

Pandemic Shocks and Household Spending*

DAVID FINCK and PETER TILLMANN

Justus Liebig University Gießen, Gießen, Germany (e-mails: david.finck@wirtschaft.uni-giessen.de; e-mail: peter.tillmann@wirtschaft.uni-giessen.de)

Abstract

We study the response of daily household spending to the surprise number of fatalities of the COVID-19 pandemic, which we label as a 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.

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 a 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

JEL Classification numbers: E21, E32, I10.

*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.

market and smooth consumption, such that consumption exhibits very little fluctuations.¹

We look at the unexpected number of fatalities of the pandemic and analyse 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 under-prediction 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 week. 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. Goolsbee and Syverson (2021) show that consumer behaviour during the pandemic is more driven by fear of infection than formal restrictions. Second, households might be barred from consumption due to a lockdown of selected activities 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.³

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

¹See Jappelli and Pistaferri (2010) for a survey of the field.

²We are grateful to an anonymous referee for this point.

³In a survey conducted early in the pandemic, Binder (2020) finds that households expect an increase in unemployment due to the pandemic.

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, Müller and Pfeifer (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.* (2020), or the feedback between the spread of the pandemic and macroeconomic aggregates, for example, Eichenbaum, Rebelo and Trabandt (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.

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, for example, workers in the service sector.⁴ 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.⁵

⁴Even 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.

⁵See Mongey, Weinberg and Pilossoph (2020) for an analysis of the effect of social distancing across workers. They find significant differences in the burden from social distancing.

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 the 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 behaviour based on innovative datasets. In an early paper, Baker *et al.* (2020) used 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. Surico, Känzig and Hacıoglu (2020) use data from a fintech company based in the UK to track the behaviour of spending. These authors also document the build-up of financial stress as well as consumption and income inequality across households. Carvalho *et al.* (2020) use six billion transactions of customers of Spain's second-largest bank to track consumption over the crisis. Coibion, Gorodnichenko and Weber (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 lockdown restrictions. They find that lockdown restrictions explain most of the fall in consumer spending since March 2020.⁶

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 for alternative state variables. Section VI illustrates potential transmission channels and section VII concludes. An online Appendix S1 contains additional results.

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.

⁶Alexander and Karger (2020) analyse consumer spending and cellphone records in the US and show the causal effect of stay-at-home orders on spending.

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>.⁷

We also have spending broken down into ZIP codes with high, middle and low median income. Below, we will refer 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 its corresponding value from 2019. Lastly, the seasonally adjusted data are indexed to its pre-pandemic level, that is, the mean of the 7-day moving average from January 8–28. Hence, our series are given in percent, such that a value of 2% in t corresponds to an increase of spending by 2% relative to its average value in January.

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, that is, 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 paid 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

⁷Chetty *et al.* (2020) collect the data on spending from Affinity Solutions Inc. This company aggregates information on credit and debit card spending. Data are 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 3) is highlighted by the red vertical lines.

financial services might be underrepresented. We should keep these limitations in mind when discussing the results.

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).⁸ 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

⁸Details on the forecasting model, including assumptions on the model parameters, are available at <https://covid19-projections.com/model-details>.

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.⁹

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 week. Formally, we regress

$$\mathbf{e}_t = F(z_{t-1})(\boldsymbol{\gamma}'\mathbf{Q}_t + \boldsymbol{\phi}'\mathbf{e}_{t-1}) + (1 - F(z_{t-1}))(\boldsymbol{\gamma}''\mathbf{Q}_t + \boldsymbol{\phi}''\mathbf{e}_{t-1}) + \boldsymbol{\epsilon}_t, \tag{2}$$

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 $\boldsymbol{\epsilon}_t$ 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.¹⁰

Table 1 reports some descriptive statistics for both the raw shock \mathbf{e}_t and the purified shock $\boldsymbol{\epsilon}_t$. 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.¹¹ 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.

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

⁹We cannot rule out that some households have different beliefs about how the virus will progress and might even mistrust the official data about new cases and deaths, e.g. Fetzner *et al.* (2020).

¹⁰We 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.

¹¹We also applied various tests to detect heteroskedasticity in our series, including Engle’s ARCH test. The test results imply that our shock series is homoscedastic.

TABLE 1

Descriptive statistics for shocks

Min	Max	Mean	Median	5th	95th	Q-stat.	P-val.
Raw shock							
-1,007	2,416	41.3	78	-559.8	688.4	313.1	0.000
Purified shock							
-653.8	1,993.2	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.




Figure 2. Raw vs. purified shock

Notes: The left panel shows the raw shock calculated as $e_t = d_t - d_{t-1}$ 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.

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 2,000 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.

Setup

The linear model of departure reads

$$y_{t+h} = \alpha_h + \beta_h \epsilon_t + \delta_h t + \gamma_h \mathbf{x}_t + \varphi_h \mathbf{D}_t + u_{t+h}, \tag{3}$$

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 β_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 γ_h contains the effects of the lagged endogenous variable and other control variables at horizon h captured in the vector \mathbf{x}_t and φ_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 equation (3) includes the stringency index provided by researchers from the University of Oxford as well as two dummy variables to account for (i) the subsequent phase of the CARES act, i.e. the Paycheck Protection Program and Healthcare Enhancement Act, and (ii) the four FOMC meetings since April.¹² 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.¹³

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

$$y_{t+h} = F(z_{t-1})(\alpha_h^I + \beta_h^I \epsilon_t + \gamma_h^I \mathbf{x}_t) + \dots \tag{4}$$

$$\dots + (1 - F(z_{t-1}))(\alpha_h^{II} + \beta_h^{II} \epsilon_t + \gamma_h^{II} \mathbf{x}_t) + \delta_h t + \varphi_h \mathbf{D}_t + u_{t+h}$$

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 new infections. Note that we condition the response of our endogenous

¹²The stringency index is meant to measure the strictness of policies restricting people’s behaviour 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.

¹³The results are insensitive to using the Akaike Information Criterion instead.

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

$$\left. \frac{\partial y_{t+h}}{\partial \epsilon_t} \right|_{z_{t-1}} = F(z_{t-1}) \beta_h^I + (1 - F(z_{t-1})) \beta_h^{II}. \quad (5)$$

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.

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)

$$F(z_t) = \frac{1}{T} \sum_{j=1}^T 1_{z_j < z_t}, \quad (6)$$

where T is the sample size and $1_{z_j < z_t} = 1$ if $z_j < z_t$ and zero otherwise. That is, $1_{z_j < z_t}$ 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.¹⁴

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.

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

¹⁴In 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 if we use the non-smoothed growth rate.




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.

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.

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 $h = 0, \dots, H$ 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 .

Finally, we follow standard practice (see Jordà, 2005) and set the maximum autocorrelation lag for the Newey–West procedure to $L = h + 1$.¹⁵

¹⁵Note 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 $h = 0, \dots, H$, 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.

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.

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 1% 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% confidence bands based on Driscoll–Kraay standard errors. For the purpose of comparison, we also report the




Figure 4. Response of total spending

Notes: The first column shows the impulse response coefficients (red-solid) β_h^I for $h = 0, \dots, H$ 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% 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 .

corresponding coefficients from the linear model (red-dashed line). Accordingly, the second column reports the corresponding values for regime II, that is, 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 $h = 0, \dots, H$, where the shaded area covers the t -critical values for a 90% confidence interval, that is ± 1.645 .

In this context, it is important to stress that a perfectly symmetric transmission of pandemic shocks would imply that $\beta_h^I = \beta_h^{II} \forall h = 0, \dots, H$. 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% on impact relative to its average value in 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, that is, 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 in Figures 10–12, 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%, 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%, 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 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, Calvet and Comon, 2003; Carroll, 2009; Carroll *et al.*, 2017; Gelman, 2021). 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.¹⁶

¹⁶Explanations include households' wealth or employment status and the accompanying heterogeneity with respect to liquidity constraints.

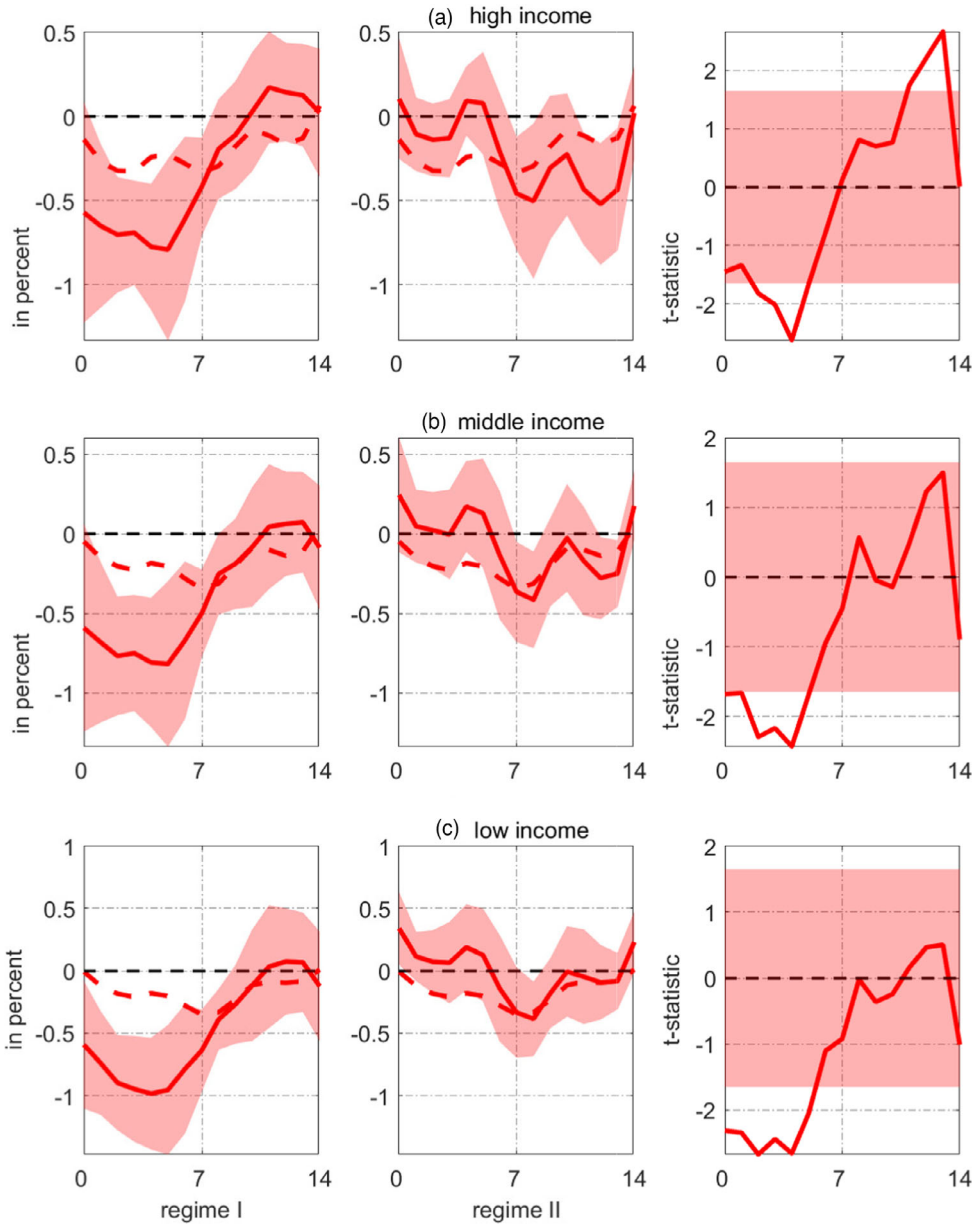


Figure 5. Response of spending by income quartile

Notes: The first column shows the impulse response coefficients β_h^I for $h = 0, \dots, H$ 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% 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 6. Response of relative spending (High Income – Low Income)

Notes: Difference of estimated coefficients $\beta_h^{\text{high},I} - \beta_h^{\text{low},I}$ in regime I and $\beta_h^{\text{high},II} - \beta_h^{\text{low},II}$ in regime II. The shaded areas cover the 5th and 95th 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 replication, for each $h = 0, \dots, H$ we then estimate the impulse response coefficients and calculate the sign of $\beta_h^{\text{high},j} - \beta_h^{\text{low},j}$ in regime $j = I, II$.¹⁷ We then use the distribution of our bootstrap and report the 5th and 95th 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, that is, 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.

¹⁷One 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.

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

$$y_{t+h} = F(z_{t-1})(\alpha_h^I + \gamma_h^I \mathbf{x}_t) + (1 - F(z_{t-1}))(\alpha_h^{II} + \gamma_h^{II} \mathbf{x}_t) + \delta t + \boldsymbol{\varphi}_h \mathbf{D}_t + u_{t+h}. \tag{7}$$

In a second step, we take the estimated forecast errors \hat{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

$$\begin{aligned} \hat{u}_{t+h} = & F(z_{t-1})(\omega_0^I \epsilon_t + \dots + \omega_h^I \epsilon_{t+h}) + \dots \\ & \dots + (1 - F(z_{t-1}))(\omega_0^{II} \epsilon_t + \dots + \omega_h^{II} \epsilon_{t+h}) + \eta_{t+h}, \end{aligned} \tag{8}$$

where ω_j^i for $j = 0, \dots, h$ and $i = I, II$ 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^2 s from a block bootstrap procedure including a bias-correction step as recommended by Gorodnichenko and Lee (2020).¹⁸

Remember that we do not have a VAR-based benchmark for our local projection-based 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^2 s 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 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

¹⁸To do so, we generate $B = 2,000$ 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.




Figure 7. FEVD by income quartiles

Notes: Explained share of forecast error variance after the bias-correction procedure (red dots) and the 5th and 95th percentiles of the distribution of the block bootstrap procedure.

above 30% after nearly two weeks, that spending for low-income households is most responsive to our pandemic shock. Interestingly, the 5th 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 are 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.¹⁹

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

¹⁹It 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.




Figure 8. Response on the state level

Notes: The first column shows the impulse response coefficients β_h^I for $h = 0, \dots, H$ 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 .

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.²⁰

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 behaviour 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

²⁰While there are no error bands shown in Figure 8, the online Appendix S1 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.

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

$$F(z_t) = \frac{\exp\left(\kappa \frac{z_t - \mu}{\sigma_z}\right)}{1 + \exp\left(\kappa \frac{z_t - \mu}{\sigma_z}\right)}, \quad (9)$$

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 behaviour of switching from one state to the other, this approach is well understood and relies on the idea of Granger and Teräsvirta (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 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.

The online Appendix S1 shows 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

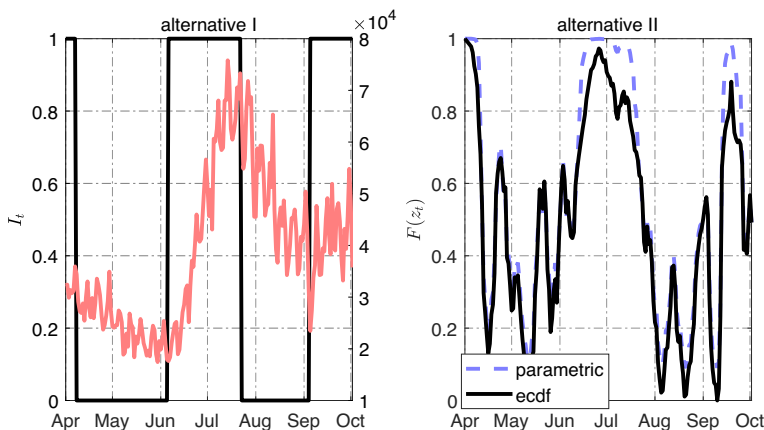


Figure 9. Alternative states

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.

like the results in the previous section. That is, total spending significantly decreases in regime I, that is, 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%. 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.

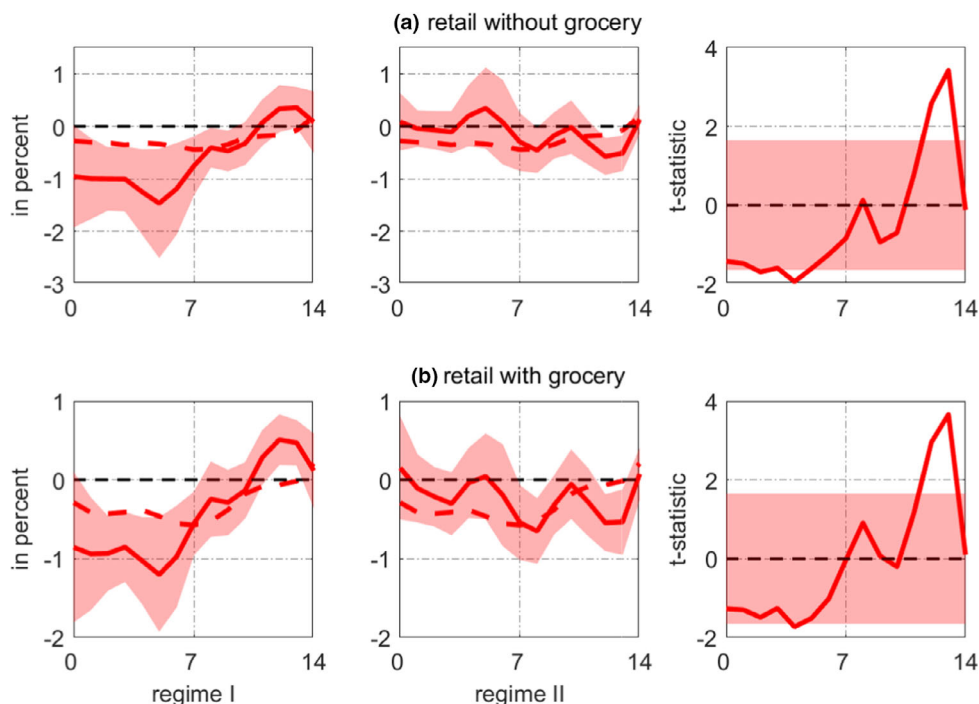


Figure 10. Response of spending: Retail

Notes: The first column shows the impulse response coefficients (red-solid) β_h^I for $h = 0, \dots, H$ 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% 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 .

Spending subcategories

The Chetty *et al.* (2020) data set allows us to dis-aggregate spending into subcategories. Figures 10–12 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 seems insensitive to the pandemic shock.

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 13 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




Figure 11. Response of spending: Subcategories I

Notes: The first column shows the impulse response coefficients (red-solid) β_h^I for $h = 0, \dots, H$ 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% 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 .

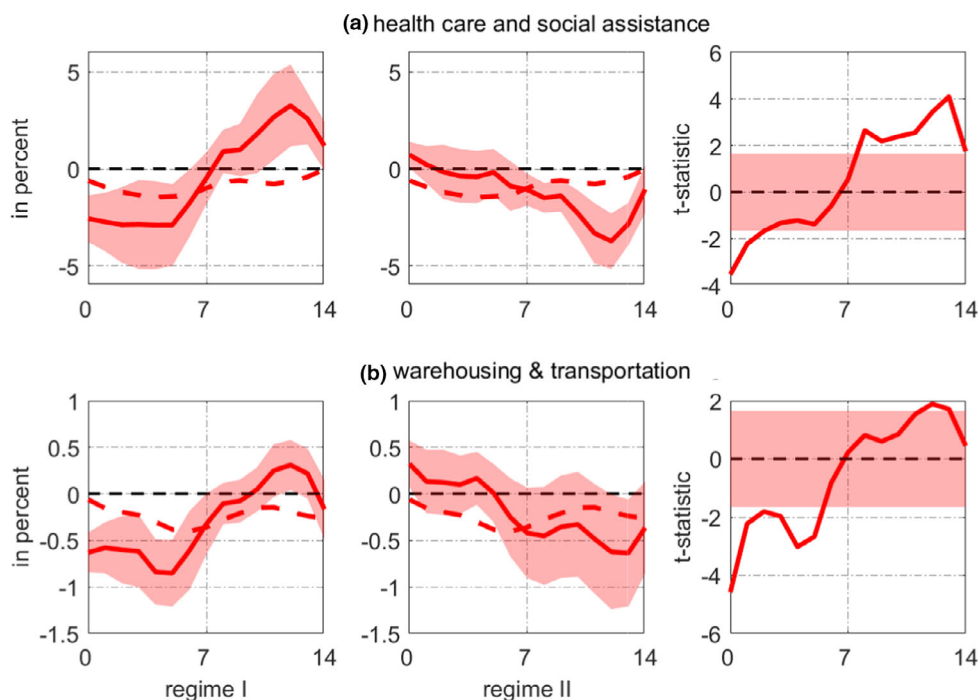


Figure 12. Response of spending: Subcategories II

Notes: The first column shows the impulse response coefficients (red-solid) β_h^I for $h = 0, \dots, H$ 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% 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 .

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.²¹ 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 14 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 consistent

²¹See <https://www.clevelandfed.org/our-research/indicators-and-data/consumers-and-covid-19.aspx>.




Figure 13. 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% 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 .

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.

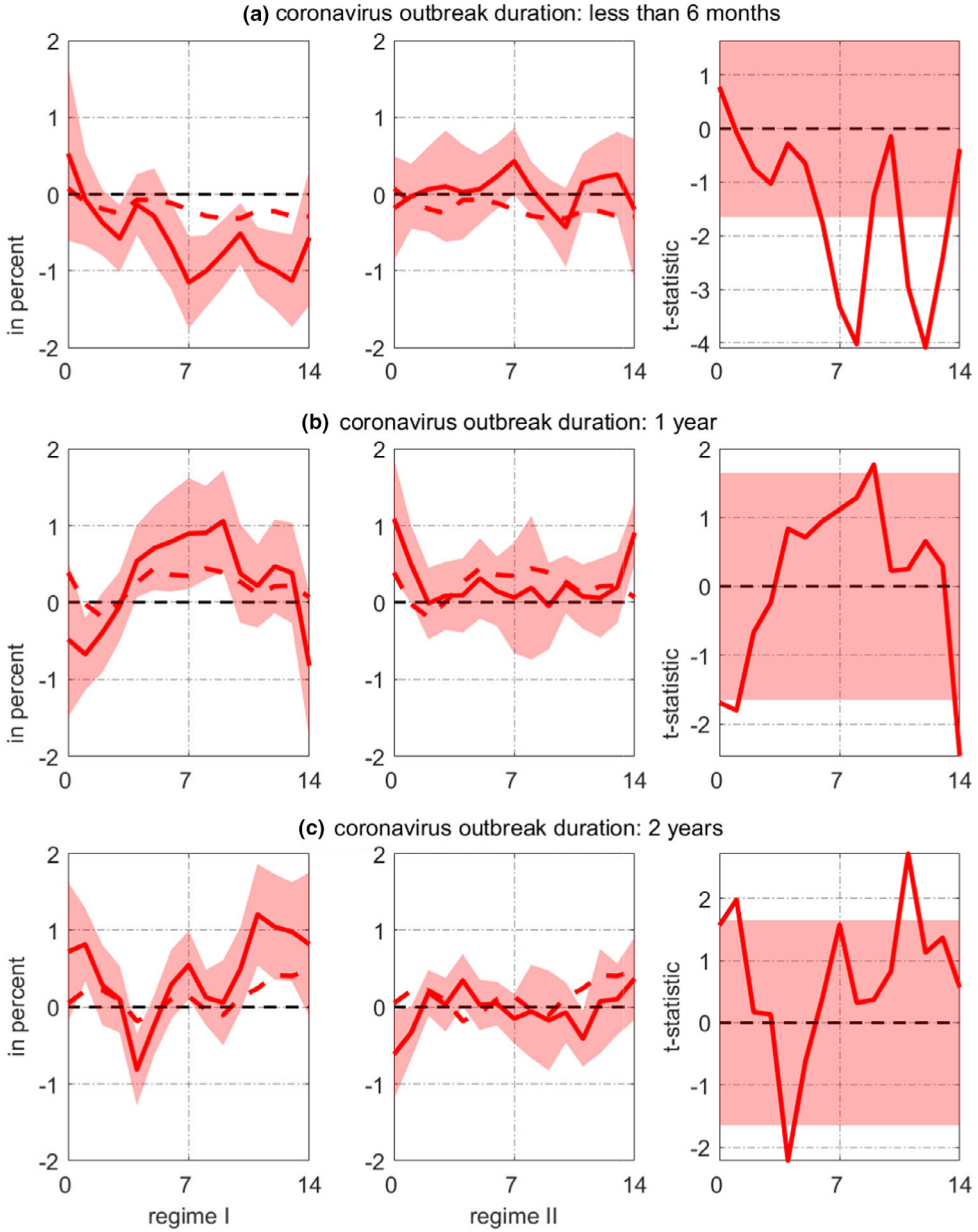


Figure 14. Response of expected duration

Notes: The first column shows the impulse response coefficients (red-solid) β_h^I for $h = 0, \dots, H$ 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% 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 .

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.

Acknowledgment

Open access funding enabled and organized by Projekt DEAL.

References

- Alexander, D. and Karger, E. (2020). Do Stay-at-Home Orders Cause People to Stay at Home? Effects of Stay-at-Home Orders on Consumer Behavior, Working Paper No. 2020-12, Federal Reserve Bank of Chicago.
- Andersen, A. L., Hansen, E. T., Johannesen, N. and Sheridan, A. (2020). Consumer Responses to the COVID-19 Crisis: Evidence From Bank Account Transaction Data, unpublished, University of Copenhagen.
- Auerbach, A. and Gorodnichenko, Y. (2012). 'Fiscal multipliers in recession and expansion', in Alesina A. and Giavazzi F. (eds), *Fiscal Policy After the Financial Crisis*, Chicago: University of Chicago Press, pp. 63–98.
- Baker, S. R., Farrokhnia, R. A., Meyer, S., Pagel, M. and Yannelis, C. (2020). 'How does household spending respond to an epidemic? Consumption during the 2020 COVID-19 pandemic', *Review of Asset Pricing Studies*, Vol. 10, pp. 834–862.
- Binder, C. (2020). 'Coronavirus fears and macroeconomic expectations', *The Review of Economics and Statistics*, Vol. 102, pp. 721–730.
- Born, B., Müller, G. and Pfeifer, J. (2020). 'Does austerity pay off?', *The Review of Economics and Statistics*, Vol. 102, pp. 323–338.
- Calvet, L. and Comon, E. (2003). 'Behavioral heterogeneity and the income effect', *The Review of Economics and Statistics*, Vol. 85, pp. 653–669.
- Carroll, C. (2009). 'Precautionary saving and the marginal propensity to consume out of permanent income', *Journal of Monetary Economics*, Vol. 56, pp. 780–790.
- Carroll, C., Slacalek, J., Tokunaka, K. and White, M. (2017). 'The distribution of wealth and the marginal propensity to consume', *Quantitative Economics*, Vol. 8, pp. 977–1020.
- Carvalho, J. R., Garcia, V. M., Hansen, S., Ortiz, A., Rodrigo, T., Roriguez Mora, J. V. and Ruiz, P. (2020). Tracking the COVID-19 Crisis with High-Resolution Transaction Data, unpublished, University of Cambridge.
- Chetty, R., Friedman, J., Hendren, N., Stepner, M. and the Opportunity Insights Team. (2020). 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 No. 27431, National Bureau of Economic Research.
- Coibion, O., Gorodnichenko, Y. and Weber, M. (2020). The Cost of the COVID-19 Crisis: Lockdowns, Macroeconomic Expectations and Consumer Spending, unpublished, University of Chicago.
- Cox, N., Ganong, P., Noel, P., Vavra, J., Wong, A., Farrell, D. and Greig, F. (2020). Initial Impacts of the Pandemic on Consumer Behavior: Evidence from Linked Income, Spending, and Savings Data, unpublished, University of Chicago.

- Dietrich, A. M., Kuester, K., Müller, G. J. and Schoenle, R. S. (2020). News and Uncertainty about COVID-19: Survey Evidence and Short-Run Economic Impact, Working Paper No. 20-12, Federal Reserve Bank of Cleveland.
- Driscoll, J. and Kraay, A. (1998). 'Consistent covariance matrix estimation with spatially dependent panel data', *The Review of Economics and Statistics*, Vol. 80, pp. 549–560.
- Eichenbaum, M. S., Rebelo, S. and Trabandt, M. (2021). 'The macroeconomics of pandemics', *The Review of Financial Studies*, Northwestern University, Vol. 34, pp. 5149–5187.
- Fetzer, T., Hensel, L., Hermle, J. and Roth, C. (2020). 'Coronavirus perceptions and economic anxiety', *Review of Economics and Statistics*, forthcoming, pp. 1–36.
- Gelman, M. (2021). 'What drives heterogeneity in the marginal propensity to consume? Temporary shocks vs. persistent characteristics', *Journal of Monetary Economics*, Vol. 117, pp. 521–542.
- Goolsbee, A. and Syverson, C. (2021). 'Fear, lockdown, and diversion: comparing drivers of pandemic economic decline', *Journal of Public Economics*, Vol. 193, pp. 104311.
- Grodnichenko, Y. and Lee, B. (2020). 'Forecast error variance decompositions with local projections', *Journal of Business & Economic Statistics*, Vol. 38, pp. 1–24.
- Granger, C. and Teräsvirta, T. (1993). *Modelling Non-Linear Economic Relationships*, Oxford University Press, Oxford.
- Gu, Y. (2020). COVID-19 Projections Using Machine Learning. <https://covid19-projections.com>. Retrieved 13 July 2020.
- Jappelli, T. and Pistaferri, L. (2010). 'The consumption response to income changes', *Annual Review of Economics*, Vol. 2, pp. 479–506.
- Jordà, O. (2005). 'Estimation and inference of impulse responses by local projections', *American Economic Review*, Vol. 95, pp. 161–182.
- Karger, E. and Rajan, A. (2020). Heterogeneity in the Marginal Propensity to Consume: Evidence from Covid-19 Stimulus Payments, Working Paper No. 2020-15, Federal Reserve Bank of Chicago.
- Mongey, S., Weinberg, A. and Pilossoph, L. (2020). Which Workers Bear the Burden of Social Distancing Policies? NBER Working Paper No. 27085, National Bureau of Economic Research.
- Ramey, V. and Zubairy, S. (2018). 'Government spending multipliers in good times and in bad: evidence from US historical data', *Journal of Political Economy*, Vol. 126, pp. 850–901.
- Surico, P., Känzig, D. and Hacıoglu, S. (2020). Consumption in the Time of COVID-19: Evidence From UK Transaction Data, CEPR Working Paper No. 14733, Centre for Economic Policy Research.
- Tenreiro, S. and Thwaites, G. (2016). 'Pushing on a string: US monetary policy is less powerful in recessions', *American Economic Journal: Macroeconomics*, Vol. 8, pp. 43–74.

Supporting Information

Additional Supporting Information may be found in the online version of this article:

Appendix S1. Appendix.