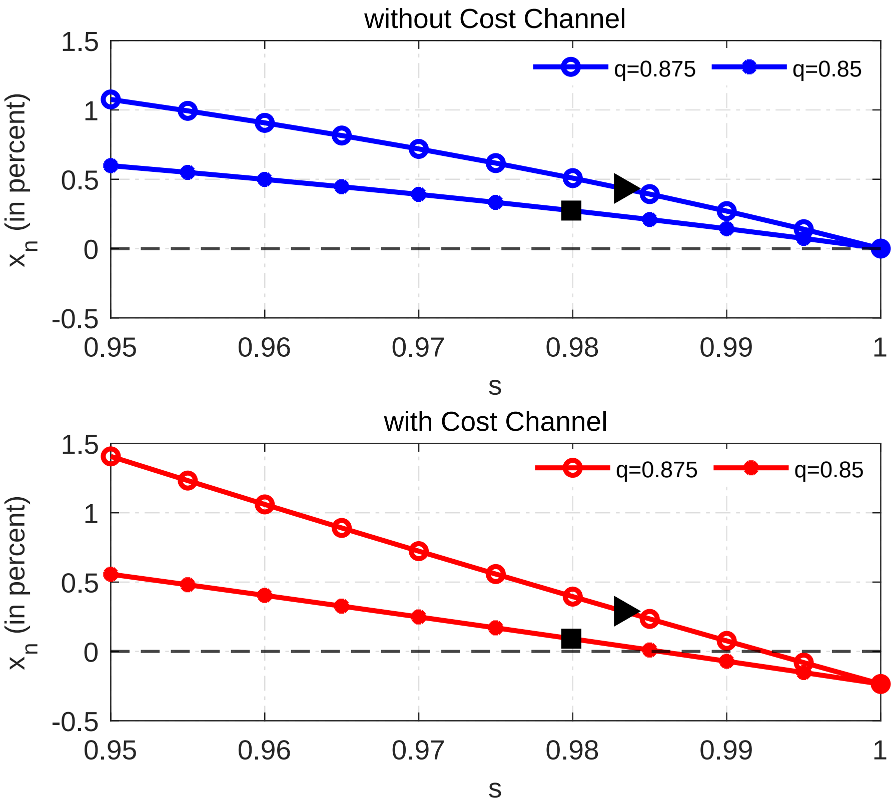

Figure 3: Outcomes for x_n in equilibrium for different combinations of s and q. The black squares (triangles) mark the corresponding equilibrium values under calibration A (B).

the optimal decision under discretion, requires a negative output gap. This is why we see a negative output gap for values of s close to one when a cost channel is present. Second, we see that under both models and for both inflation and the output gap, the equilibrium outcomes depend on s and q. That is, we observe that inflation (the output gap) decreases (increases), the more likely it is that an adverse demand shock occurs which pushes the economy back to the ZLB. The reason is that households take into account the possibility of lower income during future recessions and adjust their consumption decision for today accordingly. Not surprisingly, this effect is stronger if the expected duration of the ZLB episode increases, i.e. the higher q is.[11]

Third, and most importantly, the quantitative effects of s and q on the equilibrium outcomes of π_n and x_n strongly depend on whether a cost channel is present or not. That is, we observe that a higher probability of recurring ZLB episodes as well as a higher expected duration at the ZLB have a stronger effect under the cost channel. To show why this is the case, we use the Phillips curve and the targeting rule in state n and derive an expression for π_n that depends on π_z, i.e. equilibrium inflation at the

[11]This effect is also found in Walsh (2018).

ZLB. By doing so, we get

$$\pi_n = (1-s)\frac{\beta\lambda}{(1-\beta s)\lambda + \kappa^2(\sigma+\eta)(\sigma+\eta-\sigma\delta)}\pi_z + \frac{\kappa\lambda\delta}{(1-\beta s)\lambda + \kappa^2(\sigma+\eta)(\sigma+\eta-\sigma\delta)} i_n.$$

Given that the economy is stuck at the ZLB, the real interest rate is suboptimally high because the central bank cannot reduce the short-term interest rate below zero. The resulting negative output gap goes in hand with deflation, i.e. with $\pi_z < 0$. For a given π_z, the first term on the right-hand-side clearly shows that inflation in state n positively depends on π_z. That is, as long as there is only a marginal probability of reverting to the ZLB, the effect via the real interest rate also leads to deflation in state n. Moreover, this effect depends positively on the strength of the cost channel measured by δ.[12] The second term on the right-hand-side stems from the supply-side effect of the cost channel, i.e. the direct effect the short-term interest rate has on inflation. As opposed to the first term, this effect is positive and explains why, in equilibrium, we observe a change in both inflation (the output gap) from positive (negative) to negative (positive). This is the case when the supply-side effect is outweighs the effect from the demand side. For our calibrated model, inflation in state n is only positive when s is very close to unity, i.e. for $s > 0.985$. Importantly, the second term completely disappears when $\delta = 0$, i.e. in the absence of the cost channel.[13]

Recall that this line of arguments also holds under forward guidance. Even though things become more complicated through additional states, it is important to repeat that any different result across models we observe in the next section originates in state n, i.e. when the economy is away from the zero lower bound.

C. *Welfare Evaluation*

Throughout the paper, we follow Walsh (2018) and solve Bellman equations in order to account for the evolution of states. That is, we evaluate the equilibrium in terms of the presented value of losses. Let $l(\tau)$ be the current loss under state τ. Then, considering the whole sequence of events that will follow in the future, the Bellman equation in state τ reads

$$L_\tau = l_\tau + \beta\Big\{p(\tau|\tau)\times L(\tau) + p(\tau'|\tau)\times L(\tau')\Big\},$$

where $p(\tau|\tau)$ is the probability to remain in state τ and $p(\tau'|\tau)$ is the probability of switching to state τ', accordingly. In our case of a two state Markov structure, we

[12] In the robustness section, we try different values for δ in order to check how the strength of the cost channel affects our results.

[13] In this case, we get the same expression for π_n as a function of π_z as in Walsh (2018), namely $\pi_n = (1-s)\beta\lambda[(1-\beta s)\lambda + \kappa^2(\sigma+\eta)(\sigma+\eta)]^{-1}\pi_z$.

need one Bellman equation for each state. That is, our present discounted values L_n^d and L_z^d are given by

$$L_n^d = \frac{1}{2}\left[(\pi_n^d)^2 + \lambda(x_n^d)^2\right] + \beta s L_n^d + \beta(1-s)L_z^d \tag{11}$$

$$L_z^d = \frac{1}{2}\left[(\pi_z^d)^2 + \lambda(x_z^d)^2\right] + \beta q L_z^d + \beta(1-q)L_n^d, \tag{12}$$

where the expression in squared brackets is the current loss in either state.[14]

Figure (4) shows the corresponding present value of future discounted losses in state n, i.e. when the economy is away from the ZLB. As we would expect, for both models losses increase with the probability q of staying at the ZLB. Moreover, for both models the losses also increase the more likely the economy reverts to the ZLB, although this effect is much stronger under the cost channel. This reflects our results from before,

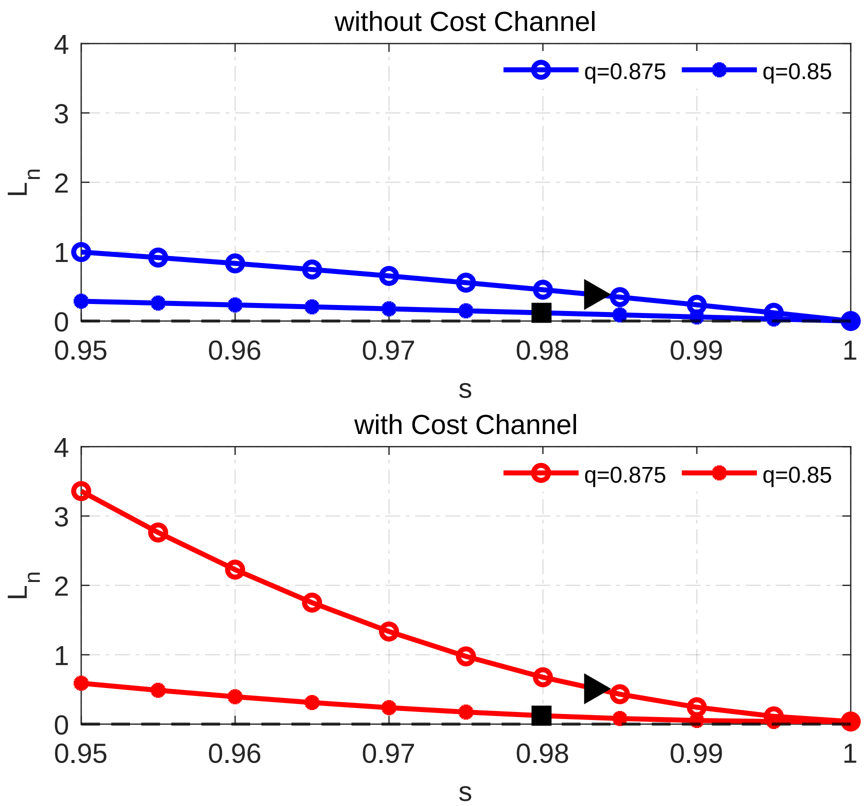

Figure 4: Loss in state n for different combinations of s and q. The black squares (triangles) mark the corresponding equilibrium values under calibration A (B). Values are expressed in percent of steady-state consumption.

as we found that the supply-side effect induced by the cost channel in combination with the more aggressive targeting rule leads to higher volatility. For instance, under calibration A, the present value of future discounted losses in state n correspond to

[14]Common practice is to express the present value of losses in terms of steady-state consumption equivalence (see Billi, 2017; Billi, 2011). Therefore, we follow Billi (2017) and Walsh (2018) and express the present value of losses L_τ^d in state $\tau = n, z$ as the share of steady-state consumption $\mu_\tau = 100 \times (1-\beta)\left[\frac{\omega\theta(1+\eta\theta)}{(1-\omega)(1-\omega\beta)}\right] L_\tau^d$, where $\omega = 0.75$.

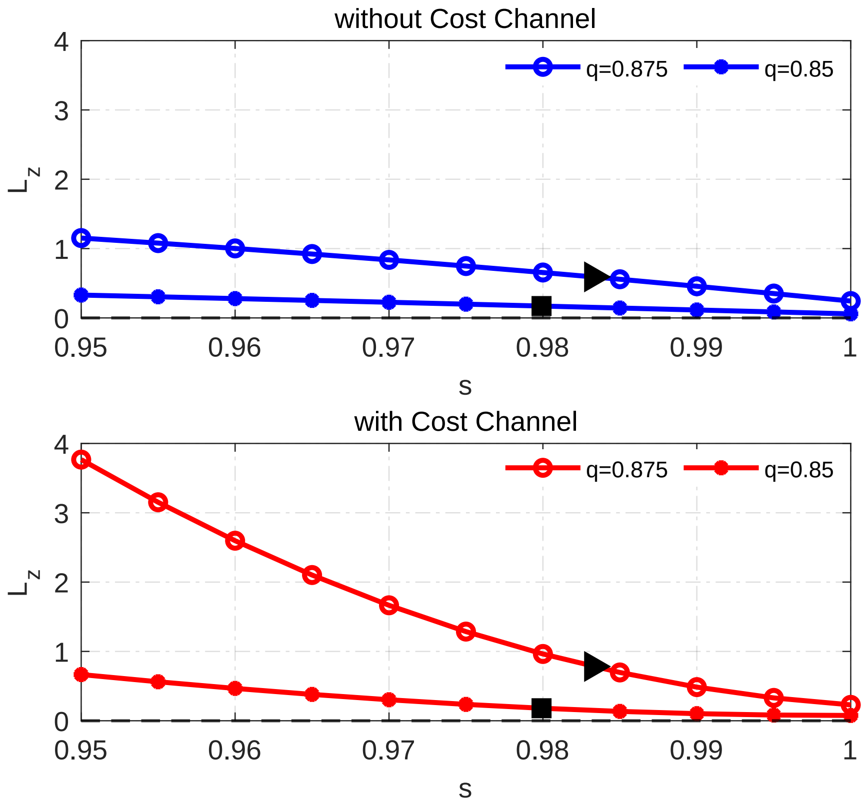

Figure 5: Loss in state z for different combinations of s and q. The black squares (triangles) mark the corresponding equilibrium values under calibration A (B). Values are expressed in percent of steady-state consumption.

0.382 percent of steady-state consumption without the cost channel and 0.510 percent when the cost channel is included, respectively. Under calibration B, with reversions to the ZLB being more likely, we get 0.119 percent and 0.122 percent, respectively.

Figure (5) shows the present value of future discounted losses at the ZLB, i.e. in state z. We see that losses are higher under the cost channel in terms of steady-state consumption. Not surprisingly, the outcomes in state z are in general higher than in state n, although the effect is small.

Figure (6) plots the difference between the results with and without the cost channel. Both the upper and bottom panels show the differences across models the same values of s and q as in Figures (4) and (5). Most strikingly, the difference in both states is always positive and shows the same patterns regarding the sensitivity to s and q as before.[15] The difference between losses at and away from the ZLB between the model with and without the cost channel very much depend on the transition probabilities. While for both $q = [0.85; 0.875]$ the difference between losses is very high, the more likely the economy reverts to the ZLB, this effect almost disappears the closer s gets to one. To sum up, these results show that the supply-side effect introduced by the cost channel changes the dynamics both at and away from the ZLB.

[15] This result is qualitatively similar to Pathberiya (2016), who studies optimal monetary policy under the cost channel at the ZLB. However, the results therein rely on a different modeling of the shock as well as on perfect foresight regarding household expectations.

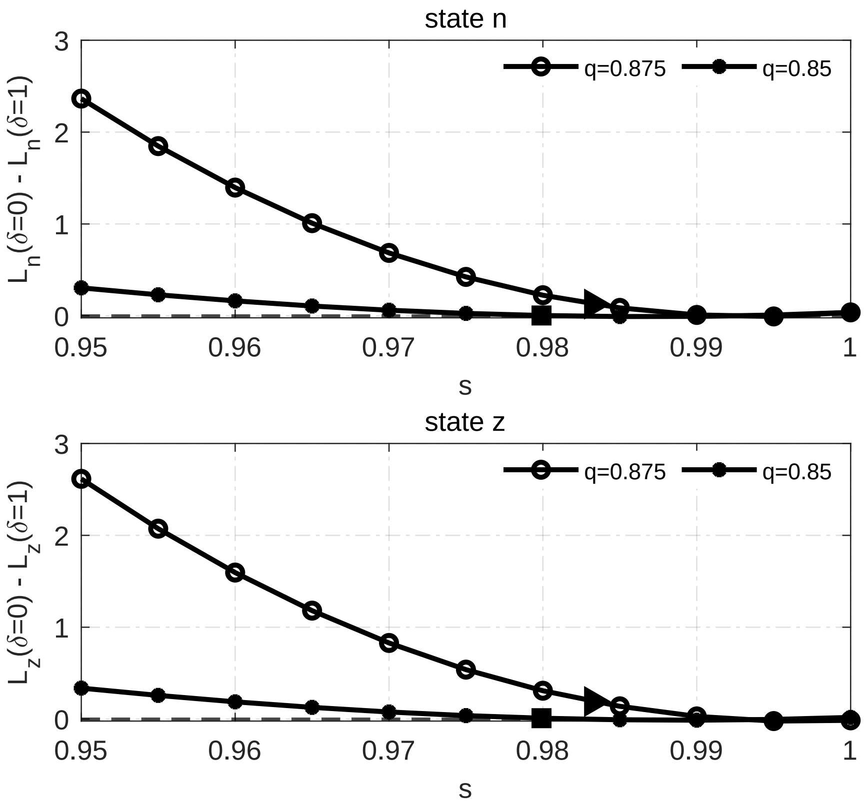

Figure 6: Difference between loss in state n and z for different combinations of s and q, with the cost channel ($\delta = 1$) and without the cost channel ($\delta = 0$). The black squares (triangles) mark the corresponding equilibrium values under calibration A (B). Values are expressed in percent of steady-state consumption.

IV Forward Guidance

Now that we have analyzed how the economy both with and without the cost channel works under optimal discretionary monetary policy, in this section we introduce forward guidance. Forward guidance amounts to the promise of the central bank to keep the nominal interest rate at zero for another k periods even when the ZLB is no longer binding in order to stimulate the economy. Having fulfilled its promise, in period $k+1$ the central bank implements the optimal discretionary policy, given that the economy has not returned to the ZLB.

For forward guidance to be credible, two conditions must be satisfied. First, losses at the ZLB under forward guidance must be small or equal to losses under optimal discretionary policy. Hence, it must hold that $L_z^{fg} \leqslant L_z^d$, where the superscript fg denotes losses under forward guidance and d denotes losses under optimal discretionary policy (i.e. $k = 0$), respectively. Second, given that the central bank keeps the nominal interest rate at zero when the ZLB is not a binding constraint, it must have no incentive to defect and implement the optimal policy instead. This incentive arises because optimal policy requires to counteract shocks accordingly. By keeping the nominal interest rate at zero instead, the central bank allows inflation and output to be different from zero. To sum up, it must hold that $L_e^{fg} < L_n^d$, where the subscript

e denotes outcomes in the exit period, i.e. the last forward guidance period. When this condition is satisfied, the central bank has no temptation to defect from promises made in the past such that fulfilling past promises is the central bank's best strategy in the exit period.[16]

In the case of one-period forward guidance, we need one additional state (state e, the exit period) consisting of two additional equations such that the joint set of equilibrium conditions reads

$$
\begin{aligned}
x_z^{fg} &= \left[q x_z^{fg} + (1-q) x_e^{fg}\right] + \sigma^{-1}\left(\left[q\pi_z^{fg} + (1-q)\pi_e^{fg}\right] + r_z\right) \\
\pi_z^{fg} &= \beta\left[q\pi_z^{fg} + (1-q)\pi_e^{fg}\right] + \kappa\,(\sigma+\eta)\,x_z^{fg} \\
x_e^{fg} &= \left[s x_n^{fg} + (1-s) x_z^{fg}\right] + \sigma^{-1}\left(\left[s\pi_n^{fg} + (1-s)\pi_z^{fg}\right] + \rho\right) \\
\pi_e^{fg} &= \beta\left[s\pi_n^{fg} + (1-s)\pi_z^{fg}\right] + \kappa(\sigma+\eta) x_e^{fg} \\
x_n^{fg} &= \left[s x_n^{fg} + (1-s) x_z^{fg}\right] - \sigma^{-1}\left(i_n^{fg} - \left[s\pi_n^{fg} + (1-s)\pi_z^{fg}\right] - \rho\right) \\
\pi_n^{fg} &= \beta\left[s\pi_n^{fg} + (1-s)\pi_z^{fg}\right] + \kappa\left[(\sigma+\eta) x_n^{fg} + \delta i_n^{fg}\right] \\
\lambda x_n^{fg} &+ \kappa\left[\sigma(1-\delta)+\eta\right]\pi_n^{fg} = 0.
\end{aligned}
$$

Notice that model-consistent expectations in state n now imply that households expect the central bank to keep the nominal interest rate at zero for one additional period, i.e. in the period the economy exits the ZLB. Hence, households place a positive weight for equilibrium values in state e which is the exit period in the case of one-period forward guidance ($k = 1$). By the same token, households realize that after keeping the nominal interest rate at zero for one period, the central bank will implement optimal discretionary policy thereafter.

The corresponding valuation equations read

$$
\begin{aligned}
L_e^{fg} &= \frac{1}{2}\left[(\pi_e^{fg})^2 + \lambda (x_e^{fg})^2\right] + \beta s L_n^{fg} + \beta(1-s) L_z^{fg} \\
L_n^{fg} &= \frac{1}{2}\left[(\pi_n^{fg})^2 + \lambda (x_n^{fg})^2\right] + \beta s L_n^{fg} + \beta(1-s) L_z^{fg} \\
L_z^{fg} &= \frac{1}{2}\left[(\pi_z^{fg})^2 + \lambda (x_z^{fg})^2\right] + \beta q L_z^{fg} + \beta(1-q) L_e^{fg},
\end{aligned}
$$

and can be jointly solved as in the case of pure discretionary policy.

A. *Results for One-Period Forward Guidance*

In this subsection, we analyze whether forward guidance is credible based on the underlying conditions stated above, i.e. we separately look at (i) the gain of forward guidance and (ii) the temptation to renege on the commitment. Recall that the former

[16] More generally, we adopt the definition of credibility as in Walsh (2018) who labels a policy plan as credible if the present value of losses obtained by implementing the policy (i.e., forward guidance in this case) is, in every state, less than or equal to the present value of losses under discretionary policy.

case corresponds to the difference of losses at the ZLB, i.e. in state z when implementing the promised path of interest rates. The latter case corresponds to the difference of losses in the exit period which controls whether fulfilling past promises is the central bank's best strategy. Figure (7) plots the gain $G = L_z^d - L_z^{fg}$ for different combinations

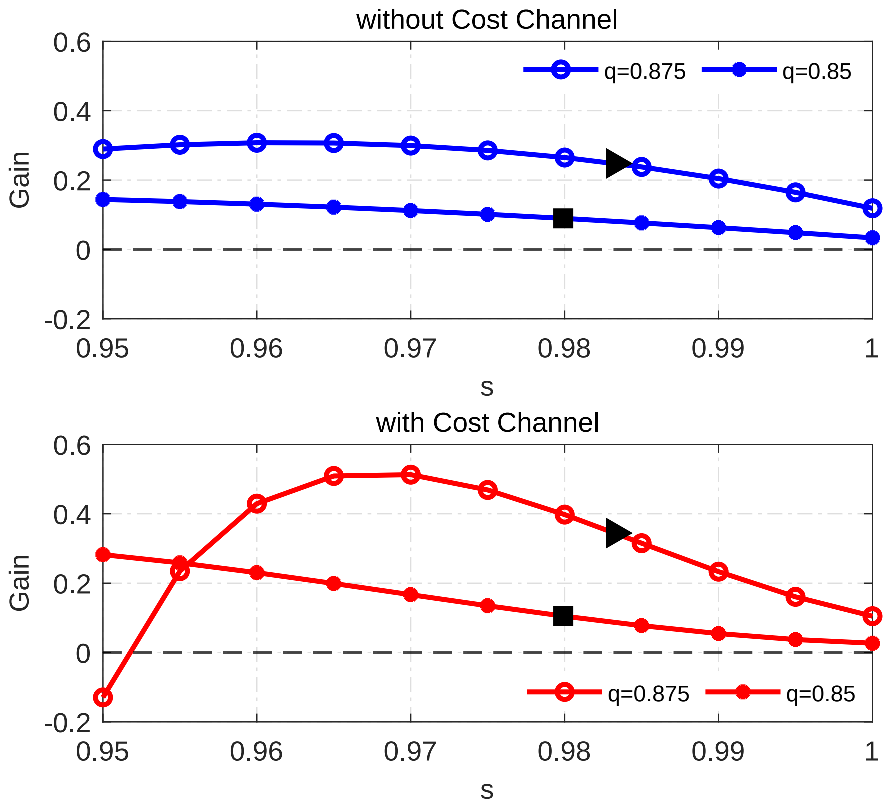

Figure 7: Gain of one-period forward guidance as $G = L_z^d - L_z^{fg}$ for different combinations of s and q.The black squares (triangles) mark the corresponding equilibrium values under calibration A (B). Values are expressed in percent of steady-state consumption.

of s and q, both for the model without the cost channel (upper panel) and with the cost channel (bottom panel).

Two things stand out. First, for $q = 0.85$, the gain of forward guidance is always positive. For our two baseline calibrations A and B, the gains are equal to 0.090 (A) and 0.248 (B) in the textbook model and 0.105 (A) as well as 0.344 (B) under the cost channel, respectively. Note that for $q = 0.85$, the gain of forward guidance increases, the more likely a reversion to the ZLB is. That is, in both models the gain monotonically increases in $1-s$.[17] Moreover, under $q = 0.85$, the gain is always higher under the cost channel than in the model without the cost channel. Intuitively, this means that the improvement of outcomes at the ZLB is larger when the cost channel is included as opposed to the case when the cost channel is absent.[18] Second, when we

[17]Note that Figures (7) and (8) mirror image the results typically found in the literature. Specifically, for $s = 1$, we find that the improvement at the ZLB delivers a positive gain for both values of q in both models. Nevertheless, the temptation is slightly positive. It follows that, in this case, the central bank is better off reneging on past promises and switch to the discretionary policy instead.

[18]Note that in the previous section, we argued that the strength of the cost channel also matters for

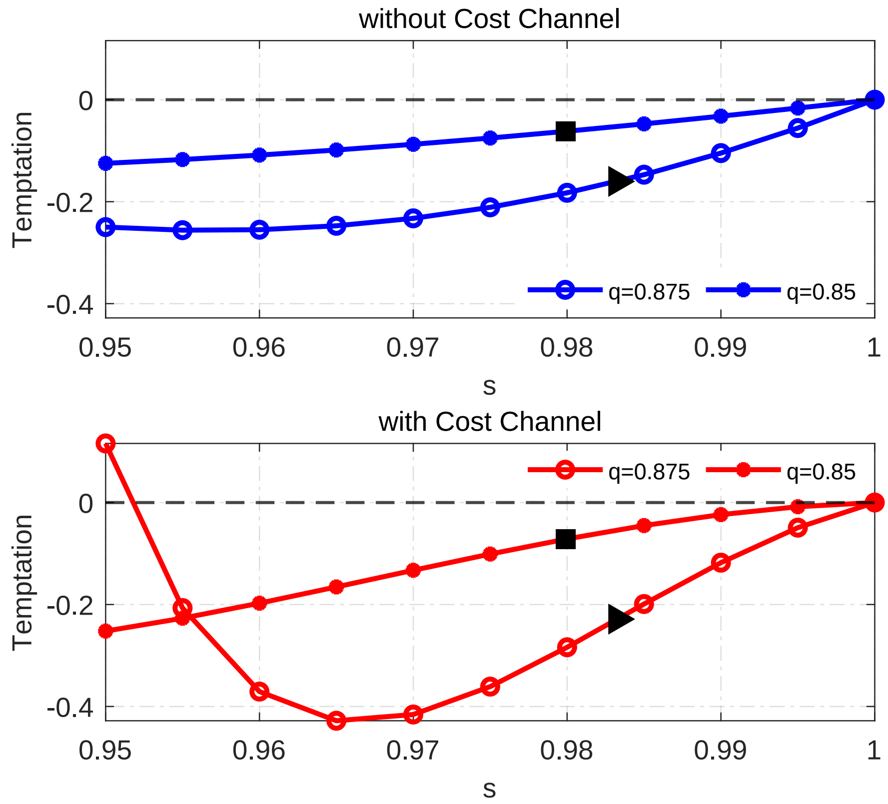

Figure 8: Temptation to defect for one-period forward guidance as $T = L_e^{fg} - L_n^d$ for different combinations of s and q. The black squares (triangles) mark the corresponding equilibrium values under calibration A (B). Values are expressed in percent of steady-state consumption.

assume a higher probability of staying at the ZLB, the central bank's promise can lack credibility under the cost channel if, at the same time, the probability of reverting to the ZLB is high. This can be seen as the gain in the bottom panel is negative for $q = 0.875$ and when s is close to 0.95. In this case, a promise to keep the nominal interest rate at zero after the economy has exited the ZLB would not be credible. However, this is not the case in the textbook model, where also under $q = 0.875$, the gain is always positive for all values of s between 0.95 and 1 and looks qualitatively very much like in the case of $q = 0.85$.

Figure (8) plots the corresponding temptation to defect. Roughly speaking, the results look qualitatively similar to the case where one would flip the paths from (7) upside-down. As a result, we find that for one-period forward guidance, there is no temptation to defect from past promises when losses at the ZLB under forward guidance are improved over outcomes under optimal discretionary policy.

Table (3) reports equilibrium outcomes for the output gap, inflation, and the nominal interest rate in all three states (except for the nominal interest rate as it is already determined in states e and z) under both, optimal discretionary policy and forward guidance. The upper part of the table reports the results for calibration A, while the

the gain of forward guidance. This can be seen in Figures (11) and (12), where we plot the gain of one-period forward guidance as a function of q and δ.

Calibration A

without Cost Channel							
k	x_n	x_e	x_z	π_n	π_e	π_z	i_n
0	1.735	1.735	−15.185	−0.260	−0.260	−2.511	0.794
1	1.320	3.008	−11.762	−0.198	−0.164	−1.911	0.844
with Cost Channel							
k	x_n	x_e	x_z	π_n	π_e	π_z	i_n
0	1.161	1.161	−17.080	−0.348	−0.348	−2.876	0.760
1	0.690	2.356	−13.336	−0.207	−0.190	−2.170	0.833

Calibration B

without Cost Channel							
k	x_n	x_e	x_z	π_n	π_e	π_z	i_n
0	1.097	1.097	−9.477	−0.164	−0.164	−1.350	0.857
1	0.753	2.557	−6.773	−0.113	−0.077	−0.927	0.902
with Cost Channel							
k	x_n	x_e	x_z	π_n	π_e	π_z	i_n
0	0.368	0.368	−10.346	−0.110	−0.110	−1.409	0.868
1	−0.021	1.845	−7.382	0.006	0.025	−0.908	0.933

Table 3: Results for one-period forward guidance. Outcomes for i are multiplied by 100, values for the π are annualized rates and values for x are in percent.

bottom part of the table reports outcomes for calibration B.

Not surprisingly, in equilibrium, all outcomes for inflation and the output gap are higher in absolute value, the longer the expected duration of the ZLB, both under optimal discretionary policy and under forward guidance. This observation holds both for the textbook model as well as under the cost channel. For example, under calibration B within the cost channel economy, deflation at the ZLB is equal to −1.409 (annualized) with a corresponding output gap of −10.346. Away from the ZLB, the equilibrium inflation rate is −0.110. The central bank adjusts the nominal interest rate in order to satisfy the targeting rule, which results in an output gap in state n of 0.368. Under calibration A and a more persistent ZLB episode, outcomes at the ZLB are almost twice as large with $\pi_z = -2.876$ and $x_z = -17.080$. Also, away from the ZLB, i.e. in state n, both the output gap and inflation are remarkably higher than before. However, when the central bank announces to keep the nominal interest rate at zero, once the economy has left the ZLB, much of the variance at and away from the ZLB can be reduced. This is where the gain of forward guidance comes from in both models, i.e. the improvement of outcomes of the inflation rate as well as the output gap at the ZLB. As we would expect, the variance of the output gap in the exit period is larger under forward guidance than under optimal discretionary policy. However, this effect is overcompensated through the improvement in inflation over discretionary policy because the output gap only has a relative weight of $\lambda = 0.003$ in the central banks' loss function. These results are qualitatively similar under the standard textbook model, although the relative differences in state n are higher under the cost channel.

Calibration A

without Cost Channel					
k	L_n	L_e	L_z	Gain	Temptation
0	0.382	0.382	0.592	0.000	0.000
1	0.222	0.222	0.344	0.248	-0.160
with Cost Channel					
k	L_n	L_e	L_z	Gain	Temptation
0	0.510	0.510	0.783	0.000	0.000
1	0.281	0.281	0.439	0.344	-0.229

Calibration B

without Cost Channel					
k	L_n	L_e	L_z	Gain	Temptation
0	0.119	0.119	0.171	0.000	0.000
1	0.056	0.057	0.082	0.090	-0.062
with Cost Channel					
k	L_n	L_e	L_z	Gain	Temptation
0	0.122	0.122	0.181	0.000	0.000
1	0.050	0.050	0.075	0.105	-0.072

Table 4: Results for one-period forward guidance. Losses are expressed in percent of steady-state consumption.

Table (4) summarizes our findings and reports the corresponding losses in each state as well as the gain and temptation under the same calibrations as in Table (3). The reported values for both, the gain and temptation imply that one-period forward guidance is credible across both models, as the gain for $k = 1$ is positive in both cases and the temptation to defect is negative.[19] Altogether, it stands out that the gain is higher under the cost channel due to lower welfare losses at the ZLB. Our results are comparable to Walsh (2018), who finds that even in a discretionary environment, forward guidance can be credible as long as the ZLB is occasionally binding. In the next section, we extend our analysis to the multiperiod case for $k > 1$.

B. *The Multiperiod Case*

Suppose the central bank promises to keep the nominal interest rate at zero for the general case of $k > 0$ periods after exiting the ZLB. Given that the economy has not reverted to the ZLB for the entire k periods, the central bank implements optimal discretionary policy in period $k+1$ by means of the targeting rule (6). If the economy instead reverts to the ZLB after $\underline{k} < k$ periods, the process starts over.

[19]Notice that for all calibrations, we also ensured that the non-negative constraint on the short-term interest rate is satisfied.

In order to jointly solve the model with multiperiod promises, we need k additional states, that is one state for each period that belongs to the promise made by the central bank. In particular, for each additional state we need two more equations because we have two additional unknowns.[20] For $k > 0$ periods, we have two equations in state z, i.e. at the ZLB, three equations in state n, i.e. when the economy is staying away from the ZLB under optimal discretionary policy and $2 \times k$ for each period (i.e. state) under forward guidance.

In this section, we focus on the credibility of forward guidance for values of $k = 0$ up to $k = 5$ under the baseline calibration for the structural parameters and the history-consistent calibrations for s and q as summarized in Table (2). However, it is worth noting that the analysis requires us to be cautious about the level of the short-term interest rate in state n. For each k under consideration we have to ensure that the interest rate in state n does not violate the non-negative constraint (3).

Calibration A

	without Cost Channel					*with Cost Channel*				
k	L_n	L_e	L_z	Gain	Temptation	L_n	L_e	L_z	Gain	Temptation
0	0.382	0.382	0.592	0.000	0.000	0.510	0.510	0.783	0.000	0.000
1	0.222	0.222	0.344	0.248	–0.160	0.281	0.281	0.439	0.344	–0.229
2	0.082	0.083	0.128	0.464	–0.300	0.085	0.085	0.136	0.647	–0.424
3	**0.004**	**0.005**	**0.007**	**0.585**	**–0.377**	**0.030**	**0.031**	**0.030**	**0.753**	**–0.478**
4	0.072	0.076	0.112	0.480	–0.307	0.426	0.433	0.563	0.220	–0.077
5	*0.469*	*0.480*	*0.726*	*-0.134*	*0.097*	*2.292*	*2.317*	*3.194*	*-2.411*	*1.807*

Table 5: Results for k-period forward guidance. Losses are expressed in percent of steady-state consumption. Rows with values in bold font indicate that under this horizon k, the gain is maximized and forward guidance is credible. Rows with values in italics correspond to horizons where forward guidance is not credible, either because (1) the gain is negative, (2) temptation is positive or (3) both at the same time.

Let us start with our results for calibration A. Table (5) reports the present value of losses in states e, n and z as well as the gain and the temptation to defect from the promised forward guidance policy for $k = 0$ (i.e. optimal discretion) up to $k = 5$. In the absence of the cost channel, promises to keep the short-term interest rate at zero for up to four periods after exiting the ZLB deliver significantly better outcomes over discretion in terms of welfare losses. As equilibrium outcomes in state n also depend on outcomes in state z, losses away from the ZLB are also being improved by promises to keeping the interest rate at zero for up to four periods. However, the present value of losses at the ZLB is minimized for $k = 3$ with a gain of 0.585. Also, note that such a promise is credible as the central bank has no temptation to defect from past promises made when the economy was stuck at the ZLB.

Under the cost channel, the qualitative results are very similar. Losses are lower under forward guidance as against discretion for up to $k = 4$. The gain is maximized for $k = 3$. In this case, the present value of losses at the ZLB is 0.030 as compared to

[20]The interest rate is obviously known under forward guidance and determined as zero.

0.783 under discretion, which results in a gain of 0.753. The corresponding value of -0.478 in the last column states that this promise is also credible.

Calibration B

	without Cost Channel					*with Cost Channel*				
k	L_n	L_e	L_z	Gain	Temptation	L_n	L_e	L_z	Gain	Temptation
0	0.119	0.119	0.171	0.000	0.000	0.122	0.122	0.181	0.000	0.000
1	0.056	0.057	0.082	0.090	−0.062	0.050	0.050	0.075	0.105	−0.072
2	0.012	0.012	0.017	0.154	−0.107	**0.016**	**0.017**	**0.019**	**0.162**	**−0.105**
3	**0.006**	**0.007**	**0.009**	**0.162**	**−0.111**	0.078	0.080	0.086	0.094	−0.042
4	0.081	0.084	0.116	0.055	−0.034	0.363	0.369	0.447	−0.267	0.247
5	*0.313*	*0.323*	*0.452*	*-0.281*	*0.205*	1.179	1.199	1.514	−1.334	1.077

Table 6: Results for k-period forward guidance. Losses are expressed in percent of steady-state consumption. Rows with values in bold font indicate that under this horizon k, the gain is maximized and forward guidance is credible. Rows with values in italics correspond to horizons where forward guidance is not credible.

What if the expected duration of ZLB episodes is shorter and reversions to the ZLB episodes are more likely? Table (6) reports the same results for calibration B. Across both models, losses at the ZLB are in general lower than in the former case, because the expected duration of ZLB episodes is shorter. Without the cost channel, forward guidance still yields improved outcomes over discretion at the ZLB for up to four periods. Also, just as in the former case, losses at the ZLB are minimized for $k = 3$, which is also a credible promise as such a promise delivers no positive temptation to defect.

However, under the cost channel, the results are qualitatively different compared to the former case. We find that, when the expected duration of ZLB episodes is lower, forward guidance delivers improved outcomes only for up to $k = 3$ as opposed to $k = 4$ for the case without the cost channel. Also, losses at the ZLB are minimized for $k = 2$ as opposed to $k = 3$ when the cost channel is absent. Recall that the inclusion of the cost channel only makes a difference in state n, as at the ZLB and under forward guidance, the nominal interest rate is equal to zero. Hence, the difference must stem from the supply-side effect of the nominal interest rate in the Phillips Curve as well as the targeting rule. Overall, it stands out that while forward guidance is less credible under the cost channel, we find that at the same time, the improvement at the ZLB can be greater under the cost channel than in its absence.

Figure (9) summarizes our results and plots the gain of forward guidance based on the two calibrations of Table (2) for $k = 0$ (discretionary policy) up to $k = 5$. Not surprisingly, across both models, the gain is higher when the expected duration of the ZLB period is longer. This is because improved outcomes for both, inflation and the output gap at the ZLB, imply expected inflation to be closer to zero upon exiting than it is under optimal discretion (see also Walsh, 2018 for this point). It stands out that in general, both models show a similar qualitative picture, saying that the gain shows an inverse U-shape, also turning from positive to negative for pretty much the same

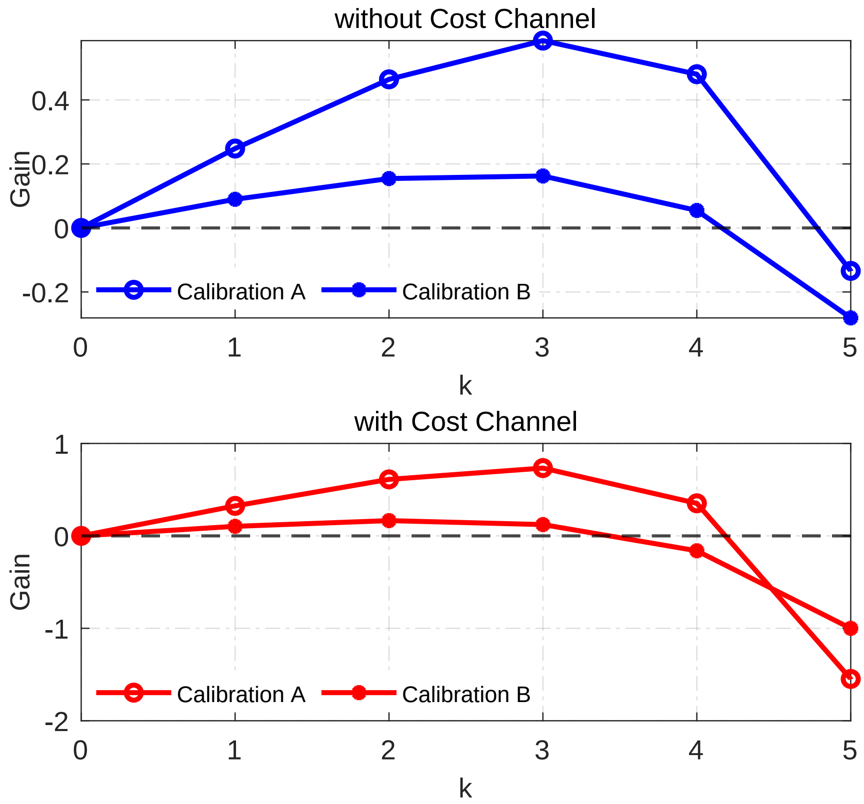

Figure 9: Gain of forward guidance as $G = L_z^d - L_z^{fg}$ for different combinations of s and q. Values are expressed in percent of steady-state consumption.

values of $k = 5$. However, with a reversion to the ZLB more likely (calibration B), the gain under the cost channel is also negative for $k = 4$ but positive in the textbook model. Figure (10) shows the temptation to defect from past promises for the same calibrations as in the former case. In all cases, the temptation shows a hump-shaped pattern and, most importantly, implies that temptation is always negative when the gain is positive at the same time. Hence, under the calibration considered, we do not have the case that the central bank has a positive temptation to defect while still being able to improve outcomes at the ZLB.

V Sensitivity Analysis

A. *The Effect of the Real Interest Rate*

In this subsection, we examine how our results depend on the effect of the real interest rate. Our benchmark calibration assigns a value of $\sigma = 2$ to the inverse elasticity of substitution. Even though this value is widely accepted in the literature, we check how our results change if we try a different value of σ, namely $\underline{\sigma} = 0.16$. This value is used for example in Woodford (2003), Adam and Billi (2006) as well as Adam and Billi (2007).[21] From an economic perspective, σ measures the inverse intertemporal

[21] In these papers (or book in the case of Woodford, 2003), the intertemporal elasticity of substitution is denoted σ rather than σ^{-1}. A value of $\sigma = 0.16$ in our model therefore refers to $\sigma = 6.25$ in the

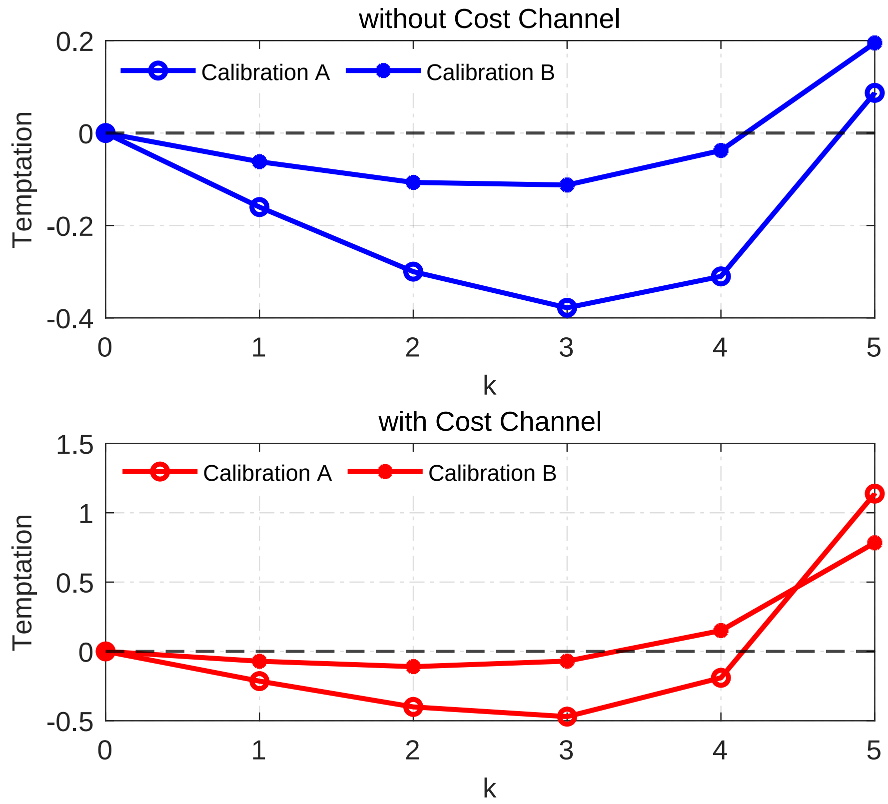

Figure 10: Temptation of forward guidance as $G = L_z^d - L_z^{fg}$ for different combinations of s and q. Values are expressed in percent of steady-state consumption.

elasticity of substitution of aggregate spending. That is, lower values of σ imply that households respond stronger to changes in the real interest rate. Table (7) reports the results for the gain and temptation, both for our calibration of $\sigma = 2$ as well as the alternative value of $\underline{\sigma} = 0.16$. Results are reported for both calibrations for s and q as in the main text.

A few things stand out. First, under calibration A, we find that a promise to keep the interest rate at zero after the economy has left the ZLB improves outcomes at the zero lower bound over discretion for a horizon up to $k = 3$. Interestingly, however, a promise to keep the interest rate at zero for three quarters yields a positive temptation, both under the cost channel as well as in the textbook model. Hence, such a promise is not credible because the central bank has an incentive to renege on past promises and switch to discretionary policy instead. Second, for $\underline{\sigma}$, the results for both the gain and temptation for all horizons are much closer across both models. Intuitively, this is because with a smaller value for σ, we increase the effect of the real interest rate on aggregate expenditure. That is, we technically reduce the relative effect of the cost channel. For both models, the highest gain is achieved when $k = 2$. This promise is also credible.

Under calibration B, the striking qualitative difference compared to the main part of the paper here is that in the previous section, promises were credible for up to $k = 4$.

papers cited above.

Calibration A

	without Cost Channel				*with Cost Channel*			
	$\sigma = 2$		$\overline{\sigma} = 0.16$		$\sigma = 2$		$\overline{\sigma} = 0.16$	
k	Gain	Temptation	Gain	Temptation	Gain	Temptation	Gain	Temptation
0	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
1	0.248	−0.160	0.358	−0.220	0.344	−0.229	0.380	−0.247
2	0.464	−0.300	**0.478**	**−0.250**	0.647	−0.424	**0.491**	**−0.280**
3	**0.585**	**−0.377**	*0.074*	*0.135*	**0.753**	**−0.478**	*0.036*	*0.134*
4	0.480	−0.307	*−1.293*	*1.275*	0.220	−0.077	*−1.432*	*1.342*
5	*−0.134*	*0.097*	*−4.195*	*3.613*	*−2.411*	*1.807*	*−4.465*	*3.774*

Calibration B

	without Cost Channel				*with Cost Channel*			
	$\sigma = 2$		$\underline{\sigma} = 0.16$		$\sigma = 2$		$\underline{\sigma} = 0.16$	
k	Gain	Temptation	Gain	Temptation	Gain	Temptation	Gain	Temptation
0	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
1	0.090	−0.062	0.512	−0.350	0.105	−0.072	0.546	−0.385
2	0.154	−0.107	**0.635**	**−0.385**	**0.162**	**−0.105**	**0.642**	**−0.423**
3	**0.162**	**−0.111**	*−0.008*	*0.164*	0.094	−0.042	*−0.061*	*0.158*
4	0.055	−0.034	*−1.826*	*1.615*	*−0.267*	*0.247*	*−1.965*	*1.665*
5	*−0.281*	*0.205*	*−5.200*	*4.260*	*−1.334*	*1.077*	*−5.403*	*4.352*

Table 7: Results for k-period forward guidance. Losses are expressed in percent of steady-state consumption. Rows with values in bold font indicate that under this horizon k, the gain is maximized and forward guidance is credible. Rows with values in italics correspond to horizons where forward guidance is not credible.

With $\underline{\sigma}$ increasing the effect of the real interest rate, this is no longer the case. Now, forward guidance is credible for up to two periods after the economy has left the ZLB. However, contrary to the previous case, all of these promises are also credible, and the central bank's best strategy is to fulfill its promises. Note that also under calibration, we find that the difference between the textbook model as well as the model featuring the cost channel became much smaller. Since, also under the benchmark calibration, the most significant improvement over discretion could be achieved for $k = 2$, this makes no qualitative difference under the cost channel.

Summing up, we find that our qualitative results do not change considerably. While the optimal forward guidance horizons are exactly the same under both models, results only change for the textbook model. Still, we find that forward guidance in both models can be credible as long as the forward guidance horizon is not too long. Most importantly, however, we find again that for the optimal horizon k, forward guidance is more powerful with the cost channel than without the cost channel.

B. *Results Without Recurring ZLB Episodes*

Throughout the paper, we mainly focused on calibrated values for s and q that match the historical data for the US. In this section, we relax this assumption and try different combinations for s and q, i.e. calibrations that do not match movements of the federal funds rate. We now try the combinations $(s, q) = (1, 0.9)$ as well as $(s, q) =$

(0.9999, 0.9). For the former case, our model collapses to the model of Eggertsson and Woodford (2003), who abstract from recurring ZLB episodes and assume that the binding constraint on the nominal interest rate is a one-off event instead. The second calibration is chosen because Walsh (2018) shows that forward guidance in a discretionary framework can be credible when there is even a marginal probability of recurring ZLB episodes. This experiment is interesting for two reasons. First, while

	without Cost Channel				*with Cost Channel*			
	$s = 0.9999$		$s = 1$		$s = 0.9999$		$s = 1$	
k	Gain	Temptation	Gain	Temptation	Gain	Temptation	Gain	Temptation
0	0.000	0.000	*0.000*	*0.000*	0.000	0.000	*0.000*	*0.000*
1	2.010	−0.020	*2.003*	*0.000*	2.010	−0.018	*1.901*	*0.000*
2	3.649	−0.035	*3.623*	*0.000*	3.440	−0.031	*3.396*	*0.000*
3	4.742	−0.045	*4.709*	*0.001*	4.364	−0.038	*4.302*	*0.001*
4	**5.060**	**−0.047**	*5.025*	*0.003*	**4.456**	**−0.036**	*3.382*	*0.008*
5	4.327	−0.036	*4.305*	*0.006*	3.412	−0.020	*0.915*	*0.015*

Table 8: Results for k-period forward guidance for different values of s. In all cases, $q = 0.9$. All values are expressed in percent of steady-state consumption. Rows with values in bold font indicate that under this horizon k, the gain is maximized and forward guidance is credible. Rows with values in italics correspond to horizons where forward guidance is not credible.

we would expect that for $s = 1$, forward guidance can never be credible in the textbook model, we do not know whether this is also the case under the cost channel. From an economic perspective, this is because under the cost channel, any demand shock always creates a trade-off between the stabilization of the output gap and inflation and leads to fluctuations even under an optimal policy, because shocks on the demand side can no longer be stabilized perfectly. Therefore, under the cost channel, $\pi_n = x_n = 0$ cannot be part of a feasible solution. Second, we found that forward guidance under the cost channel is at best as credible as under the textbook model. Therefore, we assess whether forward guidance can also be credible when the probability of reverting to the ZLB is only minuscule.

Table (8) reports the results for both models. For $s = 1$, we find exactly what we would expect for the textbook model.[22] That is, although the gain is positive and there seems to be an opportunity to improve outcomes at the ZLB, such promises are never credible because the central bank would always renege on past promises and implement discretionary policy in the exit period instead. However, also under the cost channel, the gain of forward guidance is positive, even when future ZLB episodes never reoccur. Again, we find that the central bank is better off if it defects from past promises. As a result, forward guidance is not credible in the Eggertsson and Woodford (2003) framework. Intuitively, the reason is that in the exit period, either inflation or the output gap or both are closer to zero under discretionary policy than under forward guidance, although demand shocks cannot be perfectly stabilized.

[22]In fact, these are one-by-one the results as in Walsh (2018).

When $s = 0.9999$, promises are credible for the entire horizon from $k = 1$ up to $k = 5$. In both models, the largest improvement over discretionary policy at the ZLB can be achieved for $k = 4$.

Summing up our results, we find that forward guidance can be credible under the cost channel, even if the probability of reoccurring ZLB episodes is very small (0.01%). However, we also find our usual qualitative result that forward guidance is less credible under the cost channel. This is because if we extend the horizon to $k = 7$, we find that forward guidance is credible for up to $k = 6$ in the textbook model and up to $k = 6$ under the cost channel.

C. *The Strength of the Cost Channel*

In this subsection, we assess how our results behave when we change the strength of the cost channel, which is controlled by the coefficient of δ. Values of δ that are smaller than one indicate that changes in the central bank's policy rate do not translate into a one-to-one increase of inflation through the supply side. This can be motivated by the assumption that firms only partly rely on financial intermediaries in order to pay factors of production in advance. As a result, a value of $\delta > 1$ would imply that firms have to pay a markup when borrowing from financial intermediaries. This could be motivated by the assumption that the pass-through from the central bank's policy rate to the actual lending rate is not complete. A value of $\delta = 1$ implicitly assumes that there is no friction which results in an incomplete pass-through from the central bank's policy rate to the lending rate.

Our calibration of $\delta = 1$ as our benchmark value is taken from Ravenna and Walsh (2006), who use a broad set of instruments and show that the coefficient of δ is not significantly different from one. However, the uncertainty around their estimate also covers values of δ that are far above and below one. Chowdhury et al. (2006) estimate a hybrid Phillips curve featuring a cost channel and find a value of $\delta = 1.3$ for the US. Hence, we work with two different values of δ, namely $\underline{\delta} = 0.5$ and $\overline{\delta} = 1.5$.

The results for the gain and temptation for values of $k = 0$ to $k = 5$ under our benchmark calibrations A and B are presented in Table (9) for the set of $\boldsymbol{\delta} = \left[\underline{\delta}, \delta, \overline{\delta}\right]$. Overall, we find that the calibration of δ makes a difference. Broadly speaking, we find that a lower value of δ moves the results in the direction of the results under the textbook model. While for $\underline{\delta}$, the optimal horizon for k is equal to the benchmark value of $\delta = 1$, we find that under calibration B, a promise to keep the interest rate at zero for four periods after the ZLB episode is now credible, as there is an improvement of outcomes at the ZLB and because the central bank has no temptation to defect when it comes to fulfill its promise in the exit period. In this sense, the results are qualitatively equal to the textbook model regarding the overall credibility as well as the optimal horizon for k. However, for $\overline{\delta}$ we find a different picture. Now, the optimal horizon is one quarter lower compared to the results under the baseline calibration.

Calibration A

	$\underline{\delta} = 0.5$		$\delta = 1$		$\bar{\delta} = 1.5$	
k	Gain	Temptation	Gain	Temptation	Gain	Temptation
0	0.000	0.000	0.000	0.000	0.000	0.000
1	0.277	–0.180	0.344	–0.229	0.718	–0.508
2	0.523	–0.340	0.647	–0.424	**1.147**	**–0.749**
3	**0.657**	**–0.423**	**0.753**	**–0.478**	*-2.063*	*1.972*
4	0.497	–0.309	0.220	–0.077	*-88.358*	*70.324*
5	*-0.400*	*0.299*	*-2.411*	*1.807*	*-369.642*	*286.471*

Calibration B

	$\underline{\delta} = 0.5$		$\delta = 1$		$\bar{\delta} = 1.5$	
k	Gain	Temptation	Gain	Temptation	Gain	Temptation
0	0.000	0.000	0.000	0.000	0.000	0.000
1	0.096	–0.067	0.105	–0.072	**0.077**	**–0.033**
2	**0.163**	**–0.112**	**0.162**	**–0.105**	*-0.156*	*0.208*
3	0.161	–0.107	0.094	–0.042	*-1.443*	*1.357*
4	0.010	–0.007	*-0.267*	*0.247*	*-6.953*	*6.098*
5	*-0.453*	*0.349*	*-1.334*	*1.077*	*-39.043*	*33.256*

Table 9: Results for different values of δ. All values are expressed in percent of steady-state consumption. Rows with values in bold font indicate that under this horizon k, the gain is maximized and forward guidance is credible. Rows with values in italics correspond to horizons where forward guidance is not credible.

While for calibration A, $k = 3$ periods were optimal in the benchmark case, the largest improvement can now be achieved for $k = 2$. By the same token, the optimal horizon in calibration B changed from $k = 2$ to $k = 1$.

What is also interesting is that under an incomplete interest rate pass-through, forward guidance is now far less credible than in the benchmark case as well as for a value of δ that is lower than one. We find that under calibration A, promises are only credible for up to $k = 2$, whereas, under calibration B, only one-period forward guidance can improve losses at the ZLB. Interestingly, for $k = 2$ an improvement could be achieved, but the central bank would defect from past promises and implement the time-consistent optimal policy plan instead.

Summing up our results, we find that the strength of the cost channel significantly changes our results, both quantitatively and qualitatively. A strong cost channel makes forward guidance less credible, whereas a weak cost channel pushes results in the direction of our benchmark results.

D. *Ignoring the Cost Channel*

So far, we compared the credibility of forward guidance between a standard New Keynesian model and an otherwise standard model that includes a cost channel. While the assumption that firms need to pay factors in advance and rely on financial intermediaries to do so before receiving revenues is convincing per se, this is not a standard assumption in New Keynesian models nowadays. Most researchers typically ignore the cost channel and rely on a standard textbook model, or add other frictions instead. But what happens if, instead, a cost channel is present, but mistakenly ignored by the central bank? To tackle this question, we apply a poor man's approach and assume that the central bank optimizes (4) subject to (1), (3) and a Phillips curve that reads

$$\pi_t = \beta E_t \pi_{t+1} + \kappa(\sigma + \eta) x_t,$$

where the term δi_t is now missing. However, the *true* Phillips curve is still given by (2), i.e. the cost channel actually exists but is overseen or ignored by the central bank. As a result, the corresponding targeting rule in state n is now given by

$$\pi_t = -\left(\frac{\lambda}{\kappa}\right)\frac{1}{\sigma + \eta} x_t.$$

If a present cost channel is not taken into account by the central bank, forward guidance can be less credible, because the central bank does not respond optimally in state n anymore. Because the equilibrium outcomes in the other states also depend on state n, the equilibrium values in these states will also change. We will therefore focus on the missing benefit that would have been achieved when the true model had been considered. Thus, it is straightforward to report the difference of the gain that would have been achieved, given that the central bank had behaved optimally subject to the model it perceives to be the true one. To put it differently, we actually derive the foregone benefit (improvement) of lower losses at the ZLB that would have been achieved, had the central bank taken the correct model into account.

This is done in Table (10) for our benchmark calibrations for s and q. Starting with $s = 0.98326$ and $q = 0.875$, we find that forward guidance can be credible even if the central bank responds to shocks based on a misperceived model. However, while the optimal horizon is still given by $k = 3$ which results in a gain of 0.564, the same promise would result in a gain of 0.753 in the true model. Consequently, the foregone benefit at the ZLB in terms of steady-state consumption is equal to 0.189 percent. Given that a cost channel is indeed present, this amounts to 25.1% of the gain at the ZLB. That is, while there is still an improvement over discretion through forward guidance, this benefit could have been even higher had the central bank taken the true model into account. Interestingly, we find that for all values of k that imply a credible promise in the true model, the gain at the ZLB is always lower if the central bank considers the wrong model as the true one.

Under calibration B, we find a slightly different result. With reversion to the ZLB being more likely, the wrong model yields an optimal horizon of $k = 3$, as opposed to the true model where $k = 2$ is optimal. In this case, however, the relevant foregone

Calibration A

	true model		*wrong model*		
k	Gain	Temptation	Gain	Temptation	Δ Gain
0	0.000	0.000	0.000	0.000	0.000
1	0.344	−0.229	0.244	−0.152	0.100
2	0.647	−0.424	0.446	−0.283	0.201
3	**0.753**	**−0.478**	**0.564**	**−0.354**	**0.189**
4	0.220	−0.077	0.467	−0.282	−0.244
5	*−2.411*	*1.807*	*−0.110*	*0.109*	*−2.301*

Calibration B

	true model		*wrong model*		
k	Gain	Temptation	Gain	Temptation	Δ Gain
0	0.000	0.000	0.000	0.000	0.000
1	0.105	−0.072	0.081	−0.055	0.024
2	**0.162**	**−0.105**	0.138	−0.091	0.024
3	0.094	−0.042	**0.140**	**−0.088**	−0.046
4	*−0.267*	*0.247*	0.028	−0.003	−0.295
5	*−1.334*	*1.077*	*−0.304*	*0.240*	−1.030

Table 10: Results for different values of s. In all cases, $q = 0.9$. All values are expressed in percent of steady-state consumption. Rows with values in bold font indicate that under this horizon k, the gain is maximized and forward guidance is credible. Rows with values in italics correspond to horizons where forward guidance is not credible.

benefit has to be calculated as Δ Gain=0.162–0.140=0.022, which is equal to 13.58% of the gain at the ZLB under the true model. However, overall, it stands out that given the central bank behaves optimally according to the model it believes to be the true one, we find that ignoring the cost channel might be costly in terms of aggregate steady-state consumption. This effect is, however, is much stronger, the higher the expected duration of the ZLB episode.[23]

VI Conclusion

A common finding in the literature is that the zero lower bound does not have serious consequences, given that a central bank can commit to future actions. If instead a discretionary policy environment is assumed, the zero lower bound has adverse welfare effects. Any promise of the central bank to keep the interest rate at zero once the the ZLB is no longer binding is not credible in this case because the central bank will renege on promises made in the past. This result changes if the ZLB is occasionally binding, i.e., if future ZLB episodes are possible.

In this paper, we compare the credibility of forward guidance in a standard New Keynesian textbook model and an extension of the model that allows for a cost channel.

[23] We did the same experiment with values of $\delta > 1$, i.e., under the assumption of an imperfect interest rate pass-through from the central bank's policy rate to the lending rate. In this case, the effect found above is much higher, resulting in sizable foregone improvements at the ZLB.

Since under a cost channel, any shock on the demand side always creates a trade-off between the stabilization of the output gap and inflation, we find that the supply-side effect introduced by the cost channel makes forward guidance, at best, as credible as in the absence of the cost channel. At the same time, however, we show that this trade-off makes forward guidance more powerful under the optimal horizon. Ignoring the cost channel can be costly and results in foregone steady-state consumption. If firms only partially rely on financial intermediaries when paying factors of production in advance, the cost channel does not make a difference concerning the optimal horizon of forward guidance. If, instead, the interest rate pass-through from the central bank's policy rate to the lending rate is instead incomplete, forward guidance lacks credibility when the promised length of zero interest rates is too long.

Digging deeper, one could use a more sophisticated modeling of financial frictions, in which the size of the cost channel is pinned down to agency costs (see, De Fiore, Fiorella and Tristani, 2013). Another interesting topic for future research would be to endogenize the relation between forward guidance and the financial system. In the current model, there is no such feedback. If this assumption would be relaxed instead and forward guidance changes the willingness of banks to lend, this should feed back onto the supply side.

References

Adam, Klaus and Roberto Billi, "Optimal Monetary Policy Under Commitment With a Zero Bound on Nominal Interest Rates," *Journal of Money, Credit and Banking*, 2006, *38* (7), 1877–1905.

— **and** — , "Discretionary Monetary Policy and the Zero Lower Bound on Nominal Interest Rates," *Journal of Monetary Economics*, 2007, *54* (3), 728–752.

Barth, Marvin and Valerie Ramey, "The Cost Channel of Monetary Transmission," *NBER Macroeconomics Annual*, 2001, *16*, 199–240.

Bernanke, Ben, "The New Tools of Monetary Policy," *American Economic Review*, 2020, *110* (4), 943–983.

Bilbiie, Florin, "Optimal Forward Guidance," *American Economic Journal: Macroeconomics*, 2019, *11* (4), 310–345.

Billi, Roberto, "Optimal Inflation for the US Economy," *American Economic Journal: Macroeconomics*, 2011, *3* (3), 29–52.

— , "A Note on Nominal GDP Targeting and the Zero Lower Bound," *Macroeconomic Dynamics*, 2017, *21* (8), 2138–2157.

Chattopadhyay, Siddhartha and Taniya Ghosh, "Taylor Rule Implementation of the Optimal Policy at the Zero Lower Bound: Does the Cost Channel Matter?," *Economic Modelling*, 2020, *89*, 351–366.

Chowdhury, Ibrahim, Mathias Hoffmann, and Andreas Schabert, "Inflation Dynamics and the Cost Channel of Monetary Transmission," *European Economic Review*, 2006, *50* (4), 995–1016.

Clarida, Richard, Jordi Gali, and Mark Gertler, "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature*, 1999, *37* (4), 1661–1707.

Cochrane, John H, "The New-Keynesian Liquidity Trap," *Journal of Monetary Economics*, 2017, *92*, 47–63.

De Fiore, Fiorella and Oreste Tristani, "Optimal Monetary Policy in a Model of the Credit Channel," *The Economic Journal*, 2013, *123* (571), 906–931.

Eggertsson, Gauti B, "What Fiscal Policy Is Effective at Zero Interest Rates?," *NBER Macroeconomics Annual*, 2011, *25* (1), 59–112.

— **and Michael Woodford**, "Zero Bound on Interest Rates and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, 2003, *2003* (1), 139–233.

Jung, Taehun, Yuki Teranishi, and Tsutomu Watanabe, "Optimal Monetary Policy at the Zero-Interest-Rate Bound," *Journal of Money, Credit, and Banking*, 2005, *37* (5), 813–835.

Kiley, Michael, "Policy Paradoxes in the New Keynesian Model," *Review of Economic Dynamics*, 2016, *21*, 1–15.

McKay, Alisdair, Emi Nakamura, and Jón Steinsson, "The Power of Forward Guidance Revisited," *American Economic Review*, 2016, *106* (10), 3133–3158.

Nakata, Taisuke, "Reputation and Liquidity Traps," *Review of Economic Dynamics*, 2018, *28*, 252–268.

Nakov, Anton, "Optimal and Simple Monetary Policy Rules with Zero Floor on the Nominal Interest Rate," *International Journal of Central Banking*, 2008, *104* (2), 73–127.

Pathberiya, Lasitha R C, "Optimal Monetary Policy at the Zero Lower Bound on Nominal Interest Rates in a Cost Channel Economy," *The University of Queensland, School of Economics Discussion Paper Series*, 2016, *568.*

Ravenna, Federico and Carl E Walsh, "Optimal Monetary Policy with the Cost Channel," *Journal of Monetary Economics*, 2006, *53* (2), 199–216.

Tillmann, Peter, "Do Interest Rates Drive Inflation Dynamics? An Analysis of the Cost Channel of Monetary Transmission," *Journal of Economic Dynamics and Control*, 2008, *32* (9), 2723–2744.

Walsh, Carl E, "Simple Sustainable Forward Guidance at the ELB," *unpublished, University of California, Santa Cruz*, 2018.

Woodford, Michael, *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press, 2003.

Appendix

A Solving the Model

Throughout the paper, we solve the model numerically. In the absence of forward guidance, the system of equations reads

$$
\begin{aligned}
x_n^d &= \left[s x_n^d + (1-s)x_z^d\right] - \sigma^{-1}\left(i_n^d - \left[s\pi_n^d + (1-s)\pi_z^d\right] - \rho\right)\\
\pi_n^d &= \beta\left[s\pi_n + (1-s)\pi_z\right] + \kappa\left[(\sigma+\eta)\, x_n^d + \delta i_n^d\right]\\
x_z^d &= \left[q x_z^d + (1-q)x_n^d\right] + \sigma^{-1}\left(\left[q\pi_z^d + (1-q)\pi_n^d\right] + r_z\right)\\
\pi_z^d &= \beta\left[q\pi_z^d + (1-q)\pi_n^d\right] + \kappa\,(\sigma+\eta)\, x_z^d\\
&\lambda x_n^d + \kappa\left[\sigma(1-\delta)+\eta\right]\pi_n^d.
\end{aligned}
$$

Define $\mathbf{y} = [i_n,\ x_n,\ \pi_n,\ x_z,\ \pi_z]$ and $\mathbf{c} = [\rho\sigma^{-1},\ 0,\ r_z\sigma^{-1},\ 0,\ 0]$. The appropriate (5×5) matrix with reduced-form coefficients then reads

$$
\mathbf{A} = \begin{bmatrix}
\sigma^{-1} & 1-s & -s\sigma^{-1} & -(1-s) & -(1-s)\sigma^{-1}\\
-\kappa\delta & -\kappa(\sigma+\eta) & 1-\beta s & 0 & -\beta(1-s)\\
0 & -(1-q) & -(1-q)\sigma^{-1} & 1-q & -q\sigma^{-1}\\
0 & 0 & -\beta(1-q) & -\kappa(\sigma+\eta) & 1-\beta q\\
0 & -\lambda & -\kappa\left[\eta+\sigma(1-\delta)\right] & 0 & 0
\end{bmatrix},
$$

such that the model in companion form reads

$$\mathbf{A}\mathbf{y} = \mathbf{c}.$$

The solution of the model can therefore be easily obtained as

$$\mathbf{y} = \mathbf{A}^{-1}\mathbf{c}.$$

Solving the model under forward guidance works exactly as above, except that we need two additional equations for each k, i.e. for each additional state.

B Additional Plots

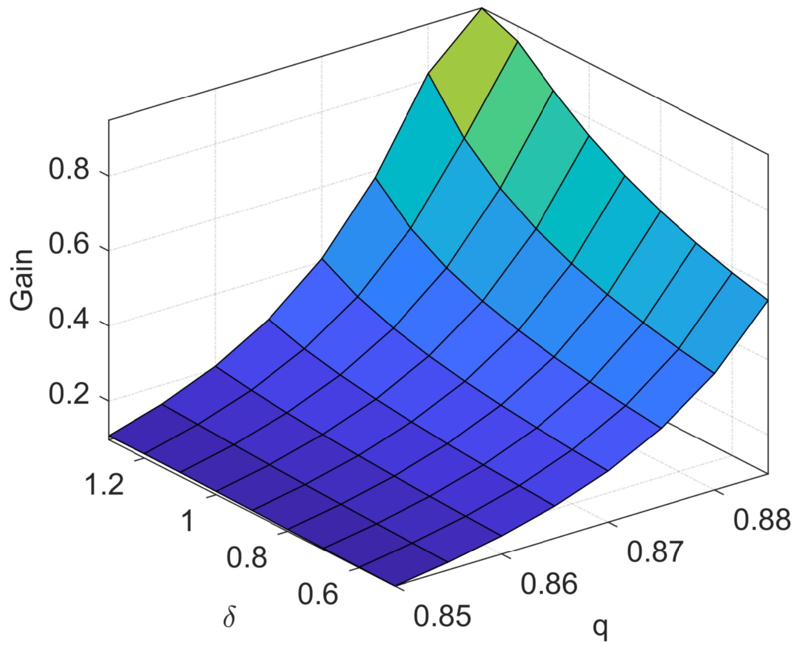

Figure 11: Gain as a function of q and δ for $s = 0.9799$. Values are expressed in percent of steady-state consumption.

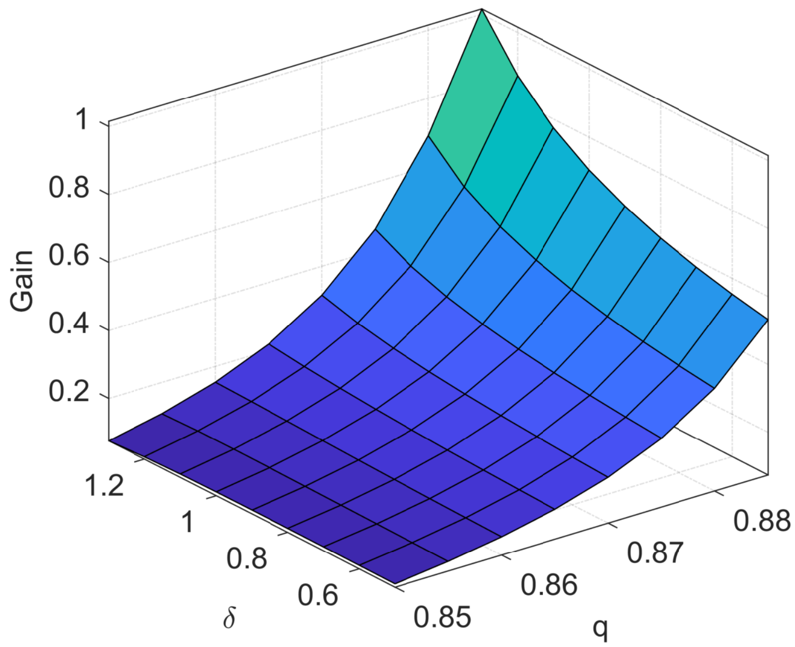

Figure 12: Gain as a function of q and δ for $s = 0.9833$. Values are expressed in percent of steady-state consumption.

Essay VI:

Pandemic Shocks and Household Spending

This paper is published as:

Finck, David and Peter Tillmann, "Pandemic Shocks and Household Spending," *Oxford Bulletin of Economics and Statistics*, 2022, *84*(2), 273–299.

This paper was presented at the following workshops and international conferences:

I: 14th International Conference on Computational and Financial Econometrics (Virtual)
Date: December 2020
Presenter: David Finck

II: 14th RGS Doctoral Conference (Virtual)
Date: March 2021
Presenter: David Finck

Acknowledgement

We thank the editor of this journal and two anonymous reviewers for very insightful comments. We are grateful to Carola Binder, Patrick Hürtgen, Daniel Grabowski, Salah Hassanin, Peter Winker and the team from Opportunity Insights for helpful discussions.

Pandemic Shocks and Household Spending

David Finck[*] Peter Tillmann[†]

Abstract

We study the response of daily household spending to the surprise number of fatalities of the COVID-19 pandemic, which we label as pandemic shock. Based on daily forecasts of the number of fatalities, we construct the surprise component as the difference between the actual and the expected number of deaths. We allow for state-dependent effects of the shock depending on the position on the curve of infections. Spending falls after the shock and is particularly sensitive to the shock when the number of new infections is strongly increasing. If the number of infections grows moderately, the drop in spending is smaller. We also estimate the effect of the shock across income quartiles. In each state, low-income households exhibit a significantly larger drop in consumption than high-income households. Thus, consumption inequality increases after a pandemic shock. Our results hold for the US economy and the key US states. The findings remain unchanged if we choose alternative state-variables to separate regimes.

Keywords: COVID-19, Consumption, Smooth-Transition Model, Consumption Inequality, State-Dependence

JEL classification: E21, E32, I10

[*]University of Giessen, email: david.finck@wirtschaft.uni-giessen.de

[†]University of Giessen, email: peter.tillmann@wirtschaft.uni-giessen.de

I Introduction

The global spread of the COVID-19 pandemic since January 2020 led to a sharp contraction of economic activity in almost all economies affected by the virus. Between January and April 2020, real personal consumption expenditures declined by more than 15%. With personal consumption expenditures accounting for 68% of US GDP in 2019, this decline in spending casts shadow on overall economic activity in 2020. Consumption recovered in May and June, partly driven by government transfers which led to an increase in real disposable income.

In this paper, we provide an analysis of the causal effect of the pandemic on household spending. Spending, such as consumption in general, should mostly be driven by unexpected shocks. According to the theory of permanent income, predictable fluctuations in future income should prompt households to tap the capital market and smooth consumption, such that consumption exhibits very little fluctuations.[1]

We look at the unexpected number of fatalities of the pandemic and analyze how it affects spending decisions. We draw on forecasts of the number of fatalities due to COVID-19 in the US provided by Gu (2020) and contrast the one day-ahead forecast with the actual number of deaths. A positive forecast error is consistent with an underprediction of the number of fatalities or a surprise in the severity of the pandemic, respectively. Importantly, the number of reported deaths exhibits transient drops on weekends, typically followed by increases during the week. We therefore purify our shock by regressing it on a set of dummies for each day of the weak. We refer to this series of unexpected deaths as a pandemic shock and use it as the key explanatory variable for household spending. As a matter of fact, this is only one possible way to model the surprising spread of the pandemic. Other surprises, such as the development of a vaccine, could also have effects on household spending.[2]

Our measure of household spending is provided by Chetty et al. (2020) and consists of debit and credit card transactions in the US. The key advantage of the data is timeliness. We can track spending on a daily frequency for the entire US economy as well as for US states. In a series of local projections, see Jordà (2005), we estimate the response of spending to a pandemic shock.

There are at least three channels through which a pandemic shock can affect spending. First, an adverse pandemic shock could prompt households to restrain consumption voluntarily. This is because the virus spreads through social interaction such as shopping in retail stores, dining or entertainment. Anxious households could reduce these activities even before official lockdown measures are in place.[3] Second, households might be barred from consumption due to a lockdown of selected activities

[1] See Jappelli and Pistaferri (2010) for a survey of the field.
[2] We are grateful to an anonymous referee for this point.
[3] Goolsbee and Syverson (2021) show that consumer behavior during the pandemic is more driven by fear of infection than formal restrictions.

or even shelter-in-place orders. An adverse pandemic shock makes these measures more likely. Third, households could perceive an unexpected change in future income and adjust their spending accordingly. Even if a household is not itself affected by the virus, the future of entire industries is at risk. Workers in the service sector, for example, cannot resort to working from home and experience a large drop in future income.[4]

While we cannot disentangle these transmission channels, we take account of an important property that all three channels have in common: the effect of a pandemic shock should be stronger if the virus spreads more rapidly. The more widespread the virus is, the larger the reluctance to shop offline, the more likely stricter lockdown measures and the more severe the drop in future income will be. Thus, the effect of the pandemic shock should depend on the position of the economy on the infection curve.

Therefore, we generalize our model and allow the pandemic shock to have regime-dependent effects. In our baseline setting, we chose the growth of the daily number of new infections as our state variable. This figure is omnipresent, especially in the media, and provides information on where the economy stands on the infection curve. The transition between states is driven by either a non-parametric model introduced by Born et al. (2020) or a parametric approach proposed by Auerbach and Gorodnichenko (2012).

We show that a pandemic shock originating when the number of new infections is growing fast has a strongly negative effect on spending. We find a significant drop of about 0.8% in spending after a pandemic shock of one standard deviation. The drop in consumption is consistent with recent macroeconomic models of the effect of income expectations and uncertainty on consumption, see Dietrich et al. (2022), or the feedback between the spread of the pandemic and macroeconomic aggregates, e.g. Eichenbaum et al. (2021). Following the strategy of Gorodnichenko and Lee (2020), a forecast error variance decomposition shows that the pandemic shock explains more than 20% of fluctuations in spending.

If the shock occurs in a situation in which the virus spreads less rapidly, spending drops by 0.4% only with the peak response occurring after one week. Throughout the paper, we find that the nexus between spending and the pandemic shock is strongly depending on the underlying regime. In almost all cases, we can reject the null hypothesis of equal spending responses across regimes. We estimate the model not only for the whole US economy, but also for the 10 largest US states. Across all states, the regime-dependent sensitivity of spending to pandemic shocks is very similar. The results remain unchanged if we use alternative state variables such as the level of new infections rather than the growth rate of infections.

[4]In a survey conducted early in the pandemic, Binder (2020) finds that households expect an increase in unemployment due to the pandemic.

We also study the spending response across income quartiles. We use data on spending for residents of ZIP codes with low, middle and high median household income. This allows us to estimate the response of household spending across income groups to a pandemic shock. The first two of the three transmission channels discussed before, voluntary and forced consumption restraint, should apply equally to high-and low-income households.

The third channel, however, should imply that low-income households reduce their spending by more compared to high-income households. This is because the drop in lifetime income should be particularly pronounced for low-productivity workers, e.g. workers in the service sector.[5] We do indeed find that in the regime with a strong growth of the number of infections, high-income households reduce their spending by 0.75%, while low-income households cut expenditures by 1%. This is remarkable because the initial fall in spending was larger for high-income households as documented by Chetty et al. (2020). Our results suggest that the economic burden of the pandemic in terms of consumption falls more on low-income households.[6] The difference in spending responses is highly statistically significant in both regimes. Thus, the pandemic contributes to a growing of consumption inequality. We also show that employment of low-income households is significantly more sensitive to the shock compared to mid-income households. High-income employment even increases after the shock. Hence, the evolution of employment supports our findings for household spending across income quartiles.

This paper contributes to the recent work on household behavior based on innovative datasets. In an early paper, Baker et al. (2020) use transaction-level data for the US in order to document the changes in consumption patterns after the outbreak of the coronavirus. Cox et al. (2020) extend this line of research and shed light on the response of consumption and saving across the income distribution. Using transaction-level data from the largest Danish bank, Andersen et al. (2020) show that the decline in spending increases in the exposure of households to the economic consequences of the pandemic. Hacioglu Hoke et al. (2020) use data from a fintech company based in the UK to track the behavior of spending. These authors also document the build-up of financial stress as well as consumption and income inequality across households. Carvalho et al. (2021) use six billion transactions of customers of Spain's second-largest bank to track consumption over the crisis. Coibion et al. (2020) estimate the effect of lockdowns on spending and household expectations based on survey data. They make use of the asynchronous timing of lockdown measures in order to identify a causal effect. The occurrence of the first corona infection is used to instrument local

[5] Even if the drop in income were equal across income groups, we expect the marginal propensity to consume (MPC) to be higher for low-income quartiles. In fact, Karger and Rajan (2020) track spending of recipients of governmental transfer payments during the COVID-19 pandemic. They find an MPC of 0.68 for hand-to-mouth consumers and 0.23 for savers.

[6] See Mongey et al. (2021) for an analysis of the effect of social distancing across workers. They find significant differences in the burden from social distancing.

lockdown restrictions. They find that lockdown restrictions explain most of the fall in consumer spending since March 2020.[7]

Most of these papers, with the exception of Coibion et al. (2020), provide descriptive evidence based on massive new datasets or estimate the response of spending to observable events. Instead, we aim at estimating the sensitivity of spending to unexpected changes in the severity of the pandemic.

The paper is organized as follows. Section II presents the data and discusses the derivation of our pandemic shock. Section III lays out our estimation strategy. Our results are discussed in Section IV. Section V presents results on the US state level. Section VI presents results for alternative state variables. Section VII illustrates potential transmission channels and Section VIII concludes.

II Data

To investigate the response of consumption to the unexpected spread of the pandemic, we rely on two data sets. The first contains information on daily household spending since the outbreak of the coronavirus and the second reports daily historical forecasts on the number of fatalities due to COVID-19.

A. *Household Spending*

Throughout the paper, the dependent variable is a measure of household spending. We have daily observations ranging from April 3 to October 4. We use the series provided by Chetty et al. (2020), which are open source and available at `https://tracktherecovery.org`.[8]

We also have spending broken down into ZIP codes with high, middle and low median income. Below, we will refer to these subgroups of households as high- and low-income households, although we do not have information on household income, but only on median income of the ZIP code of residency.

Due to the high volatility in daily spending, the publicly available series are 7-day moving averages. Furthermore, data on consumer spending exhibits strong weekly fluctuations which are autocorrelated across years. To account for this, Chetty et al. (2020) divide all spending series by their corresponding value from 2019. Lastly, the seasonally adjusted data are indexed to their pre-pandemic level, i.e. the mean of the 7-day moving average from January 8-28. Hence, our series are given in percent, such that a value of two percent in t corresponds to an increase of spending by two percent relative to its average value in January.

[7] Alexander and Karger (2020) analyze consumer spending and cellphone records in the US and show the causal effect of stay-at-home orders on spending.

[8] Chetty et al. (2020) collect the data on spending from Affinity Solutions Inc. This company aggregates information on credit and debit card spending. The data is available for nation-wide aggregates as well as for each US state.

Figure 1: Total spending and spending by income quartile

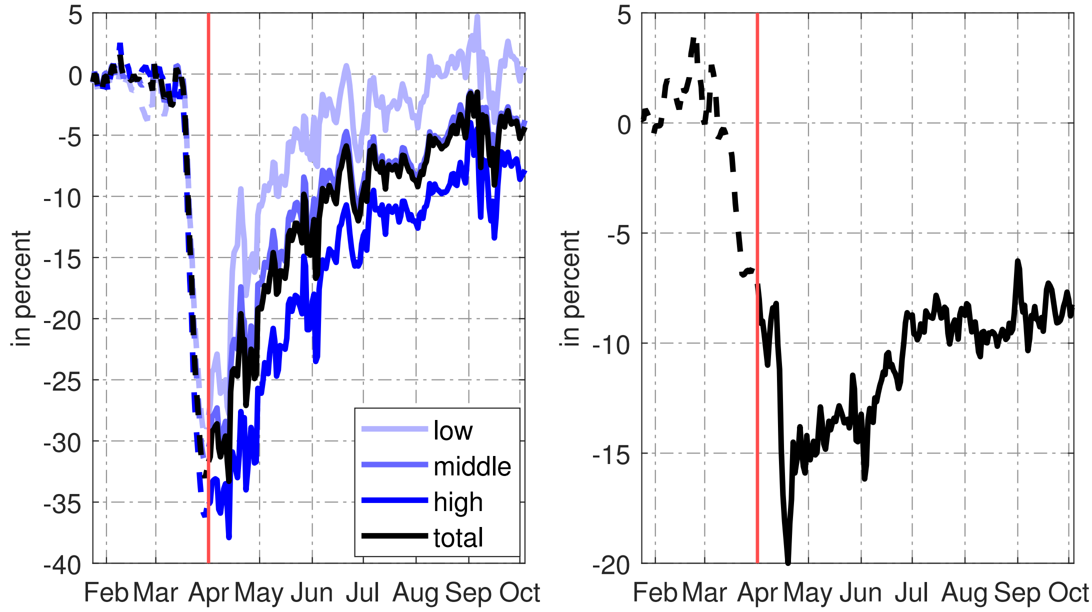

Notes: The left panel shows the difference of actual spending relative to its level in January 2020 by customers living in ZIP codes with different income classes, namely high (top quartile) median income, middle (middle two quartiles) median income as well as low (bottom quartile) median income. The right panel shows the difference in spending of customers living in ZIP codes with high median income and customers living in ZIP codes with low median income. In both samples, the start of our estimation sample (April 3rd) is highlighted by the red vertical lines.

The left panel of Figure (1) shows that for all income quartiles, household spending fell sharply in mid-March, when national emergency was declared. In early April, spending fell by 36.4% for high-income households, 32% for middle-income households and 29.8% for low-income households. The right panel shows relative spending, i.e. the difference between spending of high- and low-income households. The reversion of spending to the pre-pandemic level differs remarkably, with the level of low-income households being almost back to the pre-pandemic level. Spending from high-income households fell more and recovered less – a finding that we need to keep in mind as we show below that the sensitivity of high-income households to pandemic shocks is actually smaller than that of low-income households.

At this stage, a discussion of the limitations of the spending data is warranted. The data tracks credit and debit card transactions only. It covers approximately 10% of all credit and debit card transactions in the US. There are two main limitations compared to alternative survey data available on a much lower frequency: First, the focus on card-based transactions gives too small a weight to households with financial constraints. These households might not have access to credit card services and are not covered. This matters for the comparison of spending across income groups. Moreover, different responses of rich and poor households can also reflect different consumption baskets. Affluent households might substitute their restaurant visits by grocery shopping. While both are payed by credit card, the substitution leads to a drop in the value of transactions. Poor households, in contrast, stick to grocery shopping. Second, sectors with frequent credit card payments such as restaurants and hotels might be over-weighted, while sectors with infrequent card payments such as vehicle sales or financial services might be underrepresented. We should keep these limitations in mind when discussing the results.

B. *The Surprise Number of Fatalities*

Consumption should respond to the unexpected severity of the pandemic. Hence, in order to investigate the consumption response, we need a series of the surprise component of the spread of the pandemic. We formulate the surprise in terms of the unexpected number of fatalities due to COVID-19, i.e. the difference between expected and realized deaths.

We retrieve daily real-time projections on deaths and the unrevised reported number of deaths due to COVID-19 in the US from Gu (2020). This data is open source and can be downloaded from www.covid19-projections.com. The author takes a (machine learning) data-driven approach rooted in epidemiology to forecast infections and deaths from the coronavirus epidemic in the US (and around the world).[9] These forecasts have been covered by almost all major US media outlets. Of course, one could consider using infection numbers instead of death numbers to derive a shock. The problem, however, is that forecasting models do not try to forecast the reported numbers of infected people, but the actual number of infected people. That is, since the reported number of newly infected persons is likely to be significantly or systematically lower than the true number of newly infected persons, there is thus no forecast value to compare with the reported infected numbers. Note that this effect should not occur with the reported number of fatalities.

Importantly, we do not only have the latest forecast, but also the historical forecasts. The forecasts are updated on a daily basis. We use this data to derive a pandemic shock, i.e. the unexpected number of deaths due to COVID-19. To do so, denote $\mathbf{d}_{t|t-1}$ the forecast made in $t-1$ for deaths occurring in t. Thus, we focus on one day-ahead forecasts. Our pandemic shock is calculated as the difference between the actual outcome for t and the forecast number of deaths

$$\mathbf{e}_t = \mathbf{d}_t - \mathbf{d}_{t|t-1}. \tag{1}$$

That is, our pandemic shock is the difference between today's number of reported deaths and yesterday's forecast for today.[10]

[9]Details on the forecasting model, including assumptions on the model parameters, are available at https://covid19-projections.com/model-details

[10]We cannot rule out that some households have different beliefs about how the virus will progress and

Note that the number of reported deaths exhibits transient drops on weekends, typically followed by increases during the week. We therefore purify our shock by regressing it on a set of dummies for each day of the weak. Formally, we regress

$$\mathbf{e}_t = F(z_{t-1})\left(\gamma^I \mathbf{Q}_t + \phi^I \mathbf{e}_{t-1}\right) + (1 - F(z_{t-1}))\left(\gamma^{II} \mathbf{Q}_t + \phi^{II} \mathbf{e}_{t-1}\right) + \varepsilon_t, \quad (2)$$

where $F(z_{t-1})$ is a smooth transition function which we will explain in detail in the next section. $\mathbf{Q}_t$ contains dummies for each day of the week as well as a dummy which is equal to one until the end of the first wave of fatalities due to the coronavirus. By doing so, we account for the fact that during the first wave, uncertainty about the pandemic was high which likely leads to larger forecast errors. We therefore set this dummy to zero from June 26 onwards. This cutoff date roughly represents the end of the first wave of fatalities.

Note that the estimated residuals for ε_t can be interpreted as the pandemic shock, which cannot be explained neither by the set of dummies captured in $\mathbf{Q}_t$ nor by yesterday's forecast error.[11]

Table 1: Descriptive statistics for shocks

			RAW SHOCK				
MIN	MAX	MEAN	MEDIAN	5^{th}	95^{th}	Q-STAT.	p–VAL.
-1007	2416	41.3	78	-559.8	688.4	313.1	0.000
			PURIFIED SHOCK				
MIN	MAX	MEAN	MEDIAN	5^{th}	95^{th}	Q-STAT.	p–VAL.
-653.8	1993.2	0.00	-27.12	-358.9	299.9	11.0	0.683

Notes: Numbers are in deaths per day. Shocks are calculated based on unrevised real-time data. The last two columns report Q-statistics and p–values for a Ljung-Box test with the null hypothesis of zero autocorrelation up to 14 lags.

Table (1) reports some descriptive statistics for both the raw shock $\mathbf{e}_t$ and the purified shock ε_t. It is noteworthy that a large fraction of outliers can be explained by our daily dummies. This is because for the purified shock, the 5th and 95th percentiles are much closer to zero than for the raw shock. Also for the extreme values (minimum and maximum) of our shock, a notable fraction seems to be grounded on the seasonal patterns that is apparent in the reported number of deaths. The purified shock is almost iid and has no serial correlation for up to fourteen lags.[12] Finally, in order to interpret our shock in terms of standard deviations, we subtract the mean and divide the series by the sample standard deviation.

might even mistrust the official data about new cases and deaths, e.g. Fetzer et al. (2020).

[11] We checked whether other control variables have explanatory power, including the daily number of cases and a lag polynomial for up to seven lags. However, it turns out that the dynamics jointly explained by these variables is negligible. We therefore exclude them from our regression.

[12] We also applied various tests to detect heteroskedasticity in our series, including an Engle's ARCH test. The test results imply that our shock series is homoscedastic.

Figure 2: Raw vs purified shock

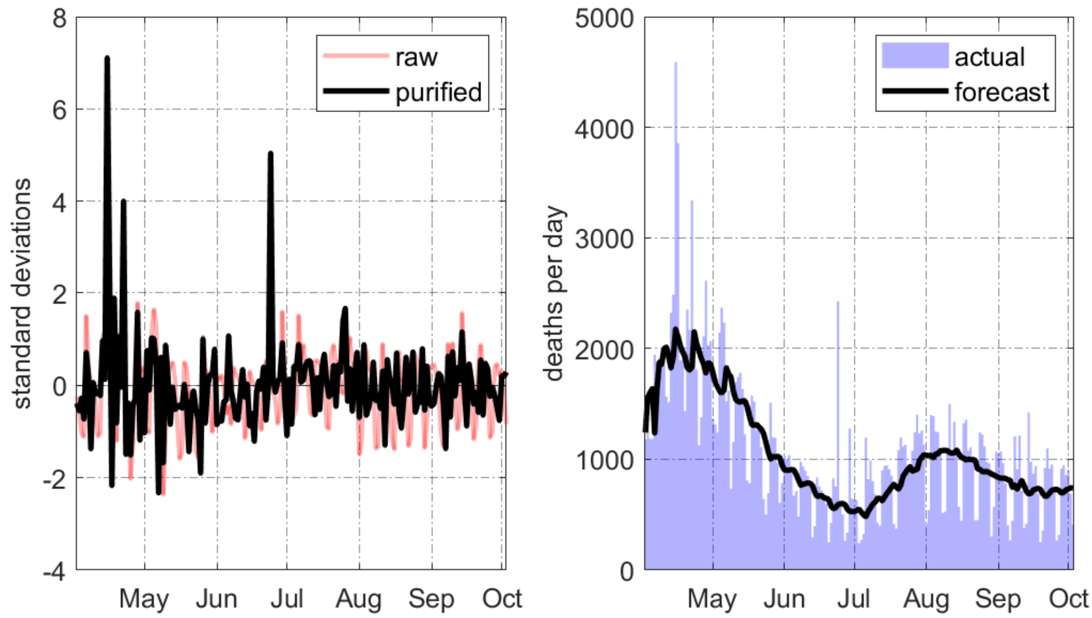

Notes: The left panel shows the raw shock calculated as $\mathbf{e}_t = \mathbf{d}_t - \mathbf{d}_{t|t-1}$ as well as the shock after our purification procedure, i.e. ε_t. The right panel shows the reported number of deaths per day (in real-time) of people infected with the coronavirus (purple bars) as well as the real-time one-step-ahead forecasts.

Figure (2) shows the underlying data we use to derive the shock as well as our shock series. Starting with the right panel, the bars show the actual daily reported number of deaths over time. The black solid line corresponds to the one-step ahead forecasts. One can immediately recognize the seasonal pattern mentioned before. The reported number of deaths increased up to 2000 per day until the end of April and started to steadily decrease afterwards, with daily deaths (on average) below 500 by the end of June. However, from early July, the number of daily deaths started to increase again until early August. Interestingly, the forecasts follow the overall direction of the actual number of reported deaths with positive forecast errors being equally likely to negative forecast errors. The left panel shows the raw and the purified shock series constructed as described above. While the raw series clearly exhibits seasonal patterns, the purified shock series now looks very much like an iid process. Moreover, especially from May onwards, we can now see that a significant fraction of the swings disappears when taking seasonality into account.

III Methodology

We investigate the effects of pandemic shocks via local projections as proposed by Jordà (2005). Local projections provide a flexible framework and are easy to implement. Moreover, they offer a straightforward way to condition the short-run effects of pandemic shocks on the state of the pandemic.

A. *Setup*

The linear model of departure reads

$$y_{t+h} = \alpha_h + \boldsymbol{\beta}_h \varepsilon_t + \delta_h t + \boldsymbol{\gamma}_h \mathbf{x}_t + \boldsymbol{\varphi}_h \mathbf{D}_t + u_{t+h}, \tag{3}$$

where y_{t+h} is the dependent variable at time $t+h$, which responds to a shock ε_t occurring in t. In our model, the dependent variable is household spending and ε_t is the pandemic shock introduced before. Our coefficient of interest is $\boldsymbol{\beta}_h$. The coefficient α_h corresponds to a fixed effect at horizon h and δ_h measures the effect of a deterministic linear trend. The vector $\boldsymbol{\gamma}_h$ contains the effects of the lagged endogenous variable and other control variables at horizon h captured in the vector $\mathbf{x}_t$ and $\boldsymbol{\varphi}_h$ contains the effects of daily dummy variables. Finally, u_{t+h} is assumed to have a zero mean and a (strictly) positive variance.

Our vector $\mathbf{D}_t$ in (3) includes the stringency index provided by researchers from the University of Oxford as well as two dummy variables to account for (1) the subsequent phase of the CARES act, i.e. the Paycheck Protection Program and Healthcare Enhancement Act, and (2) the four FOMC meetings since April.[13] In our baseline setting, $\mathbf{x}_t$ includes three lags of the endogenous variable as well as one lag of the Economic Policy Uncertainty Index (EPU). This lag structure is the recommendation of the Bayesian Schwarz Criterion.[14]

The model presented before is linear. We now generalize the model to allow for state-dependent effects, that is we condition the impact of the shock on different regimes. Our preferred version throughout this paper conditions the response on the growth rate of new infections. Therefore, we estimate a smooth transition model of the form

$$\begin{aligned} y_{t+h} = {} & F(z_{t-1})\left(\alpha_h^I + \boldsymbol{\beta}_h^I \varepsilon_t + \boldsymbol{\gamma}_h^I \mathbf{x}_t\right) + \cdots \\ & \cdots + (1 - F(z_{t-1}))\left(\alpha_h^{II} + \boldsymbol{\beta}_h^{II} \varepsilon_t + \boldsymbol{\gamma}_h^{II} \mathbf{x}_t\right) + \delta_h t + \boldsymbol{\varphi}_h \mathbf{D}_t + u_{t+h} \end{aligned} \tag{4}$$

where the fixed effects, the effects of controls and the lagged endogenous variable captured in $\mathbf{x}_t$ as well as the effect of our shock are now allowed to differ across regimes I and II at each horizon h, respectively. That is, the indicator function $F(z_{t-1})$, which lies between zero and one, determines the weight of each regime, whereby $F(z_{t-1})$ depends on outcomes of the state variable z_{t-1}, which in our case is the growth rate of

[13] The stringency index is meant to measure the strictness of policies restricting people's behavior and lies between 1 and 100. The data is available on a daily frequency at https://www.bsg.ox.ac.uk/research/research-projects/coronavirus-government-response-tracker. The index is aggregated from 17 indicators of government responses, economic policies and health system policies.

The stimulus package mentioned above provided $484 bn in additional funding to the existing Paycheck Protection Program and was signed to law by President Trump on April 24.

[14] The results are insensitive to using the Akaike Information Criterion instead.

new infections. Note that we condition the response of our endogenous variable on the state variable in $t-1$ rather than in t. By doing so, we follow the literature and rule out that our shock has a contemporaneous effect on our state variable.

In effect, the response of our endogenous variables to a shock is a weighted average of regimes I and II conditional on z_{t-1} and reads

$$\left.\frac{\partial y_{t+h}}{\partial \varepsilon_t}\right|_{z_{t-1}} = F(z_{t-1})\beta_h^I + (1 - F(z_{t-1}))\,\beta_h^{II}. \tag{5}$$

In the next subsection, we will describe the specification of $F(z_{t-1})$ in detail. However, it is important to note that our framework allows us to easily compare the sensitivity to shocks across both regimes without making explicit assumptions (as in the case of VAR models) on the economy staying in either regime I or II. That is, we can draw inference on the difference between β_h^I and β_h^{II} based on $t-$type tests.

B. *State-Dependent Dynamics*

Our approach follows Born et al. (2020) and relies on specifying the transition function $F(z_t)$ based on the empirical cumulative density function (CDF)

$$F(z_t) = \frac{1}{T}\sum_{t=1}^{T} \mathbb{1}_{z_j<z_t}, \tag{6}$$

where T is the sample size and $\mathbb{1}_{z_j<z_t} = 1$ if $z_j < z_t$ and zero otherwise. That is, $\mathbb{1}_{z_j<z_t}$ denotes the indicator function of the event $z_j < z_t$. We refer to this approach as non-parametric, as we do not need to specify parameters driving the transition.

We choose the weekly growth rate of new infections as our state variable. Since the outbreak of the pandemic, numbers on new infections are reported every day in almost all media outlets. Public policies geared towards "flattening of the curve" made this statistic particularly popular. The left panel of Figure (3) shows the weekly growth rate of new infections with the coronavirus over time. The right panel shows the transition functions $F(z_t)$ based on the empirical cumulative density function over time.[15]

Starting with the left panel, after a well-pronounced decline in the growth rate falling from 60% to almost zero in mid of April, the growth rate of new infections fluctuates around zero until the beginning of June. Since then, we observe a strong increase in the growth rate with a rise in cases of above 40% in mid June, which declines again until early August. From early to mid September onwards, the second wave of infections came with a growth rate of newly-infected people sharply increasing again until the end of our sample.

[15] In order to further eliminate the noise in the data, we take a seven-day moving average before calculating the transition function. It must be stressed, however, that we get exactly the same results

Figure 3: Deriving the state variable

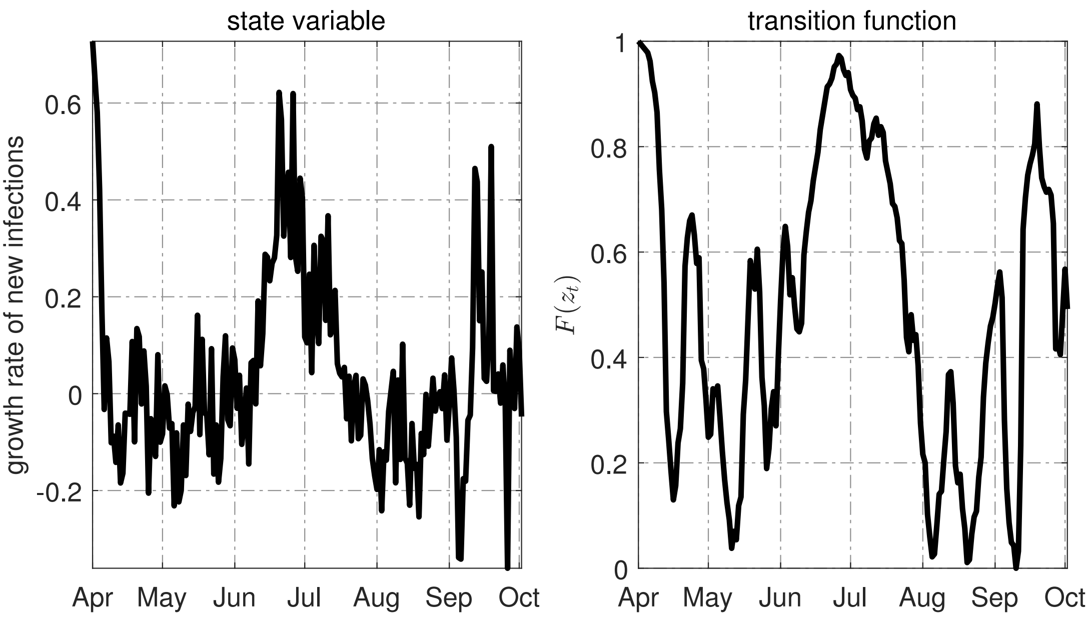

Notes: The left panel shows the weekly growth rate of daily cases of new infections with the coronavirus that causes COVID-19. The right panel shows the transition functions $F(z_t)$ based on parametric approach drawing on the empirical cumulative density function.

The right panel of Figure (3) shows the resulting transition function. While we see a sharp fall of $F(z_t)$ at the beginning of our sample, saying that the economy swiftly moves from regime I to regime II, the sudden rise in daily cases translates into a fast reversion from regime II to regime I until mid June, before it swiftly reverts to regime II until early August. At the end of our sample, we observe a sharp increase as a result of the beginning of the second wave of infections. Summing up, we observe that a high weight is attached to regime I throughout June and July as well as in mid to late September.

C. *Inference*

We regress the dependent variable at different horizons on the same set of control variables. This will likely result in autocorrelated residuals. In order to calculate standard errors that account for the possibility of serially correlated residuals both within and across equations, we follow the strategy of Tenreyro and Thwaites (2016) and Ramey and Zubairy (2018) and estimate seemingly unrelated regressions as proposed by Driscoll and Kraay (1998). That is, we estimate the parameters of interest of each equation separately and, in a second step, average the moment conditions across horizons $h = 0, ..., H$ when deriving Newey-West standard errors. As a result, Driscoll and Kraay (1998) standard errors account for autocorrelation across both, time t and horizons h.

if we use the non-smoothed growth rate.

Finally, we follow standard practice (see Jordà, 2005) and set the maximum autocorrelation lag for the Newey-West procedure to $L = h + 1$.[16]

IV Results

In this section, we first set out our baseline results. In the baseline setting, the idea is to uncover possible asymmetries across regimes in the responses of consumer spending to a standardized pandemic shock. That is, the baseline regression focuses on the effects of pandemic shocks conditional on the state of the infection curve. The sample size covers data from April 3 to October 4, consisting of 185 observations. After adjusting for leads and lags, the effective sample size consists of 168 observations. The section also reports results for different income levels as well as for the ten largest US states.

A. *Baseline Results*

Figure (4) shows the state-dependent impulse responses of total spending following a pandemic shock. Remember that throughout the paper, all spending variables are given in percentage deviation relative to the average level of January. That is, a value of one corresponds to an increase in spending of one percent relative to January. In the left column, the red-solid line depicts the impulse response coefficients in regime I following a pandemic shock. Regime I corresponds to a situation with a high growth rate of new infections. The shaded area corresponds to the 90 percent confidence bands based on Driscoll-Kraay standard errors. For the purpose of comparison, we also report the corresponding coefficients from the linear model (red-dashed line). Accordingly, the second column reports the corresponding values for regime II, i.e. the regime with a modest growth rate of new infections. The third column shows the t-statistics testing the null hypothesis $\beta_h^I - \beta_h^{II} = 0$ for adjacent horizons $h = 0, ..., H$, where the shaded area covers the t−critical values for a 90 percent confidence interval, i.e. ±1.645.

In this context, it is important to stress that a perfectly symmetric transmission of pandemic shocks would imply that $\beta_h^I = \beta_h^{II}\ \forall h = 0, ..., H$. In other words, a pandemic shock as identified in the previous section would have the same effects across both regimes. Contrary to this, we would refer to asymmetric effects when the difference between β_h^I and β_h^{II} is significantly different from zero.

Starting with the results from the left panel, i.e. the impulse response coefficients in regime I, we see a significant drop in total spending. That is, following a pandemic shock, total spending falls by about 0.6 percent on impact relative to its average value in

[16]Note that for each horizon h, our null hypothesis is $H_0 : (\beta_h^I - \beta_h^{II}) = 0$. Since we test the same null hypothesis for each $h = 0, ..., H$, one could argue that our t−statistics will result in a multiple testing problem. That is, if we test $H+1$ true null hypotheses at a significance level α, we would - on average - reject αH of them. However, as pointed out by Tenreyro and Thwaites (2016), the multiple testing problem is negligible when the t−statistics for adjacent horizons are correlated, which is what we will see when we discuss our results.

Figure 4: Response of total spending

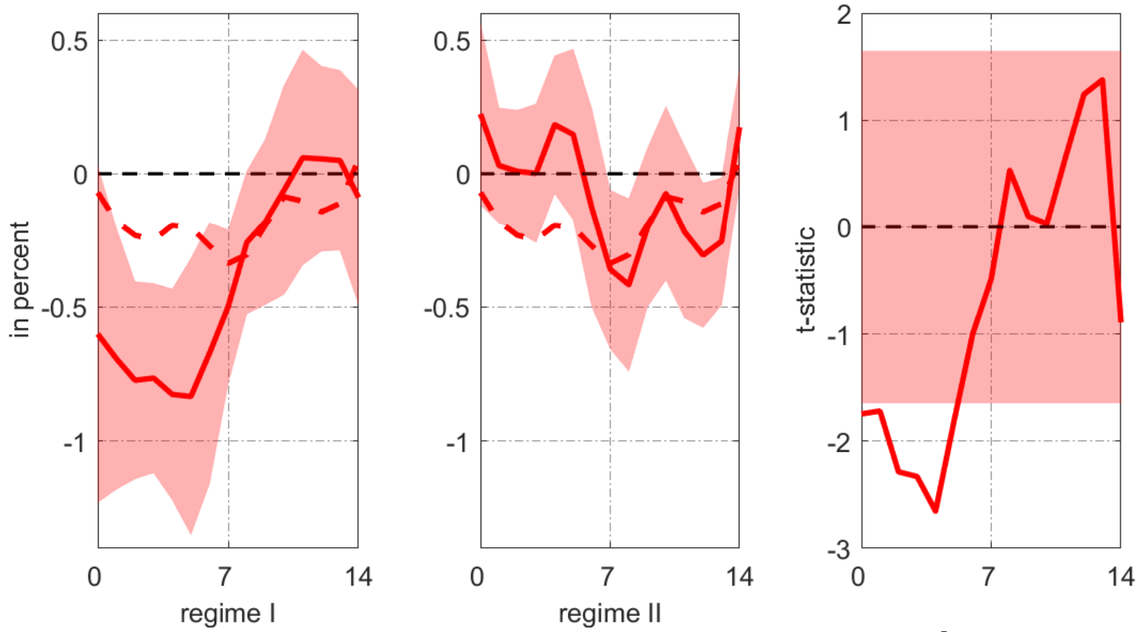

Notes: The first column shows the impulse response coefficients (red-solid) β_h^I for $h = 0, ..., H$ in regime I following a pandemic shock (one standard deviation); the second column shows the corresponding impulse response coefficients β_h^{II} in regime II. In both cases, the red-shaded area corresponds to the 90 percent confidence interval relying on Driscoll-Kraay standard errors. The red-dotted lines in the first two columns correspond to the impulse response coefficients from the linear model without allowing for state-dependent effects. The third column shows the t−statistics testing the null that $H_0 : \beta_h^I - \beta_h^{II} = 0$ for each horizon using the Driscoll-Kraay method. The red-shaded area covers the t−critical values for a 90% confidence interval, i.e. ±1.645.

January. Spending decreases even further on subsequent days and peaks at a decrease of 0.8% four to five days after the shock. Afterwards, total spending starts to steadily revert to its mean which is reached after ten days. In other words, having recognized the pandemic shock as bad news, households respond with a significant decline in aggregate spending when the reported daily number of new infections is relatively high.

However, we see a different pattern in regime II, i.e. when the growth rate of new infections is relatively small. Following a pandemic shock of the same size, spending remains unchanged for the first six days. After that, we find a drop by about 0.4% which is significant for two days. The response in regime II remains significant for all income groups. The disaggregated responses of spending, which we show below, offer an interpretation: the fall in regime II is most notable for groceries and healthcare spending. These are contact-intensive industries. For non-grocery retail and warehousing and transportation, in contrast, which can happen online, spending remains insignificant. The t−statistics in the right panel shows that the difference between the response in regime I and regime II is significantly different from zero for the first five days. This being said, we reject the null hypothesis of symmetric effects and find strong evidence for a regime-dependent response of spending to a pandemic shock.

As discussed in the introduction, the effect of the pandemic on household spending, whether it works through voluntary or forced consumption restraint or an unexpected fall in lifetime income, should increase in the spread of the pandemic. If few people are affected by the virus, the need to reduce spending, either voluntarily or through governmental restrictions, remains limited. Likewise, the drop in lifetime income remains small since a shock does not call entire industries or job profiles into question. If, in contrast, the number of infections is large, the shock should have stronger effects. Our results are consistent with this notion because the effect of the shock is significantly larger in regime I compared to regime II.

As we will see now, the impact of the shock across income quartiles is also consistent with this. Low-income households work more in contact-intensive jobs. With low-income household bearing the burden of social distancing, the future of these jobs is uncertain in a situation with many infections, while high-income households have jobs in which social distancing is possible. The impact of the shock in regime I should therefore be larger for low-income compared to high-income households.

Next, we take a closer look on the response of spending and investigate how the impulse responses differ across income quartiles. This is possible because we have data on spending by customers living in ZIP codes with different income levels, namely high (top quartile) median income, middle (middle two quartiles) median income as well as low (bottom quartile) median income. This encourages us to estimate the response of spending to our pandemic shock across different income classes.

Figure (5) reports the results for spending of all three different income categories. Starting with panel A, it stands out that, prompted by a pandemic shock, high-income households significantly reduce spending when the growth rate of new infections is relatively high (regime I). After four days, spending drops by about 0.8% and starts to revert to its mean which is reached after ten days. We also find a negative reaction in Regime II. However, we see a significant reaction much later than in regime I, namely after about one week. Overall, the reaction peaks at a drop of 0.5 percent, which is weaker than in regime I.

We have a similar picture in panels B and C. While the qualitative picture is comparable to the one of high-income households, it stands out that the response of households in regime I seems to be negatively correlated with lifetime income. That is, we see a larger drop in spending for middle-income households and an even larger drop for low-income households. In regime I, spending of low-income households drops by 1% percent, i.e. low-income households are more sensitive to the shock than high-income households.

For all income groups, the response to a pandemic shock depends strongly on where the economy is at the infections curve. Our results indicate that over the first week

Figure 5: Response of spending by income quartile

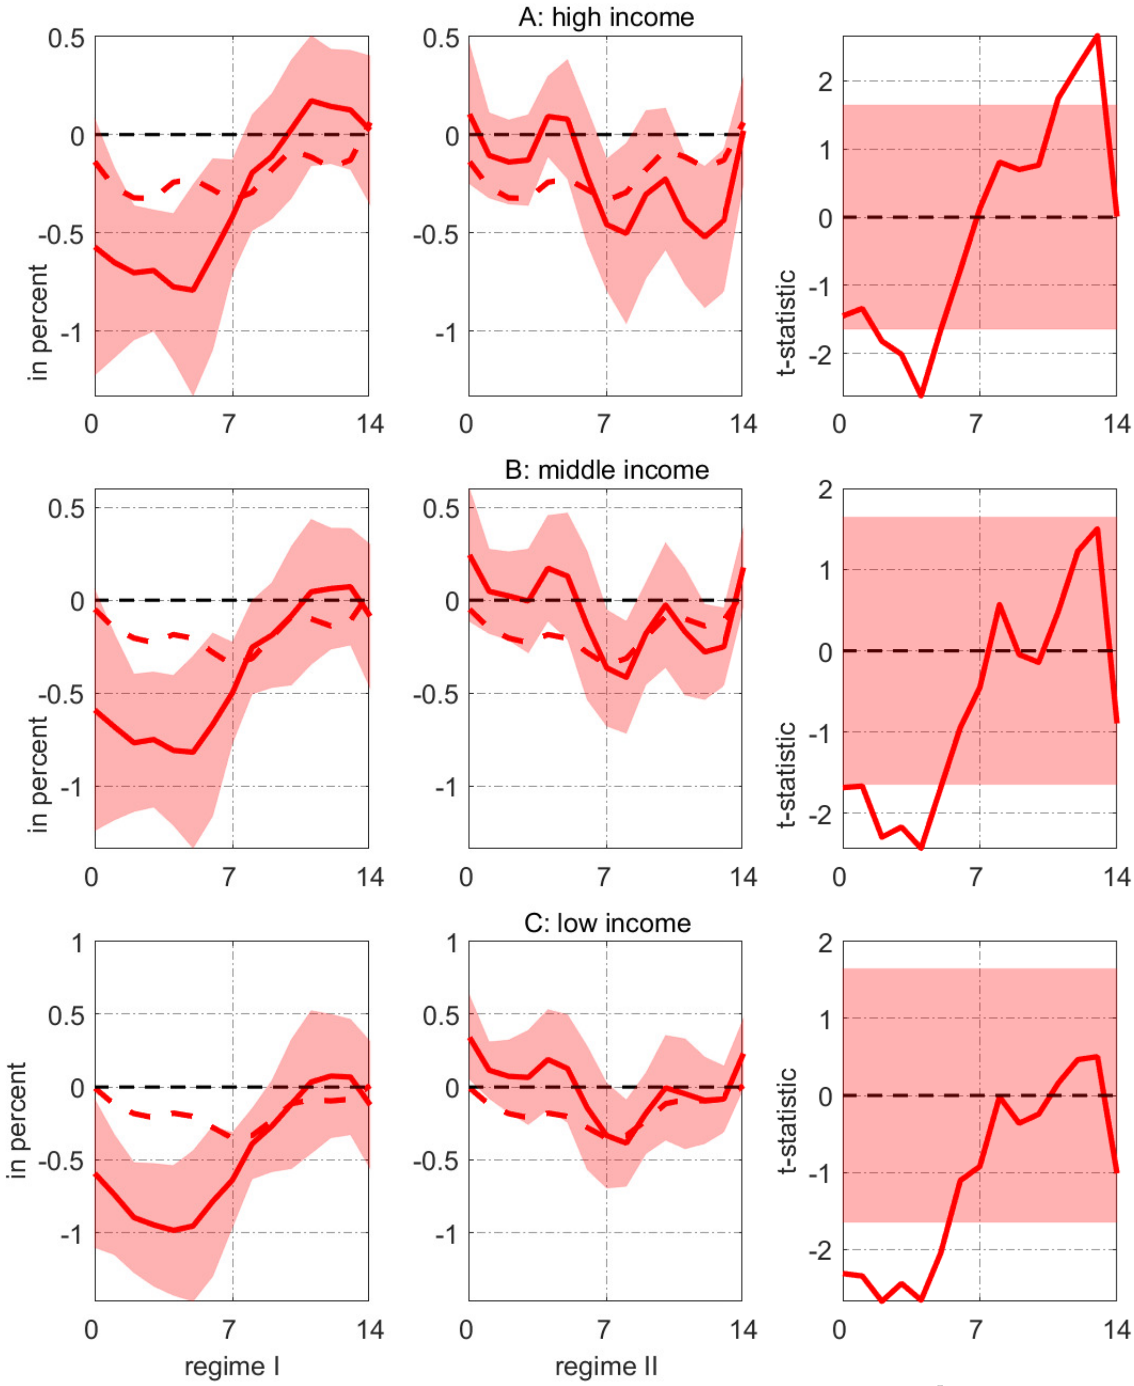

Notes: The first column shows the impulse response coefficients (red-solid) β_h^I for $h = 0, ..., H$ in regime I following a pandemic shock (one standard deviation); the second column shows the corresponding impulse response coefficients β_h^{II} in regime II. In both cases, the red-shaded area corresponds to the 90 percent confidence interval relying on Driscoll-Kraay standard errors. The red-dotted lines in the first two columns correspond to the impulse response coefficients from the linear model without allowing for state-dependent effects. The third column shows the t−statistics testing the null that $H_0 : \beta_h^I - \beta_h^{II} = 0$ for each horizon using the Driscoll-Kraay method. The red-shaded area covers the t−critical values for a 90% confidence interval, i.e. ±1.645.

after the shock the response is much stronger in regime I and is significantly different from the response in regime II. For each income quartile, we can reject the null that $\beta_h^I = \beta_h^{II}$ over the first few days considered. Hence, the response of household spending is asymmetric across regimes.

These findings can be rationalized based on the notion that the fall in lifetime income as a result of a pandemic shock is larger for low-income households. Workers in the service and hospitality sector, for example, face uncertainty about whether and when they can return to their old jobs. In addition, our results resemble what is found in the literature dealing with the nexus between household characteristics and the marginal propensity to consume (MPC). Aggregate MPC is typically found to depend on how aggregate shocks are distributed across households (see, for instance, Carroll et al., 2017; Carroll, 2009; Gelman, 2021; Calvet and Comon, 2003). In this context, higher marginal propensities to consume, as typically found in the literature, can explain why our pandemic shock has a larger impact on lower-income households. In the context of the COVID-19 pandemic, Karger and Rajan (2020) show an MPC of 0.68 for hand-to-mouth consumers and 0.23 for households with access to assets.[17]

Figure 6: Response of relative spending (high income - low income)

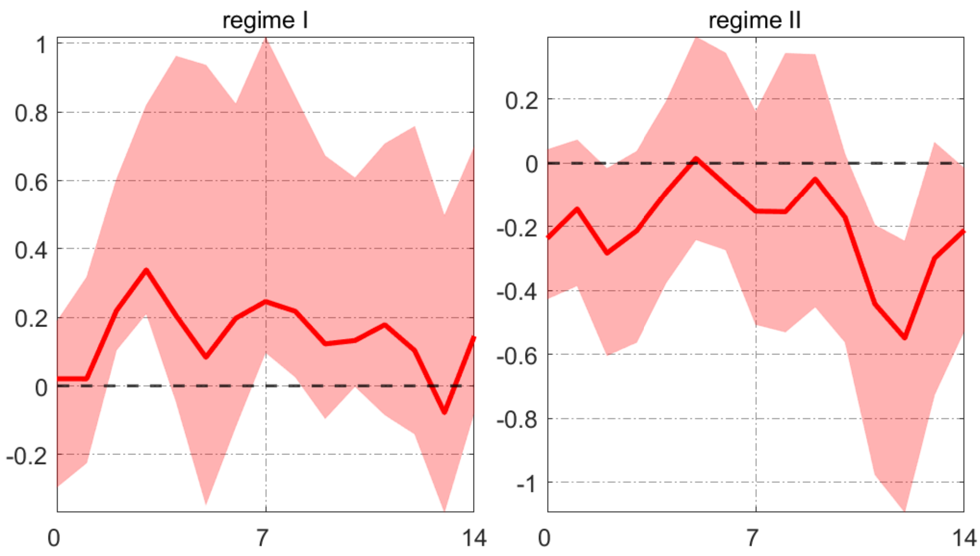

Notes: Difference of estimated coefficients $\beta_h^{high,I} - \beta_h^{low,I}$ in regime I and $\beta_h^{high,II} - \beta_h^{low,II}$ in regime II. The shaded areas cover the 5^{th} and 95^{th} percentiles from the distribution of the block bootstrap procedure as described in the text.

The previous graph revealed a significant state-dependence of the spending responses. However, we could not infer whether the response of high-income households is significantly different from low-income households. To shed light on the responses across quartiles, we proceed as follows: we generate 2,000 samples of contiguous blocks (with replacement) of four consecutive observations each. Within each

[17] Explanations include households' wealth or employment status and the accompanying heterogeneity with respect to liquidity constraints.

replication, for each $h = 0, ..., H$ we then estimate the impulse response coefficients and calculate the sign of $\beta_h^{high,j} - \beta_h^{low,j}$ in regime j =I, II.[18] We then use the distribution of our bootstrap and report the 5^{th} and 95^{th} percentiles.

We show the results in Figure (6). Starting with the left panel, i.e. regime I, we find a significantly positive difference, which peaks at about 0.3 after three days. To interpret this finding, recall that the actual response for both income quartiles was negative. Hence, the positive value means that, following a pandemic shock, low-income households reduce spending significantly more than high-income households. The results are consistent with the view that the pandemic prevents low-income households from returning to their jobs, while high-income households can reconcile their jobs with the necessary degree of social distancing. As a result, the drop in permanent income is larger for low-income households. In regime II, i.e. when the number of new infections is growing less strongly, we find that the drop in spending is stronger for high-income households. Overall, it stands out that a pandemic shock prompts an increase in consumption inequality.

B. *The Quantitative Significance of Pandemic Shocks*

So far our results imply that spending is significantly responsive to our identified pandemic shock. However, we do not yet know the overall quantitative significance of our shock. If the shock we have identified is indeed an important driver of consumer spending, this should also be reflected in the variance of the forecast errors. In this section, we therefore apply the strategy of Gorodnichenko and Lee (2020) for forecast error variance decompositions (FEVDs) within the local projection framework and assess the contribution of our pandemic shock to the variation of forecast errors at different horizons. In a first step, we estimate the same model as before, but this time we leave out the contemporaneous effect of the shock

$$y_{t+h} = F(z_{t-1})\left(\alpha_h^I + \gamma_h^I \mathbf{x}_t\right) + (1 - F(z_{t-1}))\left(\alpha_h^{II} + \gamma_h^{II}\mathbf{x}_t\right) + \delta t + \varphi_h \mathbf{D}_t + u_{t+h}. \tag{7}$$

In a second step, we take the estimated forecast errors $\widehat{u_{t+h}}$ and regress them on the shock ε_t occurring between t and $t + h$, while accounting for our regimes I and II from our baseline setting

$$\begin{aligned} \widehat{u_{t+h}} &= F(z_{t-1})\left(\boldsymbol{\omega}_0^I \varepsilon_t + \cdots + \boldsymbol{\omega}_h^I \varepsilon_{t+h}\right) + \cdots \\ \cdots &+ (1 - F(z_{t-1}))\left(\boldsymbol{\omega}_0^{II} \varepsilon_t + \cdots + \boldsymbol{\omega}_h^{II} \varepsilon_{t+h}\right) + \eta_{t+h}, \end{aligned} \tag{8}$$

[18] One difficulty in our application is that our set of control variables includes two dummies which have mostly 0-entries. It is therefore likely that inverting the matrix of right-hand side variables is not possible due to multicollinearity. To overcome this issue, we add another step and make sure that each bootstrap sample contains at least once those observations (not blocks) where the dummy variables are equal to one. From a practical point of view, this should not be a problem, since the dummy variables only improve the in-sample fit.

where ω_j^i for $j = 0, ..., h$ and $i = \text{I, II}$ measures the state-dependent effect of the pandemic shock on the estimated forecast error. Note that the coefficient of determination of this regression gives us the share of the forecast error variance which is explained by our pandemic shock. As shown by Gorodnichenko and Lee (2020), this method is a natural estimator of the population share of variance explained by the future innovations ε_t in the total variations of our endogenous variable.

Inference is based on the distribution of the R^2s from a block bootstrap procedure including a bias-correction step as recommended by Gorodnichenko and Lee (2020).[19] Remember that we do not have a VAR-based benchmark for our local projection-based

Figure 7: FEVD by income quartiles

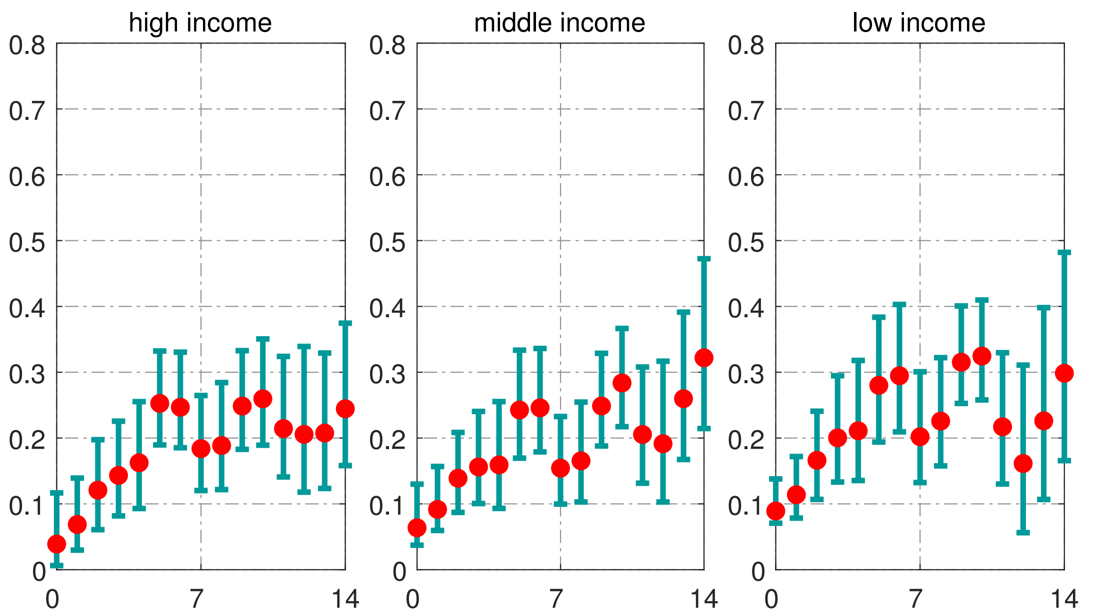

Notes: Explained share of forecast error variance after the bias-correction procedure (red dots) and the 5^{th} and 95^{th} percentiles of the distribution of the block bootstrap procedure

FEVD, which is due to the novelty of the data. However, theoretically, and based on our results so far, we expect pandemic shocks to be a major driver of fluctuations in household spending. This is what we see in Figure (7), which shows the estimated share of the forecast error variance that can be explained by our pandemic shock by income quartiles. The red dots correspond to the explained share of the forecast error variance. The green bars cover 90% of the distribution of the R^2s obtained by our bootstrap procedure. For all groups, our pandemic shock seems to be an important driver of spending. Over the first week, our shock explains about 20% of the forecast error variance. While we observe a drop in the explained share for all income quartiles

[19]To do so, we generate $B = 2000$ samples consisting of contiguous blocks of four consecutive observations each. Our bias is calculated as the difference between the mean over all bootstrap-based $R^{2,(b)}$ and the R^2 from our baseline procedure, i.e. $\text{bias}_h = B^{-1}\sum_{b=1}^{B} R^{2,(b)} - R^2$. Hence, our bias-corrected variance decomposition reads $R^{2,bc} = R^2 - \text{bias}_h$. As in the previous section, to improve the robustness of our estimates, we adjust our bootstrap algorithm and manually add the dummy observations equal to one to our bootstrap samples.

after the first week, the quantitative significance increases sharply and reaches its maximum after around 12 days. We see again, with an explained share of above 30% after nearly two weeks, that spending for low-income households is most responsive to our pandemic shock. Interestingly, the 5^{th} percentiles are above zero across all income quartiles and for all horizons considered. Our results therefore point to an important role of the pandemic shock in the variation of consumer spending.

Thus, we conclude that our pandemic shock is a significant driver for all income groups, but especially for low-income households.

V Results on the State Level

Our data on spending is also available on the level of US states. We therefore repeat our exercise from the previous subsection and now investigate the responsiveness of spending for the ten states with the largest population. The driving variable remains the nation-wide pandemic shock and the state-variable is still the nation-wide growth of new infections.[20]

Figure 8: Response on the state level

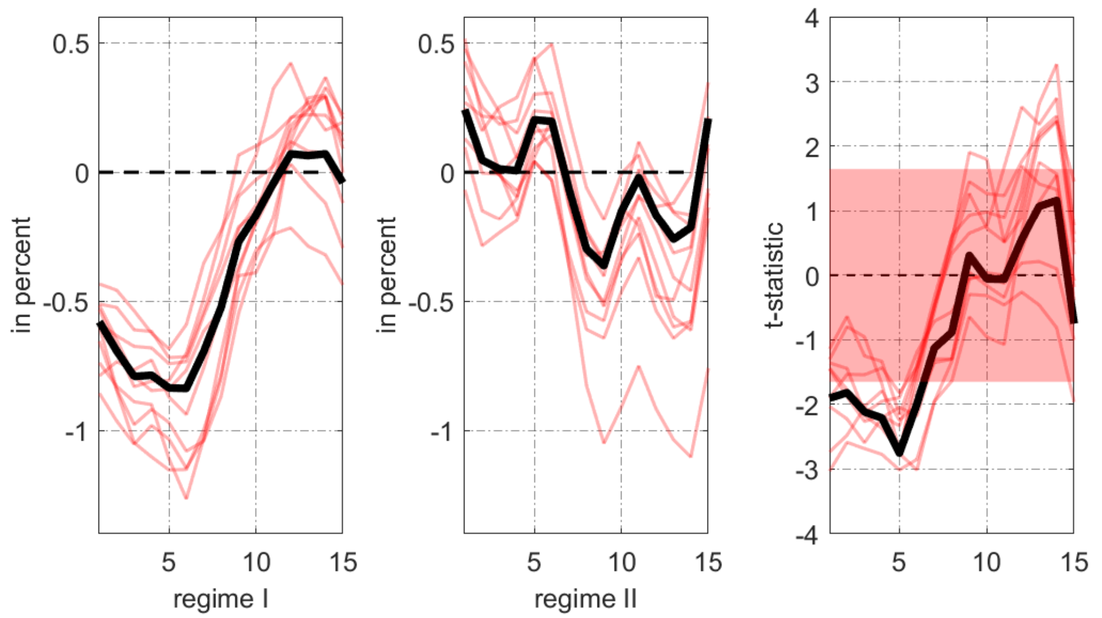

Notes: The first column shows the impulse response coefficients β_h^I for $h = 0, ..., H$ in regime I following a pandemic shock (one standard deviation), the second column shows the corresponding impulse response coefficients β_h^{II} in regime II. The third column shows the t−statistics testing the null that $H_0 : \beta_h^I - \beta_h^{II} = 0$ for each horizon using the Driscoll-Kraay method. The red-shaded area covers the t−critical values for a 90% confidence interval, i.e. ±1.645.

[20]It would be very interesting to estimate the regression of the baseline model completely state-dependent. The problem, however, is that both the reported death figures and the reported infection figures are very patchy. A shock series as in the baseline model at state level is thus unfortunately not feasible. For this reason, we have to use the nationwide shock and the nationwide state variable.

The red lines in Figure (8) show the mean impulse responses of total spending following a pandemic shock for the ten states, whereas the black line shows the mean responses on the national level. It stands out that the qualitative pattern in regime I appears to be very homogeneous across states. In regime I, we observe a sharp drop in spending in all states. Spending peaks after four or five days before it returns to its mean after two weeks. Also in regime II, the overall direction of the responses looks similar across all states. The third column shows that, for the first five days, in many cases we reject the null of equal responses across regimes. That is, also on the state level we find a stronger response of consumer spending in regime I.[21]

VI Alternative State-Variables

In our baseline setting, we choose the growth rate of daily infections as our state variable. While figures about new infections are omnipresent in the media, the drawback of this state variable is that it does not provide information on the overall number of infections. In fact, households may condition their spending behavior on the level of the infections curve, rather than the slope of the curve.

As a first alternative, we therefore repeat our estimation and use a binary indicator as our state variable. The indicator variable is 0 if the temporary peak of new infections is not yet reached and 1 if the total number of new infections decreases (alternative I). Holding anything else constant, we replace $F(z_{t-1})$ with an indicator variable $I(z_{t-1})$, where $I(z_{t-1})$ is equal to zero from April 10 to June 7 as well as from July 25 to September 6, and equal to one otherwise. These are roughly the cut-off dates that reflect a reversal of the current infection pattern.

As a second alternative, we specify $F(z_t)$ as a logistic function of the form

$$F(z_t) = \frac{\exp\left(\kappa \frac{z_t - \mu}{\sigma_z}\right)}{1 + \exp\left(\kappa \frac{z_t - \mu}{\sigma_z}\right)}, \tag{9}$$

where μ is used to control the proportion of the sample the economy spends in either state, and σ_z is the sample standard deviation of the state variable z_t. The parameter κ controls how abruptly the economy switches from one state to the other following movements of the state variable. In other words, higher values of κ mean that small movements of the state variable suffice to induce a switch from one regime to the other. However, although the parametric approach has the disadvantage that we have to make explicit assumptions about the parameters determining the behavior

[21] While there are no error bands shown in Figure (8), Figure (10) in the appendix shows the corresponding impulse response coefficients across states with ±1.645 standard deviations for selected periods, namely four, eight and twelve periods after the pandemic shock. For reasons of comparison, the transparent horizontal lines report the coefficients on the national level. Overall, it stands out that the reaction to our shock is very homogeneous across the states and Spending in most states reacts similarly to the national level.

of switching from one state to the other, this approach is well understood and relies on the idea of Granger and Terasvirta (1993) and is, among others, used in Auerbach and Gorodnichenko (2012), Ramey and Zubairy (2018) and Tenreyro and Thwaites (2016). We set $\kappa = 3$ which implies an intermediate intensity of regime-switching and set $\mu = \text{med}(z_t)$. Figure (9) shows both alternative regimes. The left panel shows

Figure 9: Alternative states

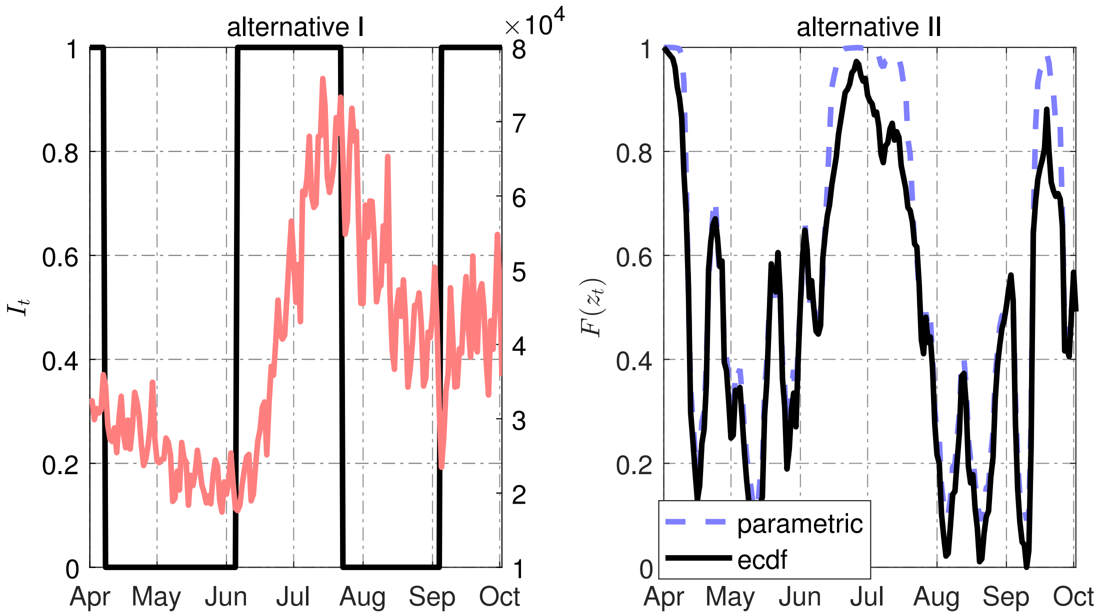

Notes: The left panel shows daily new infections (red-solid) and the corresponding alternative regimes as described in the main text. The right panel shows the alternative transition function obtained by the parametric approach as described in the text.

the actual number of daily new infections and the distinction of regimes I and II as indicated by the black vertical lines on the cut-off dates. The right panel shows the transition function based on the parametric approach (blue-dashed). The alternative transition function looks very much like our baseline transition function, although we observe a higher weight of regime I in late June and at the end of our sample.

Figures (11) to (14) in the appendix show the corresponding impulse response functions for total spending and spending across income quartiles for both alternative regime classifications. It stands out that our results from the first alternative state look exactly like the results in the previous section. That is, total spending significantly decreases in regime I, i.e. when the overall level of newly-infected people is relatively high. The peak is reached after 3-4 days with a drop of 0.6 percent. Moreover, the responses in regime II resemble those from our baseline results. Finally, we also observe a significant difference in the responses across regimes which is consistent with asymmetric effects. The responses across different income quartiles are similar to those in the benchmark model.

The results from our alternative state variable (alternative II) exhibit a similar picture. Both the shape of the impulse responses and the magnitudes of the effects remain almost unchanged. This being said, our results indicate that spending reacts more strongly when the number of new infections is high. Again, we find that the difference across regimes I and II is stronger for low-income households.

VII Transmission Mechanism

Having established heterogeneous effects on household spending across income groups and across regimes, we now want to shed light on spending subcategories. We use additional data to sketch potential transmission channels.

A. *Spending Subcategories*

The Chetty et al. (2020) data set allows us to dis-aggregate spending into subcategories. Figures (15) to (17) in the appendix show the state-dependent responses across the available segments of spending. We find that the stronger response in the first state is driven by retail spending as well as spending on health care and warehousing and transportation. Spending on groceries, in contrast, falls equally in both states, while spending on merchandise and apparel seams insensitive to the pandemic shock.

B. *Driving Forces of Consumer Spending*

As discussed before, the drop in consumer spending could reflect a host of different driving forces. One key risk households face is the loss of the job, which should affect their spending decisions. Figure (18) in the appendix depicts the response of employment across income groups to the pandemic shock. Here, we do not allow for state-dependent effects in order to isolate the effect of income. The data is again taken from Chetty et al. (2020). The results are compelling: the pandemic shock causes a significant drop in employment in low-income ZIP codes, has no significant effect in mid-income areas and a small but significant positive effect in high-income ZIP codes. Hence, the income-gap in unemployment risk widens, which is consistent with a stronger drop in spending of low-income households.

To conclude this section, we draw on alternative survey data from the Federal Reserve Bank of Cleveland.[22] The daily survey asks households, among other things, for the expected duration of the coronavirus outbreak. Differences in the expected duration could be one determinant of the significant differences in the responses across state I and state II, respectively. The data does not, however, offer a distinction across income quartiles. Figure (19) in the appendix reports the impact on the percentage of respondents expecting the coronavirus outbreak will last six months, one year or two years. In state I, the pandemic shock leads households to revise their expectations about the duration of the virus: fewer households expect the virus to disappear within six months and more households believe the virus will be more persistent. This pattern is

[22] See https://www.clevelandfed.org/our-research/indicators-and-data/consumers-and-covid-19.aspx.

consistent with the spending responses shown before. For a higher expected duration of the pandemic, households should become more reluctant to spend.

VIII Conclusion

We provided evidence on the causal effect of unexpected news about the COVID-19 pandemic on spending of US households. Our first finding is that a pandemic shock, the forecast error about the number of fatalities, has a negative effect on spending: a surprise increase in the number of deaths leads to a sharp reduction in expenditures. We also showed that this effect is depending on the position of the US economy with respect to the infection curve. With the number of new infections increasing, the effect of a shock is much stronger. If the growth rate of the number of infections is small, in contrast, the pandemic shock has almost no effect. A second finding pertains to the effect across income quartiles. If the number of infections is increasing strongly, the shock prompts a much larger adjustment of spending from low-income households compared to high-income households. Hence, the pandemic shock increases consumption inequality.

Our results have two implications for economic policies designed to stabilize aggregate economic activity. First, policy measures should target low-income households more than high-income households. Spending of low-income households is particularly sensitive to a pandemic shock, such that support packages will be more effective when targeting households with relatively low income.

Second, economic support through direct as well as indirect transfers should be conditioned on the state of the pandemic in order to stabilize consumption effectively. Transfers will be more effective when the number of infections is large, because in this state households would reduce spending the most. Due to a lag in the implementation of economic support, though, timing policy measures properly is difficult. One way to address the time lag is to commit to disburse financial support automatically once the number of infections crosses a specific threshold.

References

Alexander, Diane and Ezra Karger, "Do Stay-at-Home Orders Cause People to Stay at Home? Effects of Stay-at-Home Orders on Consumer Behavior," *The Review of Economics and Statistics*, 2020, *1–25.*

Andersen, Asger L, Emil T Hansen, Niels Johannesen, and Adam Sheridan, "Consumer Responses to the COVID-19 Crisis: Evidence From Bank Account Transaction Data," 2020. unpublished, University of Copenhagen.

Auerbach, Alan J and Yuriy Gorodnichenko, "Fiscal Multipliers in Recession and Expansion," in "Fiscal Policy After the Financial Crisis," University of Chicago Press, 2012, pp. 63–98.

Baker, Scott R, Robert A Farrokhnia, Steffen Meyer, Michaela Pagel, and Constantine Yannelis, "How Does Household Spending Respond to an Epidemic? Consumption During the 2020 COVID-19 Pandemic," *Review of Asset Pricing Studies*, 2020, *10* (4), 834–862.

Binder, Carola, "Coronavirus Fears and Macroeconomic Expectations," *The Review of Economics and Statistics*, 2020, *102* (4), 721–730.

Born, Benjamin, Gernot J Müller, and Johannes Pfeifer, "Does Austerity Pay Off?," *Review of Economics and Statistics*, 2020, *102* (2), 323–338.

Calvet, Laurent and Etienne Comon, "Behavioral Heterogeneity and the Income Effect," *Review of Economics and Statistics*, 2003, *85* (3), 653–669.

Carroll, Christopher D, "Precautionary Saving and the Marginal Propensity to Consume out of Permanent Income," *Journal of Monetary Economics*, 2009, *56* (6), 780–790.

__ , Jiri Slacalek, Kiichi Tokuoka, and Matthew N White, "The Distribution of Wealth and the Marginal Propensity to Consume," *Quantitative Economics*, 2017, *8* (3), 977–1020.

Carvalho, Vasco M, Juan R Garcia, Stephen Hansen, Álvaro Ortiz, Tomasa Rodrigo, José V Rodríguez Mora, and Pep Ruiz, "Tracking the COVID-19 Crisis with High-Resolution Transaction Data," *Royal Society Open Science*, 2021, *8* (8), 210218.

Chetty, Raj, John N Friedman, Nathaniel Hendren, Michael Stepner, and The Opportunity Insights Team, "How Did COVID-19 and Stabilization Policies Affect Spending and Employment? A New Real-Time Economic Tracker Based on Private Sector Data," NBER Working Paper 27431, National Bureau of Economic Research 2020.

Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber, "The Cost of the Covid-19 Crisis: Lockdowns, Macroeconomic Expectations, and Consumer Spending," NBER Working Paper 27141, National Bureau of Economic Research 2020.

Cox, Natalie, Peter Ganong, Pascal Noel, Joseph Vavra, Arlene Wong, Diana Farrell, Fiona Greig, and Erica Deadman, "Initial Impacts of the Pandemic on Consumer Behavior: Evidence from Linked Income, Spending, and Savings Data," *Brookings Papers on Economic Activity*, 2020, *2020* (2), 35–82.

Dietrich, Alexander M, Keith Kuester, Gernot J Müller, and Raphael Schoenle, "News and Uncertainty about COVID-19: Survey Evidence and Short-Run Economic Impact," *Journal of Monetary Economics*, 2022, pp. 35–51.

Driscoll, John C and Aart C Kraay, "Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data," *Review of Economics and Statistics*, 1998, *80* (4), 549–560.

Eichenbaum, Martin S, Sergio Rebelo, and Mathias Trabandt, "The Macroeconomics of Epidemics," *The Review of Financial Studies*, 2021, *34* (11), 5149–5187.

Fetzer, Thiemo, Lukas Hensel, Johannes Hermle, and Christopher Roth, "Coronavirus Perceptions and Economic Anxiety," *Review of Economics and Statistics*, 2020, pp. 1–36.

Gelman, Michael, "What Drives Heterogeneity in the Marginal Propensity to Consume? Temporary Shocks vs Persistent Characteristics," *Journal of Monetary Economics*, 2021, *117*, 521–542.

Goolsbee, Austan and Chad Syverson, "Fear, Lockdown, and Diversion: Comparing Drivers of Pandemic Economic Decline," *Journal of Public Economics*, 2021, *193*, 104311.

Gorodnichenko, Yuriy and Byoungchan Lee, "Forecast Error Variance Decompositions with Local Projections," *Journal of Business & Economic Statistics*, 2020, *38* (4), 921–933.

Granger, Clive W J and Timo Terasvirta, "Modelling Non-Linear Economic Relationships," *Oxford: Oxford University Press*, 1993.

Gu, Youyang, "COVID-19 Projections Using Machine Learning," *Retrieved July*, 2020, *13*, 2020.

Hoke, Sinem Hacioglu, Diego R Känzig, and Paolo Surico, "Consumption in the Time of COVID-19: Evidence From UK Transaction Data," CEPR Working Paper 14733, Centre for Economic Policy Research 2020.

Jappelli, Tullio and Luigi Pistaferri, "The Consumption Response to Income Changes," *Annual Review of Economics*, 2010, *2* (1), 479–506.

Jordà, Òscar, "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 2005, *95* (1), 161–182.

Karger, Ezra and Aastha Rajan, "Heterogeneity in the Marginal Propensity to Consume: Evidence from Covid-19 Stimulus Payments," Working Paper 2020-15, Federal Reserve Bank of Chicago 2020.

Mongey, Simon, Laura Pilossoph, and Alexander Weinberg, "Which Workers Bear the Burden of Social Distancing?," *The Journal of Economic Inequality*, 2021, *19* (3), 509–526.

Ramey, Valerie A and Sarah Zubairy, "Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data," *Journal of Political Economy*, 2018, *126* (2), 850–901.

Tenreyro, Silvana and Gregory Thwaites, "Pushing on a String: US Monetary Policy is Less Powerful in Recessions," *American Economic Journal: Macroeconomics*, 2016, *8* (4), 43–74.

Figure 10: Response of total spending on the state level

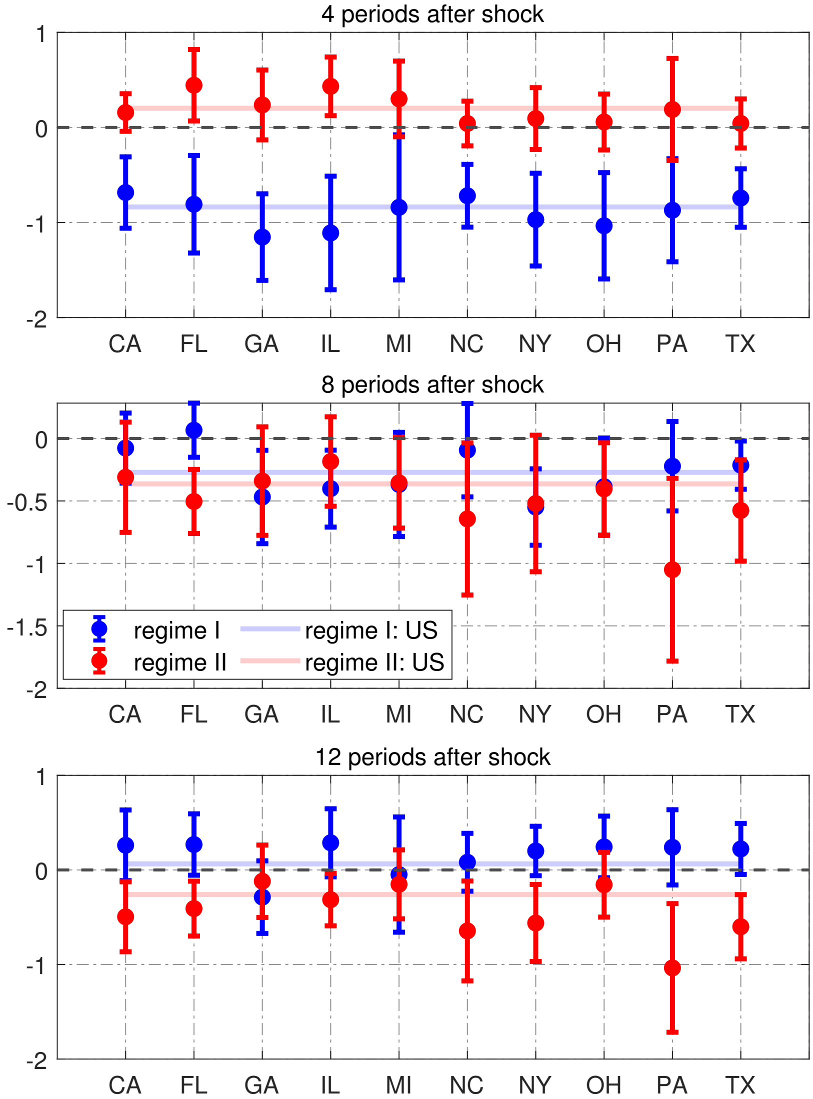

Notes: The dots correspond to the point estimates for regime I (red) and regime II (blue) after 4 periods (upper panel), 8 periods (middle panel) and 12 periods (bottom panel). The edges indicate 1.645 standard deviations in order to cover a 90% confidence interval, based on Driscoll-Kraay standard errors. The horizontal lines reflect the nation-wide effects in each regime.

Figure 11: Response of total spending: alternative I

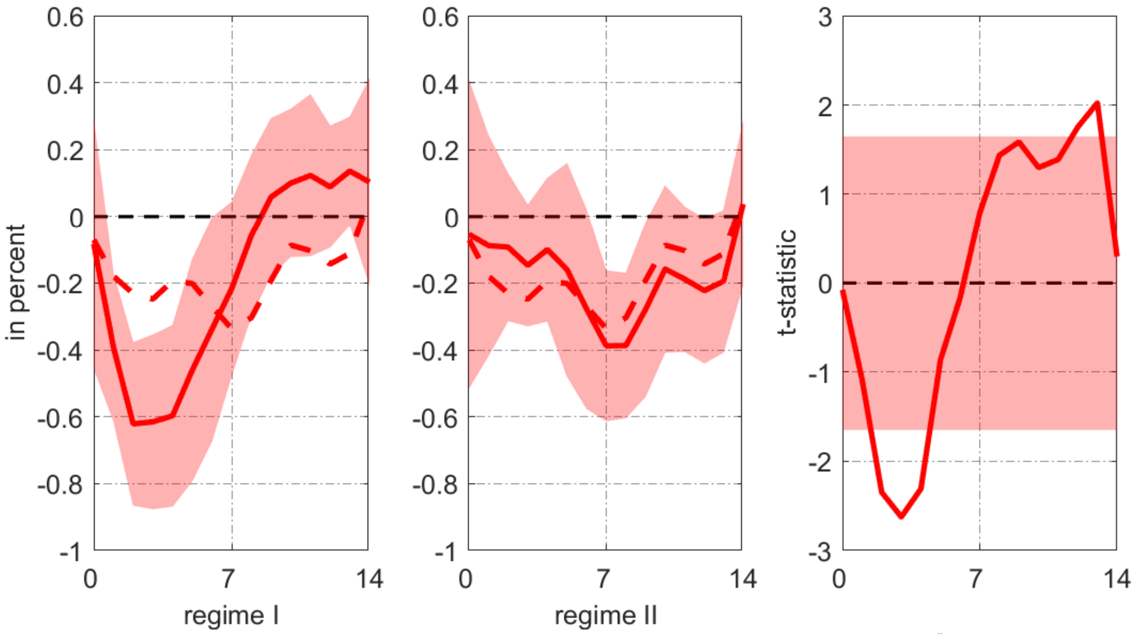

Notes: The first column shows the impulse response coefficients (red-solid) β_h^I for $h = 0, ..., H$ in regime I following a pandemic shock (one standard deviation); the second column shows the corresponding impulse response coefficients β_h^{II} in regime II. In both cases, the red-shaded area corresponds to the 90 percent confidence interval relying on Driscoll-Kraay standard errors. The red-dotted lines in the first two columns correspond to the impulse response coefficients from the linear model without allowing for state-dependent effects. The third column shows the t−statistics testing the null that $H_0 : \beta_h^I - \beta_h^{II} = 0$ for each horizon using the Driscoll-Kraay method. The red-shaded area covers the t−critical values for a 90% confidence interval, i.e. ±1.645.

Figure 12: Response of spending across income quartiles: alternative I

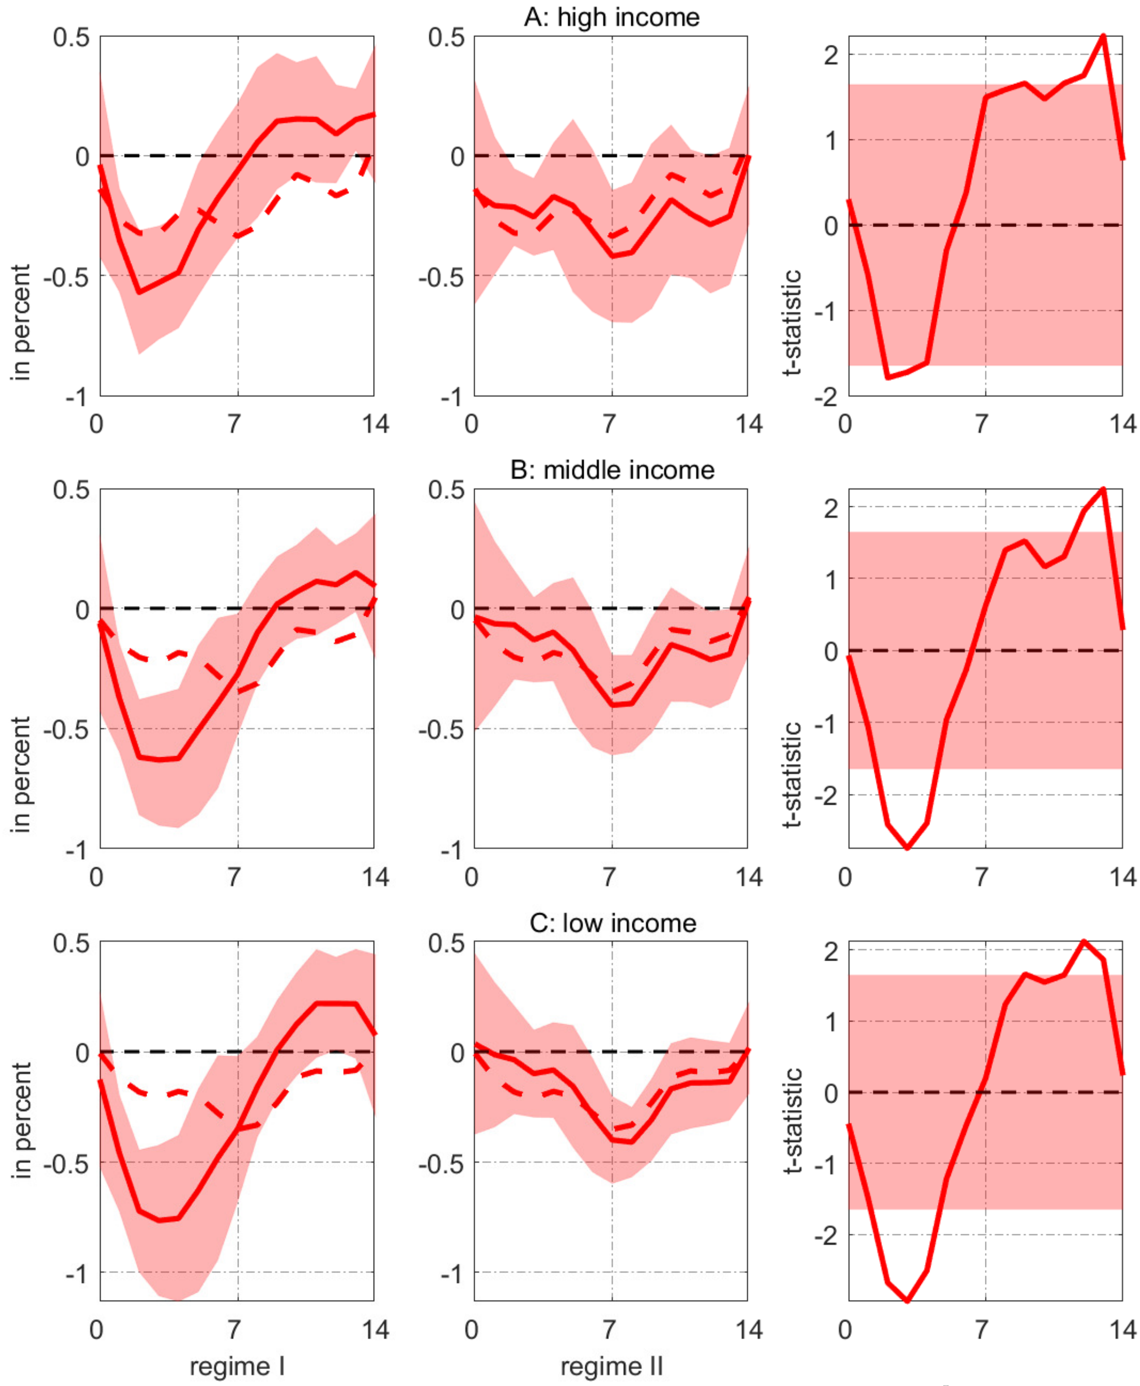

Notes: The first column shows the impulse response coefficients (red-solid) β_h^I for $h = 0, ..., H$ in regime I following a pandemic shock (one standard deviation); the second column shows the corresponding impulse response coefficients β_h^{II} in regime II. In both cases, the red-shaded area corresponds to the 90 percent confidence interval relying on Driscoll-Kraay standard errors. The red-dotted lines in the first two columns correspond to the impulse response coefficients from the linear model without allowing for state-dependent effects. The third column shows the t−statistics testing the null that $H_0 : \beta_h^I - \beta_h^{II} = 0$ for each horizon using the Driscoll-Kraay method. The red-shaded area covers the t−critical values for a 90% confidence interval, i.e. ± 1.645.

Figure 13: Response of total spending: alternative II

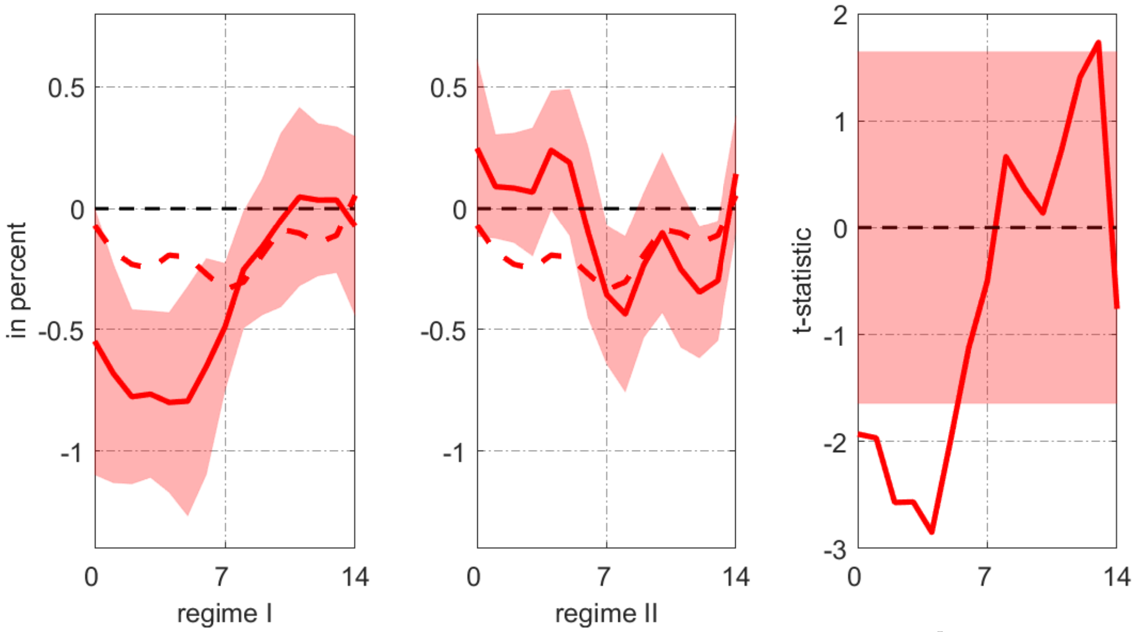

Notes: The first column shows the impulse response coefficients (red-solid) β_h^I for $h = 0, ..., H$ in regime I following a pandemic shock (one standard deviation); the second column shows the corresponding impulse response coefficients β_h^{II} in regime II. In both cases, the red-shaded area corresponds to the 90 percent confidence interval relying on Driscoll-Kraay standard errors. The red-dotted lines in the first two columns correspond to the impulse response coefficients from the linear model without allowing for state-dependent effects. The third column shows the $t-$statistics testing the null that $H_0 : \beta_h^I - \beta_h^{II} = 0$ for each horizon using the Driscoll-Kraay method. The red-shaded area covers the $t-$critical values for a 90% confidence interval, i.e. ± 1.645.

Figure 14: Response of spending across income quartiles: alternative II

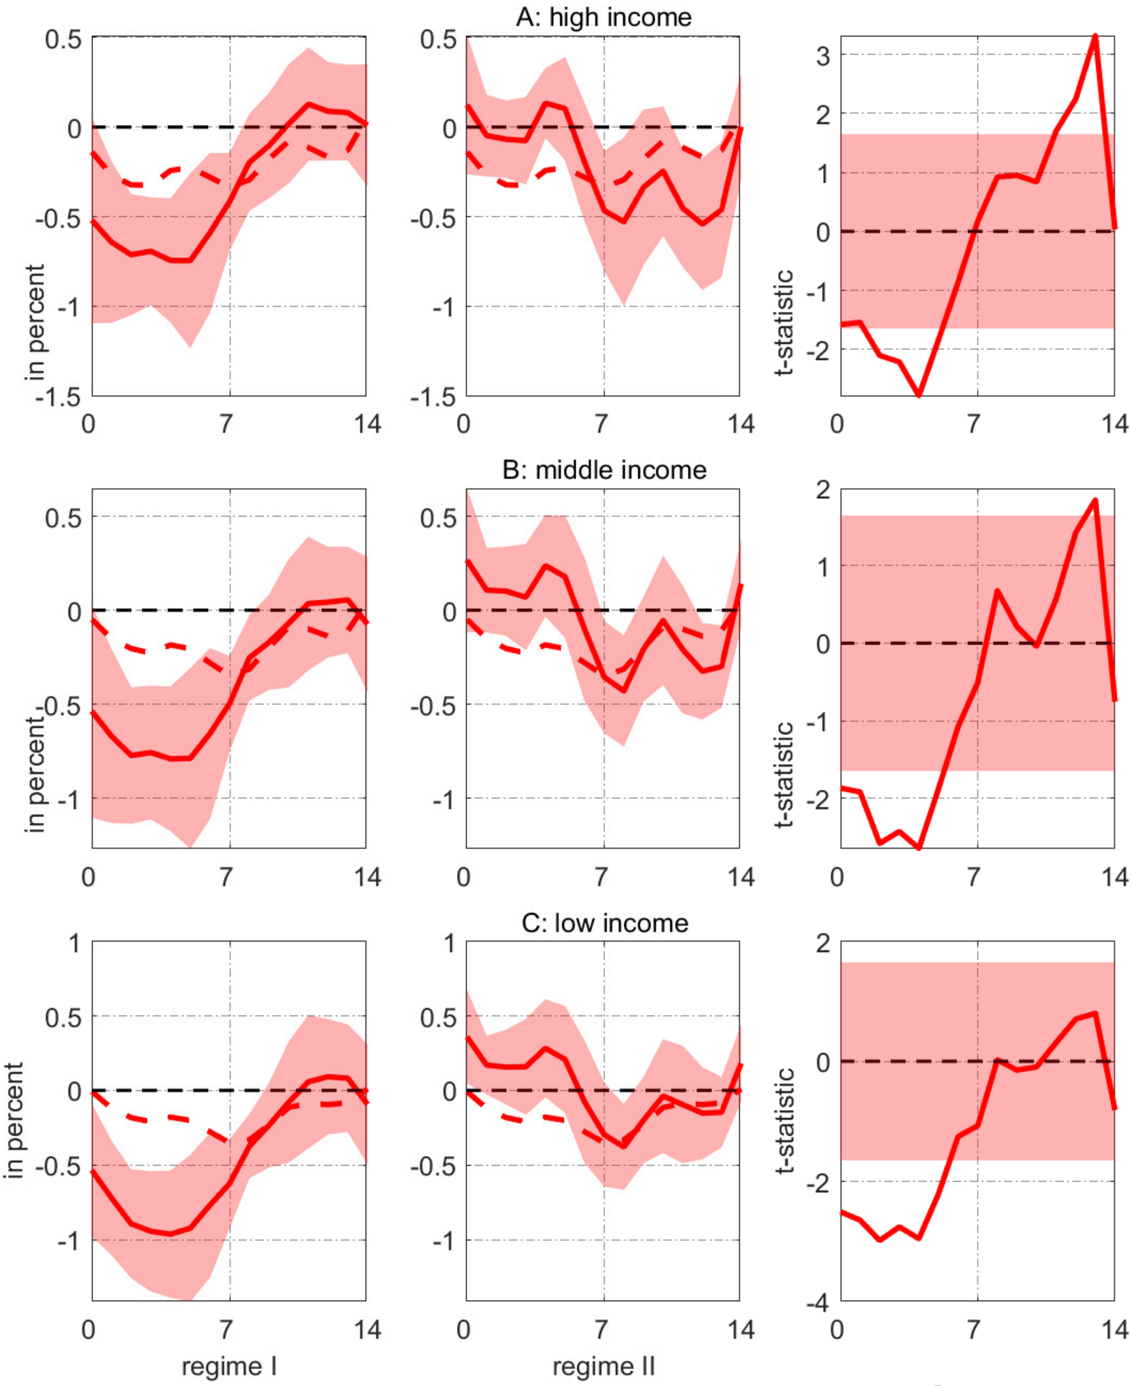

Notes: The first column shows the impulse response coefficients (red-solid) β_h^I for $h = 0, ..., H$ in regime I following a pandemic shock (one standard deviation); the second column shows the corresponding impulse response coefficients β_h^{II} in regime II. In both cases, the red-shaded area corresponds to the 90 percent confidence interval relying on Driscoll-Kraay standard errors. The red-dotted lines in the first two columns correspond to the impulse response coefficients from the linear model without allowing for state-dependent effects. The third column shows the t−statistics testing the null that $H_0 : \beta_h^I - \beta_h^{II} = 0$ for each horizon using the Driscoll-Kraay method. The red-shaded area covers the t−critical values for a 90% confidence interval, i.e. ±1.645.

Figure 15: Response of spending: retail

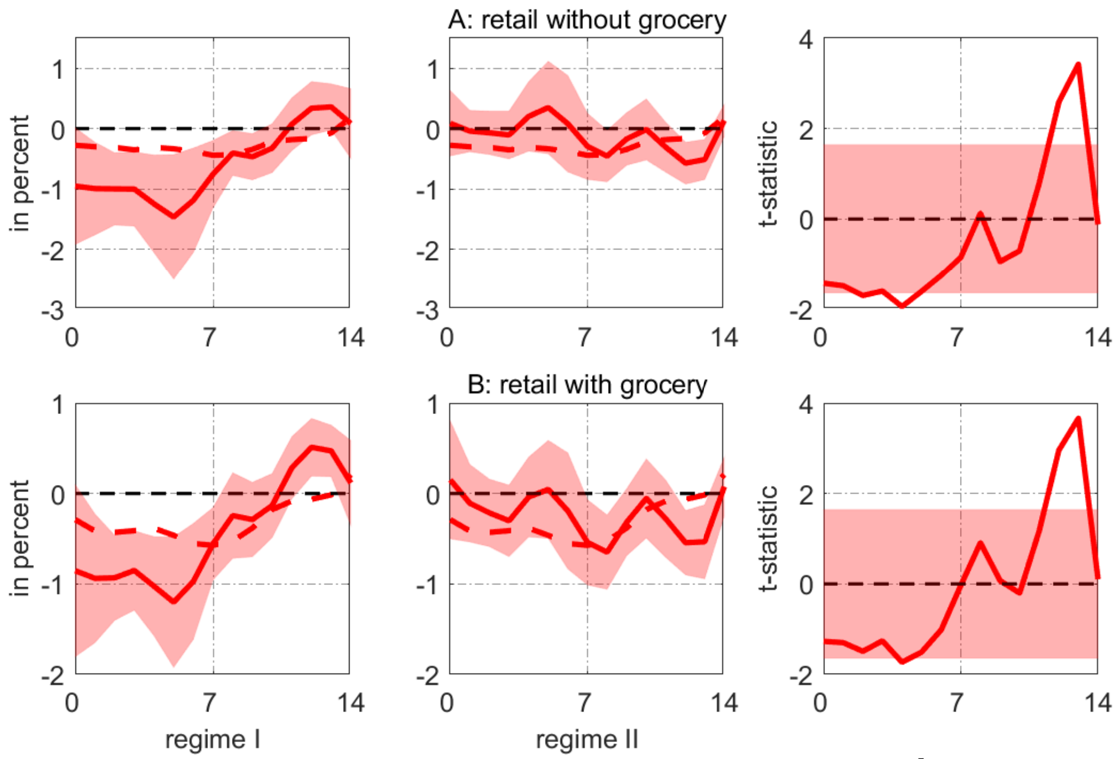

Notes: The first column shows the impulse response coefficients (red-solid) β_h^I for $h = 0, ..., H$ in regime I following a pandemic shock (one standard deviation); the second column shows the corresponding impulse response coefficients β_h^{II} in regime II. In both cases, the red-shaded area corresponds to the 90 percent confidence interval relying on Driscoll-Kraay standard errors. The red-dotted lines in the first two columns correspond to the impulse response coefficients from the linear model without allowing for state-dependent effects. The third column shows the $t-$statistics testing the null that $H_0 : \beta_h^I - \beta_h^{II} = 0$ for each horizon using the Driscoll-Kraay method. The red-shaded area covers the $t-$critical values for a 90% confidence interval, i.e. ± 1.645.

Figure 16: Response of spending: subcategories I

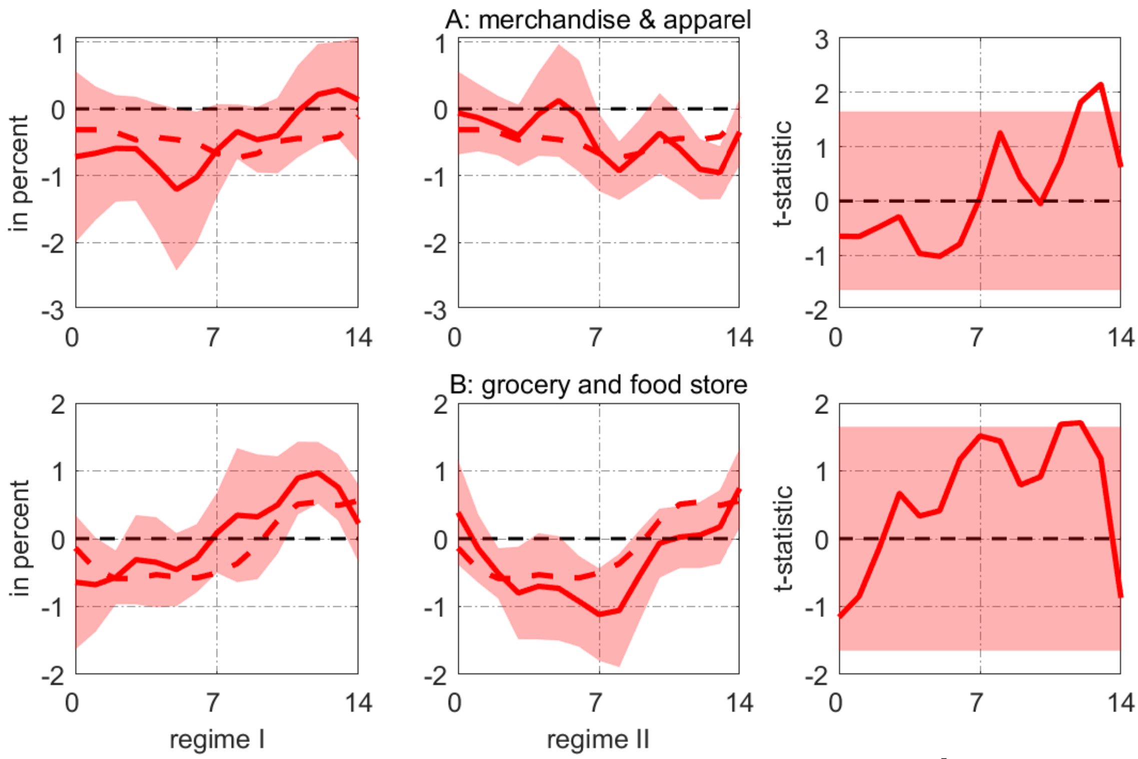

Notes: The first column shows the impulse response coefficients (red-solid) β_h^I for $h = 0, ..., H$ in regime I following a pandemic shock (one standard deviation); the second column shows the corresponding impulse response coefficients β_h^{II} in regime II. In both cases, the red-shaded area corresponds to the 90 percent confidence interval relying on Driscoll-Kraay standard errors. The red-dotted lines in the first two columns correspond to the impulse response coefficients from the linear model without allowing for state-dependent effects. The third column shows the t–statistics testing the null that $H_0 : \beta_h^I - \beta_h^{II} = 0$ for each horizon using the Driscoll-Kraay method. The red-shaded area covers the t–critical values for a 90% confidence interval, i.e. ± 1.645.

Figure 17: Response of spending: subcategories II

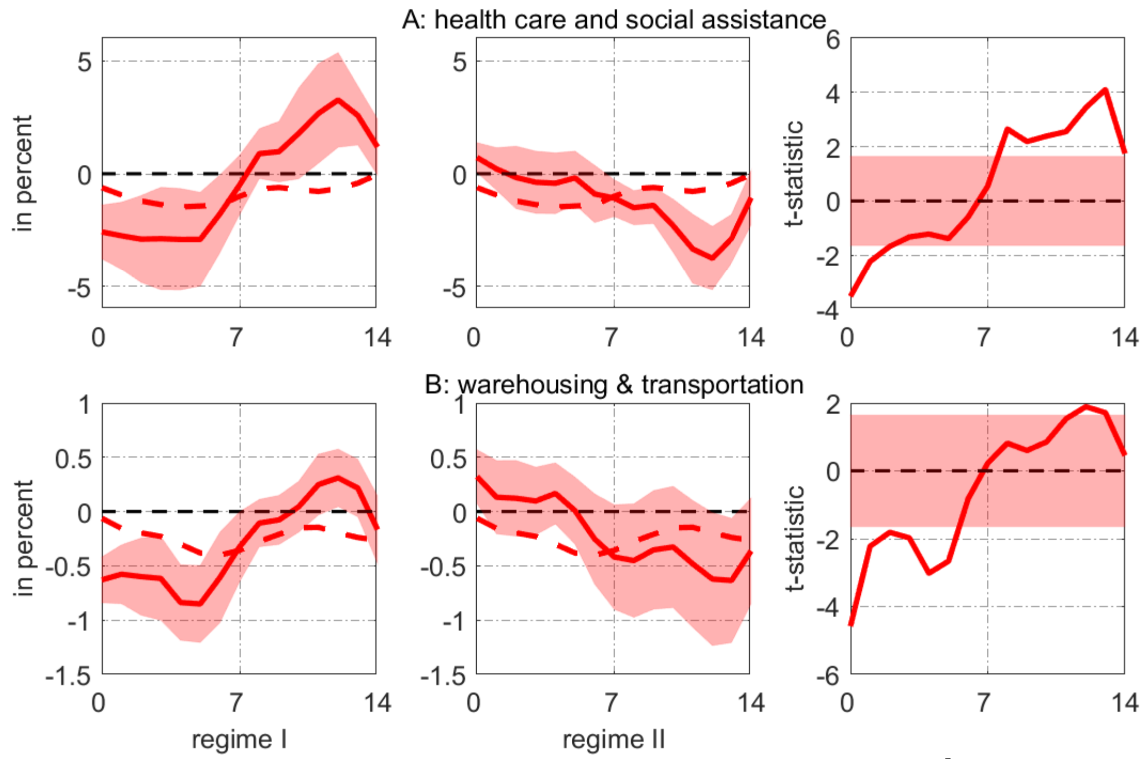

Notes: The first column shows the impulse response coefficients (red-solid) β_h^I for $h = 0, ..., H$ in regime I following a pandemic shock (one standard deviation); the second column shows the corresponding impulse response coefficients β_h^{II} in regime II. In both cases, the red-shaded area corresponds to the 90 percent confidence interval relying on Driscoll-Kraay standard errors. The red-dotted lines in the first two columns correspond to the impulse response coefficients from the linear model without allowing for state-dependent effects. The third column shows the $t-$statistics testing the null that $H_0 : \beta_h^I - \beta_h^{II} = 0$ for each horizon using the Driscoll-Kraay method. The red-shaded area covers the $t-$critical values for a 90% confidence interval, i.e. ± 1.645.

Figure 18: Response of employment by income quartile

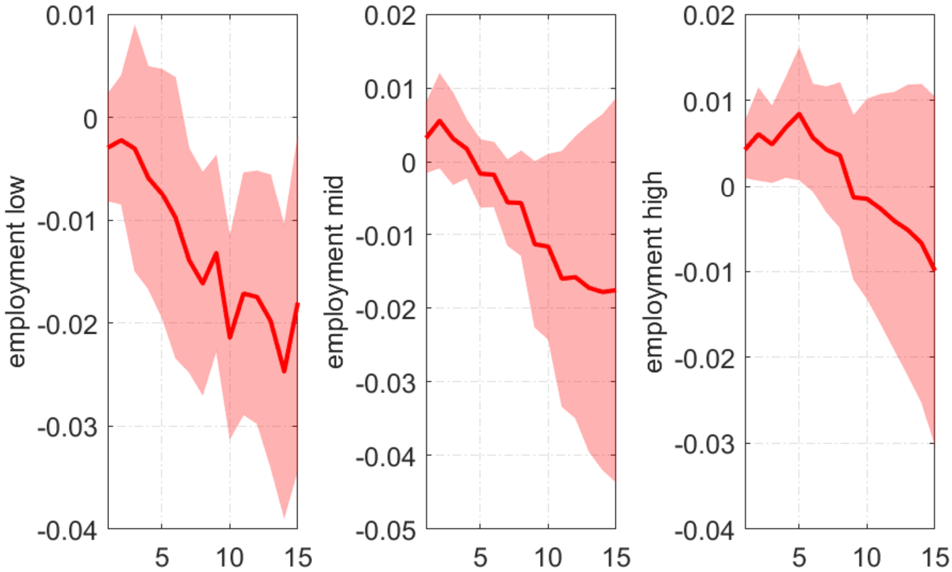

Notes: The graph shows the impulse response coefficients (red-solid) following a pandemic shock (one standard deviation) across income quartiles. In all cases, the red-shaded area corresponds to the 90 percent confidence interval relying on Driscoll-Kraay standard errors. The red-shaded area covers the $t-$critical values for a 90% confidence interval, i.e. ±1.645.

Figure 19: Response of expected duration

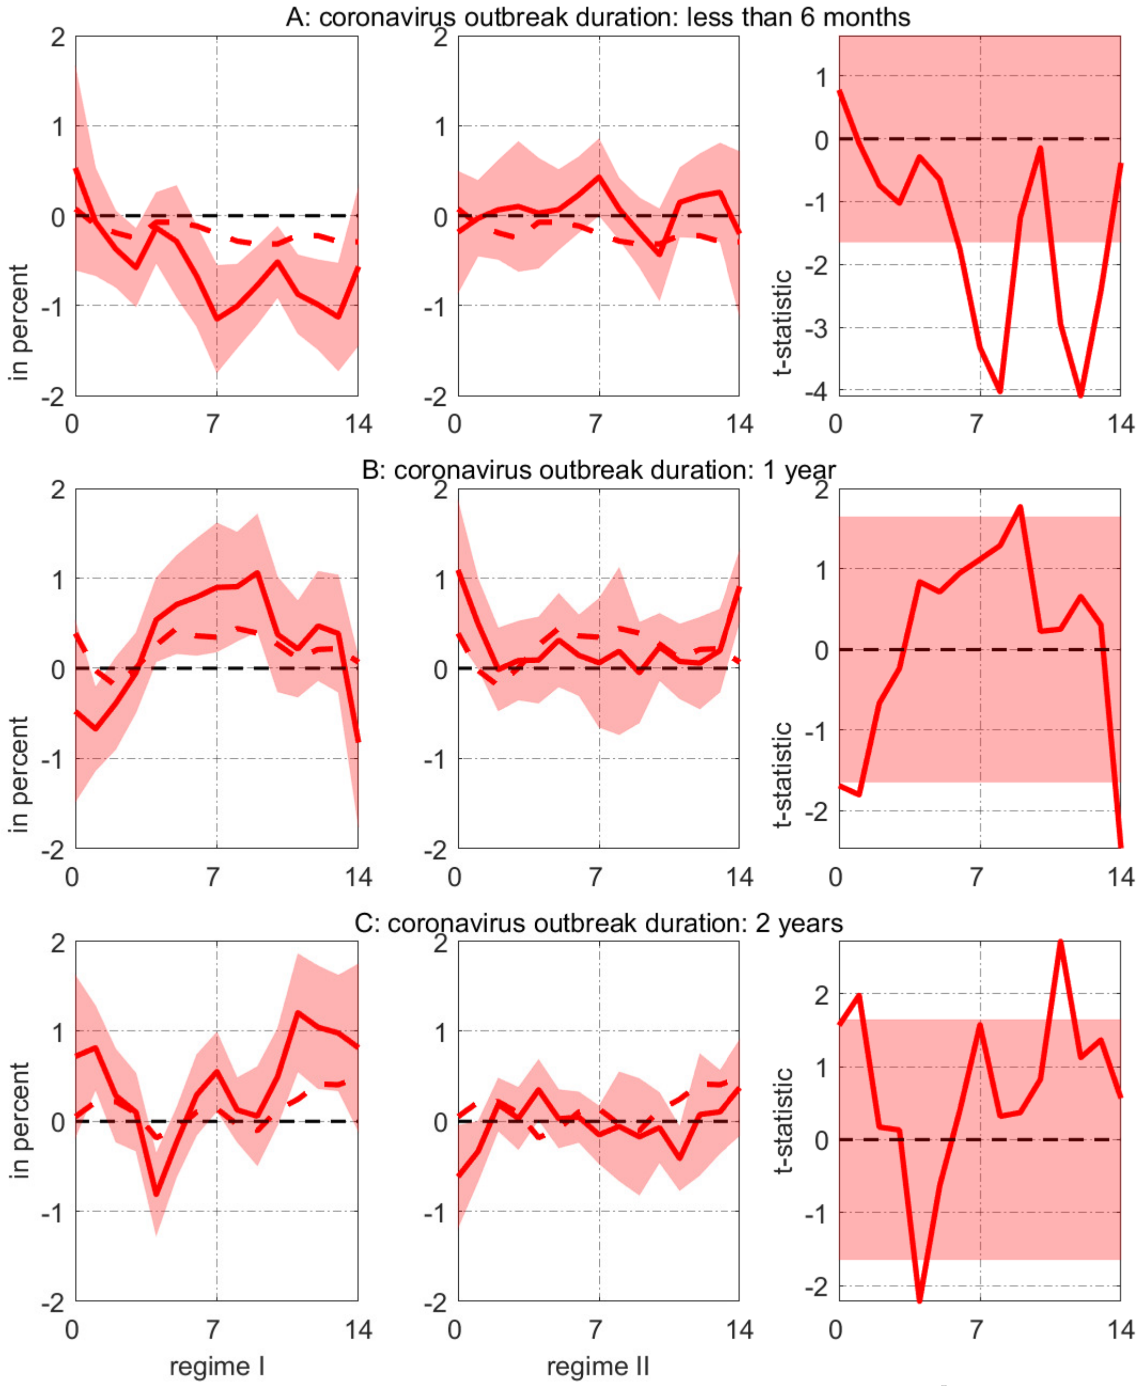

Notes: The first column shows the impulse response coefficients (red-solid) β_h^I for $h = 0, ..., H$ in regime I following a pandemic shock (one standard deviation); the second column shows the corresponding impulse response coefficients β_h^{II} in regime II. In both cases, the red-shaded area corresponds to the 90 percent confidence interval relying on Driscoll-Kraay standard errors. The red-dotted lines in the first two columns correspond to the impulse response coefficients from the linear model without allowing for state-dependent effects. The third column shows the $t-$statistics testing the null that $H_0 : \beta_h^I - \beta_h^{II} = 0$ for each horizon using the Driscoll-Kraay method. The red-shaded area covers the $t-$critical values for a 90% confidence interval, i.e. ± 1.645.

Essay VII:

The Role of Global and Domestic Shocks for Inflation Dynamics: Evidence from Asia

This paper is published as:

> **Finck, David and Peter Tillmann**, "The Role of Global and Domestic Shocks for Inflation Dynamics: Evidence from Asia," *Oxford Bulletin of Economics and Statistics*, 2022, *84*(5), 1181–1205.

This paper was presented at the following workshops and international conferences:

I: Price-setting Behavior and Inflation Dynamics in SEACEN Member Economies and Their Implications for Inflation (Kuala Lumpur, Malaysia)
Date: November 2018
Presenter: David Finck

II: Inflation and Deflation in Asia, organized by the KDI school of Public Policy (Virtual)
Date: April 2022
Presenter: David Finck

Acknowledgement

This project has been conducted while we were Visiting Research Economists at the SEACEN Centre. We thank the SEACEN Centre for its hospitality. Seminar participants at the Reserve Bank of Australia, Giessen and the SEACEN Centre provided very helpful comments. We are particularly grateful to comments from the editor of this journal and three anonymous reviewers. We also thank Peter Winker, Daniel Grabowski and Jörg Schmidt for comments and suggestions. Conference participants from Inflation and Deflation in Asia, organized by the KDI school of Public Policy and Management provided insightful comments.

The Role of Global and Domestic Shocks for Inflation Dynamics: Evidence from Asia

David Finck* Peter Tillmann†

Abstract

This paper studies inflation dynamics and the output-inflation trade-off in small open economies. We estimate a series of VAR models for a set of six Asian emerging market economies, in which we identify a battery of domestic and global shocks using sign restrictions. We find that global shocks explain large parts of inflation and output dynamics. A series of counterfactuals support these findings and suggest that the role of monetary policy is limited. We estimate reduced-form Phillips curve regressions based on alternative decompositions of output into global and domestic components. For most countries, we find a positive and significant correlation between inflation and the fraction of GDP driven by domestic shocks only. While including the output component driven by oil prices seems to 'flatten' the Phillips curve, though the effect is not significant, the component driven by global demand shocks 'steepens' the inflation-output nexus.

Keywords: Inflation Targeting, Business Cycle, Open Economy, Monetary Policy, Phillips Curve

JEL classification: E3, E5, F4

*University of Giessen, email: david.finck@wirtschaft.uni-giessen.de
†University of Giessen, email: peter.tillmann@wirtschaft.uni-giessen.de

I Introduction

Over the past two decades, many advanced and emerging economies experienced low and relatively stable inflation rates. At the same time, inflation more and more appeared to be decoupled from economic activity. The sharp drop in GDP during the Great Recession did not lead to a further drop in inflation, thus giving rise to the 'missing deflation' phenomenon. Likewise, the strong economic recovery did not go hand in hand with rising inflation rates. Based on these observations, a large literature studies the changing nature of inflation dynamics and, in particular, the shifting relationship between inflation and economic activity. This research agenda is often described in terms of a 'flattening' of the Phillips curve (see Coibion and Gorodnichenko, 2015, and others). A reduction in the slope of the Phillips curve relationship would have consequences for monetary policy. For given inflation expectations, the argument goes, a flatter Phillips curve would require a deeper recession in order to bring high inflation back to the target.

One explanation for the apparent changes in the inflation process is the ongoing global integration of financial and goods markets. However, quantifying the extent to which global forces explain inflation is not straightforward. In her survey article, Forbes (2019) argues that the role of global factors (commodity prices, measures of global slack, exchange rates, price competition) has changed over time. She finds that the relation between domestic output gaps and inflation rates has weakened and advocates that models of inflation such as the Phillips curve should incorporate changes in the global economy in order to provide a good account of the determinants of inflation. If global factors are indeed driving a substantive share of inflation, domestic monetary policy is less able to stabilize inflation and the real economy. As monetary policy primarily affects inflation through expanding or contracting domestic demand, the power of central banks to control inflation would be limited in a world in which global forces dominate.

In this paper, we add to this literature and study six Asian emerging market countries: Indonesia, Korea, Malaysia, the Philippines, Singapore and Thailand. These countries are prototypical small open economies that are well integrated into the world economy.[1] In addition, all six economies explicitly or implicitly have a monetary policy mandate for maintaining price stability.[2] Recently, IMF (2018) claims that the sensitivity of inflation rates in Asian emerging market economies with respect to real activity declined, thus leading to a flatter Phillips curve.

We shed new light on the determinants of inflation by disentangling domestic and global driving forces based on a series of counterfactuals. For example, we split GDP growth into the component that reflects domestic shocks and the part that is driven

[1] Auer and Mehrotra (2014) argue that the integration of Asian economies into global supply chains matters. The correlation of inflation rates across Asian economies increases with the extent of their bilateral trade relationships.

[2] See Volz (2015) for a discussion of the experience with inflation targeting in Asia.

by global demand shocks. These decompositions show that GDP driven by domestic shocks elicits a Phillips curve relationship that differs from the relationship between inflation and the part of GDP driven by global shock.

We proceed in three steps. First, we estimate a series of structural vector autoregressive (VAR) models, in which we use alternative sets of constraints to identify a battery of shocks. In our baseline model, we apply sign restrictions as in Corsetti et al. (2014) to identify domestic demand and supply shocks as well as global demand and supply shocks. While demand and supply shocks can be distinguished based on the responses of inflation and GDP growth being positively correlated (in the case of demand shocks) or negatively correlated (for supply shocks), we disentangle domestic from global shocks based on the relative response of domestic GDP to world GDP. Second, the identified VAR model allows to apply several structural analyses in order to address the role of the structural shocks on inflation and on the growth rate of real GDP. In doing so, we focus on four categories of shocks, i.e. global, domestic, monetary policy and residual shocks. We first decompose the variance of forecast errors. The forecast error variance decomposition (FEVD) tells us how much of the forecast error variance can be explained by exogenous shocks to other variables in the system. While the FEVD describes average movements in the data, it does not allow us to quantify the amount of how much of the observed variability is explained by specific shocks. Hence, we also decompose the history of inflation and GDP growth into the historical contributions of each shock in order to quantify the cumulative effects on these series. Our results of both the FEVDs as well as the historical decompositions suggest that global shocks play an important role for all six countries under investigation. In particular, global shocks are an important driver of inflation around the Great Recession, as they explain most of the increase and subsequent plunge in inflation rates.

We run counterfactual simulations in order to derive the hypothetical effects of shocks in the past on today's outcomes. By changing the history of selected structural shocks, this exercise summarizes the results of the historical decomposition and shows how our endogenous variables would have evolved in the absence of these shocks. In the third step, we revisit the Phillips curve relationship. The VAR model provides us with the domestic component of GDP, i.e. the fraction of GDP that is driven by all shocks other than global shocks, and global components of GDP, i.e. the components of economic activity driven by global demand and supply shocks. We study the inflation-output correlation using this decomposition of GDP.[3] The model allows different domestic and global components to enter the Phillips curve with different

[3]The growth rate of GDP is just one possible indicator of slack in the economy. Alternative indicators are the output gap, i.e. the difference between the levels of actual and potential GDP, real marginal cost or measures incorporating information from the tightness of the labor market. Krause and Lubik (2007), Faccini et al. (2013), Thomas and Zanetti (2009), Zanetti (2011) and Trigari (2009) show the relevance of slack derived from labor market conditions for inflation dynamics.

coefficients and potentially different signs. Hence, the model nests the conventional reduced-form Phillips curve specification if the coefficient on the components of real GDP growth driven by oil supply shocks and global demand shocks equal the one on the domestic component.

Our results suggest that the nature of global shocks matters. We see that global supply shocks seem to flatten the Phillips curve throughout our set of countries, though the effect is not statistically significant, while the opposite is true for global demand shocks. Importantly, we get similar results when we use an alternative identification strategy.

Our paper connects several strands of the literature: First, a recent branch of the literature studies the comovement of inflation rates across countries. Ciccarelli and Mojon (2010) find that for 22 OECD countries, a single common factor explains about 70% of the variation in inflation. They refer to this phenomenon as 'global inflation'. Unless real economic activity is equally well explained by a common factor, this implies a weakening of the relationship between domestic output and inflation. The evidence provided by Neely and Rapach (2011) and Mumtaz and Surico (2012) supports this finding. In contrast, Förster and Tillmann (2014) show evidence that is consistent with 'local inflation', that is inflation being primarily driven by domestic variables. Recently, Parker (2018) uses a very large data set with more than 200 economies to show that the global inflation hypothesis does not fit emerging and developing countries, in which only a subset of prices such as those for oil and food are driven by global shocks.

A second strand of research argues that conventional Phillips curve regressions that relate inflation to, among other variables, a measure of domestic slack such as the output gap should be augmented by measures of global slack or a 'global output gap'. Borio and Filardo (2007), using a cross-section of countries, find that the explanatory power of global factors as reflected in measures of global slack increased over time. For some countries, these authors find global factors to be the dominant drivers of inflation. Supportive evidence for advanced economies is provided by Milani (2009), Milani (2010), while Ihrig et al. (2010) cannot find evidence in favor of the 'globalization of inflation' hypothesis. Bems et al. (2018) include additional global variables into an otherwise standard New Keynesian Phillips curve estimated for 19 emerging market economies. The authors show that domestic factors are the most important drivers of inflation. Okuday et al. (2019) use sectoral data from Japan to show that an increase in the heterogeneity of shocks contributes to a weaker response of inflation to real economic activity.

The concept of global output gaps often used in the literature, however, is not without flaws (see Tanaka and Young, 2008, and Gerlach, 2011). Jašová et al. (2020) point to the fact that for a typical small open economy the domestic output gap should be highly correlated with the global output gap, i.e. the weighted gap of the economy's main trading partners. This correlation obscures the identification of the true structural driving forces of inflation dynamics. The approach taken in this paper, in contrast, identifies orthogonal domestic and global components of output based on the comovement between global and domestic variables. This procedure avoids some of the weaknesses of estimates of global slack.

A third strand uses identified time-series models to study the determinants of inflation dynamics together with other key business cycle variables. As mentioned before, Corsetti et al. (2014) and Bobeica and Jarociński (2019) propose a set of sign-restrictions that allows us to quantify the response to orthogonal domestic and global shocks, respectively. Conti et al. (2017) apply a similar identification scheme to decompose euro area inflation. All three papers attribute an important role to global driving forces of inflation. As an increasing integration of goods and financial markets should make global factors more important over time, Bianchi and Civelli (2015) allow the coefficients of their VAR model to vary over time. Their evidence suggests that global slack as a determinant of inflation does not become more important over time. Eickmeier and Kühnlenz (2018) focus on the role of China for inflation dynamics in advanced and emerging economies. Estimating a factor model for 38 countries, they find that demand and supply shocks originating in China have a significant impact on inflation in other economies.

The remainder of this paper is organized as follows. Section II introduces our empirical framework, including the data set and the identification strategy. The results, i.e. impulse responses, forecast error variance decompositions, historical decompositions and counterfactual simulations, are discussed in Section III. Section IV examines the Phillips curve trade-off based on the domestic and global components of output. Section III draws conclusions for monetary policy. An appendix contains additional results.

II Empirical Framework

The empirical analysis in this paper is based on an identified VAR model, as pioneered by Sims (1980). Much of the analysis that follows is based on the interpretation of structural shocks, i.e. disturbances that drive the dynamics of our economic variables. Therefore, we will carefully describe how the structural shocks in our analysis are identified.

A. *The VAR-Model*

Our model can be written as

$$\mathbf{y}_t = \mathbf{c} + \mathbf{B}_1\mathbf{y}_{t-1} + \cdots + \mathbf{B}_p\mathbf{y}_{t-p} + \mathbf{u}_t, \quad t = p+1, ..., T, \tag{1}$$

where $\mathbf{y}_t$ is an $n \times 1$ vector of endogenous variables, which in our case will include key macroeconomic time series. Furthermore, $\mathbf{B}_1, ..., \mathbf{B}_p$ are $n \times n$ matrices capturing the VAR-coefficients and $\mathbf{u}_t$ is an $n \times 1$ vector of residuals which is assumed to follow a multivariate normal distribution $\mathbf{u}_t \sim \mathcal{N}(0, \boldsymbol{\Sigma})$.

A major challenge when dealing with impulse responses from VAR models with $\boldsymbol{\Sigma}$ being unrestricted a-priori is that they arise from shocks that are correlated. Put differently, the variance-covariance matrix $\boldsymbol{\Sigma}$ of the reduced form VAR as in (1) is typically not diagonal. In that case, the interpretation of impulse responses is likely to be misleading given the fact that shocks typically arise simultaneously. To overcome this issue, we derive structural VARs (SVARs) for each country as they allow us to obtain the responses of variables to orthogonal shocks.
To do so, note that (1) can be formulated in a structural form that reads

$$\mathbf{A}_0\mathbf{y}_t = \boldsymbol{\mu} + \mathbf{A}_1\mathbf{y}_{t-1} + \cdots + \mathbf{A}_p\mathbf{y}_{t-p} + \varepsilon_t, \tag{2}$$

where $\varepsilon_t \sim \mathcal{N}(0, \boldsymbol{\Gamma})$ is the vector of structural shocks we are interested in. Notice that the VAR is now augmented by $\mathbf{A}_0$ such that $\boldsymbol{\Gamma}$ will be a diagonal identity matrix, i.e. the structural disturbances in ε_t are mutually independent. This is reasonable from the view that structural disturbances are uncorrelated and arise independently.

In our estimation, we rely on a Bayesian framework, where the coefficients as well as the residual variance-covariance matrix are understood as random variables and characterized by some probability distribution. The basic principle of Bayesian analysis is to combine subjective prior information with the likelihood function according to the Bayes rule in order to derive a posterior distribution which combines both sources of information. In our benchmark model, we use $p = 2$ lags. The priors for the VAR coefficients in $\mathbf{B}_1, \cdots, \mathbf{B}_p$ as well as for $\boldsymbol{\Sigma}$ are both centered around the OLS estimates. As regards the prior for the VAR coefficients, we assume a multivariate normal distribution coupled with a Minnesota prior where the autoregressive coefficients are set to 0.8. The prior for $\boldsymbol{\Sigma}$ follows an inverse Wishart distribution. As regards the choice of hyperparameters, we choose an overall tightness of $\lambda_1 = 0.1$ and a lag decay of $\lambda_3 = 2$.[4] Each country-specific benchmark estimation relies on 5000 draws, where the first 3000 draws are discarded because the first draws of the joint posterior are likely to be not representative for the target distribution we are looking for. In the appendix, we report various diagnostic statistics for the convergence properties of the sampler.

[4]We have repeated the estimation for each country and performed a grid search in which the hyperparameters are selected optimally. To do so, we specified minimum and maximum values for each hyperparameter and a step size for the grid search. In a next step, we estimated the marginal likelihood for each possible combination on the grid. We then chose the combination of hyperparameters for which the marginal likelihood is maximized. We find that our results, available on request, are unaffected by the choice of hyperparameters.

B. *Data and Shock Identification*

We estimate the model separately for six Asian emerging market economies: Indonesia (IDN), Korea (KOR), Malaysia (MYS), the Philippines (PHL), Singapore (SGP) and Thailand (THA). The vector of endogenous variables includes the oil price, real GDP, consumer prices, the short-term interest rate as a measure of monetary policy, the real effective exchange rate and the share of domestic real GDP in world real GDP. The latter variable will be particularly important in order to separate global from domestic shocks.[5] All variables other than the share in world GDP and the interest rate are expressed as year-on-year growth rates in percentage points. The data covers the sample period 2001Q1 to 2018Q1 and the frequency is quarterly.

Our choice of year-on-year growth rates is motivated by the Chinese New Year, which imposes a seasonal pattern on the data which is different compared to advanced economies. Using year-on-year rates allows us to ignore seasonal adjustment. The share in world GDP is included in differences (percentage point change from the year before). An increase in the real effective exchange rate corresponds to a real appreciation of the domestic currency.[6]

To identify structural shocks, we use two sets of alternative restrictions, see Table (1). Both sets impose alternative sign restrictions onto the variables, while the difference between both sets lies in the identification of monetary policy shocks.

I. Baseline sign restrictions

Our first identification strategy broadly follows Corsetti et al. (2014), who impose a mixture of sign and zero restrictions in order to distinguish domestic shocks from global shocks as well as supply shocks from demand shocks. We implement sign restrictions as in Arias et al. (2018). The key variable in this identification pattern is the share of real domestic GDP in world real GDP, as it allows us to distinguish disturbances that hit the global economy more than the domestic economy and vice-versa.

Both a domestic as well as a global demand shock are supposed to decrease both domestic prices as well as domestic real GDP. The imposed negative sign on the GDP share means that domestic real GDP decreases more than real GDP in the rest of the world does, i.e. the effect of a domestic demand shock has a stronger effect on domestic GDP. In contrast, a global demand shock leads to an increase in the share of domestic

[5] We use headline inflation instead of core inflation. This is because central banks typically use the growth rate of the overall price index as a target variable. Moreover, data on core inflation is not available for all six countries under consideration.

[6] The data on GDP, prices and the interest rate is taken from the CEIC data base. Oil prices are drawn from the FRED data base at the St. Louis Fed. For the real effective exchange rate we use the data provided by the BIS pertaining to a broad set of trading partners. The share in world GDP is drawn in annual frequency from the world economic outlook (2018) and interpolated (cubic spline interpolation method) to quarterly frequency.

Table 1: Identification of structural shocks

VARIABLE \ SHOCK	OIL SUPPLY	GLOBAL DEMAND	DOMESTIC DEMAND	DOMESTIC SUPPLY	MONETARY POLICY	RESIDUAL
I. BASELINE SIGN RESTRICTIONS (CORSETTI ET AL., 2014)						
OIL PRICE	+	–				0
GDP SHARE		+	–	–	0	0
REAL GDP	–	–	–	–	0	0
INFLATION	+	–	–	+	0	0
INTEREST RATE	0		–	+	+	0
EXCHANGE RATE						+
II. ALTERNATIVE SIGN RESTRICTIONS						
OIL PRICE	+	–				0
GDP SHARE		+	–	–	–	0
REAL GDP	–	–	–	–	–	0
INFLATION	+	–	–	+	–	0
INTEREST RATE	0		–	+	+	0
EXCHANGE RATE						+

NOTES: Blank cells indicate unconstrained impulse responses. A positive or negative reaction is denoted by + and −. A zero restriction is denoted by 0. All restrictions are imposed on impact only.

real GDP relative to the rest of the world GDP, implying that a global demand shock has a stronger effect on the rest of the world, though domestic real GDP and domestic consumer prices are assumed to decrease.

In order to further distinguish global demand shocks from domestic demand shocks, we also assume that a global demand shock leads to a decline of the oil price, while both the interest rate as well as the exchange rate remain unrestricted. Because we focus on small open economies, domestic demand shocks remain unrestricted with respect to the response of the oil price, but the interest-rate is assumed to decrease in order to fight deflationary pressure.

An oil supply shock, which is intended to represent a global supply shock, is restricted to decrease domestic real GDP and increase inflation, while the immediate response of the domestic short-term interest rate is restricted to zero.[7] This is due to the fact that that the central bank does not contemporaneously respond to oil shocks.[8]

In contrast to a domestic demand shock, a domestic supply shock leads to opposite responses of output and prices. In order to get distinct global and domestic supply shocks, the domestic supply shock is also assumed to decrease domestic real GDP relative to real GDP of the rest of the world. The restrictions on the monetary policy shock imply that variables other than the exchange rate respond with at least one month delay to an increase in the interest rate.

We include the exchange rate in our VAR model because we are studying small open economies. The residual shock accounts for fluctuations of the exchange rate which are not explained by all other domestic and global shocks. Hence, following Bobeica

[7]Indonesia is a net oil importing country, although the country also had net oil exports in the past. Malaysia is a (small) net oil exporting country. These potential limitations should be kept in mind when discussing the results for oil price shocks.

[8]For a further discussion, see Corsetti et al. (2014) and Bobeica and Jarociński (2019).

and Jarociński (2019), a residual shock is a surprise change in the exchange rate that contemporaneously keeps all other variables unchanged.[9]

II. Alternative sign restrictions

An alternative set of sign restrictions follows Bobeica and Jarociński (2019), who adopt a mixture of the sign restrictions proposed in Corsetti et al. (2014) and Baumeister and Benati (2013). The identification scheme differs from the baseline set of restrictions only with respect to the identification of the monetary policy shock. It is assumed that real activity and inflation immediately respond to a monetary policy shock in a way that is consistent with standard theories of monetary transmission. That is, following an unexpected increase in the policy rate, inflation and real GDP growth are assumed to fall. Since the monetary policy tightening has a stronger effect on the domestic economy than the rest of the world, the GDP share is also assumed to drop immediately after the shock hits the economy.

The restrictions are imposed on impact only. This is because we want to keep the restrictions as light as possible. Since imposing restrictions on the impact period only suffices for our purpose, we refrain from imposing addition structure on the model. All results reported throughout this paper are based on the baseline identification strategy I, i.e. the Corsetti et al. (2014) identification. The results based on identification strategy II will not be shown in order to save space but are available in the appendix. We find that most of our results remain unchanged under the alternative identification.

III Results

This section reports the main results of our paper. We first start with the results for a variance decomposition of the forecast errors. We then discuss the results from a historical decomposition and from counterfactual experiments, where the latter serve as input for our exercise in Section IV. Both for the historical decomposition and for the counterfactuals, we summarize the contributions of all shocks into four main categories: (i) monetary policy shocks, (ii) residual shocks, (iii) domestic shocks and (iv) global shocks. Thereby, we use 'domestic shocks' as the umbrella term for both domestic demand as well as domestic supply shocks. Meanwhile, 'global shocks' summarize both global demand shocks and oil supply shocks, while monetary policy shocks and residual shocks are the remainder.

A. *Forecast Error Variance Decomposition*

A structural decomposition of the forecast error variance tells us how much of the forecast error variance is due to exogenous shocks and thus indicates the amount

[9] Singapore operates a system of a managed exchange rate against a basket of currencies. Hence, a residual shock as identified here could also be interpreted as monetary policy shock.

of information each variable contributes to the other variables in the autoregressive process.

Table 2: Forecast error variance decompositions for real GDP growth

		IDN	KOR	MYS	PHL	SGP	THA
$h = 4$	*residual*	5	3	1	1	1	1
	monetary policy	4	0	1	1	0	1
	oil supply	9	10	16	35	18	18
	global demand	49	39	29	14	14	16
	domestic demand	5	39	46	14	38	36
	domestic supply	29	9	7	35	29	28
$h = 8$	*residual*	9	6	2	2	2	2
	monetary policy	6	1	2	2	1	1
	oil supply	10	13	21	37	21	19
	global demand	42	33	26	13	17	16
	domestic demand	7	35	40	16	32	33
	domestic supply	27	12	8	29	27	28
$h = 12$	*residual*	9	7	3	2	2	3
	monetary policy	6	1	3	3	2	2
	oil supply	10	13	20	36	20	19
	global demand	40	34	26	14	17	17
	domestic demand	8	34	38	17	33	32
	domestic supply	27	12	9	28	26	27

Notes: Median shares (in %) of forecast error variance for GDP growth due to structural shocks for different forecast horizons. All results rely on the Corsetti et al. (2014) identification and rounded to integers.

Table (2) reports the results of the forecast error variance decomposition for GDP growth for different horizons, i.e. for $h = 4, 8$ and 12. Our results imply that both domestic and global demand as well as domestic and global supply shocks in particular are the main drivers of real GDP growth. For our research question, the role of global shocks is of major interest. For $h = 4$, we find that oil supply shocks can explain between 10% (KOR) and 35% (PHL) of the variance of forecast errors. Global demand shocks also seem to be particularly important for our six countries, with an explanatory power ranging from 14% (PHL and SGP) to 49% (IDN). Interestingly, the explanatory power is relatively constant across the forecast horizons. For all six countries, our results imply that both the monetary policy shock as well as the residual shock play a minor role only.[10]

[10] We also estimate a Panel VAR with the alternative set of sign restrictions and the same prior specification as for our single country estimations. While details are provided in the appendix, our pooled estimates suggest that global shocks explain 30% of the fluctuation over two years. For

Table 3: Forecast error variance decompositions for inflation

		IDN	KOR	MYS	PHL	SGP	THA
$h = 4$	*residual*	1	1	1	1	1	1
	monetary policy	1	0	1	0	1	0
	oil supply	5	26	50	62	13	16
	global demand	15	21	23	13	46	27
	domestic demand	34	32	10	19	31	46
	domestic supply	43	19	16	6	8	9
$h = 8$	*residual*	2	2	2	2	2	3
	monetary policy	2	1	2	2	3	1
	oil supply	8	20	43	55	13	16
	global demand	22	24	25	15	30	24
	domestic demand	29	36	12	17	35	44
	domestic supply	36	17	17	9	17	11
$h = 12$	*residual*	3	3	2	2	3	4
	monetary policy	3	2	2	2	4	1
	oil supply	9	19	42	53	13	16
	global demand	24	23	25	15	28	24
	domestic demand	29	36	13	18	33	43
	domestic supply	33	17	17	9	20	12

Notes: Median shares (in %) of forecast error variance for inflation due to structural shocks for different forecast horizons. All results rely on the Corsetti et al. (2014) identification and rounded to integers.

For the inflation rate, see Table (3), we also find that global shocks can explain a large part of the variance of the forecast error, but the contributions are more heterogeneously distributed than for real GDP growth. For $h = 4$, our results imply that the oil supply shock is most important for Malaysia (50%) as well as for the Philippines (62%). For Indonesia, oil supply shocks can explain only 5% of the variance of the forecast error. Global demand shocks are most important for Singapore with an explanatory power of 46%. For all other countries, the values range between 13% and 27%. We also find that the contributions become more convergent as the forecast horizon becomes longer. Overall, we find that global shocks are an important driver of inflation in our small open economies.

B. *Historical Decomposition*

While structural forecast error variance decompositions and structural impulse response functions describe average movements in the data, they do not allow us to

GDP growth, global shocks are even more important. They account for 52% of fluctuations of real economic activity over the two-year horizon. See the appendix for details.

quantify how much of the historically observed fluctuations of a variable is explained by one specific shock. Even though our results so far suggest that both global and domestic shocks are important drivers of inflation and GDP growth, we do not know the effect of past (known) shocks on the fluctuation of these variables. Hence, to establish the contribution of structural shocks to the dynamics of our data series, we depart from unconditional expectations and derive the posterior distribution of structural historical decompositions for every endogenous variable. Contrary to the average contribution of our identified shocks to the variability of inflation and GDP growth from 2001 to 2018, we are now interested in the cumulative effects of past shocks. Similar to the previous section, we will only report the results for inflation and GDP growth.

We can decompose the vector of endogenous variables $\mathbf{y}_t$ into a vector of contributions from deterministic variables $\mathbf{d}^{(t)}$ and historical contributions of structural shocks. Considering variable i, the historical decomposition reads

$$\mathbf{y}_{i,t} = \mathbf{d}_i^{(t)} + \sum_{j=1}^{n} \sum_{k=0}^{t-1} \widetilde{\psi}_{k,ij} \varepsilon_{j,t-k}. \tag{3}$$

This expression states that $\mathbf{y}_{i,t}$ is the sum of the deterministic component $\mathbf{d}_i^{(t)}$ and the sum of contributions of all $j = 1, \ldots, n$ structural shocks on variable i from period $k = 0$ to $t-1$ back in the past, where $\varepsilon_{j,t-k}$ is the structural shock j in period $t-k$ and $\widetilde{\psi}_{k,ij}$ is the corresponding entry in row i and column j of the structural impulse response function matrix $\widetilde{\mathbf{\Psi}}_k$, i.e. it corresponds to the impact of shock j on variable i. The historical decomposition in (3) shows that, for example, a positive contribution of structural shock j to variable i means that shock j pushes variable i above the deterministic component, i.e. the unconditional forecast in the absence of any shocks.

Figures (1) and (2) show the historical contributions of structural shocks for the inflation rates for all countries. The black line reflects the difference between the unconditional forecast (i.e. the deterministic part) generated by the VAR and the actual data series, while the colored bars highlight the fraction of this series explained by each of the four groups of shocks.

When interpreting the historical contribution of structural shocks, it is important to note that negative values do not correspond to periods of disinflation, but negative contributions that push the inflation rate below the deterministic component which is non-negative for all countries throughout the entire sample.[11]

Four key results stand out for all countries.[12] First, while for some countries the

[11] That is, even when structural shocks contribute negatively to inflation dynamics, we can still observe positive inflation rates when the deterministic component is greater than the overall contribution of all structural shocks.

[12] It is worth noting that identification strategy II yields very similar results, although we find that the role of monetary policy shocks is more important for most countries, a finding which is qualitatively similar to the results in Bobeica and Jarociński (2019). Nevertheless, also here we find that global

Figure 1: Historical contribution of structural shocks to inflation for Indonesia, Korea and Malaysia

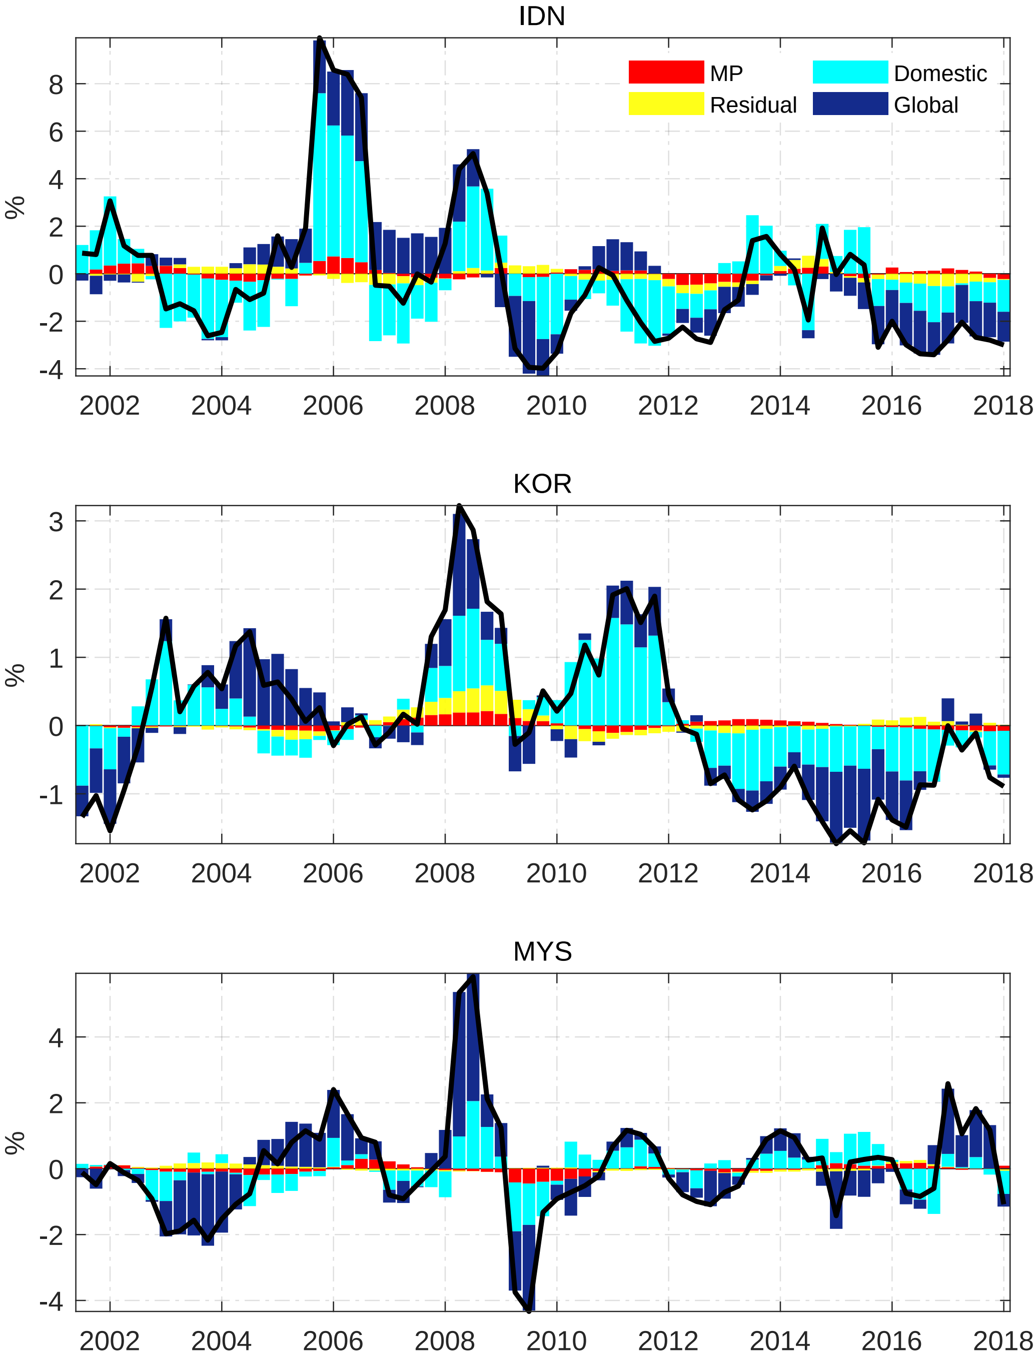

Notes: Median historical contribution of monetary policy shocks (red bars), residual shocks (yellow bars), domestic shocks (teal bars) and global shocks (blue bars) to inflation for Indonesia, Korea and Malaysia. The black path corresponds to the sum of median contributions of all structural shocks. Results rely on identification strategy I.

Figure 2: Historical contribution of structural shocks to inflation for the Philippines, Singapore and Thailand

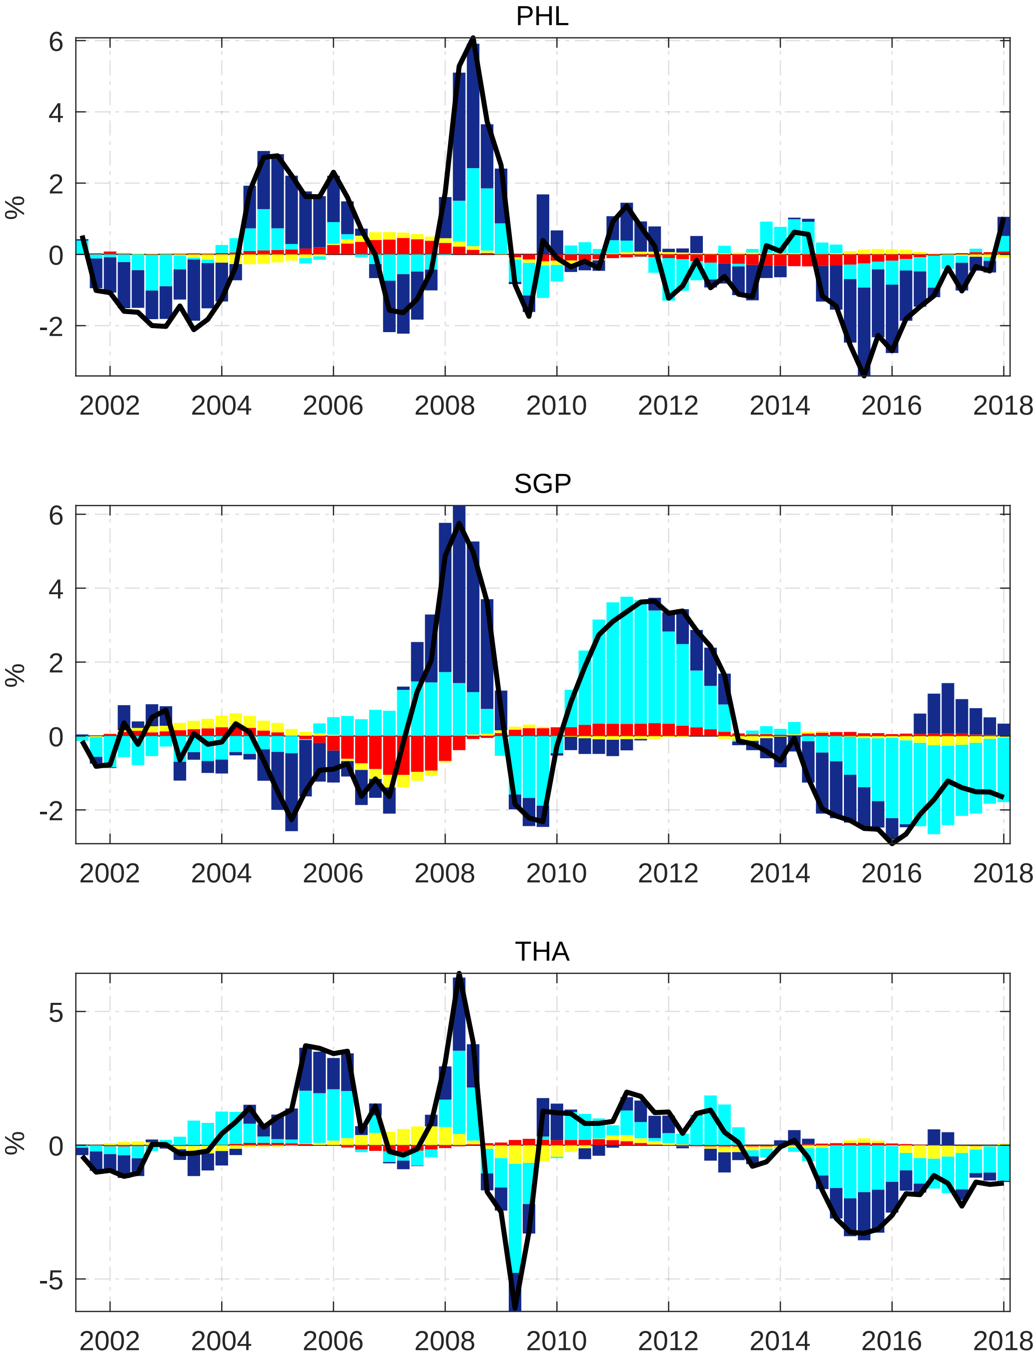

Notes: Median historical contribution of monetary policy shocks (red bars), residual shocks (yellow bars), domestic shocks (teal bars) and global shocks (blue bars) to inflation for the Philippines, Singapore and Thailand. The black path corresponds to the sum of median contributions of all structural shocks. Results rely on identification strategy I.

effect of domestic shocks dominates, global shocks (as contributions of oil supply shocks plus global demand shocks) play an important role for inflation. Both sources of inflation dynamics, i.e. domestic and global shocks, are typically positively correlated, that is they jointly push inflation up or down. There are only very few episodes in which both forces push inflation into opposite directions. If both domestic and global shocks were negatively correlated, shocks would partly offset domestic driving forces. However, the results suggest that global shocks exacerbate inflation fluctuations, thus requiring a more aggressive monetary policy response.

Second, global shocks are particularly important in 2008/9. They drive inflation up before the global financial crisis and contribute to the fall in inflation during the subsequent Great Recession. Third, the very low levels of inflation observed more recently are partly due to global forces. In Korea, Singapore and Thailand, global shocks put downward pressure on inflation after 2014.

Fourth, both residual and monetary policy shocks contribute relatively little to the fluctuation of the inflation rate. While monetary policy shocks play some noteworthy role around the Great Recession in Singapore, the Philippines and Korea, they have almost no role in the dynamics of inflation in Malaysia and Thailand. This result suggests that central banks effectively stabilize the economy with only small deviations of monetary policy from its systematic component.

Exogenous fluctuations in the real exchange rate play a minor role for inflation dynamics. This is particularly interesting in light of the strong exchange rate movements in emerging economies around the adoption and the unwinding of the Federal Reserve's Quantitative Easing. It is, however, important to keep in mind that the historical decomposition dissects inflation into structural shocks, i.e. into exogenous changes of the exchange rate. Hence, the finding that residual shocks play a small role is consistent with the notion of central banks are effective in stabilizing inflation in light of exchange rate movements.

We would expect the cross-country correlation between global shocks (both supply and demand) to be positive, since these shocks stem from abroad and, for each country, are identified by the same co-movement between domestic and foreign variables. The appendix shows that the correlation of global supply shocks is high across countries. The correlation is always positive and in most cases above 50%. It is not surprising that the correlation of global demand shocks is slightly lower than for global supply shocks. The reason is that, as for the demand shock, we separate domestic from global demand shocks by exploiting the comovement of domestic real GDP relative to real GDP of the rest of the world. This is not the case for the global supply shock which is identified as an oil supply shock. This shock is assumed to hit all countries equally and we do not restrict the response of real GDP of a single country relative to real GDP

shocks are very important for the fluctuation of both the inflation rate and the growth rate of real GDP.

of the rest of world.

Summing up, our historical decompositions support our findings from the FEVDs insofar as global and domestic shocks seem to be the main drivers of inflation across countries.[13] However, they also uncover that global shocks are primarily important in 2008/9 by explaining most of the increase in inflation in 2008 and the subsequent fall thereafter. Finally, our results in this section suggest that global shocks account for much of the recently observed low inflation rates, especially in Korea, the Philippines, Singapore and Thailand.

C. *Counterfactual Analysis*

Now that we already have first impressions of the contributions of our structural shocks, we use these results and visualize the importance of our structural shocks even further in counterfactual scenarios. We run a battery of counterfactual experiments in order to shed light on the role of alternative drivers of inflation and the business cycle.

We separately show how inflation and real GDP would have looked like in the absence of either domestic, global, monetary policy or residual shocks. The previous analysis provides us with everything we need in order derive these counterfactual paths, because these counterfactuals are the difference between the actual data and the contributions of structural shocks we have already derived before.[14] In a first scenario, we study inflation in the absence of selected structural shocks. For that purpose we suppress (1) both the global demand and oil supply shock, (2) the monetary policy shock, (3) both the domestic demand and the domestic supply shock, and (4) the residual shock.

This experiment follows, among others, Sims and Zha (2006) and can be summarized as follows: given the data, it is possible to draw all parameters from the joint posterior distribution. It is then easy to recover a sequence of unit-variance structural shocks (as described in Section II) and simulate a series that would have been observed, given the vector that contains the suppressed structural shocks. This is straightforward as we already have derived the historical decomposition.[15]

[13]Also the results from our Panel VAR support the finding that global and domestic shocks are the main drivers of inflation dynamics in our six countries. However, the Panel results suggest that global shocks are more important than domestic shocks, while also the role of both monetary policy shocks and residual shocks is more important than in our individual VARs.

[14]Note that in a historical decomposition, usually only the median (or mean) estimates are shown, so that uncertainty is ignored. However, this is usually the choice of the researcher and even in a historical decomposition one could easily show uncertainty if one wanted to.

[15]Even without deriving the historical contributions of each structural shocks, one could also construct the same counterfactual data as follows: for each draw of our estimation procedure, recover the VAR coefficients as well as the structural matrix $\mathbf{A}_0$ as Section II. Then derive the vector of structural shocks ε_t. Setting different structural shocks to zero results in a vector $\widetilde{\varepsilon}_t$ that can be used to construct the counterfactual paths. This is done by simulating the vector of counterfactual data as $\mathbf{y}_t^{cf} = \mathbf{c} + \mathbf{B}_1 \mathbf{y}_{t-1}^{cf} + ... + \mathbf{B}_p \mathbf{y}_{t-p}^{cf} + \mathbf{A}_0^{-1} \widetilde{\varepsilon}_t$. Finally, note that we are interested here in the role of structural shocks over the full sample. One can still perform interesting experiments and, for example, simulate away or change

Remember that each variable of our vector $\mathbf{y}_t$ of endogenous variables can be rewritten as

$$\mathbf{y}_{i,t} = \mathbf{d}_i^{(t)} + \sum_{j=1}^{n} \sum_{k=0}^{t-1} \widetilde{\psi}_{k,ij} \varepsilon_{j,t-k} \tag{4}$$

We can then simulate counterfactual paths by setting the sequence of an arbitrary shock to zero or, equivalently, subtracting the contribution of this shock. Suppressing shock $s \in j = 1, .., n$ over all periods therefore results in the counterfactual path

$$\mathbf{y}_{i,t}^{cf} = \mathbf{d}_i^{(t)} + \sum_{j=1, j \neq s}^{n} \sum_{k=0}^{t-1} \widetilde{\psi}_{k,ij} \varepsilon_{j,t-k} \tag{5}$$

We construct counterfactual paths for inflation and GDP growth by separately suppressing (1) both the global demand and oil supply shock, (2) the monetary policy shock, (3) both the domestic demand and the domestic supply shock, and (4) the residual shock.

In order to save space, we discuss the counterfactuals for inflation only, see Figures (3) to (8), which depict the simulated paths for inflation in hypothetical scenarios in which the aggregate global shock, the monetary policy shock, the aggregate domestic shock as well as the residual rate shock are suppressed.[16] The red solid paths correspond to the median counterfactual path over all samples, while the shaded areas enclose the 16^{th} and 84^{th} (red-shaded) percentiles as well as the 5^{th} and 95^{th} (grey-shaded) percentiles, respectively.

Take Korea as an example. Three findings are noteworthy. First, suppressing either the global shocks or the domestic shocks does make a difference. Our results suggest that in the absence of global shocks, we would have observed lower inflation rates between 2004 and 2006, saying that global shocks had inflationary pressure during that period. Global shocks also put inflationary pressure on Korea immediately before the onset of the Great Recession. We can see this from the fact that without global shocks, the inflation rate would have been significantly lower during this period. This is not surprising and reflects the boom phase immediately before the crisis. Note that this result is qualitatively very consistent with the results from Bobeica and Jarociński (2019) for the euro area as well as for the United States.

Second, we find that we would have observed lower inflation rates in the absence of domestic shocks around 2008. Thus, domestic shocks were inflationary during this period. As in the case of global shocks, we also find that adverse shocks were deflationary between 2014 and 2016. Note that this finding perfectly mirrors our result from the historical decomposition, where we saw that especially around the

only a subsequence of the shocks. In this case, the counterfactual has to be simulated as described above.

[16] The counterfactual paths for GDP growth are available in the appendix.

Figure 3: Counterfactual paths for inflation with suppressed shocks – Indonesia

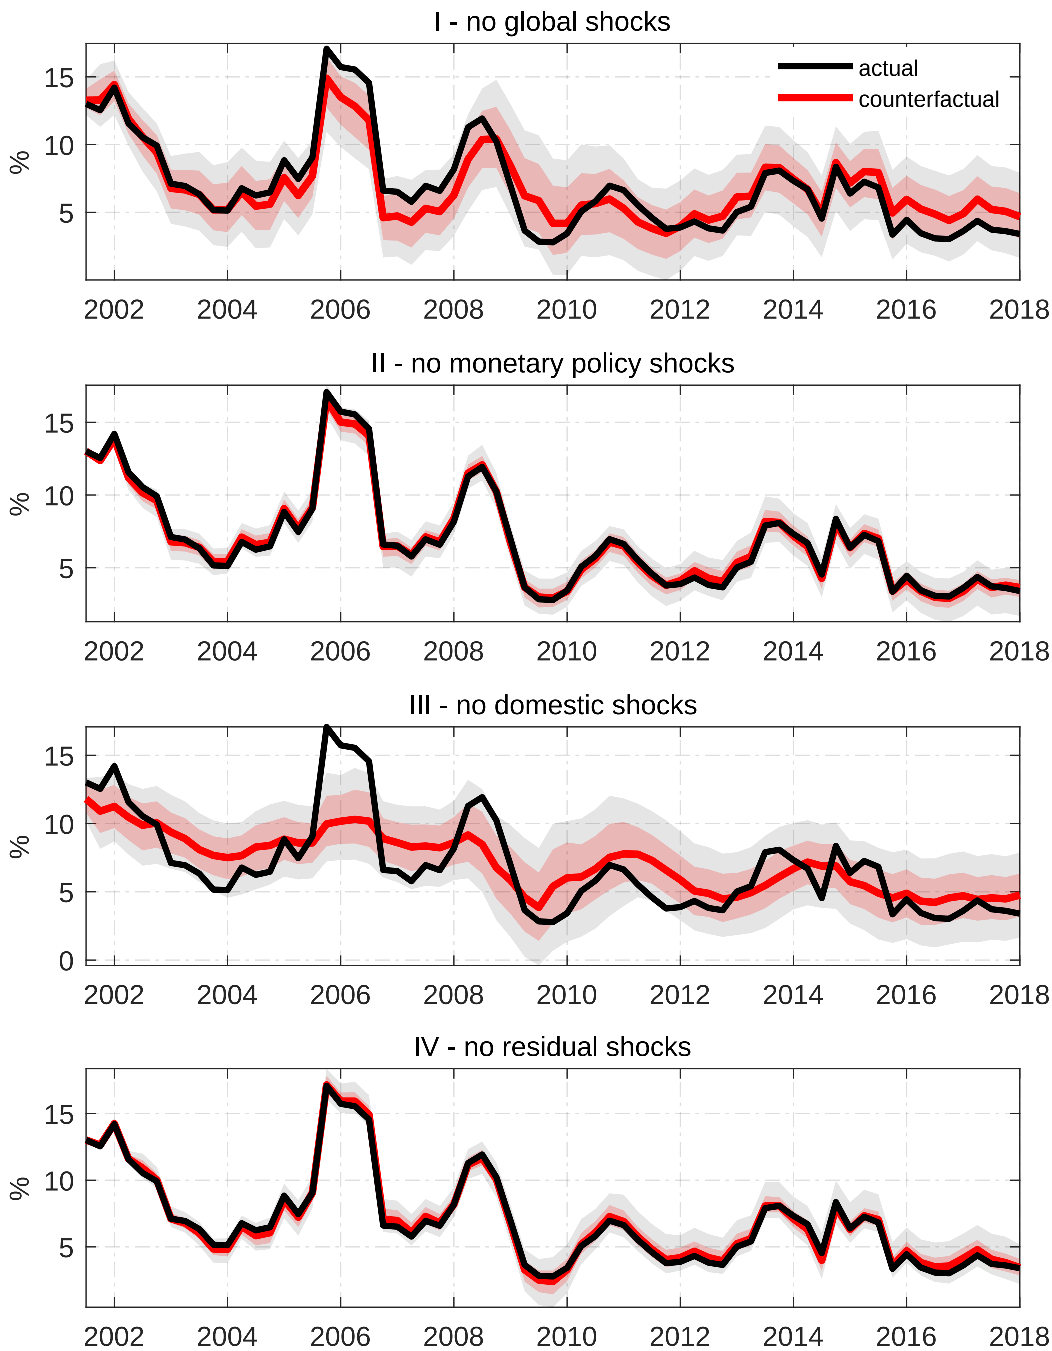

Notes: Median counterfactual paths (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area) as well as the 5^{th} and 95^{th} percentiles (grey-shaded area) for inflation (in %). In I, the counterfactual path corresponds to the inflation rate where aggregate global shocks (oil supply shocks and global demand shocks) are suppressed, while the same is done in II with monetary policy shocks, in III with domestic shocks (domestic demand and domestic supply shocks) and in IV (residual shocks). The identification of shocks relies on the Corsetti et al. (2014) identification.

Figure 4: Counterfactual paths for inflation with suppressed shocks – Korea

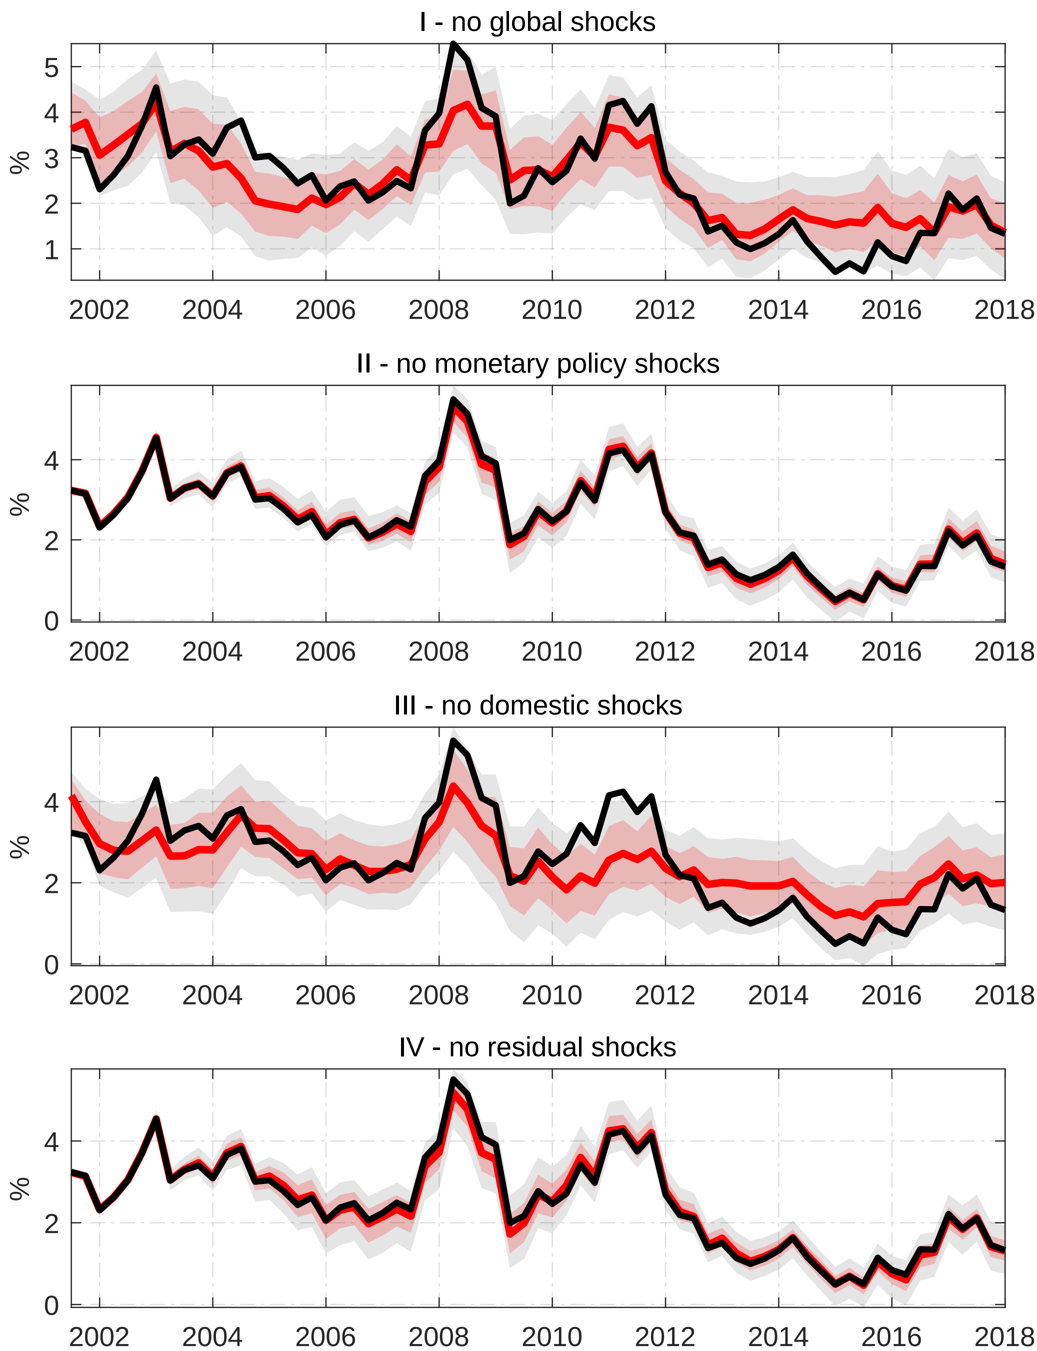

Notes: Median counterfactual paths (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area) as well as the 5^{th} and 95^{th} percentiles (grey-shaded area) for inflation (in %). In I, the counterfactual path corresponds to the inflation rate where aggregate global shocks (oil supply shocks and global demand shocks) are suppressed, while the same is done in II with monetary policy shocks, in III with domestic shocks (domestic demand and domestic supply shocks) and in IV (residual shocks). The identification of shocks relies on the Corsetti et al. (2014) identification.

Figure 5: Counterfactual paths for inflation with suppressed shocks – Malaysia

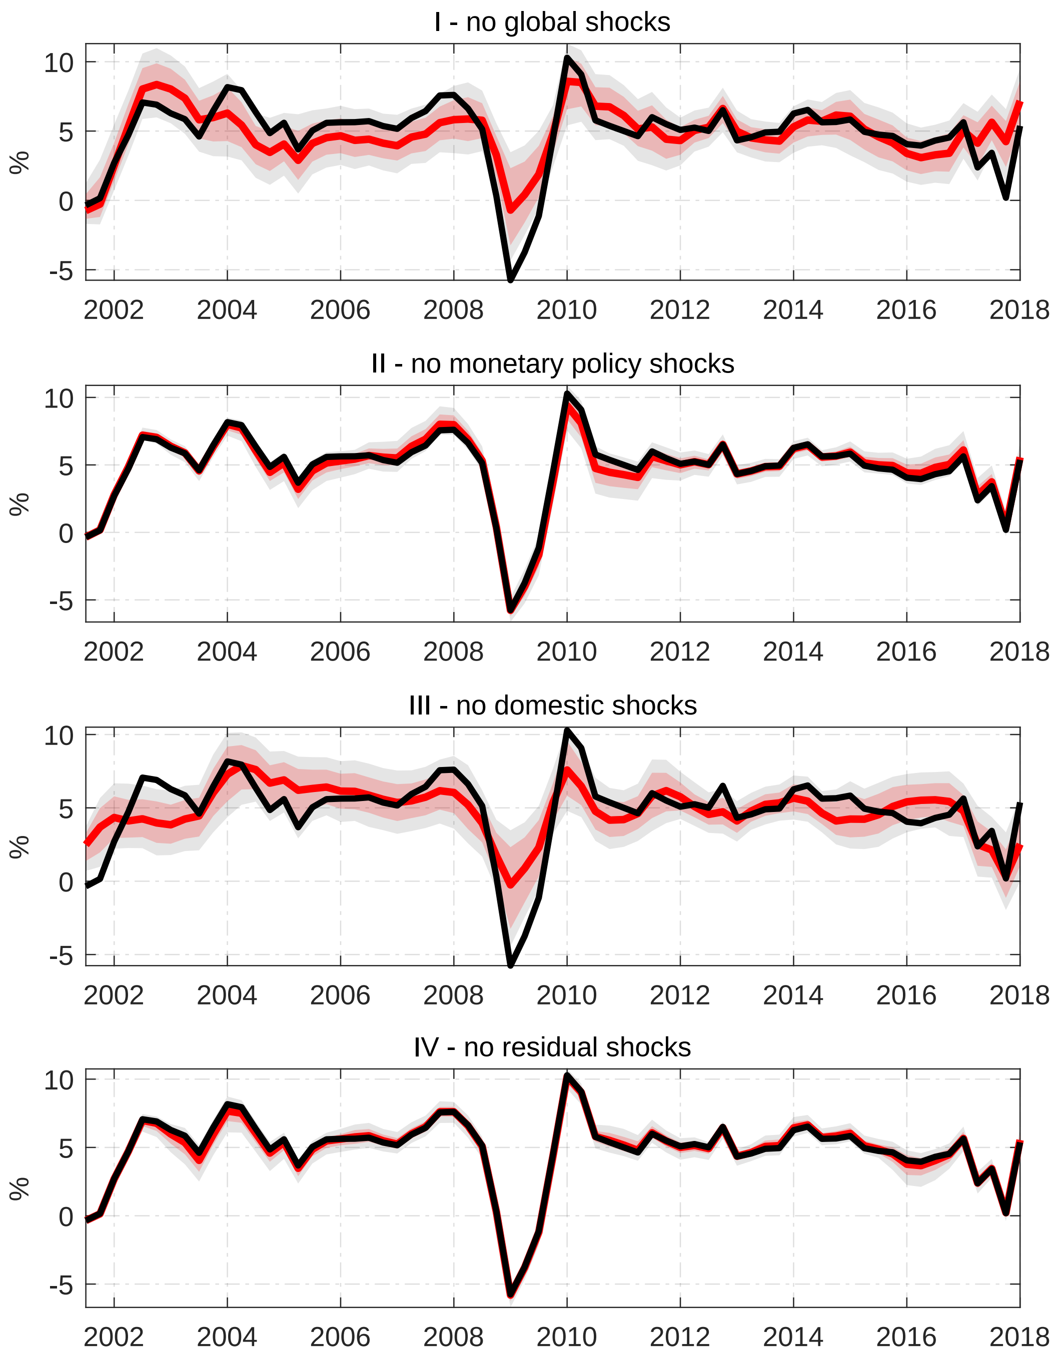

Notes: Median counterfactual paths (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area) as well as the 5^{th} and 95^{th} percentiles (grey-shaded area) for inflation (in %). In I, the counterfactual path corresponds to the inflation rate where aggregate global shocks (oil supply shocks and global demand shocks) are suppressed, while the same is done in II with monetary policy shocks, in III with domestic shocks (domestic demand and domestic supply shocks) and in IV (residual shocks). The identification of shocks relies on the Corsetti et al. (2014) identification.

Figure 6: Counterfactual paths for inflation with suppressed shocks – Philippines

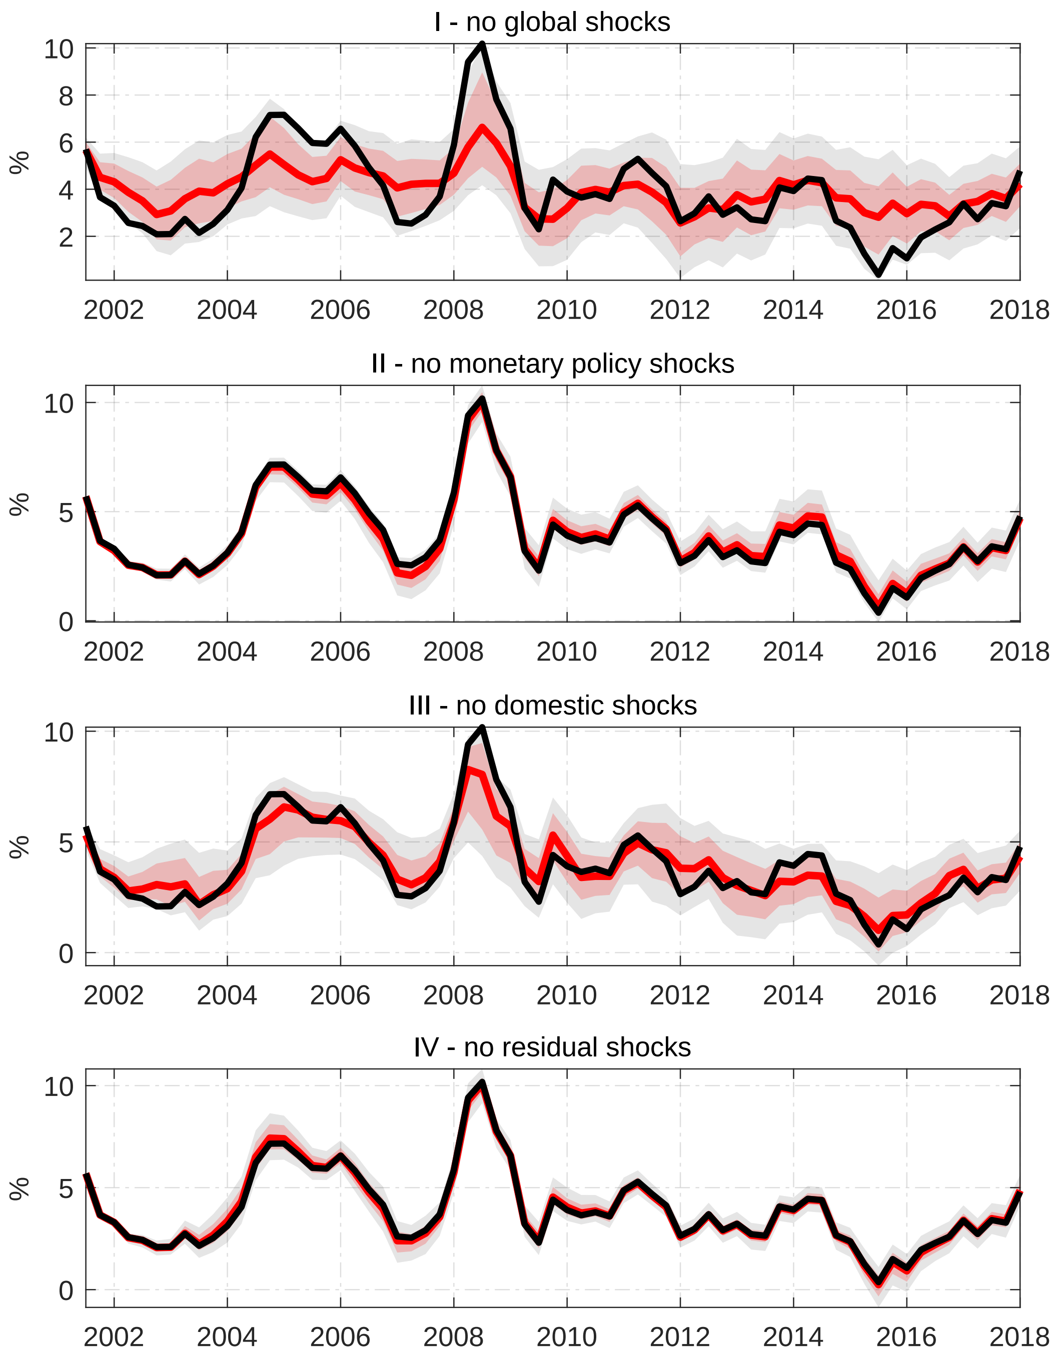

Notes: Median counterfactual paths (red solid path) with 16th and 84th percentiles (red-shaded area) as well as the 5th and 95th percentiles (grey-shaded area) for inflation (in %). In I, the counterfactual path corresponds to the inflation rate where aggregate global shocks (oil supply shocks and global demand shocks) are suppressed, while the same is done in II with monetary policy shocks, in III with domestic shocks (domestic demand and domestic supply shocks) and in IV (residual shocks). The identification of shocks relies on the Corsetti et al. (2014) identification.

Figure 7: Counterfactual paths for inflation with suppressed shocks – Singapore

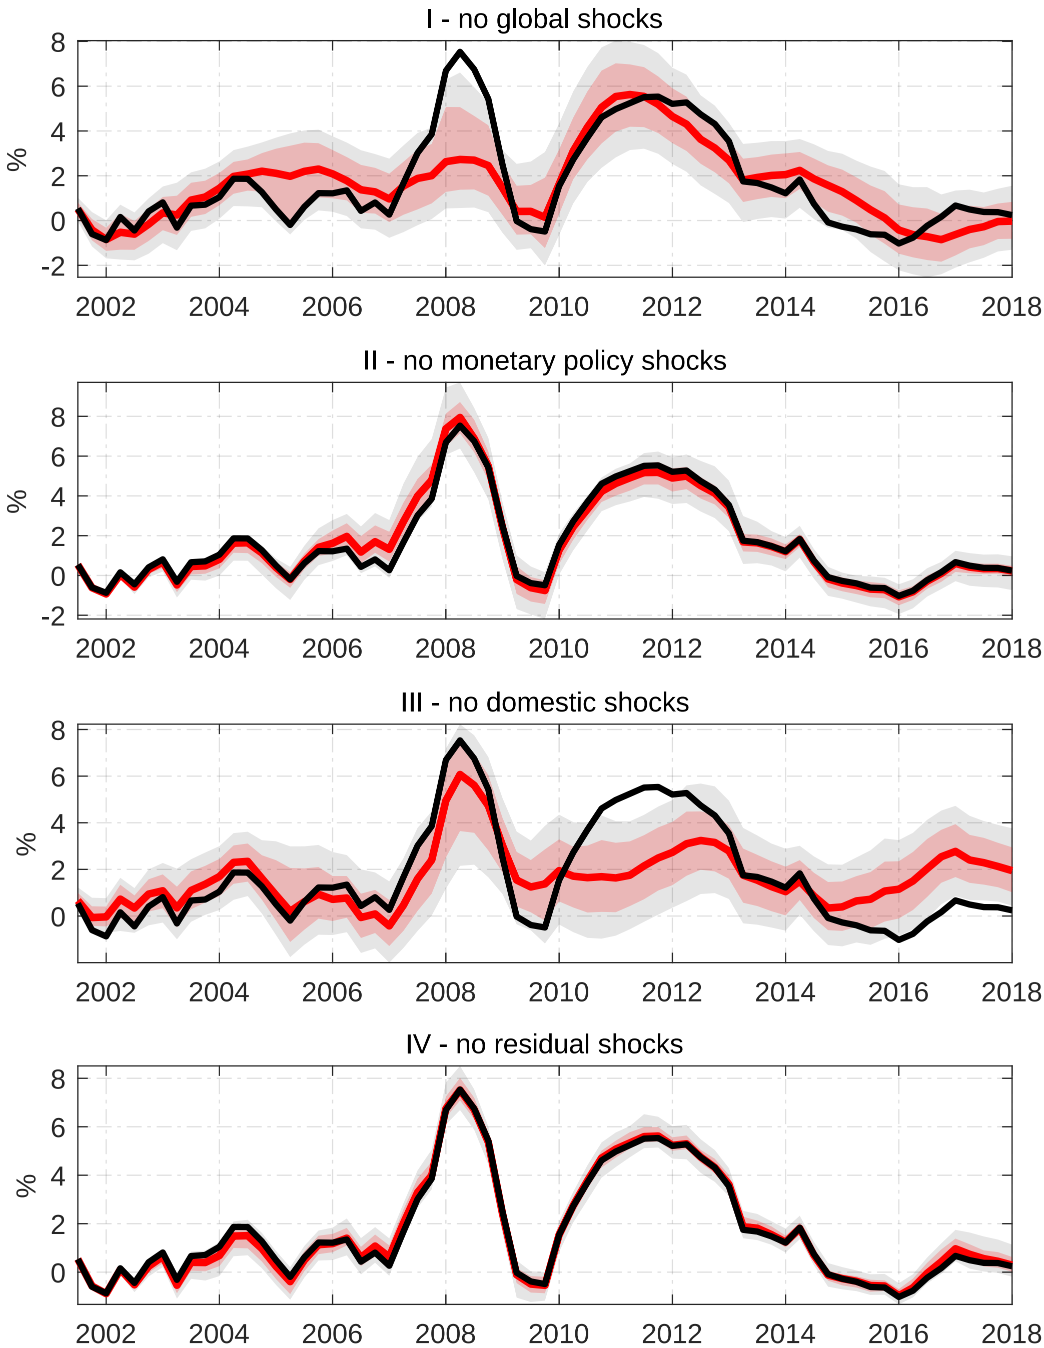

Notes: Median counterfactual paths (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area) as well as the 5^{th} and 95^{th} percentiles (grey-shaded area) for inflation (in %). In I, the counterfactual path corresponds to the inflation rate where aggregate global shocks (oil supply shocks and global demand shocks) are suppressed, while the same is done in II with monetary policy shocks, in III with domestic shocks (domestic demand and domestic supply shocks) and in IV (residual shocks). The identification of shocks relies on the Corsetti et al. (2014) identification.

Figure 8: Counterfactual paths for inflation with suppressed shocks – Thailand

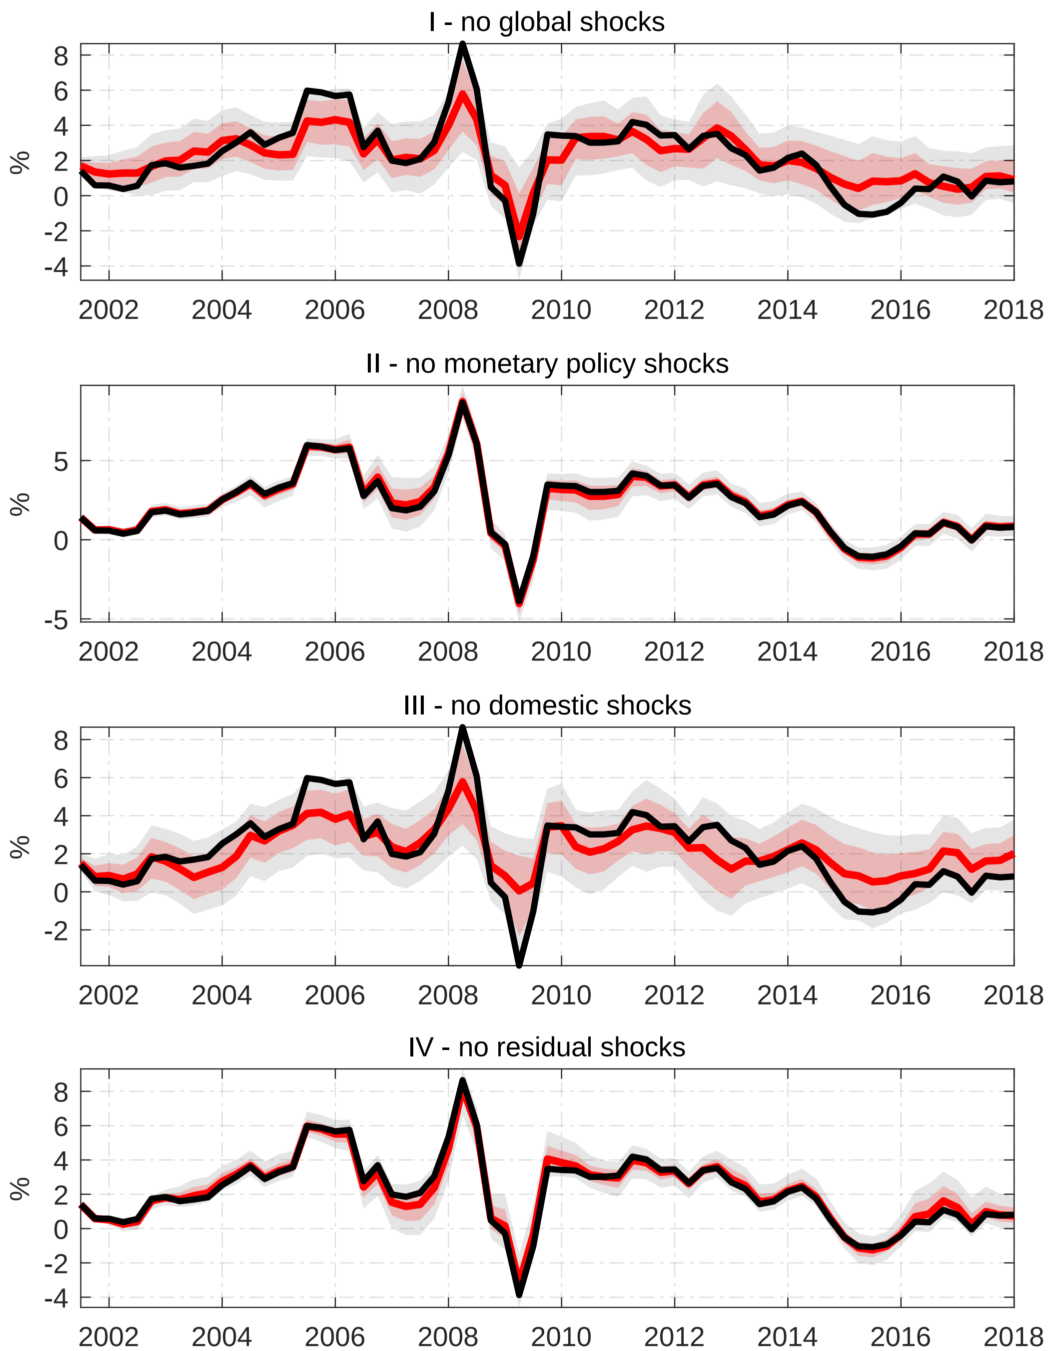

Notes: Median counterfactual paths (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area) as well as the 5^{th} and 95^{th} percentiles (grey-shaded area) for inflation (in %). In I, the counterfactual path corresponds to the inflation rate where aggregate global shocks (oil supply shocks and global demand shocks) are suppressed, while the same is done in II with monetary policy shocks, in III with domestic shocks (domestic demand and domestic supply shocks) and in IV (residual shocks). The identification of shocks relies on the Corsetti et al. (2014) identification.

Great Recession global and domestic shocks pushed inflation in the same direction.

The third finding pertains to the remaining shocks. These shocks seem to play a small role for Korean inflation as the actual inflation rate is indistinguishable from the counterfactual path for which monetary shocks (panel II) or residual shocks (panel IV) are suppressed.

Summarizing the results for Korean inflation, we conclude that global shocks as well as domestic demand seemed to play a more important role than monetary policy shocks and residual shocks. Importantly, the findings from the counterfactuals and the role of domestic and global shocks around 2008 are consistent for all countries.

D. *Robustness Check*

So far, our results suggest that monetary policy shocks have little effects on the dynamics of real activity. Therefore, we now ask whether the same is true for the systematic part of monetary policy. In order to do so, we follow, among others, Gordon and Leeper (1994) and Leeper and Zha (2003) who base the specification of monetary policy behavior on the information available to the central bank within the quarter. Recall that under our benchmark identification strategy, both structural demand and supply disturbances have simultaneous effects on the interest-rate equation. In order to impose an alternative systematic monetary policy behavior, we therefore restrict the corresponding coefficients in the structural matrix $\mathbf{A}_0$ to zero such that both demand and supply shocks do not have a contemporaneous impact on the short-term interest rate. The results (not presented) are qualitatively and quantitatively for all countries very much the same as in the benchmark case. This also confirms our previous results, i.e. that departures from the policy rule have only limited effects on inflation dynamics.

In our baseline model, we include the oil price as an endogenous variable. From the perspective of small open economies, however, the oil price is often considered exogenous such that it may not be explained by the lags of the endogenous domestic variables. We treat the oil price as endogenous as we need the oil price in order to identify global supply shocks in the VAR model. The estimated coefficients on the lagged endogenous variables in the oil price equation will be zero if oil is indeed exogenous, thus mitigating concerns of a potential misspecification.

In order to rule out remaining concerns, we re-estimate our models with the restriction that domestic variables have no effect on the oil price by means of block exogeneity. We set a prior of zero on the coefficients of all lagged domestic variables in the oil price equation in combination with a prior variance of $\lambda_5 = 0.001$.[17] This ensures that the oil price is not driven by domestic variables and that it nevertheless

[17]Note that block exogeneity is not possible with a Normal Wishart prior. The reason is that the Kronecker structure for the variance of the prior for the VAR coeffiients causes instability. Instead, we choose a Minnesota prior for this comparison, which is closest to the Normal Wishart prior and with which block exogeneity is possible. With the Minnesota prior, the results remain exactly unchanged if we compare the model with and without the assumption of block exogeneity.

can be exploited to identify a global supply shock. For all six countries, we find that the posterior distribution is very much centered around zero, i.e. the median, the 5^{th} and 95^{th} percentiles for all domestic variables are equal to 0.00. As regards our main results, we find no qualitative differences compared to our original approach. In the appendix, we can see that in this case we get the same oil supply shocks as in our baseline case. The corresponding impulse responses, not shown for space reasons but available on request, also look very similar.

IV How Global Shocks Affect the Phillips Curve

Much of the discussion about the changing nature of inflation is framed in terms of the Phillips curve relation between inflation and real activity. It is often argued that the process of inflation determination changed. Not only advanced economies, but also many emerging market economies have experienced declines in inflation that were lower than expected. A flattening of the Phillips curve could have important consequences for monetary policy as disinflation policy becomes more costly in terms of foregone economic activity.

In this section, we estimate the inflation-output nexus for our six countries under investigation and see if the Phillips curve is still 'alive'. We investigate whether the sensitivity of inflation to output depends on the nature of shocks that are driving GDP. In contrast to much of the literature, we do not add additional variables to the Phillips curve such as oil prices or measures of global output gaps in order to assess these variables' effects on the slope of the output-inflation trade-off. Instead, we decompose the observed series of output growth into components attributable to domestic and global shocks, respectively. Thus, we can show whether global and domestic factors affect the Phillips curve correlation equally.

It is important to stress that we do not offer a systematic analysis of structural shifts in the Phillips curve relationship. Our sample is too short relative to the number of VAR parameters to be estimated in order to accomplish that. Shifts in the slope of the Phillips curve could result from changes in the price-setting power of firms, changes to the conduct of monetary policy and changes in the nature of shocks that hit the economy. Occhino (2019) and Jacob and van Florenstein Mulder (2019) conduct system studies on the role of each of these factors. In contrast to these studies and in line with the first part of this paper, we focus on the global versus domestic nature of shocks driving GDP in order to decompose the inflation-output correlation. We treat the Phillips curve as an empirical regularity that describes a positive correlation between inflation and real economic activity, not as a structural relationship.

We exploit our counterfactual paths which are functions of structural (past) shocks that we identify. This being said, our poor-man's approach is easy to implement and has several advantages. First, we estimate the Phillips curve using a decomposition of GDP in terms of cumulative historical decompositions of structural shocks. That is, the model allows different domestic and global components to enter the Phillips curve with different coefficients and potentially different signs. The model therefore nests the conventional Phillips curve specification if the coefficient on the global component equals the one on the domestic component. Second, by decomposing economic growth into domestic and global components, we can avoid an econometric problem faced by studies which extend the Phillips curve by measures of global slack. Global output gaps are typically highly correlated with the domestic output gap (see Jašová et al., 2020). That is, the studies have difficulties separating the true effects from domestic and global forces. Instead, our decomposition is based on functions of orthogonal structural shocks.

Note that up to now, we bundled the contributions of oil supply shocks and global demand shocks which we referred to as the contributions of global shocks. That is, we ignore whether the contribution of oil supply shocks and global demand shocks can have different signs. We now account for this possibility by splitting up the global component into its parts, i.e. the parts that stem from oil supply shocks and global demand shocks.

We start by estimating a reduced-form Phillips curve with observable GDP growth as the determinant of inflation

$$\pi_t = c + \beta x_t + \gamma \pi_{t-1} + \varepsilon_t, \qquad \text{(model A)}$$

where π_t is the year-on-year inflation rate and x_t is the observed growth rate of real GDP. We add past inflation as a proxy for today's expectations of future inflation. Table (4) reports the results for this specification. We find a significantly positive slope of the Phillips curve for Korea, Malaysia, Singapore and Thailand. For Indonesia and the Philippines, the estimated slope coefficient is insignificantly different from zero.

Table 4: Regression results for model A

	IDN	KOR	MYS	PHL	SGP	THA
c	1.15	−0.09	−0.06	0.34	−0.58	0.02
	[0.68]	[−0.43]	[−0.16]	[0.54]	[−3.94]	[0.08]
β	−0.01	0.12	0.16	0.05	0.13	0.12
	[−0.05]	[3.52]	[2.92]	[0.61]	[7.03]	[2.45]
γ	0.83	0.83	0.71	0.84	0.95	0.75
	[11.92]	[14.36]	[7.85]	[11.49]	[24.52]	[9.98]
# obs	67	67	67	67	67	67
R^2	0.70	0.76	0.48	0.69	0.90	0.63

Notes: Estimation results for the specification with observable GDP growth. The growth rate of GDP enters with the coefficient β, while γ is the coefficient on lagged inflation. The values in square brackets correspond to the t−statistic testing the null that the corresponding coefficient is equal to zero.

We now include the counterfactual path of GDP growth. Consequently, the slope coefficient β is split into the coefficient on domestically-driven GDP and the coefficient on globally determined GDP. The baseline regression reads

$$\pi_t = c + \beta_j x_t^j + \beta_{cf(j)} x_t^{cf(j)} + \gamma \pi_{t-1} + \varepsilon_t, \qquad \text{(model B)}$$

where $x_t^{cf(j)}$ is the counterfactual path of the growth rate of domestic real GDP in which we suppress shock j. Hence, x_t^j for $j = \{oil, dem\}$ denotes the contribution of either oil supply shocks or global demand shocks to the growth rate of real GDP. Technically, this contribution corresponds to the distance between the actual data and the counterfactual path in our simulation exercise where we simulated global shocks away, see Figures (3) to (8). For the case where $\beta_{cf(j)} = \beta_j$, the distinction between domestic and global components of activity becomes obsolete. Hence, the model nests the conventional specification which regresses inflation on observable output growth.

The Phillips curve regression as explained before suffers from a generated regressors problem.[18] The reason is that the regressors x_t^j and $x_t^{cf(j)}$ are themselves results from a regression model. Hence, the estimation uncertainty surrounding these regressors must not be ignored. We address this problem as follows: instead of running a single regression using the median from the historical decomposition, we estimate the Phillips curve for each single draw we have available from our VAR estimation. By doing so, we carry uncertainty from the first stage, i.e. the VAR estimation, into the second state, i.e. the Phillips curve regression. For all coefficients, inference is based on the 16^{th} and 84^{th} percentiles of the $t-$statistics from this exercise.

Our estimation results are reported in Table (5). The upper half of the table focuses on the role of oil supply shocks, i.e. global supply shocks. For Korea, Malaysia, Singapore and Thailand, the domestic growth component, i.e. GDP growth in the absence of global supply shocks, enters with a positive coefficient. In these cases, the Phillips curve trade-off remains valid and has the expected sign. An increase in the domestic part of growth is inflationary. This stands in contrast to the inflation-impact of growth driven by global supply shocks. For four countries, GDP driven by oil supply shocks enters with a negative sign, such that the inflation-output correlation is lower once the component driven by global supply shock is included in GDP. This effect, however, remains insignificant for all countries.

Turning to global demand shocks, our results suggest that global demand shocks steepen the Phillips curve for Korea, Malaysia, Singapore and Thailand. For these countries, the estimated β_{dem} is significantly positive. For all countries, the slope coefficient on growth driven by domestic demand shocks is higher than the slope coefficient on the counterfactual growth series suppressing global demand shocks. Put differently, ignoring the fraction of GDP growth driven by global demand shocks

[18] We are grateful to an anonymous referee for this point as well as the proposed solution.

Table 5: Regression results for model B

	IDN	KOR	MYS	PHL	SGP	THA
oil supply shock						
c	1.03	−0.09	−0.05	0.26	−0.60	−0.01
	[0.03, 0.85]	[−0.93, 0.06]	[−0.31, 0.12]	[−0.81, 0.99]	[−5.08, −3.44]	[−1.36, 0.77]
β_{oil}	−0.64	−0.02	0.14	−0.27	0.07	−0.18
	[−2.21, 0.31]	[−2.19, 1.41]	[−0.90, 2.47]	[−4.64, 0.34]	[0.18, 3.42]	[−3.52, 0.13]
$\beta_{cf(oil)}$	0.01	0.13	0.16	0.11	0.13	0.14
	[−0.18, 0.62]	[3.36, 4.71]	[2.35, 3.48]	[0.69, 3.23]	[6.74, 8.48]	[2.33, 4.58]
γ	0.83	0.82	0.70	0.80	0.94	0.72
	[11.04, 12.29]	[12.26, 14.59]	[6.59, 7.94]	[9.14, 12.54]	[20.80, 26.10]	[8.69, 10.76]
# obs	67	67	67	67	67	67
R^2	0.71	0.77	0.49	0.73	0.91	0.67
global demand shock						
c	3.37	−0.02	0.11	0.53	−0.49	0.18
	[0.90, 3.09]	[−0.48, 0.73]	[−0.18, 1.10]	[0.53, 1.87]	[−4.22, −2.86]	[−0.05, 2.37]
β_{dem}	0.46	0.17	0.30	0.26	0.25	0.47
	[0.43, 2.76]	[2.17, 4.18]	[1.97, 4.27]	[0.47, 3.40]	[4.03, 9.75]	[1.93, 6.39]
$\beta_{cf(dem)}$	−0.35	0.11	0.13	0.03	0.11	0.11
	[−2.56, −0.29]	[1.60, 3.45]	[1.04, 2.87]	[−0.76, 0.62]	[5.62, 7.84]	[0.98, 2.85]
γ	0.79	0.83	0.69	0.83	0.94	0.71
	[9.23, 11.94]	[13.22, 14.62]	[7.15, 8.08]	[10.32, 11.71]	[24.01, 29.56]	[8.04, 10.57]
# obs	67	67	67	67	67	67
R^2	0.72	0.77	0.51	0.70	0.92	0.67

Notes: The values for the estimated coefficients (as well as for the coefficient of determination R^2) correspond to their median value over all 2000 draws. In the first block, β_{oil} represents the coefficient on the component of real GDP growth driven by oil supply shocks, while in the second block β_{dem} represents the coefficient on the component of real GDP growth driven by global demand shocks. In both blocks, $\beta_{cf(j)}$ for $j = \{oil, dem\}$ represents the coefficient on the counterfactual path of real GDP growth that excludes either the component driven by oil supply shocks (first block), or the component driven by global demand shocks (second block). The lagged domestic inflation rate enters with the coefficient γ. For all coefficients in both blocks, the values in square brackets correspond to the 16^{th} and 84^{th} percentile of the t−statistic testing the null that the corresponding coefficient is equal to zero. All results rely on the Corsetti et al. (2014) identification.

leads to a flatter Phillips curve. Consequently, the coefficient on the counterfactual in the lower half of Table (5) is smaller than the coefficient on observable growth reported in Table (4).

As a third specification, we include both global components of GDP separately, i.e. x_t^{oil} and x_t^{dem}. The counterfactual $x_t^{cf(glo)}$ reflects GDP growth in the absence of both types of global shocks

$$\pi_t = c + \beta_{oil}x_t^{oil} + \beta_{dem}x_t^{dem} + \beta_{cf(glo)}x_t^{cf(glo)} + \gamma\pi_{t-1} + \varepsilon_t. \quad \text{(model C)}$$

Table (6) reports the estimated coefficients. Again, the fraction of GDP driven by oil supply shocks remains an insignificant driver of inflation. The occurrence of global demand shocks, in contrast, contributes to a steeper inflation-output nexus for Korea, Malaysia, Singapore and Thailand.

Table 6: Regression results for model C

	IDN	KOR	MYS	PHL	SGP	THA
c	3.04	−0.03	0.13	0.46	−0.51	0.13
	[0.42, 3.23]	[−0.94, 1.00]	[−0.31, 1.20]	[−0.49, 2.28]	[−4.89, −2.33]	[−0.99, 2.61]
β_{oil}	−0.91	−0.04	0.09	−0.34	0.05	−0.19
	[−2.71, 0.14]	[−2.55, 1.32]	[−1.49, 2.33]	[−5.40, −0.55]	[−0.81, 3.53]	[−4.02, 0.16]
β_{dem}	0.50	0.18	0.31	0.36	0.26	0.48
	[0.35, 3.02]	[2.2, 4.73]	[1.87, 4.62]	[0.62, 4.66]	[4.15, 10.98]	[1.99, 7.19]
$\beta_{cf(glo)}$	−0.29	0.12	0.13	0.08	0.12	0.13
	[−2.62, 0.18]	[1.78, 4.08]	[1.00, 3.04]	[−0.08, 2.79]	[5.53, 8.68]	[1.26, 4.46]
γ	0.78	0.80	0.67	0.76	0.92	0.66
	[8.59, 11.94]	[11.46, 14.75]	[6.501, 8.12]	[8.17, 12.84]	[21.07, 30.24]	[7.38, 11.09]
# obs	67	67	67	67	67	67
R^2	0.73	0.78	0.54	0.76	0.93	0.73

Notes: The values for the estimated coefficients (as well as for the coefficient of determination R^2) correspond to their median value over all 2000 draws. β_{oil} represents the coefficient on the component of real GDP growth driven by oil supply shocks, while β_{dem} represents the coefficient on the component of real GDP growth driven by global demand shocks. $\beta_{cf(glo)}$ represents the coefficient on the counterfactual path of real GDP growth that excludes the component driven by both global supply shocks and global demand shocks. The lagged domestic inflation rate enters with the coefficient γ. For all coefficients, the values in square brackets correspond to the 16^{th} and 84^{th} percentile of the t−statistic testing the null that the corresponding coefficient is equal to zero. All results rely on the Corsetti et al. (2014) identification strategy.

We can conclude that the effect of global driving forces on the Phillips curve trade-off critically depends on the nature of these forces. While global factors in terms of oil supply shocks do not affect the Phillips curve, global demand shocks steepen the inflation-output correlation.[19]

V Conclusions

This paper adds to the discussion about the changing nature of inflation dynamics in six Asian emerging market economies. We estimate a series of VAR models, in which we identify a battery of demand and supply shocks shocks using sign restrictions. Focusing on the co-movement between domestic and global variables, our identification strategy also allows us to distinguish between global and domestic shocks. Relying on forecast error variance decompositions and historical decompositions, we find that (1)

[19]Our results are qualitatively very similar when using the counterfactuals from the alternative identification strategy. These regression results are available in the appendix.

global factors play an important role for both inflation and the growth rate of real GDP across all countries under consideration and (2) the role of monetary policy is limited. While global factors can explain the sharp increases and the subsequent plunges around the Great Recession, they also contribute much to the low inflation rates that have been recently observed. Since global factors are driving a substantive share of inflation, domestic monetary policy is increasingly less able to stabilize inflation and the real economy.

We also revisit the Phillips curve relation between inflation and real activity. This is particularly important for policymakers as monetary policy in the short-run induces movements along the Phillips curve by controlling domestic demand. By decomposing the observed growth rates of domestic real GDP into components attributable to domestic and global shocks, we investigate whether the sensitivity of inflation to output depends on the nature of shocks driving the real economy.

Our results suggest that including the components of growth due to oil price shocks and global demand shocks, respectively, changes the inflation-output correlation. GDP growth due to oil supply shocks seem to flatten the Phillips curve in most countries, though the effect is not statistically significant. The contrary is true for the fraction of GDP due to global demand shocks. The Phillips curve correlation increases once we include the part of GDP driven by global demand. Hence, we show that global integration affects the Phillips curve and that the nature of global shocks determines whether the curve steepens. It should be noted that the focus of this paper is on the role of global vs domestic drives of inflation dynamics and their effects on the slope of the Phillips curve. In principle, we could use a similar framework to address the role of demand vs supply shocks to study the slope of the Phillips curve conditional on demand shocks only. We leave this for future research.

Our results highlight the difficulties facing inflation targeting central banks in the region. While monetary policy affects domestic demand, global demand, which drives the bulk of inflation, is not under the control of monetary policy.

References

Arias, Jonas E, Juan F Rubio-Ramírez, and Daniel F Waggoner, "Inference Based on Structural Vector Autoregressions Identified With Sign and Zero Restrictions: Theory and Applications," *Econometrica*, 2018, *86* (2), 685–720.

Auer, Raphael A and Aaron Mehrotra, "Trade Linkages and the Globalisation of Inflation in Asia and the Pacific," *Journal of International Money and Finance*, 2014, *49*, 129–151.

Baumeister, Christiane and Luca Benati, "Unconventional Monetary Policy and the Great Recession: Estimating the Macroeconomic Effects of a Spread Compression at the Zero Lower Bound," *International Journal of Central Banking*, 2013, *9* (2), 165–212.

Bems, Rudolfs, Francesca G Caselli, Francesco Grigoli, Bertrand Gruss, and Weicheng Lian, "Is Inflation Domestic or Global? Evidence from Emerging Markets," *IMF Working Paper No. 18/241*, 2018, *International Monetary Fund.*

Bianchi, Francesco and Andrea Civelli, "Globalization and Inflation: Evidence from a Time-Varying VAR," *Review of Economic Dynamics*, 2015, *18* (2), 406–433.

Bobeica, Elena and Marek Jarociński, "Missing Disinflation and Missing Inflation: A VAR Perspective," *International Journal of Central Banking*, 2019, *15* (1), 199–232.

Borio, Claudio and Andrew Filardo, "Globalisation and Inflation: New Cross-Country Evidence on the Global Determinants of Domestic Inflation," *BIS Working Paper No. 227*, 2007, *Bank for International Settlements.*

Ciccarelli, Matteo and Benoit Mojon, "Global Inflation," *The Review of Economics and Statistics*, 2010, *92* (3), 524–535.

Coibion, Olivier and Yuriy Gorodnichenko, "Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation," *American Economic Journal: Macroeconomics*, 2015, *7* (1), 197–232.

Conti, Antonio Maria, Stefano Neri, and Andrea Nobili, "Low Inflation and Monetary Policy in the Euro Area," *Working Paper No. 2005*, 2017, *European Central Bank.*

Corsetti, Giancarlo, Luca Dedola, and Sylvain Leduc, "The International Dimension of Productivity and Demand Shocks in the US Economy," *Journal of the European Economic Association*, 2014, *12* (1), 153–176.

Eickmeier, Sandra and Markus Kühnlenz, "China's Role in Global Inflation Dynamics," *Macroeconomic Dynamics*, 2018, *22* (2), 225–254.

Faccini, Renato, Stephen Millard, and Francesco Zanetti, "Wage Rigidities in an Estimated Dynamic, Stochastic, General Equilibrium Model of the UK Labour Market," *The Manchester School*, 2013, *81*, 66–99.

Forbes, Kristin J, "Has Globalization Changed the Inflation Process?," *BIS Working Paper No. 791*, 2019, *Bank for International Settlements.*

Förster, Marcel and Peter Tillmann, "Reconsidering the International Comovement of Inflation," *Open Economic Review*, 2014, *25* (5), 841–863.

Gerlach, Petra, "The Global Output Gap: Measurement Issues and Regional Disparities," *BIS Quarterly Review, June*, 2011.

Geweke, John, "Evaluating the Accuracy of Sampling-Based Approaches to the Calculations of Posterior Moments," *Bayesian Statistics*, 1992, *4*, 641–649.

Gordon, David B and Eric M Leeper, "The Dynamic Impacts of Monetary Policy: An Exercise in Tentative Identification," *Journal of Political Economy*, 1994, *102* (6), 1228–1247.

Ihrig, Jane, Steven B Kamin, Deborah Lindner, and Jaime Marquez, "Some Simple Tests of the Globalization and Inflation Hypothesis," *International Finance*, 2010, *13* (3), 343–375.

IMF, "Low Inflation in Asia: How Long Will it Last?," *Regional Economic Outlook Asia-Pacific*, 2018, *April 2018.*

Jacob, Punnoose and Thomas van Florenstein Mulder, "The Flattening of the Phillips Curve: Rounding up the Suspects," *Analytical Notes No. 2019/06*, 2019, *Reserve Bank of New Zealand.*

Jašová, Martina, Richhild Moessner, and Előd Takáts, "Domestic and Global Output Gaps as Inflation Drivers: What Does the Phillips Curve Tell?," *Economic Modelling*, 2020, *87*, 238–253.

Krause, Michael U and Thomas A Lubik, "The (Ir)Relevance of Real Wage Rigidity in the New Keynesian Model with Search Frictions," *Journal of Monetary Economics*, 2007, *54* (3), 706–727.

Leeper, Eric M and Tao Zha, "Modest Policy Interventions," *Journal of Monetary Economics*, 2003, *50* (8), 1673–1700.

Milani, Fabio, "Does Global Slack Matter more than Domestic Slack in Determining US Inflation?," *Economics Letters*, 2009, *102* (3), 147–151.

— , "Global Slack and Domestic Inflation Rates: A Structural Investigation for G-7 Countries," *Journal of Macroeconomics*, 2010, *32* (4), 968–981.

Mumtaz, Haroon and Paolo Surico, "Evolving International Inflation Dynamics: World and Country-Specific Factors," *Journal of the European Economic Association*, 2012, *10* (4), 716–734.

Neely, Christopher J and David E Rapach, "International Comovements in Inflation Rates and Country Characteristics," *Journal of International Money and Finance*, 2011, *30* (7), 1471–1490.

Occhino, Filippo, "The Flattening of the Phillips Curve: Policy Implications Depend on the Cause," *Economic Commentary No. 2019-11*, 2019, *Federal Reserve Bank of Cleveland.*

Okuday, Tatsushi, Tomohiro Tsurugaz, and Francesco Zanetti, "Imperfect Information, Shock Heterogeneity, and Inflation Dynamics," *IMES Discussion Paper Series No. 19-E-15*, 2019, *Institute for Monetary and Economic Studies, Bank of Japan.*

Parker, Miles, "How Global is "Global Inflation"?," *Journal of Macroeconomics*, 2018, *58*, 174–197.

Primiceri, Giorgio E, "Time Varying Structural Vector Autoregressions and Monetary Policy," *The Review of Economic Studies*, 2005, *72* (3), 821–852.

Raftery, Adrian E and Steven Lewis, "How Many Iterations in the Gibbs Sampler?," *Bayesian Statistics*, 1992, *4*, 115–130.

Sims, Christopher A, "Macroeconomics and Reality," *Econometrica*, 1980, *48* (1), 1–48.

— **and Tao Zha**, "Were There Regime Switches in US Monetary Policy?," *American Economic Review*, 2006, *96* (1), 54–81.

Tanaka, Misa and Chris Young, "The Eonomics of Global Output Gap Measures," *Bank of England Quarterly Bulletin*, 2008, p. Q3.

Thomas, Carlos and Francesco Zanetti, "Labor Market Reform and Price Stability: An Application to the Euro Area," *Journal of Monetary Economics*, 2009, *56* (6), 885–899.

Trigari, Antonella, "Equilibrium Unemployment, Job Flows, and Inflation Dynamics," *Journal of Money, Credit and Banking*, 2009, *41* (1), 1–33.

Volz, Ulrich, "On the Future of Inflation Targeting in East Asia," *Review of Development Economics*, 2015, *19* (3), 638–652.

Zanetti, Francesco, "Labor Market Institutions and Aggregate Fluctuations in a Search and Matching Model," *European Economic Review*, 2011, *55* (5), 644–658.

Appendix

A Convergence Checks

In order to assess the convergence of the sampler, we apply different metrics.

First, we calculate the inefficiency factors for convergence. These are the inverse of the relative numerical efficiency measure of Geweke (1992) and defined by $(1+2\sum_{k=1}^{\infty}\rho_k)$, where ρ_k is the k^{th} order autocorrelation for the underlying parameter. Inefficiency factors of around 20 or below are regarded as satisfactory. Table (7) reports the minimum, the maximum, the median and the average inefficiency factor for each country. For all countries, the inefficiency factors are far

Table 7: Distribution of inefficiency factors

	B draws					
	IDN	KOR	MYS	PHL	SGP	THA
median	3.02	2.95	2.98	2.93	3.04	2.96
mean	3.00	3.00	2.98	2.96	3.04	2.97
min	2.67	2.65	2.69	2.66	2.63	2.61
max	3.39	3.43	3.34	3.48	3.45	3.38
	Σ draws					
	IDN	KOR	MYS	PHL	SGP	THA
median	3.00	2.99	2.99	3.03	2.90	3.08
mean	3.01	3.00	2.95	3.04	2.90	3.07
min	2.85	2.72	2.74	2.72	2.69	2.82
max	3.32	3.21	3.22	2.27	3.10	3.33

Notes: The elements of **B** are the reduced-form coefficients, and the elements of **Σ** are the stacked entries in the covariance matrix.

below 20 for all parameters in **B** and in **Σ**. As a second check, we follow Raftery and Lewis (1992) and calculate the number of draws required to a achieve a given level of precision in our posterior samples. We use common values and set our quantiles of the marginal posteriors to 2.5% and 97.5%, the probability of attaining the required accuracy to 95% and the desired accuracy to 2.5%.[20] Table (8) shows that the required number of draws is always far below the number of iterations in our application.

As a final check, we take a look at the 1st and 5th order autocorrelation between draws for all parameters in **B** and **Σ**. We find that the sample autocorrelation is very low, ranging between -0.05 and 0.05 for all countries.

In sum, we find that all convergence checks are satisfactory.

[20]See, for instance, Raftery and Lewis (1992) and Primiceri (2005).

Table 8: Distribution of Raftery-Lewis statistics

	IDN	KOR	MYS	PHL	SGP	THA
B DRAWS						
MEDIAN	604	614	601	603	619	604
MEAN	615	617	607	607	618	609
MIN	570	570	560	570	570	570
MAX	712	712	706	672	700	739
$\boldsymbol{\Sigma}$ DRAWS						
	IDN	KOR	MYS	PHL	SGP	THA
MEDIAN	599	594	598	599	604	594
MEAN	608	605	608	613	605	594
MIN	570	570	570	570	560	570
MAX	682	686	672	720	650	665

Notes: The elements of **B** are the reduced-form coefficients, and the elements of $\boldsymbol{\Sigma}$ are the stacked entries in the covariance matrix.

B Further Results for the Baseline Specification

A. *Impulse Responses*

We use the BEAR toolbox to estimate our VAR models.[21] The algorithm to derive structural impulse responses by means of sign restrictions follows Arias et al. (2018). Figures (9) to (20) show, for each country, the impulse responses of the inflation rate and the growth rate of real GDP to all six shocks. The shock size is normalized to one standard deviation. Since the SVAR model is heavily restricted, we do not emphasize the interpretation of the responses too much. Note, however, that most responses are persistent, though the restrictions on the sign of the responses are imposed on impact only. Take Korea as an example. The responses of inflation, see Figure (11), to domestic and global demand shocks are quite persistent. Inflation returns to its mean three or four years after the shock. Hence, the sign restrictions do not seem to overly stretch the data.

[21] See https://www.ecb.europa.eu/pub/research/working-papers/html/bear-toolbox.en.html

Figure 9: Impulse responses of inflation - Indonesia

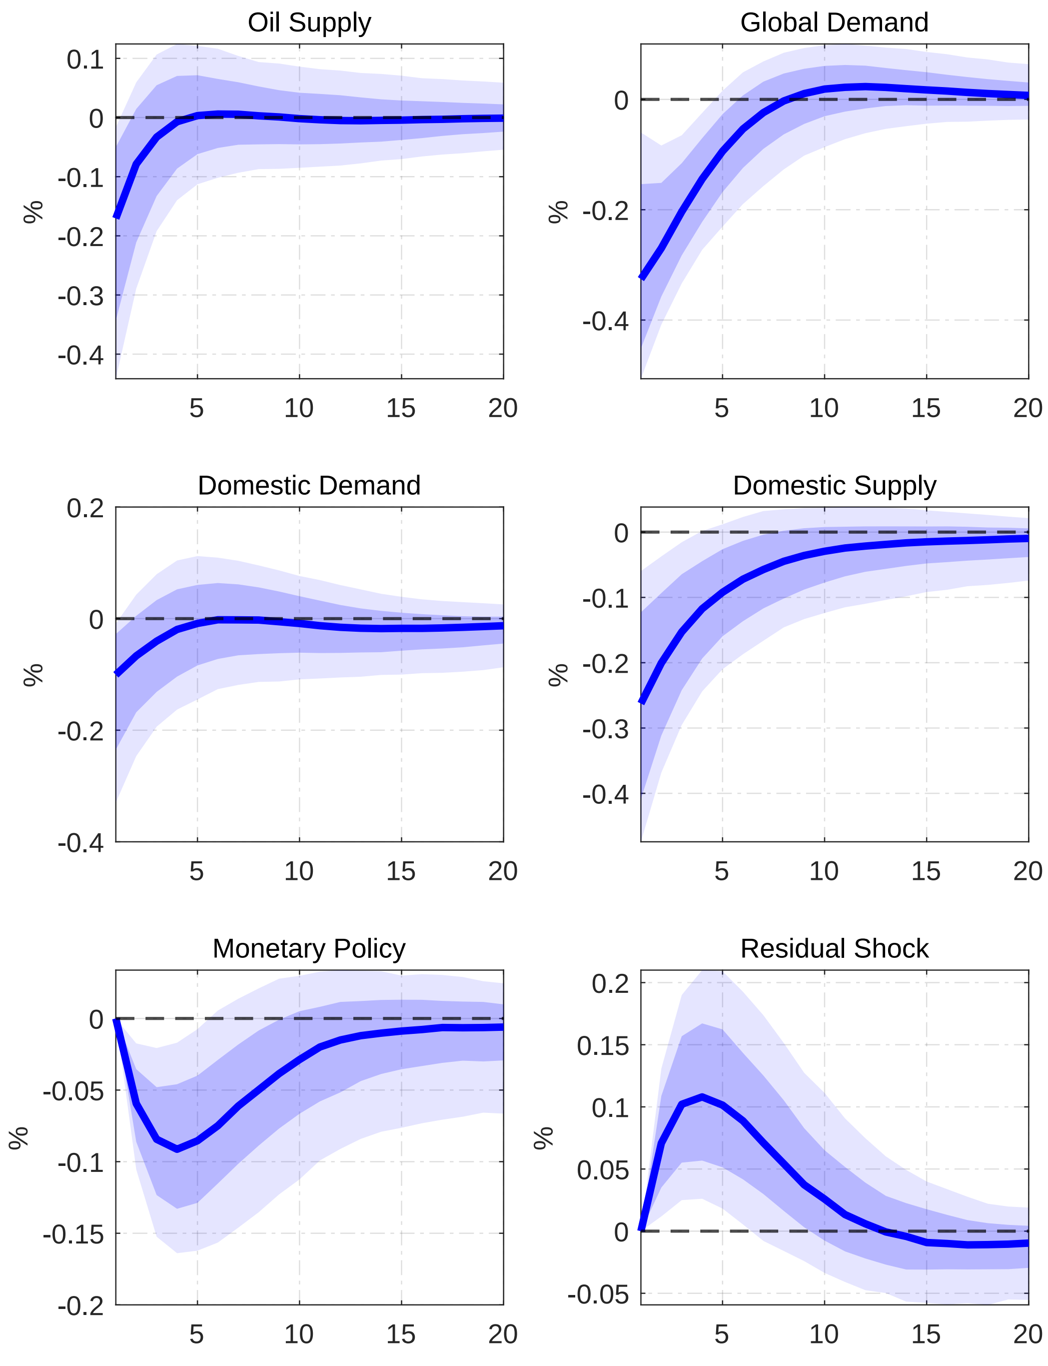

Notes: Median impulse response (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area) as well as 5^{th} and 95^{th} percentiles (grey-shaded area). The sign-restrictions are imposed on impact only. The identification of shocks relies on the Corsetti et al. (2014) identification.

Figure 10: Impulse responses of real GDP growth - Indonesia

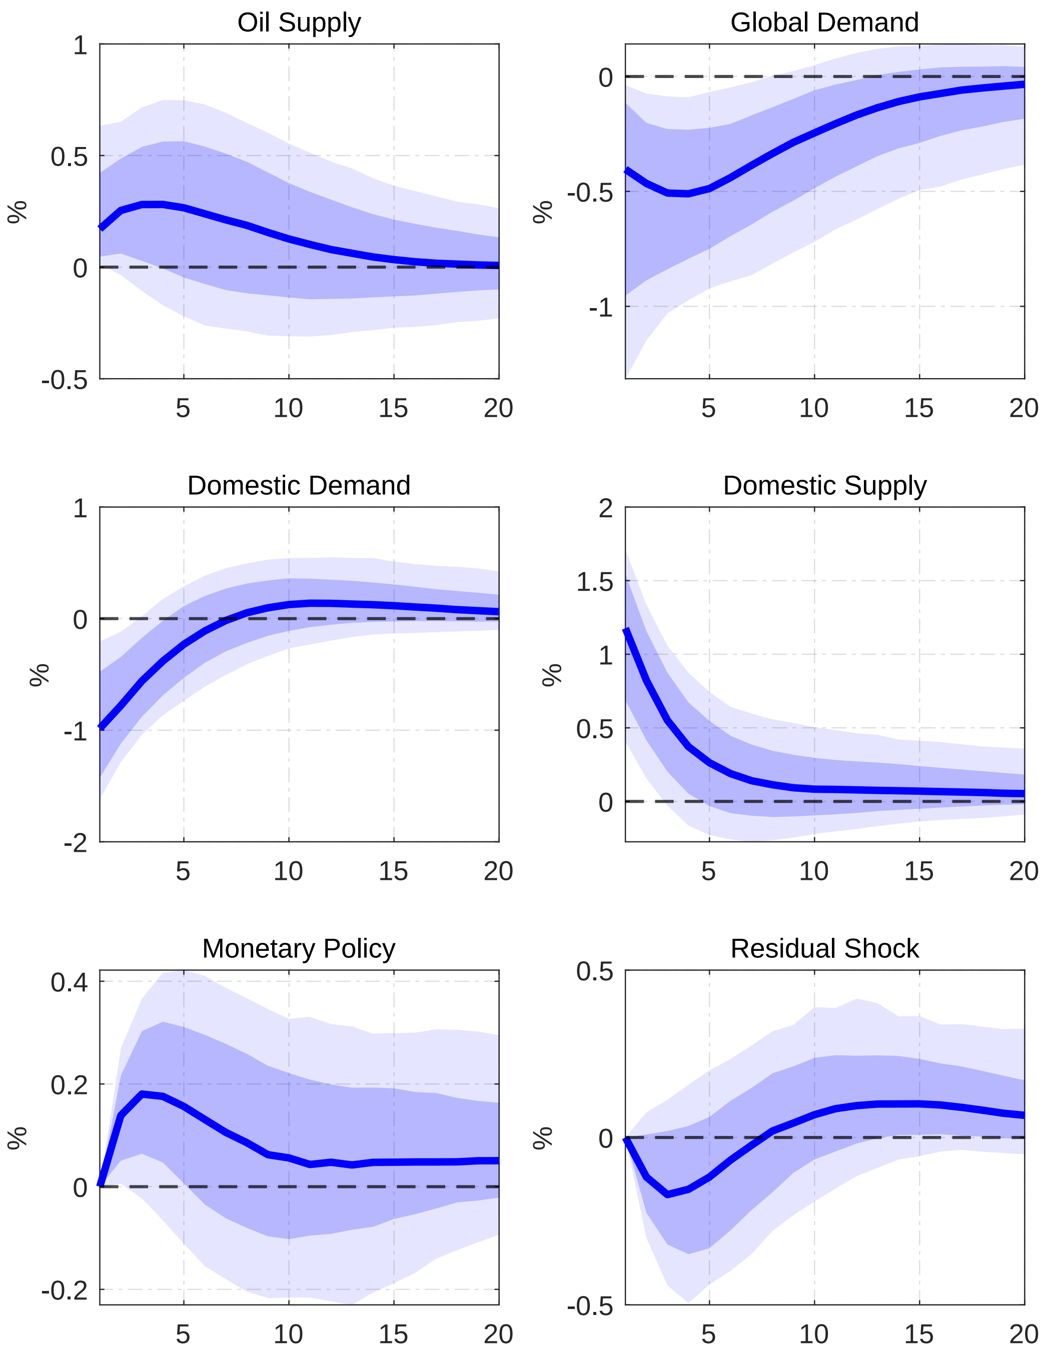

Notes: Median impulse response (red solid path) with 16th and 84th percentiles (red-shaded area) as well as 5th and 95th percentiles (grey-shaded area). The sign-restrictions are imposed on impact only. The identification of shocks relies on the Corsetti et al. (2014) identification.

Figure 11: Impulse responses of inflation - Korea

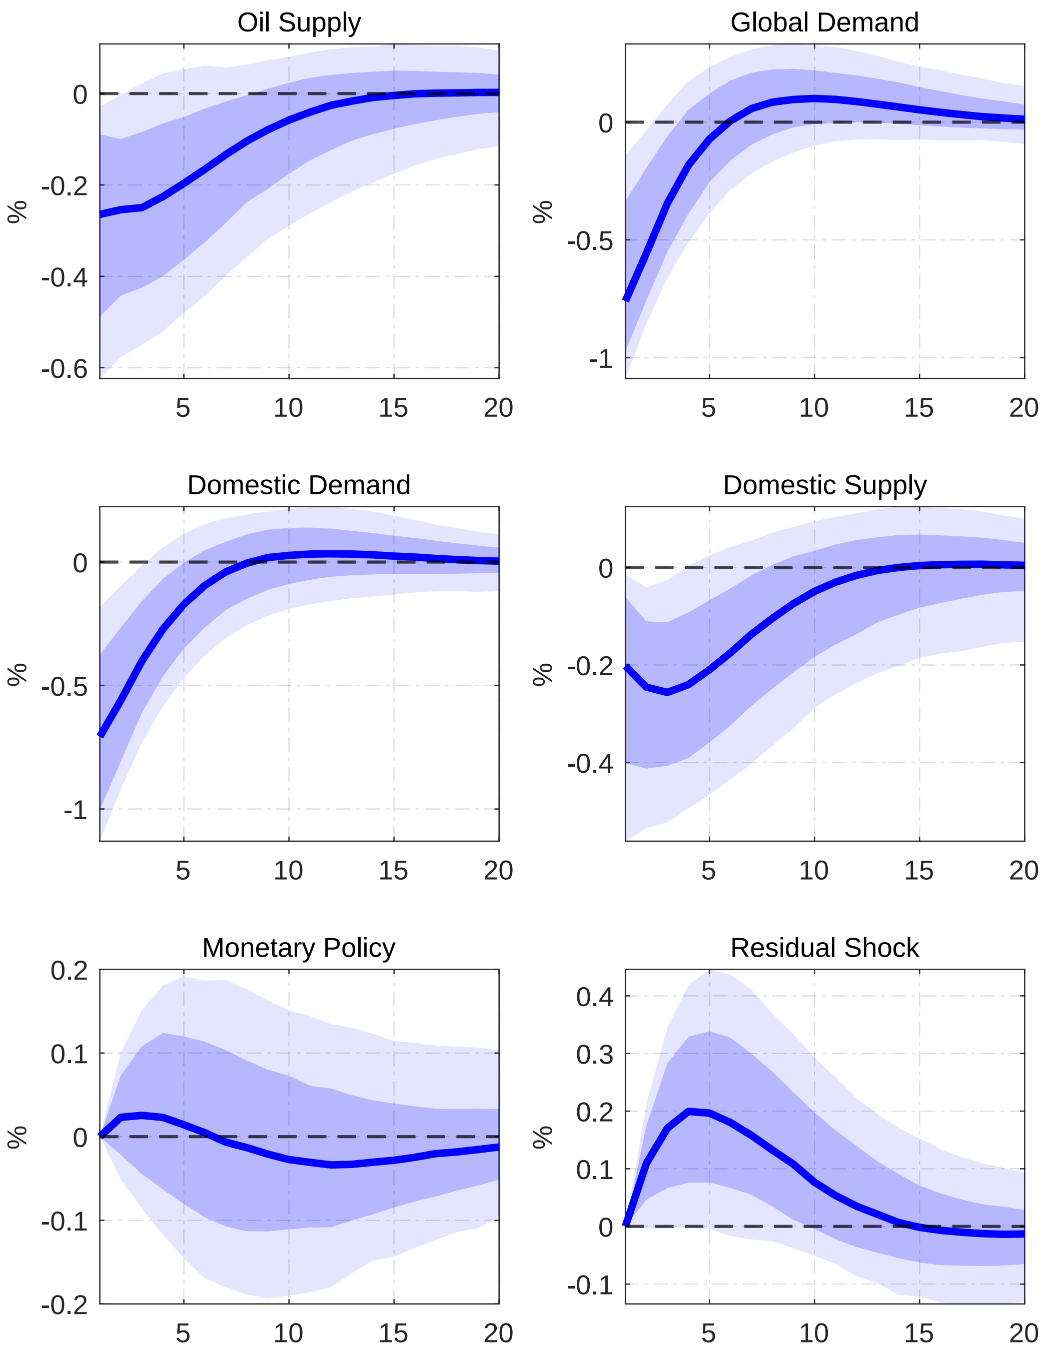

Notes: Median impulse response (red solid path) with 16th and 84th percentiles (red-shaded area) as well as 5th and 95th percentiles (grey-shaded area). The sign-restrictions are imposed on impact only. The identification of shocks relies on the Corsetti et al. (2014) identification.

Figure 12: Impulse responses of real GDP growth - Korea

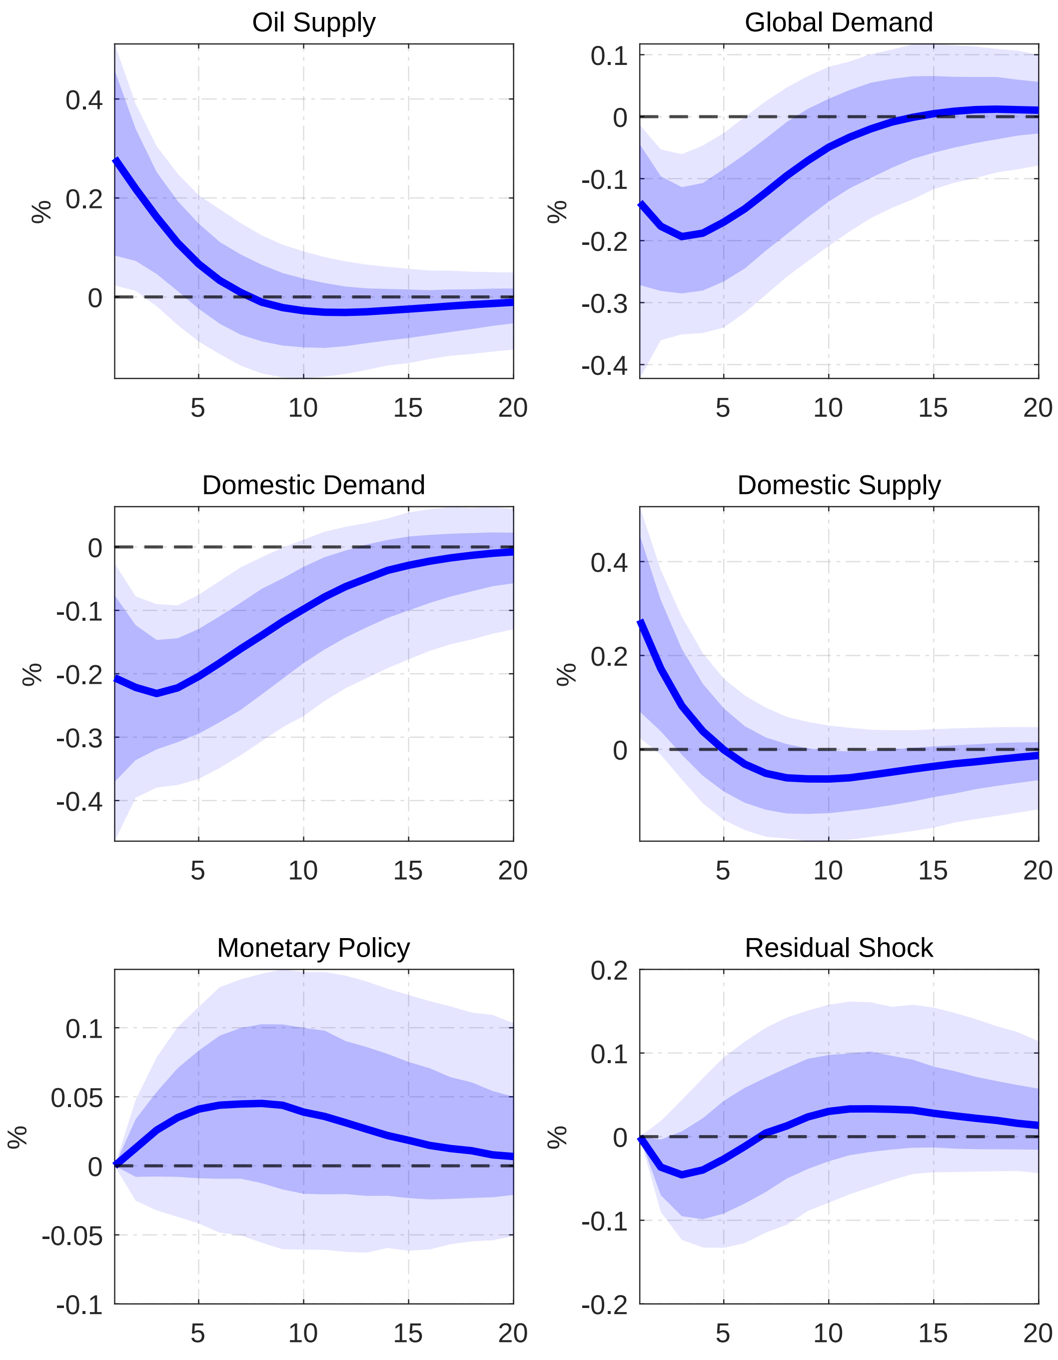

Notes: Median impulse response (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area) as well as 5^{th} and 95^{th} percentiles (grey-shaded area). The sign-restrictions are imposed on impact only. The identification of shocks relies on the Corsetti et al. (2014) identification.

Figure 13: Impulse responses of inflation - Malaysia

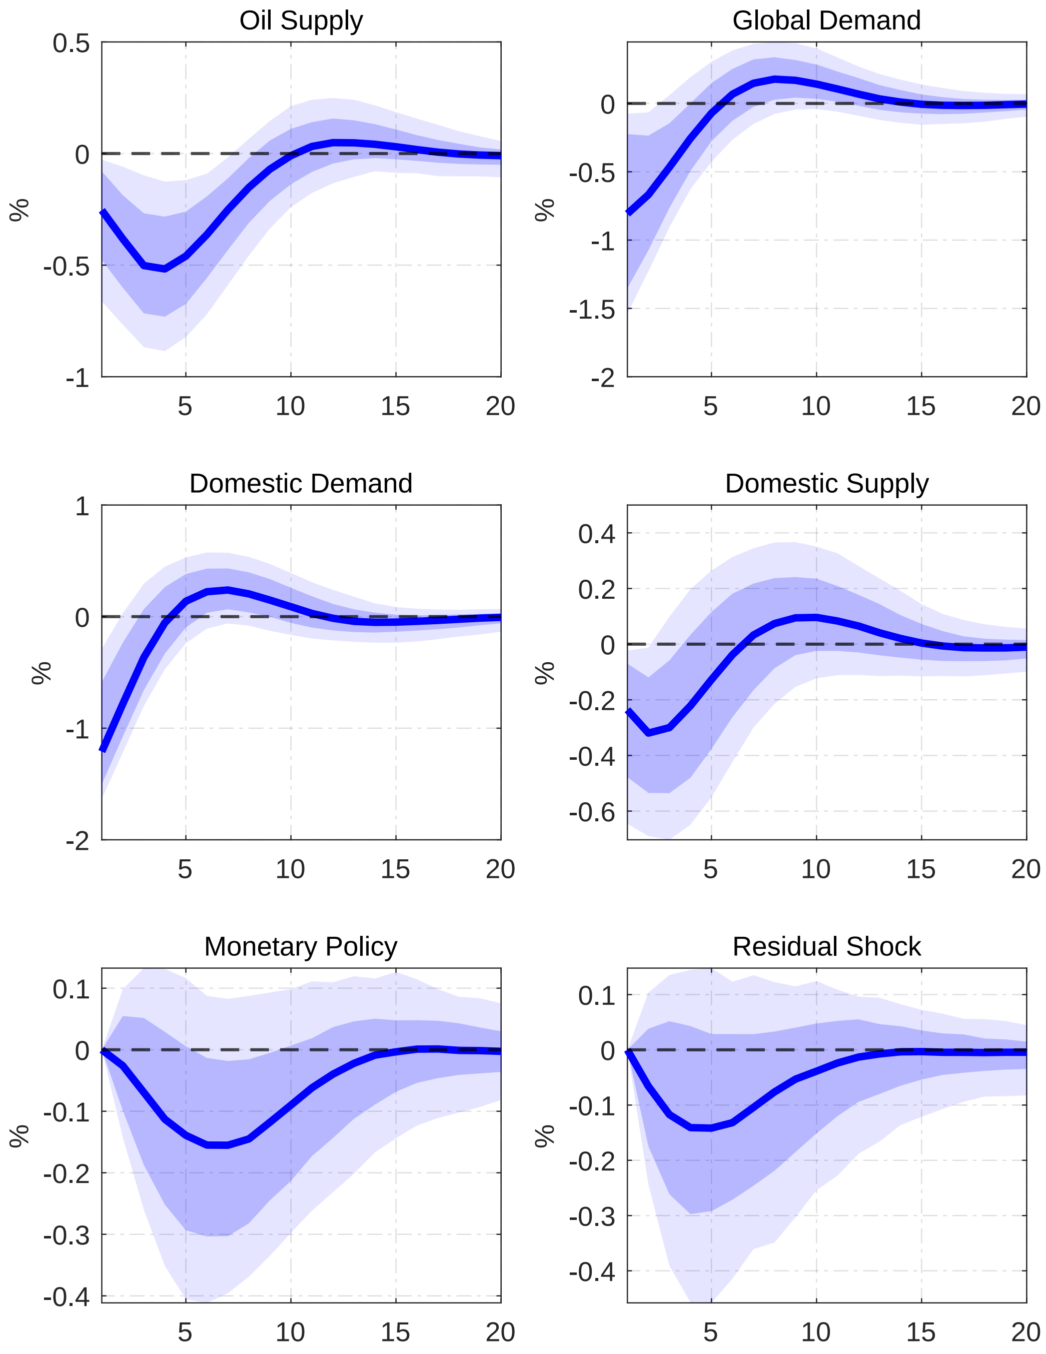

Notes: Median impulse response (red solid path) with 16th and 84th percentiles (red-shaded area) as well as 5th and 95th percentiles (grey-shaded area). The sign-restrictions are imposed on impact only. The identification of shocks relies on the Corsetti et al. (2014) identification.

Figure 14: Impulse responses of real GDP growth - Malaysia

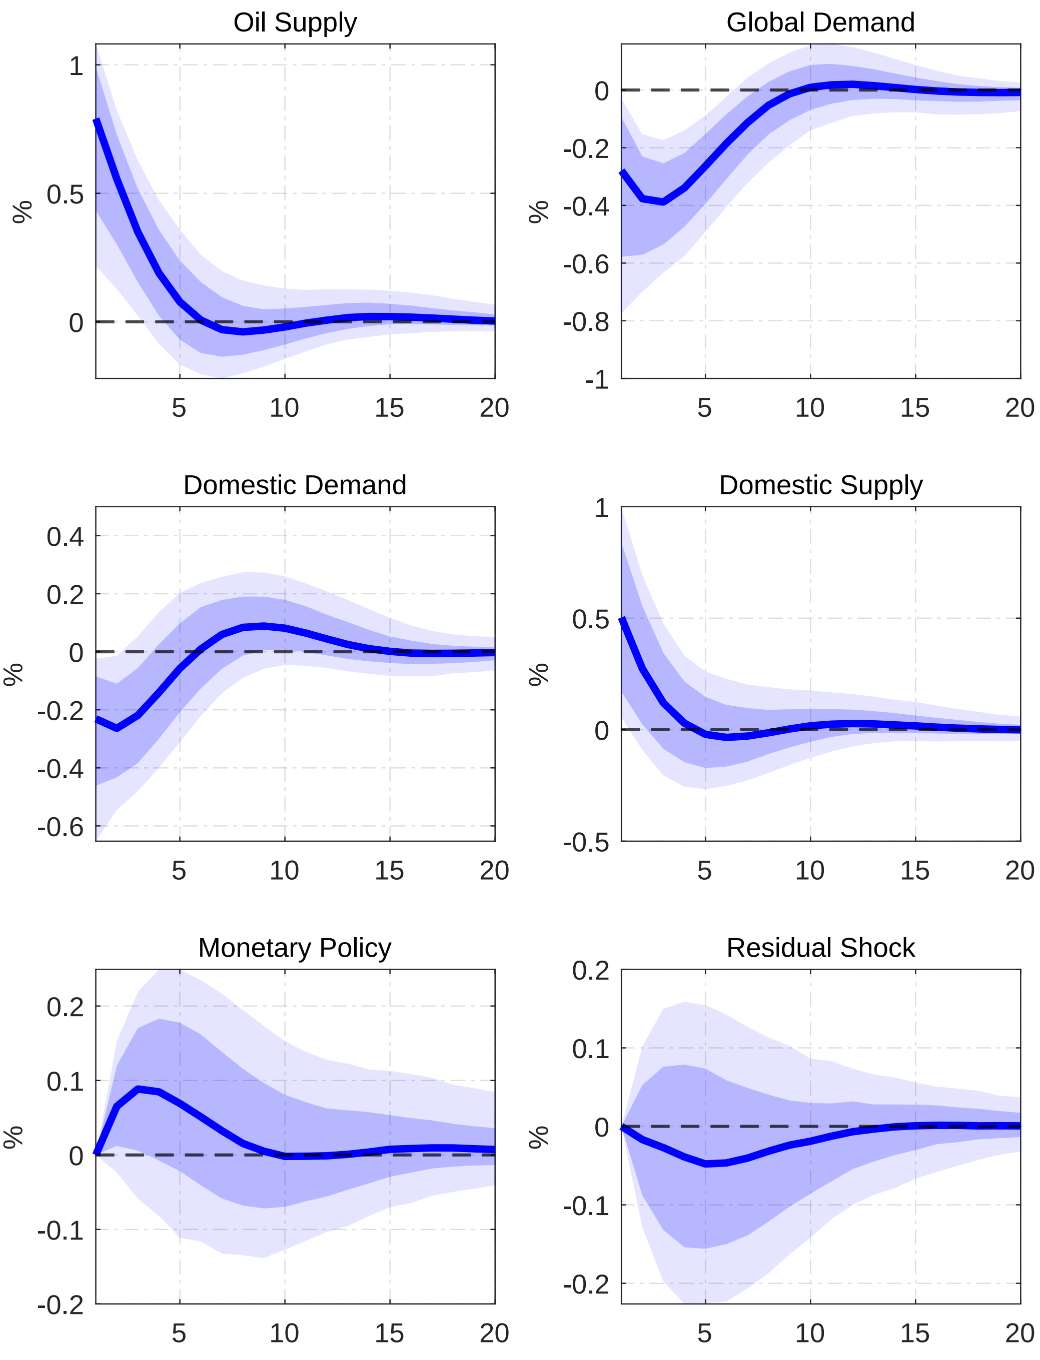

Notes: Median impulse response (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area) as well as 5^{th} and 95^{th} percentiles (grey-shaded area). The sign-restrictions are imposed on impact only. The identification of shocks relies on the Corsetti et al. (2014) identification.

Figure 15: Impulse responses of inflation - Philippines

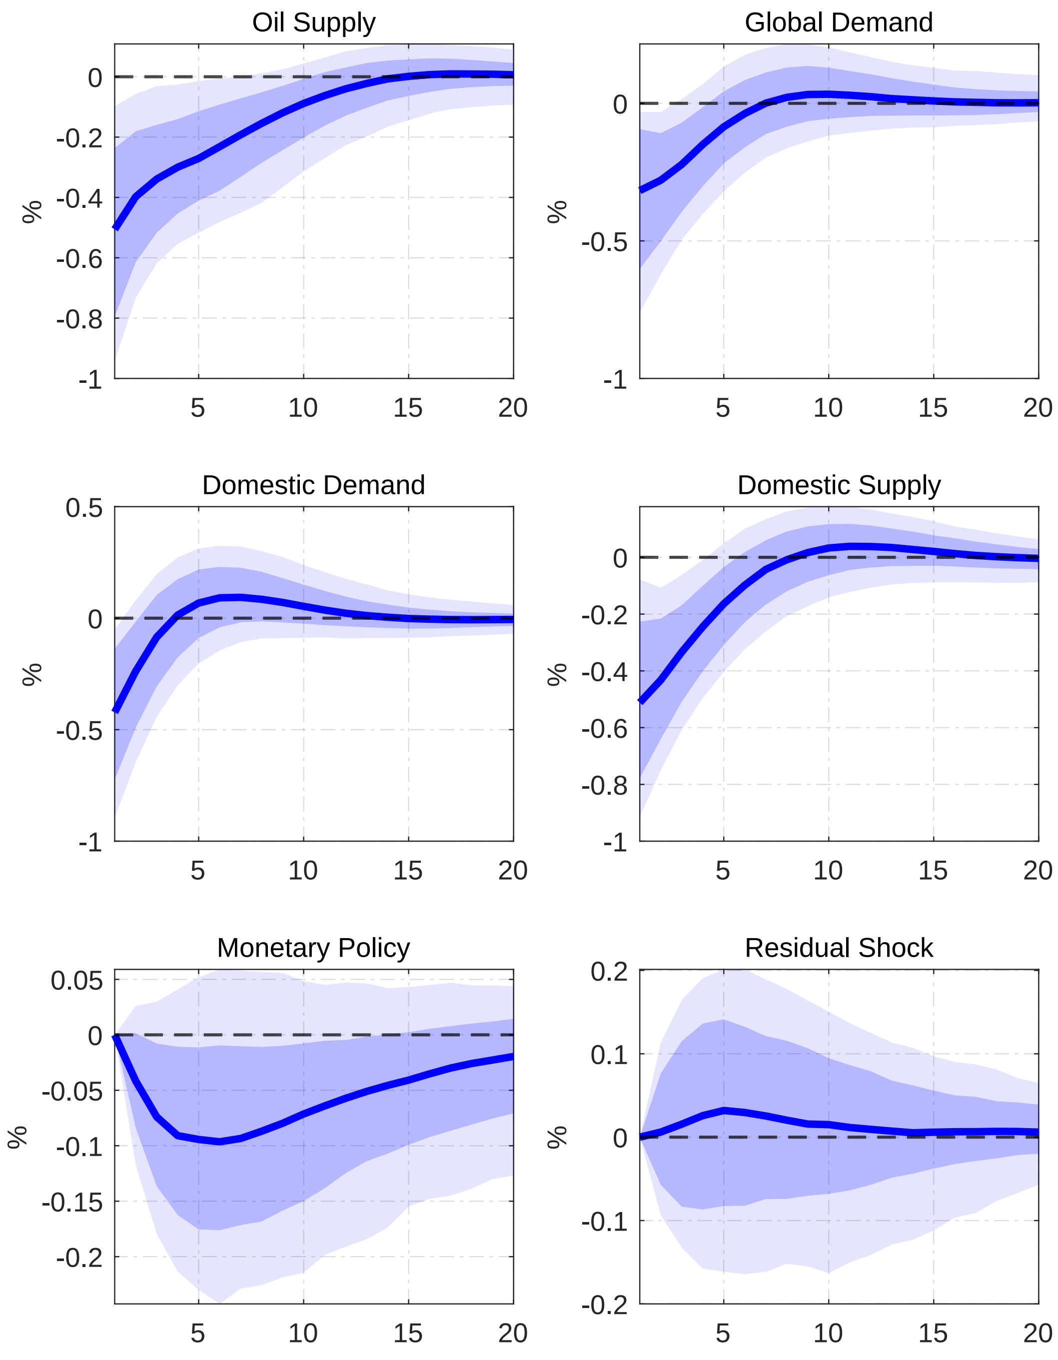

Notes: Median impulse response (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area) as well as 5^{th} and 95^{th} percentiles (grey-shaded area). The sign-restrictions are imposed on impact only. The identification of shocks relies on the Corsetti et al. (2014) identification.

Figure 16: Impulse responses of real GDP growth - Philippines

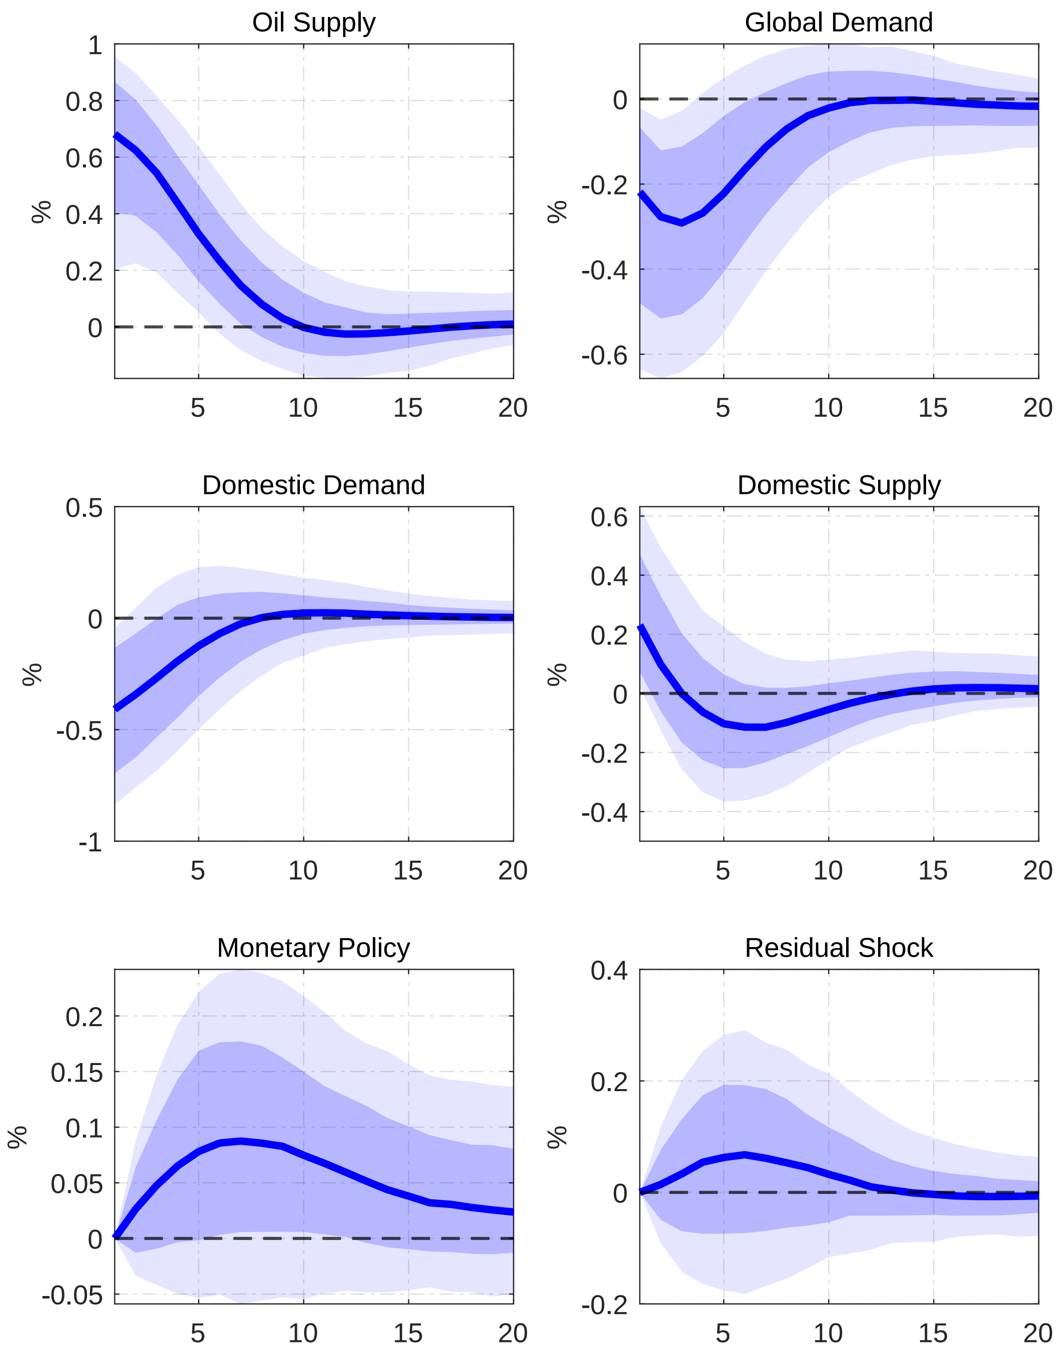

Notes: Median impulse response (red solid path) with 16th and 84th percentiles (red-shaded area) as well as 5th and 95th percentiles (grey-shaded area). The sign-restrictions are imposed on impact only. The identification of shocks relies on the Corsetti et al. (2014) identification.

Figure 17: Impulse responses of inflation - Singapore

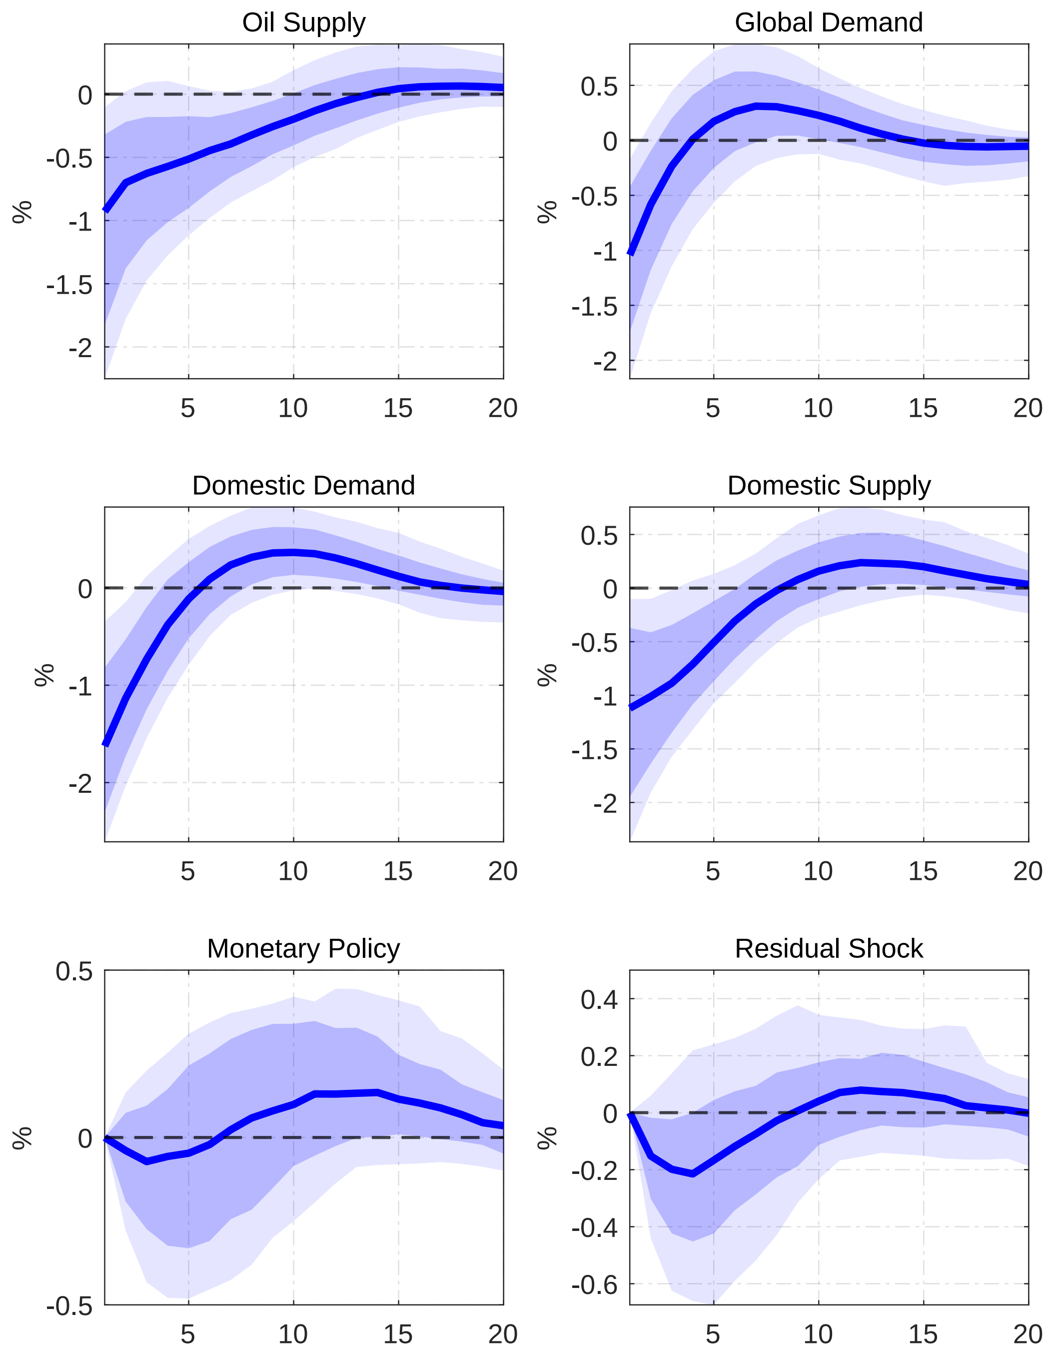

Notes: Median impulse response (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area) as well as 5^{th} and 95^{th} percentiles (grey-shaded area). The sign-restrictions are imposed on impact only. The identification of shocks relies on the Corsetti et al. (2014) identification.

Figure 18: Impulse responses of real GDP growth - Singapore

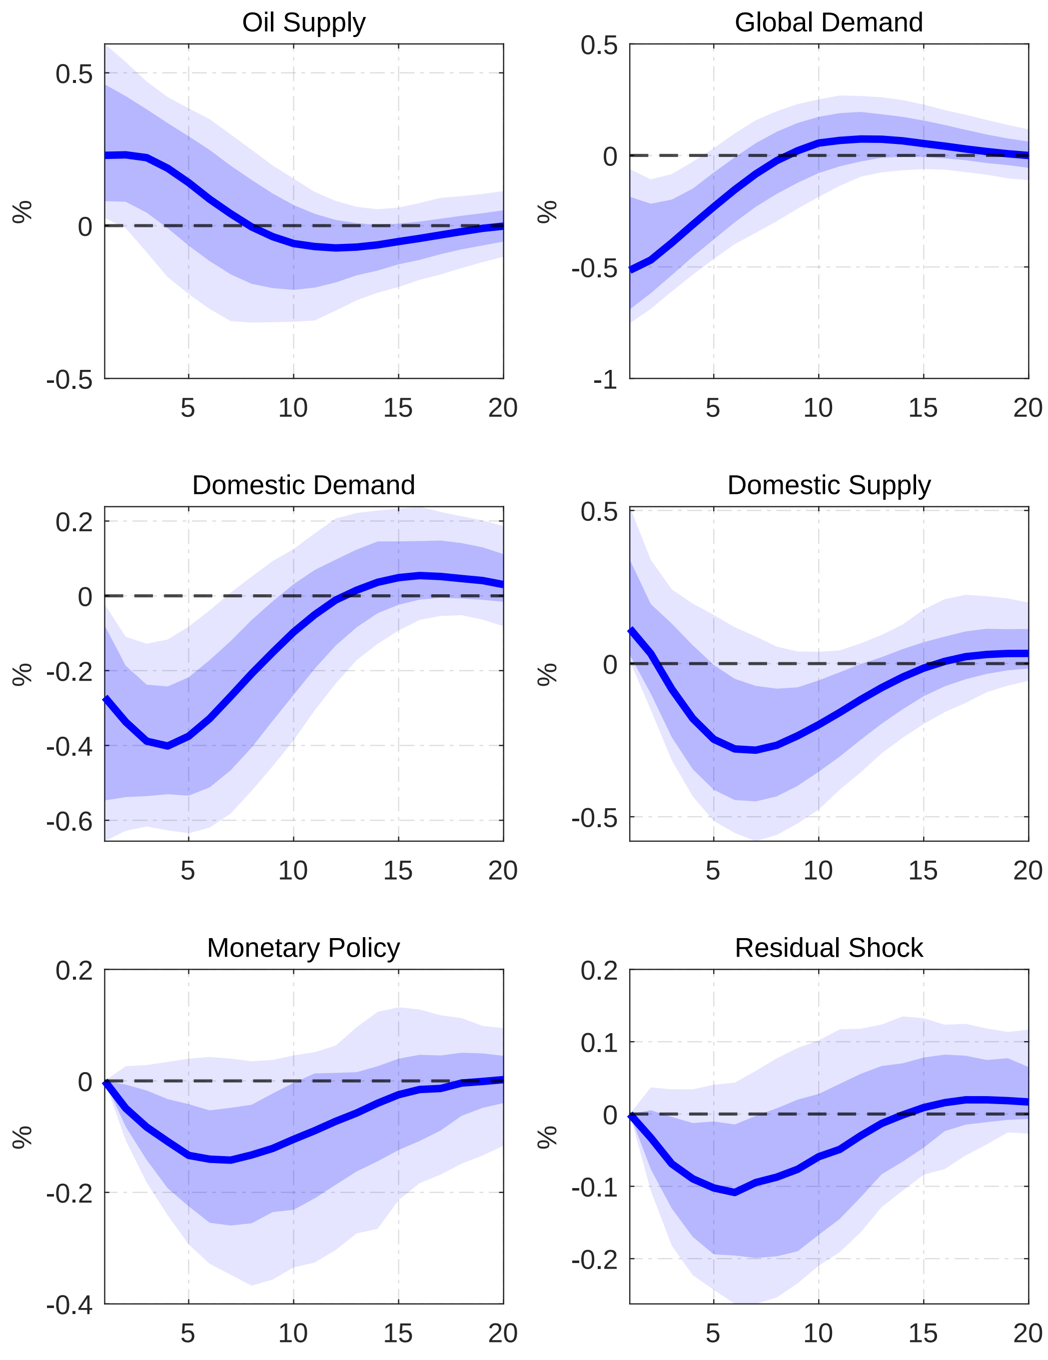

Notes: Median impulse response (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area) as well as 5^{th} and 95^{th} percentiles (grey-shaded area). The sign-restrictions are imposed on impact only. The identification of shocks relies on the Corsetti et al. (2014) identification.

Figure 19: Impulse responses of inflation - Thailand

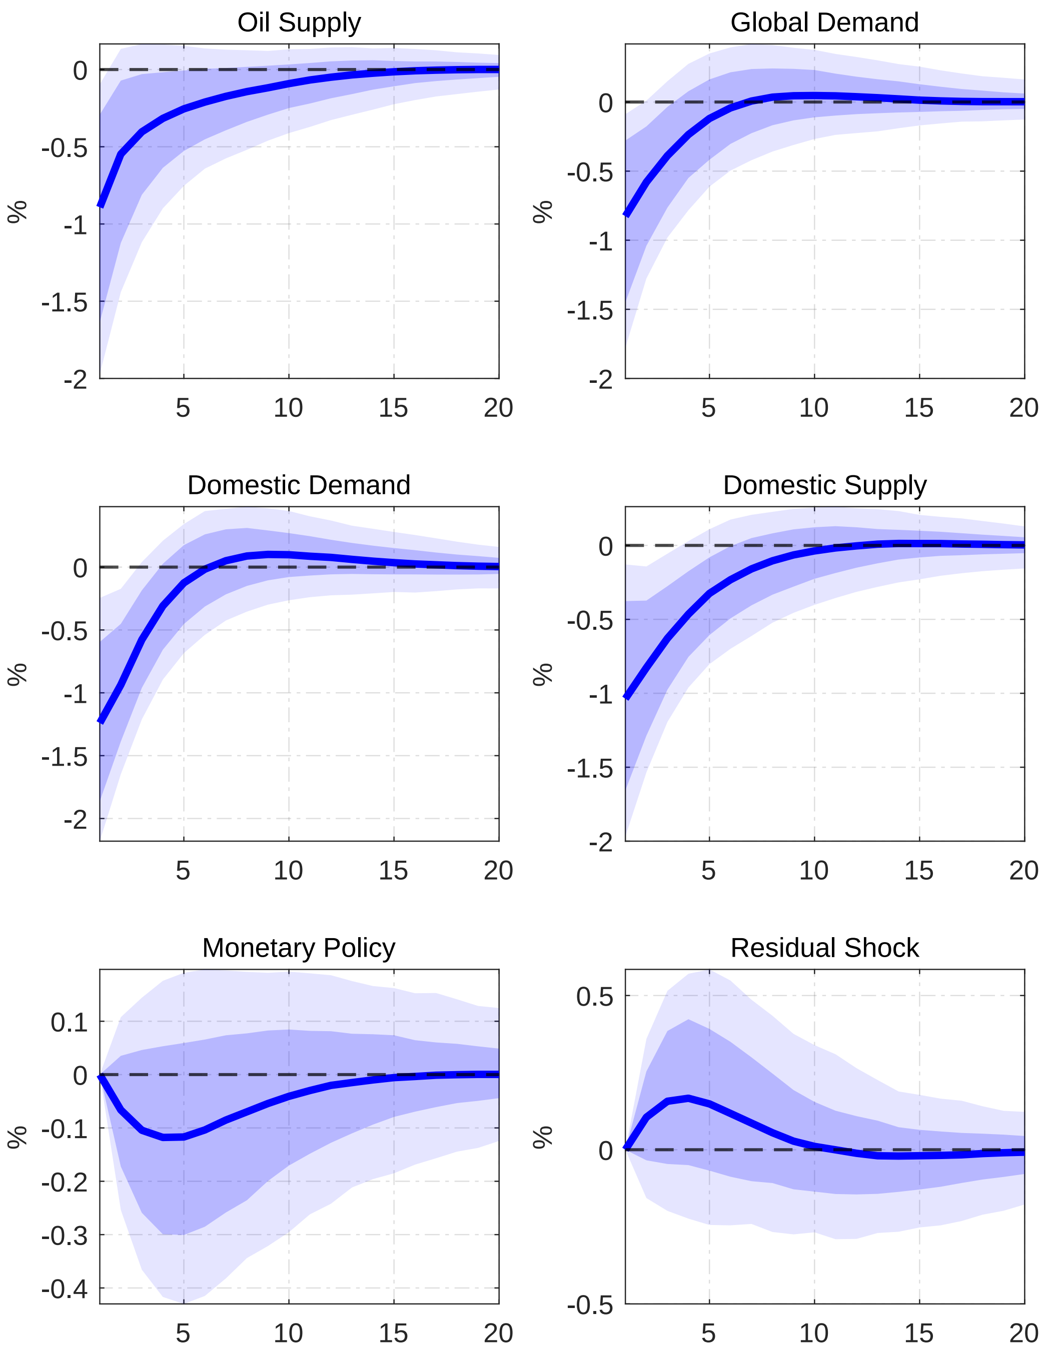

Notes: Median impulse response (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area) as well as 5^{th} and 95^{th} percentiles (grey-shaded area). The sign-restrictions are imposed on impact only. The identification of shocks relies on the Corsetti et al. (2014) identification.

Figure 20: Impulse responses of real GDP growth - Thailand

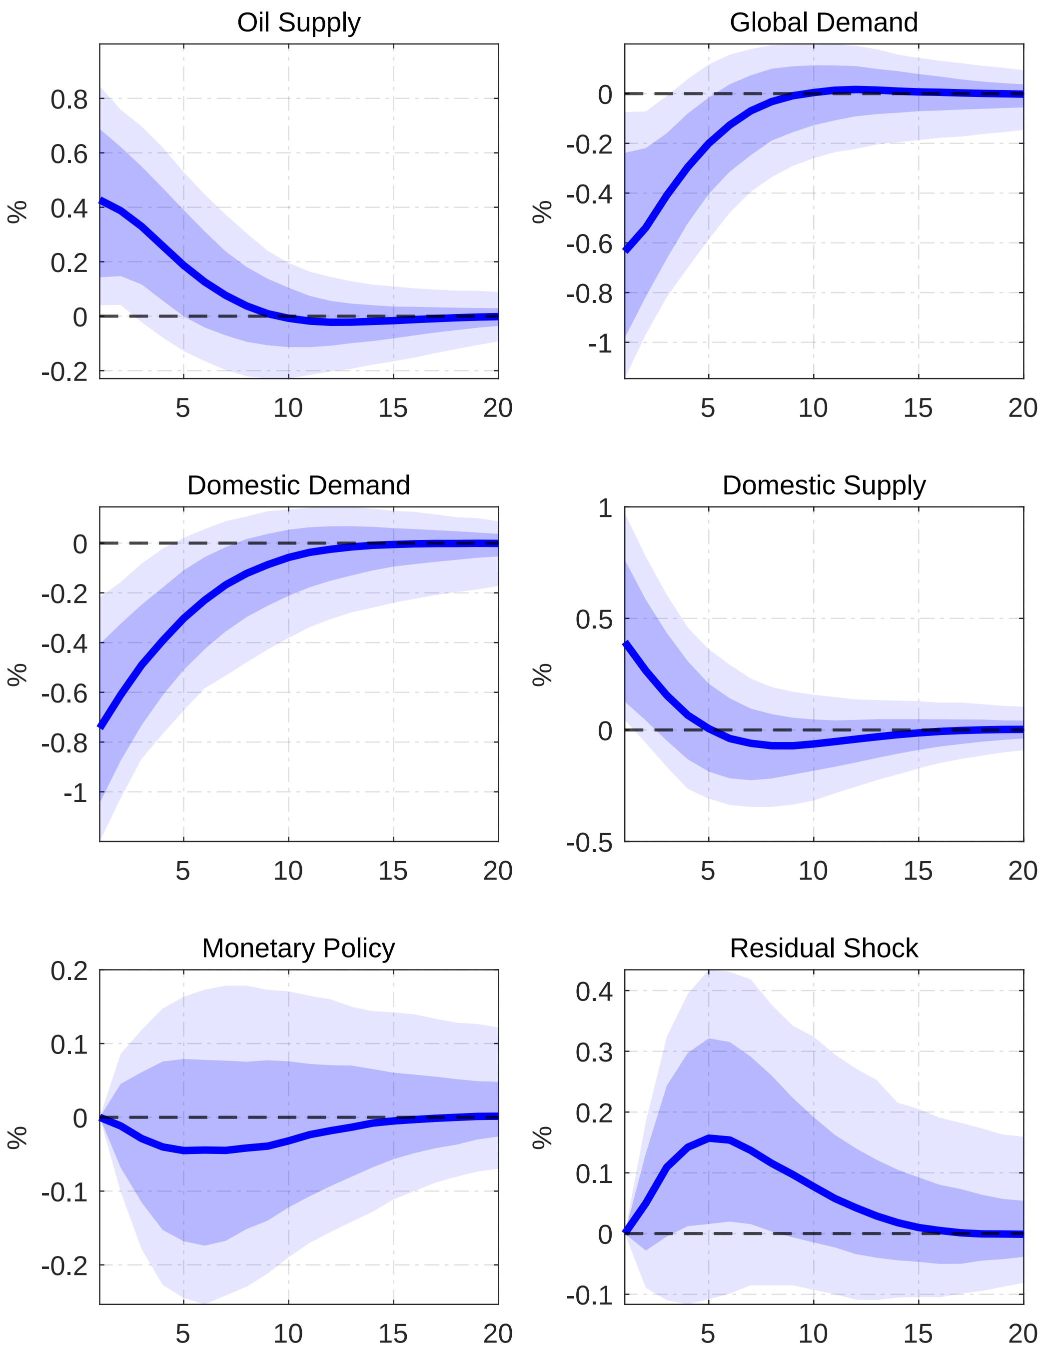

Notes: Median impulse response (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area) as well as 5^{th} and 95^{th} percentiles (grey-shaded area). The sign-restrictions are imposed on impact only. The identification of shocks relies on the Corsetti et al. (2014) identification.

B. *Counterfactual Simulations*

In the main text, we present counterfactual simulations where we suppress the historical contributions of certain shocks. The idea can be summarized as follows: having derived the historical decomposition for each sample, counterfactual paths with suppressed shocks can be simulated as the difference between the actual path minus the historical decompositions of one single specific shock or groups of structural shocks for that sample.[22] Formally, variable i in t can be decomposed into the deterministic component (i.e. initial conditions) and the historical contribution of each structural shock, that is

$$\mathbf{y}_{i,t} = \mathbf{d}_i^{(t)} + \sum_{j=1}^{n} \sum_{k=0}^{t-1} \widetilde{\psi}_{k,ij} \varepsilon_{j,t-k}$$

The first expression states that $\mathbf{y}_{i,t}$ is the sum of the deterministic component $\mathbf{d}_i^{(t)}$ and the sum of contributions of all $j = 1, .., n$ structural shocks on variable i from period $k = 0$ to $t - 1$ back in the past, where $\varepsilon_{j,t-k}$ is the structural shock j in period $t - k$ and $\widetilde{\psi}_{k,ij}$ is the corresponding entry in row i and column j of the structural impulse response function matrix $\widetilde{\mathbf{\Psi}}_k$, i.e. it corresponds to the impact of shock j on variable i. The counterfactual path with suppressed shock $s \in j = 1, ..., n$ is then given by

$$\mathbf{y}_{i,t}^{cf} = \mathbf{d}_i^{(t)} + \sum_{j=1, j \neq s}^{n} \sum_{k=0}^{t-1} \widetilde{\psi}_{k,ij} \varepsilon_{j,t-k},$$

where the historical contribution of shock s on variable i is subtracted. While we show the counterfactual paths of inflation for all countries in the main text, we present the counterfactual paths for the growth rate of real GDP in this appendix. This is interesting because we use the counterfactual paths for real GDP as input for the regression in Section IV in the main paper.

[22] Note that throughout the paper as well as here in the appendix, we report the 5^{th} percentiles, the 16^{th} percentiles, the median, the 84^{th} percentiles as well as the 95^{th} percentiles over all 2000 samples/draws.

Figure 21: Counterfactual paths for real GDP growth with suppressed shocks – Indonesia

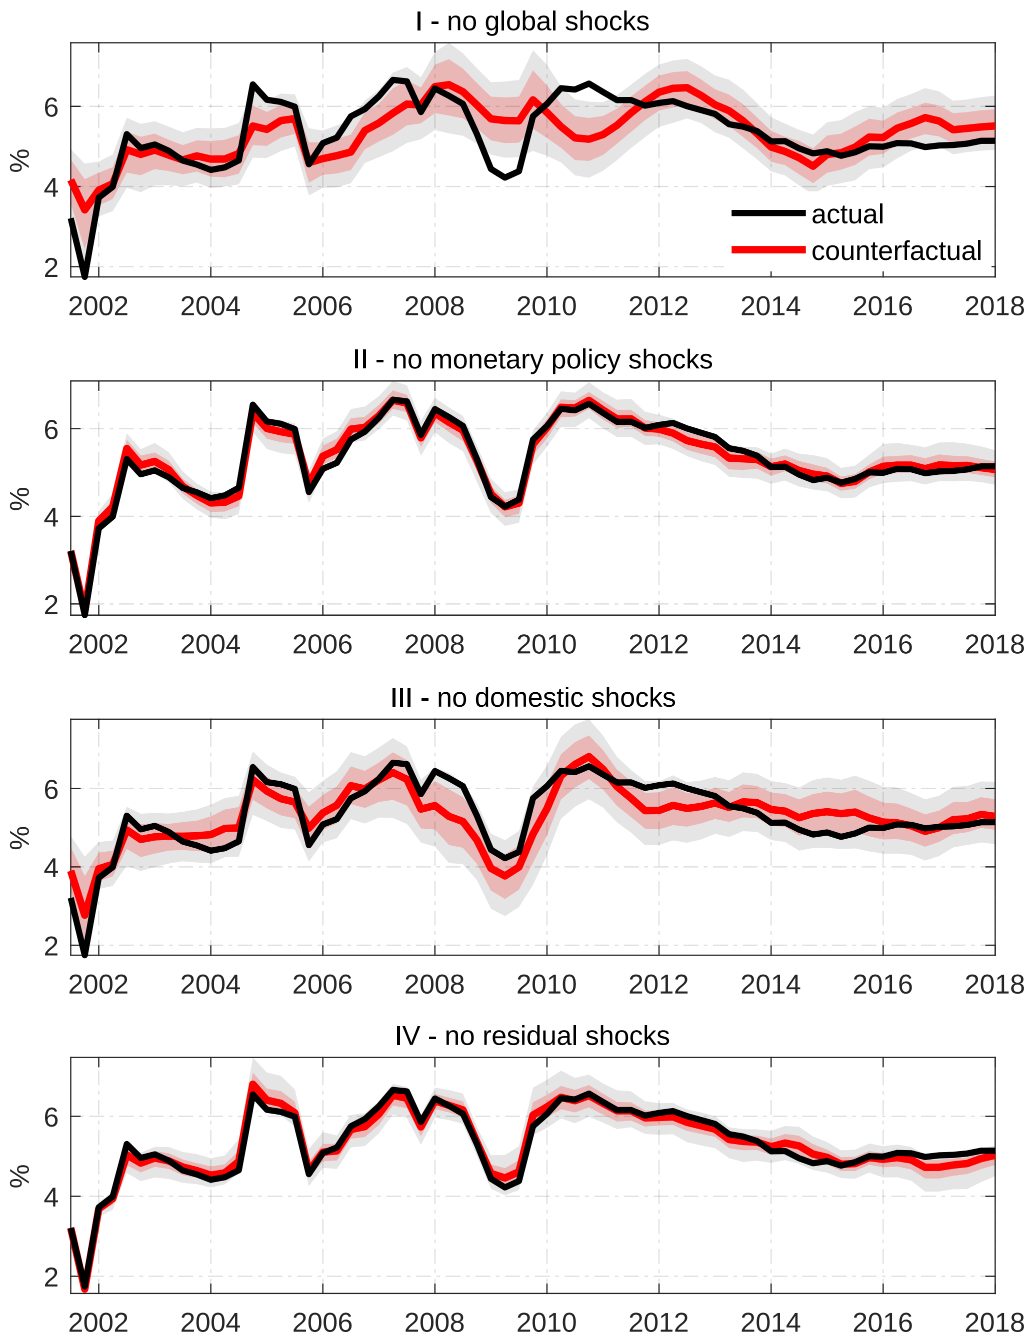

Notes: Median counterfactual paths (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area), and 5^{th} and 95^{th} percentiles (grey-shaded area) for the growth rate of real GDP (in %). In I, the counterfactual path corresponds to the growth rate of real GDP where aggregate global shocks (oil supply shocks and global demand shocks) are suppressed, while the same is done in II with monetary policy shocks, in III with domestic shocks (domestic demand and domestic supply shocks) and in IV (residual shocks). The identification of shocks relies on the Corsetti et al. (2014) identification.

Figure 22: Counterfactual paths for real GDP growth with suppressed shocks – Korea

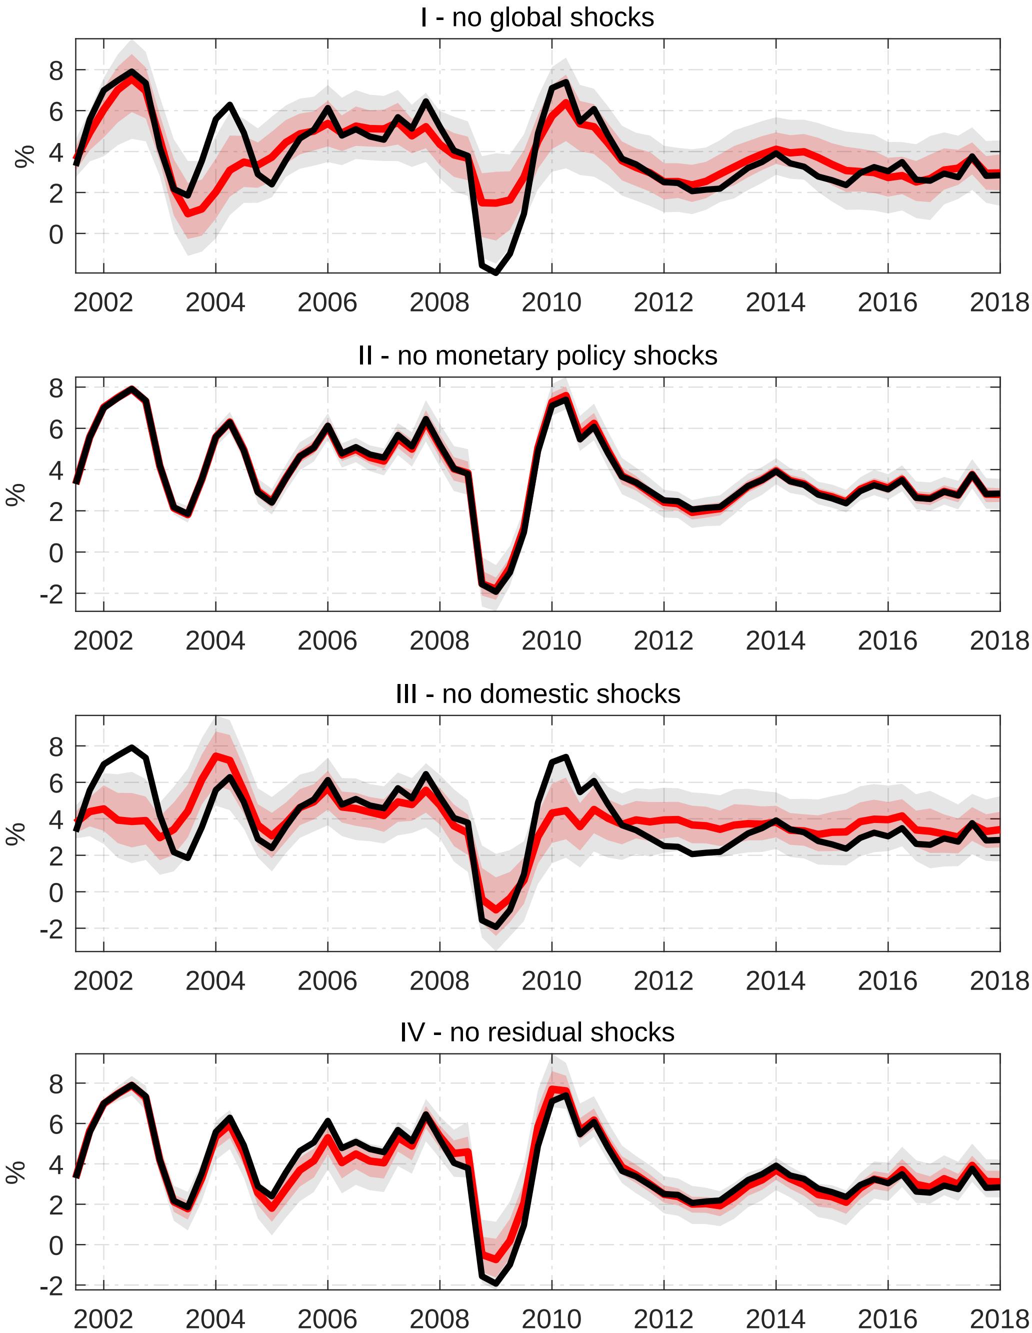

Notes: Median counterfactual paths (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area), and 5^{th} and 95^{th} percentiles (grey-shaded area) for the growth rate of real GDP (in %). In I, the counterfactual path corresponds to the growth rate of real GDP where aggregate global shocks (oil supply shocks and global demand shocks) are suppressed, while the same is done in II with monetary policy shocks, in III with domestic shocks (domestic demand and domestic supply shocks) and in IV (residual shocks). The identification of shocks relies on the Corsetti et al. (2014) identification.

Figure 23: Counterfactual paths for real GDP growth with suppressed shocks – Malaysia

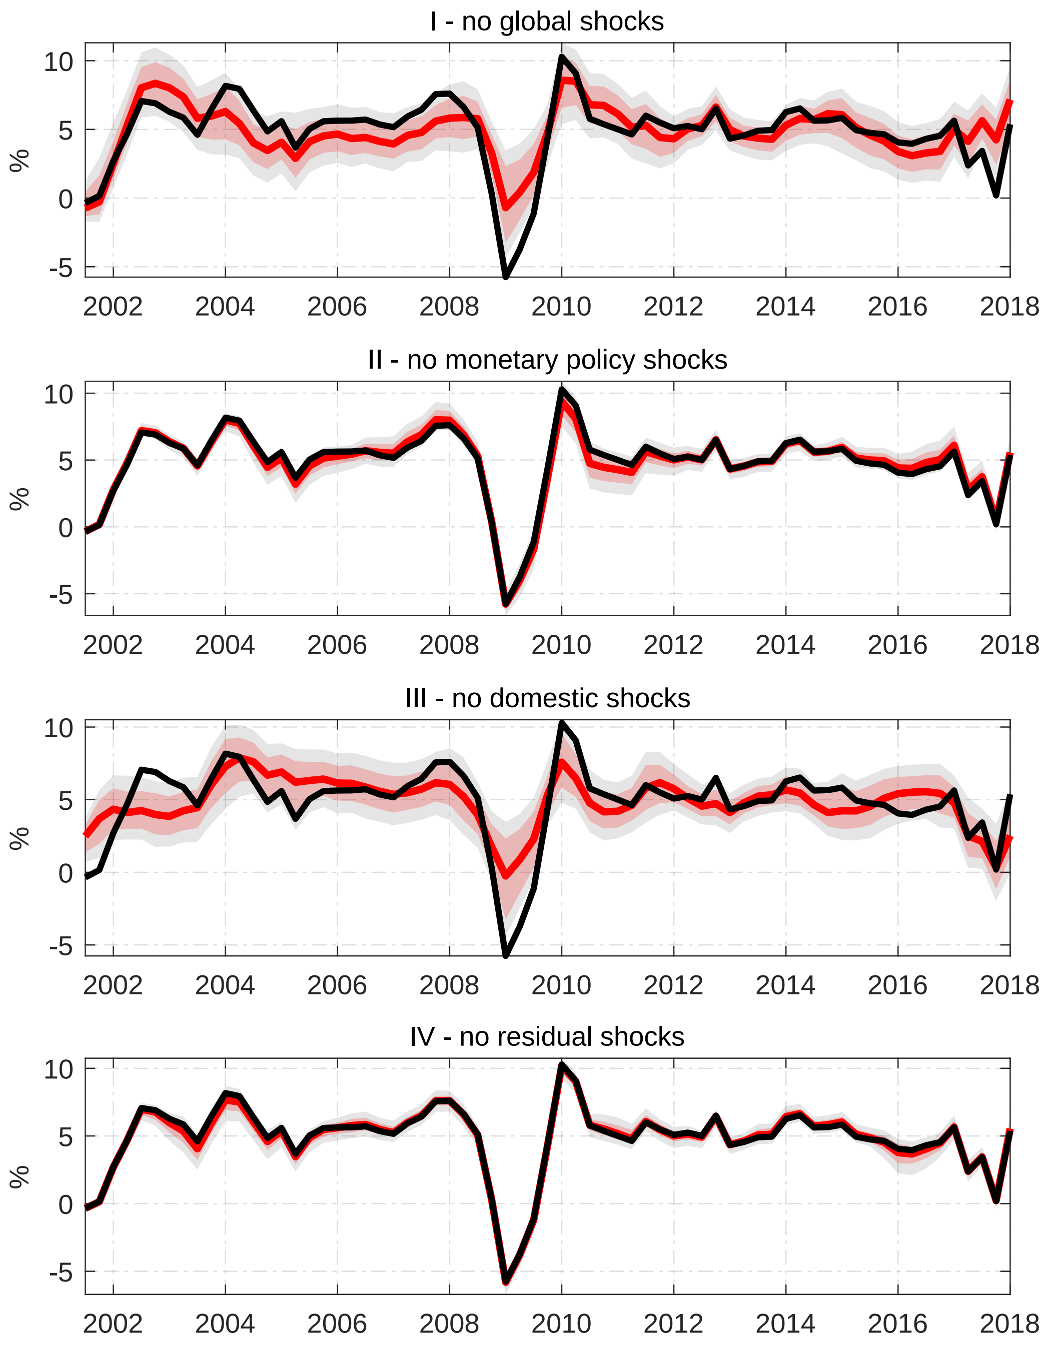

Notes: Median counterfactual paths (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area), and 5^{th} and 95^{th} percentiles (grey-shaded area) for the growth rate of real GDP (in %). In I, the counterfactual path corresponds to the growth rate of real GDP where aggregate global shocks (oil supply shocks and global demand shocks) are suppressed, while the same is done in II with monetary policy shocks, in III with domestic shocks (domestic demand and domestic supply shocks) and in IV (residual shocks). The identification of shocks relies on the Corsetti et al. (2014) identification.

Figure 24: Counterfactual paths for real GDP growth with suppressed shocks – Philippines

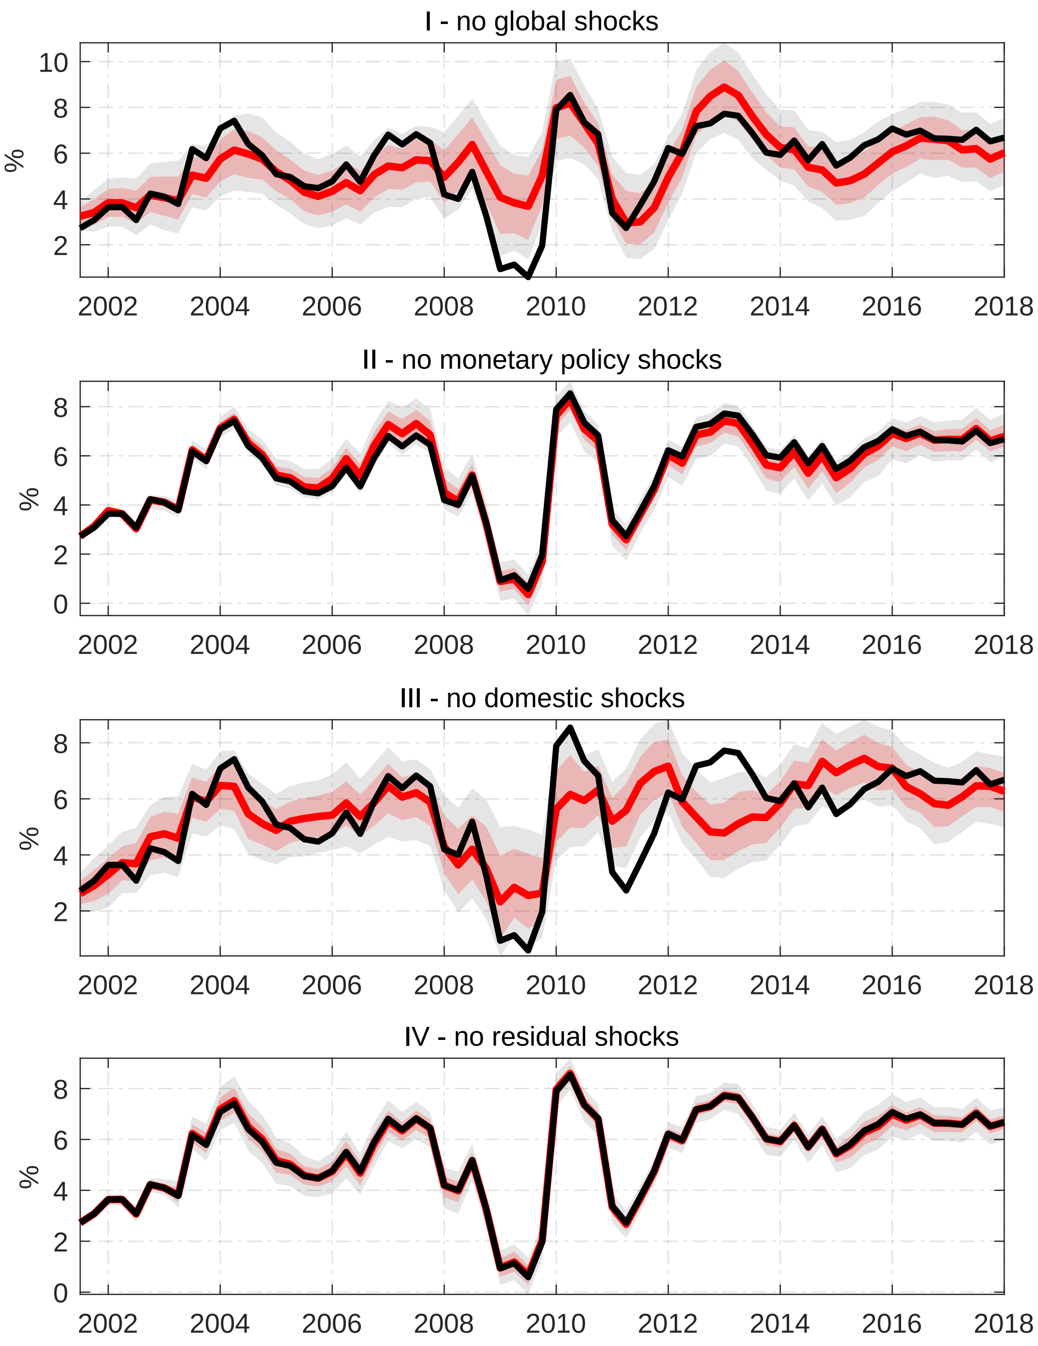

Notes: Median counterfactual paths (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area), and 5^{th} and 95^{th} percentiles (grey-shaded area) for the growth rate of real GDP (in %). In I, the counterfactual path corresponds to the growth rate of real GDP where aggregate global shocks (oil supply shocks and global demand shocks) are suppressed, while the same is done in II with monetary policy shocks, in III with domestic shocks (domestic demand and domestic supply shocks) and in IV (residual shocks). The identification of shocks relies on the Corsetti et al. (2014) identification.

Figure 25: Counterfactual paths for real GDP growth with suppressed shocks - Singapore

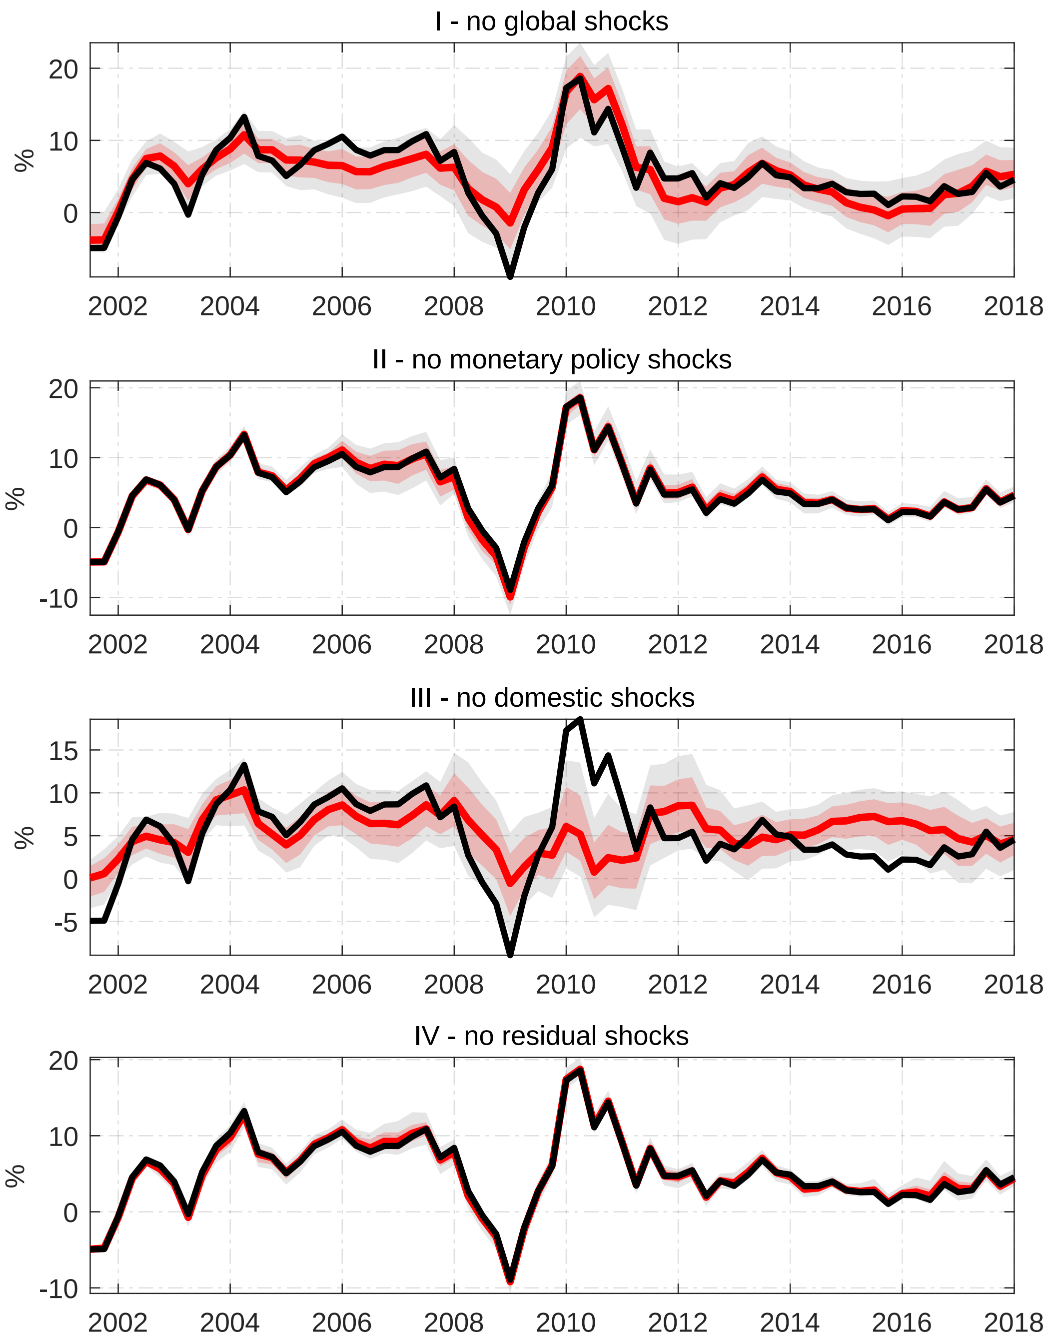

Notes: Median counterfactual paths (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area), and 5^{th} and 95^{th} percentiles (grey-shaded area) for the growth rate of real GDP (in %). In I, the counterfactual path corresponds to the growth rate of real GDP where aggregate global shocks (oil supply shocks and global demand shocks) are suppressed, while the same is done in II with monetary policy shocks, in III with domestic shocks (domestic demand and domestic supply shocks) and in IV (residual shocks). The identification of shocks relies on the Corsetti et al. (2014) identification.

Figure 26: Counterfactual paths for real GDP growth with suppressed shocks – Thailand

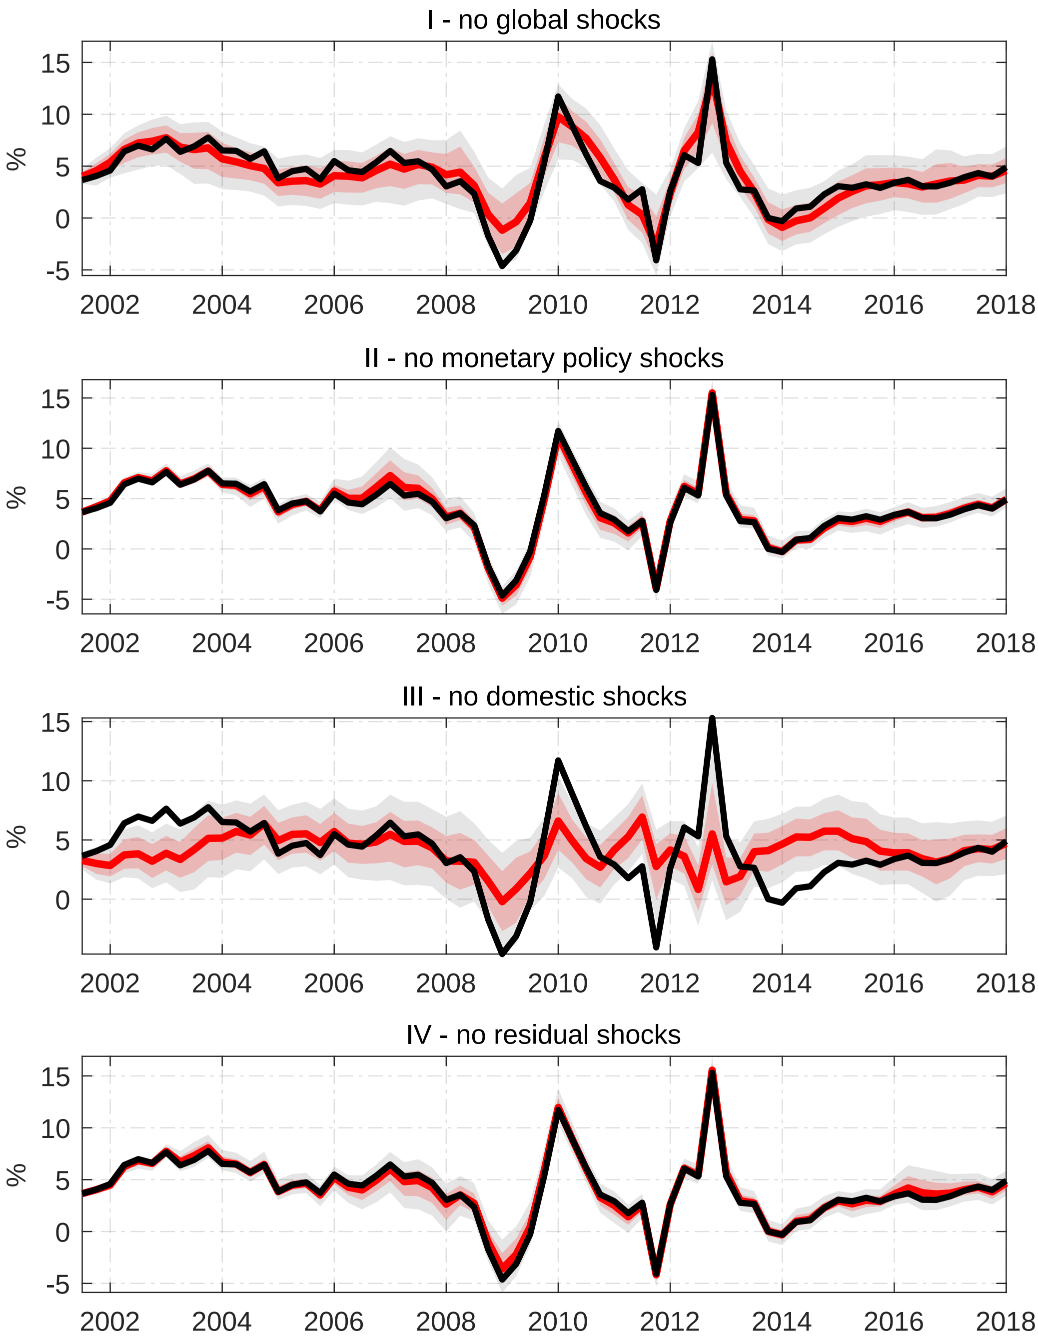

Notes: Median counterfactual paths (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area), and 5^{th} and 95^{th} percentiles (grey-shaded area) for the growth rate of real GDP (in %). In I, the counterfactual path corresponds to the growth rate of real GDP where aggregate global shocks (oil supply shocks and global demand shocks) are suppressed, while the same is done in II with monetary policy shocks, in III with domestic shocks (domestic demand and domestic supply shocks) and in IV (residual shocks). The identification of shocks relies on the Corsetti et al. (2014) identification.

C Further Robustness Checks

This section reports the results for additional robustness checks. These include (1) results for an alternative identification strategy and (2) results for the case where the oil price is restricted to not respond to domestic variables.

A. *Results for the Alternative Identification*

In this subsection, we report several results from an alternative identification of structural shocks. Table (9) reports the identification strategy. In short, the difference between the two identification strategies is the identification of the monetary policy shock. Here we assume that the growth rate of real GDP, the share of real GDP relative to GDP of the rest of the world as well as the inflation rate no longer react with a delay, but react on impact. That is, an unexpected increase in the policy rate leads to a drop in the inflation rate and the growth rate of real GDP. Since real GDP reacts more strongly after a domestic monetary policy shock than real GDP of the rest of the world, the GDP share also falls on impact.

Table 9: Alternative identification of structural shocks

VARIABLE \ SHOCK	OIL SUPPLY	GLOBAL DEMAND	DOMESTIC DEMAND	DOMESTIC SUPPLY	MONETARY POLICY	RESIDUAL
OIL PRICE	+	–				0
GDP SHARE		+	–	–	–	0
REAL GDP	–	–	–	–	–	0
INFLATION	+	–	–	+	–	0
INTEREST RATE	0		–	+	+	0
EXCHANGE RATE						+

Notes: Blank cells indicate unconstrained impulse responses. A positive or negative reaction is denoted by + and –. A zero restriction is denoted by 0. All restrictions are imposed on impact only.

Before we get to the main results, it is worth taking a look at the correlation of the identified structural shocks between the identification strategy from the main part of the paper as well as the strategy documented in Table (9). For this purpose, Table (10) reports for each country the correlation of the median for both approaches for all six shocks.

Table 10: Correlation of shocks across approaches

SHOCK \ COUNTRY	IDN	KOR	MYS	PHL	SGP	THA
MONETARY POLICY	0.71	0.65	0.74	0.70	0.64	0.66
OIL SUPPLY	0.89	0.98	0.92	0.97	0.88	0.98
GLOBAL DEMAND	0.96	0.96	0.98	0.98	0.87	0.99
DOMESTIC DEMAND	0.63	0.70	0.76	0.53	0.61	0.68
DOMESTIC SUPPLY	0.96	0.84	0.92	0.90	0.82	0.93
RESIDUAL SHOCK	1.00	1.00	1.00	1.00	1.00	1.00

Notes: Correlation of estimated structural shocks from the baseline identification strategy and the alternative identification strategy. All values are based on the median values.

Overall, we find that the correlation is very high for all six shocks. Except for the monetary policy shock and the domestic demand shock, we find that for all six countries, the other four shocks almost overlap. For the monetary policy shock, the correlation between the approaches is noticeably lower, which is due to the different identification strategy, but we also find that the general pattern is very similar.

Figure 27: Historical contribution of structural shocks to inflation for Indonesia, Korea and Malaysia

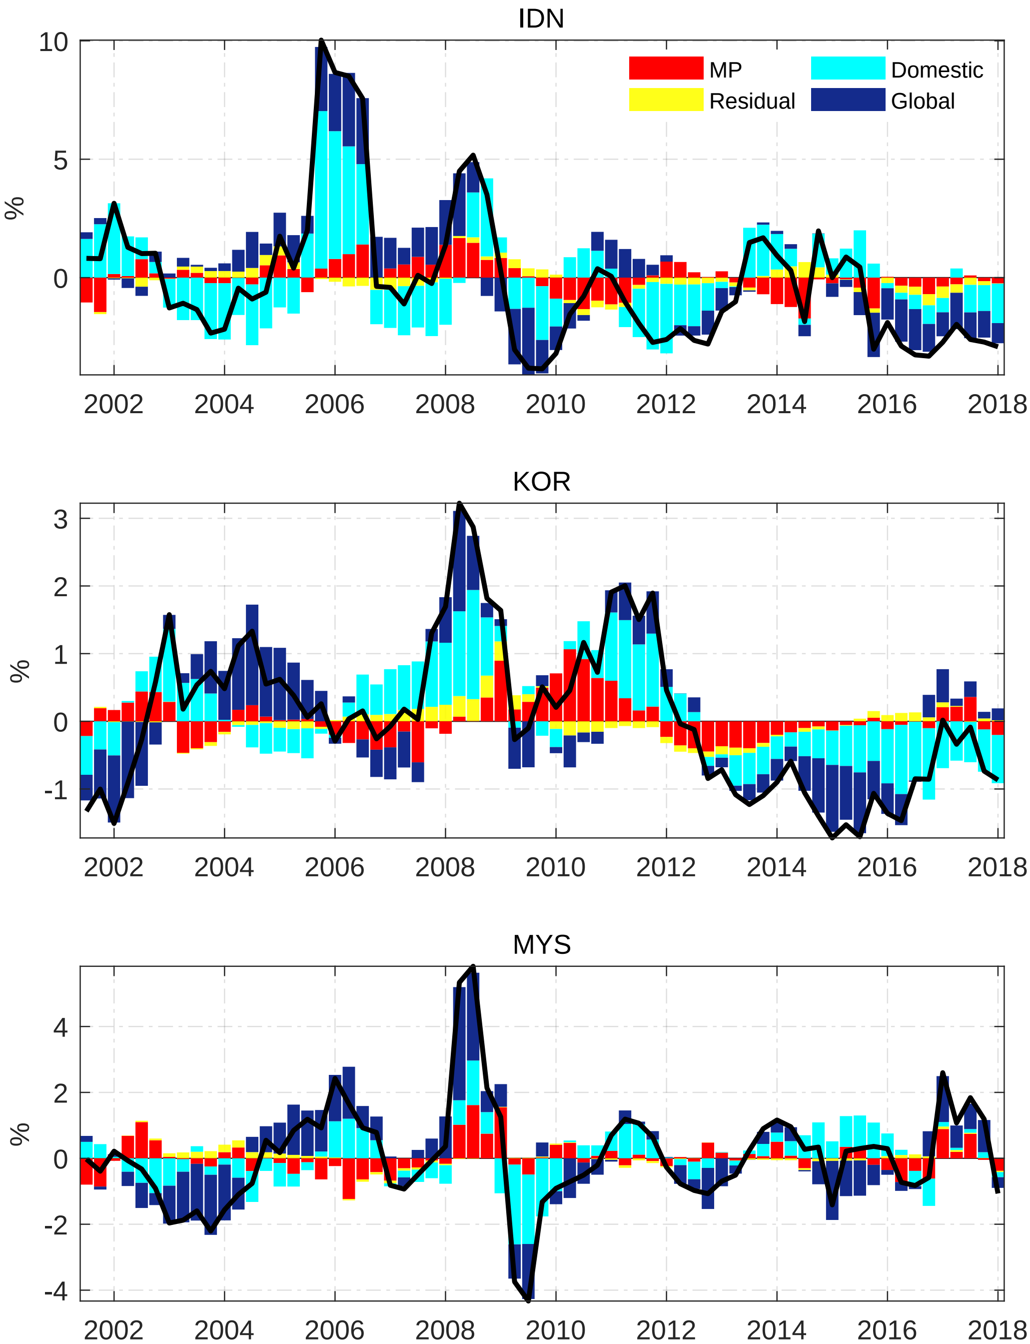

Notes: Median historical contribution of monetary policy shocks (red bars), residual shocks (yellow bars), domestic shocks (teal bars) and global shocks (blue bars) to inflation for Indonesia, Korea and Malaysia. The black path corresponds to the sum of median contributions of all structural shocks. Results rely on the alternative identification strategy.

Figure 28: Historical contribution of structural shocks to inflation for the Philippines, Singapore and Thailand

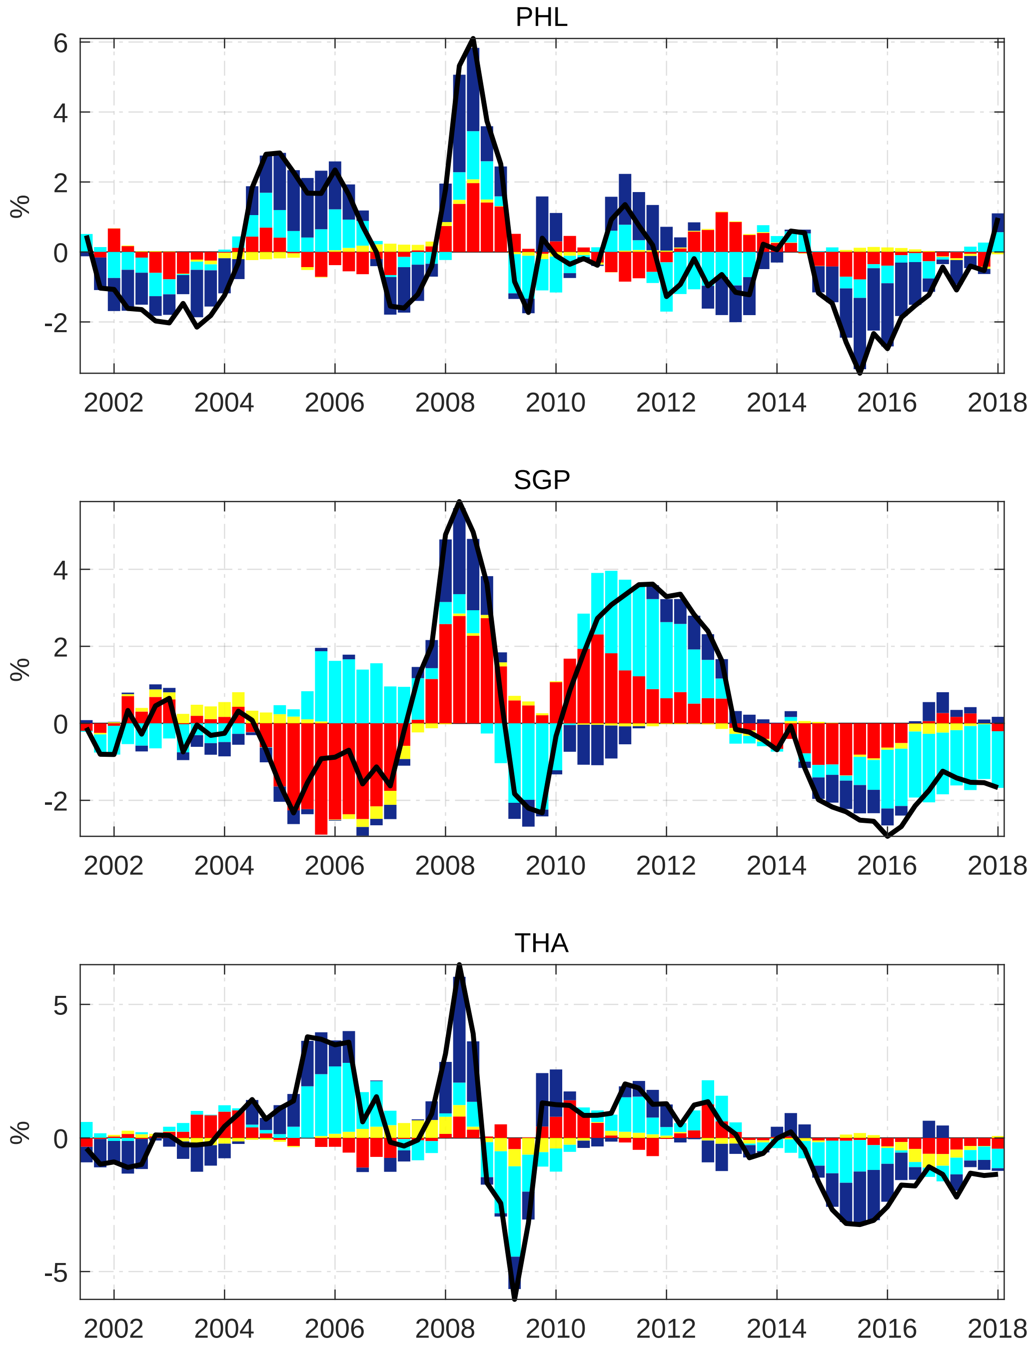

Notes: Median historical contribution of monetary policy shocks (red bars), residual shocks (yellow bars), domestic shocks (teal bars) and global shocks (blue bars) to inflation for the Philippines, Singapore and Thailand. The black path corresponds to the sum of median contributions of all structural shocks. Results rely on the alternative identification strategy.

Figure 29: Historical contribution of structural shocks to real GDP growth for Indonesia, Korea and Malaysia

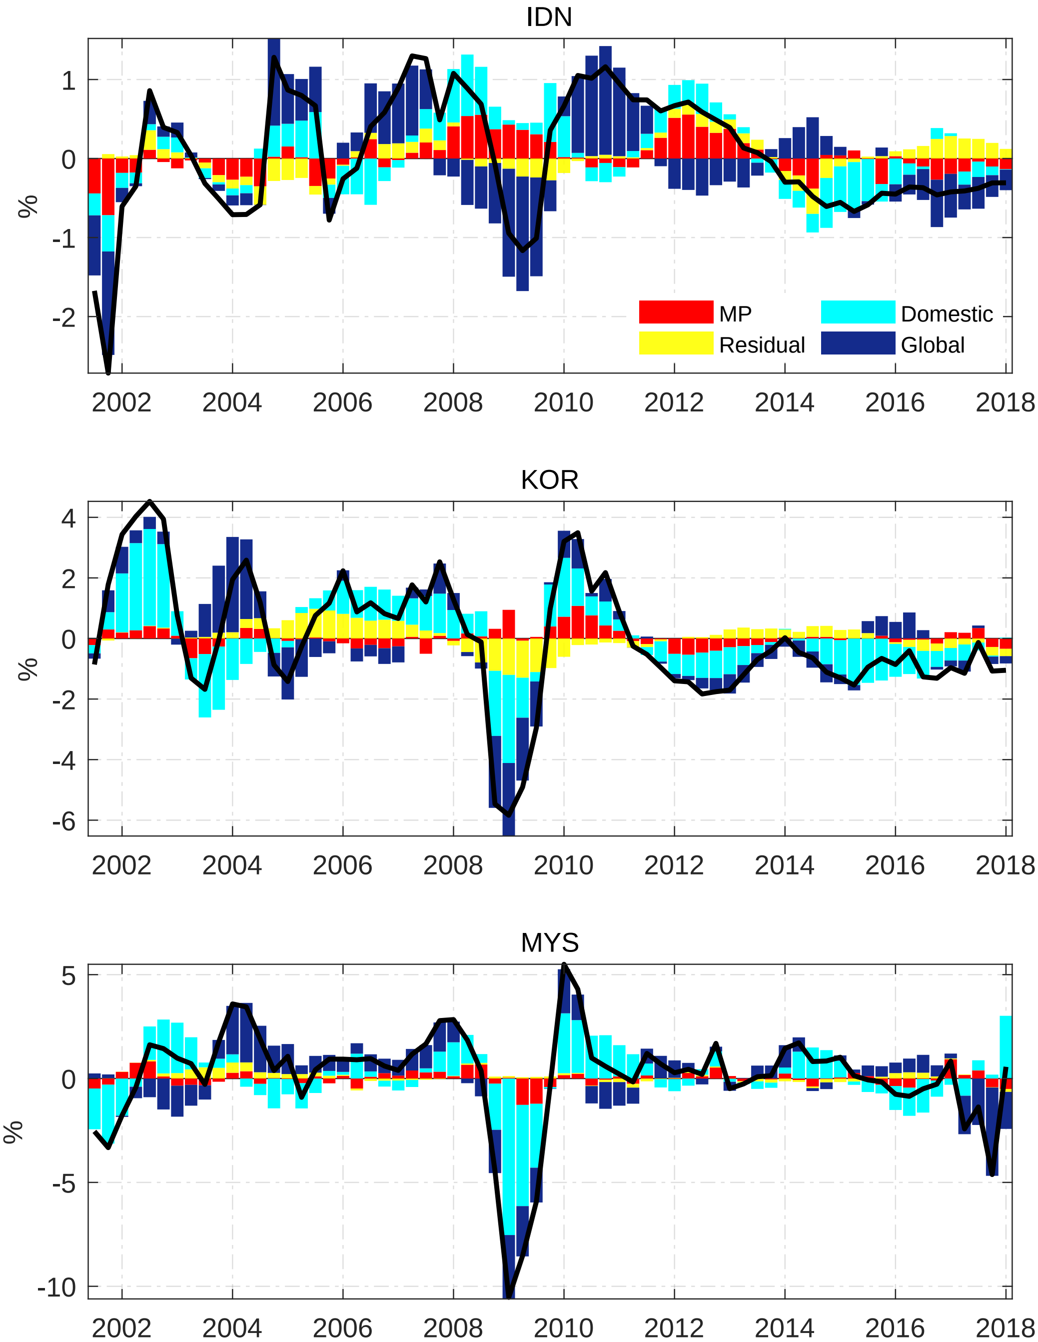

Notes: Median historical contribution of monetary policy shocks (red bars), residual shocks (yellow bars), domestic shocks (teal bars) and global shocks (blue bars) to the growth rate of real GDP for Indonesia, Korea and Malaysia. The black path corresponds to the sum of median contributions of all structural shocks. Results rely on the alternative identification strategy.

Figure 30: Historical contribution of structural shocks to real GDP growth for the Philippines, Singapore and Thailand

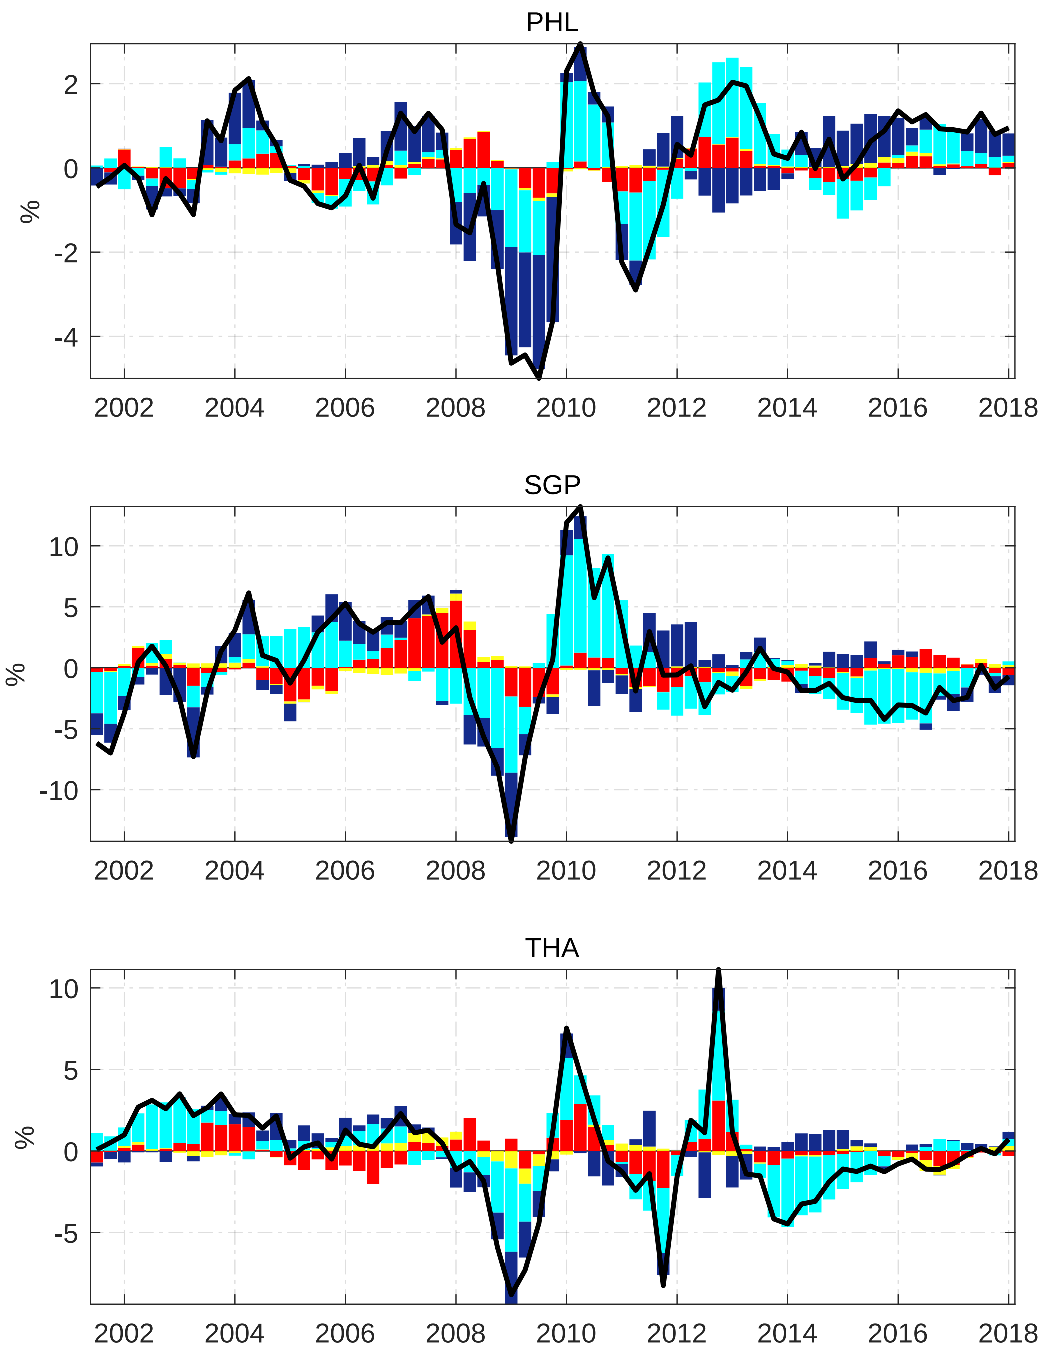

Notes: Median historical contribution of monetary policy shocks (red bars), residual shocks (yellow bars), domestic shocks (teal bars) and global shocks (blue bars) to the growth rate of real GDP for the Philippines, Singapore and Thailand. The black path corresponds to the sum of median contributions of all structural shocks. Results rely on the alternative identification strategy.

b. *Counterfactual Simulations*

Figure 31: Counterfactual paths for inflation with suppressed shocks – Indonesia

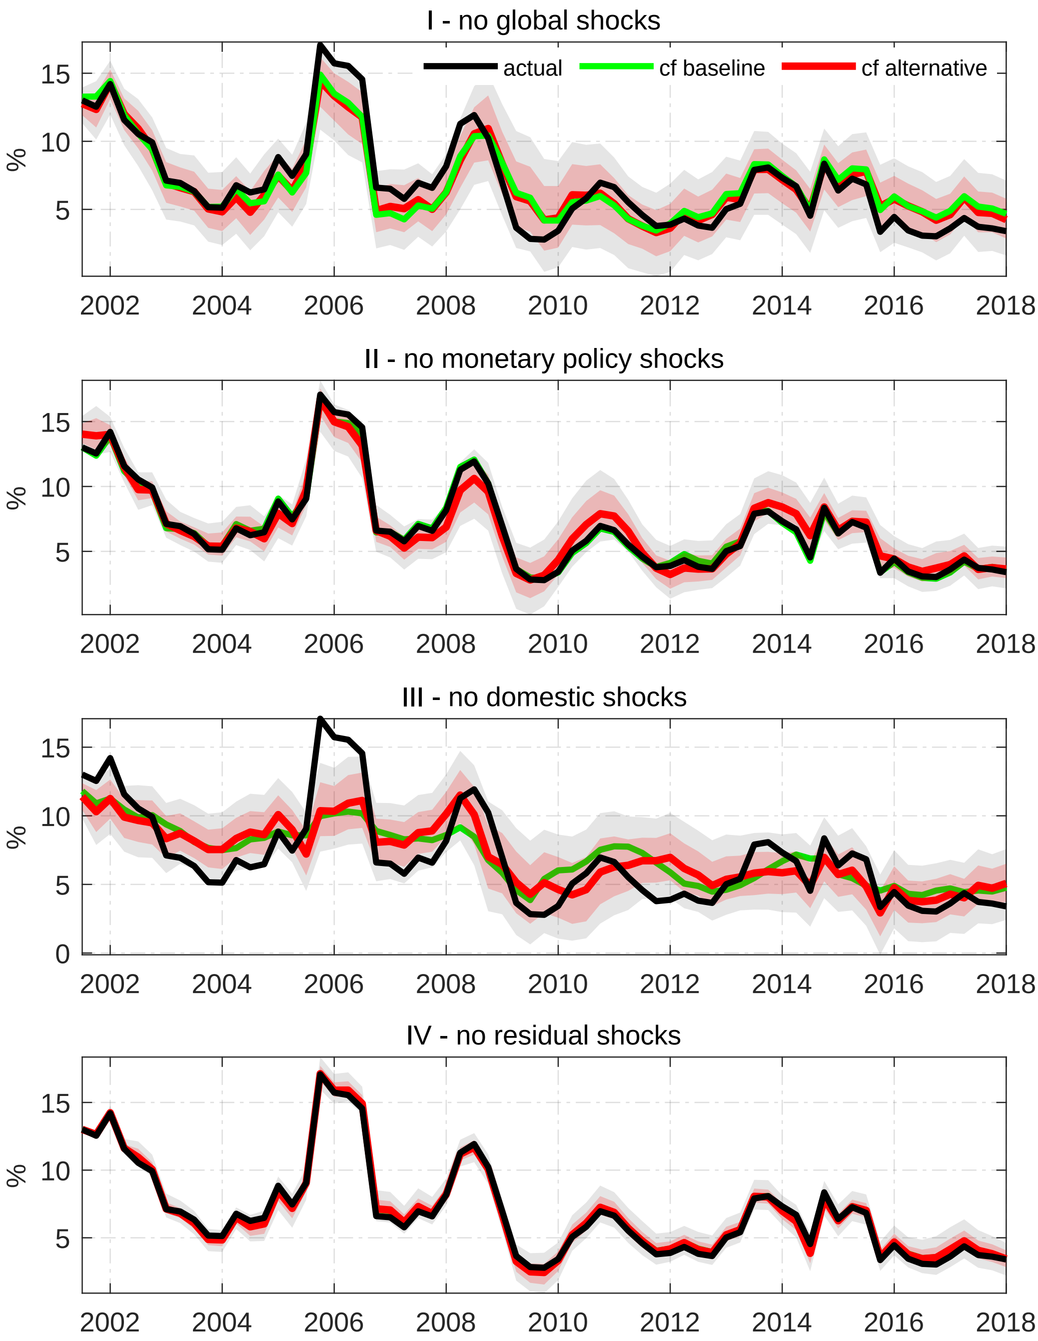

Notes: Median counterfactual paths (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area), and 5^{th} and 95^{th} percentiles (grey-shaded area) for the inflation rate (in %). In I, the counterfactual path corresponds to the inflation rate where aggregate global shocks (oil supply shocks and global demand shocks) are suppressed, while the same is done in II with monetary policy shocks, in III with domestic shocks (domestic demand and domestic supply shocks) and in IV (residual shocks). The identification of shocks relies on the alternative identification strategy as in Table (9). The green solid path shows the median counterfactual path for the baseline identification strategy.

Figure 32: Counterfactual paths for real GDP growth with suppressed shocks – Indonesia

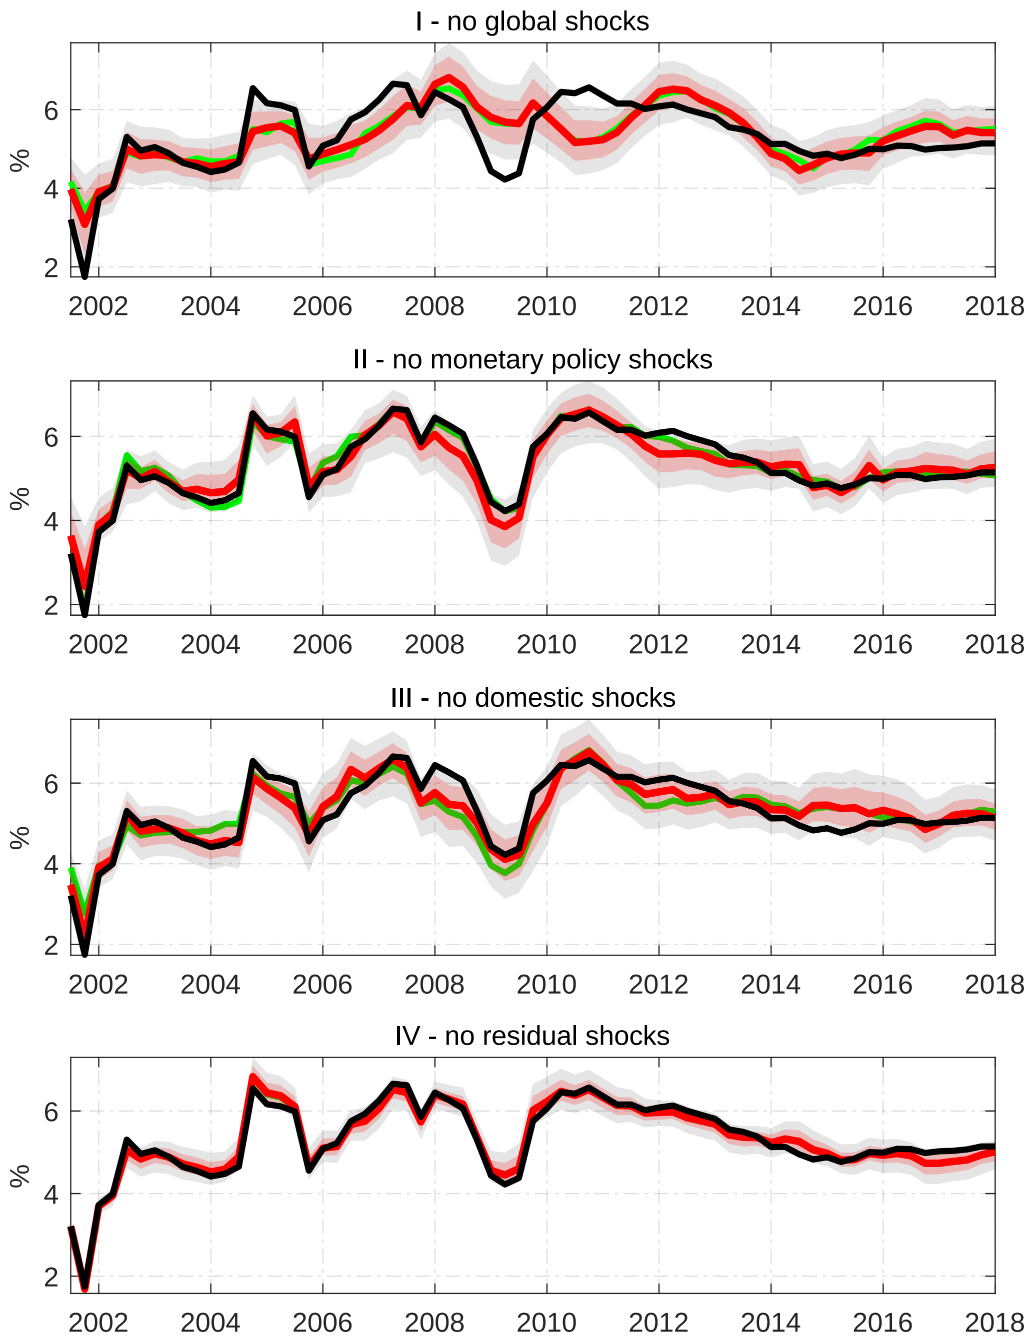

Notes: Median counterfactual paths (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area), and 5^{th} and 95^{th} percentiles (grey-shaded area) for the growth rate of real GDP (in %). In I, the counterfactual path corresponds to real GDP growth rate where aggregate global shocks (oil supply shocks and global demand shocks) are suppressed, while the same is done in II with monetary policy shocks, in III with domestic shocks (domestic demand and domestic supply shocks) and in IV (residual shocks). The identification of shocks relies on the alternative identification strategy as in Table (9). The green solid path shows the median counterfactual path for the baseline identification strategy.

Figure 33: Counterfactual paths for inflation with suppressed shocks – Korea

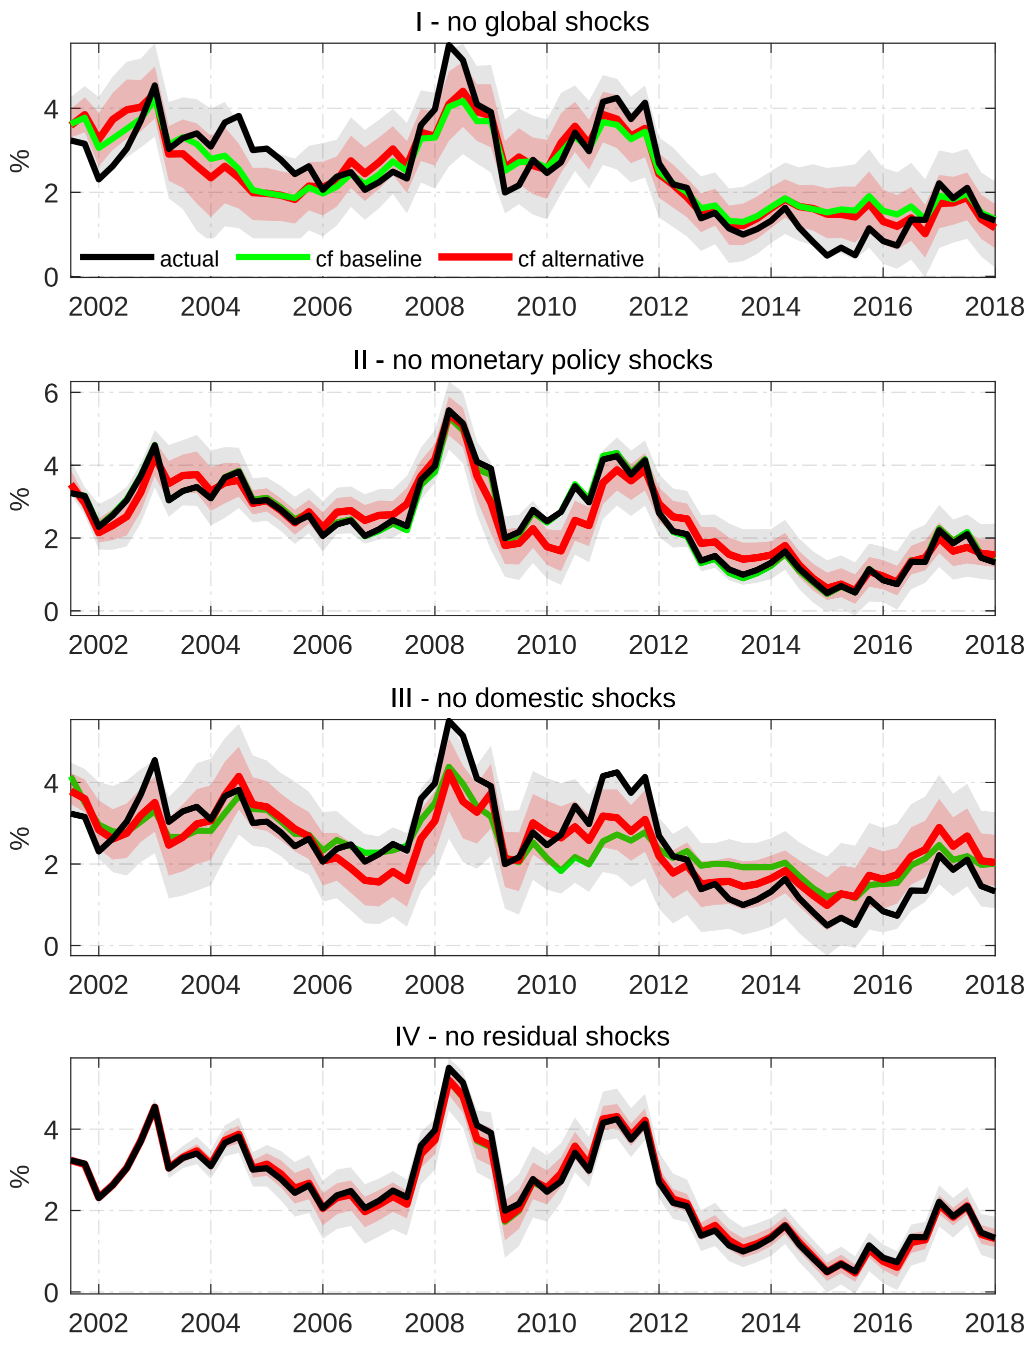

Notes: Median counterfactual paths (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area), and 5^{th} and 95^{th} percentiles (grey-shaded area) for the inflation rate (in %). In I, the counterfactual path corresponds to the inflation rate where aggregate global shocks (oil supply shocks and global demand shocks) are suppressed, while the same is done in II with monetary policy shocks, in III with domestic shocks (domestic demand and domestic supply shocks) and in IV (residual shocks). The identification of shocks relies on the alternative identification strategy as in Table (9). The green solid path shows the median counterfactual path for the baseline identification strategy.

Figure 34: Counterfactual paths for real GDP growth with suppressed shocks – Korea

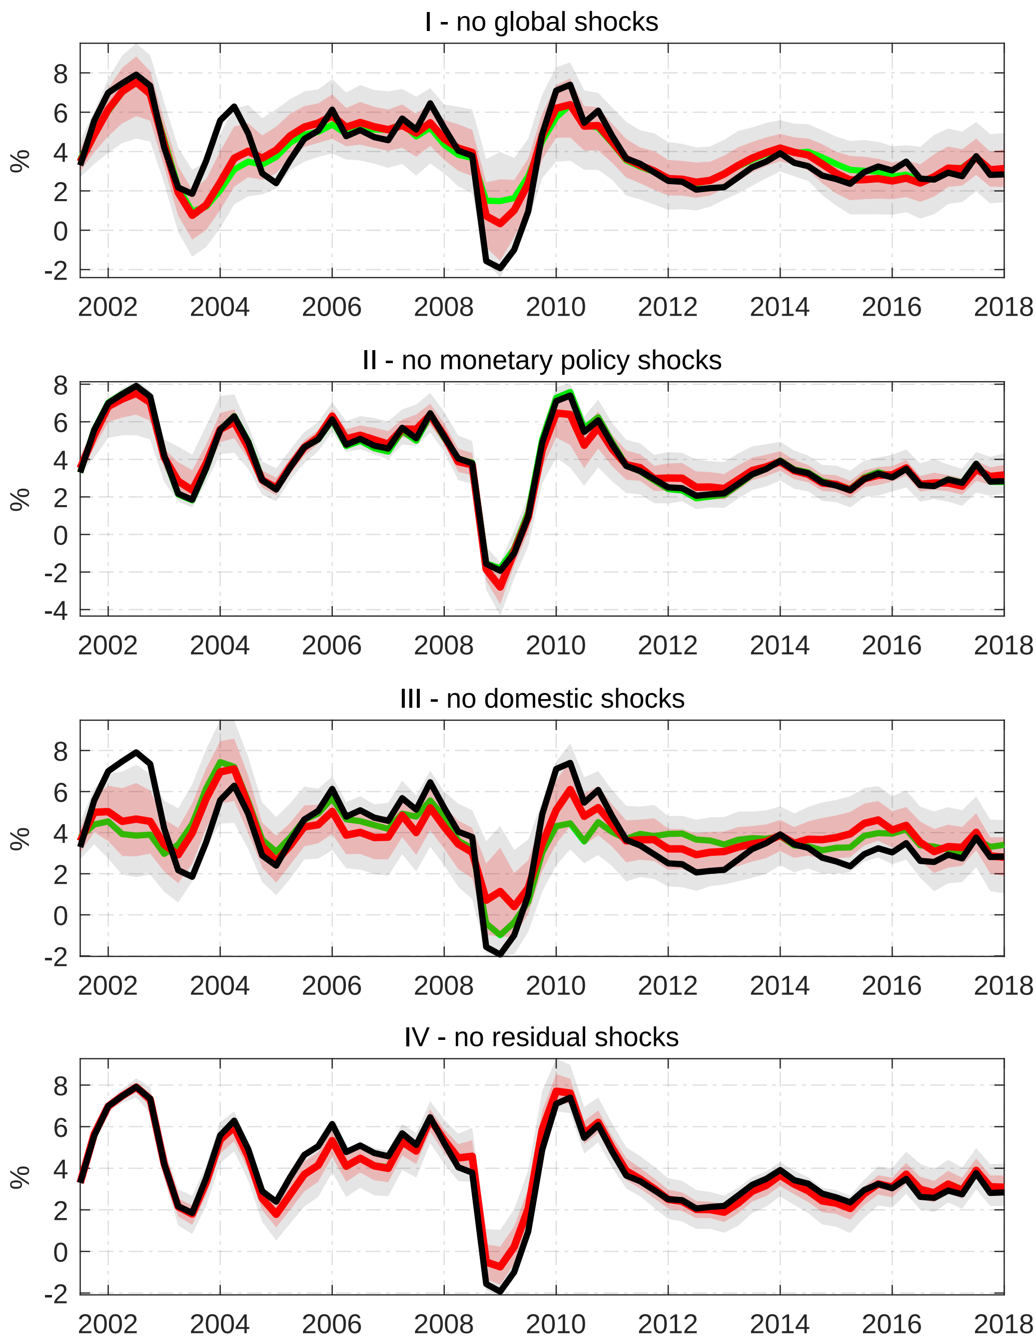

Notes: Median counterfactual paths (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area), and 5^{th} and 95^{th} percentiles (grey-shaded area) for the growth rate of real GDP (in %). In I, the counterfactual path corresponds to real GDP growth rate where aggregate global shocks (oil supply shocks and global demand shocks) are suppressed, while the same is done in II with monetary policy shocks, in III with domestic shocks (domestic demand and domestic supply shocks) and in IV (residual shocks). The identification of shocks relies on the alternative identification strategy as in Table (9). The green solid path shows the median counterfactual path for the baseline identification strategy.

Figure 35: Counterfactual paths for inflation with suppressed shocks – Malaysia

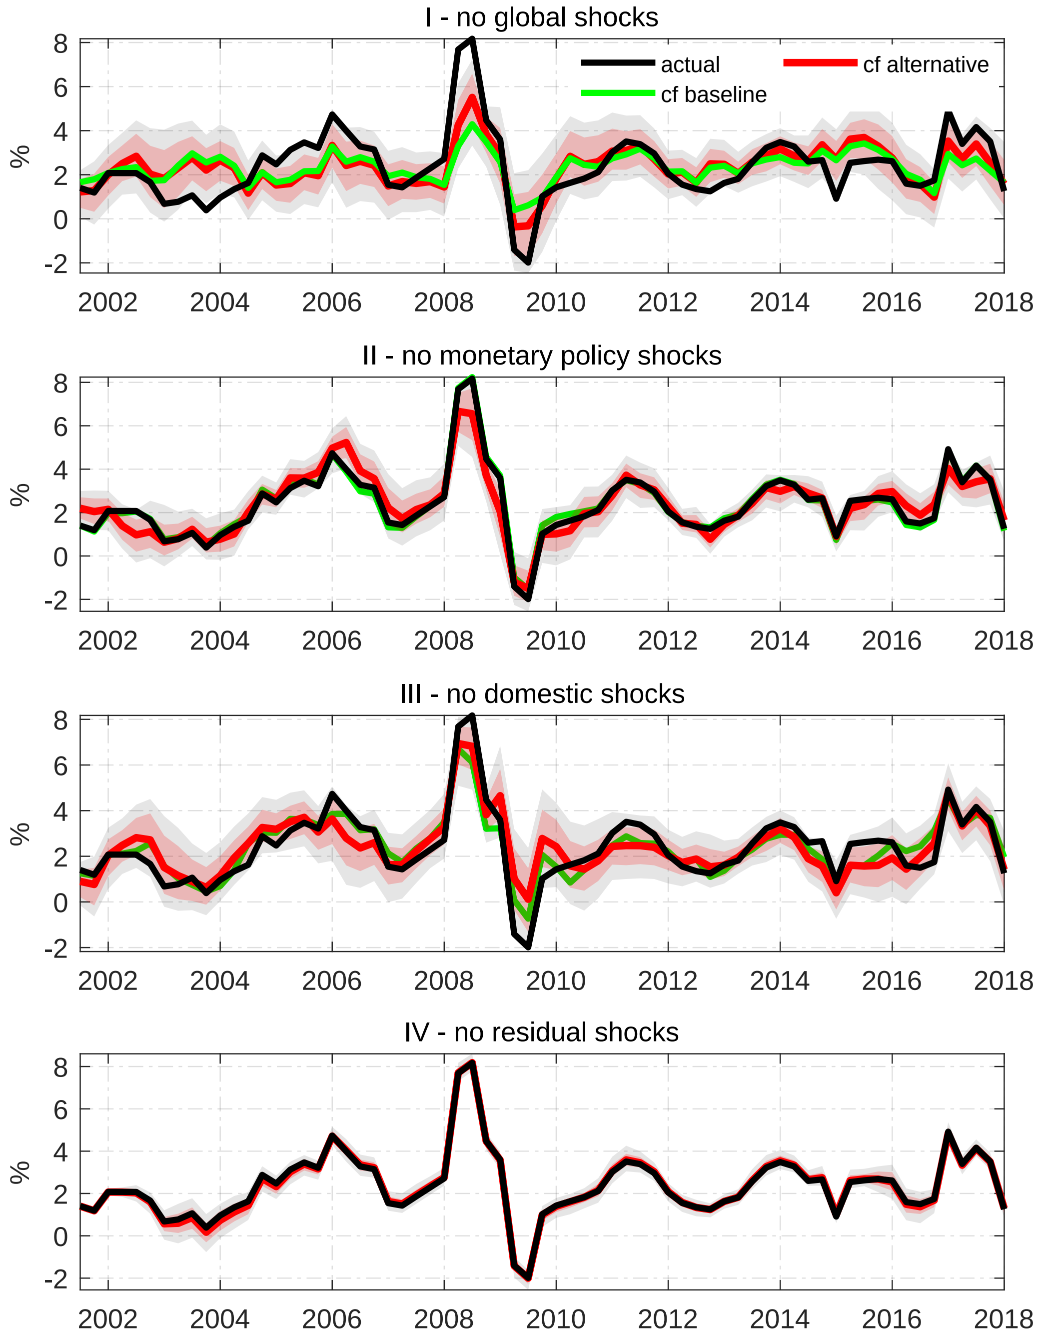

Notes: Median counterfactual paths (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area), and 5^{th} and 95^{th} percentiles (grey-shaded area) for the inflation rate (in %). In I, the counterfactual path corresponds to the inflation rate where aggregate global shocks (oil supply shocks and global demand shocks) are suppressed, while the same is done in II with monetary policy shocks, in III with domestic shocks (domestic demand and domestic supply shocks) and in IV (residual shocks). The identification of shocks relies on the alternative identification strategy as in Table (9). The green solid path shows the median counterfactual path for the baseline identification strategy.

Figure 36: Counterfactual paths for real GDP growth with suppressed shocks – Malaysia

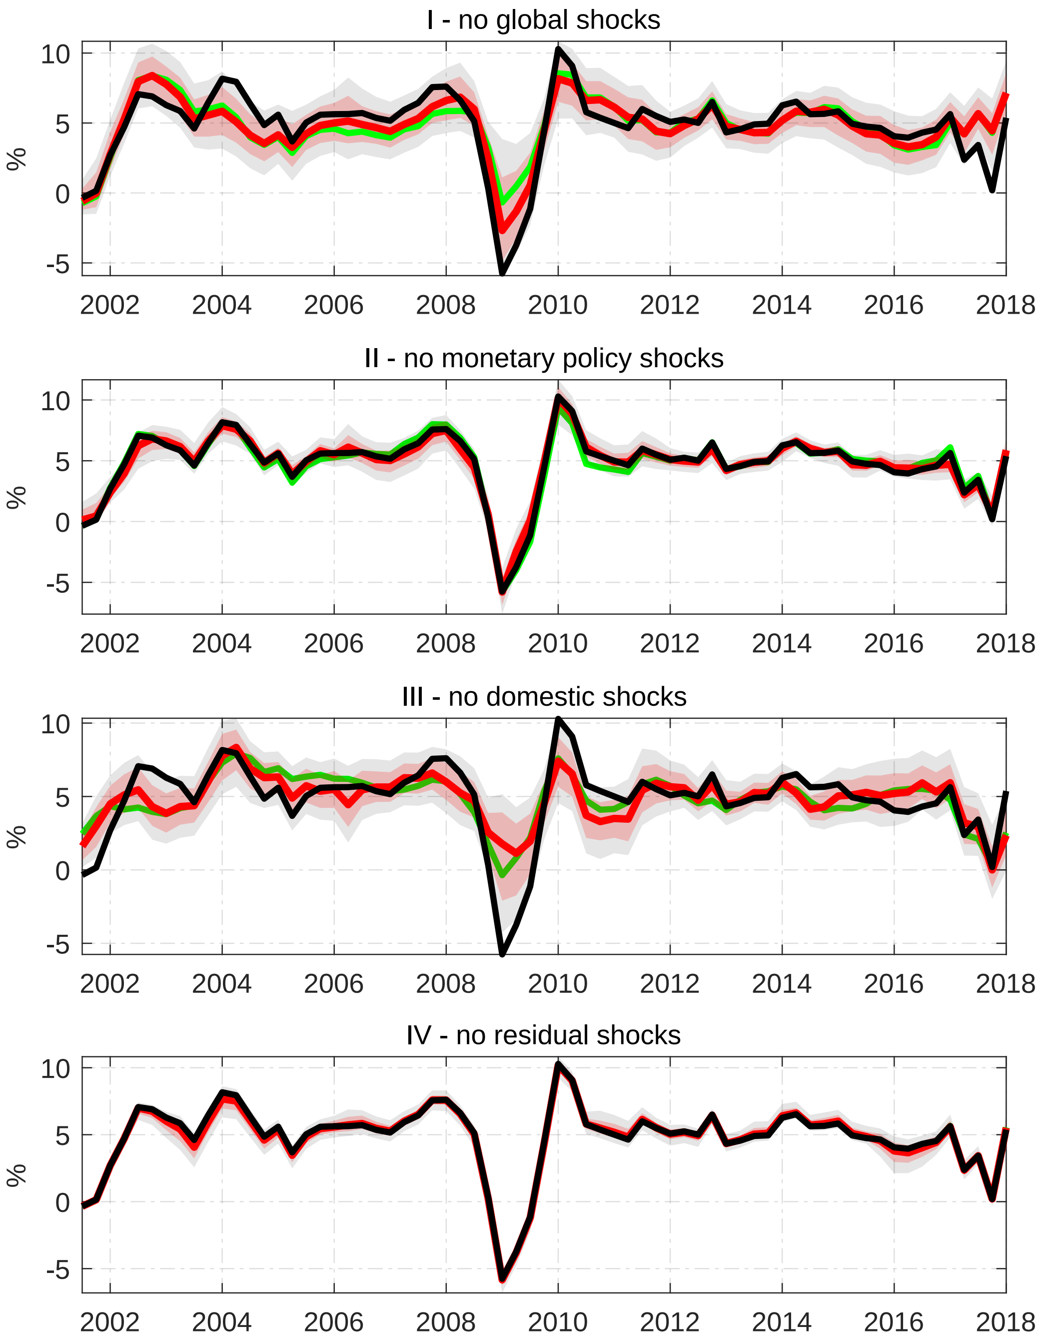

Notes: Median counterfactual paths (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area), and 5^{th} and 95^{th} percentiles (grey-shaded area) for the growth rate of real GDP (in %). In I, the counterfactual path corresponds to real GDP growth rate where aggregate global shocks (oil supply shocks and global demand shocks) are suppressed, while the same is done in II with monetary policy shocks, in III with domestic shocks (domestic demand and domestic supply shocks) and in IV (residual shocks). The identification of shocks relies on the alternative identification strategy as in Table (9). The green solid path shows the median counterfactual path for the baseline identification strategy.

Figure 37: Counterfactual paths for inflation with suppressed shocks – Philippines

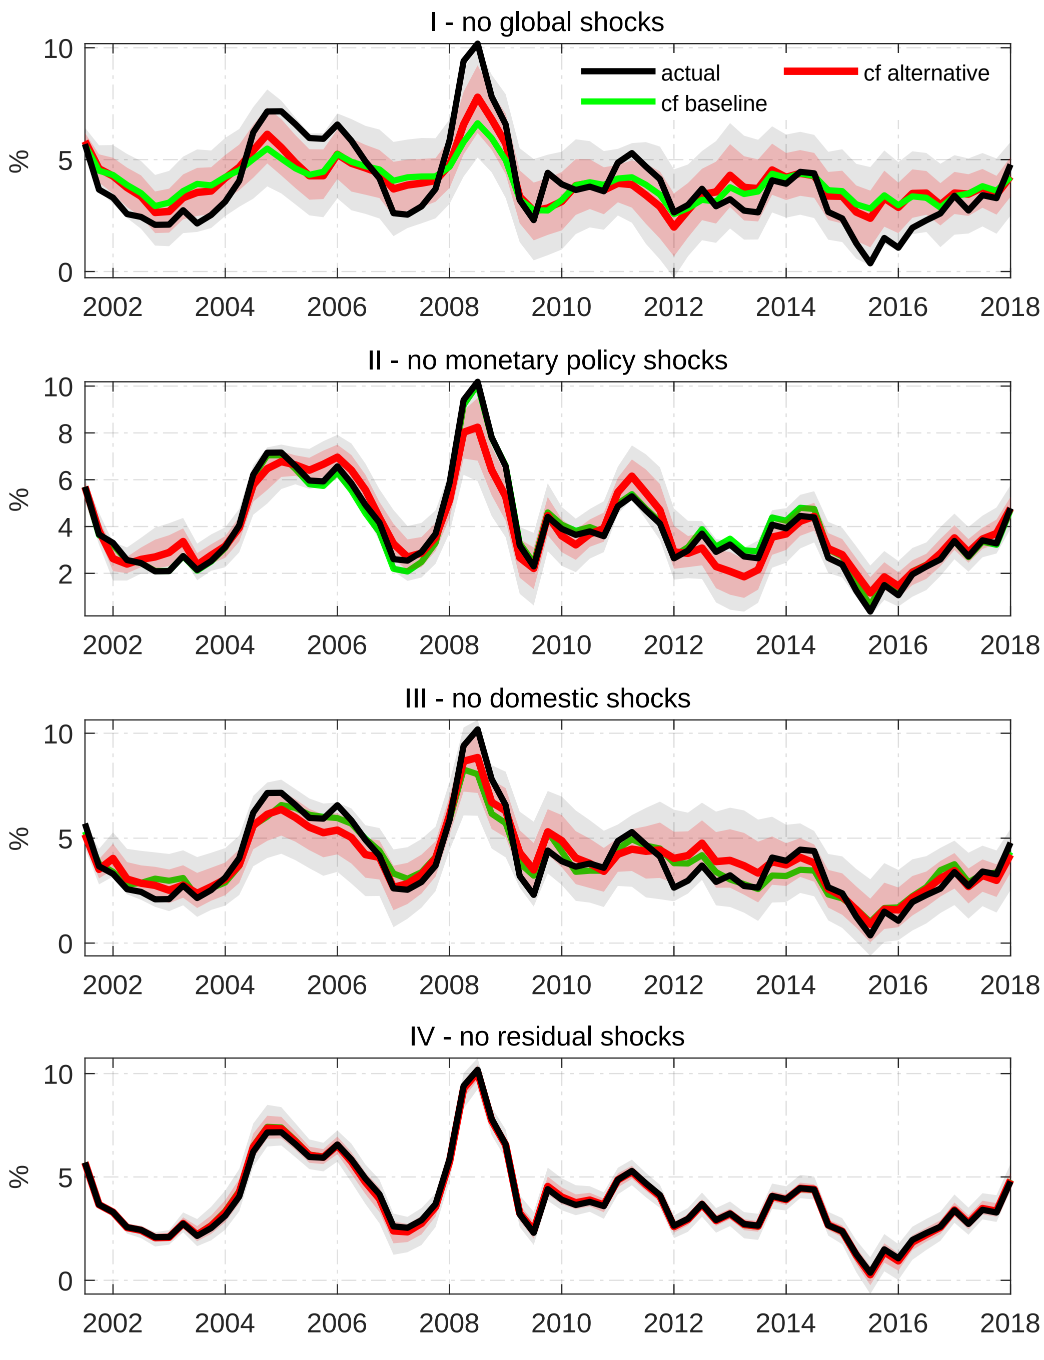

Notes: Median counterfactual paths (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area), and 5^{th} and 95^{th} percentiles (grey-shaded area) for the inflation rate (in %). In I, the counterfactual path corresponds to the inflation rate where aggregate global shocks (oil supply shocks and global demand shocks) are suppressed, while the same is done in II with monetary policy shocks, in III with domestic shocks (domestic demand and domestic supply shocks) and in IV (residual shocks). The identification of shocks relies on the alternative identification strategy as in Table (9). The green solid path shows the median counterfactual path for the baseline identification strategy.

Figure 38: Counterfactual paths for real GDP growth with suppressed shocks - Philippines

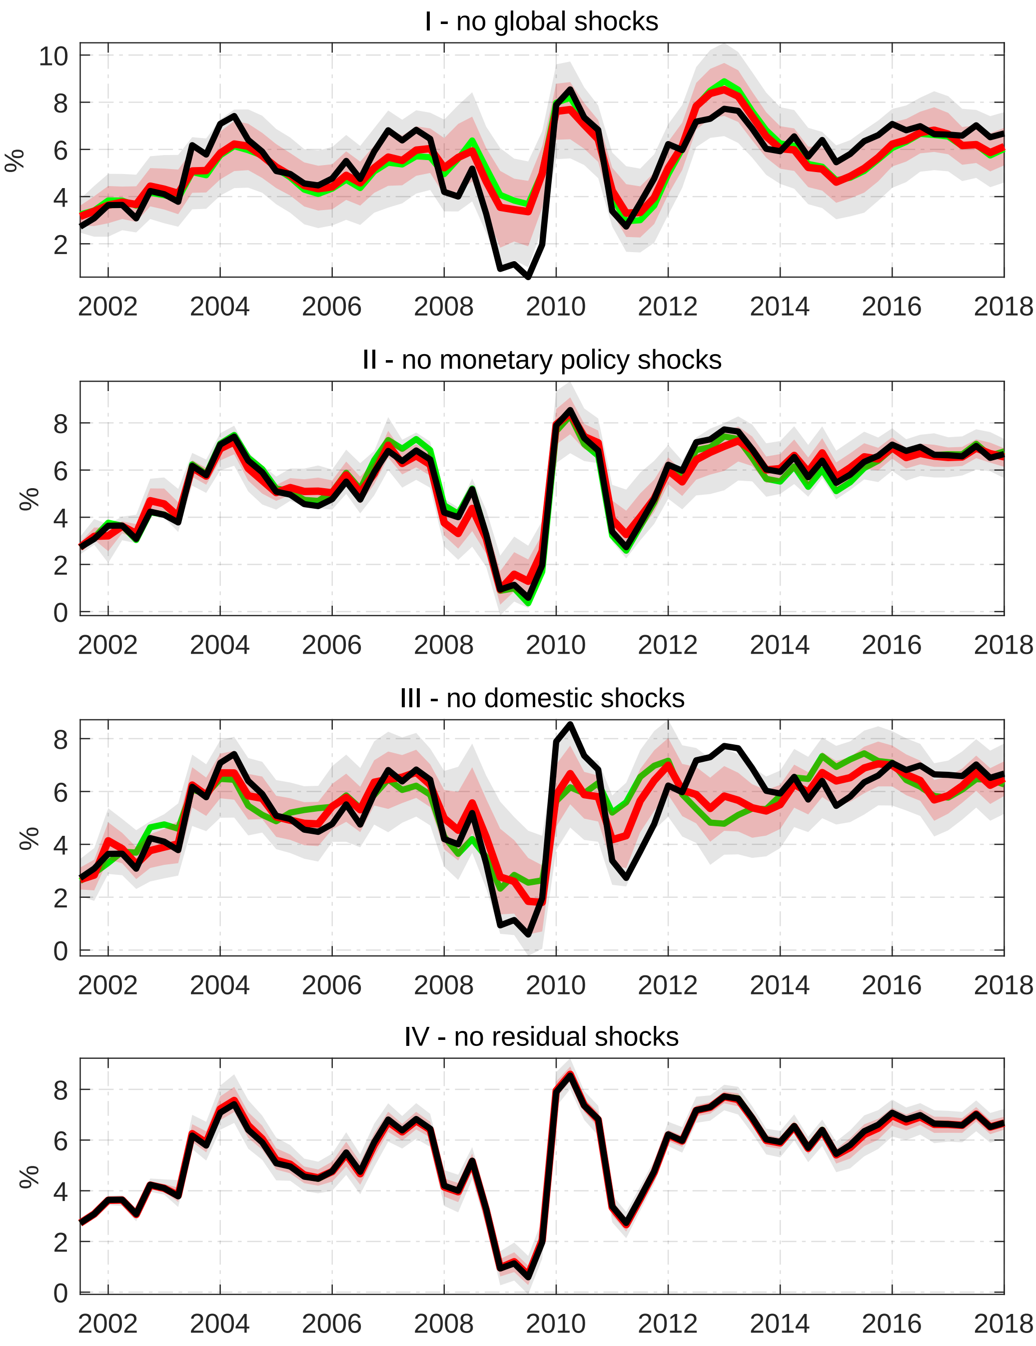

Notes: Median counterfactual paths (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area), and 5^{th} and 95^{th} percentiles (grey-shaded area) for the growth rate of real GDP (in %). In I, the counterfactual path corresponds to real GDP growth rate where aggregate global shocks (oil supply shocks and global demand shocks) are suppressed, while the same is done in II with monetary policy shocks, in III with domestic shocks (domestic demand and domestic supply shocks) and in IV (residual shocks). The identification of shocks relies on the alternative identification strategy as in Table (9). The green solid path shows the median counterfactual path for the baseline identification strategy.

Figure 39: Counterfactual paths for inflation with suppressed shocks – Singapore

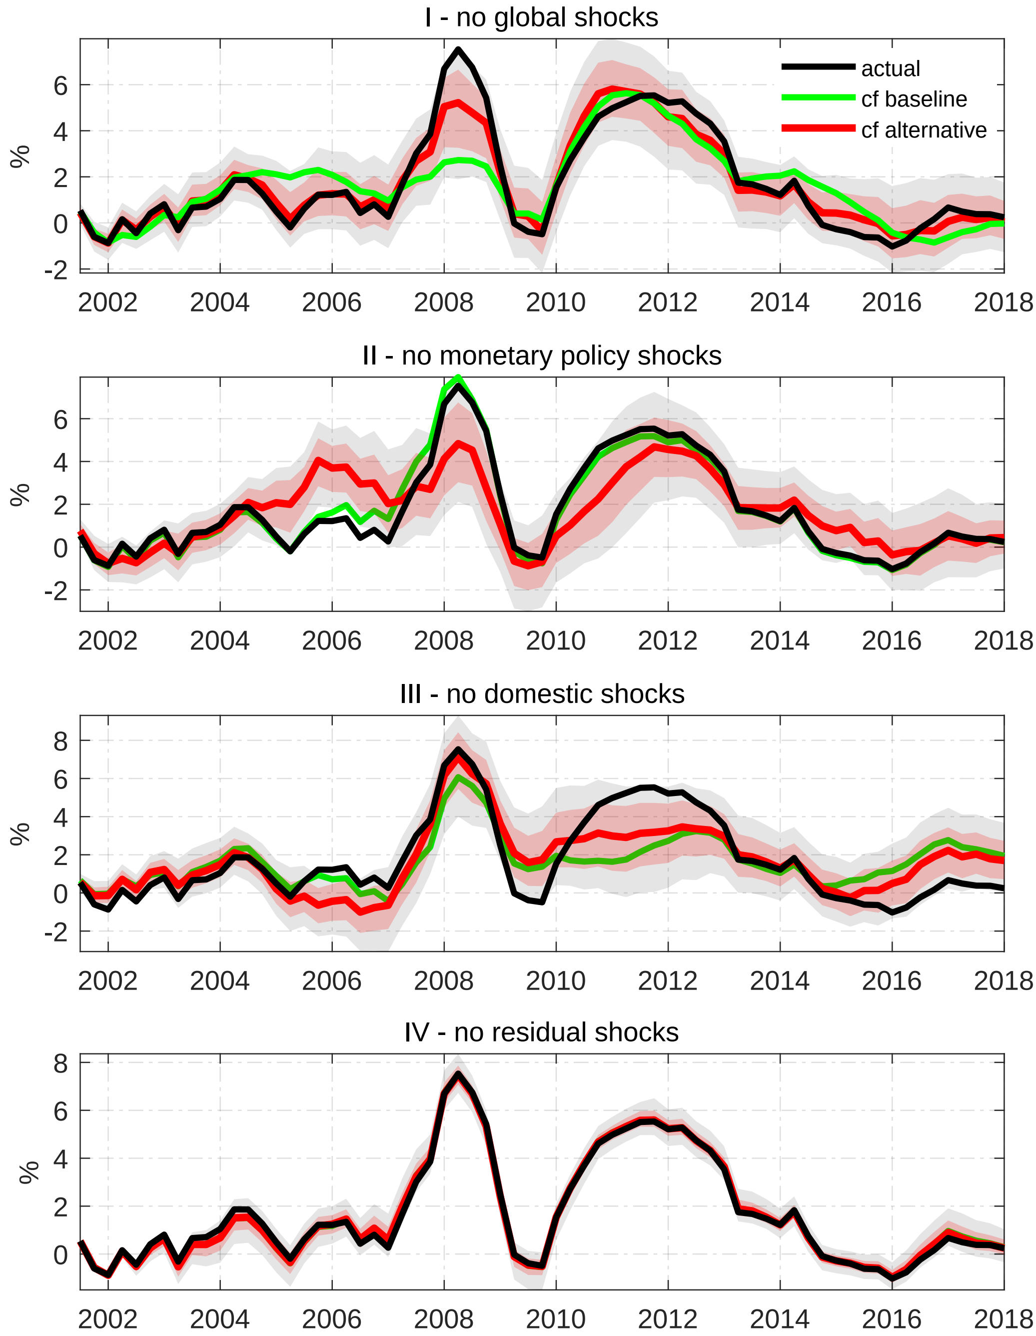

Notes: Median counterfactual paths (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area), and 5^{th} and 95^{th} percentiles (grey-shaded area) for the inflation rate (in %). In I, the counterfactual path corresponds to the inflation rate where aggregate global shocks (oil supply shocks and global demand shocks) are suppressed, while the same is done in II with monetary policy shocks, in III with domestic shocks (domestic demand and domestic supply shocks) and in IV (residual shocks). The identification of shocks relies on the alternative identification strategy as in Table (9). The green solid path shows the median counterfactual path for the baseline identification strategy.

Figure 40: Counterfactual paths for real GDP growth with suppressed shocks – Singapore

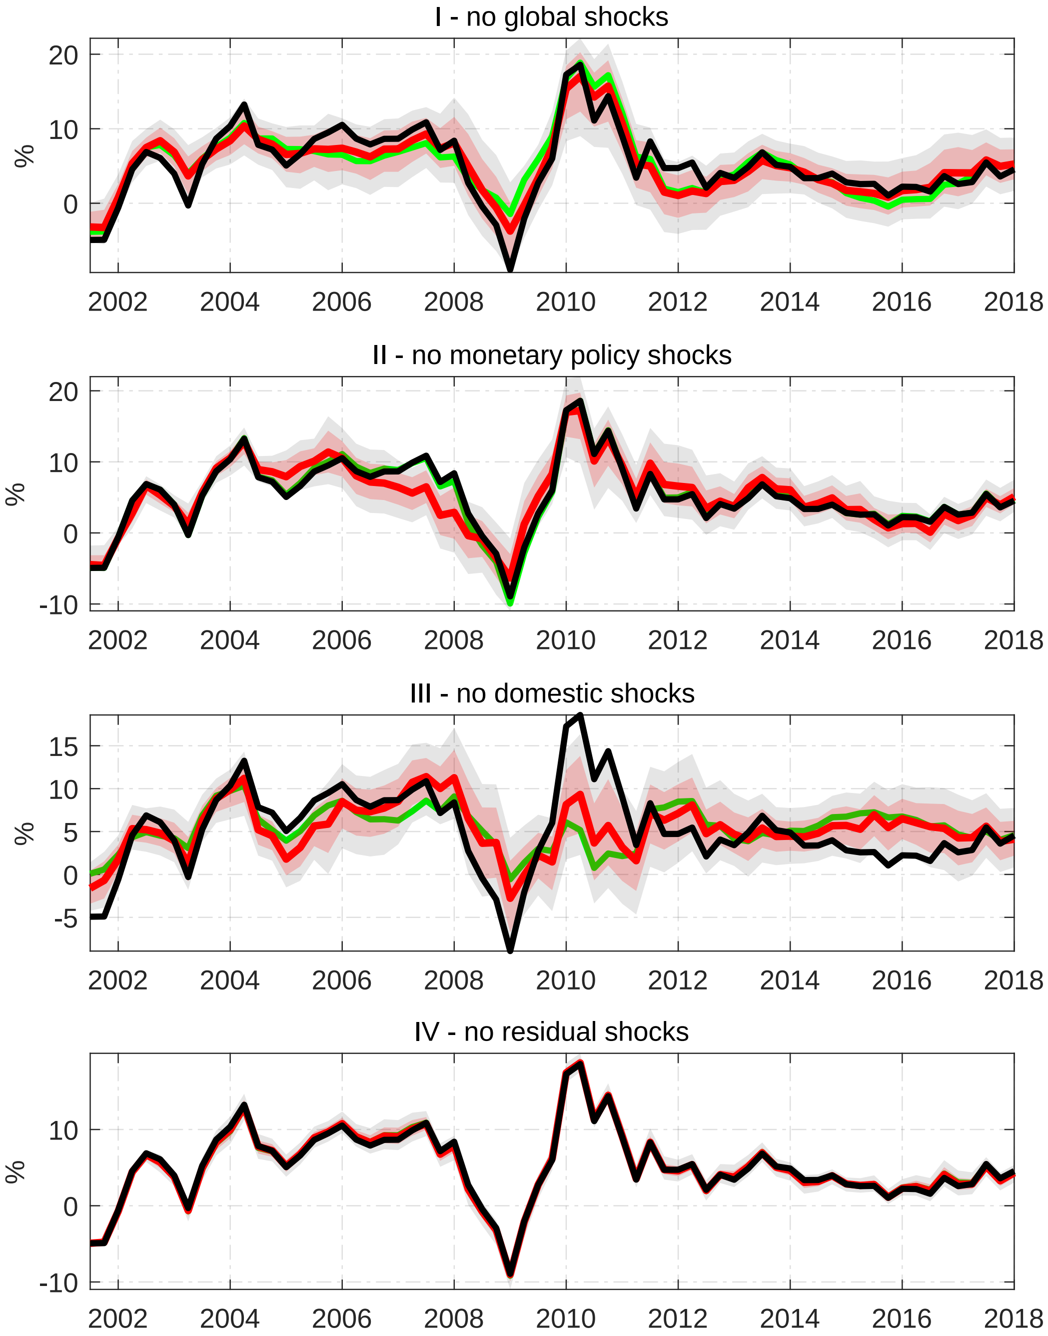

Notes: Median counterfactual paths (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area), and 5^{th} and 95^{th} percentiles (grey-shaded area) for the growth rate of real GDP (in %). In I, the counterfactual path corresponds to real GDP growth rate where aggregate global shocks (oil supply shocks and global demand shocks) are suppressed, while the same is done in II with monetary policy shocks, in III with domestic shocks (domestic demand and domestic supply shocks) and in IV (residual shocks). The identification of shocks relies on the alternative identification strategy as in Table (9). The green solid path shows the median counterfactual path for the baseline identification strategy.

Figure 41: Counterfactual paths for inflation with suppressed shocks – Thailand

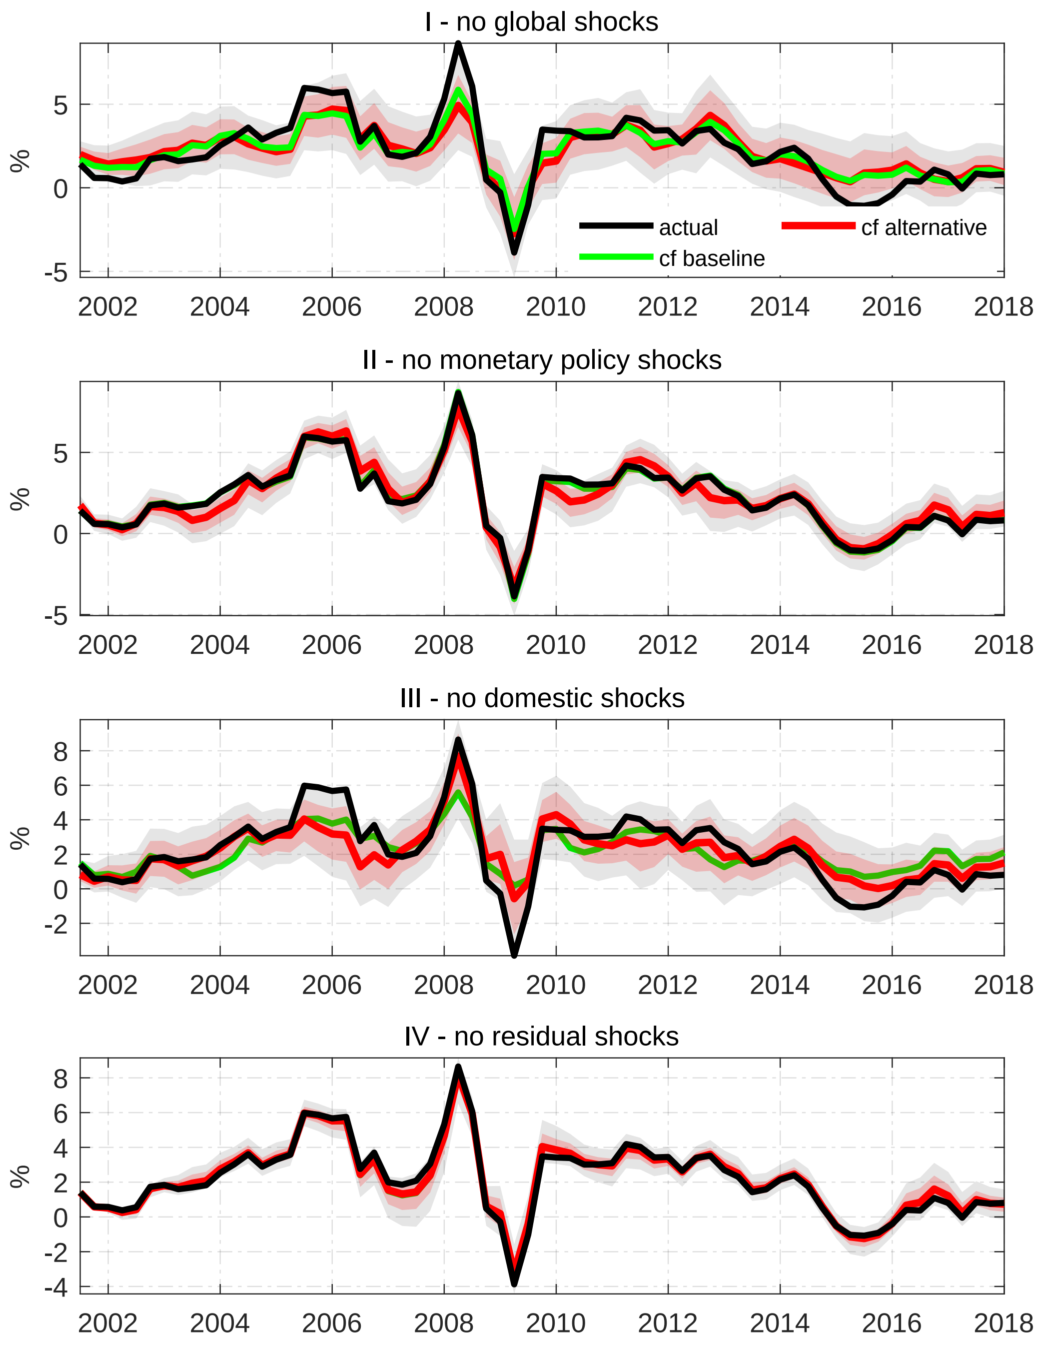

Notes: Median counterfactual paths (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area), and 5^{th} and 95^{th} percentiles (grey-shaded area) for the inflation rate (in %). In I, the counterfactual path corresponds to the inflation rate where aggregate global shocks (oil supply shocks and global demand shocks) are suppressed, while the same is done in II with monetary policy shocks, in III with domestic shocks (domestic demand and domestic supply shocks) and in IV (residual shocks). The identification of shocks relies on the alternative identification strategy as in Table (9). The green solid path shows the median counterfactual path for the baseline identification strategy.

Figure 42: Counterfactual paths for real GDP growth with suppressed shocks – Thailand

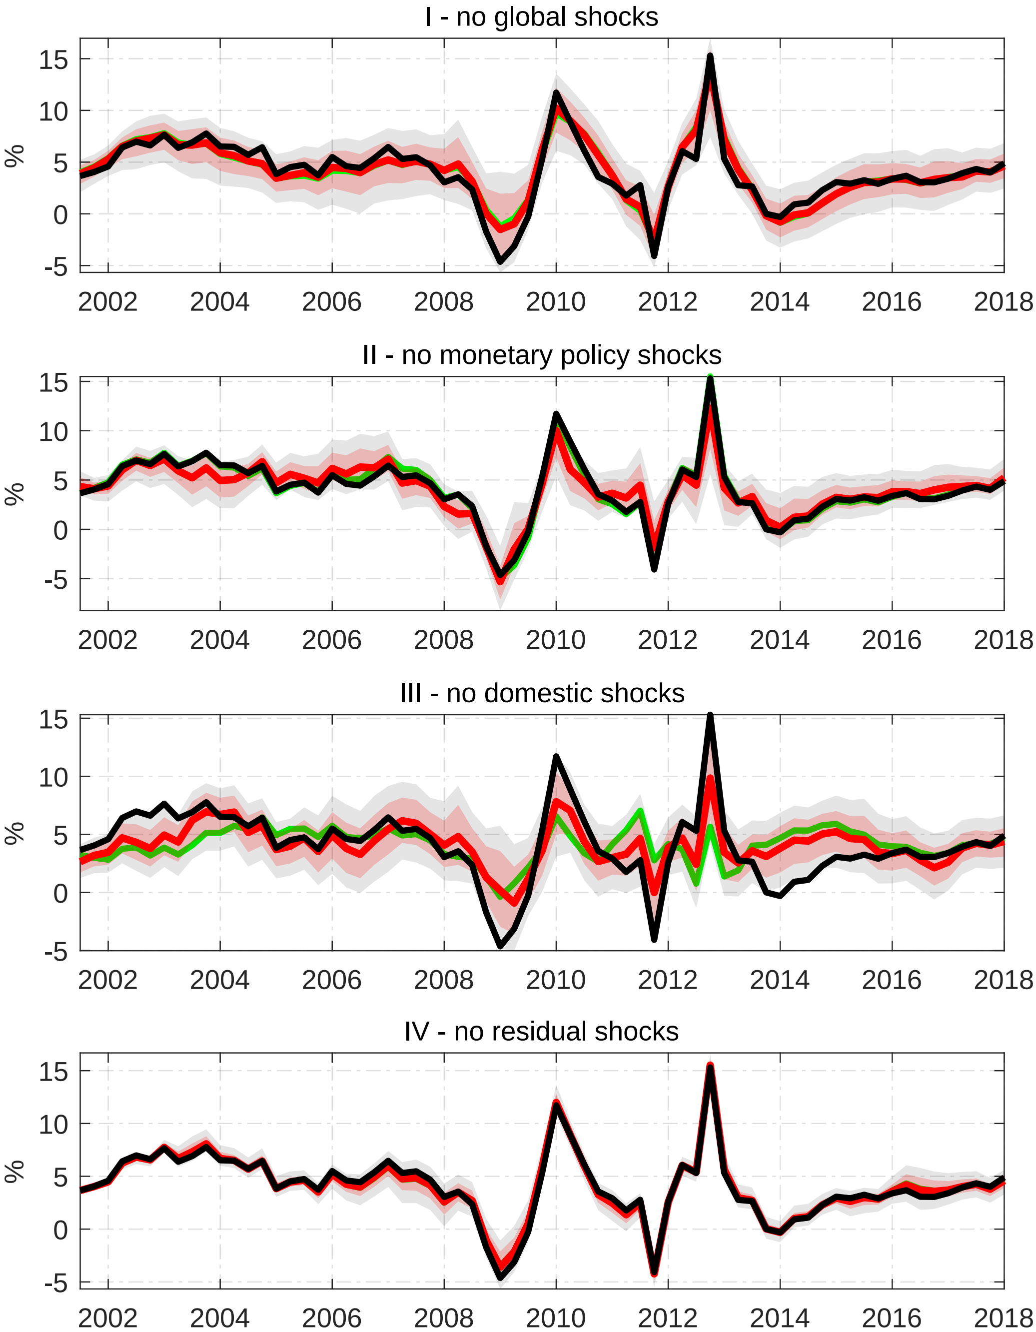

Notes: Median counterfactual paths (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area), and 5^{th} and 95^{th} percentiles (grey-shaded area) for the growth rate of real GDP (in %). In I, the counterfactual path corresponds to real GDP growth rate where aggregate global shocks (oil supply shocks and global demand shocks) are suppressed, while the same is done in II with monetary policy shocks, in III with domestic shocks (domestic demand and domestic supply shocks) and in IV (residual shocks). The identification of shocks relies on the alternative identification strategy as in Table (9). The green solid path shows the median counterfactual path for the baseline identification strategy.

c. *Forecast Error Variance Decompositions*

In order to save space, we show the results of the forecast variance decomposition for the alternative identification strategy in two tables only, one for the growth rate of real GDP and one for inflation. Overall, we again find that global factors are important for all countries.

Table 11: Forecast error variance decompositions for real GDP growth

		IDN	KOR	MYS	PHL	SGP	THA
$h = 4$	*residual*	6	4	1	1	1	1
	monetary policy	22	7	6	6	13	22
	oil supply	22	16	23	40	26	15
	global demand	29	20	21	14	14	17
	domestic demand	4	42	37	13	24	28
	domestic supply	18	12	12	26	23	17
$h = 8$	*residual*	9	7	2	2	1	2
	monetary policy	21	8	8	7	17	20
	oil supply	19	18	25	39	24	17
	global demand	27	20	19	14	14	17
	domestic demand	6	34	35	14	21	26
	domestic supply	19	13	12	24	23	18
$h = 12$	*residual*	9	7	2	2	2	3
	monetary policy	21	8	8	8	19	20
	oil supply	18	18	23	38	23	17
	global demand	26	22	19	14	14	17
	domestic demand	7	33	35	16	21	25
	domestic supply	18	13	12	23	22	18

Notes: Median shares (in %) of forecast error variance for the growth rate of real GDP due to structural shocks for different forecast horizons. All values are rounded to integers. Results rely on the alternative identification strategy.

However, we also find that for all countries, both for real GDP growth rate and especially the inflation rate, the explanatory power of monetary policy shocks for the variance of the forecast error is more important than under the baseline identification strategy. This is particularly true for the inflation rate in Singapore.

Table 12: Forecast error variance decompositions for inflation

		IDN	KOR	MYS	PHL	SGP	THA
$h = 4$	*residual*	2	1	1	1	1	1
	monetary policy	8	27	23	15	56	16
	oil supply	11	18	29	46	7	29
	global demand	16	22	26	17	13	25
	domestic demand	20	21	12	10	13	19
	domestic supply	44	11	10	11	9	10
$h = 8$	*residual*	2	2	2	2	2	4
	monetary policy	9	22	21	13	44	16
	oil supply	12	15	26	43	9	27
	global demand	23	22	27	17	12	23
	domestic demand	19	28	13	13	15	20
	domestic supply	34	12	11	13	18	11
$h = 12$	*residual*	3	3	2	2	3	4
	monetary policy	11	21	21	13	39	16
	oil supply	12	15	25	41	10	26
	global demand	24	21	26	17	12	22
	domestic demand	19	29	14	14	15	20
	domestic supply	30	12	12	13	22	12

Notes: Median shares (in %) of forecast error variance for inflation due to structural shocks for different forecast horizons. All values are rounded to integers. Results rely on the alternative identification strategy.

d. *Regression Results*

Table 13: Regression results for the alternative identification strategy (model B)

	IDN	KOR	MYS	PHL	SGP	THA
oil supply shock						
c	0.49	−0.10	−0.06	0.11	−0.62	−0.02
	[−0.58, 0.74]	[−1.27, −0.04]	[−0.60, 0.21]	[−1.18, 0.82]	[−5.79, −3.71]	[−1.41, 0.81]
β_{oil}	−1.02	−0.01	−0.05	−0.32	0.06	−0.24
	[−3.56, −0.18]	[−1.88, 1.46]	[−2.68, 1.47]	[−4.69, −0.52]	[−0.16, 3.58]	[−4.48, −0.02]
$\beta_{cf(oil)}$	0.11	0.14	0.18	0.14	0.14	0.15
	[−0.89, 1.32]	[3.41, 4.89]	[2.80, 4.55]	[0.73, 3.57]	[6.89, 9.21]	[2.40, 4.88]
γ	0.83	0.82	0.66	0.80	0.95	0.71
	[11.27, 13.05]	[12.71, 14.87]	[6.05, 7.85]	[9.92, 12.93]	[22.83, 27.60]	[8.68, 11.21]
# obs	67	67	67	67	67	67
R^2	0.72	0.77	0.51	0.74	0.91	0.68
global demand shock						
c	2.50	−0.03	0.05	0.53	−0.55	0.13
	[0.81, 2.31]	[−0.49, 0.56]	[−0.19, 0.91]	[0.52, 1.76]	[−4.32, −3.15]	[−0.07, 1.77]
β_{dem}	0.52	0.21	0.33	0.31	0.22	0.40
	[0.45, 2.42]	[2.11, 4.22]	[1.86, 4.08]	[0.48, 3.46]	[2.66, 7.39]	[1.65, 5.37]
$\beta_{cf(dem)}$	−0.22	0.11	0.14	0.03	0.12	0.11
	[−1.69, −0.19]	[2.13, 3.55]	[1.33, 2.94]	[−0.54, 0.63]	[6.19, 7.82]	[1.33, 2.78]
γ	0.80	0.82	0.70	0.83	0.94	0.72
	[9.98, 11.98]	[13.46, 14.82]	[7.35, 8.13]	[10.19, 11.77]	[23.22, 26.82]	[8.47, 10.46]
# obs	67	67	67	67	67	67
R^2	0.71	0.77	0.51	0.70	0.91	0.66

Notes: Estimation results for the regression of model B in Section IV in the main paper. The estimated coefficients (as well as for the coefficient of determination R^2) are median value over all 2000 draws. The lagged domestic inflation rate enters with the coefficient γ. In the first block, β_{oil} represents the coefficient on the component of real GDP growth driven by oil supply shocks, while in the second block, β_{dem} represents the coefficient on the component of real GDP growth driven by global demand shocks. In both blocks, $\beta_{cf(j)}$ represents the coefficient on the counterfactual path of real GDP growth that excludes either the component driven by oil supply shocks (first block) or the component driven by global demand shocks (second block). For all coefficients in both blocks, the values in square brackets correspond to the 16^{th} and 84^{th} percentile of the $t-$statistic testing the null that the corresponding coefficient is equal to zero. All results rely on the alternative identification strategy, see Table (9).

Table 14: Regression results for the alternative identification strategy (model C)

	IDN	KOR	MYS	PHL	SGP	THA
c	1.70	−0.06	0.06	0.30	−0.60	0.09
	[−0.23, 2.05]	[−1.18, 0.68]	[−0.46, 1.00]	[−0.91, 1.76]	[−5.67, −3.18]	[−1.08, 2.06]
β_{oil}	−1.33	−0.02	−0.08	−0.35	0.06	−0.24
	[−3.99, −0.50]	[−2.15, 1.35]	[−2.98, 1.27]	[−5.19, −0.61]	[−0.15, 3.81]	[−4.86, −0.03]
β_{dem}	0.69	0.14	0.35	0.38	0.22	0.41
	[0.54, 3.07]	[3.41, 4.89]	[1.85, 4.60]	[0.57, 4.35]	[2.77, 7.87]	[1.74, 6.44]
$\beta_{cf(glo)}$	−0.08	0.14	0.16	0.11	0.14	0.14
	[−1.37, 0.95]	[3.41, 4.89]	[1.71, 4.07]	[0.15, 3.25]	[6.39, 9.23]	[1.79, 4.71]
γ	0.80	0.81	0.64	0.77	0.94	0.67
	[9.62, 12.83]	[11.98, 15.09]	[5.82, 7.98]	[9.01, 13.02]	[22.08, 28.87]	[7.65, 11.43]
# obs	67	67	67	67	67	67
R^2	0.74	0.79	0.55	0.77	0.92	0.73

Notes: Estimation results for the regression of model C in Section IV in the main paper. The estimated coefficients (as well as for the coefficient of determination R^2) are median values over all 2000 draws. The lagged domestic inflation rate enters with the coefficient γ. β_{oil} represents the coefficient on the component of real GDP growth driven by oil supply shocks, while β_{dem} represents the coefficient on the component of real GDP growth driven by global demand shocks. $\beta_{cf(glo)}$ represents the coefficient on the counterfactual path of real GDP growth that excludes the component driven by both oil supply shocks and global demand shocks. For all coefficients, the values in square brackets correspond to the 16^{th} and 84^{th} percentile of the t−statistic testing the null that the corresponding coefficient is equal to zero. All results rely on the alternative identification strategy, see Table (9).

B. *Results under Block Exogeneity*

The fact that we use the oil price as an endogenous variable in our paper may seem problematic as we are considering small open economies. The oil price, the argument goes, should be insensitive to fluctuations in domestic variables of these small economies. This raises the question whether we should not consider the oil price as an exogenous variable in our VAR model.[23]

To address this critique, we would like to point out two aspects: first, we need the oil price as an endogenous variable in order to identify a global supply shock (oil supply shock) in our SVAR, because such a shock is defined by the combination of an increase in the oil price and the co-movement of other variables (see, for example, Bobeica and Jarociński, 2019).[24] Since the oil price does not rise in the case of a domestic supply shock, we thus separate the domestic supply shock from the global supply shock.

Second, the estimated coefficients on the lags of the domestic variables in the VAR equation for the oil price should reflect the endogenous or exogenous nature of oil prices, respectively. Thus, the estimated coefficients will be zero if oil prices are indeed

[23] We thank an anonymous referee for this point.

[24] Along these lines, the ratio of domestic GDP in global GDP is modeled as an endogenous variable in our model, saying that implicitly the GDP of the rest of the world is also explained model-wise by domestic variables. However, this is only to exploit the co-movement between the endogenous variables to separate domestic shocks from foreign shocks, see Bobeica and Jarociński (2019).

exogenous for a specific county. Indeed, we find that the oil price can be explained mainly by its own lagged observations.

Nevertheless, to test whether it is problematic if the oil price can react to fluctuations in the domestic variables, we repeat our estimation for all countries under the restriction that the coefficients of all domestic variables in the oil price equation are restricted to zero by means of block exogeneity.[25] This ensures that the oil price is explained only by itself and by the real GDP growth from the rest in the world. To do so, we use a prior of zero with a very small prior variance for all coefficients of the domestic variables, such that these coefficients are estimated to be zero.

We find for all countries that the impulse responses, the historical decompositions as well as the counterfactuals under block exogeneity almost perfectly match our baseline results.[26] Figures (43) and (44) highlight the results from this exercise and show the median oil supply shock under the baseline assumption (red solid path) with 16^{th} and 84^{th} percentiles (red-shaded area), the 5^{th} and 95^{th} percentiles (grey-shaded area) as well as the median oil supply shock under block exogeneity (black path) for all countries. For comparison, the green line shows the median oil price shock without block exogeneity but with the same Minnesota prior as in the case of block exogeneity. As can be seen, there is no remarkable difference between the two approaches. Interestingly, with the Minnesota prior, the results remain exactly unchanged if we compare the model with and without the assumption of block exogeneity.

[25] Note that in this case, we have to rely on a Minnesota prior. The reason is that when using the prior from our baseline specification, the Kronecker structure for the variance of the prior for the VAR coefficients causes instability.

[26] All these results are available on request.

Figure 43: Oil supply shock under block exogeneity for Indonesia, Korea and Malaysia

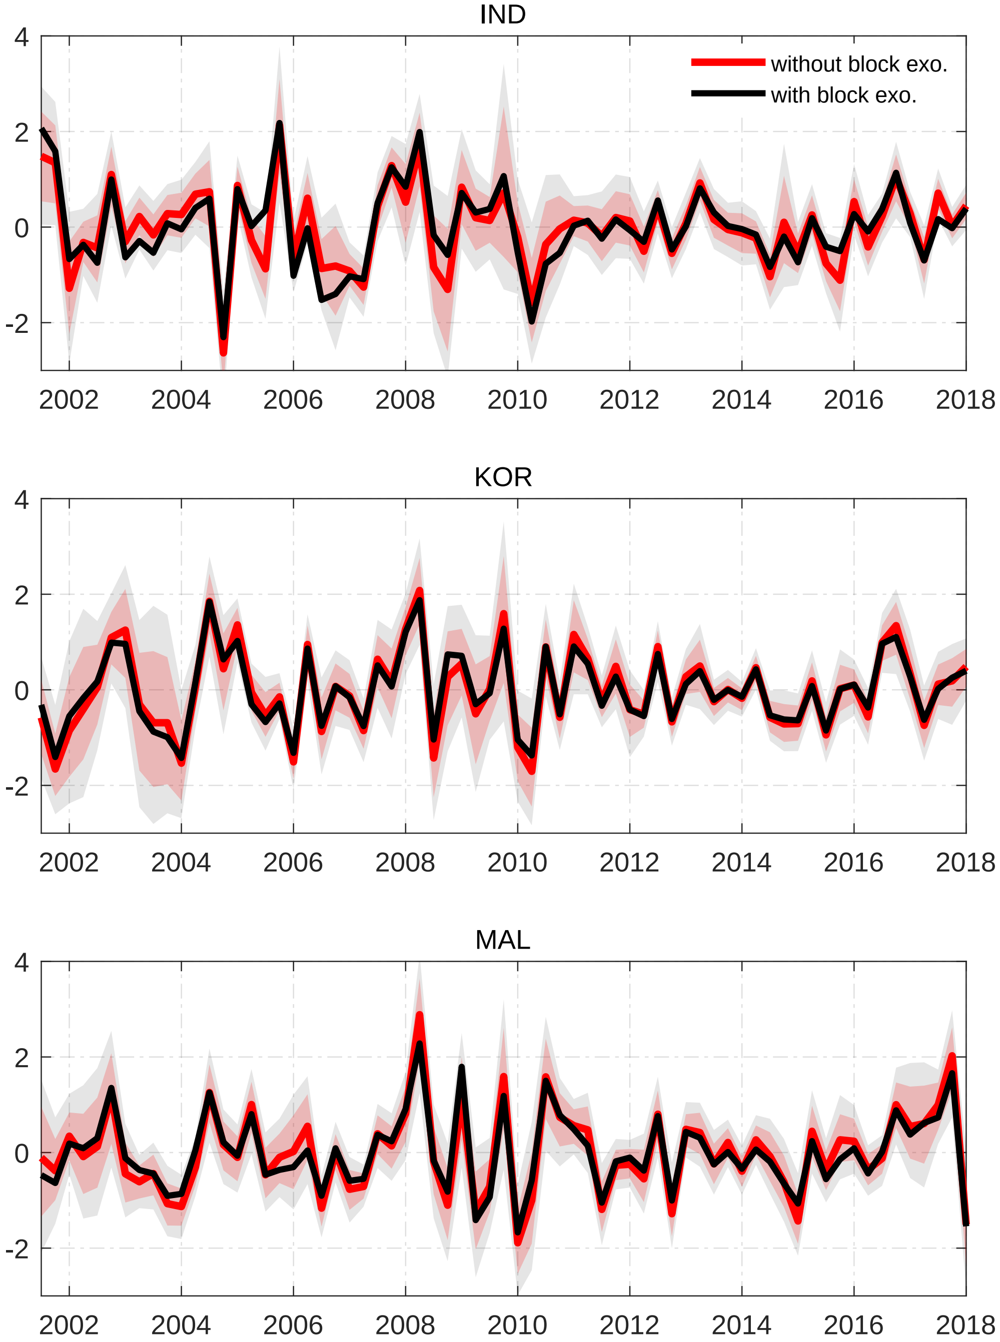

Notes: Baseline refers to the oil price shock from the baseline model. The shock from the model with (without) block exogeneity estimated with the Minnesota prior is denoted "block exo" ("no block exo").

Figure 44: Oil supply shock under block exogeneity for the Philippines, Singapore and Thailand

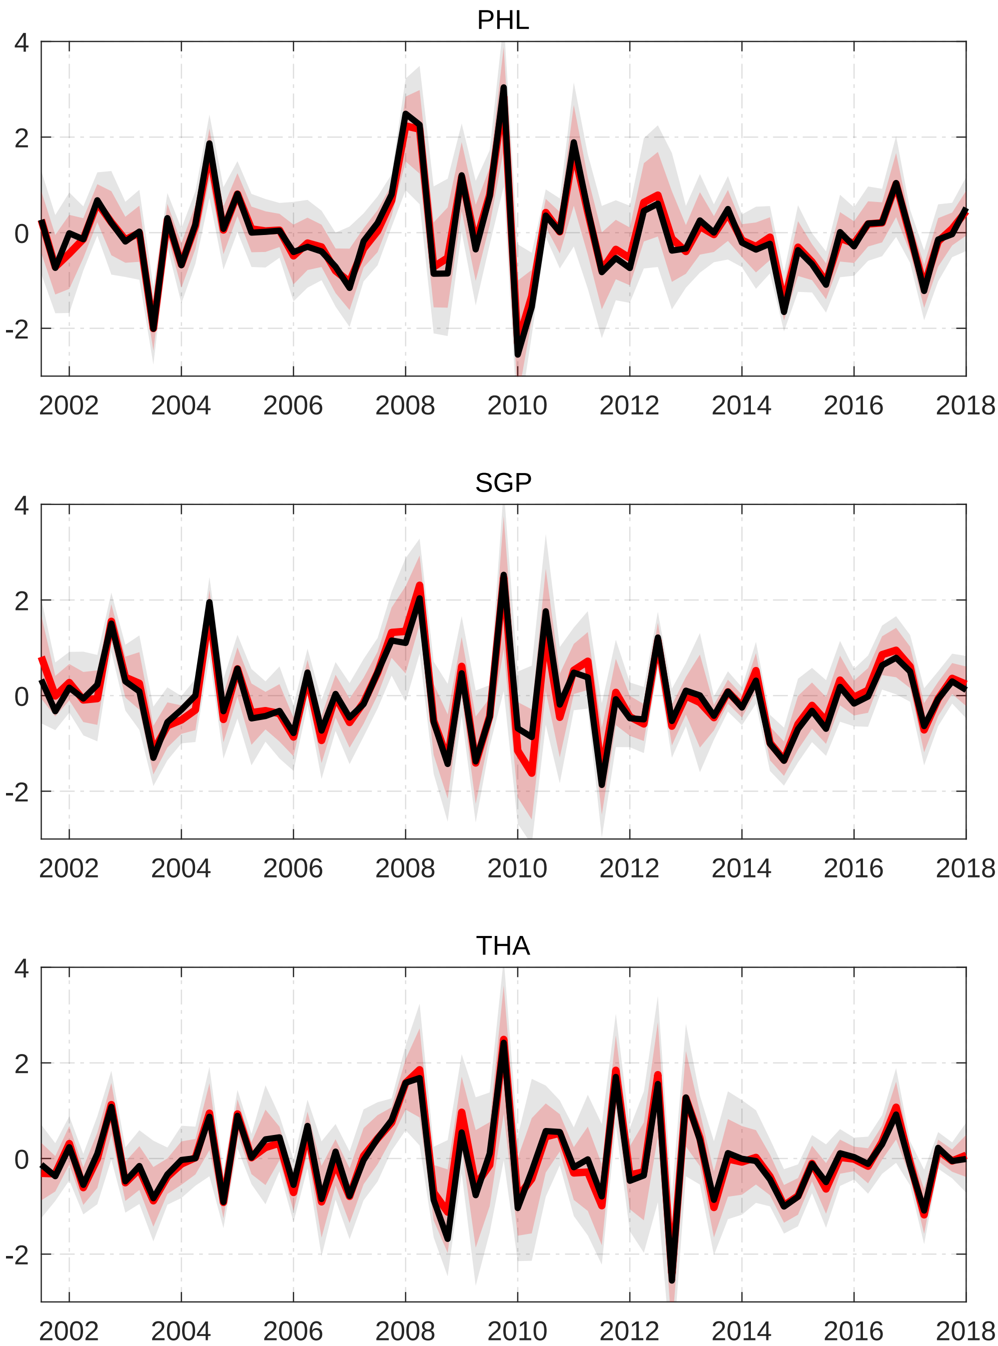

Notes: Baseline refers to the oil price shock from the baseline model. The shock from the model with (without) block exogeneity estimated with the Minnesota prior is denoted "block exo" ("no block exo").

D Results from a Panel Regression

In this section, we present the results from a panel VAR, i.e. we report our results of pooled estimates. The reason is that via pooling the countries we could substantially improve the precision of the estimates. This is supported by the fact that the results reported so far seem stable across countries. Since we assume that dynamic coefficients are homogeneous across units, our panel VAR reads

$$\begin{bmatrix} \mathbf{y}_{1,t} \\ \mathbf{y}_{2,t} \\ \vdots \\ \mathbf{y}_{N,t} \end{bmatrix} = \begin{bmatrix} \mathbf{A}_1 & 0 & \cdots & 0 \\ 0 & \mathbf{A}_1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mathbf{A}_1 \end{bmatrix} \begin{bmatrix} \mathbf{y}_{1,t-1} \\ \mathbf{y}_{2,t-1} \\ \vdots \\ \mathbf{y}_{N,t-1} \end{bmatrix} + \ldots$$
$$+ \begin{bmatrix} \mathbf{A}_p & 0 & \cdots & 0 \\ 0 & \mathbf{A}_p & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mathbf{A}_p \end{bmatrix} \begin{bmatrix} \mathbf{y}_{1,t-p} \\ \mathbf{y}_{2,t-p} \\ \vdots \\ \mathbf{y}_{N,t-p} \end{bmatrix} + \begin{bmatrix} \mathbf{C} \\ \mathbf{C} \\ \vdots \\ \mathbf{C} \end{bmatrix} \mathbf{x}_t + \begin{bmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \\ \vdots \\ \varepsilon_{N,t} \end{bmatrix},$$

where $\mathbf{E}\left(\varepsilon_{i,t}\varepsilon'_{j,t}\right) = 0$ and $\mathbf{E}\left(\varepsilon_{i,t}\varepsilon'_{i,t}\right) = \mathbf{\Sigma}_{ii,t} = \mathbf{\Sigma}_c$, meaning that $\mathbf{\Sigma}_c$ is common to all countries. Hence, we get

$$\mathbf{\Sigma}_t = \begin{bmatrix} \mathbf{\Sigma}_c & 0 & \cdots & 0 \\ 0 & \mathbf{\Sigma}_c & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mathbf{\Sigma}_c \end{bmatrix} = I_n \otimes \mathbf{\Sigma}_c$$

Since such a pooled VAR is just a conventional VAR, we use a Normal-Wishart prior as in our benchmark case. Also the derivation of our structural shocks follows the same sign-restriction pattern as in the main text. In the following, we do not report the impulse responses but only the forecast error variance decomposition.

Table 15: Forecast error variance decompositions for the panel regression

		REAL GDP	INFLATION
$h = 4$	*residual*	5	1
	monetary policy	4	1
	oil supply	9	5
	global demand	49	15
	domestic demand	5	34
	domestic supply	29	43
$h = 8$	*residual*	9	2
	monetary policy	6	2
	oil supply	10	8
	global demand	42	22
	domestic demand	7	29
	domestic supply	27	36
$h = 12$	*residual*	9	3
	monetary policy	6	3
	oil supply	10	9
	global demand	40	24
	domestic demand	8	29
	domestic supply	27	33

Notes: Median shares (in %) of the forecast error variance of the growth rate of real GDP and the inflation rate due to structural shocks for different forecast horizons. All results rely on the alternative identification strategy, see Table (9).

Figure 45: Historical contribution of structural shocks to inflation for Indonesia, Korea and Malaysia

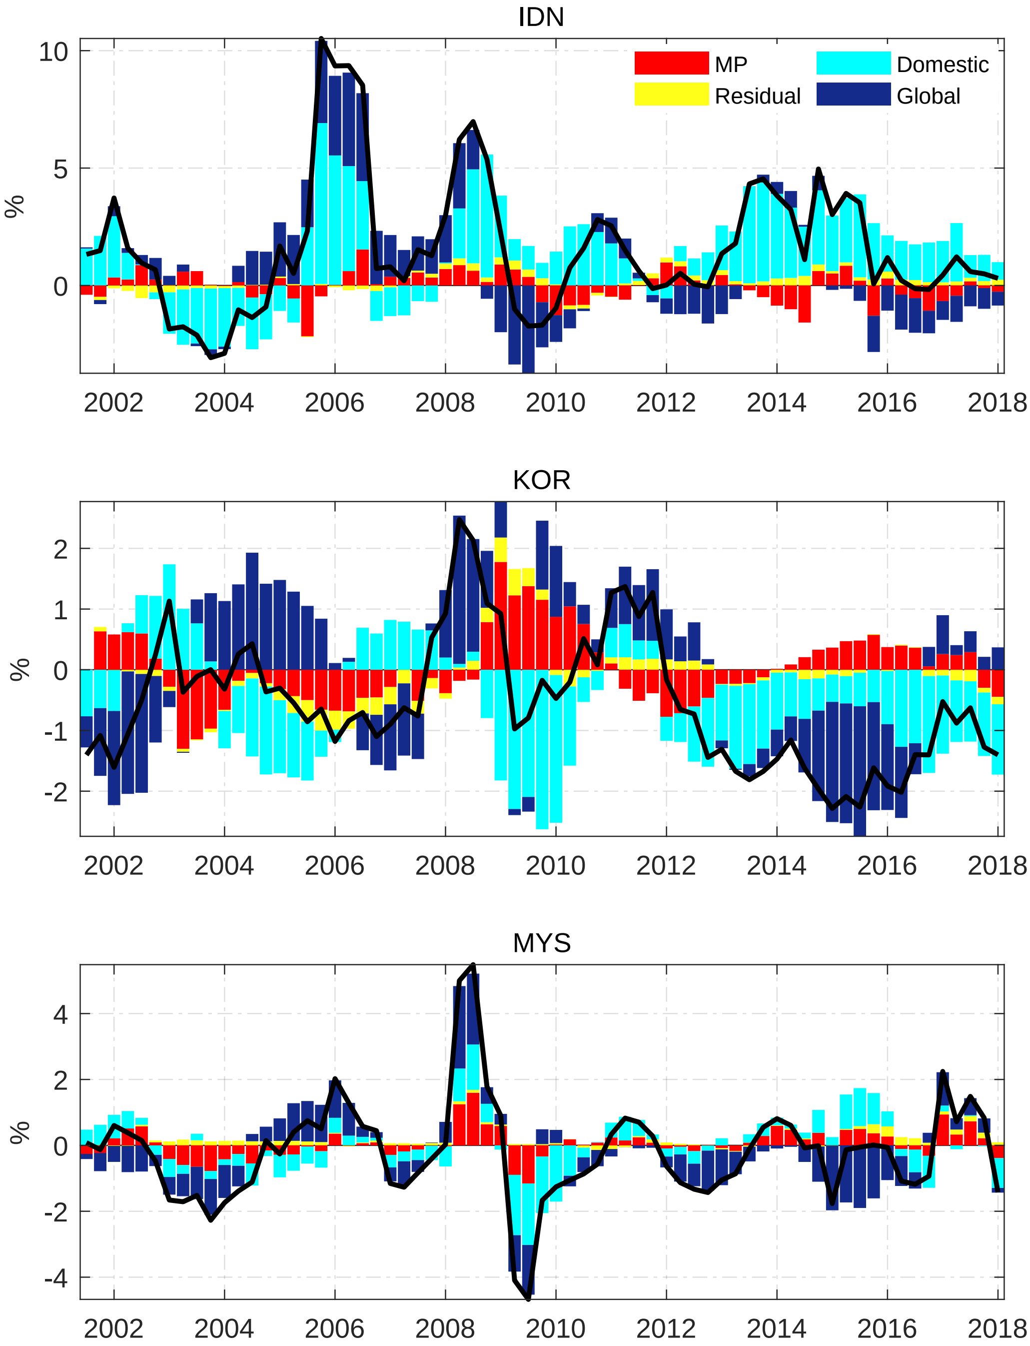

Notes: Median historical contribution of monetary policy shocks (red bars), residual shocks (yellow bars), domestic shocks (teal bars) and global shocks (blue bars) to inflation for Indonesia, Korea and Malaysia. The black path corresponds to the sum of the median contributions of all structural shocks. Results rely on the alternative identification strategy, see Table (9).

Figure 46: Historical contribution of structural shocks to inflation for the Philippines, Singapore and Thailand

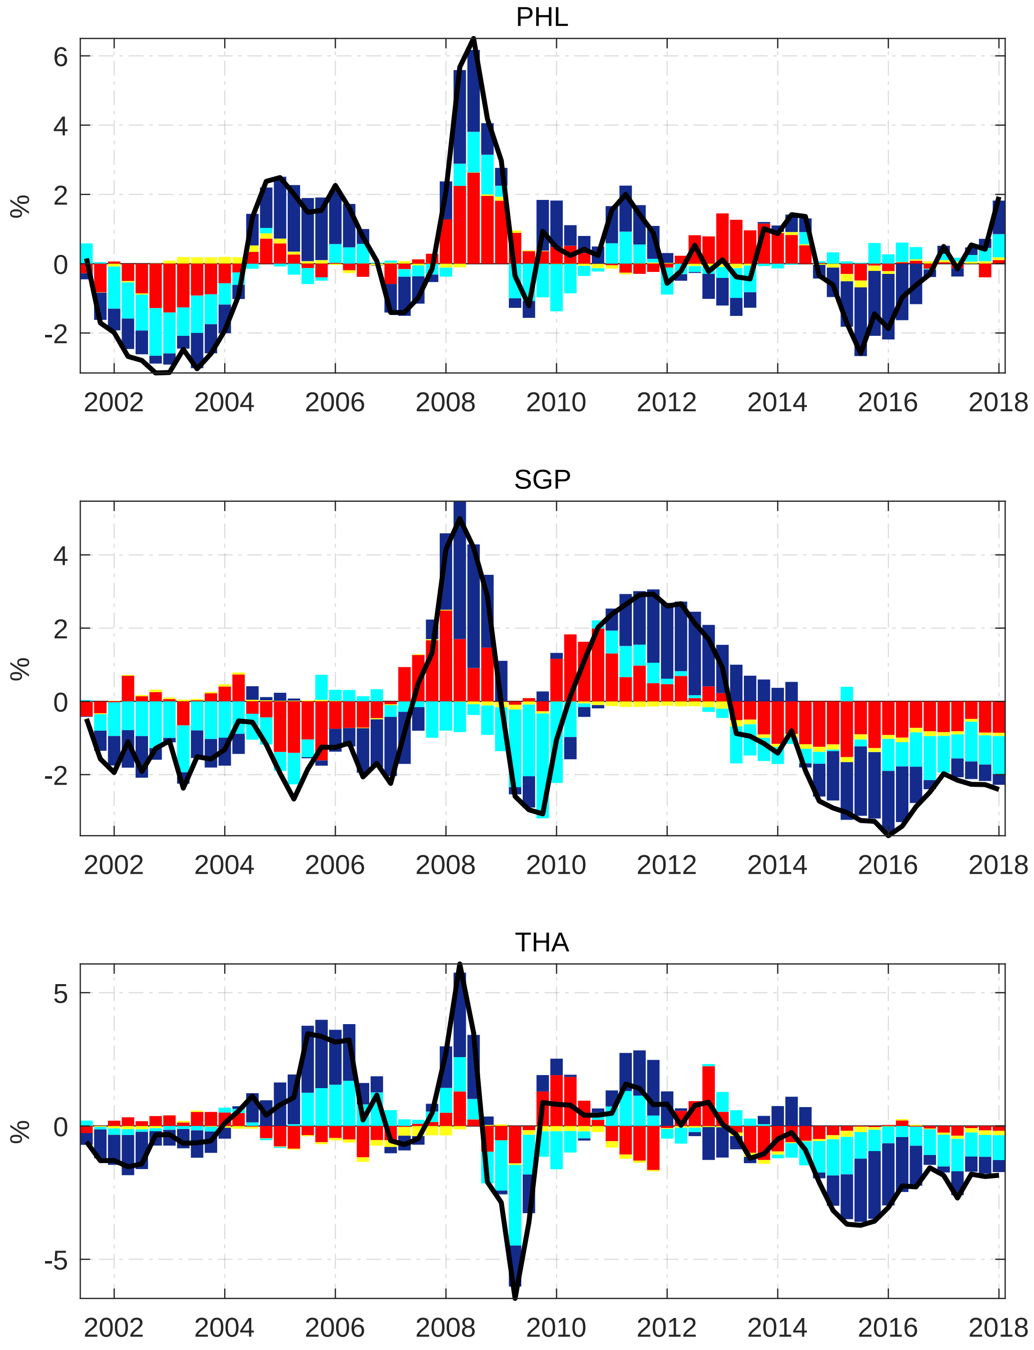

Notes: Median historical contribution of monetary policy shocks (red bars), residual shocks (yellow bars), domestic shocks (teal bars) and global shocks (blue bars) to inflation for the Philippines, Singapore and Thailand. The black path corresponds to the sum of the median contributions of all structural shocks. Results rely on the alternative identification strategy, see Table (9).

Figure 47: Historical contribution of structural shocks to real GDP growth for Indonesia, Korea and Malaysia

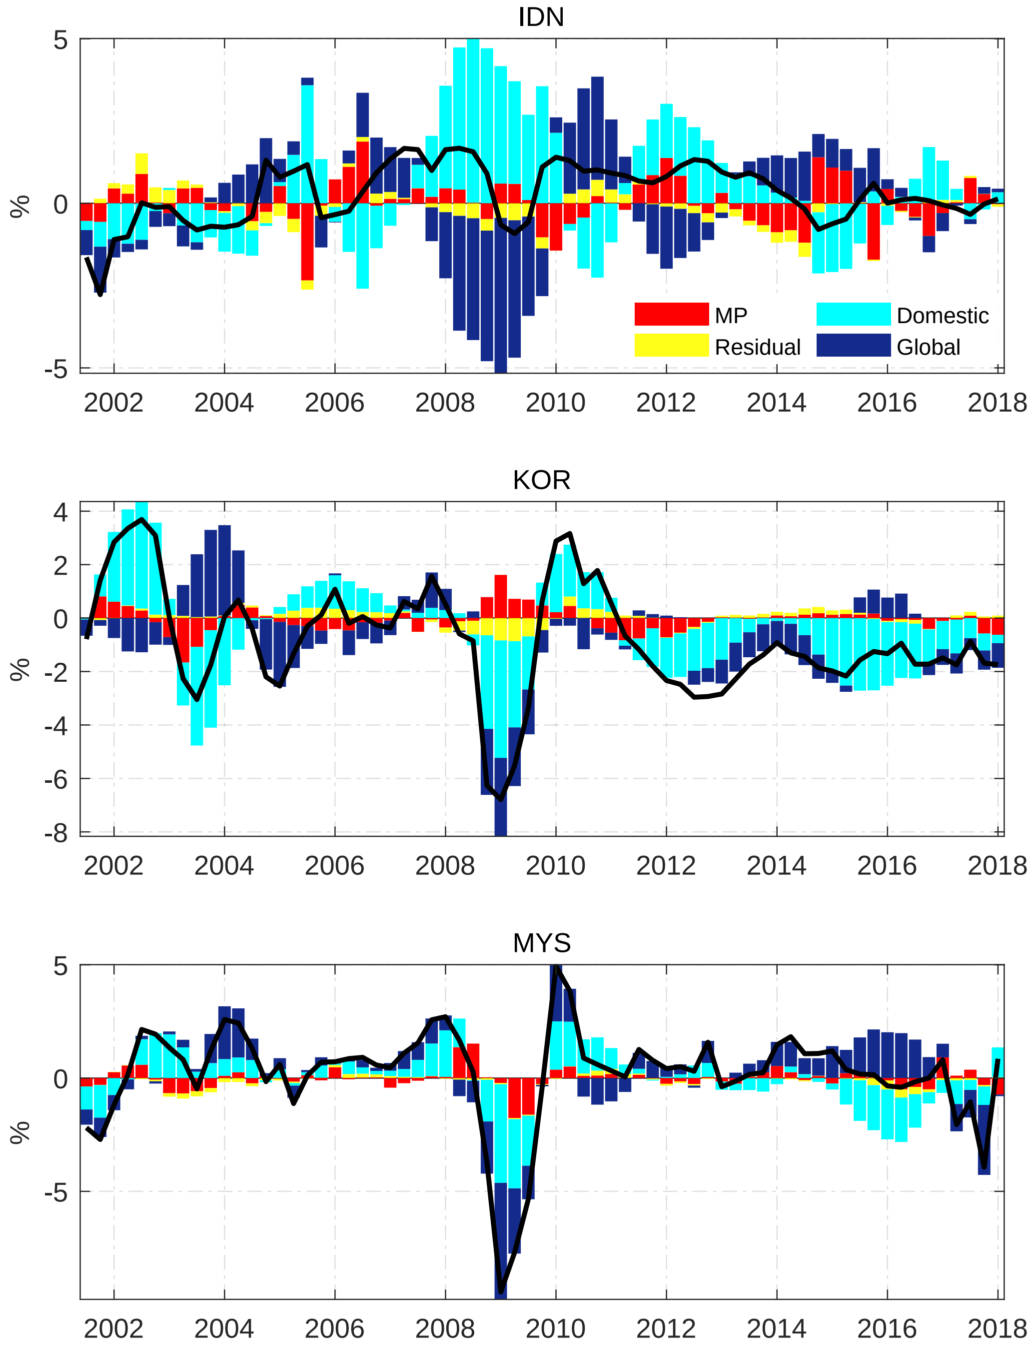

Notes: Median historical contribution of monetary policy shocks (red bars), residual shocks (yellow bars), domestic shocks (teal bars) and global shocks (blue bars) to the growth rate of real GDP for Indonesia, Korea and Malaysia. The black path corresponds to the sum of the median contributions of all structural shocks. Results rely on the alternative identification strategy, see Table (9).

Figure 48: Historical contribution of structural shocks to real GDP growth for the Philippines, Singapore and Thailand

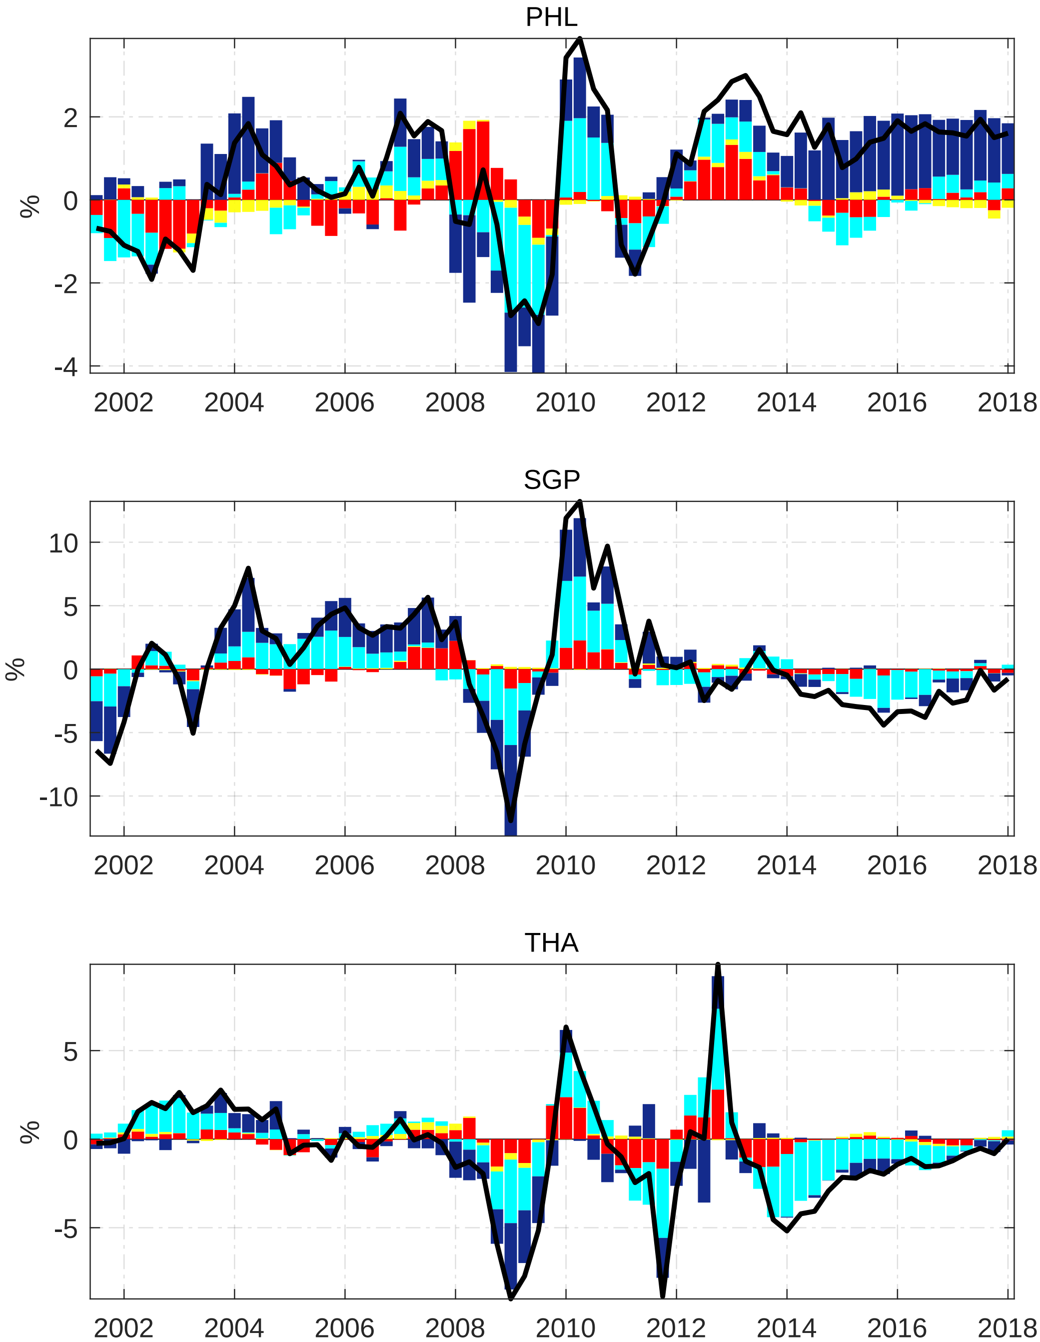

Notes: Median historical contribution of monetary policy shocks (red bars), residual shocks (yellow bars), domestic shocks (teal bars) and global shocks (blue bars) to the growth rate of real GDP for the Philippines, Singapore and Thailand. The black path corresponds to the sum of the median contributions of all structural shocks. Results rely on the alternative identification strategy, see Table (9).

Essay VIII:

The Macroeconomic Effects of Global Supply Chain Disruptions

This is an updated version of a previous draft which is currently under revision at *BOFIT Discussion Papers*.

This paper was presented at the following workshops and international conferences:

I: BOFIT Seminar (Helsinki, Finland)
Date: September 2022
Presenter: Peter Tillmann

II: 2nd Lille-Reading Workshop on International Finance (Virtual)
Date: November 2022
Presenter: David Finck

III: Franco-German Fiscal Policy Seminar 2022 (Berlin, Germany)
Date: November 2022
Presenter: Peter Tillmann

Acknowledgement

Sönke Maatsch and Torsten Schmidt kindly provided data on the RWI/ISL Container Throughput Index. Gianluca Beningno provided regional data from the Global Supply Chain Pressure Index. Efrem Castelnuovo, Georgios Georgiadis and Ben Schumann kindly shared their shock series. We thank Lutz Kilian, Luca Fosso, Mathias Hoffmann, Patrick Hürtgen, Paul Rudel, Sebastian Rüth, Dennis Finck, and Karsten Kucharczyk for helpful comments. Seminar participants at BOFIT/Bank of Finland, participants at 2nd Lille-Reading Workshop on International Finance, participants at the Franco-German Fiscal Policy Seminar 2022, and seminar participants at the University of Giessen provided valuable feedback.

The Macroeconomic Effects of Global Supply Chain Disruptions

David Finck[*] Peter Tillmann[†]

Abstract

Highly interconnected global supply chains make countries vulnerable to supply chain disruptions. This paper estimates the macroeconomic effects of global supply chain shocks for the euro area. Our empirical model combines business cycle variables with data from international container trade. Using a novel identification scheme, we augment conventional sign restrictions on the impulse responses by narrative information about three episodes: the Tōhoku earthquake in 2011, the Suez Canal obstruction in 2021, and the Shanghai backlog in 2022. We show that a global supply chain shock causes a drop in euro area real economic activity and a strong increase in consumer prices. Over a horizon of one year, the global supply chain shock explains about 30% of inflation dynamics. We also use regional data on supply chain pressure to isolate shocks originating in China. Our results show that supply chain disruptions originating in China are an important driver for unexpected movements in industrial production, while disruptions originating outside China are an especially important driver for the dynamics of consumer prices.

Keywords: Container Trade, Supply Chain, Inflation, Narrative Identification, Sign Restrictions

JEL classification: E32, F14, F62

[*]University of Giessen, email: david.finck@wirtschaft.uni-giessen.de
[†]University of Giessen, email: peter.tillmann@wirtschaft.uni-giessen.de

I Introduction

The rise of global value chains has been a pervasive feature of globalization. Antràs (2020) defines global value chains as "a series of stages involved in producing a product or service that is sold to consumers, with each stage adding value, and with at least two stages being produced in different countries". In his survey, Antràs (2020) further summarizes the case for participation in global value chains, which "allow countries to benefit from the comparative advantage of other countries not only at the sectoral level but also at the stage level within sectors". Hence, global supply chains are a source of welfare gains.

However, interconnected global supply chains come with a drawback: the tight network of global sourcing makes countries vulnerable to disruptions of global value chains. This danger manifested itself during the Covid-19 pandemic between 2020 and 2022. Even small disruptions in production and logistics cascaded into sizable macroeconomic shocks when authorities imposed lockdowns to contain the spread of the pandemic. For instance, the zero-Covid policy pursued by authorities in Shanghai in the spring of 2022 led to standstills in manufacturing, port closures, and large delays in international container trade. These disruptions have macroeconomic consequences for highly integrated advanced economies. At the time of writing, the strong increase in inflation in the euro area but also in other advanced countries are in some parts attributed to global supply chain disruptions (e.g., Tenreyro, 2021; Lane, 2022; Reis, 2022).

The euro area is closely embedded in international supply chains.[1] Figure (1) documents the extent to which selected euro area countries participate in cross-border supply chains.[2] In several member countries, more than 50% of the output of the manufacturing sector directly or indirectly crosses more than one border.[3] Thus, disruptions to the flow of goods along the supply chain should be a major source of business cycle fluctuations in the euro area.

This paper aims to quantify the macroeconomic effects of disruptions to global supply chains. We estimate the impact of global supply chain shocks on the business cycle in the euro area within a structural vector autoregression (VAR) model identified using a combination of sign restrictions and narrative restrictions (Antolín-Díaz and Rubio-Ramírez, 2018). Our key contribution is, first, a new identification scheme that isolates exogenous distortions of global supply chains and, second, an estimated VAR model that shows the effect of such shocks on real economic activity and inflation in the euro area.

This paper uses data from international container trade to measure the disruption of

[1] In this paper, we use the expressions global value chains and global supply chains interchangeably.

[2] These countries are: Austria (AT), Belgium (BE), Estonia (EE), Finland (FI), France (FR), Germany (DE), Greece (GR), Ireland (IE), Italy (IT), Latvia (LV), Lithuania (LT), the Netherlands (NL), Portugal (PT), Slovakia (SK), Slovenia (SI), Spain (ES).

[3] See Gunnella et al. (2019) for a discussion of the integration of the euro area into global supply chains.

Figure 1: Manufacturing output related to global value chains for selected euro area countries

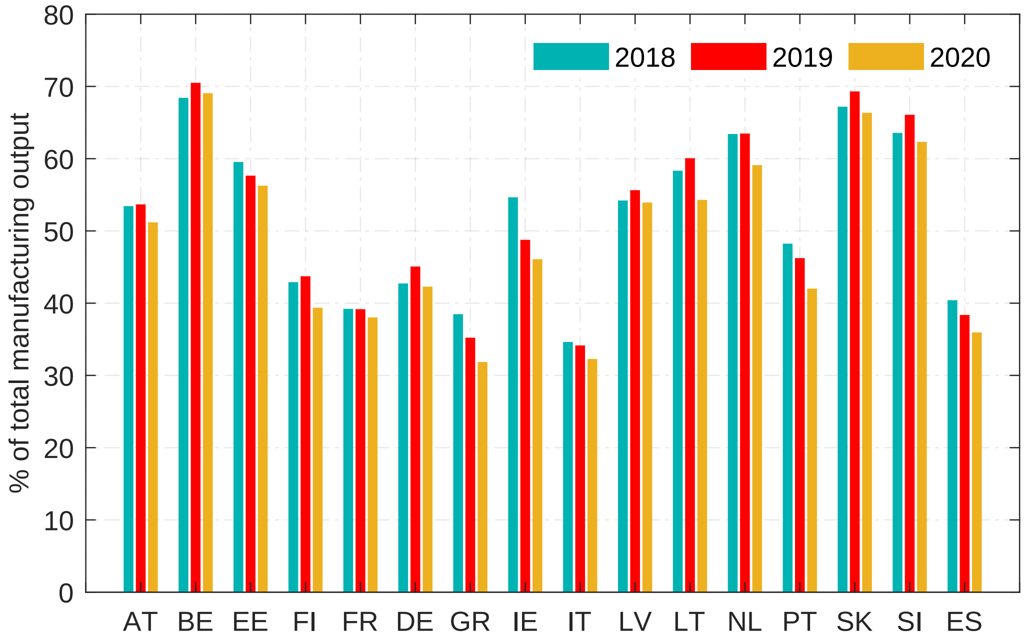

Notes: Share of manufacturing output that directly or indirectly crosses more than one border. The data is taken from the WITS database of the World Bank, see https://wits.worldbank.org/link.

global supply chains. Container shipping is the backbone of global trade and reflects disruptions to global sourcing. The key rationale behind our identification is that an adverse global supply chain shock should temporarily lead to a smaller number of containers being processed and an increase in the costs of shipping a container. Hence, an adverse shock tightens the supply of shipment capacities. In contrast, a favorable supply chain shock generates excess capacity in the market for container trade which increases the container throughput and lowers container freight rates.

We proceed as follows. First, we estimate a VAR model that includes standard euro area variables such as industrial production and consumer prices, as well as variables summarizing container shipping, which reflects global supply chains. The latter comprise the HARPEX index of global container freight rates, the RWI/ISL index of container throughput in the most important container ports of the euro area, and the Global Supply Chain Pressure Index provided by Benigno et al. (2022). We employ restrictions on the sign of the impulse response functions following Arias et al. (2018) in order to identify a global supply chain shock. In particular, we assume that a contractionary supply chain shock raises container rates and lowers the number of processed containers. We further assume that these effects are accompanied by an increase in the Global Supply Chain Pressure Index. Importantly, we leave the macroeconomic variables unrestricted.

It turns out that the shock identified through sign restrictions does not match the widely accepted historical narratives with respect to three key episodes. Each of these periods is characterized by exogenous disruptions of global supply chains. Specifically, we want that the global supply chain shock exhibits a positive, i.e. restrictive, realization during the following episodes: the first is the Tōhoku earthquake and the following tsunami in March 2011, which distorted global supply chain inter-linkages. Boehm et al. (2019) and Carvalho et al. (2021) use this event in order to identify the supply chain disruptions with firm-level data. The second event is the obstruction of the Suez Canal in March 2021. When Ever Given, one of the largest container ships in the world, blocked the channel for six days, 200 ships had to wait on either entrance to the channel. The obstruction caused massive delays in container trade. Furceri et al. (2022) use this event as an instrument to identify supply chain shocks. The third episode is the zero-Covid policy imposed by authorities in Shanghai in April 2022. As a consequence of the ultra-restrictive lockdown, hundreds of ships could not be processed in time in the port of Shanghai, causing a severe backlog of container trade.

Therefore, we follow Antolín-Díaz and Rubio-Ramírez (2018) and impose constraints using narrative information. Therefore, we rely on narrative restrictions meant to constrain the admissible set of structural parameters by ensuring that the structural shocks and historical decompositions align with the established narrative around these three historical events. Specifically, we assume that the global supply chain shock was positive during the three episodes mentioned before. Besides these restrictions on the sign of the shock, we also impose the constraint that in March 2011, the global supply chain shock was the dominant driving force of the Global Supply Chain Pressure Index. Our results imply that the narrative information from the three episodes of supply chain distortions is indeed pivotal for identification.

Our key result is that a global supply chain shock strongly affects the euro area business cycle. A shock of one standard deviation causes a fall in industrial production by about one percent and an increase in consumer prices by 0.3 percent. This effect is highly persistent and peaks after 17 months. Import prices in the euro area are even more sensitive to supply chain disruptions and increase by one percent. The global supply chain shock explains a large part of business cycle dynamics. Our results imply that over a horizon of one year, supply chain disruptions account for 9% of the fluctuations of industrial production and 29% of the adjustment of consumer prices. In a counterfactual analysis, we show the contribution of the supply chain shock associated with the Tōhoku earthquake on economic activity and prices. We also use sectoral price indices to show that import and producer prices of manufactured goods and intermediate goods are particularly sensitive to supply chain shocks, in which prices of consumer goods respond less.

These responses are highly statistically significant once we impose sign and narrative restrictions. If we relax the narrative restrictions, the significant effects on the euro area business cycle disappear. Hence, identifying a structural shock that matches the established historical narrative is crucial. Our results also remain unchanged if we change the implementation of the narrative restrictions and distinguish the global supply chain shock from an increase in geopolitical risk. The identified shock series is uncorrelated with oil supply and other global shocks, which we take from prominent contributions to the literature. Hence, the supply chain shock is a separate and very powerful driving force of economic activity that has not yet been studied by the literature on business cycles in open economies.

Finally, we shed light on the geographic origin of the global supply chain shock. We distinguish supply chain disruptions originating in China from supply chain disruptions in the rest of the world. We find that both supply chain disruptions have qualitatively similar effects on industrial production. However, consumer prices in the euro area are much more sensitive to supply chain shocks from the rest of the world than those from China. Decomposing the dynamics of the two variables over time shows that Chinese shocks are more important in driving industrial production than shocks from the rest of the world. In contrast, supply chain disruptions emanating from the rest of the world are more important for the dynamics of consumer prices.

A small number of papers study the macroeconomic effects of supply chain shocks.[4] Kilian et al. (2021) study the role of frictions in container trade for recovering the U.S. economy from the Covid-19 pandemic. The authors construct a series of container trade to and from North America, but also use the RWI/ISL container index for robustness. They estimate a three-variable VAR model with real personal consumption, industrial production, and the number of containers. The VAR model is identified recursively, thus imposing the constraint that shocks to container trade need at least a month to affect the U.S. economy. Their results support the notion that frictions in container trade explain a large share of the incomplete recovery after the pandemic. In our paper, we choose an alternative identification based on narrative sign restrictions alongside traditional sign restrictions. One advantage is that we do not need to impose any constraint on the adjustment of the macroeconomic variables to the global shipping variables. Most importantly, the identification scheme in Kilian et al. (2021) does not allow for a clear distinction between supply-side and demand-side disruption in container trade. In addition, we concentrate on the impact of the shock on prices. Since Kilian et al. (2021) do not include prices in their model, they cannot say anything about the inflationary impact of shocks.

In an application of their Global Supply Chain Pressure Index, Benigno et al. (2022) use local projections in order to estimate the response of inflation in the U.S. and the euro area to supply chain pressure as well as a global demand and an oil price shock. They adopt a recursive identification scheme and show that supply bottlenecks contribute strongly to inflationary pressure since the outbreak of the pandemic.

LaBelle and Santacreu (2022) investigate the effect of supply chain disruptions on

[4]Auer et al. (2017) study the impact of global value chains on inflation dynamics. They argue that an expansion of value chains makes domestic inflation more sensitive to measures of global economic slack.

U.S. producer price inflation. The authors let the variation of global sourcing across industries interact with measures of global supply chain pressure. Their results imply that supply chain pressure explained up to 20 percentage points of the producer price inflation in November 2021. Isaacson and Rubinton (2022) estimate the pass-through from shipping costs to import prices at the level of good types. In their baseline model, a one percent increase in freight rates raises import prices by two basis points. Hence, the effect is relatively small.

Furceri et al. (2022) use the Baltic Dry Index to measure global shipping costs and estimate its effect on inflation for a large panel of countries from 1992 to 2021. They estimate panel local projections with the Baltic Dry Index as the driving variable. In a separate specification, they use the Suez Canal obstruction in 2021 to instrument shipping costs. The authors find a highly significant and economically large response of consumer, producer, and import prices.

Capolongo et al. (2022) estimate a VAR model including macroeconomic variables as well as container prices and delivery times for the sample period 2002 to 2021. Imposing sign restrictions, they identify two supply and two demand shocks. Shocks to supply chains have a strong impact on prices. The increase in delivery times during the Covid-19 pandemic was largely due to demand rather than supply disruptions. The authors impose restrictions on all variables, including the macroeconomic time series. In this paper, we identify a global supply chain shock while leaving real economic activity and prices completely unrestricted.

Khalil and Weber (2022) identify disruptions to Chinese supply chains using a sign-restricted VAR model. Their key identifying assumption is that U.S. manufacturing firms substitute Chinese manufactured imports by imported goods from the rest of the world if the supply chain shock emanates from China. They find large effects of this shock on manufacturing production and prices.

These findings can also be rationalized in a theoretical model. Alessandria et al. (2022) present a two-country model with heterogeneous firms and a rich set of supply chain frictions to understand the aggregate effects of supply chain shocks. They show that delays in shipping can be highly contractionary for the whole economy. Di Giovanni et al. (2022) also use a calibrated multi-sector model to show the quantitative impact of supply chain bottlenecks during the pandemic.

The remainder of this paper is structured as follows: Section II introduces our empirical model, and Section III explains the data and the identification scheme. The main results are discussed in Section IV, while V includes a battery of robustness checks and further results. Section VI is devoted to the role of China. Section VII draws conclusions. An online appendix contains additional material.

II Methodology

This section outlines our baseline model specification and discusses how we implement the sign and narrative restrictions.

A. *Structural VAR Model*

We are interested in the structural vector autoregression of the form

$$\mathbf{y}_t'\mathbf{A}_0 = \mathbf{c} + \mathbf{y}_{t-1}'\mathbf{A}_1 + \ldots + \mathbf{y}_{t-p}'\mathbf{A}_p + \varepsilon_t', \tag{1}$$

where $\mathbf{y}_t$ is an $n \times 1$ vector which contains the endogenous variables, $\mathbf{A}_1, \ldots, \mathbf{A}_p$ are $n \times n$ matrices of parameters and $\mathbf{c}$ is a $1 \times n$ vector of parameters. ε_t is an $n \times 1$ vector of structural shocks and $\mathbf{A}_0$ is an invertible $n \times n$ matrix which contains the contemporaneous relationships among the endogenous variables. The model described above can be rewritten in compact form as

$$\mathbf{y}_t'\mathbf{A}_0 = \mathbf{x}_t'\mathbf{A}_+ + \varepsilon_t', \tag{2}$$

where $\mathbf{x}_t$ is a $(np+1) \times 1$ vector given as $\mathbf{x}_t' = \left[1, \mathbf{y}_{t-1}', \ldots, \mathbf{y}_{t-p}'\right]$ and $\mathbf{A}_+' = \left[\mathbf{c}' \ \mathbf{A}_1' \ldots \mathbf{A}_p'\right]$ is of the dimension $(np+1) \times n$. Finally, the reduced-form version of the model we estimate is given by

$$\mathbf{y}_t' = \mathbf{x}_t'\mathbf{B} + \mathbf{u}_t', \tag{3}$$

where $\mathbf{B} = \mathbf{A}_+\mathbf{A}_0^{-1}$, $\mathbf{u}_t' = \varepsilon_t'\mathbf{A}_0^{-1}$ and $E[\mathbf{u}_t\mathbf{u}_t'] = \left(\mathbf{A}_0\mathbf{A}_0'\right)^{-1} = \mathbf{\Sigma}$. While $\mathbf{A}_0$ and $\mathbf{A}_+$ contain the structural parameters, summarized as $\mathbf{\Theta} = (\mathbf{A}_0, \mathbf{A}_+)$, $\mathbf{B}$ and $\mathbf{\Sigma}$ contain the reduced-form parameters.

B. *Baseline Sign Restrictions*

The structural form in (1) is yet not identified. Hence, we need to impose restrictions on the set of structural parameters. While the Cholesky approach is one natural candidate to do so, this approach comes with probably too restrictive assumptions in the case of partial identification because it pins down the entire structural matrix and, among other things, imposes restrictions on the responses following shocks in that we are not interested. One common alternative since Faust (1998) and Canova and De Nicolo (2002), which is nowadays well understood, builds on a handful of uncontroversial sign and zero restrictions on either the impulse response functions or on the admissible set of structural parameters themselves (see, for instance, Arias et al., 2019). The advantage is that such minimalistic restrictions are generally weaker and less restrictive than in the Cholesky (or triangular factorization) approach. The literature provides several algorithms for Bayesian Inference based on SVARs with sign and zero restrictions. Throughout the paper, we closely follow Arias et al. (2018), who present very efficient algorithms implementing a mixture of sign and zero restrictions. Moreover, their algorithms guarantee that identification solely comes from the sign and zero restrictions imposed by the researcher. [5]

Denote by $\mathbf{e}_{j,n}$ the jth column of $\mathbf{I}_n$ and the set of structural parameters $\Theta = (\mathbf{A}_0, \mathbf{A}_+)$. We can then implement sign restrictions on the impulse responses as

$$\Gamma(\Theta) = \left(\mathbf{e}_{1,n}'\mathbf{F}(\Theta)'\,\mathbf{S_1}', \cdots, \mathbf{e}_{n,n}'\mathbf{F}(\Theta)'\,\mathbf{S_n}'\right) > 0, \tag{4}$$

where appropriate matrices for $\mathbf{S}_j$ and $\mathbf{F}(\Theta)$ result in a formalized set of sign restrictions on the IRFs. To do so, it is convenient to vertically stack the structural impulse responses that are subject to a restriction into a matrix $\mathbf{F}(\Theta)$ across horizons. Hence, in order to impose sign restrictions, one defines an $s_j \times r_j$ restriction matrix $\mathbf{S_j}$ containing entries of zero, one, and minus one across both the variables and the horizons over which the restrictions shall be imposed.

Note that, throughout the paper, we are only interested in partial identification, i.e. we are not interested in the identification of all the n structural shocks.

C. *Narrative Restrictions*

We follow Antolín-Díaz and Rubio-Ramírez (2018) and implement narrative restrictions. These restrictions are meant to constrain the admissible set of structural parameters by ensuring that around selected historical events, the structural shocks and the historical decomposition (or both) align with the established narratives, which will be explained in the next section. Antolín-Díaz and Rubio-Ramírez (2018) show that even a single narrative sign restriction may dramatically sharpen and even change the inference of SVARs originally identified via traditional sign restrictions.

We rely on two types of narrative sign restrictions. The first type restricts the sign of the structural shocks, while the second type restricts the historical decomposition of the endogenous variables by putting the absolute historical contribution of the shock of interest in proportion to the historical contributions of the other shocks.

Note that, conditional on the structural parameters in $\Theta = (\mathbf{A}_0, \mathbf{A}_+)$, the orthogonal structural shocks ε_t which, in contrast to the reduced-form innovations $\mathbf{u}_t$, have an economic interpretation, are obtained as

$$\varepsilon_t'(\Theta) = \mathbf{y}_t'\mathbf{A}_0 - \mathbf{x}_t'\mathbf{A}_+. \tag{5}$$

Regarding the sign of the shock, we might want to impose the restrictions that the sign of a structural shock j for the periods $t_1, \ldots, t_{s_j}$ is positive. Following Antolín-Díaz

[5]These algorithms, therefore, overcome several pitfalls with other popular approaches, e.g. the penalty function approach. See Arias et al. (2018) for an intensive discussion.

and Rubio-Ramírez (2018), this class of narrative sign restrictions can be imposed as

$$\mathbf{e}'_{j,n}\varepsilon_{t_v}(\Theta) > 0, \qquad 1 \leqslant v \leqslant s_j, \tag{6}$$

where $\mathbf{e}'_{j,n}$ is the jth column of the identity matrix $\mathbf{I}_n$ and s_j is the number of restrictions.

As for the second type of narrative restrictions, we use narrative information that narrows the set of admissible parameters via information about the contribution of shocks to the unexpected movement of endogenous variables. This type of restriction is based on narrative information for which we know that the absolute contribution of shock j to variable i from periods $t_v + h_v$ was more important than the absolute contribution of any other shock to the same variable for the same period. Formally, this restriction can be expressed as

$$\begin{aligned}&|H_{i_v,j,t_v,t_v+h_v}\left(\Theta, \varepsilon_{t_v}\left(\Theta\right), ..., \varepsilon_{t_v+h_v}\left(\Theta\right)\right)|\\ &\qquad - \max_{j'\neq j} |H_{i_v,j',t_v,t_v+h}\left(\Theta, \varepsilon_{t_v}\left(\Theta\right), ..., \varepsilon_{t_v+h_v}\left(\Theta\right)\right)| > 0,\end{aligned} \tag{7}$$

where $|H_{i_v,j,t_v,t_v+h_v}(\cdot)|$ denotes the historical contribution of the jth shock to the ith variable from periods t_v to $t_v + h_v$ for $1 \leq v \leq s_j$.

III Data and Identification

This section describes the set of endogenous variables and our identification strategy. We achieve identification through a combination of conventional restrictions on the sign of the impulse responses following a supply chain shock and additional narrative information on the role of this shock in selected episodes.

A. *Data*

The vector of endogenous variables includes six variables. We include three variables that reflect the business cycle in the euro area. These three variables consist of industrial production (including construction), the Harmonized Index of Consumer Prices, and the index of import prices. In Section IV, we replace the index of import prices with industry-specific price levels.

The remaining three variables reflect international container shipping and global supply chains. The first is the RWI/ISL container throughput index provided by the Leibniz-Institut für Wirtschaftsforschung (RWI) in Essen, Germany, and the Institute of Shipping Economics and Logistics (ISL) in Bremen, Germany. The index reports the (seasonally adjusted) number of processed containers in the North Range, i.e. the ports of Le Havre, Zeebrugge, Antwerp, Rotterdam, Bremen/Bremerhaven, and Hamburg. These are the most important ports for container trade in the euro area. In December 2007, the last month for which we have disaggregated data on some ports, North Range throughput accounted for 70% of the total container throughput in euro area ports. As a caveat, the index does not contain information on the value or the volume of cargo.[6]

The second shipping variable is the HARPEX PETERSEN Charter Rates Index (HARPEX) which reflects the worldwide price development on the charter market for container ships. Note that the HARPEX is different from other price indices for shipping costs that are used in the literature, e.g. the Baltic Dry Index, which records the freight rates of raw materials, i.e. input used in an early stage of the production process. The HARPEX tracks prices for container shipment of semi-finished or finished products created from raw materials.

The last variable we include in our VAR is the Global Supply Chain Pressure index (GSCPI) provided by the Federal Reserve Bank of New York (Benigno et al., 2022). These authors construct a summary indicator of global supply chain pressure based on Purchasing Managers' Index (PMI) surveys for manufacturing firms in China, the euro area, Japan, Korea, Taiwan, the UK, and the US, measure of transportation costs such as the HARPEX and the Baltic Dry Index and indicator of airfreight costs.

The inclusion of PMI data allows the authors to isolate the supply-side conditions of global value chains. Benigno et al. (2022) regress the country-specific supply chain measures from the PMI survey (delivery times, backlogs, purchased stocks) on the "new orders" component of the PMI surveys. The residual of this regression enters the construction of the GSCPI to ensure that it reflects supply-side pressure. The authors also regress the transportation cost proxies on the "new orders" and the "quantities purchases" components of the PMI survey. The residuals reflect transportation costs net of demand-side effects. After controlling for demand effects, Benigno et al. (2022) extract the first principal components of the series, which is then expressed in standard deviations from its mean.

Figure (2) shows the evolution of the three shipping variables, i.e. the HARPEX, the GSCPI, and the Northrange container throughput. The figure illustrates the strong increase in global supply chain pressure since 2020, which is also reflected in the sharp rise in the HARPEX index of freight rates. Between 2012 and 2017, we see a long phase of below-average supply chain pressure going hand-in-hand with stable freight rates and container throughput.

In our baseline specification, we use six lags. All variables except the GSCPI are included in logs. Hence, the impulse response functions for industrial production, consumer prices, import prices, the HARPEX, and container throughput are understood as percentage changes. The GSCPI instead enters in standard deviations. In order to account for the large fluctuations since the outbreak of the pandemic, we include a dummy from 2020M3 onward. The sample is 2007M1 to 2022M7, and the

[6]See Döhrn and Maatsch (2012) and Döhrn (2019) for more information on the RWI/ISL index. The authors show that the index closely reflects the dynamics of global trade.

Figure 2: The shipping variables in our VAR model

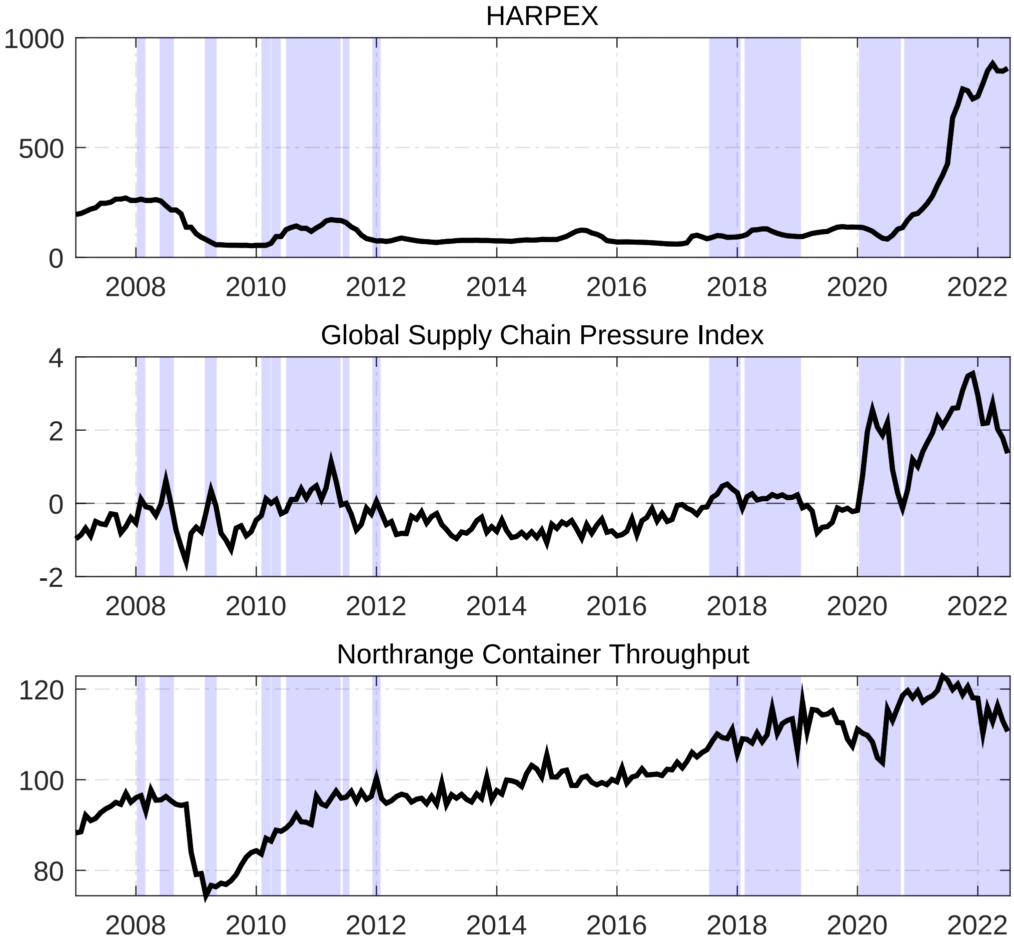

Notes: The figures show the series HARPEX, the GSCPI, and the Northrange Container Throughput index, which we use in our VAR model. The HARPEX and the Container Throughput Index are measured in index points (2015 = 100). The GSCPI is shown in standard deviations from the sample average. The shaded areas represent episodes in which the GSCPI is above zero, i.e., when supply chain pressure is tight relative to the average.

frequency is monthly.

B. *Conventional Sign Restrictions*

We now explain how we identify a global supply chain shock using sign restrictions. As can be seen in Table (1), we restrict three of our six endogenous variables, while the responses of the remaining three variables are left unrestricted. Importantly, we do not restrict the adjustment of the three variables reflecting the euro area business cycle.

Our identification of a supply chain shock is based on the idea that a disruption of global supply chains changes prices and quantities of container shipment. We assume that a supply chain shock leads to an increase in the HARPEX, i.e. to an increase in charter rates (ship rentals) for container vessels. Elementary theory suggests

Table 1: Baseline sign restrictions

HARPEX	GSCPI	Container Throughput	IP	Consumer Prices	Import Prices
+	+	–			

Notes: All restrictions hold on impact and for two consecutive months.

that a shock to supply should move to opposite movements of prices and quantities. Hence, this increase is accompanied by a reduced number of processed containers. This negative co-movement should identify a supply shock. In order to distinguish our supply chain shock from any other supply shock, such as technology or oil price shocks, we assume that the shock also increases supply chain pressure. Therefore, we additionally restrict the response of the global supply chain pressure index (GSCPI) and assume that this index rises after a supply chain shock. Importantly, an increase in the GSCPI is supply-side driven by design, as the GSCPI is purged of demand-driven factors.

We deliberately leave the reactions of industrial production, consumer prices as well as import prices unrestricted. Notice that a negative co-movement of the responses of consumer prices and industrial production is a common identifying assumption for a conventional supply shock. Instead, we identify a supply *chain* shock and look at whether and how domestic consumer prices, import prices, and industrial production react to such supply chain disruptions.

C. *The Narrative Information*

We now discuss the information we will use to elicit the narrative sign restrictions. In our preferred specification, we rely on three episodes for which narrative information clearly suggests the presence of exogenous, contractionary disruptions to global supply chains. All three episodes feature local exogenous shocks to manufacturing and logistics that cascaded into global supply chain disruptions. Importantly, the events are selected by the unambiguous nature of the supply chain disruption according to the established narrative rather than the magnitude of the disruption. Put differently, these events are clear cases of supply chain shocks, but not necessarily the largest. In the following, we describe these episodes in detail and explain which narrative restrictions we derive from them.

Tōhoku Earthquake and Tsunami in 2011

On March 11, 2011, the most powerful earthquake ever recorded in Japan shocked the Pacific Ocean off the coast of the Tōhoku region on the Japanese Island of Honshu. The earthquake caused a tsunami, which led to the Fukushima Daiichi nuclear disaster and flooded large areas along the coastline. With almost 20,000 deaths, the disaster also destroyed a large part of the physical capital stock in the region.

The Tōhoku earthquake is highly suitable for our purpose because it constitutes a major disruption of global supply chains. As argued by Boehm et al. (2019), the disaster had only a minor impact on the physical infrastructure of the largest Japanese ports (Yokohama, Tokyo, and Kobe). However, the earthquake hit a highly industrialized region, thus leading to a drop in industrial production of about 40% in the disaster-stricken prefectures. In contrast, aggregate industrial production in Japan fell by 13.4% only (Carvalho et al., 2021). This shortfall in manufacturing output propagated through highly entwined global value chains. Boehm et al. (2019) show that the shock led to a drop of one percent in total manufacturing and a two percent fall in the production of durable goods in the U.S. economy. Carvalho et al. (2021) show that the original shock propagated both upstream and downstream through global supply chains. As a consequence, it not only affected direct customers and suppliers of firms in the disaster-stricken region but also had an indirect effect on a much larger set of firms.[7] Thus, the Tōhoku earthquake is an exogenous disruption to global supply chains. From this, we infer our first narrative restrictions:

Narrative Restriction 1. *The supply chain shock takes a positive value in March 2011.*

Undoubtedly, the earthquake was the dominant shock affecting global supply chains in March 2011, which allows us to impose our second narrative restriction:

Narrative Restriction 2. The supply chain shock is the most important driver for the global supply chain pressure index in March 2011.

This restriction implies that in March 2011, the absolute value of the contribution of the supply chain shock is larger than the absolute value of the contribution of any other structural shock.

Suez Canal Obstruction in 2021

On March 23, 2021, Ever Given, one of the largest container ships in the world en route from China to Rotterdam, blocked the Suez Canal. Strong winds wedged the 400-meter-long ship owned by a Japanese company between both canal banks. Consequently, the most important bottleneck for the shipping routes between Europe and Asia was impassable for more than six days before the Ever Given could be freed in the afternoon of March 29, 2021. *Bloomberg* reports that the Ever Given alone transported $1 billion worth of cargo.[8] According to information provided by *Reuters*, more than 200 ships were backed on either entrance of the canal.[9] As reported by *The Financial Times*, the resulting backlog in global supply chains needed weeks to clear.[10]

[7]See also Freund et al. (2022) for an analysis of the adjustment of supply chains in the aftermath of the Tōhoku earthquake.
[8]See https://www.bloomberg.com/link.
[9]See https://graphics.reuters.com/link.
[10]See https://www.ft.com/link.

Furceri et al. (2022) use the Suez Canal obstruction in 2021 as in instrument for the identification of shipping shocks, thus exploiting the exogenous variation in shipping conditions associated with the blockage.[11]

We thus assume that the global supply chain shock in March 2021 has a positive sign:

Narrative Restriction 3. The supply chain shock takes a positive value in March 2021.

Shanghai Backlog in 2022

The outbreak of the Covid-19 pandemic hit the world economy in the spring and summer of 2020. Across countries, authorities imposed lockdowns on public life to contain the virus's spread. These lockdowns led to a sharp deterioration of economic activity in 2020 and a recovery in late 2020 and 2021.[12] In early 2022, China adopted a particularly strict "zero-Covid" strategy for the city of Shanghai, determined to reduce the (relatively small) number of new infections of the Omicron variant of the virus at any cost. As a consequence, economic activity and public life came to a standstill. The lockdown led to a disruption of logistics in the port of Shanghai, China's most important port for container trade. The resulting long delays in container traffic made headline news abroad, and images of the huge backlog of ships went viral on Western social media. Delays in container trade caused severe stress to global manufacturing and "exposed global supply chain strains" (*The Financial Times*).[13] For our purpose, it is important to stress that the strictness of the lockdown imposed in Shanghai was unexpected as authorities responded to the small number of Covid-19 cases with an exceptional determination that far exceeded the intensity of containment measures known from other countries during the pandemic. Hence, the port disruptions can be interpreted as an exogenous variation to global supply chains.[14]

Narrative Restriction 4. The supply chain shock takes a positive value in April 2022.

D. *Inference*

Throughout the paper, our results rely on 10000 draws that satisfy the baseline sign restrictions. We use the same algorithm and the same priors as in Antolín-Díaz and Rubio-Ramírez (2018). The algorithm makes independent draws from a uniform-normal-inverse Wishart posterior of the reduced-form parameters conditional on the baseline and the narrative sign restrictions.

[11] Frohm et al. (2021) discuss the consequences of the Suez Canal incident on global shipping and trade.

[12] Chen and Tillmann (2022) estimate the nature and the magnitude of cross-country spillovers of lockdowns during the Covid-19 pandemic.

[13] See https://www.ft.com/link.

[14] We estimated the model several times with different combinations of the narrative restrictions. For example, we omitted the fourth narrative restriction or the third narrative restriction times. Our results (not reported) are very robust even in these cases.

IV Results

In this section, we discuss our baseline results. We first discuss the responses of the endogenous variables to the supply chain shock identified before. We then look at how informative the narrative information we chose is by computing and comparing the rejection rates for each narrative restriction. Furthermore, we conduct a counterfactual experiment in which we suppress the supply chain shock for the same periods on which we impose our narrative restrictions.

Figure 3: Posterior of the identified supply chain shock

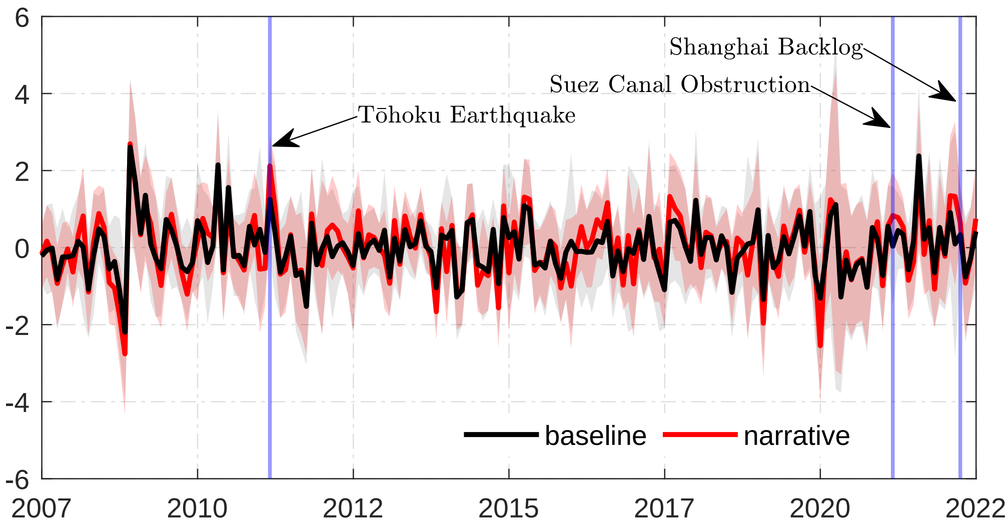

Notes: Posterior of the identified supply chain shock for the baseline model and the baseline model plus narrative restrictions. We also show 90 percent credible bands.

A. *Responses to a Supply Chain Shock*

Figure (3) shows the time series of our identified supply disruptions shock, both for the models that satisfy the baseline sign restrictions and the models that also meet the narrative restrictions. Note that we identify a restrictive shock. The figure also highlights the episodes on which we impose the narrative restrictions. Remember that the three episodes correspond to the clearest cases of adverse supply disruptions, not necessarily the largest realizations of these disruptions.

Figure (4) shows the responses of the endogenous variables to a supply chain shock. The black-solid lines correspond to the medians across all models that satisfy the baseline sign restrictions, while the red-solid lines correspond to the medians across all models that additionally satisfy our narrative restrictions. The light-shaded areas correspond to the 90 percent credible bands for the baseline restrictions, while the red-shaded areas correspond to the 90 percent credible bands that additionally satisfy the narrative restrictions.

Figure 4: The responses to a supply chain shock

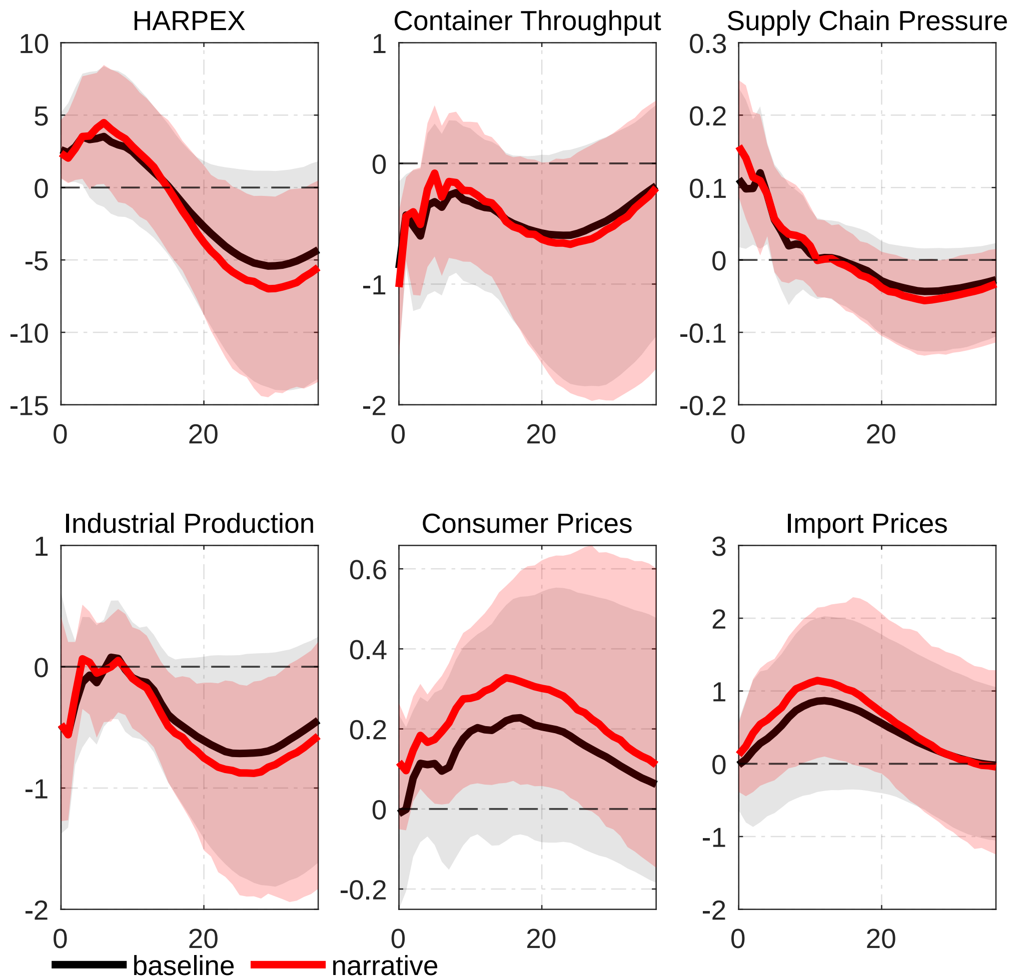

Notes: The black-solid line corresponds to the median of the baseline results, while the red-solid line corresponds to the median of the models that satisfy both the baseline sign restrictions and the narrative restrictions. The light-shaded areas correspond to the 90 percent credible bands for the baseline restrictions, while the red-shaded areas correspond to the 90 percent credible bands that additionally satisfy the narrative restrictions.

Let us first look at the responses of the three variables on which we imposed restrictions. On impact, container prices rise by two percent and continue to rise until they peak after about half a year, with prices rising by three percent. After that, prices for shipping rentals gradually fall below the mean, though this effect is insignificant. We see that the number of containers being processed falls by about one percent on impact and only returns to its mean at the end of the forecast horizon. It should be noted that this effect is significant for the first three periods only. The co-movement of shipping rentals and processed containers complicates the delivery of pre-products necessary to manufacture final products. Consequently, we see that the pressure on supply chains jumps on impact. Supply chain pressure decreases gradually and only returns to its mean after about 15 months.

The information from the narrative restrictions matters: the dynamic adjustment of the HARPEX index becomes statistically significant once we impose the narrative restrictions on top of the conventional sign restrictions. Likewise, the impact response of the GSCPI is greater in magnitude if we impose narrative restrictions on the model.

We now turn to the euro area business cycle variables, i.e. industrial production, consumer, and import prices. We start by discussing the impulse responses based on models that satisfy the baseline sign restrictions, i.e. the traditional sign restrictions. Although we left the responses of these variables unrestricted, the median responses for all variables point in the direction we would expect. Following a supply chain disruption, industrial production in the euro area falls. Both consumer and import prices rise, with the latter increasing almost five times as much as the former. However, all these effects are not significant, as the credible bands include the zero line in all cases. Hence, a supply chain shock identified with conventional sign restrictions does not cause significant business cycle dynamics.

This finding fundamentally changes as soon as we look at the impulse responses based on the draws/models that meet our narrative restrictions in addition to the traditional sign restrictions. The median responses are slightly stronger than those based on models that only satisfy the traditional sign restrictions. In addition, the credible bands are narrower. Most importantly, we now find that a supply chain shock triggers a significant adjustment in industrial production, consumer prices, and import prices. After two years, industrial production is about one percent lower than in the absence of the supply chain shock. The delayed response of industrial activity, which becomes apparent after about one year, could be because, on average, firms can still draw on stocks of intermediate products before supply bottlenecks become binding.

A supply chain shock of one standard deviation increases consumer prices by about 0.3 percent. It is important to highlight that the shock leads to a contraction in economic activity and an increase in prices. Hence, the shock resembles the conventional notion of a supply shock that triggers opposite responses of quantities and prices and supports our interpretation of the shock as a *supply*-side shock. The restrictions on the three shipping variables ensure that the shock is not a conventional supply shock but a supply chain shock identified through international container shipping. The shock strongly raises import prices by one percent. Both price levels increase more strongly when the narrative restrictions are imposed in addition to the sign restrictions.

Overall, we see that a supply chain shock has qualitatively similar effects on consumer prices and industrial production as a conventional supply shock, usually identified by a rise in prices and a fall in production. It should be noted, however, that these variables are unrestricted in our case. Looking only at the traditional sign restrictions, these responses are insignificant. We find that our narrative restrictions shrink the set of admissible parameters, so uncertainty around the impulse responses decreases remarkably. In the following, we shed light on the importance of the events on which our narrative restrictions are based.

B. *Importance of Events*

The baseline results show that narrative restrictions on only three historical events remarkably narrow the credible bands around the impulse responses by shrinking the admissible set of structural parameters. These findings are based on all narrative restrictions jointly. We now assess the importance of each narrative sign restriction for each historical event. For that purpose, we provide rejection probabilities, study the posterior distributions of the shock in the three episodes on which we impose narrative constraints, and simulate counterfactual paths for selected variables.

We evaluate the role of the events by calculating rejection probabilities, i.e. the proportion of draws that satisfy the traditional sign restrictions but result in a rejection of one of the narrative restrictions. The higher this number, the more restrictive the constraints are. Put differently, a higher number for an event implies that this event is particularly informative for our identification. Table (2) reports the probabilities of violating the narrative restrictions.

Table 2: Probabilities of violating the narrative restrictions

Restriction	Tōhoku Earthquake	Suez Canal Obstruction	Shanghai Backlog	Any Event
Sign of the Shock	9	48	32	65
Historical Decomposition	72	–	–	72
Any Restriction	72	48	32	86

Notes: All probabilities are calculated based on the rejection rates from the baseline estimation. Values are in percent and rounded to whole numbers.

It is striking that our baseline restrictions imply the correct sign of the shock during the Tōhoku earthquake for almost all models. Here, only 9% of the draws imply a negative sign. Consequently, 91% of all draws imply the sign we expected. Restricting the historical decomposition for this episode leads to a rejection of 72% of all draws. Interestingly, both restrictions combined have the same probability of violating any of the two narrative restrictions. This means that almost all draws that satisfy our restriction on the historical decomposition also imply the correct sign for the structural shock.

We find that the Shanghai backlog from April 2022 is the most informative event. 65% of all draws, while satisfying the conventional sign restrictions, result in a negative sign of the supply chain shock for this period.

Overall, 14% of all models that satisfy the conventional sign restrictions also satisfy all narrative restrictions. As noted by Antolín-Díaz and Rubio-Ramírez (2018), high probabilities of violating narrative restrictions should not be interpreted as evidence against their validity. In fact, the opposite is true, as high rejection rates tell us that using the baseline specification alone results in many models or structural parameters, respectively, that should be rejected. Thus, we see that although we use only three historical events, they are very informative for achieving identification.

Figure (5) offers a different perspective on the information content of the three episodes. In this figure, we show the posterior distributions for the realization of the supply chain shock in the three episodes that satisfy the baseline sign restrictions and, additionally, the narrative restrictions. We find that the restriction on the size of the shock in March 2011 (Tōhoku earthquake) is innocuous: even if we impose the conventional sign restrictions only, most of the probability mass lies in the positive region. The same is true for the distribution during the Suez Canal obstruction. For the Shanghai backlog, the narrative information concerning the sign of the shock is more restrictive as we force the entire probability mass into the positive region, which would otherwise be more or less symmetrically distributed around zero. This is consistent with the information from the rejections rates, where we showed that the Shanghai backlog is the most informative event.

Figure 5: Supply chain shock for selected episodes

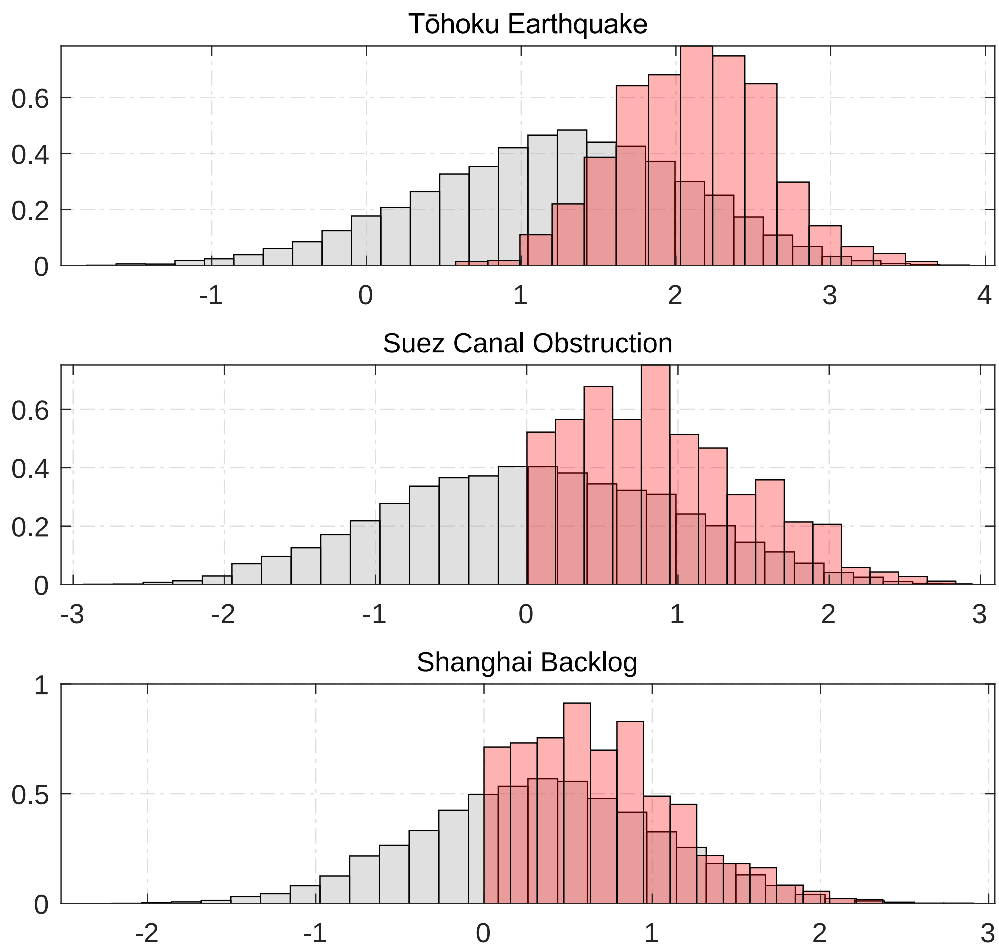

Notes: The light-grey histograms plot the posterior distribution for selected episodes that satisfy the baseline identification, while the light-red histograms plot the posterior distributions that additionally satisfy the narrative sign restrictions.

It is straightforward to examine the extent to which each event contributes to the dynamics of each variable in the past. We address this question by conducting counterfactual simulations. Recall that two of the three events occured near the end of the sample period. Since we have already seen in the impulse responses that a supply chain shock sometimes has significant effects on the endogenous variables only after about one year, such a counterfactual only makes sense for the first narrative event, the Tōhoku earthquake in March 2011. To do so, we assume there is no supply chain shock in March 2011 and re-simulate our endogenous variables conditional on the estimated parameters and the counterfactual shock series.

Figure 6: Suppressing the Tōhoku earthquake

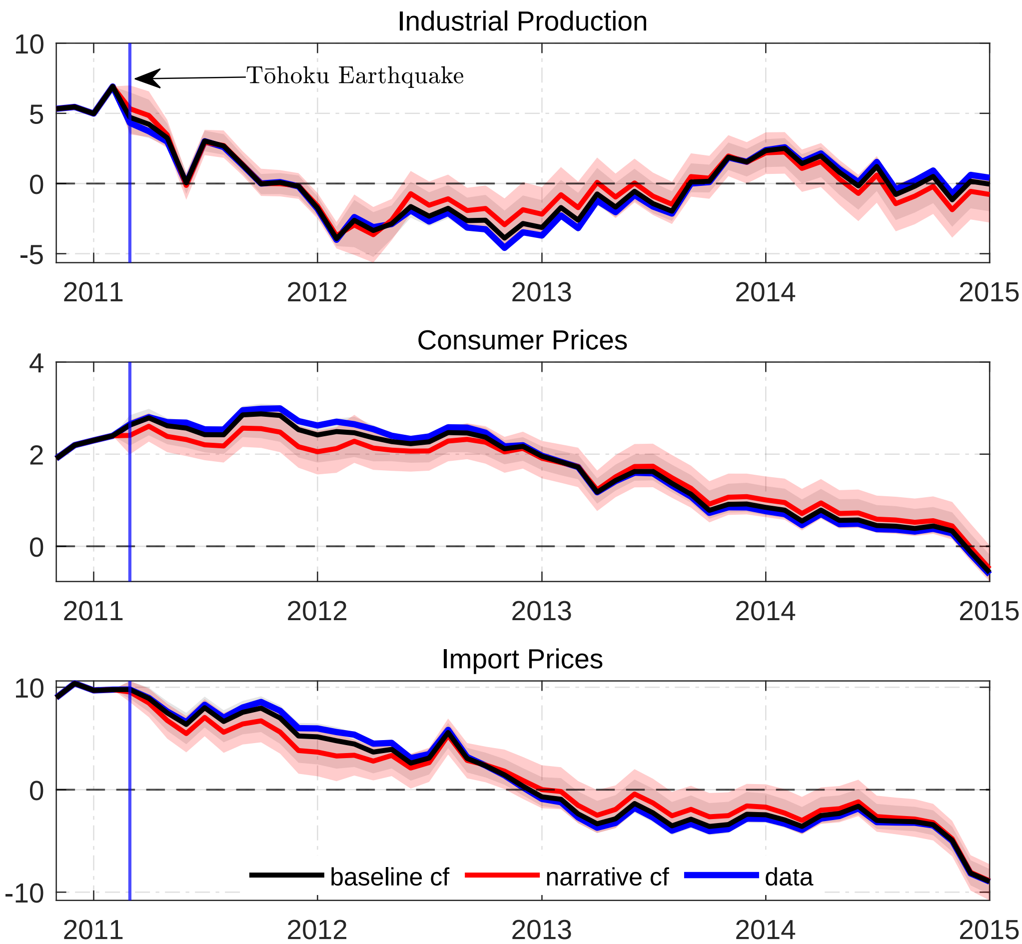

Notes: The black-solid lines correspond to the medians of all counterfactuals that satisfy the baseline sign restrictions, while the red-solid lines correspond to the medians of the models that additionally satisfy the narrative restrictions. The shaded areas correspond to the 90% credible bands. The blue-solid lines correspond to the observed data. The vertical lines mark the occurrence of the Tōhoku earthquake.

Figure (6) shows the counterfactual simulations for industrial production, consumer prices, and import prices for this exercise. The black-solid lines correspond to the medians of all counterfactuals that satisfy the conventional sign restrictions, while the red-solid lines correspond to the medians of the models that additionally satisfy the narrative restrictions. The shaded areas correspond to the 90% credible bands. The blue-solid lines correspond to the observed data. The Tōhoku earthquake is marked by the vertical lines. Note that, since the time series enter our model in logs, we transformed the observed time series and the counterfactuals into year-on-year growth rates to visualize the effect.

If we conduct the counterfactual analysis based on the draws that satisfy the conventional sign restrictions only, we cannot find a sizable effect. Most of the time after March 2011, the actual data and the counterfactual path satisfying the sign restrictions overlap. This changes once we use the draws that also meet the narrative restrictions. The counterfactual path lies above the actual path for industrial production in late 2012 and throughout 2013. Hence, the economic contraction would have been less severe in the absence of the global supply chain shock. Furthermore, in 2012 the counterfactual inflation path lies below the actual path. Hence, the global supply chain shock in March 2011 contributed to higher inflationary pressure.

C. *Other Price Indices*

Our baseline model includes aggregate consumer and import prices for the euro area. We now want to compare the inflationary effect of the supply chain disruption across alternative price level indicators. We re-estimate the baseline model six times and, each time, replace the import price index with one of six alternative price series. In Figure (7), we only show these series' responses. We use import prices as well as producer prices (all in natural logs) for three NACE Rev. 2 categories: manufacturing, intermediate goods, and consumer goods, excluding food.[15]

We find that import and producer prices respond similarly to supply chain disruptions for each sectoral category. Interestingly, in each of the six cases, we see that adding the narrative restrictions leads to a stronger response. This means that the baseline sign restrictions, which lead to results that are partially inconsistent with the established narratives, underestimate the effect of supply chain disruptions.

D. *The Explanatory Power of the Supply Chain Shock*

The previous section sheds light on the importance of the realization of the supply chain shock selected in selected episodes for the evolution of the business cycle in the euro area. We now want to generalize the analysis and quantify the contribution of the supply chain shock for the six endogenous variables globally, i.e. outside specific events. For that purpose, we decompose the forecast error variance of each endogenous variable into the contribution of the supply chain shock and all the remaining shocks. Table (3) reports this decomposition for selected forecast horizons. The shock

[15]The online appendix to this paper contains the references to each data series in the ECB's Statistical Data Warehouse.

Figure 7: The responses to a supply chain shock: alternative price indices

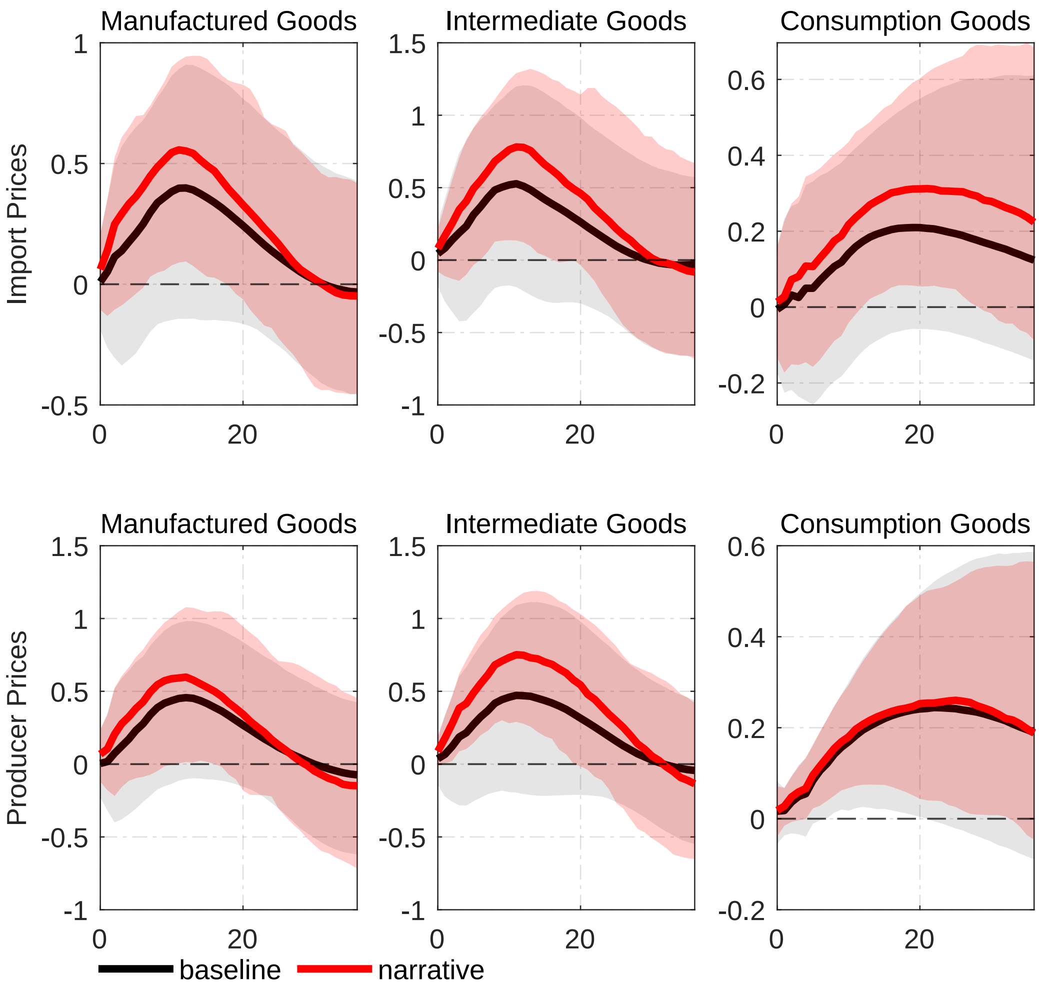

Notes: The black-solid line corresponds to the median of the baseline results, while the red-solid line corresponds to the median of the models that satisfy both the baseline sign restrictions and the narrative restrictions. The light-shaded areas correspond to the 90 percent credible bands for the baseline restrictions, while the red-shaded areas correspond to the 90 percent credible bands that additionally satisfy the narrative restrictions. All units are in percent.

explains a large share of the supply chain pressure index. This finding supports our identification of a supply chain shock. The explanatory power for the supply chain pressure index over a horizon of one year increases from 19% to 24% once we add narrative restrictions on top of conventional sign restrictions. This also supports our choice of historical episodes from which we infer information for identification. Over a horizon of one year, the shock explains 11% of the dynamics of industrial production and even 13% (12%) of the dynamics of consumer prices (import prices). Put differently, the shock explains a large fraction of business cycle dynamics. In addition, the explanatory power increases if we add narrative restrictions. This is particularly true for the two price series, for which the explanatory power almost doubles when narrative restrictions are imposed.

Table 3: Forecast error variance decomposition

Baseline Restrictions

h	HARPEX	GSCPI	Container Throughput	IP	Consumer Prices	Import Prices
0	15	13	19	9	8	7
3	10	17	16	11	12	8
6	10	19	12	10	13	9
12	10	19	11	11	13	12
24	13	19	17	18	15	14
36	16	20	18	19	15	15

Baseline plus Narrative Restrictions

h	HARPEX	GSCPI	Container Throughput	IP	Consumer Prices	Import Prices
0	12	25	28	7	9	4
3	9	24	18	8	18	7
6	11	25	11	8	21	11
12	12	24	10	9	29	19
24	16	25	16	21	32	22
36	22	26	18	24	30	23

Notes: Share of the forecast error variance at horizon h, which can be explained by shocks identified through traditional sign restrictions (upper block) and through traditional sign restrictions that also satisfy the narrative restrictions (lower block). Values are rounded to whole numbers.

V Sensitivity Analysis

This section examines the robustness of our findings and presents a battery of sensitivity checks. First, we estimate the model for the G6 economies. In a second check, we compare our global supply shock with prominent structural shocks from the literature. In a third check, we investigate whether geopolitical risk drives our results. We also investigate how monetary policy reacts to global supply chain distortions. In a final robustness exercise, we relax (tighten) the minimum duration of the sign restrictions.

A. *Identifying a Global Supply Chain Shock for the G6 Economies*

We identify a global supply chain shock based on a VAR model that, among other variables, includes container throughput, industrial production and consumer prices from the euro area. Does the inclusion of euro area information confounds the interpretation of the shock as a *global* shock? In order to address this concern, we re-estimate the baseline model but replace euro area variables by the corresponding variables from the G6 economies, i.e. the U.S., Canada, Japan, Germany, France and Italy. Due to the lack of container data for the UK, we cannot estimate the model for the full set of the G7 economies and have to leave out the UK.[16] We add container throughput for all ports of the G6 and use real GDP-weights to calculate industrial production and consumer prices for the G6. The series are seasonally adjusted. All restrictions on the model remain identical to the estimation of the baseline model.

The resulting shock series exhibits a correlation of 0.70 (baseline restrictions) and 0.61 (narrative restrictions) with the corresponding series of the model estimated on euro area data. The high correlation supports the interpretation of our benchmark shock as *global*. Not surprisingly, the correlation is less than perfect as the narrative restrictions are not equally tight for all six economies.

The impulse responses shown in Figure (8) are very similar to our findings for the euro area. In the G6 economies, the global supply chain shock causes a significant fall in production and a strong and persistent increase in consumer prices. Both responses are significant if we impose the narrative restrictions. The similarity of the responses of the G6 economies and the euro area corroborates the notion that we indeed identify a global supply chain shock.

B. *Comparing the Global Supply Chain Shock with Other Shocks*

This paper identifies a global supply chain shock using information from global container shipping. It is important that our shock does not accidentally pick up other supply-side distortions, which might also affect shipping. In order to understand the properties of our shock series, we now compare it with alternative shocks. We select a range of shock series from the recent literature and compute the correlation with our estimated shock. In particular, we focus on oil supply shocks, a natural candidate for global supply-side disturbances. We take these series from Kilian and Murphy (2012), Baumeister and Hamilton (2019), and Känzig (2021), which are the most prominent papers on oil supply shocks in the literature.

Table (4) contains the contemporaneous correlation with our supply chain shock. We find that all the key oil supply shocks from the literature are uncorrelated with our supply chain shock. This supports the notion that our shock reflects specific distortions to the global economy which are not attributable to unexpected changes in oil supply. The only significant correlation at the 5% level is with the oil inventory demand shock from Baumeister and Hamilton (2019).

Finally, one point deserves special emphasis. Many of the shocks in Table (4) have, similar to our supply chain disruptions, the same causal consequences as classical supply shocks, i.e., a negative co-movement of output and prices. For instance, we expect an increase in HARPEX and a decrease in container throughput when an oil supply shock hits the economy. These effects would also lead to an increase in consumer prices and a drop in industrial production. Importantly, however, for all shocks in Table (4), the signs of the respective shocks do not consistently match the established narratives laid

[16]As we could not find import prices for all G6 economies, our model consists of five variables only.

Figure 8: The responses to a supply chain shock: G6 economies

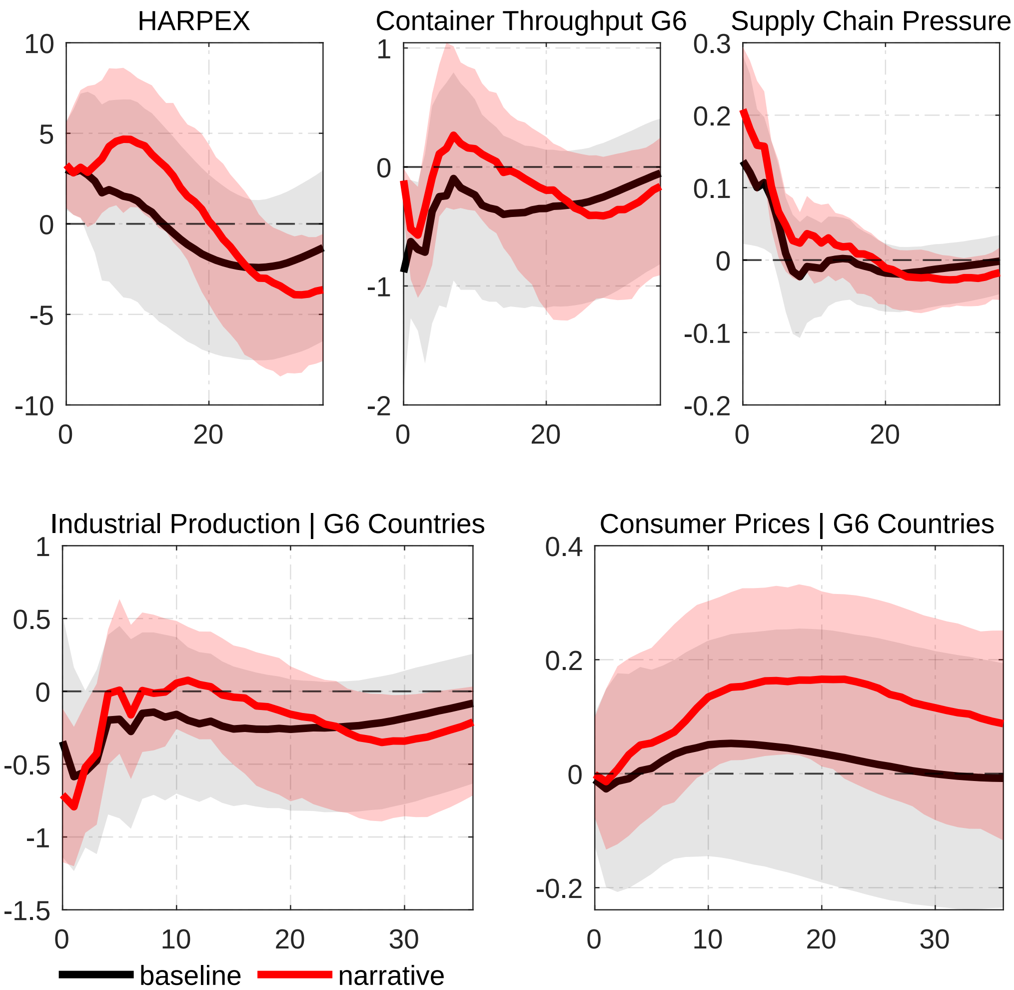

Notes: The black-solid line corresponds to the median of the baseline results, while the red-solid line corresponds to the median of the models that satisfy both the baseline sign restrictions and the narrative restrictions. The light-shaded areas correspond to the 90 percent credible bands for the baseline restrictions, while the red-shaded areas correspond to the 90 percent credible bands that additionally satisfy the narrative restrictions.

out in Section III. Hence, they are not just uncorrelated with the supply chain shock, but also have the wrong sign in the three episodes we use for identification.

C. *A Note on Estimation Uncertainty*

As is standard in the econometric literature, we use the posterior median response as our response function. Inoue and Kilian (2022) show that conventional impulse response estimators such as the posterior median response function or the posterior mean response function are not necessarily the Bayes estimator of the impulse response vector. They derive the Bayes estimator of vectors of structural VAR impulse responses under a range of alternative loss functions.

In this subsection, we use the algorithm of Inoue and Kilian (2022) under an additively separable absolute loss function and calculate the Bayes estimator of our impulse

Table 4: Correlation of the global supply chain shock with other shocks

Shock	Source	ρ	p−value	n	Sample
Oil Supply	Kilian and Murphy (2012)	0.087	0.284	154	Jul 2007 - Apr 2020
	Kilian and Murphy (2012), Antolín-Díaz and Rubio-Ramírez (2018)	0.138	0.090	154	Jul 2007 - Apr 2020
	Baumeister and Hamilton (2019)	−0.005	0.950	176	Jul 2007 - Feb 2022
	Känzig (2021)	0.083	0.279	180	Jul 2007 - Jun 2022
Oil-Specific Demand	Kilian and Murphy (2012)	−0.074	0.362	154	Jul 2007 - Apr 2020
	Kilian and Murphy (2012), Antolín-Díaz and Rubio-Ramírez (2018)	−0.070	0.391	154	Jul 2007 - Apr 2020
Global Risk	Georgiadis et al. (2021)	−0.067	0.426	144	Jul 2007 - Jun 2019
Global Uncertainty	Caggiano and Castelnuovo (2022)	0.027	0.747	142	Jul 2007 - Apr 2019
Oil Consumption Demand	Baumeister and Hamilton (2019)	0.030	0.689	176	Jul 2007 - Feb 2022
Oil Inventory Demand	Baumeister and Hamilton (2019)	0.162	0.032	176	Jul 2007 - Feb 2022
Economic Activity	Baumeister and Hamilton (2019)	−0.056	0.461	176	Jul 2007 - Feb 2022

Notes: Correlation of the global supply chain shock with selected shock from the literature, where ρ is the Pearson correlation coefficient. The p−value corresponds to the test whether the correlation is equal to zero and n is the sample size. From Känzig (2021), we take the series of oil supply surprises. As for the Kilian and Murphy (2012) oil supply shock and the Kilian and Murphy (2012) oil demand shock, we updated the estimation sample until April 2020 and estimated the model with the same specification (same lag length and same identification strategy) as in the original paper. That is, we use the month-on-month growth rate of global crude oil production, an index of real economic activity, and the log of the real oil price. We also estimated the Kilian and Murphy (2012) model for the updated sample using the same narrative sign restrictions as in Antolín-Díaz and Rubio-Ramírez (2018), except for restrictions on the elasticity bounds.

response vectors accordingly. Figure (9) shows our baseline results as in Figure (4) and additionally plots the Bayes estimator (dotted lines) under the loss function mentioned above.

Overall, it stands out that the Bayes estimator nearly matches the pointwise median response function. We find a similar picture for all the other impulse response functions presented in this paper.

Finally, Inoue and Kilian (2022) argue that conventional pointwise quantile error bands are not a valid measure of the estimation uncertainty about the impulse response vector, as they ignore the mutual dependence of the responses, thus understating the estimation uncertainty about the impulse response vector.

Using their algorithm to derive joint credible regions under the same loss function as above, we find that we have the same significant results as in our main results only on the 68% credible region.[17] However, on the 90% credible region, we only have boarderline significant effects. As described in Inoue and Kilian (2022), we thus also find that pointwise quantile error bands seem to understate the estimation uncertainty about the respective impulse response vector.

D. *Supply Chain Disruptions and Geopolitical Risk*

Geopolitical tensions such as military threats, outright wars, or terrorist attacks could also interrupt supply chains. Thus, we need to ensure that we do not mislabel changes in geopolitical risk as an exogenous supply chain shock. To help us separate these two dimensions, we draw on the work of Caldara and Iacoviello (2022). These authors construct an index of geopolitical risk based on an extensive analysis of a century of

[17] In order to save space, we do not present the results here. The results are available upon request.

Figure 9: The responses to a supply chain shock

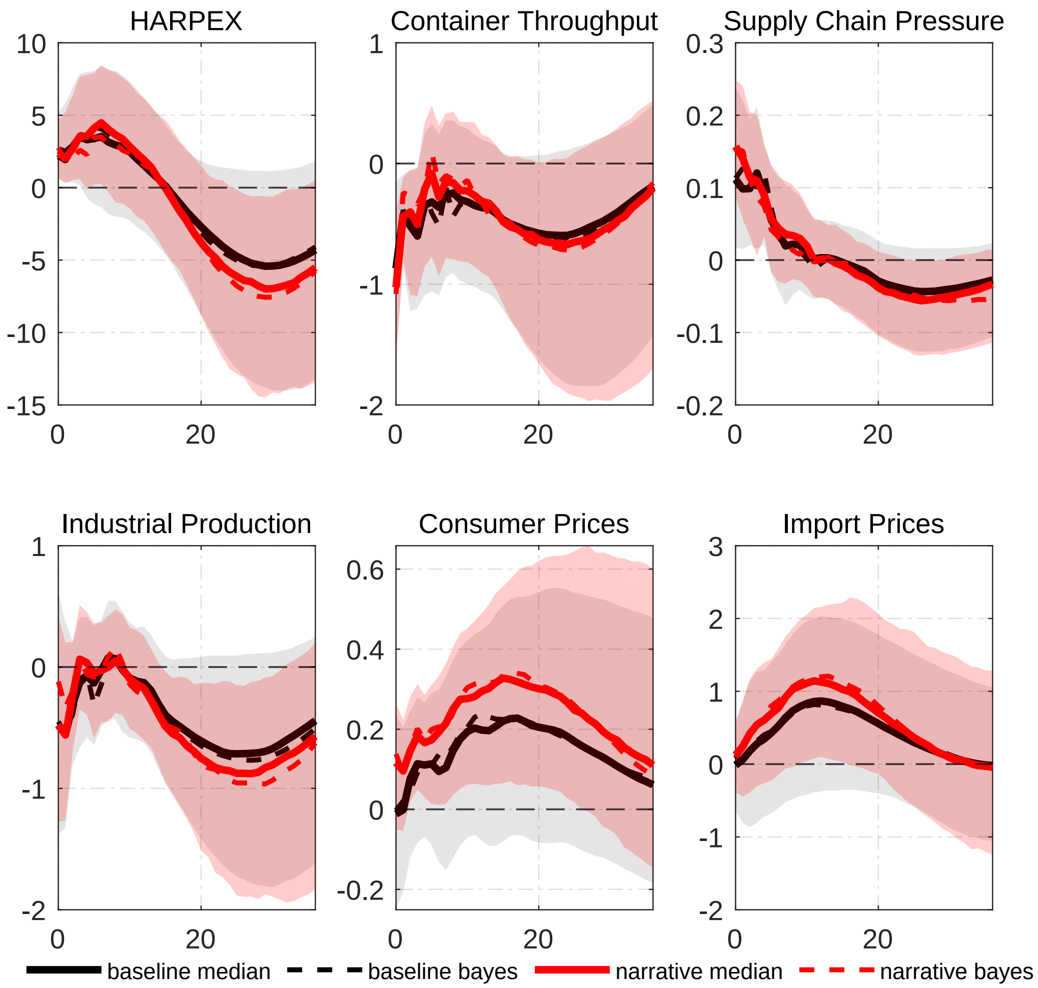

Notes: The black-solid line corresponds to the median of the baseline results, while the red-solid line corresponds to the median of the models that satisfy both the baseline sign restrictions and the narrative restrictions. The red (black) solid line corresponds to the Bayes estimator using the algorithm of Inoue and Kilian (2022) under an additively separable absolute loss function. The light-shaded areas correspond to the 90 percent credible bands for the baseline restrictions, while the red-shaded areas correspond to the 90 percent credible bands that additionally satisfy the narrative restrictions.

newspaper coverage. We re-estimate our baseline model but replace import prices with the log of the geopolitical risk index. In addition, we impose a zero restriction on the response of this index to our shock. The supply chain shock raises the HARPEX and the GSCPI, reduces the number of containers processed in the euro area but leaves geopolitical risk unchanged on impact. The latter restriction distinguishes a supply chain disruption from a spike in geopolitical risk.

The estimated impulse responses, see Figure (10), are virtually unchanged compared to our previous findings. The responses of industrial production and consumer prices in the euro area are still highly significant. As geopolitical risk remains constant on impact, we are not confusing supply chain disruptions with geopolitical tensions. Nevertheless, we see that geopolitical risk rises after a shock, which is another adverse

Figure 10: The responses to a supply chain shock: the role of geopolitical risk

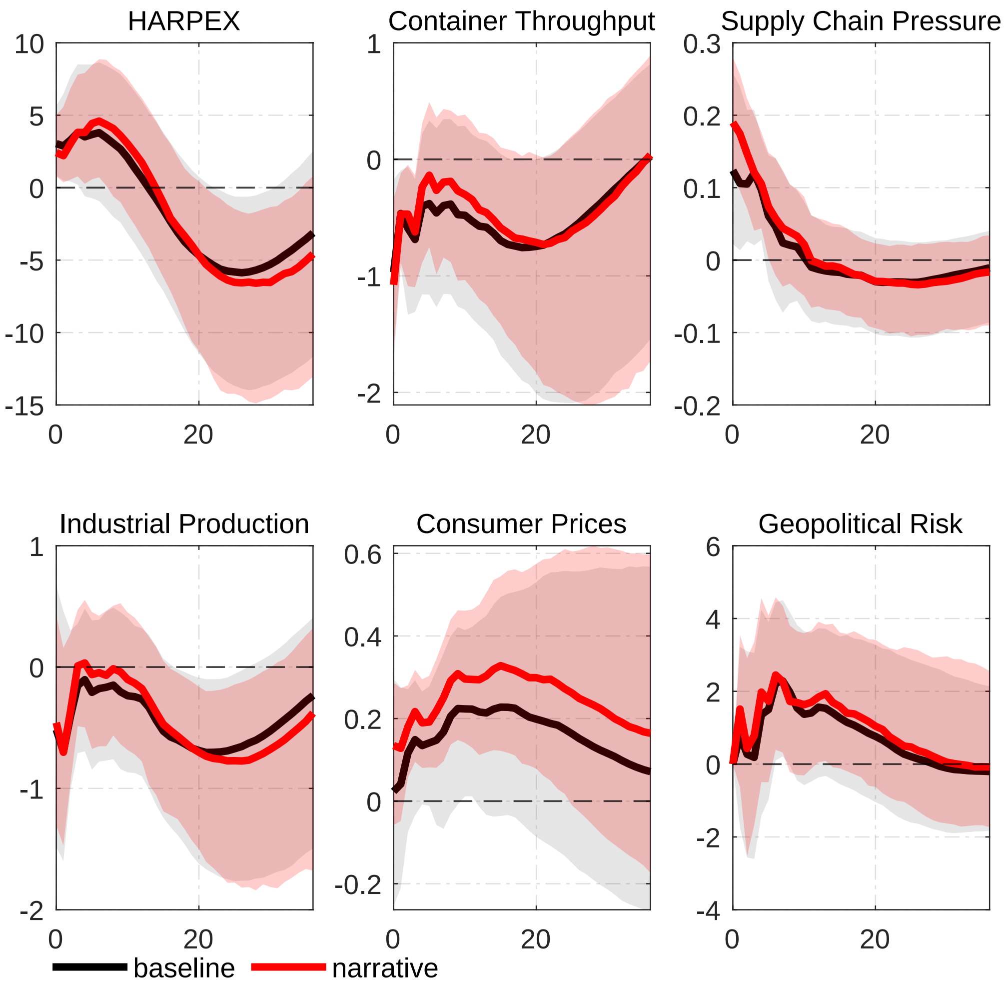

Notes: The black-solid line corresponds to the median of the baseline results, while the red-solid line corresponds to the median of the models that satisfy both the baseline sign restrictions and the narrative restrictions. The light-shaded areas correspond to the 90 percent credible bands for the baseline restrictions, while the red-shaded areas correspond to the 90 percent credible bands that additionally satisfy the narrative restrictions.

consequence of disrupted supply chains.

E. *The Response of Monetary Policy*

The supply chain shock identified in this paper elicits opposite responses of economic activity and prices, a textbook feature of supply-side disruptions. We now study the European Central Bank's (ECB) response to the increase in prices and the economic contraction. As monetary policy primarily affects the demand side of the economy, the response to supply-side disturbances is intricate. Importantly, we do not take a stand on how monetary policy *should* respond, but rather estimate the ECB's *actual* response. To so do, we replace the sixth variable in our VAR model, the series of import prices, with the ECB's policy instrument. Since the ECB adopted a range of unconventional policies, such as asset purchases and forward guidance when the euro area was at the effective lower bound, we use the shadow short-rate calculated by Wu and Xia (2016) as a summary measure of both conventional and unconventional monetary policy. All restrictions on the model remain unchanged.

Figure 11: The responses to a supply chain shock: the role of monetary policy

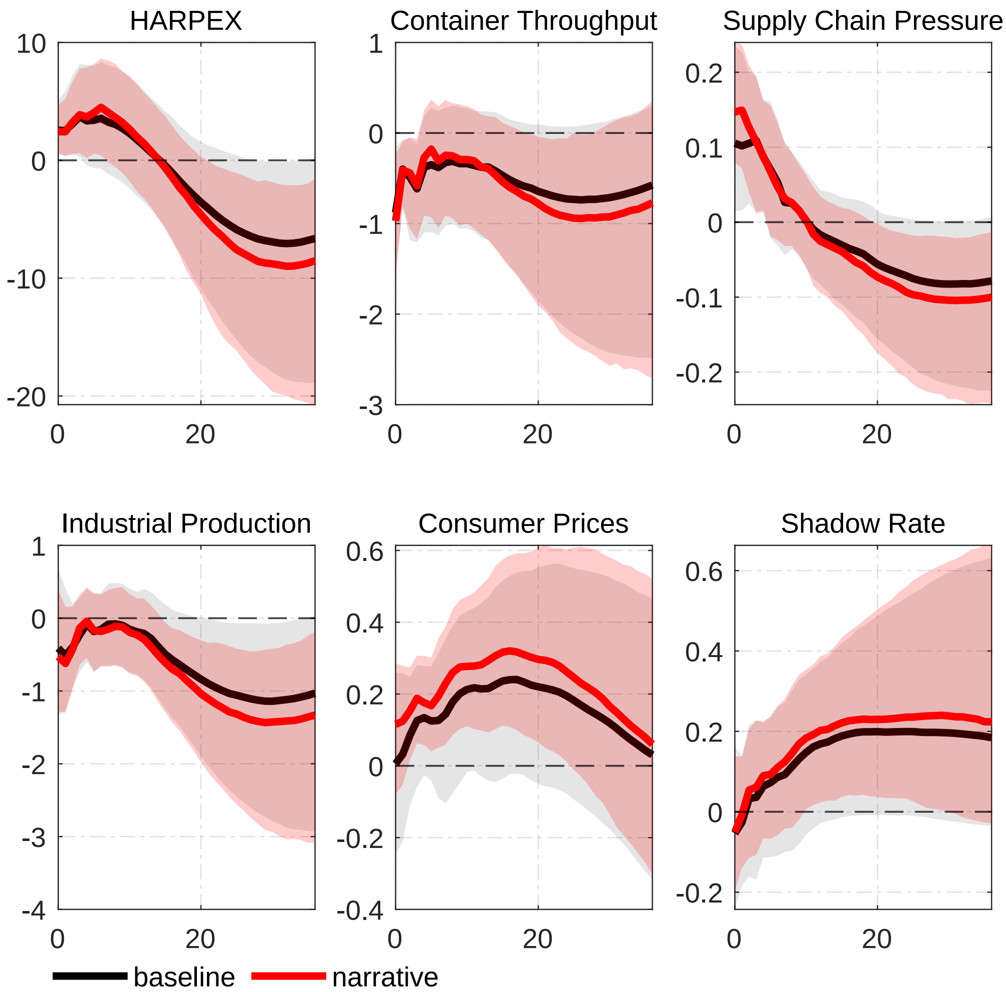

Notes: The black-solid line corresponds to the median of the baseline results, while the red-solid line corresponds to the median of the models that satisfy both the baseline sign restrictions and the narrative restrictions. The light-shaded areas correspond to the 90 percent credible bands for the baseline restrictions, while the red-shaded areas correspond to the 90 percent credible bands that additionally satisfy the narrative restrictions.

Figure (11) documents the estimated impulse response functions. We find that the ECB tightens monetary conditions by about 20 basis points. This response is relatively persistent even as inflationary pressure subsides.

F. *Changing the Duration of the Sign Restrictions*

In our baseline specification, we restrict the effect of a supply chain shock to last for at least one quarter, i.e. on impact and for two consecutive months. However, since the minimum duration of a shock effect is generally difficult to justify on an ad-hoc basis, we now show alternative specifications for the number of periods at which the sign restrictions apply. We try each combination ranging from *only on impact* up to *on impact and for the five consecutive months*, i.e. for two full quarters. That is, we estimate five additional models with different durations of sign restrictions, keeping anything else identical to the benchmark case.

Figure 12: The responses to a supply chain shock: alternative restriction horizon

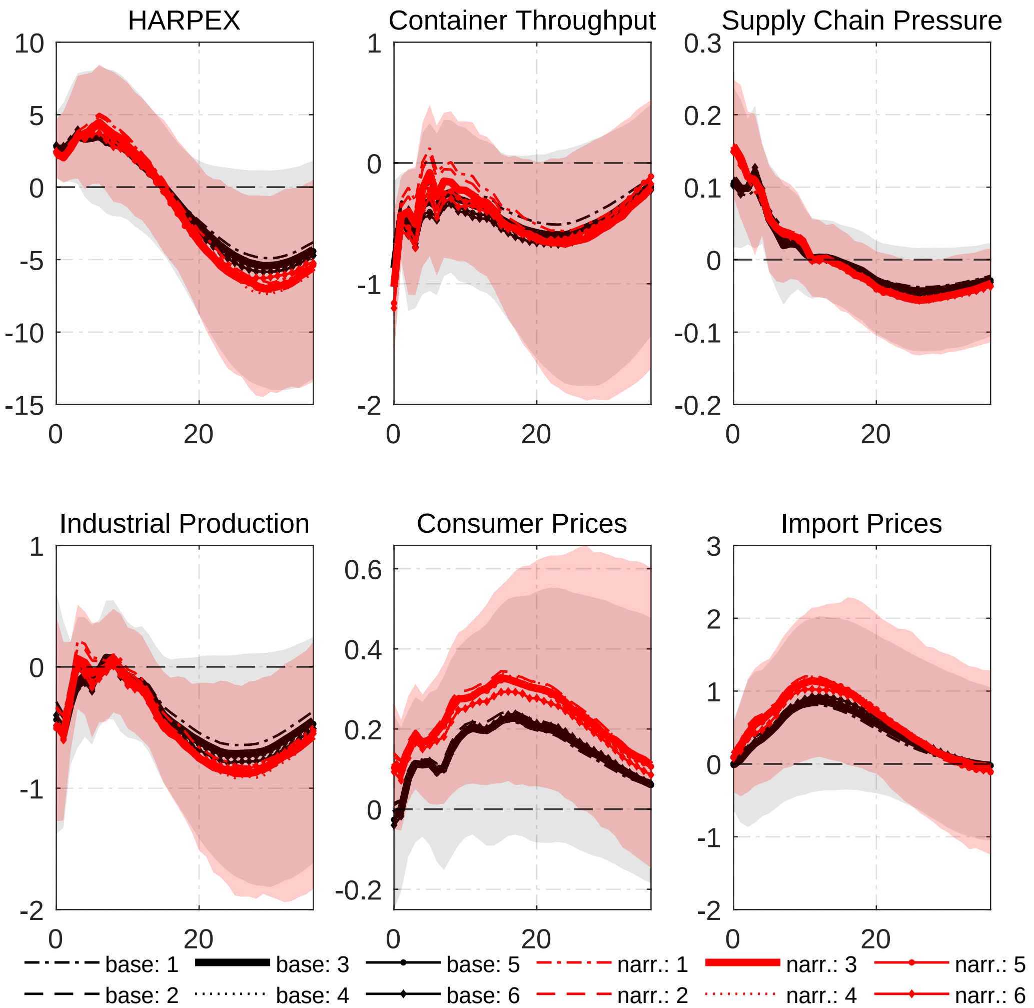

Notes: The black-solid lines correspond to the medians of the baseline results for different restriction horizons, while the red-solid lines correspond to the medians of the models that satisfy both the baseline sign restrictions and the narrative restrictions. The light-shaded areas correspond to the 90 percent credible bands for the baseline restrictions (baseline specification), while the red-shaded areas correspond to the 90 percent credible bands that additionally satisfy the narrative restrictions.

Figure (12) shows our baseline results for comparison and additionally plots the medians for different specifications regarding the duration of the sign restrictions. As can be seen, changing the duration does not make a big difference. For all variables, we find that the medians are very close to each other both for the baseline restrictions and for the baseline restrictions that also satisfy narrative restrictions. We only show the error bands for the baseline specification as our results are almost identical in terms of significance.

VI Supply Chain Shocks from China and the Rest of the World

So far, we have looked at a global supply chain shock, i.e. an unexpected disruption in global supply chains. Hence, our baseline model does not allow us to pin down the regional origin of the shock. One region appears particularly interesting when analyzing supply chains: China. This is motivated by the fact that China is the manufacturing powerhouse of the world economy and central to global sourcing and production chains. This is also inspired by the recent frictions in container trade in Chinese ports in the spring of 2022, which are the result of the zero-Covid policy adopted by authorities in Shanghai. In this section, we, therefore, go one step further and distinguish a supply chain shock emanating from China from a supply chain shock originating in the rest of the world (RoW).

In order to distinguish a Chinese supply chain disruption from a disruption originating in the rest of the world, we estimate a modified VAR model. We include the HARPEX, euro area container throughput, industrial production, and consumer prices. Furthermore, we split the GSCPI into two components and include each of them in the VAR model. The first component is the regional supply chain pressure index for China, which we take from Benigno et al. (2022). The second component measures supply chain pressure in the rest of the world, i.e. the world economy excluding China. Denote sp_t^{Global} the Global Supply Chain Pressure Index and sp_t^{China} the regional supply chain pressure index for China. We obtain the RoW-index as the estimated residual from the following regression

$$sp_t^{Global} = c + \omega sp_t^{China} + \varepsilon_t.$$

The RoW-index is orthogonal to the Chinese index by construction. In the online appendix, we derive an alternative RoW-index where we subtract the explained variation of the Chinese subindex in the Global Supply Chain Pressure Index, thus allowing for common factors.

The complete set of sign restrictions is summarized in Table (5). Importantly, a supply chain shock emanating from China and outside China should raise the HARPEX and lower container throughput in the euro area. The idea of separating a supply chain disruption from China from that from the rest of the world is based on exploiting the co-movement of the respective supply chain pressure on impact. We, therefore, impose a zero restriction on the RoW-index. For instance, a supply chain disruption emanating from China should raise the Chinese supply chain pressure index on impact, while it should not affect the respective RoW-index. By the same token, a supply chain disruption originating outside of China should raise the RoW-index, while the shock should not contemporaneously affect the Chinese supply chain pressure index. We impose a zero restriction on the Chinese supply chain pressure index in this case. As we allow for two shocks, we need to modify the set of narrative restrictions.

Specifically, we assign Narrative Restrictions 1 to 3 to the RoW-supply chain shock. Narrative Restriction 4 applies to the Chinese supply chain shock only.

Table 5: Baseline sign restrictions: Rest of the World vs. China

shock	HARPEX	Throughput	SCPI RoW	SCPI China	IP	Consumer Prices
Rest of the World	+	–	+	0		
China	+	–	0	+		

Notes: The restrictions on the HARPEX and the Northrange container throughput hold on impact and for two consecutive months. The non-negativity restriction on the regional supply chain pressure indices are also imposed on impact an for two consecutive months. The zero restrictions are imposed on impact only.

The resulting impulse response functions are shown in Figures (13) and (14). A few things stand out. First, a supply chain disruption originating either in China or the rest of the world has very similar effects on both HARPEX and container throughput. Interestingly, these effects are also very similar to the results we have already found in our baseline model, where we did not control the shock's origin. Moreover, the impulse responses of euro area industrial production look very similar in both cases and similar to the results from our baseline model.

Second, there is a noticeable difference in the response of consumer prices. We find no significant effect on prices at the 90 percent confidence level after a shock emanating from China.[18] After a supply chain shock originating outside of China, however, we find that consumer prices increase significantly. It stands out that this reaction looks both quantitatively and qualitatively very similar to the reaction found in the baseline model in Section IV.

Third, and perhaps most notable, we find that narrative information helps us identify supply chain disruptions, in particular, that emanate from the rest of the world. One possible reason is that three out of four narrative sign restrictions refer to the shock originating outside China. Consequently, the influence of these restrictions on the set of admissible parameters is most noticeable here.

The responses of our endogenous variables are derived from the estimated parameters by assuming a supply chain disruption of one standard deviation. It is also worth looking at an alternative way to present the results, which is not just based on the parameter estimates, but also the innovations themselves. Therefore, we first look at the forecast error variance decomposition and see how much of the variance in forecast errors is explained by Chinese supply chain disruption and disruption emanating

[18] On the 68 percent level, however, we find that consumer prices in the euro area react significantly to a Chinese supply chain disruption.

Figure 13: The responses to a supply chain shock from the rest of the world

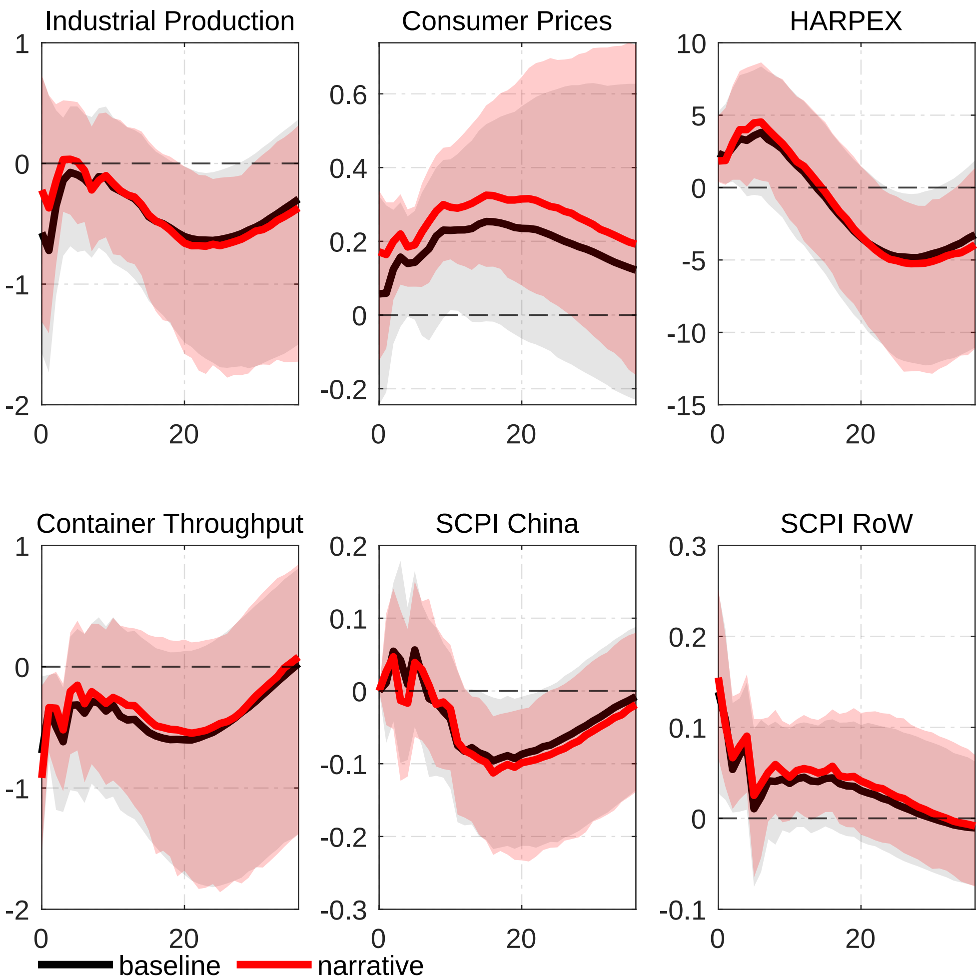

Notes: The black-solid line corresponds to the median of the baseline results, while the red-solid line corresponds to the median of the models that satisfy both the baseline sign restrictions and the narrative restrictions. The light-shaded areas correspond to the 90 percent credible bands for the baseline restrictions, while the red-shaded areas correspond to the 90 percent credible bands that additionally satisfy the narrative restrictions.

from the rest of the world, respectively. The upper panel of Table (6) reports the forecast error variance decompositions for those models that satisfy the baseline sign restrictions, while the bottom panel reports the results for the models that additionally satisfy the narrative sign restrictions. Let us concentrate on the bottom half of the table, which shows the explanatory power when both traditional and narrative sign restrictions are imposed.

Again, three things stand out. First, both Chinese shocks and shocks originating outside China seem equally important for container price fluctuations and container throughput in the euro area. For instance, our results imply that over the horizon of one year, shocks originating in China (outside of China) contribute to 10% (8%) of the fluctuations of the HARPEX and 8% (7%) of container throughput in the euro area, respectively. Unsurprisingly, we find that Chinese supply chain disruptions explain

Figure 14: The responses to a supply chain shock from China

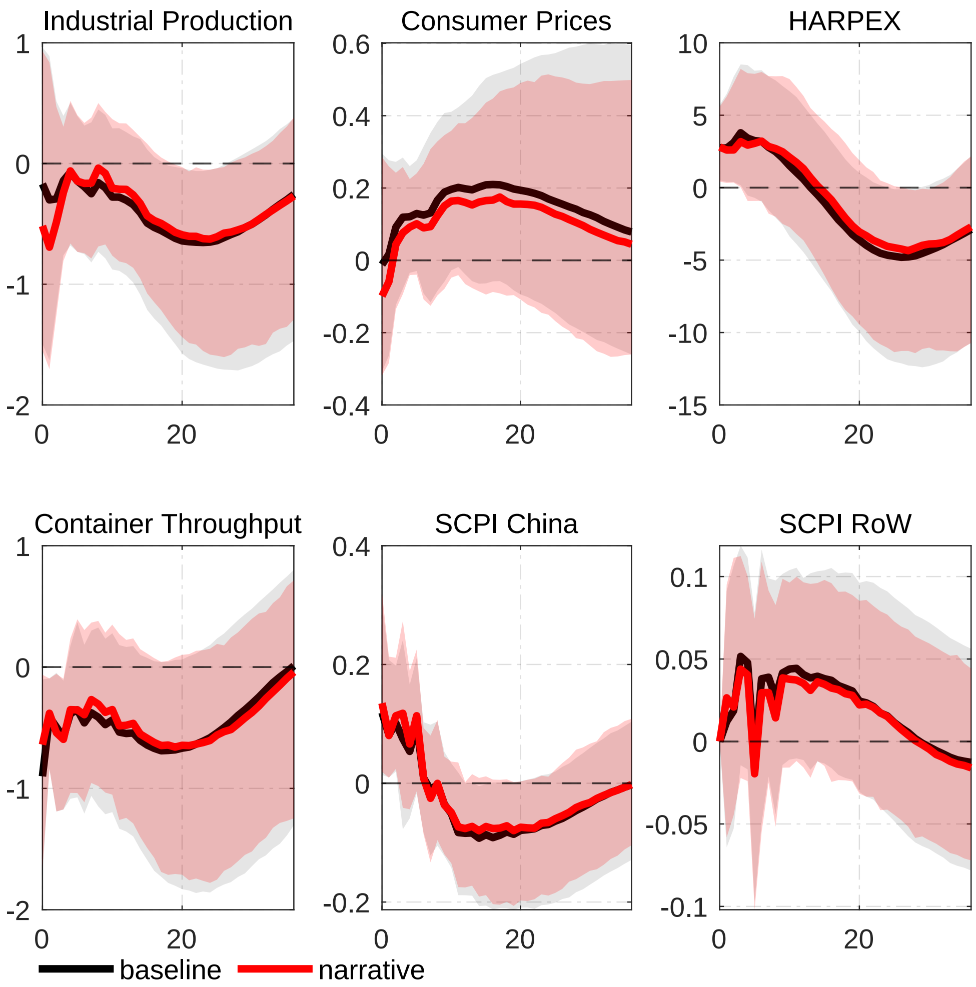

Notes: The black-solid line corresponds to the median of the baseline results, while the red-solid line corresponds to the median of the models that satisfy both the baseline sign restrictions and the narrative restrictions. The light-shaded areas correspond to the 90 percent credible bands for the baseline restrictions, while the red-shaded areas correspond to the 90 percent credible bands that additionally satisfy the narrative restrictions.

significantly more of the fluctuation in the Chinese supply chain pressure index than shocks originating outside China and vice versa.

Second, and more important for our analysis, is the apparent divergence in the importance of our two shocks for the fluctuation of industrial production and consumer prices, respectively. We find that Chinese supply chain disruptions are a significantly more important driver of unexpected movements in industrial production than shocks originating in the rest of the world. Over the horizon of one year, Chinese supply chain disruptions explain 16% of the fluctuation in industrial production in the euro area, whereas shocks originating outside China explain only 7%. One possible explanation is that intermediate goods from China are essential for production in the euro area. When the supply of these goods is distorted, firms cannot substitute Chinese intermediate goods by others, even at higher prices, and thus have to cut

Table 6: Forecast error variance decomposition: RoW vs China

Baseline Restrictions												
h	HARPEX		SCPI RoW		SCPI China		Container Throughput		IP		Consumer Prices	
	RoW	China	RoW	China	RoW	China	RoW	China	RoW	China	RoW	China
3	9	9	14	4	2	8	14	13	7	18	12	15
6	10	8	17	6	4	10	10	10	7	15	16	16
12	9	8	19	8	6	14	9	9	9	14	18	17
24	12	13	19	9	11	16	13	15	14	20	18	15
36	15	15	18	10	12	16	15	16	16	20	17	13

Baseline plus Narrative Restrictions												
h	HARPEX		SCPI RoW		SCPI China		Container Throughput		IP		Consumer Prices	
	RoW	China	RoW	China	RoW	China	RoW	China	RoW	China	RoW	China
3	10	9	17	4	2	11	15	10	5	25	22	12
6	11	7	20	6	3	12	10	7	5	20	28	12
12	10	8	23	7	6	15	8	7	7	16	34	11
24	14	12	23	8	13	16	12	13	13	19	31	10
36	18	14	22	9	14	15	13	15	17	19	28	9

Notes: Share of the forecast error variance for horizon *h*, which can be explained by shocks identified through traditional sign restrictions (upper block) and through traditional sign restrictions that also satisfy the narrative restrictions (lower block). Values are rounded to whole numbers.

production. When the supply of goods from other parts of the world is distorted, firms can substitute the missing inputs, although at a higher price, such that production falls only moderately. As a consequence, consumer prices do not respond much in the first case, but increase in the latter.

Third, we see again that the narrative sign restrictions concerning the Tōhoku earthquake in 2011, the Suez Canal obstruction in 2021, and the Shanghai backlog in 2022 strongly affect our results. Once these established narratives are taken into account for inference, we see that the baseline restrictions alone lead to results that (1) underestimate (overestimate) the effects shocks originating in the rest of the world (China) have on consumer prices in the euro area. Our results imply that, over a horizon of one year, shocks emanating from the RoW account for 34% of the fluctuation of consumer prices, while supply chain disruptions originating in China only account for 11%.

Figure (15) presents the historical decomposition of the two variables we are mostly interested in, euro area industrial activity and consumer prices, into the contributions of global and Chinese supply chain shocks. Interestingly, the contributions of both shocks to industrial production usually have the same sign: if global supply chain disruptions contain industrial activity, so do disruptions from China. However, from 2018 until 2020 and in late 2021, we see that the contributions have different signs. Interestingly, we observe that Chinese supply chain disruptions seem more important than disruptions in the RoW. This fits with the results from forecast error variance decomposition.

Also, for prices, the contributions of the two shocks mostly have the same sign. It is particularly striking that Chinese shocks exerted downward pressure on consumer

Figure 15: Historical decomposition: RoW vs China

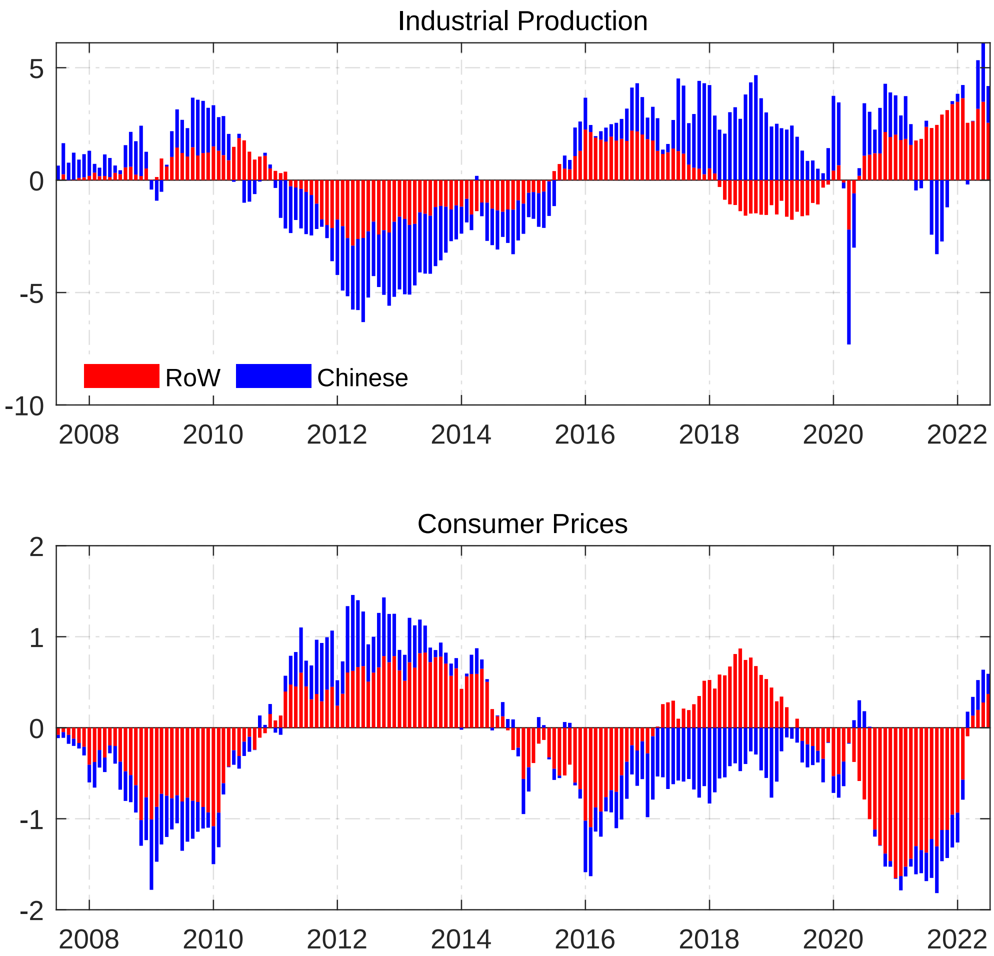

Notes: All bars rely on the median over all draws.

prices in the euro area from 2008-2010 and from 2016 to 2019. In 2020 and 2021, i.e. during the pandemic, both forces had a large negative effect on prices, with RoW-shocks being the dominant force. In 2022, both types of supply chain shocks push inflation upwards. At first sight, it appears puzzling that the contribution of the two shocks is still moderate at the end of the sample period despite strong inflationary pressure. One should keep in mind, however, that at each point in time, the decomposition shows the contribution of the contemporaneous shock as well as all past shocks. As the effect of the two shocks on prices is very persistent, the decomposition in 2021 and 2022 is still affected by the disinflationary nature of shocks occurring in 2020. With additional observations becoming available over time, we would certainly find much more pronounced inflationary contributions of both shocks.

VII Conclusion

Closely integrated global supply chains are a key feature of globalized economies. As a potential side effect, temporary distortions of supply chains propagate globally and affect economic activity and prices around the globe. In this paper, we quantified the effect of global supply chain shocks on the business cycle in the euro area.

We propose an estimated VAR model identified through a combination of sign restrictions and the imposition of narrative information as a tool to model the aggregate effects of supply chain disruptions. We obtain three key findings. First, global supply chain shocks are a main driver of real economic activity and prices in the euro area. An adverse supply chain disruption causes a drop in industrial production and increased consumer prices. Producer prices are five times more sensitive to the shock than consumer prices. These findings are very robust with respect to changing the model specification and the timing of our narrative restrictions. Second, over a horizon of one year, the global supply chain shock explains about 30% of inflation dynamics. Hence, the identified shock is a major determinant of inflationary pressure in the euro area. Over the one-year horizon, the shock accounts for 24% of the variation of the Global Supply Chain Pressure Index, a finding that supports the interpretation of our shock as a global supply chain shock. Third, global supply chain shocks originating in China have a quantitatively similar effect on industrial production as supply shocks from the rest of the world. However, consumer prices are much more responsive to supply chain shocks from the rest of the world than from China.

In light of rising geopolitical tensions, global supply chain diversification seems warranted to avoid being overly dependent on authoritarian regimes. In addition, climate change could disrupt global supply chains and highlights the need for a diversified network of value chains. Our results lend support to this notion from a perspective of stabilization policy. A more broad-based global sourcing structure should reduce the exposure to global supply chain shocks of the type assessed in this paper. Our results call for a diversification of value chains, not a re-shoring of production. The latter would clearly increase business cycle fluctuations.

References

Alessandria, George, Shafaat Yar Khan, Armen Khederlarian, Carter Mix, and Kim J Ruhl, "The Aggregate Effects of Global and Local Supply Chain Bottlenecks: 2020-2022," *unpublished, University of Rochester*, 2022.

Antolín-Díaz, Juan and Juan F Rubio-Ramírez, "Narrative Sign Restrictions for SVARs," *American Economic Review*, 2018, *108* (10), 2802–2829.

Antràs, Pol, "Conceptual Aspects of Global Value Chains," *The World Bank Economic Review*, 2020, *34* (3), 551–574.

Arias, Jonas E, Dario Caldara, and Juan F Rubio-Ramirez, "The Systematic Component of Monetary Policy in SVARs: An Agnostic Identification Procedure," *Journal of Monetary Economics*, 2019, *101*, 1–13.

—, Juan F Rubio-Ramírez, and Daniel F Waggoner, "Inference Based on Structural Vector Autoregressions Identified With Sign and Zero Restrictions: Theory and Applications," *Econometrica*, 2018, *86* (2), 685–720.

Auer, Raphael, Claudio Borio, and Andrew J Filardo, "The Globalisation of Inflation: The Growing Importance of Global Value Chains," *BIS Working Papers No. 602*, 2017.

Baumeister, Christiane and James D Hamilton, "Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks," *American Economic Review*, 2019, *109* (5), 1873–1910.

Benigno, Gianluca, Julian di Giovanni, Jan J Groen, and Adam I Noble, "The GSCPI: A New Barometer of Global Supply Chain Pressures," *Federal Reserve Bank of New York Staff Report No. 1017*, 2022.

Boehm, Christoph E, Aaron Flaaen, and Nitya Pandalai-Nayar, "Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 Tōhoku Earthquake," *Review of Economics and Statistics*, 2019, *101* (1), 60–75.

Caggiano, Giovanni and Efrem Castelnuovo, "Global Financial Uncertainty," *unpublished, University of Padova*, 2022.

Caldara, Dario and Matteo Iacoviello, "Measuring Geopolitical Risk," *American Economic Review*, 2022, *112* (4), 1194–1225.

Canova, Fabio and Gianni De Nicolo, "Monetary Disturbances Matter for Business Fluctuations in the G-7," *Journal of Monetary Economics*, 2002, *49* (6), 1131–1159.

Capolongo, Angela, Michael Kuehl, and Vlad Skovorodov, "Global Supply-Side Disruptions and Longer Delivery Times in the Euro Area," *Available under* https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4022960, 2022.

Carvalho, Vasco M, Makoto Nirei, Yukiko U Saito, and Alireza Tahbaz-Salehi, "Supply Chain Disruptions: Evidence from the Great East Japan Earthquake," *The Quarterly Journal of Economics*, 2021, *136* (2), 1255–1321.

Chen, Hongyi and Peter Tillmann, "Lockdown Spillovers," *unpublished, University of Giessen*, 2022.

Di Giovanni, Julian, Şebnem Kalemli-Özcan, Alvaro Silva, and Muhammed A Yildirim, "Global Supply Chain Pressures, International Trade, and Inflation," *NBER Working Paper Series No. 30240*, 2022.

Döhrn, Roland, "Sieben Jahre RWI/ISL-Containerumschlag-Index–ein Erfahrungsbericht," *Wirtschaftsdienst*, 2019, *99* (3), 224–226.

— and Sönke Maatsch, "Der RWI/ISL-Containerumschlag-Index: Ein neuer Frühindikator für den Welthandel," *Wirtschaftsdienst*, 2012, *92* (5), 352–354.

Faust, Jon, "The Robustness of Identified VAR Conclusions about Money," *Carnegie-Rochester Conference Series on Public Policy*, 1998, *49*, 207–244.

Freund, Caroline, Aaditya Mattoo, Alen Mulabdic, and Michele Ruta, "Natural Disasters and the Reshaping of Global Value Chains," *IMF Economic Review*, 2022, *70*, 590–623.

Frohm, Erik, Vanessa Gunnella, Michele Mancini, and Tobias Schuler, "The Impact of Supply Bottlenecks on Trade," *ECB Economic Bulletin 6/2021, European Central Bank*, 2021.

Furceri, Davide, Yan Carriere-Swallow, Pragyan Deb, Daniel Jimenez, and Jonathan David Ostry, "Shipping Costs and Inflation," *IMF Working Paper 2022/061, International Monetary Fund*, 2022.

Georgiadis, Georgios, Gernot J Müller, and Ben Schumann, "Global Risk and the Dollar," *ECB Working Paper No. 2021/2628*, 2021.

Gunnella, Vanessa, Alexander Al-Haschimi, Konstantins Benkovskis, Francesco Chiacchio, François de Soyres, Benedetta Di Lupidio, Michael Fidora, Sebastian Franco-Bedoya, Erik Frohm, Katerina Gradeva et al., "The Impact of Global Value Chains on the Euro Area Economy," *ECB Occasional Papers Series No. 221*, 2019.

Inoue, Atsushi and Lutz Kilian, "Joint Bayesian Inference about Impulse Responses in VAR Models," *Journal of Econometrics*, 2022, *231* (2), 457–476.

Isaacson, Maggie and Hannah Rubinton, "Shipping Prices and Import Price Inflation," *Federal Reserve Bank of St. Louis Working Paper No. 2022-017*, 2022.

Känzig, Diego R, "The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements," *American Economic Review*, 2021, *111* (4), 1092–1125.

Khalil, Makram and Marc-Daniel Weber, "Chinese Supply Chain Shocks," *unpublished, Deutsche Bundesbank*, 2022.

Kilian, Lutz and Daniel P Murphy, "Why Agnostic Sign Restrictions Are Not Enough: Understanding the Dynamics of Oil Market VAR Models," *Journal of the European Economic Association*, 2012, *10* (5), 1166–1188.

__ , Nikos K Nomikos, and Xiaoqing Zhou, "Container Trade and the US Recovery," *International Journal of Central Banking, forthcoming*, 2021.

LaBelle, Jesse and Ana Maria Santacreu, "Global Supply Chain Disruptions and Inflation During the Covid-19 Pandemic," *Federal Reserve Bank of St. Louis Review*, 2022, *104* (2), 78–91.

Lane, Philip R', "Bottlenecks and Monetary Policy," *Blog post, The ECB Blog, February 10*, 2022.

Raftery, Adrian E and Steven Lewis, "How Many Iterations in the Gibbs Sampler?," *Bayesian Statistics*, 1992, *4*, 115–130.

Reis, Ricardo, "The Burst of High Inflation in 2021–22: How and Why Did We Get Here?," *unpublished, London School of Economics*, 2022.

Tenreyro, Silvana, "International Trade, Global Supply Chains and Monetary Policy," *Speech given at a CEPR Webinar, October 25*, 2021.

Wu, Cynthia and Dora Xia, "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound," *Journal of Money, Credit, and Banking*, 2016, *48* (2-3), 253–291.

Appendix

A Convergence Diagnostics

In this section, we use two common metrics to investigate our posterior draws' properties. In a first exercise, we calculate the inefficiency factors. Inefficiency factors are calculated as $1 + 2\sum_{j=1}^{\infty} \rho_j$, where ρ_j is the autocorrelation of j^{th} order for the underlying parameter. Values of around 20 are regarded as satisfactory. Table (7) reports the inefficiency factors of our entire parameter space for both the draws from the baseline identification and the draws that also satisfy the narrative restrictions. As can be seen, all inefficiency factors are far below 20.

Table 7: Distribution of inefficiency factors

	Mean	Median	Min	Max	10^{th}	90^{th}
	Baseline Restrictions					
B	4.00	3.98	3.33	4.75	3.66	4.34
$\mathbf{\Sigma}$	4.03	3.96	3.45	5.21	3.68	4.46
	Baseline plus Narrative Restrictions					
B	2.96	2.96	2.49	3.38	2.78	3.19
$\mathbf{\Sigma}$	2.95	2.90	2.71	3.42	2.74	3.18

Notes: **B** corresponds to the reduced-form coefficients, and $\mathbf{\Sigma}$ are the elements of the estimated variance-covariance matrix.

Table 8: Distribution of Raftery-Lewis statistics

	Mean	Median	Min	Max	10^{th}	90^{th}
	Baseline Restrictions					
B	163	154	145	372	149	163
$\mathbf{\Sigma}$	154	152	146	167	151	158
	Baseline plus Narrative Restrictions					
B	156	152	139	211	143	171
$\mathbf{\Sigma}$	165	160	143	213	147	203

Notes: **B** corresponds to the reduced-form coefficients, and $\mathbf{\Sigma}$ are the elements of the estimated variance-covariance matrix.

As a second check, we look at the Raftery and Lewis (1992) diagnostic. They use a two-state Markov chain assumption to construct a univariate diagnostic that aims to report the recommended number of iterations needed for a given level of precision in posterior samples.[19] Table (8) reports the corresponding distribution of the recommended minimum number of draws. Overall, we see that we have far more than the required number of draws, conditional on our desired accuracy goal.

To sum up, the convergence diagnostics seem satisfactory.

B Data Sources

In the main paper, we estimate the responses of the Harmonized Index of Consumer Prices as well as import prices and producer prices. Table (9) reports the mnemonics from the ECB's Statistical Data Warehouse for these alternative price series.

Table 9: Mnemonics from the ECB's Statistical Data Warehouse for different price indices

Import price index	
Specification	*Identifier*
Total import price index	STS.M.I8.N.IMPX.NS0020.4.000
Manufacturing	STS.M.I8.N.IMPR.2C0000.4.000
Intermediate goods	STS.M.I8.N.IMPR.NS0040.4.000
Consumer goods	STS.M.I8.N.IMPR.NS0081.4.000
Producer price index	
Specification	*Identifier*
Total producer price index	STS.M.I8.N.PRIN.NS0020.4.000
Manufacturing	STS.M.I8.N.PRIN.2C0000.4.000
Intermediate goods	STS.M.I8.N.PRIN.NS0040.4.000
Consumer goods	STS.M.I8.N.PRIN.NS0081.4.000

C Alternative Decomposition of the Global Supply Chain Pressure Index

In this subsection, we lay out an alternative model for comparing the roles of Chinese supply chain shocks and shocks emanating from the rest of the world (RoW). While the identification strategy is similar to the one presented in Section VI in the main paper, the difference lies in how we derive a supply chain pressure index for the rest of the world. Denote sp_t^{global} the Global Supply Chain Pressure Index and sp_t^j the subindexes for j = *USA, Euro Area, Japan, Korea, Taiwan, China, Great Britain*. In

[19] We follow the most common values (see, for instance, Raftery and Lewis, 1992) and set our quantiles of the marginal posteriors to 2.5% and 97.5%, the minimum probability needed to achieve our stationary posterior distribution of 95% and the desires accuracy of 2.5%, such that our reporting based on a 95% interval result in the true posterior values with a probability lying between 92.5%–97.5%.

a first step, we estimate

$$sp_t^{Global} = c + \sum_{j=1}^{6} \omega_j sp_t^j + \varepsilon_t,$$

where c is a constant and ω_j are the regression coefficients for each subindex of the Global Supply Chain Pressure Index under the restriction that $\sum_{j=1}^{6} \omega_j = 1$. In a second step, we derive the supply chain pressure index for the rest of the world index as

$$sp_t^{RoW} = sp_t^{Global} - \widehat{\omega}_{China} sp_t^{China},$$

where $\widehat{\omega}_{China}$ is the estimated coefficient for the supply chain pressure index of China. Hence, we subtract the part of the Global Supply Chain Pressure Index, which is explained by the Chinese supply chain pressure index.

Figure 16: Different RoW supply chain pressure indexes

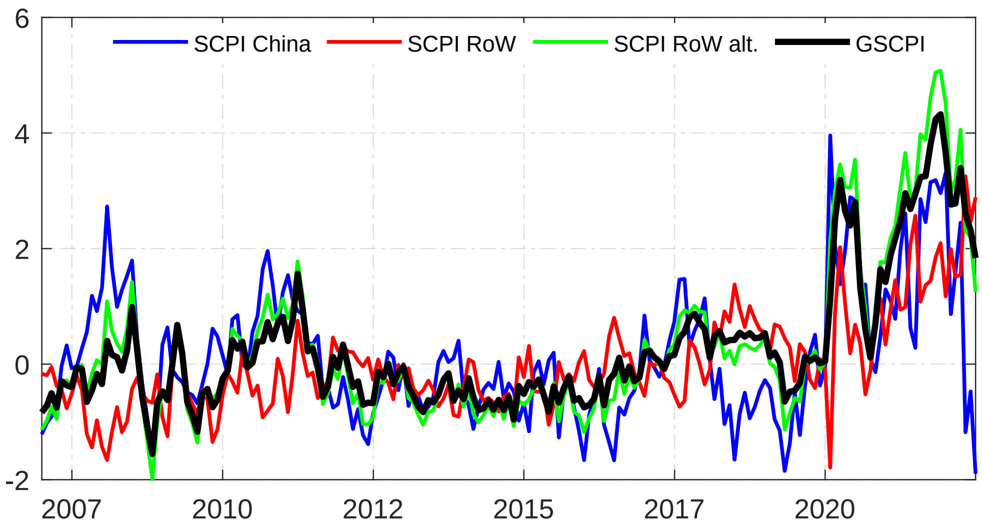

Notes: The alternative RoW supply chain pressure index is calculated as $sp_t^{RoW} = sp_t^{Global} - \widehat{\omega}_{China} sp_t^{China}$.

This can make sense if one assumes that common factors drive the global supply chain pressure index and the Chinese index. The alternative RoW index accounts for this possibility and only subtracts the part of the global supply chain pressure index explained by the Chinese supply chain pressure index. Interestingly, the correlation of the shocks emanating in the rest of the world across both approaches is very high, with 99.7% for the medians from those models that satisfy the baseline sign restrictions and 97.2% for those models which also satisfy the narrative sign restrictions.

Figure (16) plots the corresponding RoW supply chain pressure index from both this approach and the approach in the main paper against the Chinese supply chain pressure index and the Global Supply Chain Pressure Index. Note that while the correlation of both RoW supply chain pressure indexes is high with 62%, the crucial difference is that in this alternative approach, the RoW index is not orthogonal to the global supply chain pressure index.

Figure 17: The responses to a supply chain shock from the rest of the world: alternative specification

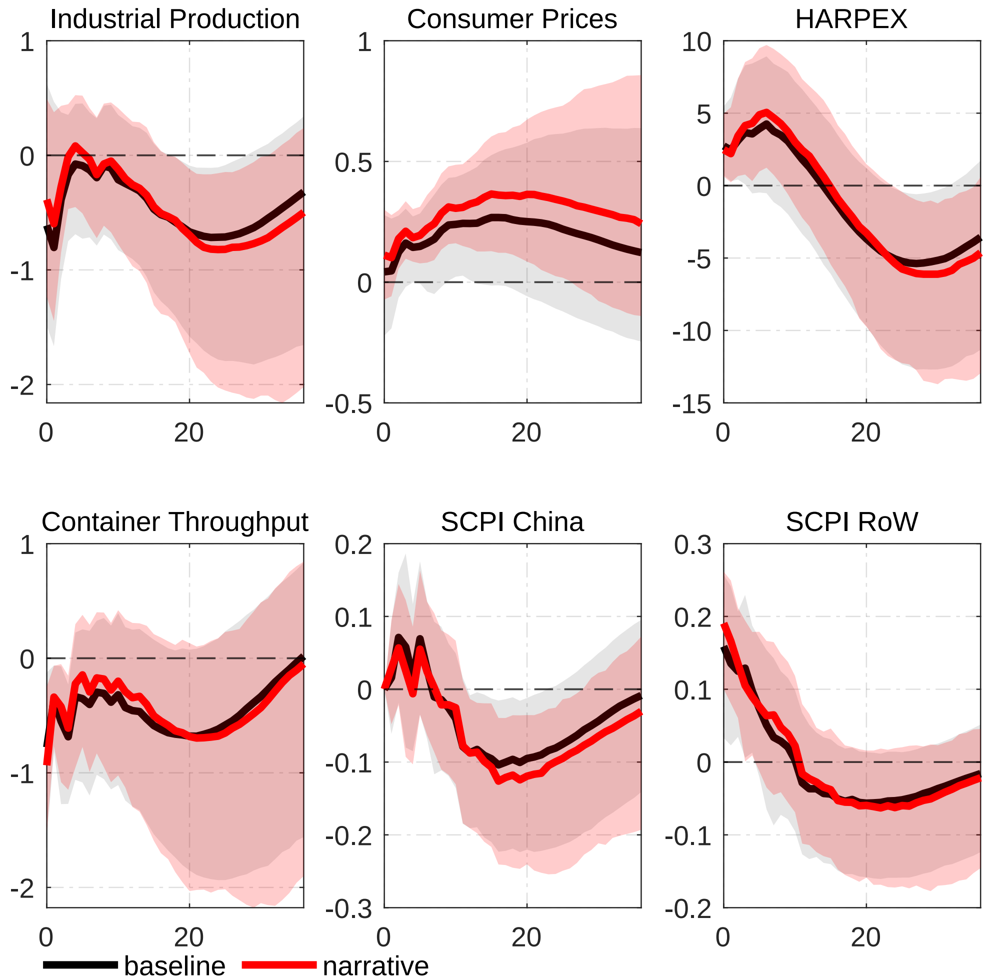

Notes: The black-solid line corresponds to the median of the baseline results, while the red-solid line corresponds to the median of the models that satisfy both the baseline sign restrictions and the narrative restrictions. The light-shaded areas correspond to the 90 percent credible bands for the baseline restrictions, while the red-shaded areas correspond to the 90 percent credible bands that additionally satisfy the narrative restrictions.

Figures (17) and (18) show the corresponding impulse response functions of our endogenous variables from shocks emanating from the rest of the world and from China, respectively. It stands out that the results are very similar to the results from Section VI in the main paper, although the responses of industrial production seem to be somewhat smoother. Similar to the main paper, we see that consumer prices do not react significantly to Chinese supply chain disruptions, but the median seems to be shifted downwards compared to the approach used in the main paper. However, it stands out that our narrative sign restrictions in the alternative approach seem to

Figure 18: The responses to a supply chain shock from China: alternative specification

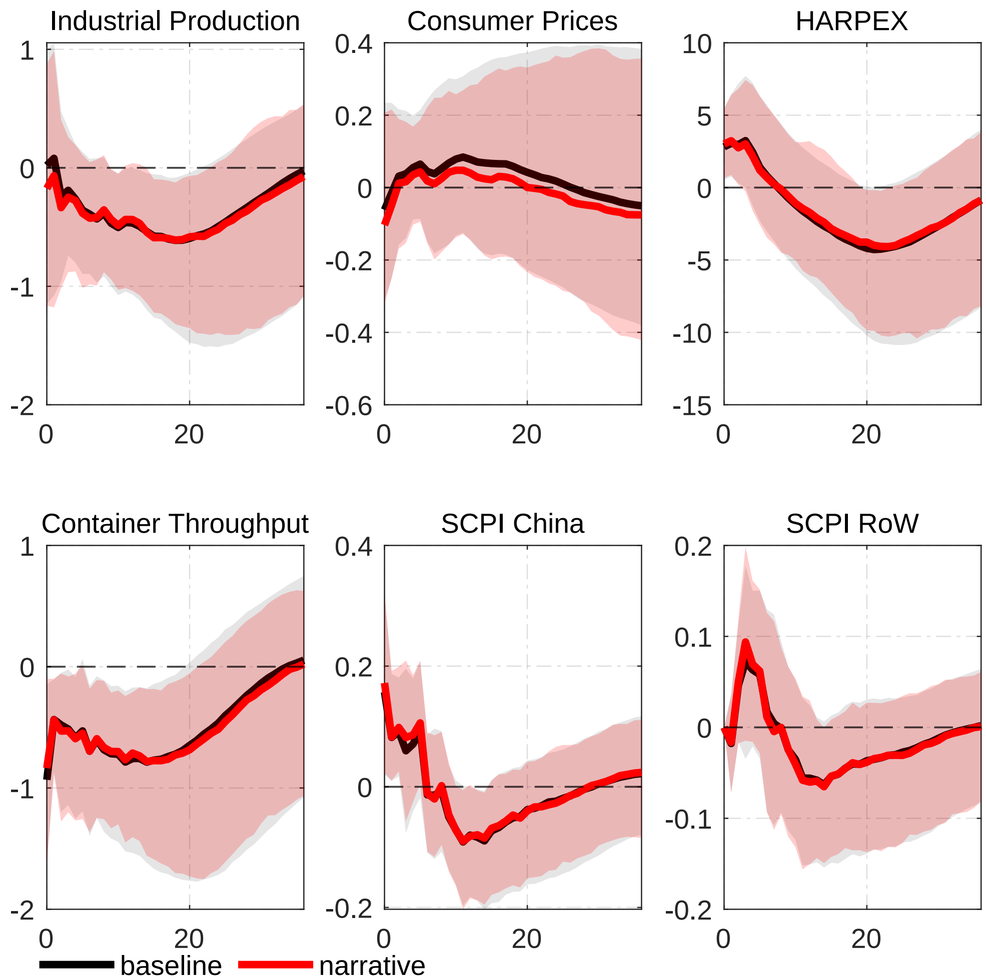

Notes: The black-solid line corresponds to the median of the baseline results, while the red-solid line corresponds to the median of the models that satisfy both the baseline sign restrictions and the narrative restrictions. The light-shaded areas correspond to the 90 percent credible bands for the baseline restrictions, while the red-shaded areas correspond to the 90 percent credible bands that additionally satisfy the narrative restrictions.

be highly informative for identification. That is, following a supply chain shock from the rest of the world, the response of consumer prices is shifted upward. It becomes highly significant once the established narratives are taken into account.

Essay IX:

On the Empirical Relevance of the Exchange Rate as a Shock Absorber at the Zero Lower Bound

This paper is currently under revision at *Deutsche Bundesbank Discussion Paper Series* and also available as:

> **Finck, David, Mathias Hoffmann, and Patrick Hürtgen**, "On the Empirical Relevance of the Exchange Rate as a Shock Absorber at the Zero Lower Bound," *MAGKS Joint Discussion Paper Series in Economics*, 2022, No. 34-2022.

This paper was presented at the following workshops and international conferences:

I: Ausschuss für Außenwirtschaft from the Verein für Socialpolitik (Zurich, Switzerland)
Date: May 2022
Presenter: Mathias Hoffmann

II: 4*th* Behavioral Macroeconomic Workshop (Bamberg, Germany)
Date: June 2022
Presenter: David Finck

Disclaimer

The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Deutsche Bundesbank or the Eurosystem.

Acknowledgement

We thank Elena Bobeica, Zeno Enders, Pierre-Olivier Gourinchas, Philipp Harms, Gernot Müller, Paul Rudel, Sebastian Rüth, Frank Schorfheide, Peter Tillmann, Peter Winker and Christian Wolf for many helpful comments and suggestions. We are also grateful for insightful comments from participants at the Ausschuss für Außenwirtschaft from the Verein für Socialpolitik, participants at the 4*th* Behavioral Macroeconomic Workshop in Bamberg, seminar participants at the Deutsche Bundesbank as well as seminar participants at the University of Giessen.

On the Empirical Relevance of the Exchange Rate as a Shock Absorber at the Zero Lower Bound

David Finck[*] Mathias Hoffmann[†] Patrick Hürtgen[‡]

Abstract

The open economy New Keynesian model with flexible exchange rates postulates that the real exchange rate appreciates in response to an asymmetric negative demand shock in a zero lower bound (ZLB) scenario and exacerbates the adverse macroeconomic effects. However, when monetary policy is able to accommodate the adverse effects of the negative demand shock via unconventional measures, the model can generate a real depreciation at the ZLB. This paper examines these counteracting exchange rate channels empirically. We estimate the effect of a negative asymmetric demand shock on the real exchange rate and inflation expectations as well as output and prices by employing state-dependent and sign-restricted local projection methods for the euro area vis-à-vis the United States, Canada, and Japan. We find that the real exchange rate depreciates when interest rates are not at the ZLB but also when they are. Furthermore, our empirical results show that the real exchange rate can absorb considerable variations in output, confirming its shock-absorbing capacity before but also during the ZLB episode. The stabilizing role of the exchange rate is accompanied by a significant expansion of the ECB's balance sheet in the ZLB period, while it remained unaffected in the pre-ZLB period. Overall, our empirical results favor the open economy New Keynesian model with unconventional measures when interest rates are at the ZLB.

Keywords: Zero Lower Bound, Exchange Rate, Local Projections, State-dependent Effects

JEL classification: F31, E31, E37, C54

[*]University of Giessen, email: david.finck@wirtschaft.uni-giessen.de
[†]Deutsche Bundesbank, email: mathias.hoffmann@bundesbank.de
[‡]Deutsche Bundesbank, email: patrick.huertgen@bundesbank.de

I Introduction

There is a traditional view starting with Friedman (1953) that regimes of flexible exchange rates allow the real exchange rate to depreciate in response to asymmetric negative demand disturbances.[1] This real depreciation then delivers efficient macroeconomic stabilization. However, when monetary policy is constrained by the ZLB, the New Keynesian argument is that flexible exchange rate regimes cannot stabilize cyclical developments in response to adverse demand shocks (see, for instance, Cook and Devereux, 2013; Cook and Devereux, 2016). However, when monetary policy is able to accommodate the adverse effects of the negative demand shock via unconventional measures, the model can generate a real depreciation when interest rates are at the ZLB.

Given these opposing predictions, this paper studies the behavior of the exchange rate when interest rates are at the ZLB. To present the argument, we lay out a two-country New Keynesian model. The countries are highly integrated via financial markets but less than perfectly integrated in goods markets, so relative price adjustments across countries are required. In each country, firms set their prices in their own currencies, and the nominal exchange rate floats. The central banks follow a Taylor rule unless the ZLB binds. We examine the case of a severe global recession where either (i) both countries or (ii) one country is in a liquidity trap due to recessionary asymmetric demand shocks. We show that the driving forces for the real appreciation in a ZLB scenario are the strongly falling (relative) inflation expectations. Since the central bank cannot lower the policy rate, the falling inflation expectations cause a rise in the real interest rate differential in the ZLB-constrained country and, hence, a real appreciation.

In the light of the ZLB, many central banks have adopted unconventional policy measures to counteract the negative effects of the severe global recession and to stabilize inflation expectations.[2] Therefore, in a second step, we investigate the exchange rate response when interest rates are at the ZLB to an asymmetric demand shock in our two-country model when the central bank can stabilize inflation expectations via forward guidance. We show that when the liquidity-trapped country's monetary policy is sufficiently accommodative, inflation expectations are stabilized, and the real exchange rate can depreciate and absorb cyclical developments in response to an adverse demand shock during a ZLB period. Based on these two opposing exchange rate outcomes, in our third and main step, we aim to assess the empirical relevance of the exchange rate as a shock absorber when interest rates are at the ZLB. We utilize monthly data over the time horizon from 1999 to 2020. We estimate the effect of a

[1] When the zero lower bound (ZLB) is not binding, the central bank reduces its nominal interest rate in response to a negative demand shock. This interest rate cut causes a real depreciation, implying an expenditure switching towards cheaper goods.

[2] For a precise definition of unconventional monetary policy see also the speech by Smaghi (2009), which he gave as a member of the Executive Board of the ECB.

Figure 1: Policy rates from 1999-2020

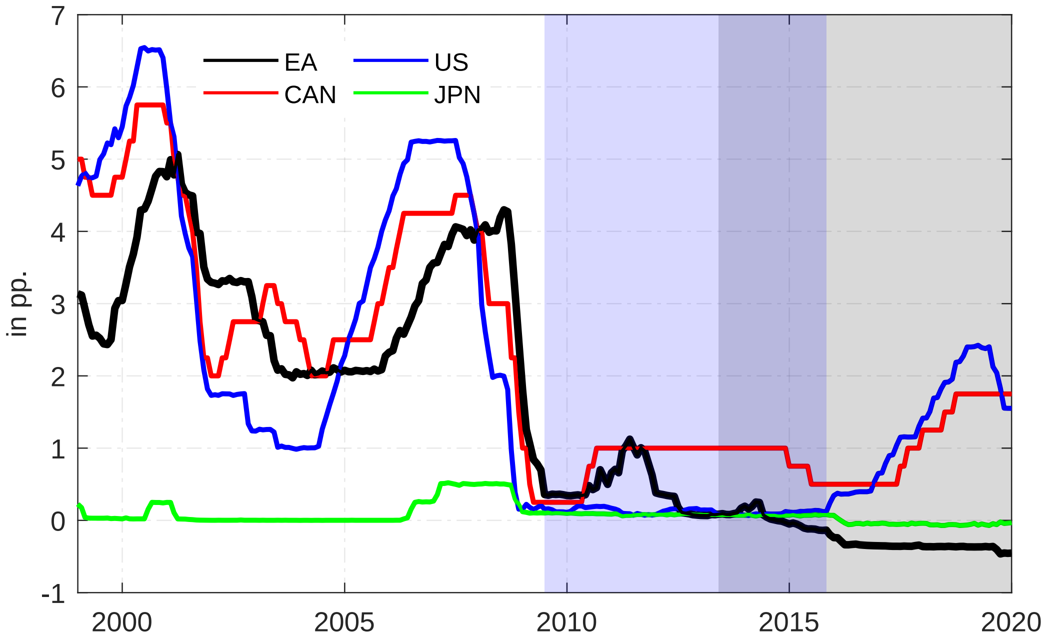

Notes: Gray and blue-shaded areas indicate ZLB periods of the euro area and the United States, respectively.

negative asymmetric demand shock on the real exchange rate and inflation expectations as well as output and prices in the euro area vis-à-vis Canada, Japan, and the US. We have chosen those countries to control for different foreign monetary policies when analyzing the empirical effects of a negative euro area demand shock during our estimation horizon. Figure (1) illustrates that Canada was the only economy among the G7 countries analyzed whose policy rate was not constrained by the ZLB over our estimation horizon. In contrast, Japan's interest rates have remained at the ZLB throughout the last two decades. Figure (1) also shows that since the Great Recession, the US experienced times when the policy rate was constrained at the ZLB or remained unconstrained. At the same time, interest rates in the euro area reached the ZLB shortly after those in the US did and have remained there ever since. Thus, analyzing the euro area vis-à-vis Canada, Japan, and the US allows us to empirically account for different foreign monetary policy stances when assessing the empirical relevance of the exchange rate as a shock absorber when interest rates are at the ZLB.

We employ state-dependent and sign-restricted local projection (LP) methods. As pointed out by Plagborg-Møller and Wolf (2021), LPs are conceptually not different from vector autoregressions (VARs), which are the most popular empirical approach in macroeconometrics to study the propagation of structural shocks. Instead, they are different linear projection techniques sharing the same estimand (in population) under different finite sample properties. Specifically, they show that VARs and LPs estimate the same impulse responses under an unrestricted lag structure. This seminal result implies that VAR-based structural identification, including short-run and sign restrictions, can be implemented equivalently within an LP framework. We extend the Plagborg-Møller and Wolf (2021) framework and allow the impulse responses obtained via sign restrictions to be different across the states of the economy. Our states are specified with respect to the euro area's monetary policy stance. In particular, we compare times where the ECB's policy rate was either unconstrained or constrained by the ZLB and assess the effect of a negative asymmetric demand shock on the real exchange rate, inflation expectations, output, and prices in the euro area vis-à-vis Canada, Japan, and the US.

In this respect, one could interpret our findings as evidence of how well the central bank's unconventional monetary policy reacted to a severe negative demand disturbance when the ZLB constraint was binding. If the exchange rate would depreciate in a ZLB scenario in response to such a shock, monetary policy would have been sufficiently accommodative to cushion the deflationary side-effects of the ZLB constraint on economic activity.

Our empirical findings show that the euro real exchange rate not only depreciates prior to the ZLB period but also during the ZLB episode. At the same time, inflation expectations do not react more strongly when the ZLB is binding.[3] We also find that output in the euro area vis-à-vis Canada, Japan, and the US does not fluctuate more in response to asymmetric negative demand disturbances than before the euro area's ZLB episode. Hence, our empirical results contrast with those of the standard New Keynesian model, which predicts that the real exchange rate appreciates in response to a negative demand shock during a ZLB period and that inflation and expected inflation are substantially more volatile when interest rates are at the ZLB. Nonetheless, they are in line with an extended version of the two-country New Keynesian model, which accounts for the ability of the central bank to allow for unconventional measures in a ZLB period to stabilize inflation expectations. Indeed, our empirical findings show that the ECB's balance sheet expanded significantly in response to the negative demand shock in the ZLB period, while it remained unaffected in the pre-ZLB period. Hence, the ECB's unconventional policies during the ZLB period supported a real depreciation, which helped to absorb cyclical developments in response to the adverse demand shock. Counterfactual variance decompositions show that, depending on the country pair, movements in the real exchange rate absorb at least eight percentage points of the variation in output. In this respect, the exchange rate plays an important role in stabilizing the economy when interest rates are at the ZLB.

We conduct a set of robustness checks of our empirical findings: first, we also estimate our model for the euro area vis-à-vis a broad set of trading partners, for which time-varying trade weights between the euro area and these countries are available all the way back to the launch of the euro. That is, we summarize the information

[3] Corsetti et al. (2014) assess positive demand shocks for the US prior to the ZLB period, finding a real appreciation.

for output and exchange rates for each of these countries and treat them as if they were one single country. Our results for this exercise are very similar to our main findings in that the trade-weighted euro real exchange rate depreciates in response to a negative demand shock both when interest rates are not at the ZLB and when they are. Second, note that the mere possibility of the ZLB being reached might affect economic decisions and expectations even before the ZLB becomes binding. Put differently, economic agents anticipate that the policy rate might reach the ZLB in the future, and their inflation expectations might already respond today. Hence, the consequences of the ZLB can be effective even before the short rate reaches zero. Moreover, there is no generally defined start date for the ZLB period. To account for these concerns, we also estimate our model in a smooth transition LP framework with the policy rate as a state variable so that the states are functions of the policy rate in the euro area. We find the same results as in our main specification. Third, we find that the results from a state-dependent VAR are qualitatively very similar to our benchmark results, although the VAR-based responses are smoother compared to the LP-based responses.

The work by Debortoli et al. (2020), Müller et al. (2022) and Stavrakeva and Tang (2020) is closely related to our analysis. The latter two papers also assess the importance of expectations for the exchange rate. However, the authors focus on the role of monetary policy shocks or broker-dealer relationships in the exchange rate market. Hence, the authors do not consider asymmetric demand disturbances across countries and do not assess the possible state-dependent macroeconomic effects due to the ZLB episode. The work by Debortoli et al. (2020) provides empirical evidence supporting the irrelevance hypothesis by illustrating that US output, inflation, and long-term rates are not affected by the ZLB episode. In contrast to their work, we assess the international dimension of liquidity traps and show empirically that the real exchange rate can depreciate in response to negative asymmetric demand shocks when interest rates are at the ZLB. In this respect, we can support the argument of flexible exchange rates by Friedman (1953), Mundell (1961), Obstfeld and Rogoff (2000), and Obstfeld and Rogoff (2002), i.e. that flexible exchange rates act as a shock absorber to country-specific shocks, such as negative demand shocks. We show that this is the case even if the ZLB constraint is binding but monetary policy is sufficiently accommodative. Thus, at the ZLB, a flexible exchange rate can allow for the adjustment of relative prices so that output and prices are stabilized in the ZLB-constrained country.

Our paper is also connected to a strand of the literature that examines exchange rate policies when the ZLB has been reached. Amongst these, Amador et al. (2020) assess an exchange rate policy that is inconsistent with interest rate parity because of a binding ZLB constraint. Coenen and Wieland (2004) investigate the effectiveness of an exchange rate peg and price level targeting regime in stimulating an economy in a liquidity trap. In a related work, Svensson (2001) argues that price-level targeting, a devaluation of the currency, and a temporary exchange rate peg can be employed to escape a liquidity trap. In relation to this work, we show that the ECB's unconventional monetary policy measures seem to have been sufficiently accommodative in the ZLB period for the real exchange rate to depreciate. We show that this real depreciation then supported a cushioning of the deflationary side-effects of the negative demand shock in the ZLB period.

The paper proceeds as follows: in Section II, we lay out a two-country open economy model with sticky prices to show, in a nutshell, the theoretical considerations of a real appreciation or depreciation in a ZLB period. Based on this, in Section III we present our empirical specification to assess the hypothesis of possibly opposing exchange rate outcomes in response to asymmetric demand shocks when interest rates are at the ZLB. Section IV outlines the empirical findings. Section V contains robustness checks. Section VI concludes.

II Assessing the Theoretical Mechanism

To depict the opposing real exchange rate outcomes of an appreciation versus a depreciation when interest rates are at the ZLB, we explore the two-country New Keynesian model as in Clarida et al. (2002). Following Engel (2011), we allow for less than perfectly integrated goods markets and extend the model by accounting for the ZLB. Here, we briefly outline the model to set the basis for the empirical analysis. The main ingredients of the model are two equally sized countries, home (H) and foreign (F), which are connected via trade in goods and state-contingent assets. In each of the economies there is a continuum of i households indexed by $i \epsilon [0, 1]$. Each household i consumes domestic and foreign goods, aggregated by a Cobb-Douglas technology, and faces a time preference shock ξ_i, which we also refer to as a demand shock. Our analysis focuses on a negative shock to ξ_i, which implies that households are willing to consume more in the future rather than today. Therefore, they increase savings and reduce their demand for goods. The time preference shock $\xi_{i,t}$ is unanticipated and follows a stochastic decay. In every period $t \geq 1$ it holds that $\xi_{i,t} = \xi_{i,t-1}$ with probability $\mu < 1$.[4] With probability $(1 - \mu)$ the time preference shock $\xi_{i,t}$ returns to zero. This occurs at the same period T for all households in the two countries. A continuum of monopolistically competitive sticky-price firms sells their goods in their own currency to home and foreign households. The monetary authorities decide on the underlying nominal interest rates while the country-specific fiscal authorities collect either lump-sum taxes or pay transfers to their residents. The full model and all optimality conditions are laid out in Hoffmann and Hürtgen (2021). In the following, we focus on the relevant equilibrium conditions to present the argument of a real

[4]For $0 < \mu < 1$, the ZLB will expire in expectations, see Eggertson and Woodford (2003). This ensures that inflation today is pinned down by expectations that the stable manifold will determine future inflation.

appreciation versus depreciation when interest rates are at the ZLB.

A. *The Model's Equilibrium Relationships*

We log-linearize the model around its steady-state values. Lower case letters reflect log deviations from the variable X_t in the steady state X: $x_t = \log(X_t) - \log(X)$. We start with the evolution of consumption and output. We then discuss the determination of inflation, policy rates, and real exchange rates. In equilibrium, all households are identical and the individual's expectation equals average expectations $E_t^i = \overline{E}_t$. Then

$$c_t = \overline{E}_t \left[c_{t+1} - \frac{1}{\sigma} \left(\xi_{t+1} - \xi_t + r_t - \pi_{t+1} \right) \right] \tag{1}$$

is the Euler equation, with $\overline{E}_t \left[r_t - \pi_{t+1} \right]$ reflecting household's consumption-based real interest rate, which depends on consumer price inflation (CPI), π_t, as well as the policy rate, r_t. The intertemporal elasticity of substitution equals $1/\sigma$. A similar condition in the foreign economy is indexed by $*$. The Markov property of the time preference (i.e. demand) shocks ξ and ξ^* implies that under independent monetary policy and flexible exchange rates, there are no predetermined state variables. In expectation, all endogenous variables in the world economy will inherit the same persistence as the shock itself.

International linkages are expressed by the variable $x_t^R = \left(x_t - x_t^* \right)/2$, which denotes relative world variables. From the firms' resource constraints and households' demand conditions, aggregate output becomes

$$y_t = c_t + \left(1 - \frac{v-1}{\delta} \right) y_t^R - \frac{v(2-v)}{\delta} \xi_t^R. \tag{2}$$

The intensity of the home bias equals $0 \leq (v-1)/\delta \geq 1$, whereby $\delta \geq 1$ is a function of σ and v, with $v \geq 1$. If there is no home bias, $v = 1$ and $\delta = \sigma$. Given that goods markets are only imperfectly integrated, we set $v > 1$. From (1) and (2), a relation between interest rates, CPI, and output is obtained

$$r_t^R = \overline{E}_t \left[\pi_{t+1}^R + \sigma \frac{v-1}{\delta} \Delta y_{t+1}^R - \frac{(v-1)^2}{\delta} \Delta \xi_{t+1}^R \right], \tag{3}$$

with $\pi^R = \pi_H^R - (2-v) \left(\pi_H^R - r_{-1}^R \right)$ denoting relative CPI and π_H^R defining relative domestic price inflation. Equation (3) describes relative output in response to a demand shock. Expected CPI inflation equals

$$\overline{E}_t \left[\pi_{t+1} \right] = \overline{E}_t [\pi_{Ht+1}] + (2-v) \, \overline{E}_t [r_t^R - \pi_{Ht+1}^R], \tag{4}$$

with $\pi_{H,t} = \frac{\kappa}{2} y_t - \kappa_{(y-y^*)} y_t^R + \kappa_{(c-c^*)} \xi_t^R + \beta E_t \left[\pi_{Ht+1} \right]$ denoting gross domestic price inflation.[5] It holds that $\kappa > \kappa_{(y-y^*)} \geq \kappa_{(c-c^*)} \geq 0$, for $v \geq 1$. κ defines the responsiveness of domestic inflation to domestic output. The responsiveness to relative output is given by $\kappa_{(y-y^*)}$ and captures how strongly inflation adjusts to changes in relative output. The response to relative time preference conditions is determined by $\kappa_{(c-c^*)}$. The foreign country has similar conditions, with the second and third terms of the right-hand side of domestic price inflation taking opposite signs.

The monetary authorities adopt the following monetary policies by following a non-linear Taylor rule

$$r_t = \max\left\{-\ln(1/\beta), \phi\pi_{Ht}\right\} \text{ and } r_t^* = \max\left\{-\ln(1/\beta), \phi\pi_{Ft}^*\right\}, \tag{5}$$

with $r = \ln(1/\beta)$ and $\beta < 1$ being the household's discount factor. π_{Ht} denotes gross home producer price inflation. The Taylor principle holds, and the reaction on inflation is given by $\phi > 1$. Hence, the central banks' nominal interest rates react by more than the actual price change, pushing real interest rates in the desired direction. However, when interest rates are at the ZLB, the monetary authorities can lower the nominal rate up to $r_t = -\ln(1/\beta)$ and movements in the real interest rate would only depend on expected inflation.

From equations (1)-(3) and the respective foreign counterparts, we obtain a relationship between policy rates, relative inflation rates, and the real exchange rate q,

$$(v-1)\, r_t^R = \overline{E}_t\left[(v-1)\,\pi_{Ht+1}^R + \frac{\Delta q_{t+1}}{2}\right], \tag{6}$$

which mirrors the real UIP condition. From (6) it follows that changes in the real exchange rate are due to movements in the real interest rate differential between the home and foreign country, $\overline{E}_t\left[r_t^R - \pi_{t+1}^R\right]$. From the real UIP condition also becomes clear that when the Taylor rule (5) determines the nominal rate, monetary policy and inflation expectations determine the real exchange rate.

B. *The Effects of a Negative Asymmetric Demand Shock*

We examine the effects of negative demand shocks based on these equilibrium relationships. To set the stage, we focus on situations where interest rates are not at the ZLB and those where they are. This lays out the main mechanisms at work and explains how a real appreciation in a ZLB period occurs due to the decline in inflation expectations. Then, we show that a real appreciation can also occur if the ZLB is only binding in one country. We use this example to illustrate that with unconventional measures, such as forward guidance, inflation expectations are stabilized, and a real depreciation in a ZLB period can occur.

We consider equilibria where the variables are constant from the period the prefer-

[5]This follows Calvo (1983). For details see appendix A.1. in Hoffmann and Hürtgen (2021).

ence shock hits until the shock reverts back to zero, and the economy is in its non-stochastic steady state.[6] Since the time preference shocks ξ_i and ξ_i^* are unanticipated and follow a stochastic decay, expected consumption, output, inflation, and the real exchange rate inherit the Markov property of the demand shock. They take on the same values as long as the shock lasts and will revert to zero once the shock disappears.[7] Therefore, the time subscript t is replaced by the state subscript s. Then we can express the real UIP condition (6) by

$$\frac{q_{\mathsf{s}}}{2} = -\frac{(v-1)}{(1-\mu)}\left(r_{\mathsf{s}}^R - \left[\mu\pi_{H\mathsf{s}}^R\right]\right). \tag{7}$$

Thus, when the real interest rate differential $(v-1)\,(r_{\mathsf{s}}^R - [\mu\pi_{H\mathsf{s}}^R])$ falls, the real exchange rate depreciates, $q_{\mathsf{s}} > 0$, and vice versa. Consider now an asymmetric negative demand shock, which hits the home country more severely. This will lead to a (relative) fall in consumption, output, and, hence, prices as well as inflation expectations, $[\mu\pi_{H\mathsf{s}}^R] < 0$.

ZLB is (not) binding in both countries. When the ZLB is not binding, the central banks can use their policy rates to respond to the asymmetric negative demand shock.[8] Following the monetary policy rule (5), the policy rate will be reduced by more than the fall in inflation in response to the negative demand shock. This will cause a decline in the relative real interest rate by $(v-1)\left(\phi-\mu\right)[\mu\pi_{H\mathsf{s}}^R]/\mu < 0$. It follows from (7) that the decline in the real interest rate differential will lead to a real depreciation

$$\frac{q_{\mathsf{s}}}{2} = -\frac{(v-1)}{(1-\mu)}\frac{\phi-\mu}{\mu}\left[\mu\pi_{H\mathsf{s}}^R\right] > 0. \tag{8}$$

When the asymmetric demand shock is so severe that the ZLB is binding, the central banks cannot sufficiently use their policy rates to respond to it.[9] It follows that the fall in the relative inflation expectations will directly translate into a rise in the real interest rate differential, $-(v-1)\,[\mu\pi_{H\mathsf{s}}^R] > 0$, and, hence, a real appreciation

$$\frac{q_{\mathsf{s}}}{2} = \frac{v-1}{1-\mu}\left[\mu\pi_{H\mathsf{s}}^R\right] < 0. \tag{9}$$

Equations (8) and (9) illustrate that, without further monetary policy action, the real exchange rate appreciates in response to an asymmetric demand shock when interest rates are at the ZLB and that (relative) inflation expectations matter.

ZLB is binding only in the home country. Even in an environment where only the home country's interest rates are at the ZLB, and the foreign economy can use

[6] This builds on the work by Eggertson and Woodford (2003) and Christiano et al. (2011).
[7] Thus, in every period $0 \leq t \leq T-1$, the variables are constant but depend on the shocks.
[8] It holds that $\xi^{crit} < \xi_{\mathsf{s}} < 0$ and $\xi^{*crit} < \xi_{\mathsf{s}}^* $ with $\xi_{\mathsf{s}} < \xi_{\mathsf{s}}^*$, see Hoffmann and Hürtgen (2021).
[9] These conditions are satisfied for $\xi_{\mathsf{s}} < \xi^{crit}$ and $\xi_{\mathsf{s}}^* < \xi^{*crit}$, but $\xi_{\mathsf{s}} < \xi_{\mathsf{s}}^*$.

monetary policy to counteract the consequences of the home country's asymmetric negative demand shock, perverse responses of the real exchange rate can occur.[10] Based on the monetary policy rules (5) and the real UIP condition (7), the real exchange rate evolves by

$$\frac{q_{\mathrm{S}}}{2} = \frac{v-1}{1-\mu}\left(\left[\mu\pi^R_{H\mathrm{S}}\right] + \frac{\ln(1/\beta) + \phi\pi^*_{F\mathrm{S}}}{2}\right). \tag{10}$$

Now not only relative inflation expectations $[\mu\pi^R_{H\mathrm{S}}]$ matter for the real interest rate differential and the real exchange rate, but also the foreign monetary policy and its ability to affect foreign inflation, $\phi\pi^*_{F\mathrm{S}}$. For example, a rise in foreign inflation $\pi^*_{F\mathrm{S}}$ would cause a rise in the real exchange rate compared to (9), as this is accompanied by a rise in the foreign real interest rate, $(\phi-\mu)\pi^*_F > 0$.[11] But when the ZLB in the home country is binding, relative inflation (expectations) deteriorates, so that $[\mu\pi^R_{H\mathrm{S}}] < (\ln(1/\beta) + \phi\pi^*_{F\mathrm{S}})$. Hence, the real interest rate differential widens, and the real exchange rate appreciates. This situation occurs in a ZLB scenario when the difference between the home and foreign demand shock is sufficiently large.

Furthermore, Hoffmann and Hürtgen (2021) show that in all of the described situations, the decline in inflation expectations will be more severe when the ZLB is binding. The reason is that when interest rates are at the ZLB, firms want to lower their prices in response to the decline in consumption and output. However, since prices are sticky, some price adjustments will only occur in the future. Households will understand this and adjust their inflation expectations downwards accordingly. Without any monetary policy intervention, this will lead to a stronger rise in today's real interest rate and a real appreciation, amplifying the decline in consumption, output, and, hence, inflation expectations.

Unconventional measures: Forward guidance as an example. Smaghi (2009) argues that in stressful times for the global financial system, easing monetary policy by lowering policy rates towards the ZLB would not be enough to counteract the recessionary consequences of adverse demand shocks. Therefore, unconventional tools are needed. Via forward guidance – one manifestation of unconventional monetary policy – policymakers can directly influence expectations about future interest rates by resorting to a conditional commitment to maintaining policy rates at the lower bound for a significant period of time (e.g. Eggertson and Woodford, 2003). To align the future interest rate level in the market, the central bank needs to operate in the market for financial assets, such as government bonds. Those operations will then be reflected

[10] In this case, it holds that $\xi_{\mathrm{S}} < \xi^{crit} < 0$, but $\xi^{*crit} < \xi^*_{\mathrm{S}}$.

[11] A rise in foreign inflation occurs when the foreign demand shock is sufficiently positive, i.e. $\xi^*_{\mathrm{S}} > 0$. We will maintain this assumption. If the foreign demand shock were negative, i.e. $\xi^{*crit} < \xi^*_{\mathrm{S}} < 0$, foreign inflation would also fall. From (10), a real appreciation occurs unambiguously in the home country.

in a larger central bank balance sheet due to an expansion of monetary liabilities.

In this section, we focus on the effects of forward guidance on the real exchange rate. However, given the simplicity of our model, we refrain from assessing the effects on the central bank's balance sheet here but will do so in Section IV.B of our empirical analysis in more detail.

In line with our above analysis, we maintain the environment where the ZLB constrains only the home country. The foreign country pursues a policy of positive interest rates, following the Taylor rule, as described in (5). The home country's central bank announces in the ZLB state a path of the *future state*.[12] In particular, it will set the nominal interest rate in periods $t \geq T$ to achieve an inflation target of $\overline{\pi}_H > 0$, to raise inflation expectations in periods $t < T$. We assess the following scenario: At time $t = 1$, there is an asymmetric negative demand shock to the home country, which drives its policy rate to the ZLB. For illustrative purposes, we focus on a special case where the central bank has perfect foresight and knows that the preference shock will last for T periods. We assume that the shock persists for $T = 4$ periods, while the home country's central bank will set the $\overline{\pi}_H > 0$ to 40 basis points annually for $t = 8$ periods. We compare this scenario to the one where the home country does not follow any monetary policy action in a ZLB scenario, as described above.

The blue dashed (red dashed) lines in Figure (2) show the responses of the home (foreign) economy when the home country's interest rates are at the ZLB but it does not pursue forward guidance, and the foreign country's monetary policy is unconstrained. In such a situation, the real exchange rate would appreciate, as shown by the top-right panel of Figure (2). Furthermore, the figure shows in the upper panel a severe decline in output and inflation in the home economy in response to the negative asymmetric demand shock at home. These findings illustrate our discussion from above, namely that without further monetary policy intervention, the real exchange rate appreciates, and output and inflation decline strongly in the country that is constrained by the ZLB.

However, the solid blue (red) lines in Figure (2) show the responses of the home (foreign) country when the home country adopts forward guidance. Then, when the home country's policymakers commit to the future path of interest rates, it follows that a real exchange rate depreciation (blue-solid line) can be generated by announcing higher inflation today. This also impacts positively on output and inflation by mitigating its otherwise severe decline. The real depreciation also allows for an expenditure switching effect toward home-produced goods, as can be seen by the relatively stronger decline in foreign output by the bottom left panel of Figure (2), when the

[12]Note that many papers also examine the credibility of forward guidance announcements (see, for instance, Nakata and Sunakawa, 2022; Walsh, 2018; Finck, 2020). In these papers, forward guidance announcements are credible (sustainable) when fulfilling past promises is the central banks' best strategy at any point until the promised lift-off date. A typical result in all these papers is that forward guidance can be credible only if recurrent ZLB episodes are possible, as assumed in this paper. We do not address possible credibility issues and assume no time inconsistency problem exists.

Figure 2: The ZLB-constrained home country with and without forward guidance

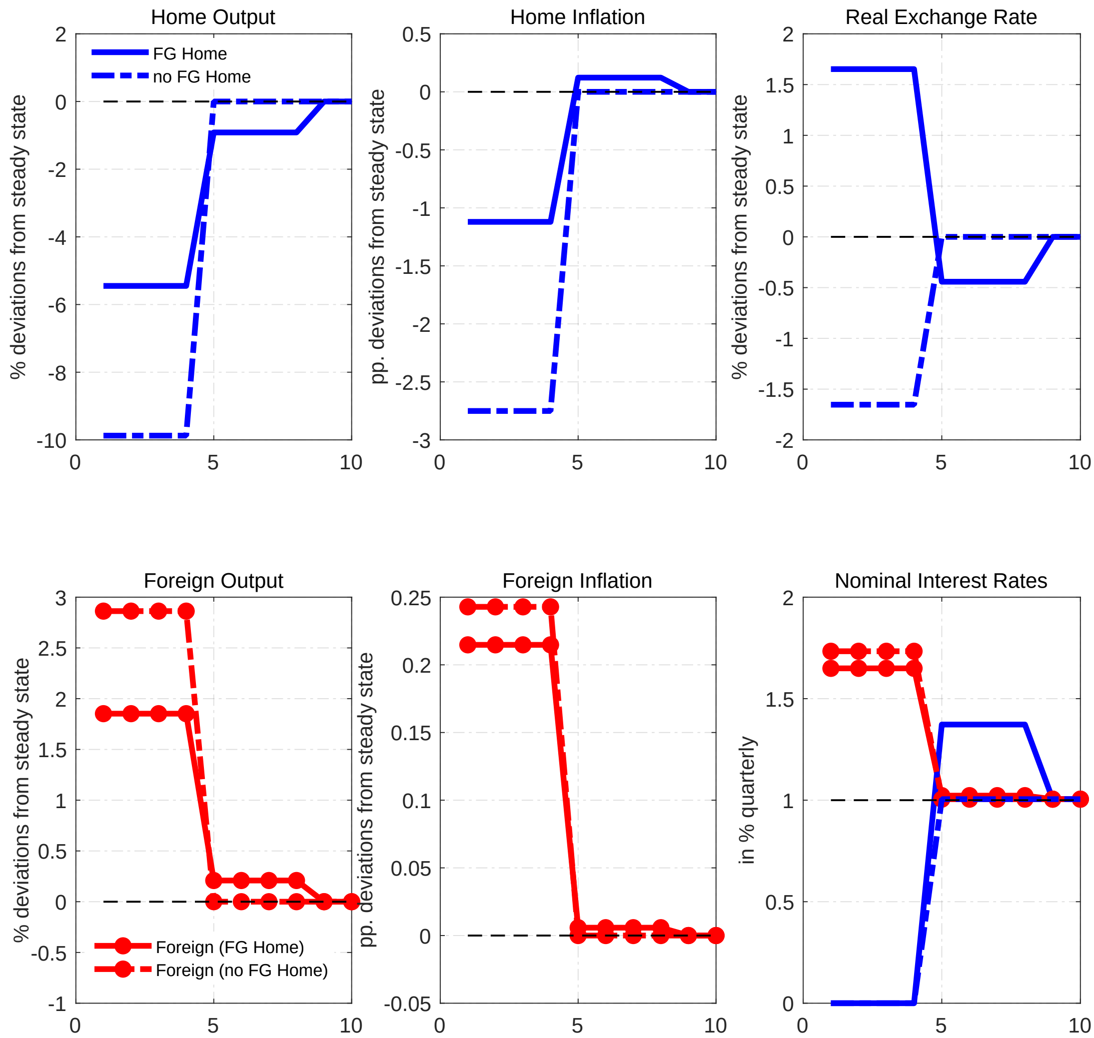

Notes: It holds that $\xi_S < \xi^{crit} < 0$ and $\xi^{*crit} < 0 < \xi^*_S$ and $\mu = 0.8$, whereby $\sigma = 2$, $\upsilon = 1.5$, $\delta = 1.75$, $\kappa = 2.5$, $\kappa_{(y-y^*)} = 1.64$ and $\kappa_{(c-c^*)} = 0.25$.

home country is at the ZLB, but it pursues a monetary policy of forward guidance.

In summary, when the ZLB is binding in at least one country and monetary policy cannot sufficiently accommodate negative demand shocks, no expenditure switching is possible via the real exchange rate since a real appreciation occurs for the country in which the ZLB is binding. This effect is caused by a strong decline in (relative) inflation expectations in the country which is hit most severely by the asymmetric negative demand shock. The model also predicts that output, inflation (expectations), and the real exchange rate are substantially more volatile when interest rates are at the ZLB. However, our model also shows that when monetary policy can stabilize inflation expectations in a ZLB scenario via unconventional measures, such as forward guidance, the real exchange rate can depreciate and stabilize movements in output and prices.[13] The next section assesses these opposing predictions empirically.

III State-Dependent Local Projections

This section empirically explores the exchange rate arguments by estimating a sign-restricted state-dependent local projection model. First, we lay out the econometric model and explain its state-dependent structure. Second, we discuss the data used to estimate our empirical model and show how we differentiate the two monetary states of a non-binding and binding ZLB. Third, we explain our identification strategy. Finally, the statistical inference is laid out to draw the structural impulse responses.

A. *Specification*

We set up a non-linear model based on sign restrictions to distinguish the possible state-dependent real exchange rate effects. In particular, we estimate a state-dependent LP model to empirically investigate the role of the exchange rate as a shock absorber in a ZLB scenario.

As pointed out by Plagborg-Møller and Wolf (2021), LPs are conceptually not different from VARs. In fact, they are different linear projection techniques sharing the same estimand under different finite sample properties. This result implies that VAR-based structural identification, including short-run and sign restrictions, can be implemented equivalently within an LP framework.[14] We empirically utilize this finding when assessing the propagation of asymmetric demand shocks and extend the idea of Plagborg-Møller and Wolf (2021) to a non-linear framework. More precisely, we condition the effect of the demand shock on whether the ZLB is binding or not. To fix notation, first, consider the linear model

$$y_{i,t+h} = \alpha_{i,h} + \beta_{i,h} y_t' + \gamma_{i,h} x_t' + u_{i,h,t},$$

where $y_{i,t+h}$ is the ith endogenous variable in the vector y_t at horizon $t+h$ and $\alpha_{i,h}$, $\beta_{i,h}$, $\gamma_{i,h}$ contain the projection coefficients for the control variables in y_t and x_t, respectively. Specifically, the $n \times p + 1$ vector $\gamma_{i,h} = \left[\phi_{i,h,1}, \ldots, \phi_{n,h,1}, \phi_{i,h,2}, \ldots, \phi_{n,h,p}, \delta_{i,h}\right]$ contains the stacked coefficients for the covariates in $x_t = \left[y_{t-1}, \ldots, y_{t-p}, t\right]$. Finally, $u_{i,h,t}$ corresponds to the projection residual of variable i at horizon h in t with (strictly) positive variance. Based on this, we set up a non-linear model to distinguish between the possible state-dependent effects of the ZLB. Thus, we allow for different projection

[13] When accounting for the so-called forward guidance puzzle, Hoffmann and Hürtgen (2021) show that the effects would remain but in a mitigated manner.

[14] LPs could also accommodate other identification schemes such as narrative sign restrictions as in Antolín-Díaz and Rubio-Ramírez (2018), or long-run restrictions as in Blanchard and Quah (1989).

coefficients across two states for all variables:

$$y_{i,t+h} = (1 - \lambda_t)\left[\alpha^{I}_{i,h} + \beta^{I}_{i,h} y_t + \gamma^{I}_{i,h} x_t\right] + \lambda_t\left[\alpha^{II}_{i,h} + \beta^{II}_{i,h} y_t + \gamma^{II}_{i,h} x_t\right] + u_{i,h,t}.$$

The superscripts *I* and *II* distinguish between state *I=pre-ZLB* and *II=ZLB*, respectively. Hence, λ_t is equal to zero when in t the economy is in state *I=pre-ZLB* and equal to one when the economy is in state *II=ZLB*.

It has to be noted that, even though VARs might be as robust to misspecifications as LPs *in population*, this is not the case when one is forced to use a finite lag structure, saying that the impulse responses might in general differ *in sample*. However, the consensus in the literature is that impulse responses in an LP framework under a finite lag length are more robust to misspecification than in a VAR framework. This is because iterated forecasts in a VAR can increase the misspecifications as the horizon expands. This is not the case in an LP framework since we estimate the reduced-form coefficients separately for each $h = 0, ..., H$.

Moreover, VARs implicitly assume no changes in the state of the economy (see, for instance, Tenreyro and Thwaites, 2016; Alpanda and Zubairy, 2019; and Alpanda et al., 2021). That is, based on the fact that VAR-based impulse responses rely on iterated forecasts, one implicitly assumes that once the shock occurs in state j, the economy will stay forever in this state. This is not the case within a state-dependent LP framework, where the reduced-form impulse response coefficients reflect the average effect of shocks as a function of the state of the economy within the same period the shock hits. Hence, it also comprises the average effect of the shock on future changes in the state of the economy.[15]

B. *Data*

We estimate our baseline model for the euro area vis-à-vis Canada, the US, and Japan. We use six variables for each country pair to estimate our empirical model. The first four variables only pertain to the euro area: (1) industrial production (excl. construction), (2) the Harmonised Index of Consumer Prices (HICP), (3) one-year-ahead inflation expectations from Consensus Economics, (4) and the (shadow) short rate.[16] The (shadow) short rate is included for the following three reasons: Firstly, in the euro area, there is no generally defined start date for the ZLB period. Therefore, we assess when the Wu and Xia (2016) shadow interest rate for the euro area is below

[15]Note that this point is even more important when one specifies a model with frequent changes between the states. As far as the presence of the ZLB as a state is concerned, this point is more relevant than before, at least since the COVID-19 pandemic. This is based on the observation that many central banks lowered their policy rates in response to the pandemic and returned to the ZLB for a second time after the Great Recession.

[16]Consensus Economics forecasts for the euro area have only been available since December 2002, so we approximate the forecasts from January 1999 to November 2002 by real GDP-weighted forecasts from Germany, the Netherlands, Spain, France, and Italy. For the period in which the forecasts for the euro area are also available, we find a very similar pattern with a correlation of 88.4%.

0 and set the indicator variable in our local projection model for the ZLB-state λ_t to 1 from June 2013 onward. Secondly, the shadow short rate is also meant to capture the unconventional policy measures conducted by the ECB, such as forward guidance and quantitative easing, which in our theoretical model play an important role in generating a possible real depreciation when interest rates are at the ZLB. Finally, the (shadow) short rate will help us to distinguish demand from monetary policy shocks when identifying a country-specific negative demand shock. The last two variables contain both domestic (euro area) and foreign information: (5) the bilateral real exchange rate and (6) relative industrial production, i.e. domestic production relative to foreign industrial production. The latter is also needed to identify a domestic demand shock.[17]

Industrial production, consumer prices, and the exchange rate are included in logs. The impulse responses for these variables are therefore understood as percentage changes. Both inflation expectations and the short rate are expressed in percentage points. Our sample starts in January 1999 and ends in February 2020 to avoid possible asymmetries in the propagation of the extremely large structural shocks during the ongoing pandemic.[18] Our baseline specification includes $p = 2$ lags of the endogenous variables. However, in the robustness section, we show that our results are robust to different lag choices.

C. *Identification of a Country-Specific Demand Shock*

We identify a domestic demand shock by means of sign restrictions, which are summarized in Table (1).

Table 1: Sign restrictions of a country-specific demand shock

ip	*prices*	*short rate*	*relative ip*	*exp. inflation*	*real exchange rate*
−	−	−	−		

Notes: The restrictions on industrial production, prices, and the (shadow) short rate hold on impact and for five consecutive months. The restriction on relative industrial production holds on impact.

Our theoretical model of Section II suggests, as illustrated in Figure (2), that an adverse demand shock leads to a decline in domestic production and prices. However, we would expect the same reactions of production and prices following a contractionary monetary policy shock. To separate demand shocks from monetary policy shocks, we include in our identification strategy the (shadow) short rate in our empirical model to account for a reduction in the policy rate to counteract deflationary pressure from a negative demand shock.

[17] In this paper, the real exchange rates are derived such that an increase refers to a depreciation, while a decrease refers to an appreciation.

[18] See Lenza and Primiceri (2020) for a thorough discussion and suggestions on how to handle a sequence of extreme observations such as those recorded during the COVID-19 pandemic.

Moreover, as shown in Figure (2), our model predicts that production in the euro area should fall more strongly than in the foreign country in response to a negative asymmetric euro area demand shock. This finding also allows for separating domestic from foreign demand shocks. The distinction can be especially important when two open economies are particularly strongly connected, and a shock from abroad also has effects on domestic variables (see Bobeica and Jarociński, 2019 and Corsetti et al., 2014). In other words, with this identification strategy, we ensure that the reaction of domestic variables is not a spillover effect due to foreign shocks. As shown in Table (1), we leave the response of inflation expectations and the real exchange rate unrestricted, which are our key variables of interest.

Figure 3: Shocks across different country pairs

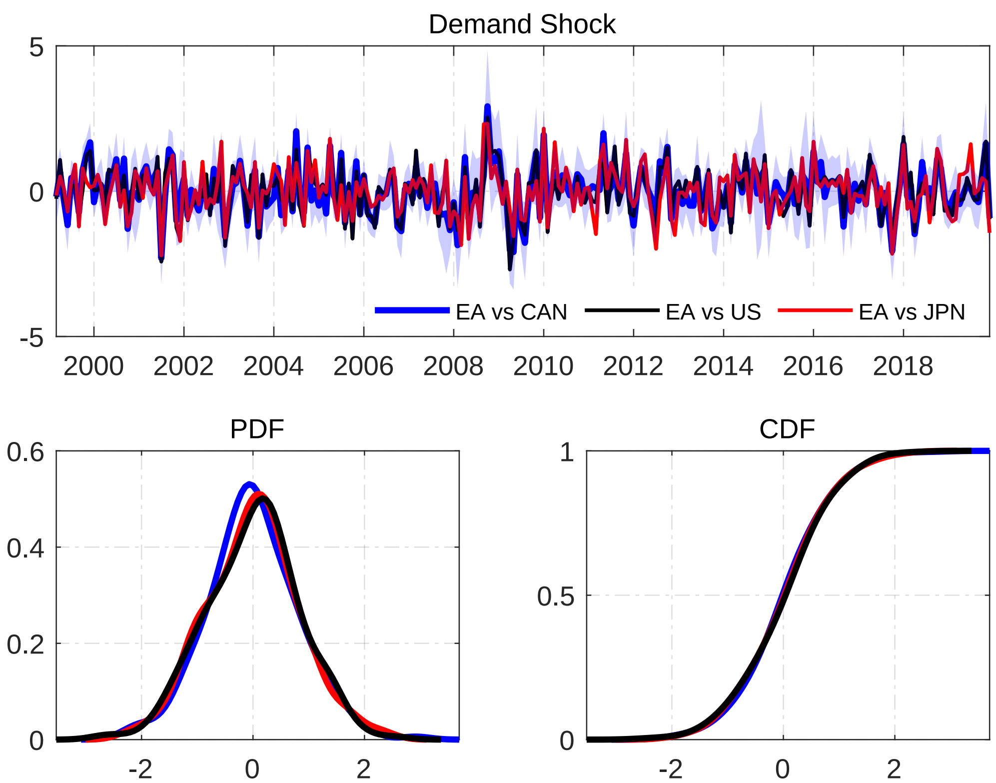

Notes: The upper panel shows the median shocks for all three country pairs: euro area vis-à-vis Canada (blue line), euro area vis-à-vis the US (black line), and euro area vis-à-vis Japan (red line). The blue-shaded areas correspond to the 5*th* and 95*th* percentiles for the euro area vis-à-vis Canada model. For all three country pairs, the lower panels show the cumulative distribution functions (left panel) and the probability density functions (right panel) over the full sample.

Figure (3) shows the shock distribution of the euro area demand shock across the different country pairs. The upper panel illustrates the median shock for the euro area vis-à-vis Canada (blue line), the US (black line), and Japan (red line). The figure shows that the obtained asymmetric demand shock across the three country pairs, which originates in the euro area, are well aligned. This is confirmed by the second panel of Figure (3), which shows that their cumulative distributions and the probability densities are very similar. Based on this, we now turn to the inference of the structural impulse responses.

D. *Inference*

This section outlines our procedure to draw structural impulse response functions from our sign-restricted state-dependent LP model. Sign and zero restrictions are easy to implement and well understood in VAR frameworks (see, for instance, Rubio-Ramirez et al., 2010; Arias et al., 2018; Uhlig, 2005). Plagborg-Møller and Wolf (2021) show that sign and zero restrictions can also be easily implemented within an LP framework. The idea relies on the finding that the projection coefficients from an LP framework are the reduced-form impulse responses of y_t with respect to VAR-based Wold innovations $e_t = y_t - E\left(y_t | \{y_\tau\}_{\tau<t}\right)$ at horizon h and that the projection residuals $u_{1,1,t} \dots u_{n,1,t}$ equal these Wold innovations. As a result, the variance-covariance matrix Σ obtained within an LP framework contains the same information as obtained from a VAR framework. As regards the implementation of zero and sign restrictions within an LP framework, the procedure can be sketched as follows. In a first step, for each horizon $h = 0 \dots H$, we estimate the model and store the coefficients $\beta^I_{i,h}$ for state I=*pre-ZLB* and $\beta^{II}_{i,h}$ for state II=*ZLB* in appropriate vectors $C^I_h = \left[\beta^I_{1,h}\ \beta^I_{2,h} \dots \beta^I_{n,h}\right]$ and $C^{II}_h = \left[\beta^{II}_{1,h}\ \beta^{II}_{2,h} \dots \beta^{II}_{n,h}\right]$, respectively. Plagborg-Møller and Wolf (2021) show that structural impulse functions at horizon h can be derived as

$$\Theta^j_h\left(Q, C^j_h, f(\Sigma)\right) = C^j_h f(\Sigma) Q,$$

where $f(\Sigma)$ is an appropriate decomposition (e.g. Cholesky) of the horizon-1 projection residuals $Var(u_{1,1,t} \dots u_{n,1,t}) = f(\Sigma) f(\Sigma)'$ and Q is an orthogonal matrix with $QQ' = Q'Q = I_n$. The boundaries of the identified set for the impulse responses for the ith variable at horizon h can be easily approximated numerically through random draws of orthogonal matrices Q as in Arias et al. (2018), which are subject to

$$\begin{aligned} \mathbf{S}_k \Theta^j\left(Q, C^j, f(\Sigma)\right) e_k &\geq 0 \\ \mathbf{Z}_k \Theta^j\left(Q, C^j, f(\Sigma)\right) e_k &= 0, \end{aligned}$$

where $\Theta^I = \left[\Theta'^I_0\ \Theta'^I_1 \dots \Theta'^I_H\right]$ and $\Theta^{II} = \left[\Theta'^{II}_0\ \Theta'^{II}_1 \dots \Theta'^{II}_H\right]$ are the $n(H+1) \times n$ matrices of stacked impulse response coefficients in state I and state II, respectively, and the $n(H+1) \times n(H+1)$-dimensional matrices $\mathbf{S}_k$ and $\mathbf{Z}_k$ are set as in Rubio-Ramirez et al. (2010) with e_k being the kth column of the identity matrix.

Inference on the impulse responses is based on various percentiles over all draws that satisfy our zero and sign restrictions. As pointed out by Fry and Pagan (2011), this procedure does not report sampling (estimation) uncertainty as is typically done in a conventional LP approach using robust standard errors. Instead, it reports the distribution across the models that satisfy the set of zero and sign restrictions. Note that our algorithm ensures we use the same draws for Q in both states. In other words, we discard the draw when a candidate for Q satisfies all sign and zero restrictions in state *I=pre-ZLB* but not in state *II=ZLB*, and vice-versa.

IV Results

In this section, we assess the responses of the euro area to a negative domestic demand shock. In the first step, we present the results from our main specification by country pairs, as discussed in the previous section. In a second step, we assess the reaction of the ECB's balance sheet to the negative domestic demand shock to further understand the role of unconventional policy measures in accommodating the negative disturbance when interest rates are at the ZLB. Finally, we explore the importance of the real exchange rate response to the negative demand shock for output through forecast error variance decompositions.

A. *Main Results*

Since the Canadian economy was not bound by the ZLB, we can draw a clear inference regarding the effects of the euro area's ZLB period on the real exchange rate and inflation expectations. We then compare the results to those of the euro area vis-à-vis Japan and the US.

Figure (4) shows our results for the euro area vis-à-vis Canada. Across all panels, the solid blue line shows the responses for the state in which the ZLB is binding, while the solid red line shows the responses for the state in which it is not. The blue (red) shaded areas cover the 5*th* and 95*th* percentiles. Hence, we report 90% confidence bands.

A few things stand out. First, we find that, following an unexpected demand shock, industrial production falls and reverts much faster to its expected value when the economy's interest rates are at the ZLB compared to when they are not. More precisely, after an initial drop in industrial production of about 0.7 percent, it takes about seven months until industrial production reverts to its mean when the ZLB is binding, while the mean reversion takes 13 months when the ZLB is not binding, therefore resulting in a higher overall cumulative effect in state *I*. This can probably explain why the reaction of the (shadow) short rate is found to be stronger in state *I* than in state *II*, as the central bank cuts the short rate by more (in absolute terms) to compensate for the larger drop in industrial production.[19]

[19]Note that the confidence bands around the response of industrial production are somewhat wider in state *I* than in state *II*. A candidate explanation would be the high volatility around the Great Recession, for which we do not explicitly account. We have therefore re-estimated all models and included the OECD-based Recession Indicators as an additional explanatory variable in the model to increase the in-sample fit in periods marked as recessions. However, we find that even in this case, the uncertainty in state *I* is still noticeably higher than in state *II*.

Figure 4: Results for the euro area vis-à-vis Canada

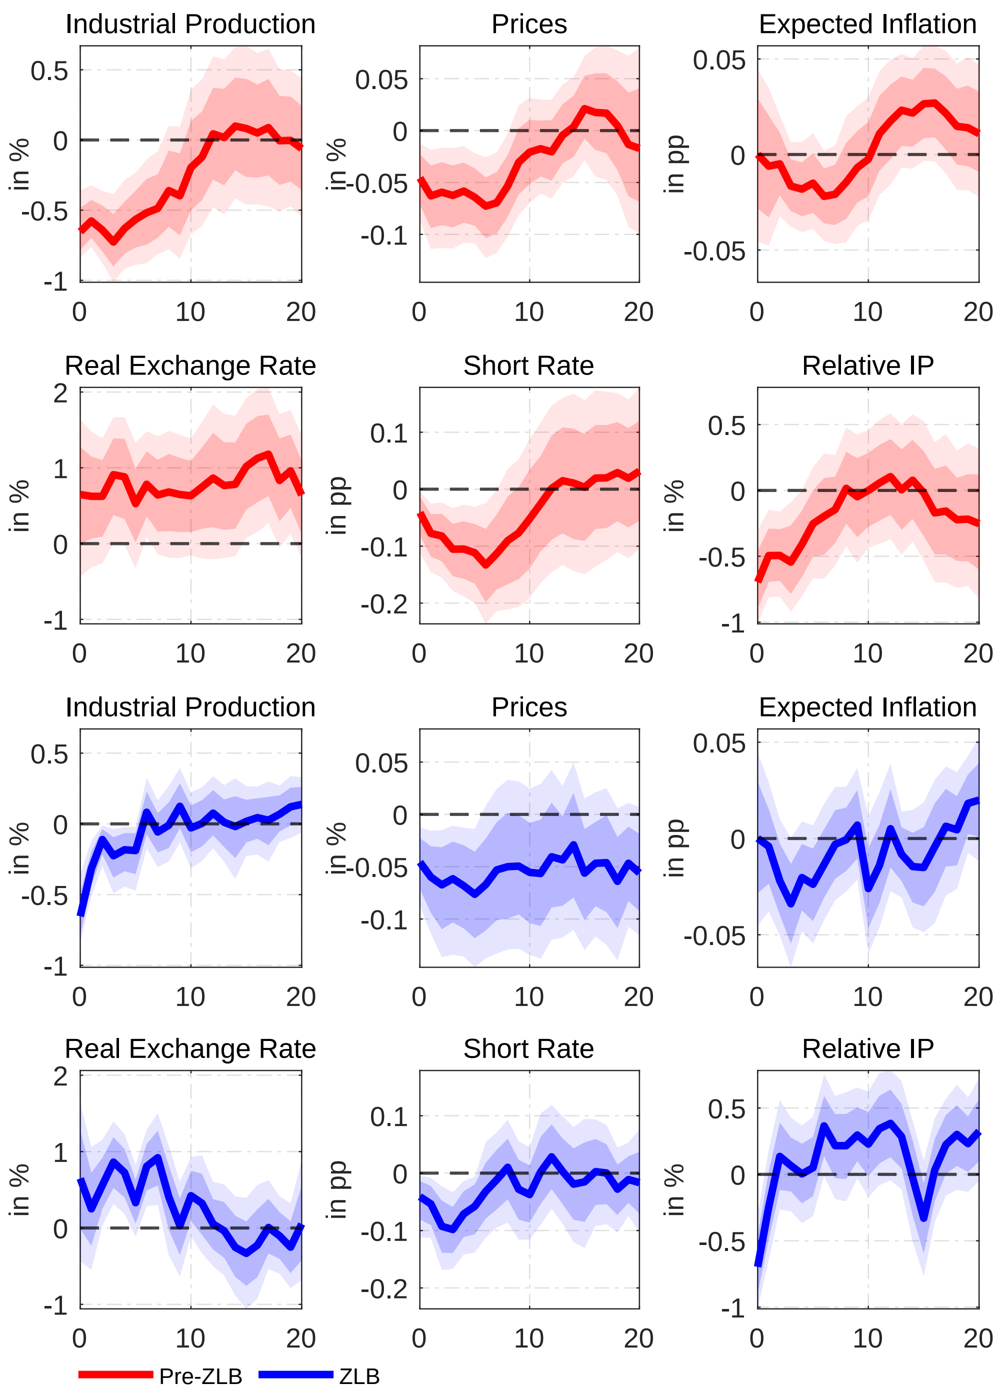

Notes: Median impulse response for the pre-ZLB period (red line) and the ZLB period (blue line). For both states, the shaded areas correspond to the 5*th* (95*th*) and the 16*th* (84*th*) percentiles.

Second, we find that the reaction of both prices and one-year-ahead inflation expectations is very similar across both states. Although the response of expected inflation is somewhat noisier in a ZLB scenario than in a non-ZLB scenario, both impulse responses show a similar picture. Interestingly, this implies that the decline in inflation expectations in the ZLB period is not stronger than in the pre-ZLB period. We find that the decline in both sub-samples is around two basis points within the first quarter and that, in both states, inflation expectations return to their mean after ten months or so. Note that these findings are more in line with the results of the open economy New Keynesian model in Section II, which predicts that both inflation and expected inflation are less volatile in a ZLB scenario when monetary policy can counteract the negative demand shock through unconventional monetary policy measures.

Third, the real exchange rate depreciates significantly in both samples, i.e. both before the ZLB has been reached and once the ZLB has been reached. The size of the real depreciation is very similar and amounts to one percent in both samples after the first quarter. As before, we find that the uncertainty for the impulse responses in a non-ZLB period is larger than for the responses in a ZLB period. Moreover, we again find that mean reversion occurs much faster when interest rates are at the ZLB than when the ZLB is not binding. The real exchange rate depreciation in a ZLB scenario is contrary to the appreciation predicted by the model in Section II, where monetary policy would not respond beyond the ZLB constraint. However, the real depreciation is in line with the model extension, where monetary policy can counteract the negative demand shock by unconventional policy measures. Similar to the reactions of prices and expected inflation, the reaction of the real exchange rate when interest rates are at the ZLB is not stronger than when interest rates are not at the ZLB.

Figure (5) shows our findings of the euro area vis-à-vis Japan. Throughout our estimation period from January 1999 until February 2020, Japan's interest rates remained at the ZLB, as illustrated by Figure (1). Therefore, Japan offers a clean counterpart to our assessment of the euro area vis-à-vis Canada. Note that in this case, too, we find that the reactions of both industrial production and the (shadow) short rate show a pattern that is very similar to the first country pair, although the difference in the responses of industrial production across both regimes is less pronounced. Most importantly, we find again that a negative euro area demand shock leads to a decline in both prices and expected inflation, as well as a real depreciation of the euro exchange rate. This real depreciation occurs within both monetary policy regimes i.e. in the pre-ZLB period and once interest rates are at the ZLB. Figure (5) illustrates that prices, expected inflation, and the real exchange rate move very closely in both monetary policy regimes for the first 12 months. Importantly, however, the decline in prices and inflation expectations is not more accentuated when interest rates are at the ZLB. At the same time, we find that the real exchange rate does not react more strongly when

Figure 5: Results for the euro area vis-à-vis Japan

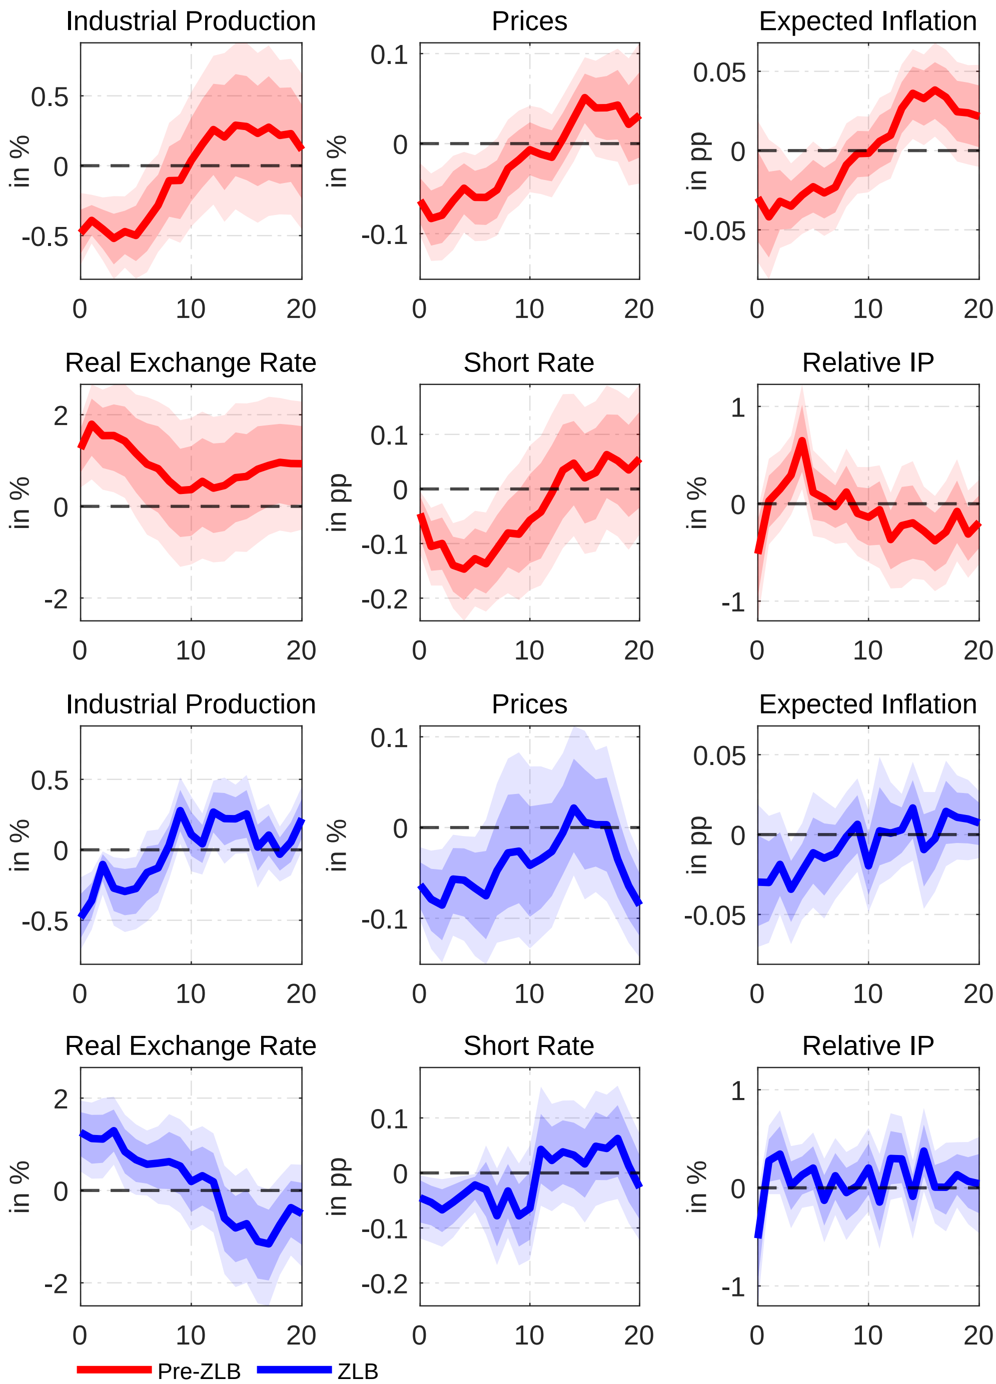

Notes: Median impulse response for the pre-ZLB period (red line) and the ZLB period (blue line). For both states, the shaded areas correspond to the 5*th* (95*th*) and the 16*th* (84*th*) percentiles.

the economy's interest rates are at the ZLB. Figure (6) provides the impulse responses

Figure 6: Results for the euro area vis-à-vis the US

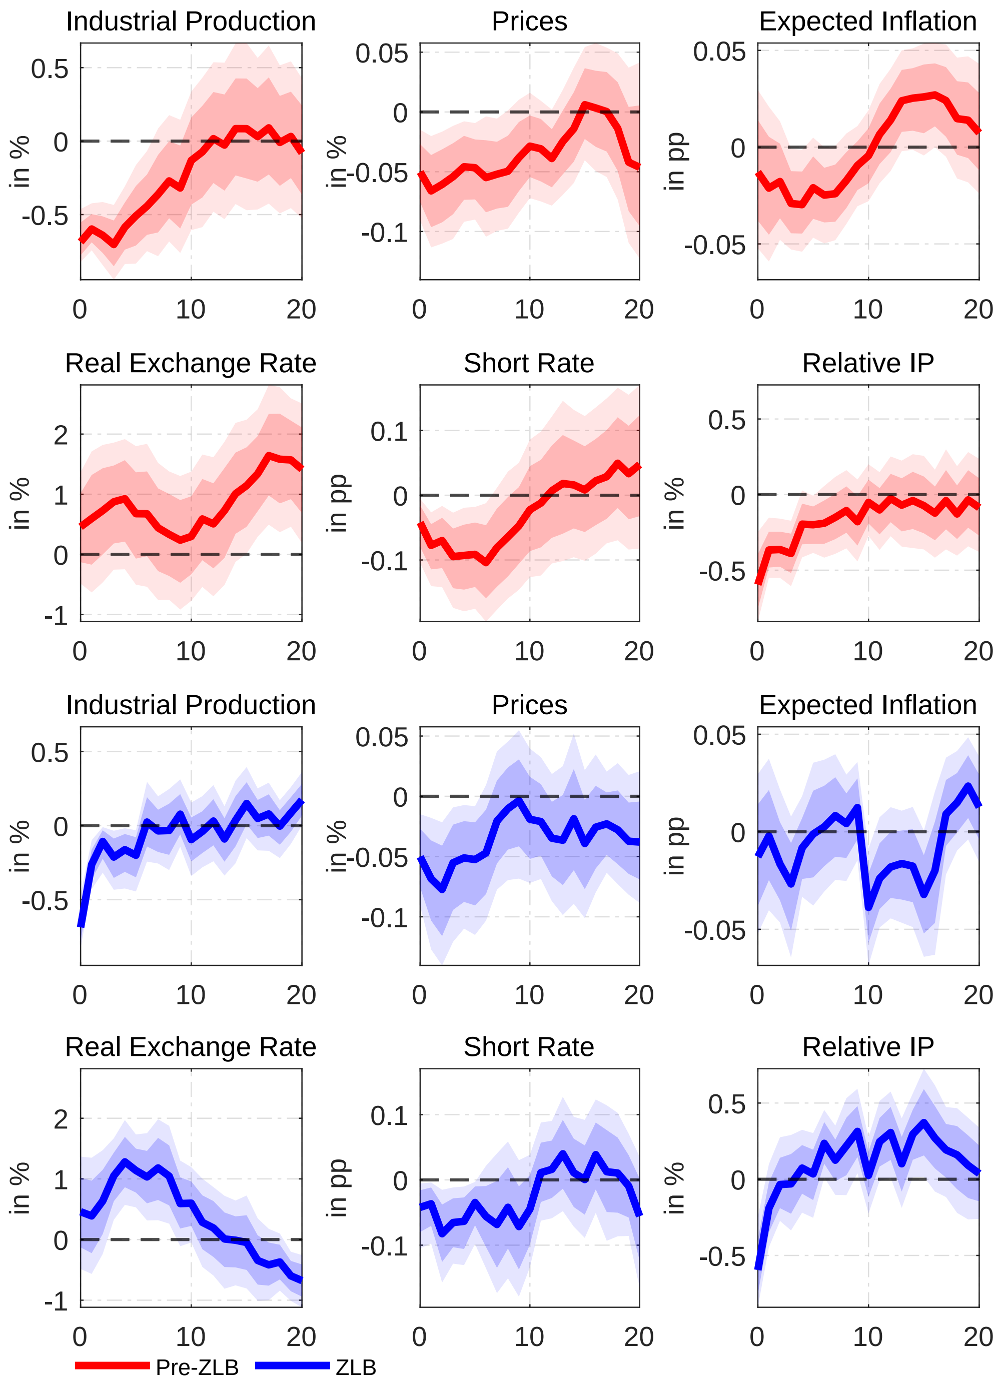

Notes: Median impulse response for the pre-ZLB period (red line) and the ZLB period (blue line). For both states, the shaded areas correspond to the 5*th* (95*th*) and the 16*th* (84*th*) percentiles.

of the euro area vis-à-vis the US. In contrast to Japan, the US economy switched policy regimes. Since the Great Recession, the US has experienced times when the policy rate was constrained by the ZLB or remained unconstrained, as shown by Figure (1). To keep things short, our main findings are qualitatively robust for this country pair. Regardless of the euro area's state of being constrained or not constrained by the ZLB, we find a depreciation of the real exchange rate and the responses of prices and inflation expectations are within the two regimes' confidence bands.

B. *Unconventional Monetary Policy*

In the next step, we assess the ECB's balance sheet response to understand the role of unconventional policy measures in accommodating the negative disturbance when interest rates are at the ZLB. The motivation for this assessment can be seen in Figure (11) in the appendix. Before the ZLB was binding, the ECB's balance sheet grew modestly over time. However, from June 2014 onwards, the ECB's balance sheet increased tremendously. The reason was the ECB's large-scale asset purchases, which were mirrored in the ECB's balance sheet as part of the ECB's forward guidance, as outlined in Section II.

We re-estimate our model for each country pair by adding the ECB's balance sheet as a seventh variable (in logs), keeping everything else to the benchmark specification. Importantly, we also leave the balance sheet response unrestricted. As discussed in Section II, the large-scale asset purchases were conducted during the ZLB period to reduce long-term interest rates by driving down the yields on the securities the ECB was purchasing to meet their forward guidance. This should lead to lower interest rates throughout the economy and stimulate economic activity. Therefore, in our empirical assessment, we should expect a significant rise in the ECB's balance sheet to counteract the ZLB constraint during our ZLB state, while we would only expect an unincisive reaction of the balance sheet in the pre-ZLB period.

Figure (7) reports the results for this exercise. Since all the remaining impulse responses look similar to those in our baseline specification, we show the balance sheet response for all three country pairs. It stands out that only when interest rates are at the ZLB do we see a significant, expansionary balance sheet response. More specifically, the balance sheet rises significantly in all three cases after a demand shock, peaking at around two percent after about ten months.

In non-ZLB scenarios, on the other hand, we do not observe any significant balance sheet reaction. The responses fluctuate around zero for all three country pairs and even show negative values after about ten months. However, these responses are not significant.

Overall, the balance sheet responses during the ZLB period are one good explanation for the accommodative decline of the shadow short rate in our baseline model in reaction to the negative demand disturbance.

C. *Quantifying the Role of the Real Exchange Rate as a Shock Absorber*

In this section, we quantify the potential of the exchange rate as a shock absorber. More specifically, we assess by how much the depreciation of the real exchange rate

Figure 7: State-dependent response of the ECB's balance sheet

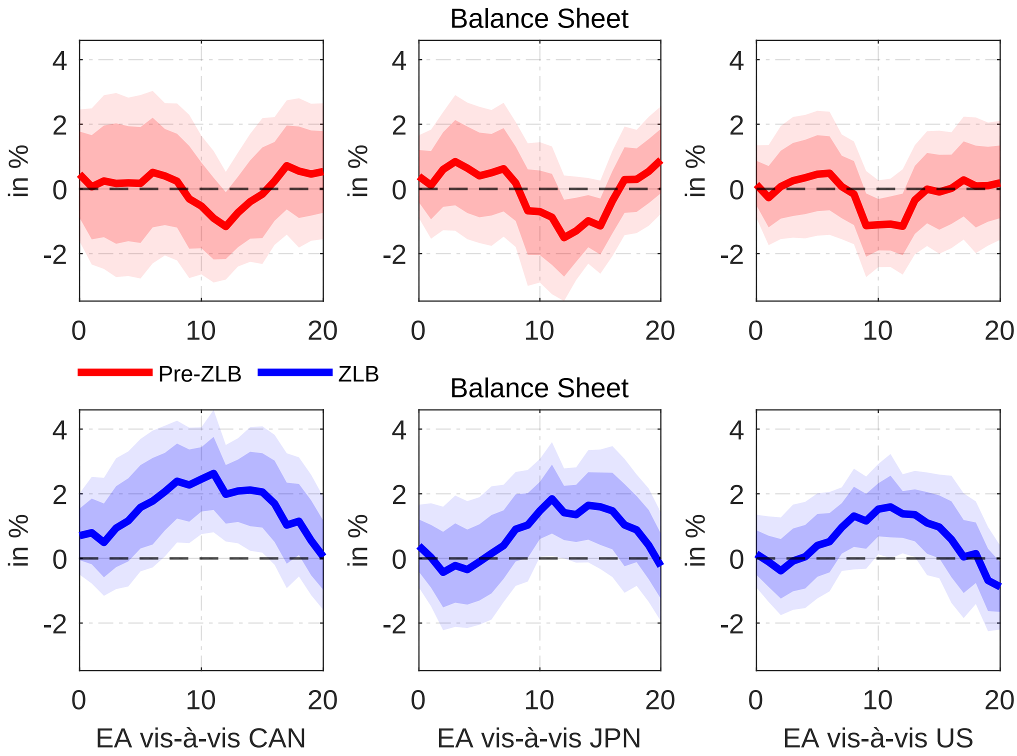

Notes: Median impulse response for the pre-ZLB period (red line) and the ZLB period (blue line). For both states, the shaded areas correspond to the 5*th* (95*th*) and the 16*th* (84*th*) percentiles.

cushions the effects of the adverse demand shocks on output and prices.

To do so, we proceed in two steps. Firstly, we simulate a counterfactual scenario by extracting the feedback effect of the response of the exchange rate on the response of the other endogenous variables. An appropriate environment for such an exercise relies on restricting the lag polynomials in the exchange rate equation to zero, except for its own lags. Thus, we assume that the exchange rate does not respond to past movements of all remaining endogenous variables. This, in addition to a zero restriction on the response on impact, guarantees that the exchange rate does not react to the demand shock over the entire projection horizon $h = 0, 1 \ldots H$. Our approach is in line with Bachmann and Sims (2012), who analyze the role of consumer confidence in the propagation of government spending shocks by isolating the feedback effects from the response of confidence on the response of the remaining endogenous variables.

Secondly, we calculate the forecast error variance decomposition (FEVD) for the baseline specification and the counterfactual scenario. Note that the FEVD tells us how important the demand shock is in explaining the variations of the variables in the model. It is important to note that the differences between the FEVDs must come from the suppressed feedback effect of the exchange rate on the other variables. Hence, the difference in the FEVDs tells us how much the depreciation of the real exchange

Table 2: Forecast error variance decompositions

		euro area vis-à-vis Canada				
		Pre-ZLB			*ZLB*	
	benchmark	*counterfactual*	*difference*	*benchmark*	*counterfactual*	*difference*
$h = 0$	67.2	77.7	10.5	67.2	77.7	10.5
$h = 3$	63.3	71.7	8.4	39.3	41.0	1.7
$h = 6$	52.9	60.7	7.8	30.3	30.8	0.5
$h = 12$	29.6	35.3	5.7	21.0	20.2	-0.8
		euro area vis-à-vis Japan				
		Pre-ZLB			*ZLB*	
	benchmark	*counterfactual*	*difference*	*benchmark*	*counterfactual*	*difference*
$h = 0$	46.7	75.5	28.8	46.7	75.5	28.8
$h = 3$	45.8	55.0	9.2	30.3	37.5	7.2
$h = 6$	38.4	39.5	1.1	24.7	23.2	-1.5
$h = 12$	20.7	22.1	1.4	20.4	15.9	-4.6
		euro area vis-à-vis the US				
		Pre-ZLB			*ZLB*	
	benchmark	*counterfactual*	*difference*	*benchmark*	*counterfactual*	*difference*
$h = 0$	73.7	81.2	7.5	73.7	81.2	7.5
$h = 3$	58.8	70.4	1.6	39.2	44.5	5.3
$h = 6$	54.0	52.6	-1.4	27.5	30.3	2.7
$h = 12$	26.6	24.8	-1.8	19.8	21.2	1.4

Notes: All values are based on the median estimates and are rounded to one decimal point.

rate cushions the effects of the adverse demand shocks.

Table (2) reports the results for this exercise. A few things stand out. First, the exchange rate can cushion much of the adverse effects of the demand shock, especially in the first few months after the shock occurs. We find for all three country pairs that the explanatory power of the demand shock on industrial production would be much higher if the exchange rate could not respond. For example, for the country pair euro area vis-à-vis the US, we find that in the baseline specification on impact, 73.7 percent of the forecast errors are explained by the demand shock in both states. In the counterfactual scenario, on the other hand, this value is 81.2 percent, i.e. around eight percentage points more. We also find a qualitatively similar picture and an even higher absorbing effect of the real exchange rate for the other two country pairs. Second, the absorbing effect of the exchange rate seems to be relatively short-lived. We find for all country pairs that the absorbing effect of the exchange rate fades after six months.

In conclusion, our empirical results in Table (2) show that the real exchange rate can absorb a considerable amount of variations in output, confirming its shock-absorbing capacity before and during the ZLB episode.

V Robustness Checks

In this section, we assess the robustness of our findings. To do so, we proceed as follows: First, we re-estimate our model for a broad set of 16 countries. Second, we investigate how our results change when we allow for smooth transitions between the states of the economy. Third, we compare our results from the LP to the empirical evidence from a state-dependent VAR. Finally, we compare the effects from different lag structures.

A. *Considering a Broad Set of Countries*

Our main results rely on three country pairs representing different monetary policy stances. In this subsection, we re-estimate our model for the euro area vis-à-vis a broader set of countries, namely 16 different trading partners for which trade weights are available. These countries consist of the 15 trading partners considered by the Bank of International Settlement (BIS) for the calculation of the *narrow index of the real effective* exchange rate, as well as China.[20] Based on the time-varying trade weights, as reported in Table (3), we calculate the real effective exchange rate for this set of countries.

Moreover, we use the same trade weights to calculate a weighted time series for industrial production in these countries. Hence, we summarize the information for all 16 countries and treat them as a single country. Finally, we calculate the share of industrial production in the euro area to this weighted industrial production series. By doing so, we are fully consistent with our baseline identification strategy.

We calculate the trade-weighted industrial production for the set of countries as

$$\overline{ip}_t = \chi_t^{1,ea} ip_{1,t} + \chi_t^{2,ea} ip_{2,t} + \ldots + \chi_t^{16,ea} ip_{16,t},$$

where $\chi_t^{j,ea}$ is the trade weight for country $j = 1, \ldots, 16$ in the real effective exchange rate, and $ip_{j,t}$ refers to the seasonally adjusted industrial production index of country j in period t.[21] Figure (8) shows the impulse responses for this exercise. Two things stand out. First, the impulse responses look similar to those presented in Section IV. In particular, we observe a very similar pattern for the responses of industrial production, prices, expected inflation, and the (shadow) short rate. Second, we find again that the exchange rate depreciates both when interest rates are at the ZLB and when they are not. More precisely, the exchange rate response looks very similar to the responses of our model for the euro area vis-à-vis the US. That is, the exchange rate significantly

[20]These countries comprise Australia, Canada, Taiwan, Denmark, Hong Kong, Japan, Korea, Mexico, New Zealand, Norway, Singapore, Sweden, Switzerland, the United Kingdom, and the United States. Importantly, we rescaled the trade weights for the 15 countries included in the narrow index of the real effective exchange rate such that the trade weights of these countries and China sum up to 100.

[21]Note that for Australia, New Zealand, and Switzerland, industrial production is only available on a quarterly basis. We, therefore, interpolate these series in order to get monthly data.

Figure 8: Results for the euro area vis-à-vis 16 trading partners

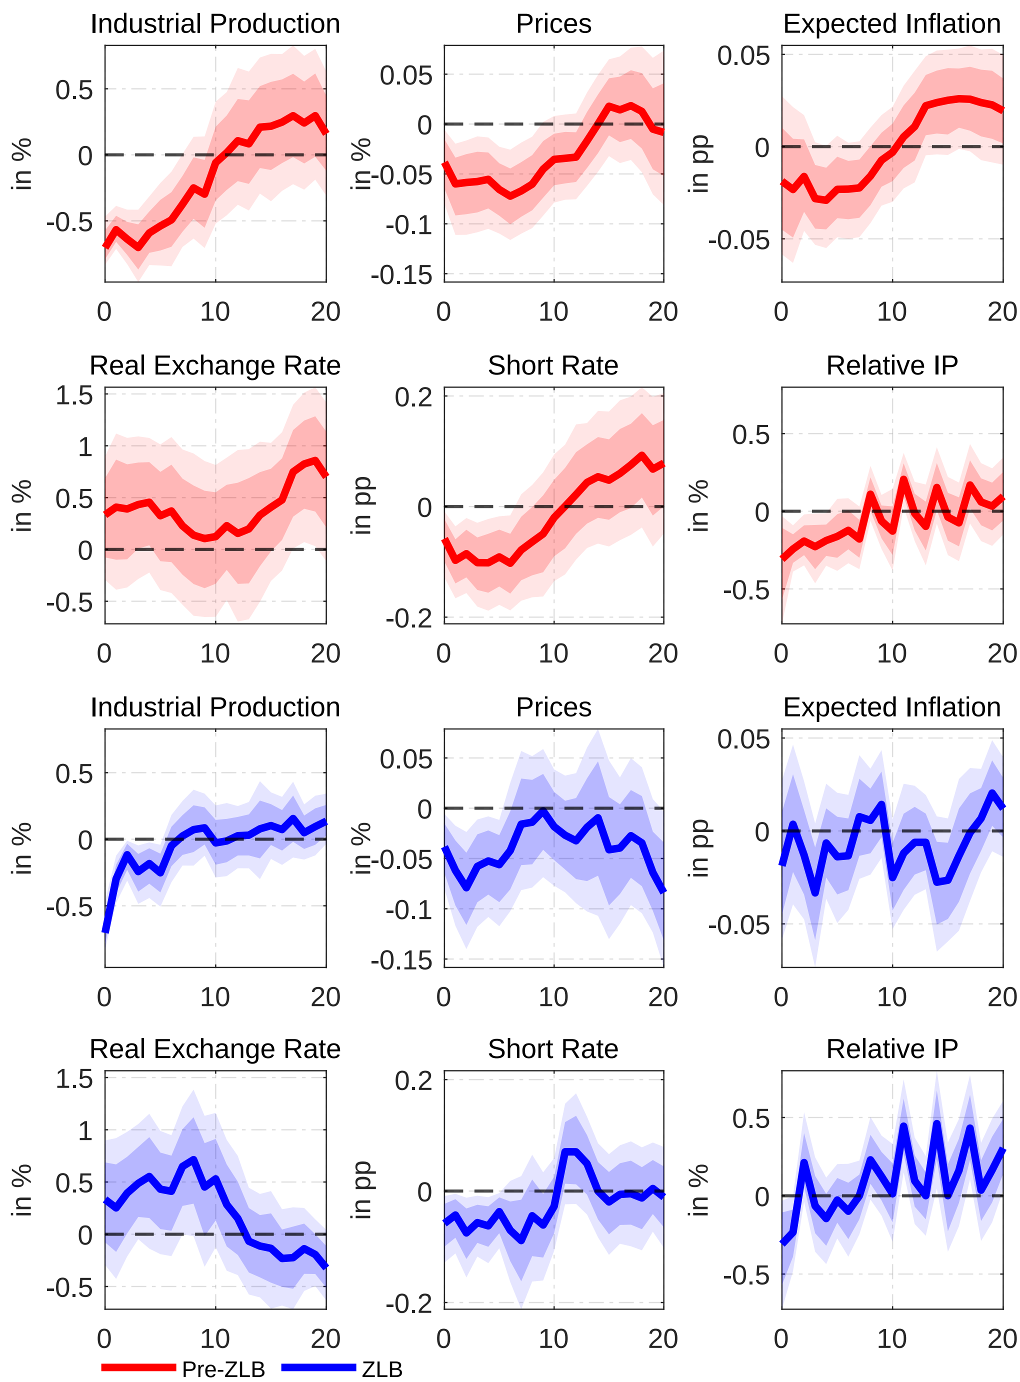

Notes: Median impulse response for the pre-ZLB period (red line) and the ZLB period (blue line). For both states, the shaded areas correspond to the 5*th* (95*th*) and the 16*th* (84*th*) percentiles.

Table 3: Trade weights

Country \Years	**1999–2001**	**2002–2004**	**2005–2007**	**2008–2010**	**2011–2013**	**2014–2020**
			ASIA			
China	7.5	12.0	17.5	23.6	26.2	26.9
Hong Kong	0.8	0.6	0.5	0.4	0.5	0.4
Japan	11.5	10.3	9.2	8.4	7.7	6.5
Korea	3.2	3.7	4.3	4.2	4.1	4.0
Singapore	2.1	2.0	2.0	1.9	2.0	1.8
Taiwan	2.9	2.6	2.5	2.4	2.3	2.2
			AMERICA			
Canada	2.2	2.1	2.1	2.0	2.0	1.9
Mexico	1.6	1.6	1.6	1.7	2.0	2.1
United States	26.3	23.6	21.4	19.9	19.5	21.3
			EUROPE			
Denmark	3.0	3.1	3.1	2.9	2.5	2.5
Norway	1.5	1.6	1.7	1.7	1.5	1.3
Sweden	5.5	5.6	5.6	5.2	5.1	4.6
Switzerland	8.1	8.4	8.0	8.7	8.6	8.1
United Kingdom	22.8	21.7	19.3	15.9	14.7	15.3
			OCEANIA			
Australia	0.9	1.0	1.0	1.1	1.1	0.9
New Zealand	0.2	0.2	0.2	0.2	0.2	0.2
Sum Σ	100	100	100	100	100	100

Notes: All values are taken from the Bank of International Settlement. Values are rounded to one decimal point.

depreciates after half a year or so in the ZLB period, while the same is true after one year when the ZLB is not binding.

Table (4) reports the results from the same counterfactual exercise as in Section IV. As can be seen, the exchange rate also absorbs a considerable part of the adverse effects

Table 4: FEVD for the euro area vis-à-vis 16 trading partners

		Pre-ZLB			*ZLB*	
	benchmark	*counterfactual*	*difference*	*benchmark*	*counterfactual*	*difference*
$h = 0$	78.7	85.3	6.6	78.7	85.3	6.6
$h = 3$	74.2	76.6	2.4	43.1	46.9	3.8
$h = 6$	63.0	62.9	–0.1	31.4	32.6	1.2
$h = 12$	35.7	35.1	–0.6	18.0	18.2	0.2

Notes: All values are based on the median estimates and are rounded to one decimal point.

caused by the demand shock. However, as for the three country pairs in Section IV, this effect lasts only a few months. In particular, our results indicate that the effect disappears after six months and the values for the baseline specification and the counterfactual scenario are almost identical.

Overall, we find that both the impulse responses and the role of the exchange rate as a shock absorber are very similar to the results from the previous section.

B. *Allowing for Smooth Transitions*

An important assumption in our analysis is that the economy jumps instantly from one state to another. However, the possibility of the ZLB being reached might affect economic decisions and expectations even before the ZLB becomes binding. If economic agents anticipate that the policy rate might reach the ZLB in the future, their inflation expectations might already react today. Hence, the economic consequences of the ZLB can already be effective even before the short rate reaches zero. We acknowledge this by allowing for a smooth transition between the two regimes *I=pre-ZLB* and *II=ZLB*. The model then reads as

$$y_{i,t+h} = (1 - F(z_t))\left(\alpha^{I}_{i,h} + \beta^{I}_{i,h} y_t + \gamma^{I}_{i,h} x_t\right) + F(z_t)\left(\alpha^{II}_{i,h} + \beta^{II}_{i,h} y_t + \gamma^{II}_{i,h} x_t\right) + u_{i,h,t},$$

where everything is kept similar to the benchmark case, except that we replace λ_t with a logistic transition function $F(z_t)$ of the form

$$F(z_t) = \frac{\exp\left(\kappa \frac{z_t - \mu}{\sigma_z}\right)}{1 + \exp\left(\kappa \frac{z_t - \mu}{\sigma_z}\right)},$$

where μ is used to control the proportion of the sample the economy spends in either state, and σ_z is the sample standard deviation of the state variable z_t. Finally, the parameter κ controls how abruptly the economy switches from one state to the other following movements of the state variable z_t. As regards the choice of μ, we choose the

Figure 9: Evolution of the state variable over time

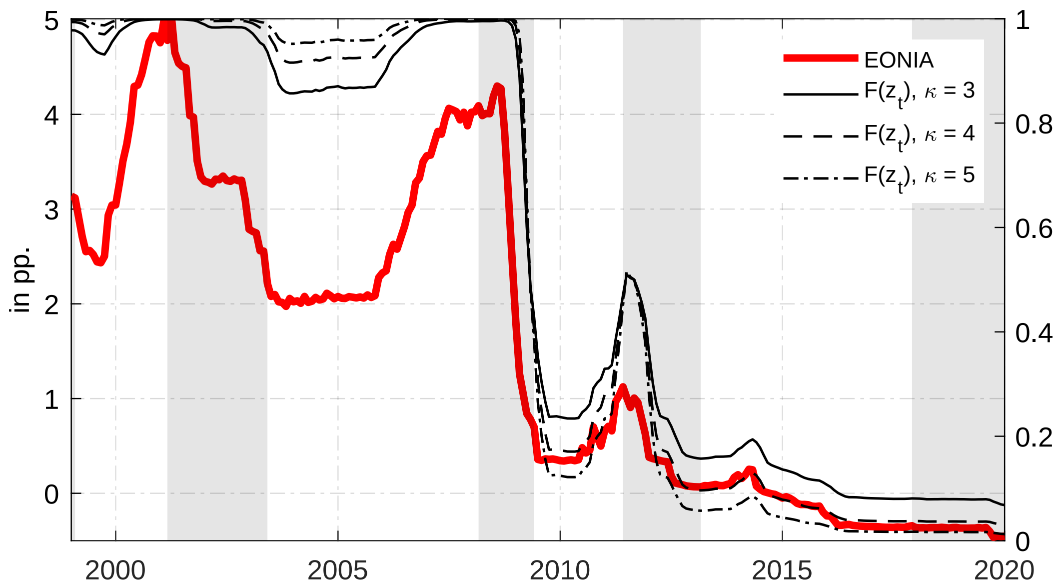

Notes: The red line corresponds to the EONIA short-term rate (left axis), while the black lines correspond to the evolution of the state variable for $\kappa \in [3, 4, 5]$ (right axis).

median value of the policy rate. This is based on the fact that, within our observation sample, the euro area policy rate was below one percentage point for 49.6% of the time. We choose three different values for κ, namely $\kappa \in [3, 4, 5]$. These values indicate an intermediate degree of intensity of the regime-switching behavior (see, for instance, Auerbach and Gorodnichenko, 2012; Tenreyro and Thwaites, 2016). Figure (9) shows the evolution of the state variable for all three parameters for κ over time. As can be seen, there is now a smooth transition to the ZLB state rather than the abrupt switch as in the baseline specification.

Figures (15) to (17) in the appendix show the results for this exercise together with the baseline results. Two things stand out. First, the results for all three country pairs are similar to those from the baseline specification with quasi-observable states. Second, the choice of κ does not seem to play a major role either, as the differences between impulse responses for different κ are negligible. While we do not plot the confidence bands for all individual impulse responses in order to avoid an unnecessarily crowded graph, we can make similar statements about significance as in our main results. This being said, here, too, we find a significant exchange rate depreciation after an adverse demand shock in both states.

Finally, one has to recognize that this result is also interesting in that the determination of the exact timing of the ZLB is not clearly defined. The methodology in this exercise circumvents this problem by dividing the proportions of states into periods in which the policy rate is either below or above one percentage point.

C. *Evidence From a State-Dependent SVAR*

As a further robustness check, for each country pair, we estimate a state-dependent SVAR. We use the same endogenous variables in $\mathbf{y}_t$ and rely on the same identification strategy as in the baseline model. The state-dependent VAR, in its reduced form, reads

$$\mathbf{y}_t = (1 - \lambda_t)\,\mathbf{B}^I\mathbf{x}_t + \lambda_t\mathbf{B}^{II}\mathbf{x}_t + \mathbf{e}_t,$$

where $\mathbf{y}_t$ is an $n \times 1$ vector of endogenous variables, $\mathbf{x}_t = \left[\mathbf{y}_{t-1}, ..., \mathbf{y}_{t-p}, 1, t\right]$ is an $n \times np + 2$ matrix of control variables, $\mathbf{B}^j = \left[\phi_1^j, ..., \phi_p^j, \mathbf{c}, \boldsymbol{\delta}\right]$ contains the reduced-form VAR coefficients and $\mathbf{e}_t$ is an $n \times 1$ vector of reduced-form error terms. Similar to the LP framework, the vector λ_t is equal to zero when the economy in period t is in state I and equal to one when the economy is in state II, respectively. We use the same identification strategy as in the LP framework, and for each model, we use $p = 2$ lags. This lag length is preferred by the AIC.

We find that the impulse responses are qualitatively similar to our benchmark results, although the VAR-based responses are somewhat smoother than the LP-based responses. On the one hand, this shows that our results are robust to the model specification, but on the other hand, it also shows that LPs and VARs produce qualitatively similar results even in a state-dependent framework (see, for instance, Plagborg-Møller and Wolf, 2021). Figure (10) illustrates those findings for the real exchange rate

Figure 10: Results from a state-dependent SVAR

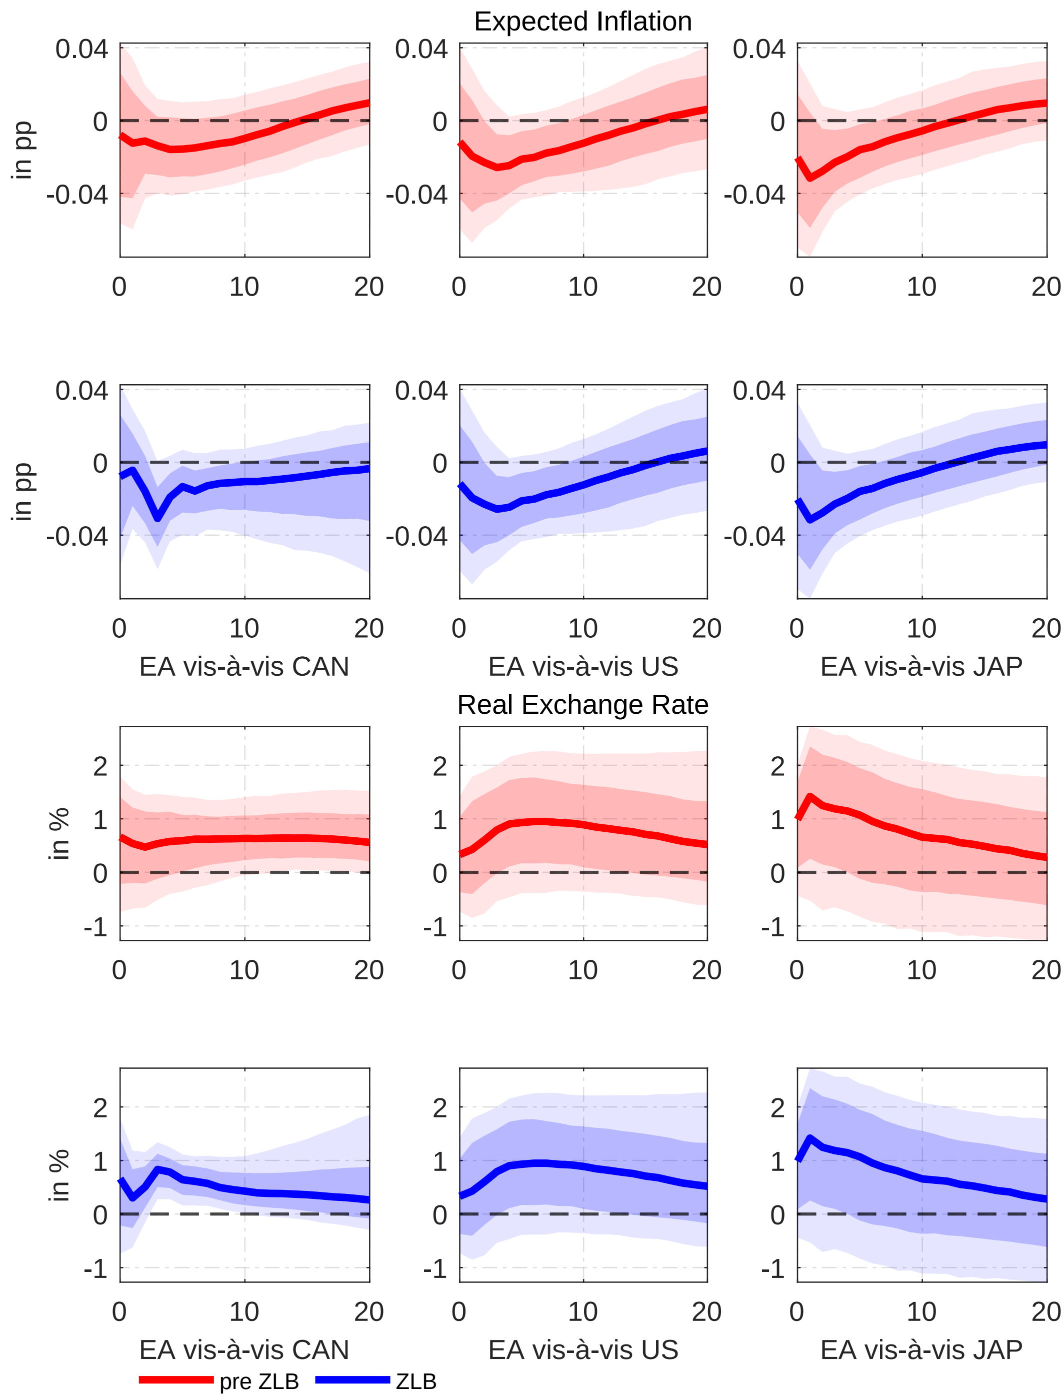

Notes: Median impulse response for the pre-ZLB period (red line) and the ZLB period (blue line). For both states, the shaded areas correspond to the 5*th* (95*th*) and the 16*th* (84*th*) percentiles.

and inflation expectations for the euro area vis-à-vis Canada, Japan, and the US. From the VAR-based assessment, we highlight two main findings: first, the real exchange rate depreciates significantly in both samples, i.e. prior to the ZLB being reached and once it has been reached for all country pairs. As for the LP analysis, the real exchange rate depreciation in the ZLB period contrasts with the predictions of the New Keynesian model outlined in Section II. Secondly, the decline in inflation expectations in the ZLB period is not stronger than in the pre-ZLB period, thereby again confirming our findings from the state-dependent LP. As mentioned above, the findings align with the predictions of the extended New Keynesian model of Section II, which allows the central bank to counteract the ZLB constraint through unconventional measures, in contrast to the version of the model in which unconventional policy measures do not feature.

The full set of results for state-dependent SVAR can be found in Figures (18) and (20).[22] They also confirm our findings from Section IV concerning the evolution of industrial production and prices across the three country pairs.

D. *Different Lag Specifications*

It is well known that the dynamic properties of impulse responses may depend on the lag order of the underlying model fitted to the data, both for VAR and LP frameworks. This subsection reports our baseline results for different lag-length specifications.[23] For each country pair, we, therefore, show impulse responses for the case where the baseline model is estimated for $p = 1$ to $p = 4$ lags. To best compare the results, each model starts with the fifth observation, including the baseline specification estimated with two lags.

Figures (12) to (14) in the appendix report the corresponding results for this exercise. Note that, for all country pairs and all endogenous variables, the estimated impulse responses for $p = 1, 3, 4$ follow our baseline results very closely and always lie within the confidence bands of our baseline specification of $p = 2$. We conclude that our results are robust concerning the chosen lag length.

VI Conclusion

According to the standard New Keynesian argument, the exchange rate cannot work as a shock absorber in response to an adverse demand shock when interest rates are at

[22] These figures also show the impulse responses from the baseline specifications obtained by local projections. Importantly, the impulse responses from the VAR and the LPs mostly overlap for the restricted variables over the restricted projection horizon. Thus, we can confirm the finding of Plagborg-Møller and Wolf (2021) also for a state-dependent model.

[23] We also re-estimated our models where we take the past value of the state variable λ_t, as is often recommended in the literature. The results do not change in any way. Concerning state-dependent local projections using external shocks, Gonçalves et al. (2022) discuss conditions under which the state-dependent local projections estimator can be expected to recover the true population impulse responses. They show that in the case of endogenous state variables, the state variable has to be a function of the number of lags of the endogenous variable being equal to the projection horizon. In our approach, we identify the structural shock within the model instead of using external shocks. We also do not explicitly model the state variable based on current and past realizations of the endogenous variables.

the ZLB. However, allowing for forward guidance as an additional policy tool suggests that the exchange rate could, in fact, also depreciate in a ZLB scenario. Our main contribution is to test this hypothesis empirically. We identify an adverse demand shock using sign restrictions and show that the real exchange rate depreciates both when interest rates are at the ZLB and when they are not. We find that the exchange rate fulfills its role as a shock absorber even when interest rates are at the ZLB, and output and inflation (expectations) are not significantly more volatile.

Our findings are robust for different country pairs: euro area vis-à-vis Canada, Japan, and the US. These countries represent different monetary stances among G7 countries. Furthermore, our empirical results show that the real exchange rate can absorb considerable variations in output, confirming its shock-absorbing capacity before and during the ZLB episode. Our findings are accompanied by a significant expansion of the ECB's balance sheet during the ZLB spell, while it remained unaffected in the pre-ZLB period. Overall, our empirical results support an open economy New Keynesian model with unconventional measures when interest rates are at the ZLB.

References

Alpanda, Sami and Sarah Zubairy, "Household Debt Overhang and Transmission of Monetary Policy," *Journal of Money, Credit and Banking*, 2019, *51* (5), 1265–1307.

— , **Eleonora Granziera, and Sarah Zubairy**, "State Dependence of Monetary Policy Across Business, Credit and Interest Rate Cycles," *European Economic Review*, 2021, *140*, 103936.

Amador, Manuel, Javier Bianchi, Luigi Bocola, and Fabrizio Perri, "Exchange Rate Policies at the Zero Lower Bound," *The Review of Economic Studies*, 2020, *87* (4), 1605–1645.

Antolín-Díaz, Juan and Juan F Rubio-Ramírez, "Narrative Sign Restrictions for SVARs," *American Economic Review*, 2018, *108* (10), 2802–2829.

Arias, Jonas E, Juan F Rubio-Ramírez, and Daniel F Waggoner, "Inference Based on Structural Vector Autoregressions Identified With Sign and Zero Restrictions: Theory and Applications," *Econometrica*, 2018, *86* (2), 685–720.

Auerbach, Alan J and Yuriy Gorodnichenko, "Measuring the Output Responses to Fiscal Policy," *American Economic Journal: Economic Policy*, 2012, *4* (2), 1–27.

Bachmann, Rüdiger and Eric R Sims, "Confidence and the Transmission of Government Spending Shocks," *Journal of Monetary Economics*, 2012, *59* (3), 235–249.

Blanchard, Olivier J and Danny Quah, "The Dynamic Effects of Aggregate Demand and Supply Disturbances," *American Economic Review*, 1989, *79* (4), 655–673.

Bobeica, Elena and Marek Jarociński, "Missing Disinflation and Missing Inflation: A VAR Perspective," *International Journal of Central Banking*, 2019, *15* (1), 199–232.

Calvo, Guillermo A, "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 1983, *12* (3), 383–398.

Christiano, Lawrence, Martin Eichenbaum, and Sergio Rebelo, "When Is the Government Spending Multiplier Large?," *Journal of Political Economy*, 2011, *119* (1), 78–121.

Clarida, Richard, Jordi Galı, and Mark Gertler, "A Simple Framework for International Monetary Policy Analysis," *Journal of Monetary Economics*, 2002, *49* (5), 879–904.

Coenen, Günter and Volker W Wieland, "Exchange-Rate Policy and the Zero Bound on Nominal Interest Rates," *American Economic Review*, 2004, *94* (2), 80–84.

Cook, David and Michael B Devereux, "Sharing the Burden: Monetary and Fiscal Responses to a World Liquidity Trap," *American Economic Journal: Macroeconomics*, 2013, *5* (3), 190–228.

— **and** — , "Exchange Rate Flexibility Under the Zero Lower Bound," *Journal of International Economics*, 2016, *101*, 52–69.

Corsetti, Giancarlo, Luca Dedola, and Sylvain Leduc, "The International Dimension of Productivity and Demand Shocks in the US Economy," *Journal of the European Economic Association*, 2014, *12* (1), 153–176.

Debortoli, Davide, Jordi Galí, and Luca Gambetti, "On the Empirical (Ir)Relevance of the Zero Lower Bound Constraint," *NBER Macroeconomics Annual*, 2020, *34* (1), 141–170.

Eggertson, Gauti B and Michael Woodford, "The Zero Interest Bound and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, 2003, *Act 1*, 139–233.

Engel, Charles, "Currency Misalignments and Optimal Monetary Policy: A Reexamination," *American Economic Review*, 2011, *101* (6), 2796–2822.

Finck, David, "Forward Guidance Under the Cost Channel," *MAGKS Joint Discussion Paper Series in Economics No. 04-2020*, 2020.

Friedman, Milton, "The Case for Flexible Exchange Rates," *Essays in Positive Economics*, 1953, *157* (203), 33.

Fry, Renee and Adrian Pagan, "Sign Restrictions in Structural Vector Autoregressions: A Critical Review," *Journal of Economic Literature*, 2011, *49* (4), 938–960.

Gonçalves, Sílvia, Ana María Herrera, Lutz Kilian, and Elena Pesavento, "When Do State-Dependent Local Projections Work?," *CEPR Discussion Paper No. DP17265*, 2022.

Hoffmann, Mathias and Patrick Hürtgen, "Do Exchange Rates Absorb Demand Shocks at the ZLB?," *Deutsche Bundesbank Discussion Paper*, 2021, *No. 13/2021.*

Lenza, Michele and Giorgo E Primiceri, "How to Estimate a VAR after March 2020," *National Bureau of Economic Research*, 2020, *No. w27771.*

Müller, Gernot J, Martin Wolf, and Thomas Hettig, "Exchange Rate Undershooting: Evidence and Theory," *CEPR Discussion Paper No. DP13597*, 2022.

Mundell, Robert A, "A Theory of Optimum Currency Areas," *American Economic Review*, 1961, *51* (4), 657–665.

Nakata, Taisuke and Takeki Sunakawa, "Credible Forward Guidance," *IMES Discussion Paper*, 2022, *2022* (E-6).

Obstfeld, Maurice and Kenneth Rogoff, "New Directions for Stochastic Open Economy Models," *Journal of International Economics*, 2000, *50* (1), 117–153.

__ **and** __ , "Global Implications of Self-Oriented National Monetary Rules," *The Quarterly Journal of Economics*, 2002, *117* (2), 503–535.

Plagborg-Møller, Mikkel and Christian K Wolf, "Local Projections and VARs Estimate the Same Impulse Responses," *Econometrica*, 2021, *89* (2), 955–980.

Rubio-Ramirez, Juan F, Daniel F Waggoner, and Tao Zha, "Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference," *The Review of Economic Studies*, 2010, 77 (2), 665–696.

Smaghi, Lorenzo B, "Conventional and Unconventional Monetary Policy," *Speech at the Center for Monetary and Banking Studies, Geneva*, 2009, *April 28.*

Stavrakeva, Vania and Jenny Tang, "Deviations from FIRE and Exchange Rates: A GE Theory of Supply and Demand," Technical Report, London Business School and Federal Reserve Bank Boston 2020.

Svensson, Lars E O, "The Zero Bound in an Open Economy: A Foolproof Way of Escaping from a Liquidity Trap," *Monetary and Economic Studies*, 2001, *19* (S1), 277–312.

Tenreyro, Silvana and Gregory Thwaites, "Pushing on a String: US Monetary Policy Is Less Powerful in Recessions," *American Economic Journal: Macroeconomics*, 2016, *8* (4), 43–74.

Uhlig, Harald, "What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure," *Journal of Monetary Economics*, 2005, *52* (2), 381–419.

Walsh, Carl E, "Simple Sustainable Forward Guidance at the ELB," *unpublished, University of California, Santa Cruz*, 2018.

Wu, Jing Cynthia and Fan Dora Xia, "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound," *Journal of Money, Credit and Banking*, 2016, *48* (2-3), 253–291.

Figure 11: Consolidated balance sheet of the Eurosystem

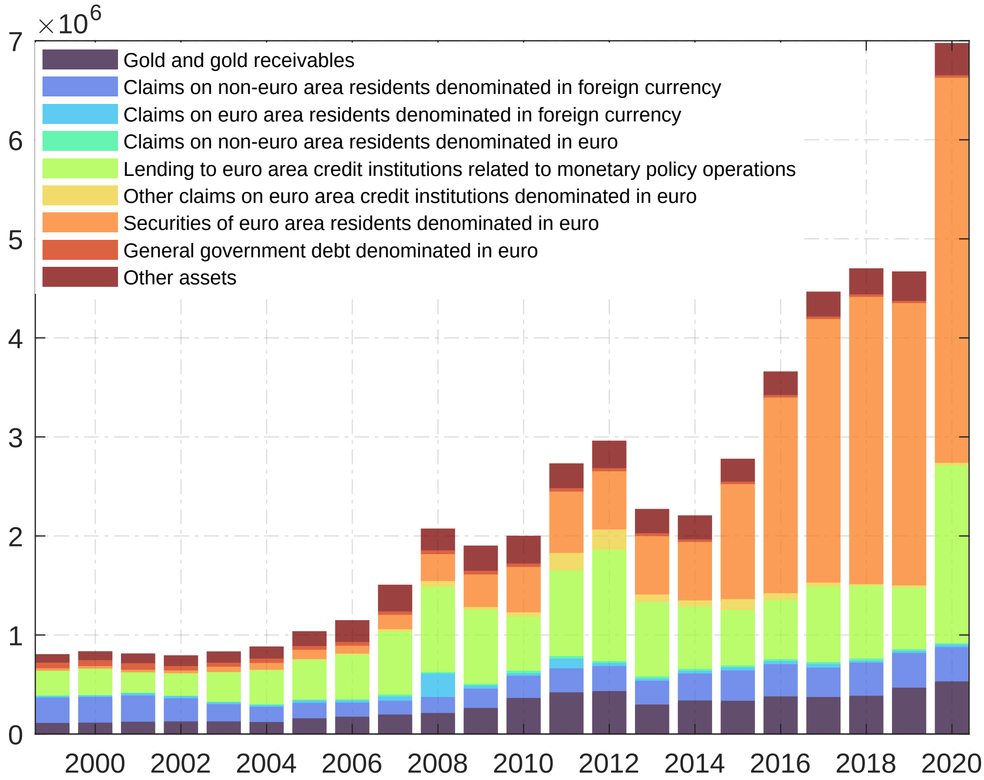

Notes: In million euros.

Figure 12: Different lag structure | euro area vis-à-vis Canada

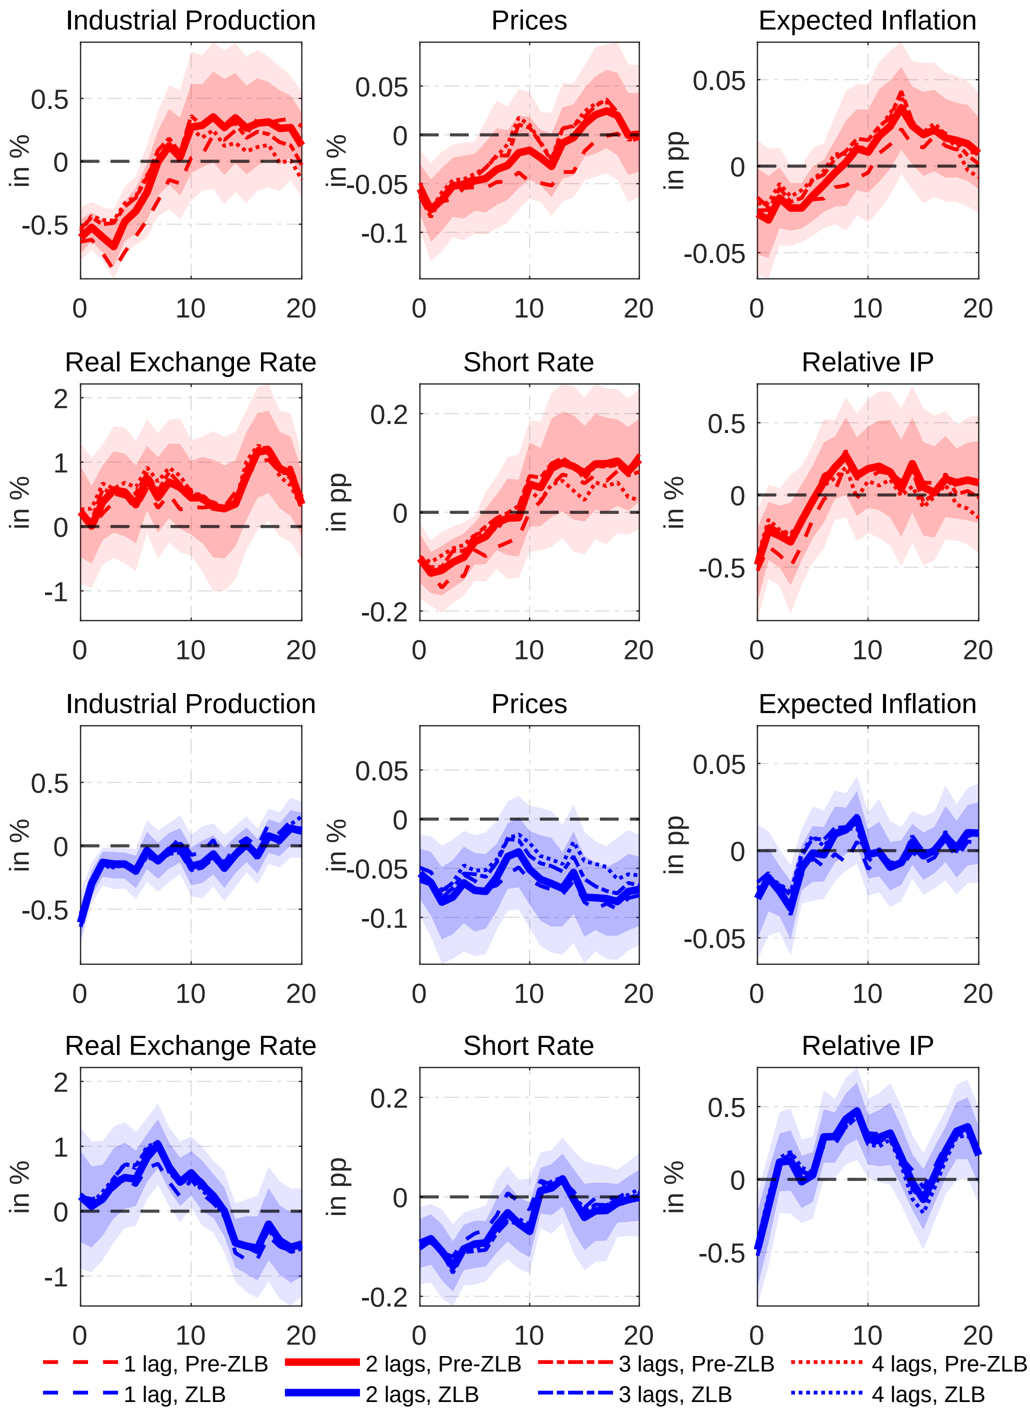

Notes: Median impulse response coefficients for different lag specifications. For both states, the shaded areas correspond to the 5*th* (95*th*) and the 16*th* (84*th*) percentiles of the baseline specification with $p = 2$ lags.

Figure 13: Different lag structure | euro area vis-à-vis the US

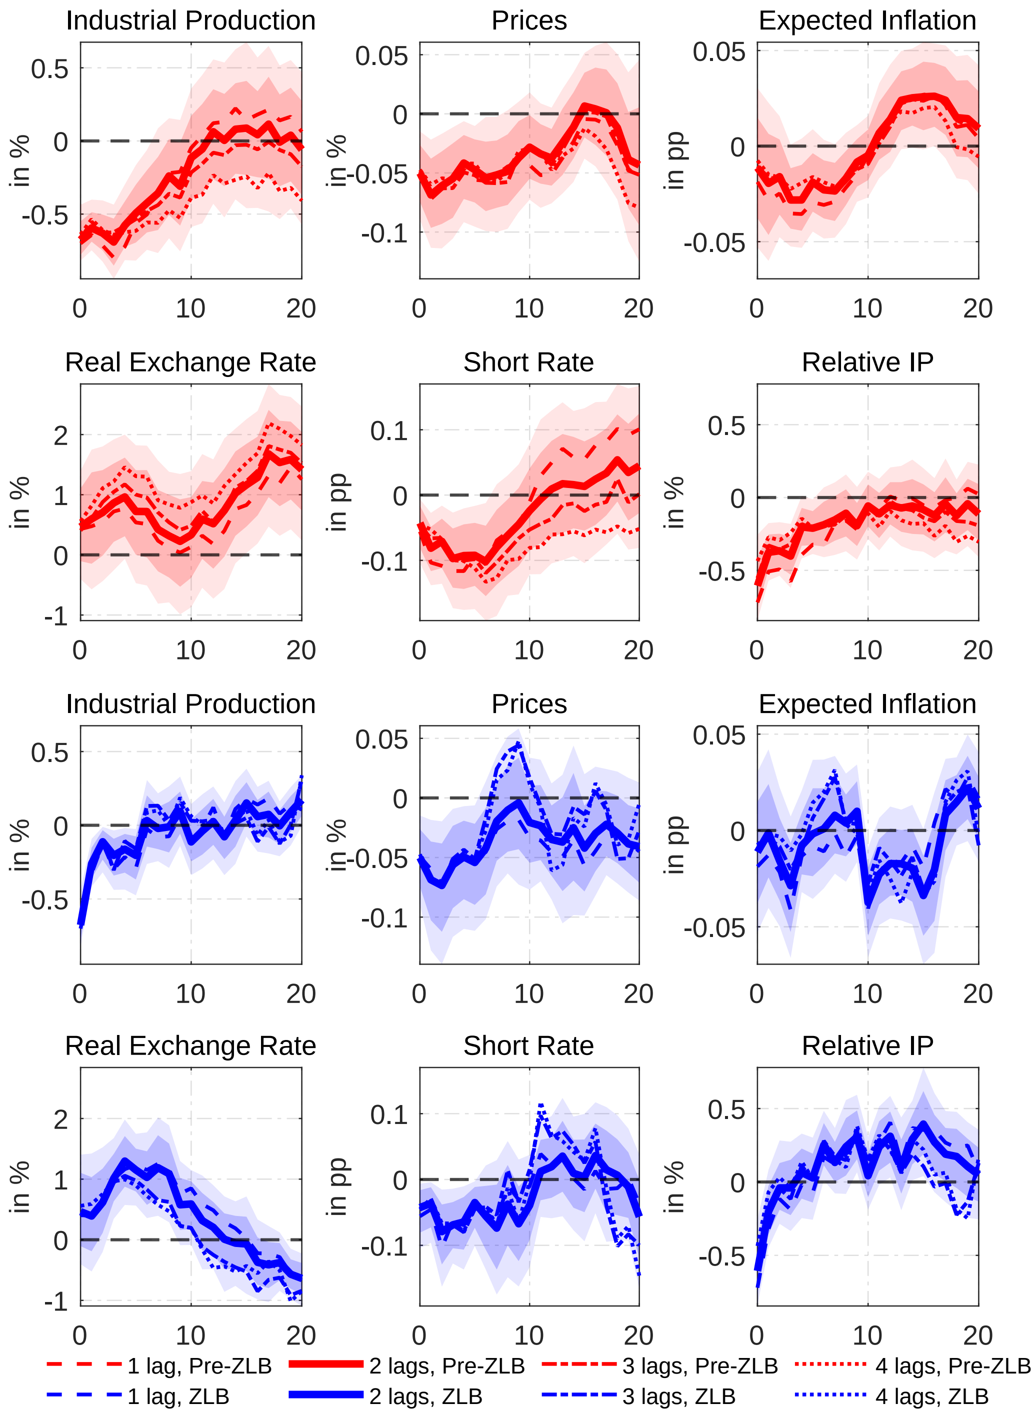

Notes: Median impulse response coefficients for different lag specifications. For both states, the shaded areas correspond to the 5*th* (95*th*) and the 16*th* (84*th*) percentiles of the baseline specification with $p = 2$ lags.

Figure 14: Different lag structure | euro area vis-à-vis Japan

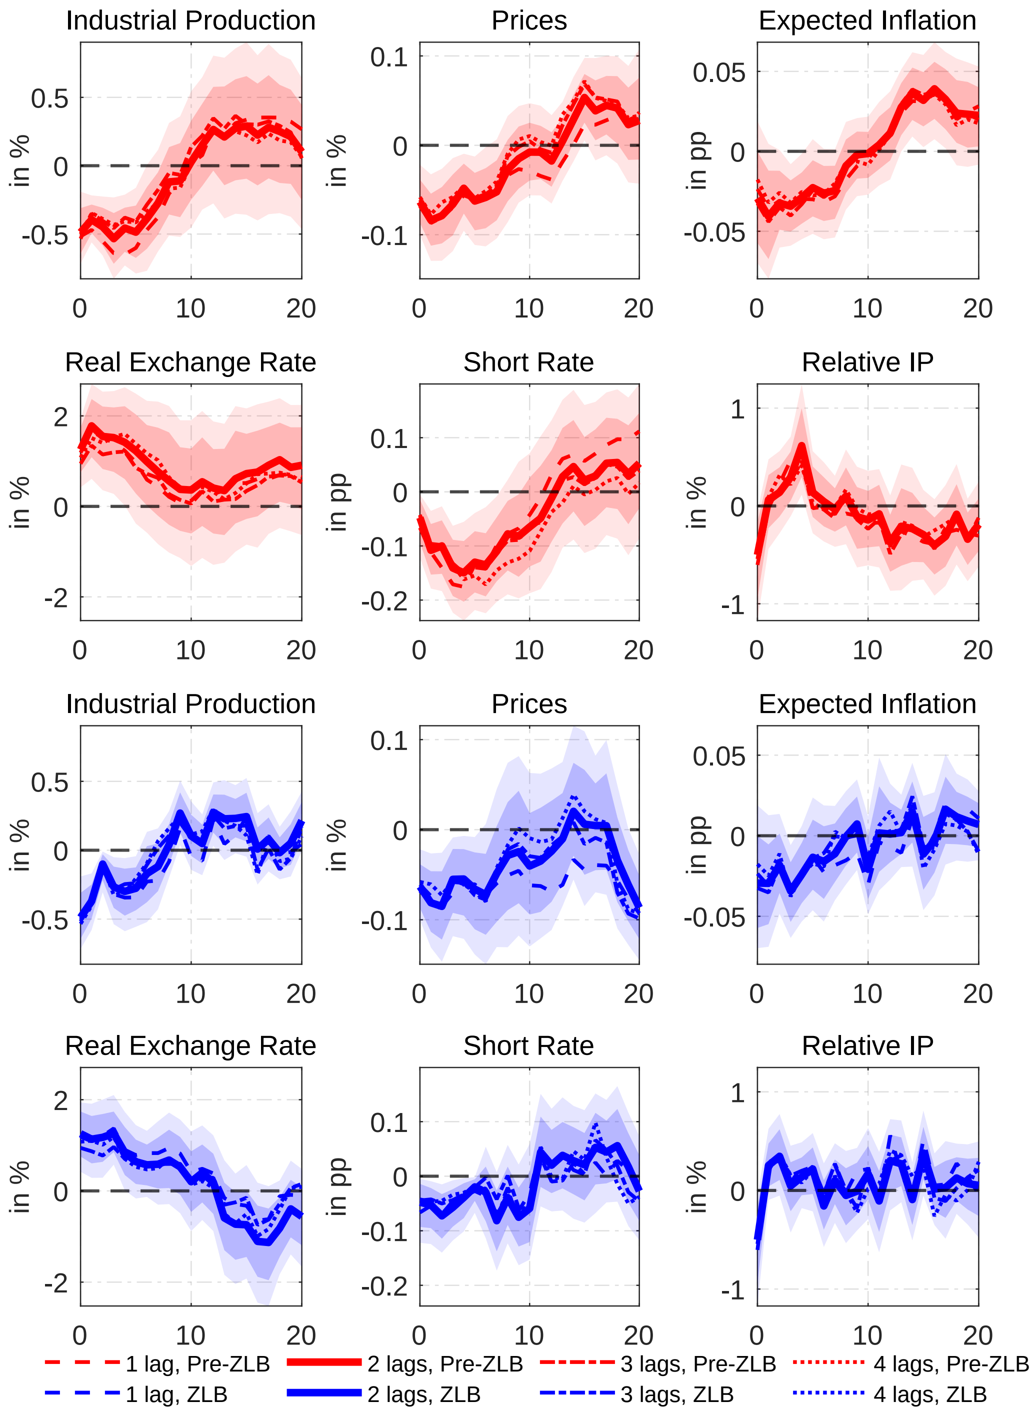

Notes: Median impulse response coefficients for different lag specifications. For both states, the shaded areas correspond to the 5*th* (95*th*) and the 16*th* (84*th*) percentiles of the baseline specification with $p = 2$ lags.

Figure 15: Smooth transitions | euro area vis-à-vis Canada

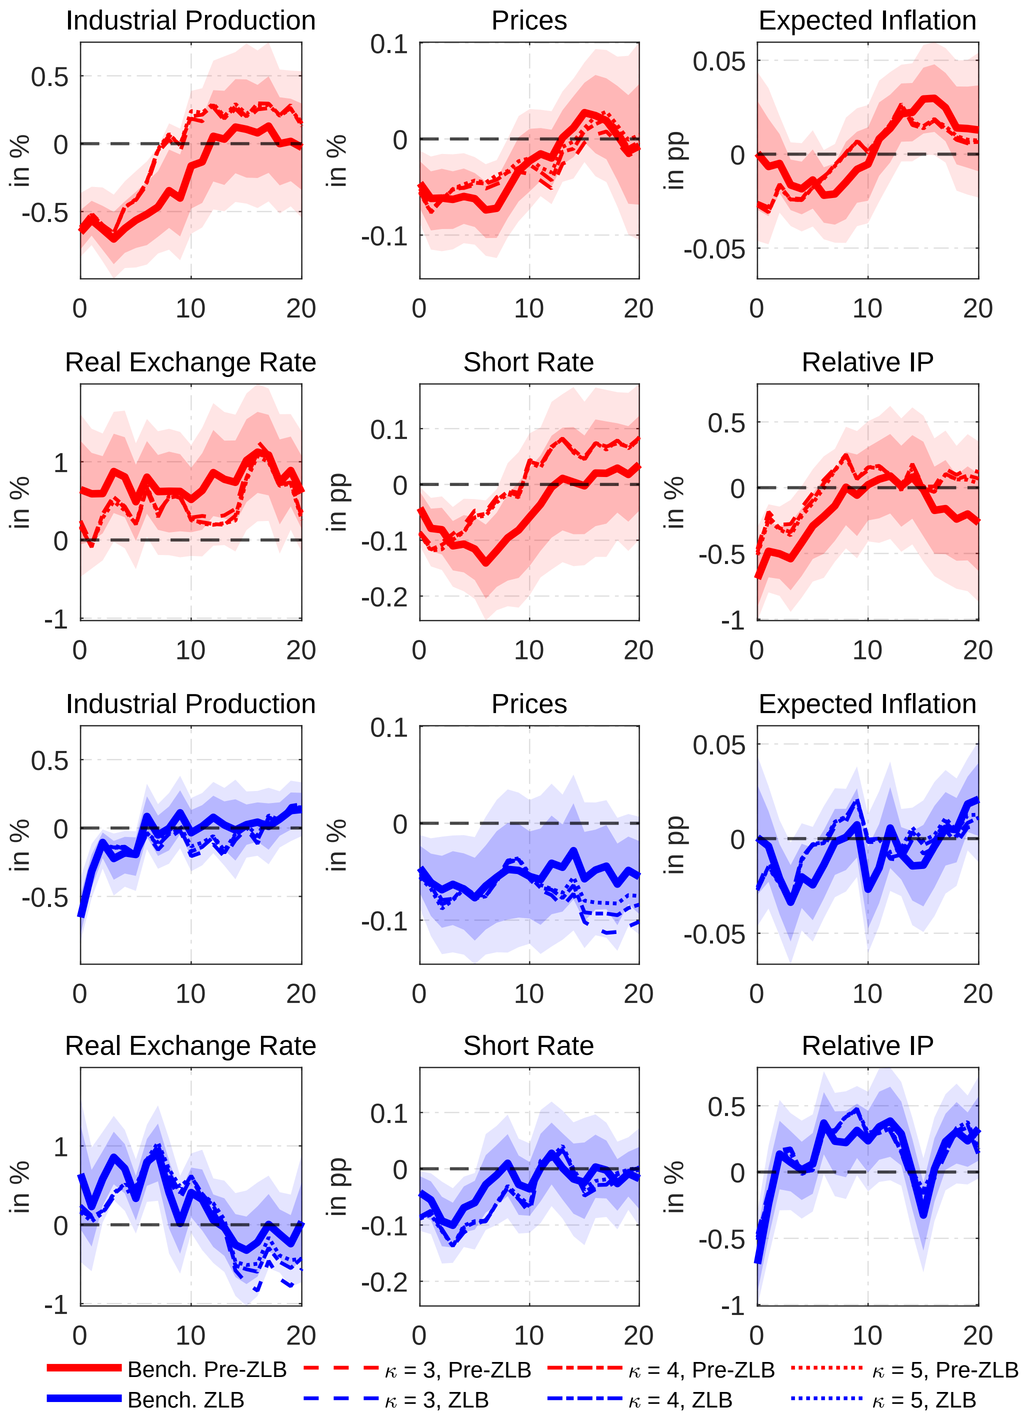

Notes: Median impulse response coefficients for different values for κ. For both states, the shaded areas correspond to the 5*th* (95*th*) and the 16*th* (84*th*) percentiles of the baseline specification.

Figure 16: Smooth transitions | euro area vis-à-vis the US

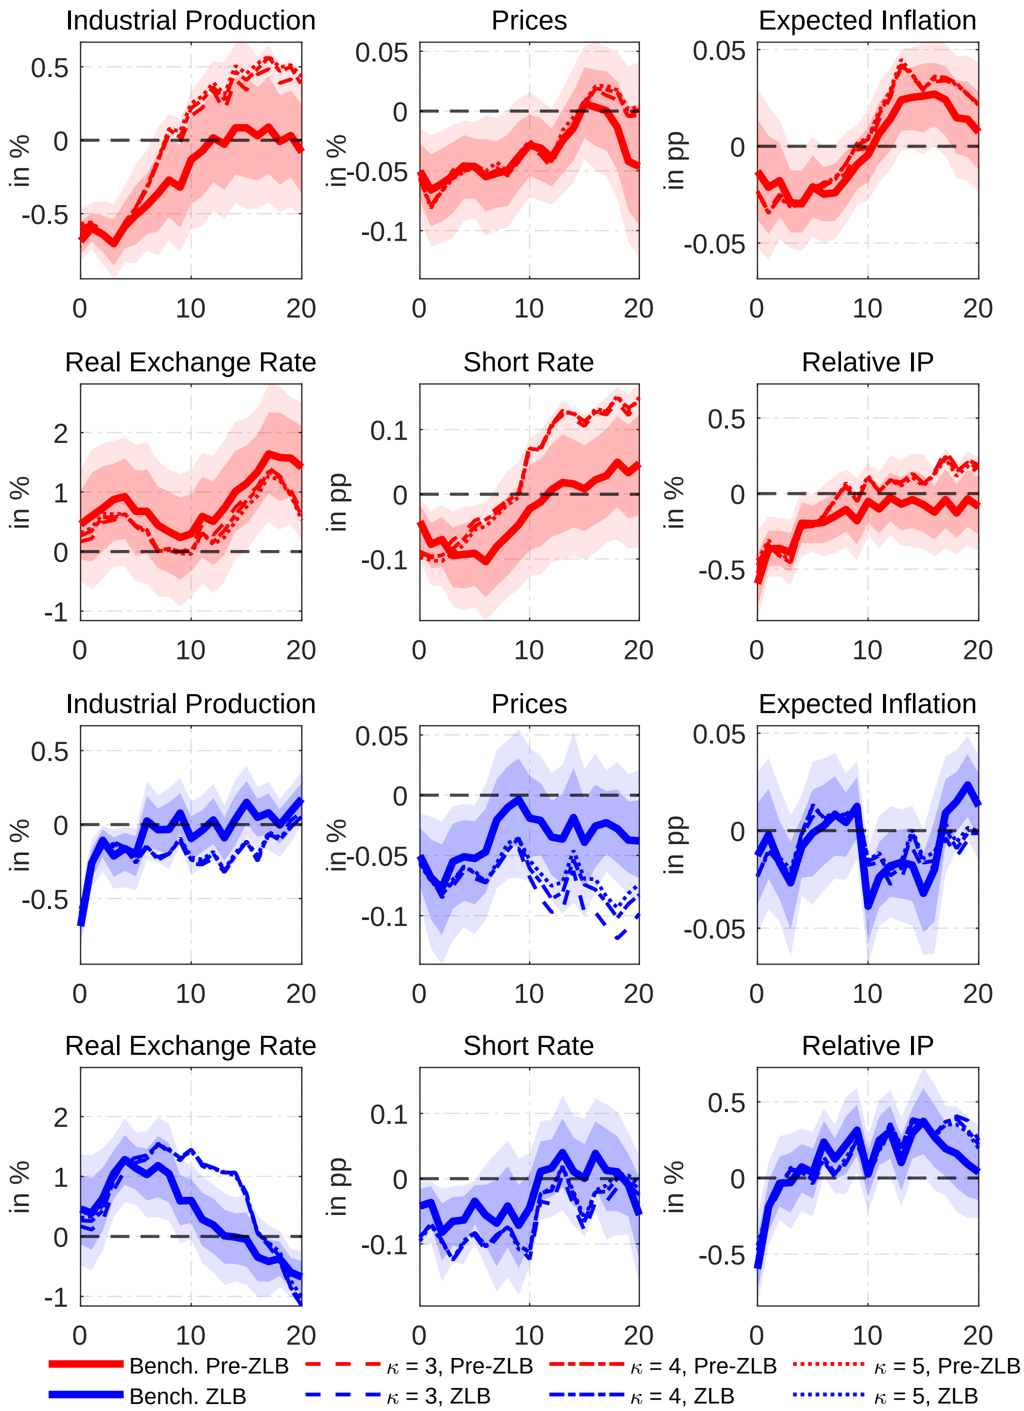

Notes: Median impulse response coefficients for different values for κ. For both states, the shaded areas correspond to the 5*th* (95*th*) and the 16*th* (84*th*) percentiles of the baseline specification.

Figure 17: Smooth transitions | euro area vis-à-vis Japan

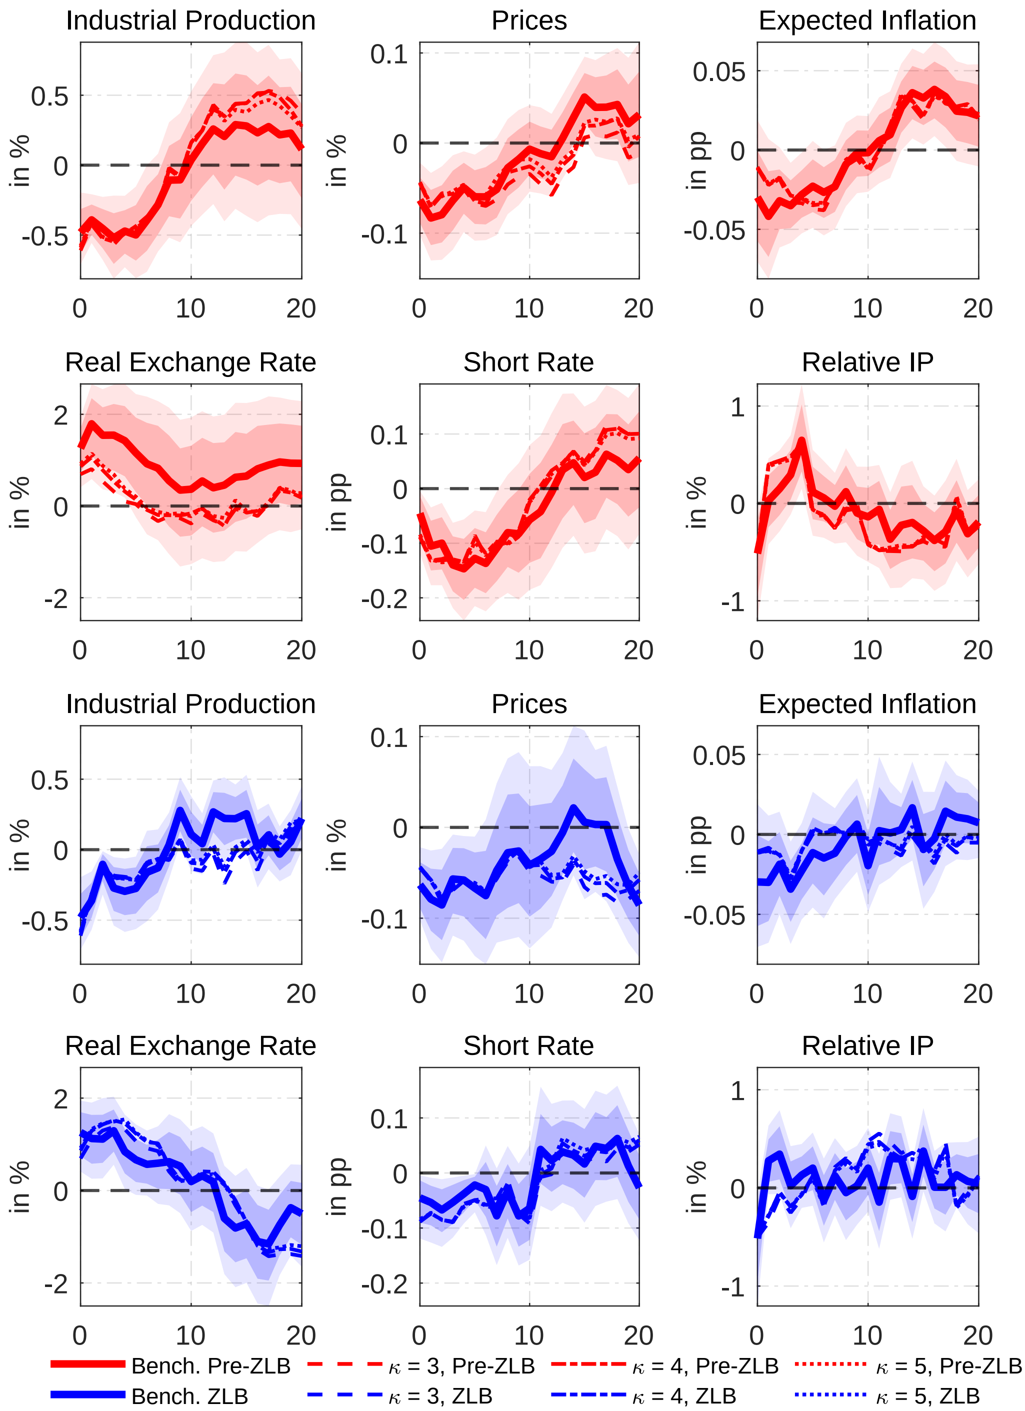

Notes: Median impulse response coefficients for different values for κ. For both states, the shaded areas correspond to the 5*th* (95*th*) and the 16*th* (84*th*) percentiles of the baseline specification.

Figure 18: Results from a state-dependent SVAR for the euro area vis-à-vis Canada

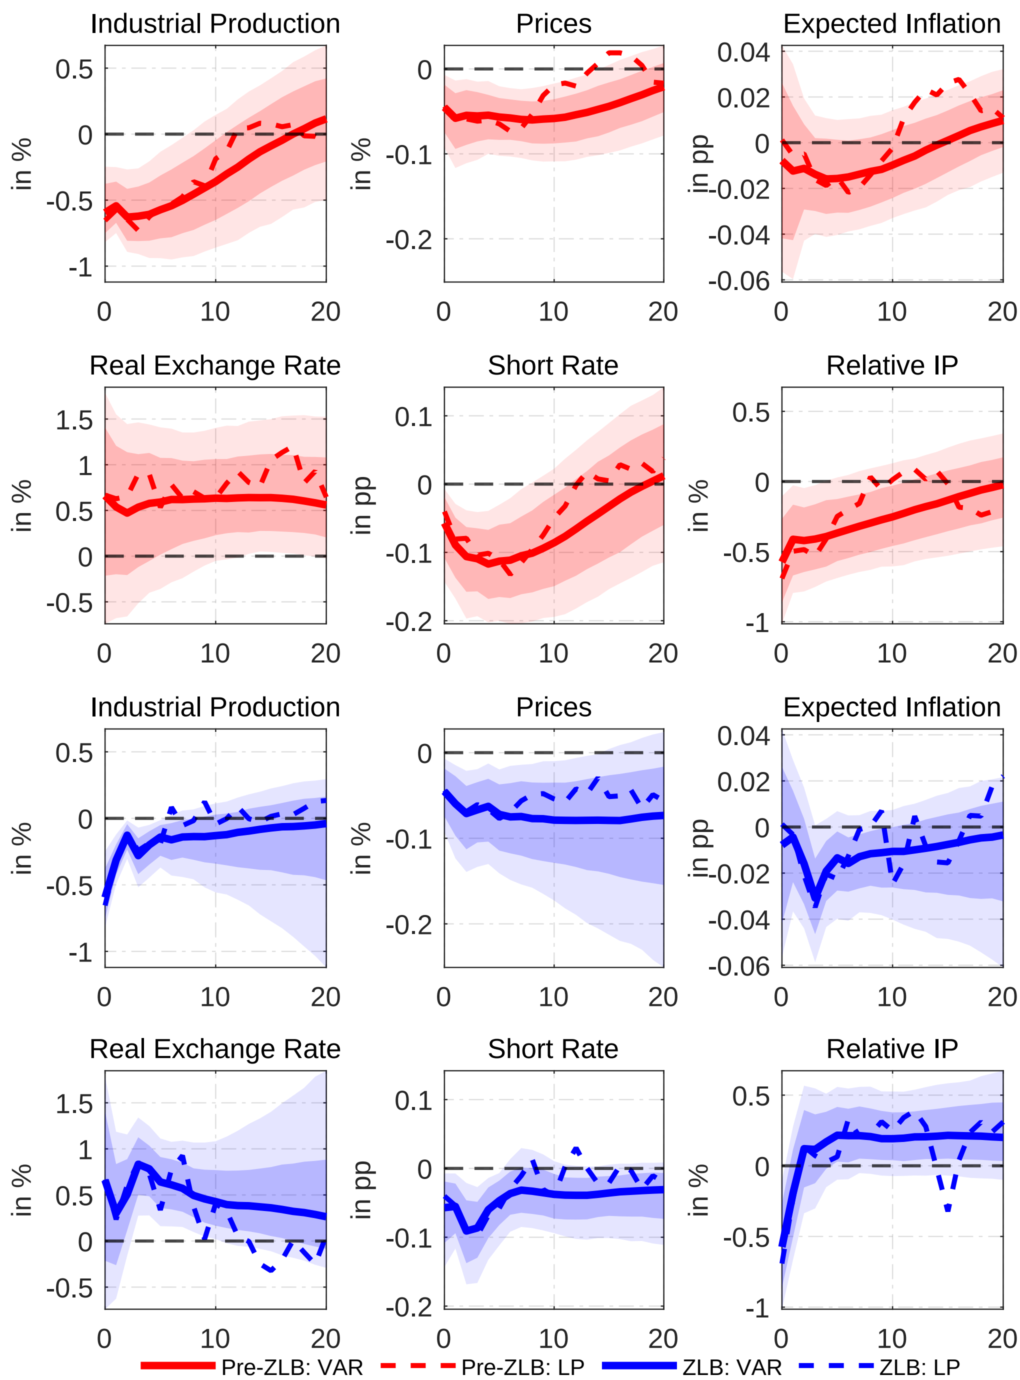

Notes: Median impulse response for the pre-ZLB period (red line) and the ZLB period (blue line). For both states, the shaded areas correspond to the 5*th* (95*th*) and the 16*th* (84*th*) percentiles.

Figure 19: Results from a state-dependent SVAR for the euro area vis-à-vis Japan

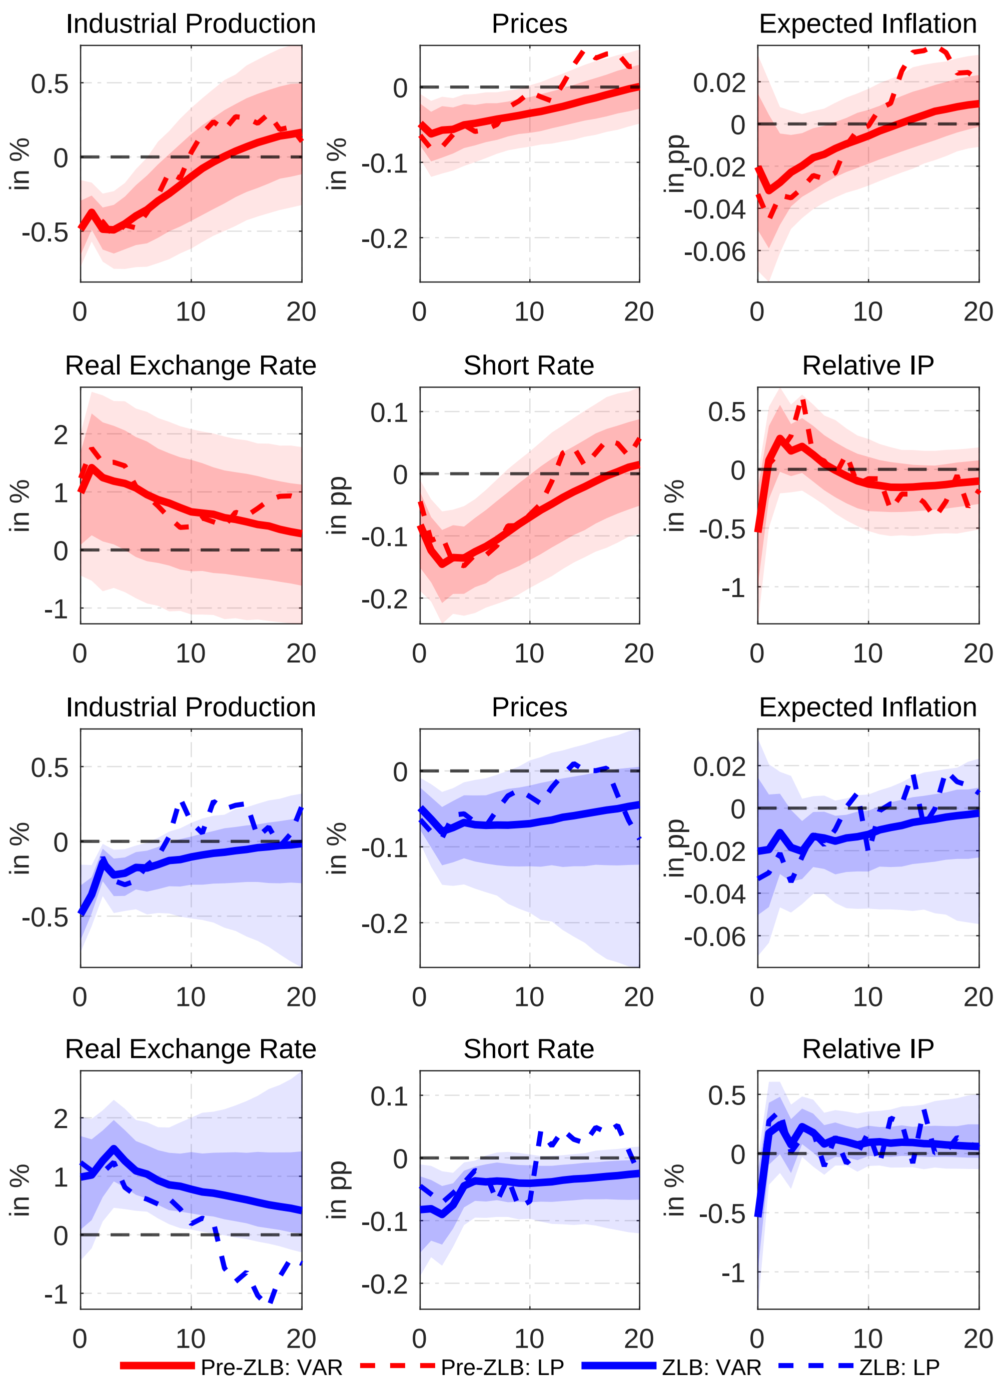

Notes: Median impulse response for the pre-ZLB period (red line) and the ZLB period (blue line). For both states, the shaded areas correspond to the 5*th* (95*th*) and the 16*th* (84*th*) percentiles.

Figure 20: Results from a state-dependent SVAR for the euro area vis-à-vis the US

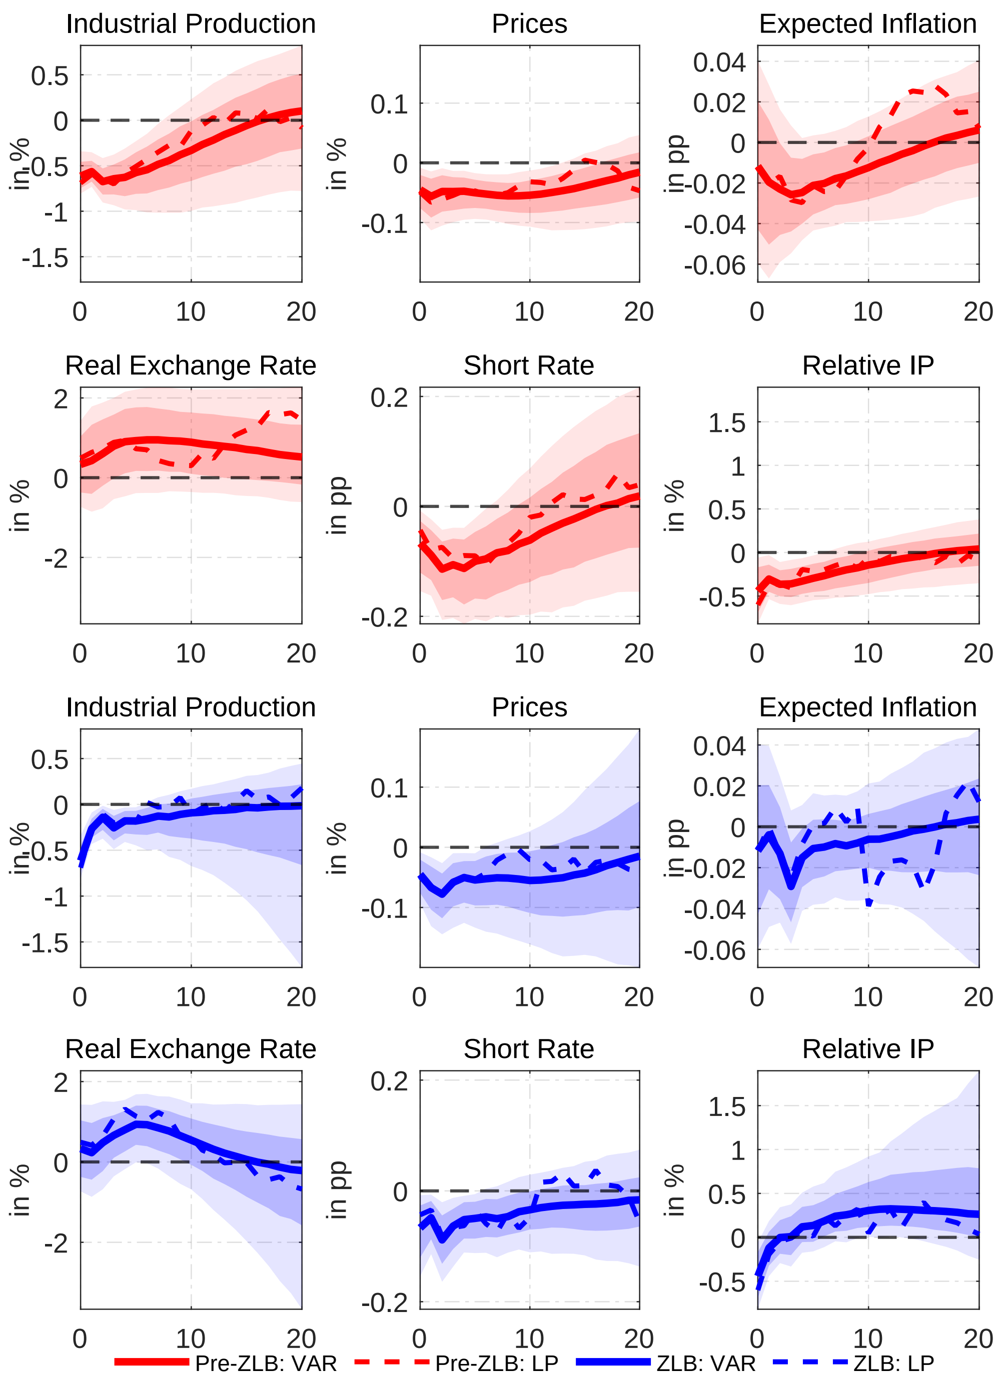

Notes: Median impulse response for the pre-ZLB period (red line) and the ZLB period (blue line). For both states, the shaded areas correspond to the 5*th* (95*th*) and the 16*th* (84*th*) percentiles.

Affidavit

Ich erkläre hiermit, dass ich die vorgelegten und nachfolgend aufgelisteten Aufsätze selbstständig und nur mit den Hilfen angefertigt habe, die im jeweiligen Aufsatz angegeben oder zusätzlich in der nachfolgenden Liste aufgeführt sind. In der Zusammenarbeit mit den angeführten Koautoren war ich mindestens anteilig beteiligt. Bei den von mir durchgeführten und in den Aufsätzen erwähnten Untersuchungen habe ich die Grundsätze guter wissenschaftlicher Praxis, wie sie in der Satzung der Justus-Liebig-Universität Gießen zur Sicherung guter wissenschaftlicher Praxis niedergelegt sind, eingehalten.

David Finck
24. November 2022

I hereby declare that I completed the papers submitted and listed hereafter independently and only with those forms of support mentioned in the relevant paper. When working with the authors listed, I contributed no less than a proportionate share of the work. In the analyses that I have conducted and to which I refer in the papers, I have followed the principles of good academic practice, as stated in the Statute of Justus Liebig University Giessen for ensuring good scientific practice.

David Finck
November 24, 2022

Submitted Papers

1. **Finck, David**, "Has Monetary Policy Really Become Less Effective in the Euro Area? A Note", *Applied Economics Letters*, 2019, *26*(13), 1087–1091.

2. **Finck, David, Jörg Schmidt, and Peter Tillmann**, "Mortgage Debt and Time-Varying Monetary Policy Transmission," *Macroeconomic Dynamics*, *forthcoming*, 2021.

3. **Finck, David and Paul Rudel**, "Do Credit Supply Shocks Have Asymmetric Effects?," *Empirical Economics*, *forthcoming*, 2022.

4. **Finck, David**, "Optimal Monetary Policy Under Heterogeneous Beliefs", *MAGKS Joint Discussion Paper Series in Economics*, 2022, No. 43-2022.

5. **Finck, David**, "Forward Guidance Under the Cost Channel," *MAGKS Joint Discussion Paper Series in Economics*, 2020, No. 04–2020.

6. **Finck, David and Peter Tillmann**, "Pandemic Shocks and Household Spending," *Oxford Bulletin of Economics and Statistics*, 2022, *84*(2), 273–299.

7. **Finck, David and Peter Tillmann**, "The Role of Global and Domestic Shocks for Inflation Dynamics: Evidence from Asia," *Oxford Bulletin of Economics and Statistics*, 2022, *84*(5), 1181–1205.

8. **Finck, David and Peter Tillmann**, "The Macroeconomic Effects of Global Supply Chain Disruptions," *unpublished, University of Giessen*, 2022.

9. **Finck, David, Mathias Hoffmann, and Patrick Hürtgen**, "On the Empirical Relevance of the Exchange Rate as a Shock Absorber at the Zero Lower Bound," *MAGKS Joint Discussion Paper Series in Economics*, 2022, No. 34–2022.