

JUSTUS LIEBIG UNIVERSITY GIESSEN

DISSERTATION

**ESSAYS ON TECHNOLOGICAL
CHANGE AND ITS ENVIRONMENTAL
IMPACTS**

Submitted in fulfilment of the requirements for the degree of

DOCTOR RERUM POLITICARUM (Dr. rer. pol.)

in the

Faculty of Economics and Business Studies

Chair of Economics of Digitalisation

by

Janna Axenbeck

on

April 19, 2023

Supervisors:

Prof. Dr. Irene Bertschek (first supervisor), Chair of Economics of Digitalisation,
Faculty of Economics and Business Studies, Justus Liebig University Giessen,

Prof. Dr. Peter Winker (second supervisor), Chair of Statistics and Econometrics,
Faculty of Economics and Business Studies, Justus Liebig University Giessen.

Acknowledgements

First and foremost, I would like to express my deepest gratitude to my supervisor Prof. Dr. Irene Bertschek. Words cannot express how extremely thankful I am for the freedom you gave me in my research, the opportunities you opened up for me, your open-mindedness, and your empathy. I am also very grateful for your knowledge, expertise, and guidance, which not only helped me to succeed in this project but also prepared me for future challenges. Furthermore, I would like to extend my sincere gratitude to my second supervisor Prof. Dr. Peter Winker. I greatly benefited from our discussions, your encouragement to question things critically, as well as your belief in my abilities. Moreover, I would like to express my deepest appreciation to the members of my dissertation committee for their time and valuable feedback. Also, I would like to thank the ZEW – Leibniz Centre for European Economic Research (ZEW Mannheim) and the Justus Liebig University Giessen very much for the profound education I received and all the opportunities that were given to me. I am confident that I will greatly benefit on my future journey from the knowledge and skills I developed.

In addition, I am very grateful for the amazing co-authors who accompanied me on my way. Anne, I do not know whether writing a paper without ever meeting each other in person would have worked with anybody else. At least, it would never have been so much fun. Thomas, thank you so much for your time, support, and patience, as well as for strengthening my belief in my own abilities. I have profited incredibly much from your experience and your sympathetic ear. Patrick, special thanks for your support, solving so many (methodological) challenges with me, and all the great discussions we had. Daniel, thank you for the spirit I needed to finish the last chapter. I would also like to thank my dear colleagues at the Digital Economy Department at ZEW Mannheim for making the time of writing this thesis so pleasant. I will always remember your support and kindness. Special thanks goes also to the Dig Docs, the Boulder Gang, and the Neckarstadt Girls.

Moreover, I want to deeply thank my family and friends, especially Gabi, Rüdiger, Robin, Malu, Imogen, Esther, Maike, and Valeska for their unconditional support and unshaken belief in me. Finally, I would like to thank Moritz, who has supported me in many ways during this time, especially for his encouragement, his interest in my work, and his strong shoulder I could lean on in difficult times.

Preface

Given the climate crisis and the sharp rise in energy costs, discussions on the potential of new technologies to reduce energy consumption and mitigate environmental damage are frequently part of the public debate. In the thesis at hand, which is submitted in partial fulfilment of the requirements for the academic degree of *doctor rerum politicarum*, I contribute to this discourse by shedding new light on the measurement of current technological progress and its impact on energy use patterns, as well as changes in mobility. To this end, I apply econometric, text mining, and machine learning methods to firm-level data. My dissertation consists of four self-contained essays, which I wrote between December 2018 and April 2023 as an external doctoral candidate at the Chair of Economics of Digitalisation at the Justus Liebig University Giessen. During this time, I was employed as a researcher at ZEW Mannheim. The thesis is divided into six chapters.

Chapter 1 serves as an introduction to the thesis. It presents the problems that are addressed and introduces fundamental definitions and concepts that are applied. The chapter ends with a (brief) outline of each essay.

Accurately observing technological change is a prerequisite for determining whether new technologies can help tackle current socioeconomic crises. In Chapter 2 (Essay 1), I leverage natural language processing techniques to analyse the potential of firm websites to detect innovation efforts (in co-authorship with Patrick Breithaupt). We find that firm websites contain useful information that can be linked to traditional, survey-based indicators of firm-level innovation activities. As web-based information can be quickly updated, is available at a very granular regional level, and is less expensive than information from questionnaire-based surveys, we conclude that web-based innovation indicators are a useful complement to traditional data sources. My relative contribution to the chapter is 50%.

In the other three essays of this thesis, I scrutinise the potential of current technological change to enhance environmental sustainability. Here, I focus on information and communications technologies (ICT), which have been fundamentally reshaping our socioeconomic system for several decades. I empirically examine how the ongoing diffusion of ICT affects environmental outcomes.

In Chapter 3 (Essay 2), I analyse the link between firm digitalisation and reductions in mobility during the Covid-19 pandemic. In particular, I examine whether a potential ICT-enabled decline in mobility persisted after the end of most Covid-19 restrictions in Germany, as this may reduce the environmental burden of transportation in the long term. My co-authors Irene Bertschek, Patrick Breithaupt, Daniel Erdsiek, and I show that the variation in the degree of firm digitalisation between German districts can be robustly linked to a greater decrease in mobility (measured via mobile network data) during Covid-19 waves. However, mobility reductions diminished after the severity of the pandemic declined in March 2022, suggesting no persistent environmental improvements. I contributed 50% of the chapter's total content.

In the following two chapters, I analyse the relationship between digital technologies and energy consumption. In both chapters, I use administrative panel data on manufacturing firms retrieved from German statistical offices.

In Chapter 4 (Essay 3), Thomas Niebel and I analyse the link between digital technologies and energy intensity improvements in manufacturing industries at the firm level. Results indicate no substantial energy intensity improvements in connection with ICT adoption. My relative contribution to the chapter is 60%.

In Chapter 5 (Essay 4), I investigate, in co-authorship with Anne Berner and Thomas Kneib, the heterogeneity of the link between ICT adoption and absolute energy consumption. To this end, we apply flexible tree-based machine learning. The results indicate that the relationship is heterogeneous but suggest that ICT adoption relates more frequently to an increase in absolute energy use than to a decline. I contributed 45% of the chapter's total content.

In Chapter 6, the findings from each essay are synthesised, leading to final conclusions and an outlook for future research.

Contents

Acknowledgements	i
Preface	iii
1 Introduction	1
1.1 Technological Change, Digitalisation, and Environmental Outcomes – New Opportunities for Measuring Relationships	1
1.2 Contribution	3
2 Innovation Indicators Based on Firm Websites – Which Website Charac- teristics Predict Firm-Level Innovation Activity?	8
2.1 Introduction	8
2.2 Literature Review	10
2.3 Data	12
2.3.1 Text-based Features	15
2.3.2 Meta Information Features	16
2.3.3 Network Features	17
2.4 Descriptive Analysis	17
2.5 Methodology	22
2.6 Results	26
2.7 Discussion	31
2.8 Conclusion	34
3 Firm Digitalisation and Mobility – Do Covid-19-Related Changes Per- sist?	36
3.1 Introduction	36
3.2 Related Literature	38
3.3 Data	41
3.3.1 Mobility	41
3.3.2 Firm Digitalisation	43
3.3.3 Control Variables	45
3.4 Descriptive Insights	45
3.5 Econometric Approach and Results	48

3.6	Robustness	53
3.7	Discussion & Conclusion	55
4	Digital Technology Adoption and Energy Intensity in Manufacturing – Firm-Level Insights	57
4.1	Introduction	57
4.2	Related Literature	60
4.3	Theoretical Frameworks	62
4.3.1	Translog Cost Function	63
4.3.2	CES Production Function	64
4.4	Data	65
4.4.1	Data Sources	65
4.4.2	Variable Description	66
4.4.3	Additional Descriptive Statistics	69
4.5	Econometric Analysis	71
4.5.1	Translog Cost Function	71
4.5.2	Reduced Form CES Production Function	74
4.5.3	Properties of the Link Between ICT and Energy Relating to Aggregation Issues	76
	Differences Across Industries	76
	Differences Between and Within Firms	80
4.6	Discussion	82
4.7	Conclusion and Policy Implications	84
5	What Drives the Relationship Between Digitalisation and Industrial Energy Demand? Exploring Firm-Level Heterogeneity	86
5.1	Introduction	86
5.2	Digitalisation and Energy Use in Manufacturing	89
5.3	Methodology	92
5.3.1	Measuring Heterogeneous Relationships	93
5.3.2	Generalised Random Forests	93
	Identifying Assumptions	95
	Orthogonalisation	97
	Panel Structure	98
5.4	Data	98
5.4.1	Microdata on the German Manufacturing Sector	98
5.4.2	Variable Description	99
5.5	Results	102
5.5.1	Conditional Average Treatment Effects	102
5.5.2	Analysing Effect Heterogeneity	105

5.5.3	Group Differences in the Light of Current Policies	108
5.6	Robustness	111
5.7	Summary and Conclusion	112
6	Concluding Remarks	114
	Bibliography	141
	Appendices	142
	Appendix A Innovation Indicators Based on Firm Websites – Which Website Characteristics Predict Firm-Level Innovation Activity?	143
A.1	Comparison of the Distribution Between the MIP and the Applied Subsample	143
A.2	List of Emerging Technology Terms Used in the Conducted Keyword Search	145
A.3	Detailed Information on the Calculation of Features	147
Text-Based Features	147
Meta Information Features	149
Network Features	150
A.4	Most Relevant Features for Each ‘All’ Feature Model	152
A.5	Learned Hyperparameters for Random Forest Models Using Different Feature Sets and Target Variables	154
A.6	AUC Values for Different Splits Between Training and Test Sample	155
	Appendix B Firm Digitalisation and Mobility – Do Covid-19-Related Changes Persist?	156
B.1	Variable Description	156
B.2	Additional Descriptive Statistics	161
B.3	Results With Respect to the Average Effect	164
B.4	Results of Robustness Checks	165
B.4.1	Parallel Trends	165
B.4.2	Digitalisation Observed in December 2022	166
B.4.3	Further Robustness Checks	168
B.4.4	Firm-Level Link	169
B.4.5	Link to a District’s WFH Potential	171
	Appendix C Digital Technology Adoption and Energy Intensity in Manufacturing	172
C.1	Additional Data	172
C.2	Perpetual Inventory Method (PIM)	175

C.3	Additional Descriptive Statistics	177
C.3.1	Number of Observations per Year	177
C.3.2	Distribution of Energy Prices	177
C.3.3	Details on the Distribution of S_E	177
C.3.4	Plausibility of Differences in Software Usage with Respect to Industry and Regional Characteristics	178
C.4	Calculation of Energy Cost Savings per Software Investment in the Year of Investment	180
C.5	Technological Illustration of the Heterogeneity Bias	181
C.6	Additional Econometric Estimations	183
C.6.1	Robustness Checks with Respect to a Potential Measurement Error	183
	Calculation of Software Capital Stocks	183
	“Zero” Software Capital Stocks	186
	Additional Robustness Checks	188
C.6.2	IV Estimates	190
C.6.3	Pooled OLS, FE, and Mundlak	192
C.6.4	Industry-Specific Effects – Regression Table	193
Appendix D What Drives the Relationship Between Digitalisation and In-		
dustrial Energy Demand?		194
D.1	Description of Variables	194
D.2	Descriptive Statistics	198
D.3	Evaluating Assumptions and the Fit of the Causal Forest	201
D.4	Variable Importance	206
D.5	Robustness Analysis	207
Appendix E Link Between ICT Adoption and Changes in Total Energy		
Consumption, Output, as well as Energy Intensity		210
Declaration of Authorship		212

List of Figures

2.1	Average occurrence of different emerging technology terms on firm websites with and without product innovations	20
2.2	Differences in the topic share of the top ten topics with the strongest average correlation with MIP-based innovation indicators	22
2.3	Feature importance values for ‘all’ feature models	30
3.1	Transfer learning approach for measuring firm digitalisation	44
3.2	Regional distribution in January 2020 and in December 2022 of the web-based firm digitalisation indicator	45
3.3	Average weekly change in mobility per district over the observed time frame.	47
3.4	Monthly change in mobility associated with digitalisation	49
4.1	Software capital intensity by firms’ PC usage	69
4.2	Software capital intensity by maximum data transmission rate	69
4.3	Industry-specific estimations by Equation (4.8)	78
4.4	Differences with respect to the average level of energy intensity estimated by Equation (4.8)	79
5.1	Illustration of Causal Forest partitioning	94
5.2	Distribution of the conditional average treatment effect (CATE) for the three different outcome variables: energy use, electricity use, and non-electric fossil fuel use	103
5.3	Bivariate distributions and smoothed regression lines for ICT-related changes in energy consumption and selected variables (total energy use)	105
5.4	Difference between energy prices with respect to the 20% of firms with the highest predicted ICT-related change in energy, electricity, or fossil fuel consumption and the 20% with the lowest predicted ICT-related change	107
5.5	Group average treatment effects (GATE) grouped by the economic strength of the corresponding region for different quintiles of labour (number of employees; L), tangible capital (K) and output (Q)	109

5.6	Group average treatment effects (GATE) grouped by energy-intensive (Divisions: 10-12, 17,19, 20, 23, 24) and remaining industries for different quintiles of labour (number of employees; L), tangible capital (K) and output (Q).	110
A.1	Firm distribution based on number of employees	143
A.2	Firm distribution for economic sectors based on two-digit NACE codes	144
A.3	AUC values for different splits between training and test sample . . .	155
B.1	Change in mobility compared to 2019 by distance in % (7-day average)	162
B.2	Comparison between quintiles	163
B.3	Analysis of parallel trends before the Covid-19 pandemic (estimated β^m coefficients)	165
B.4	Monthly change in mobility associated with digitalisation	166
C.1	Distribution of P_E	177
C.2	Distribution of P_E [external].	177
C.3	Average software capital intensity by industry between 2009 and 2017	178
C.4	Average software capital intensity by region between 2009 and 2017 . .	179
D.1	Distribution of propensity scores between treatment and control group	201
D.2	Inverse-propensity weighted histograms for treated and untreated observations (Part I)	203
D.3	Inverse-propensity weighted histograms for treated and untreated observations (Part II)	204
D.4	Inverse-propensity weighted histograms for treated and untreated observations (Part III)	205
D.5	Variable Importance for the three Causal Forests with the outcome variables total energy use, electricity use, and fossil fuels	206
D.6	Bivariate distributions of the predicted treatment effect and production factors	208
D.7	Comparison between OOB predictions and predictions of the test sample	209

List of Tables

1.1	Contribution table	7
2.1	Summary statistics for product innovators, process innovators, innovators, as well as firms with innovation expenditures	14
2.2	Features related to text, meta information, and network measures . . .	15
2.3	Descriptive statistics for selected variables	18
2.4	Content of the LDA topics with the strongest relationship to MIP-based innovation indicators	21
2.5	Results for Random Forest classification models using different feature sets and target variables	27
3.1	DiD results providing insights into changes in the link between mobility reductions and firm digitalisation for different phases of the pandemic	52
4.1	Summary statistics of selected variables	70
4.2	First-difference estimation results of Equation (4.8)	72
4.3	First-difference results of Equation (4.9)	75
4.4	Correlation between average energy intensity at the industry (two-digit NACE level) as well as at the firm level and the variance of the software capital growth rate, respectively	80
4.5	Comparison of software coefficients for Pooled OLS, FE, and Mundlak	82
5.1	Summary of the steps of the Causal Forest algorithm with orthogonalisation and honesty	97
5.2	Variable overview	99
5.3	Best Linear Predictor Test for the forest with total energy use as outcome	104
A.1	Most relevant features for product innovators	152
A.2	Most relevant features for process innovators	152
A.3	Most relevant features for innovators	153
A.4	Most relevant features for innovation expenditure	153
A.5	Learned hyperparameters for Random Forest models using different feature sets and target variables	154

B.1	Description of variables	156
B.2	Overview of descriptive statistics	161
B.3	Average decrease in mobility associated with digitalisation considering mobility changes over the entire day, daytime mobility changes, nighttime mobility changes, as well as differences between working days and weekends during the two pandemic years	164
B.4	DiD results providing insights into changes in the link between mobility reductions and firm digitalisation with respect to different phases of the pandemic using digitalisation observed in 2022	167
B.5	Further robustness checks	168
B.6	Link between firm digitalisation and WFH at the firm level	169
B.7	Link between firm digitalisation and increased e-commerce activity at the firm level	170
B.8	Equation B.1 with digitalisation replaced by a district's WFH potential	171
C.1	Description of additional data sources	172
C.2	Number of observations per year	177
C.3	Detailed descriptive statistics on the distribution of S_E	177
C.4	Equation (4.8) with software capital stocks modified by different depreciation rates	184
C.5	Equation (4.8) with software capital stocks modified by different lengths of periods considered for the initial capital stock calculation	185
C.6	Robustness checks with respect to "zero" software capital stocks	187
C.7	Further robustness checks	189
C.8	IV estimates	191
C.9	Pooled OLS, FE, and Mundlak (all model coefficients)	192
C.10	Differences across industries – Regression results	193
D.1	Description of variables	194
D.2	Averages and standard errors of firm characteristics for treated and untreated firms (2010 to 2017)	198
D.3	Best Linear Predictor Test for the forest with all outcomes	202
D.4	Robustness tests	207
E.1	Link between ICT adoption and changes in total energy consumption, output, and energy intensity	211

List of Abbreviations

AFiD	Amtliche Firmendaten für Deutschland
AI	Artificial intelligence
AIPW	Augmented inverse probability weighting
a.m.	Ante meridiem
ARGUS	Automated Robot for Generic Universal Scraping
ATE	Average treatment effect
AUC	Area under the curve
CART	Classification and regression trees
CATE	Conditional average treatment effect
CES	Constant elasticity of substitution
CIS	Community Innovation Survey
COD	Chemical oxygen demand
CO₂	Carbon dioxide
Covid-19	Coronavirus disease 2019
Destatis	Statistisches Bundesamt (transl. Federal Statistical Office)
DiD	Difference-in-differences
EEG	Erneuerbare-Energien-Gesetz (transl. Renewable Energy Sources Act)
EU	European Union
EU ETS	European Union Emissions Trading System
FE	Fixed effects
5G	Fifth generation of mobile communications
GATE	Group average treatment effect
GDP	Gross domestic product
GPT	General purpose technologies
GRF	Generalized Random Forest
Gt	Gigatons
GWh	Gigawatt hours
HHI	Herfindahl–Hirschman Index
HTML	Hypertext markup language
ICT	Information and communications technologies
ID	Identification
IV	Instrumental variable

KLE	Capital, labour, and energy inputs
KLEMS	Capital, labour, energy, materials, and service inputs
km	Kilometre
kWh	Kilowatt hours
LASSO	Least absolute shrinkage and selection operator
LDA	Latent Dirichlet Allocation
Mbit/s	Megabits per second
MDI	Mean decrease in impurity
MIP	Mannheim Innovation Panel
ms	Millisecond(s)
MUP	Mannheimer Unternehmenspanel (Mannheim Enterprise Panel)
NACE	Statistical classification of economic activities in the European community
OECD	Organisation for Economic Co-operation and Development
OLED	Organic light-emitting diode
OLS	Ordinary least squares
OOB	Out-of-bag
PATSTAT	Worldwide patent statistical database
PC	Personal computer
PIM	Perpetual inventory method
p.m.	Post meridiem
pp	Percentage points
R&D	Research and development
RDC	Research Data Centres of the Statistical Offices of the Federation and the Federal States
ROC	Receiver operating characteristic curve
SEO	Search engine optimisation
SME	Small and medium-sized enterprise
SO₂	Sulfur dioxide
SUTVA	Stable Unit Treatment Value Assumption
3D	Three-dimensional
TOBI	Text data based output indicators as base of a new innovation metric
UNFCCC	United Nations Framework Convention on Climate Change
URL	Uniform resource locator
US	United States
VPN	Virtual private network
WFH	Work(ing) from home
ZEW	ZEW – Leibniz Centre for European Economic Research

Chapter 1

Introduction

1.1 Technological Change, Digitalisation, and Environmental Outcomes – New Opportunities for Measuring Relationships

Technological change is considered key to economic growth (Solow 1956, 1957). In particular, it is believed that general purpose technologies (GPT), such as the steam engine or semiconductors, lead to large leaps in technological progress and productivity gains (Bresnahan & Trajtenberg 1995). To this day, digital technologies, which rely on the encoding of information into binary code (zeroes and ones), are regarded as one of the most significant GPT (Jovanovic & Rousseau 2005). In this thesis, the terms ‘information and communication technologies’ (ICT) and ‘digital technologies’ are used interchangeably and the term ‘digitalisation’ refers to the widespread deployment of these technologies. GPT have in common that they allow for pervasive implementation, have an inherent potential for technical improvements, and spark complementary innovations (Bresnahan & Trajtenberg 1995). As digital technologies possess these attributes, they can be associated with fundamental changes in the way we live and work. Examples here are the rise of online shopping and remote work enabled by communication technologies, as well as the growing adoption of automated decision-making processes, which are supported by artificial intelligence (AI).

Alongside the digital transformation, a shift towards a more environmentally sustainable society is currently (slowly) taking place to combat climate change. As a result, there is a growing policy focus on a twin transition towards a green and digital future (Muench et al. 2022).

Generally speaking, technological change is pivotal for reaching climate targets, as emerging technologies can mitigate or replace current activities that harm the environment.¹ However, new technologies can also aggravate pollution (Jaffe et al.

¹For instance, clean energy technologies are necessary for the decarbonisation of our energy system.

2003).² Hence, technological change can help or hinder environmental progress and the relationship appears to be intricate in nature. This ambivalence also applies to digital technologies (Lange et al. 2020). On the one hand, digital technologies can contribute to improved environmental outcomes by increasing energy and resource efficiency and providing dematerialised solutions (Berkhout & Hertin 2001). For instance, big data and artificial intelligence can help to prevent excess production and reduce error rates by providing more precise information. Also, virtual meetings require less energy than in-person ones. On the other hand, the production, use, and disposal of digital technologies also consume energy as well as resources and it is estimated that the direct emissions of digital technologies currently account for 2.1% to 3.9% of global greenhouse gas emissions (Freitag et al. 2021). Also, ICT-induced economic growth may spark energy consumption, as more goods can be consumed due to increased wealth. Consequently, it is unclear whether the green and the digital transitions complement each other or if the diffusion of digital technologies hinders the realisation of climate goals.

The complexity of the relationship may be a reason why previous empirical studies on the net effect of digital technologies on environmental outcomes, such as energy consumption, energy intensity, and carbon emissions, show inconsistent results (Zhang & Wei 2022). While some studies confirm environmental improvements (e.g. Schulte et al. 2016, Taneja & Mandys 2022, Wang et al. 2022, Xu et al. 2022), other studies find an increase in environmental harms (e.g. Sadorsky 2012, Alataş 2021, Ren et al. 2021). Most of these studies have in common that they rely on aggregated data and assume a linear relationship between digital technologies and environmental outcomes. Thus, diverging results may stem from some inherent shortcomings. Firstly, aggregation tends to bias results (e.g. Robinson 1950, Solow 1987, Koetse et al. 2008, Haller & Hyland 2014). Secondly, due to its complexity, the link between digitalisation and environmental outcomes is likely to be heterogeneous, i.e., it may depend on specific factors (cf. Horner et al. 2016, Lange et al. 2020). Hence, standard linear regression models may fall short of fully uncovering the intricate nature of the link. As technological progress is critical for economic growth and environmental outcomes, however, unbiased insights are of great public interest.

Furthermore, a prerequisite for the thorough analysis of the impact of technological progress on socioeconomic outcomes is to adequately measure the direction and diffusion of technological change. Standard innovation and digitalisation indicators are often based on questionnaire-based surveys (cf. Peters & Rammer 2013), which suffer from additional drawbacks, such as a lack of timeliness and regional granularity, as well as high non-response rates. In consequence, researchers are limited to the

²For example, the increasing use of rechargeable batteries to support renewable energies by auxiliary energy storage systems can contribute to water, soil, and air pollution if not disposed of properly (Mrozik et al. 2021).

information that is collected and very current analyses are rarely possible, which is especially an issue for the analysis of fast-changing technological trends. Measuring digitalisation poses some further challenges, as the term does not relate to a single but to a group of technologies and there is no universally agreed-upon method for measuring it.

Just as the digital transformation is reshaping our economy, it also offers new opportunities to address previous shortcomings in measuring technological change and analysing the impact of technological change on environmental outcomes.³ For instance, digital technologies enable the collection and usage of larger, more granular, and more diverse data sets (Einav & Levin 2014). One example are secure remote connections, which provide online access to large firm-level administrative data sets, facilitating analyses at the disaggregated level. This may allow for more precise insights into the link between digital technologies and environmental outcomes. Moreover, web data, i.e., information that is accessible via the internet, may pave the way for new approaches to measure technological change. Web data can be analysed in real-time (Choi & Varian 2012), which provides an advantage over survey-based innovation indicators that tend to have substantial time lags. However, web data is usually unstructured, which imposes additional challenges for empirical economists. Statistical learning algorithms (the most dominant form of AI)⁴ provide a remedy, as they allow to structure the data and to extract meaningful information (Mullainathan & Spiess 2017, Gentzkow et al. 2019). For instance, website texts can be converted into a format that is suitable for regression analyses by applying natural language processing tools. In addition, statistical learning enhances the accuracy of predictions for counterfactual outcomes, leading to an improved estimation of individualised treatment effects (Einav & Levin 2014, Chernozhukov et al. 2018, Athey & Wager 2019). As a result, differences in effects, i.e., heterogeneous relationships, can be observed with greater accuracy, allowing for new insights into the intricate link between digital technologies and environmental outcomes.

1.2 Contribution

In this dissertation, I leverage the new opportunities created by the digital transformation for empirical economists, such as (1) access to larger as well as more granular, diverse, and timely data sets, (2) statistical learning tools for structuring data, and (3) innovative methods for measuring heterogeneous relationships, to shed novel light on the measurement of technological change and the link between digitalisation and environmental outcomes.

³Please note that the list of advantages is not exhaustive.

⁴Statistical learning and machine learning are used as synonyms in this thesis.

The first essay of this thesis (Chapter 2) deals with the measurement of technological change. Here, my co-author Patrick Breithaupt and I explore the potential of firm websites in combination with text and data mining techniques to improve current survey-based measurement of firm-level innovation activities.

Today, nearly every firm has a website, which includes direct or indirect information about firm-level innovation activities, such as information on new products, key personnel decisions, and firm strategies (Gök et al. 2015). This information can be accessed in real-time and at a large scale, however, little is known about its accuracy. For instance, firms could present themselves in an unrealistically positive way. To provide insights into how useful this information is for measuring innovation activities, we use data on 4,487 firms from the Mannheim Innovation Panel (MIP) 2019, which is a large-scale questionnaire-based survey. Using this survey as a benchmark, we analyse which firm website characteristics perform best as predictors for survey-based innovation indicators. Website characteristics are measured by several text and data mining methods and are used as features in different Random Forest classification models. Our results show that the most relevant website characteristics are textual content in general, the percentage of text written in English language, the number of subpages, and the number of characters on a website. Furthermore, when all website characteristics are used together, the accuracy increases by up to 18 percentage points compared to a baseline prediction based on the sample mean. Results also indicate a better performance for the prediction of product innovators and firms with innovation expenditures than for the prediction of process innovators. Hence, we conclude that especially web-based information on product innovations and innovation expenditures is a useful complement to survey-based information. This essay was written as part of the research project “TOBI - Text Data Based Output Indicators as Base of a New Innovation Metric” (funding ID: 16IFI001, Dr. Georg Licht) and is published in the peer-reviewed journal *PLoS One*. In total, my contribution to this essay amounts to 50%.

All remaining essays focus on the link between digitalisation on environmental outcomes. In the second essay (Chapter 3), my coauthors Irene Bertschek, Patrick Breithaupt, Daniel Erdsiek, and I analyse the impact of digitalisation on changes in mobility over the course of the Covid-19 pandemic. Transportation is responsible for a large share of global carbon emissions (IEA 2020). However, the Covid-19 pandemic has sparked hope that firm digitalisation through remote work and online services will lead to long-lasting reductions in mobility, thereby persistently improving environmental outcomes. We scrutinise this belief by leveraging information on firm websites as in Chapter 2. However, in this chapter, we apply a novel text-mining approach based on transfer learning that allows for measuring the extent to which firms write about digitalisation on their websites. We use this information

as an indicator for firm digitalisation. This approach has the advantages that it allows for the timely measurement of digitalisation for a large number of firms and it does not focus on specific digital technologies but rather aims at measuring digitalisation in general. We link the derived firm digitalisation indicator to changes in mobility (observed via mobile network data) within German districts between January 2020 and December 2022. Our results indicate a decrease in mobility associated with firm digitalisation compared to the pre-crisis level during the first two years of the pandemic. However, the decrease in mobility diminished after most Covid-19 restrictions were lifted, suggesting that environmental improvements are not long-lasting. This essay did not receive any funding and I contributed 50% of the total content.

The ongoing digital transformation has also raised hopes for ICT-based climate protection within manufacturing industries, but restricted access to administrative micro-data has made it difficult, in the past, to determine how digital technologies affect firm-level energy use. Previous studies that examine manufacturing industries tend to find substantial improvements in energy intensity but rely on aggregated data, such as Bernstein & Madlener (2010), Schulte et al. (2016), and Taneja & Mandys (2022). In the last two essays of my thesis, I contribute to closing this research gap and explore administrative panel data on more than 25,000 manufacturing firms collected by German statistical offices between 2009 and 2017. Thoroughly calculated software capital stocks serve as an indicator for the firm-level degree of digitalisation. Both essays have been written as part of the research project “Climate Protection Potential of Digital Transformation (CliDiTrans): Micro- and Macroeconomic Evidence on the Role of Demand Effects and Production Relocation” (funding ID: 01LA1818B, Dr. Thomas Niebel).

The first of both essays (Chapter 4) addresses the shortcoming of a potential aggregation bias in previous empirical literature on the climate protection potential of digital technologies in manufacturing. For this purpose, my co-author Thomas Niebel and I use firm-level data to re-estimate production/cost function approaches previously used to measure the link between digitalisation and energy intensity with aggregated data. Our results confirm a robust link between software capital and energy intensity improvements. However, the effect size is rather small. For the average firm, we find that a one per cent increase in software capital relates to a decrease in energy intensity by 0.003%. Hence, in contrast to previous industry-level results, we do not observe substantial energy intensity improvements in connection with ICT adoption. Moreover, we find that the relationship between ICT and en-

ergy intensity exhibits properties that can lead to an aggregation bias. My relative contribution to this essay is 60%.

The second of both essays (Chapter 5) extends existing literature by focusing on the heterogeneous impacts of digital technologies on absolute energy consumption in manufacturing firms and sheds light on the most important drivers of differences in effects. My co-authors Anne Berner, Thomas Kneib, and I apply non-linear, flexible tree-based statistical learning. Our results strongly indicate firm-level heterogeneity, but suggest that digital technologies are more frequently related to an increase in absolute energy use than to a decline. Multiple characteristics such as energy prices and firms' energy mix explain differences in the effect. I contributed 45% of this essay's content.

In summary, my research shows that (1) the use of web data and text mining tools can improve measurement of technological change, and (2) leveraging large administrative data sets as well as web data in combination with advanced statistical learning methods provides a new understanding of the link between digitalisation and environmental outcomes. In particular, my results with respect to environmental outcomes indicate that digitalisation, at least at its current stage, does not necessarily lead to significant environmental improvements. This result is especially relevant for policymakers, as it challenges the belief that a green and digital future will occur naturally without political guidance.

All essays were originally written for publication in peer-reviewed journals. In Table 1.1, I list these essays, referring to the co-authors, the publication status, and my contribution. The remainder of this thesis is structured as follows: In Chapter 2, the first essay is presented, Chapter 3 contains the second essay, Chapter 4 the third essay, and Chapter 5 the fourth essay. In Chapter 6, I link the research strands, discuss my results, and provide avenues for future research.

Table 1.1: **Contribution table.**

Paper	Innovation Indicators Based on Firm Websites – Which Website Characteristics Predict Firm-Level Innovation Activity?
Co-authors	Patrick Breithaupt
Status	Published in <i>PLoS One</i>
Own key contributions	My contribution is 50%: <ul style="list-style-type: none"> • I conceptualised the paper together with Patrick Breithaupt. • I was responsible for the data curation and data analysis jointly with Patrick Breithaupt. • I wrote the manuscript jointly with Patrick Breithaupt.
Paper	Firm Digitalisation and Mobility – Do Covid-19-Related Changes Persist?
Co-authors	Irene Bertschek, Patrick Breithaupt, and Daniel Erdsiek
Status	Working paper
Own key contributions	My contribution is 50%: <ul style="list-style-type: none"> • I conceptualised the paper together with Irene Bertschek and Daniel Erdsiek. • I conceptualised the digitalisation indicator together with Patrick Breithaupt. • I was responsible for the data curation and econometric analysis. • I mainly wrote the manuscript.
Paper	Digital Technology Adoption and Energy Intensity in Manufacturing – Firm-Level Insights
Co-authors	Thomas Niebel
Status	Working paper
Own key contributions	My contribution is 60%: <ul style="list-style-type: none"> • I conceptualised the paper together with Thomas Niebel. • I developed the empirical strategy and I was responsible for the data analysis. • I mainly wrote the manuscript.
Paper	What Drives the Relationship Between Digitalisation and Industrial Energy Demand? Exploring Firm-Level Heterogeneity
Co-authors	Anne Berner and Thomas Kneib
Status	Working paper
Own key contributions	My contribution is 45%: <ul style="list-style-type: none"> • I conceptualised the paper together with Anne Berner. • I developed the empirical strategy and implemented the econometric modelling jointly with Anne Berner. • I wrote major parts of the manuscript.

Chapter 2

Innovation Indicators Based on Firm Websites – Which Website Characteristics Predict Firm-Level Innovation Activity?

joint work with Patrick Breithaupt

2.1 Introduction

Innovation, defined as the implementation of either new or significantly improved products or processes as well as combinations thereof (OECD/Eurostat 2019), brings vast benefits to consumers and businesses. Moreover, technological progress is considered a main driver of economic growth (Solow 1957). It is, therefore, a matter of public interest to analyse and understand innovation dynamics, as it is conducted in several studies (e.g., Crepon et al. 1998, Klomp & Van Leeuwen 2001, Belderbos et al. 2004, Hall et al. 2005, Griffith et al. 2006, Frenz & Ietto-Gillies 2009, Kogan et al. 2017).

A prerequisite for the analysis of innovation-related questions is to correctly measure firm-level innovation activities. However, it should be noted that no universally accepted measurement approach exists. For example, firm-level innovation indicators are traditionally constructed with data from large-scale questionnaire-based surveys like the biennial European CIS or the annual MIP (see Rammer et al. 2019, Peters & Rammer 2013), which is also the German contribution to the CIS. However, these innovation indicators suffer from some major drawbacks (cf. Mairesse & Mohnen 2010, Pukelis & Stanciauskas 2019, Kinne & Axenbeck 2020). For instance, the MIP annually covers around 18,000 firms, which corresponds to only a fractional share of the total number of German firms. As a result, the survey may lack regional granularity and comprehensive coverage. In addition, questionnaire-based surveys

– especially on a large scale – have the added disadvantages of being costly and lacking timeliness. Also, most surveys require firm participation and as a consequence, surveys such as the MIP suffer from low response rates (Mairesse & Mohnen 2010). Besides, firm-level innovation can also be studied by patent or publication analysis. However, respective indicators only cover technological progress for which legal protection is sought (Archibugi & Planta 1996, Arundel & Kabla 1998) and not every innovation can be patented. For example, due to the German regulatory framework it is quite difficult to patent software, i.e., digital innovations.

Issues, however, could be solved by adding web-based information: Advances in computing power, statistical learning methods, as well as natural language processing tools enable, e.g., researchers to extract website information on a large scale. This makes it technically possible to complement traditional innovation indicators with information from scraped firm websites. Nowadays, almost every firm has an online presence. Firm websites can include information about new products, key personnel decisions, firm strategies, and relationships with other firms (Gök et al. 2015). Those pieces of information might be directly or indirectly related to a firm's innovation status. By using this information, it is possible to conduct an automatic, timely, and comprehensive analysis of firm-level innovation activities, as measurements can be carried out faster and in shorter intervals in comparison to traditional indicators.

The contribution of this paper to the question of whether web-based innovation indicators are feasible is threefold. First, we analyse to what extent firm websites improve predictions of firm-level innovation activity. Second, we assess which characteristics of a website relate most to a firm's innovation status. Third, we examine which characteristics are appropriate for predicting different forms of innovation activity. We test the latter by additionally comparing the predictive power for different innovation indicators related either to product innovations, process innovations or innovation expenditures. We assume differences between indicators, for example, because firms with process innovations may have a smaller incentive to announce respective innovation activity on their websites. This may be due to the fact that new processes are less relevant for most website visitors.

For our analysis, data on 4,487 German firms from the MIP 2019 is used. We extract their websites' text and hyperlink structure by applying the ARGUS web-scrapers (Kinne & Axenbeck 2020). Several methods including topic modelling and other natural language processing tools are applied to generate features that potentially relate to the firm-level innovation status. Furthermore, we extract information related to a website's technical maturity, such as how fast it is responding and whether a version for mobile end user devices is available. After extracting and calculating a wide variety of features, we divide them into three different feature sets:

I) text-based features including, e.g., words, document-topic probabilities derived from a topic modelling algorithm, and the share of English language, II) meta information features including, e.g., website size-related features, availability of a mobile version and loading time, and III) network features including, e.g., hyperlinks to social networks, as well as incoming and outgoing hyperlinks. Based on these three feature groups, we analyse which website characteristics best predict a firm's innovation status reported in the MIP 2019 by using a Random Forest classifier.

Our results show that predictions based on website characteristics can perform significantly better than a random prediction based on the sample mean. Consequently, firm websites include information that relates to firm-level innovation activity. In addition, our website characteristics better predict firms with product innovations and innovation expenditures than with process innovations. Moreover, text features make the biggest contribution to our prediction performance.

Evaluating the predictive power of single variables across feature sets by means of the mean decrease in impurity (MDI) reveals that the language of a website and website size measured by the number of subpages as well as the total amount of characters are always relevant in the models with the highest predictive power for all considered innovation indicators. Moreover, there are characteristics that are highly important only for specific indicators, e.g., the verb "to develop" is more important for innovation expenditures and product innovators than for process innovators. The remainder of this paper is structured as follows: Previous literature is reviewed in Section 2.2. In Section 2.3, we present our data and in Section 2.4 the descriptive statistics. Section 2.5 describes the methodology and Section 2.6 shows the results, which are discussed in Section 2.7. This paper concludes in Section 2.8.

2.2 Literature Review

The use of text data to generate innovation-related indicators has already been tested in previous studies. For example, Kelly et al. (2021) show that the significance, i.e., relevance, of a patent is higher when its textual content is very distinct to previous patents but similar to subsequent ones. Lenz & Winker (2020) generate innovation-related topics from 170,000 technology news articles using a Paragraph Vector Topic Model. They analyse the diffusion of the identified topics within the text corpus. Their results suggest that technology trends can be assessed by measuring the importance of topics over time. Using PATSTAT data, Tacchella et al. (2020) show that context similarity of technological codes relates to innovative events. The likelihood that new combinations of technological codes appear in one patent can be predicted by their context similarity in patents where they have been used before.

Remarkable work is also conducted by Bellstam et al. (2021). In this study, a Latent Dirichlet Allocation (LDA) model is fitted with analyst reports of firms included

in the S&P 500 index. The LDA topic that has the lowest Kullback-Leibler divergence to the wording of a mainstream economic textbook on innovation is chosen as an innovation indicator. The authors show that firms have patents with greater impact (i.e., more citations per patent) if the innovation topic has a larger share in their analyst report. However, analyst (or also annual) reports are not available for every firm and smaller firms are particularly underrepresented. In contrast, firm websites are available for a large share of small and medium-sized firms.

Furthermore, previous literature shows that information produced online can be used to construct frequent real-time estimates (Gentzkow et al. 2019). Famous ‘nowcasting’ examples that utilize web-based information are Ginsberg et al. (2009), who use Google search queries to accurately predict influenza activity in the United States (US). Choi & Varian (2012) claim that search engine query indices are also often correlated with economic activities and enable to generate frequent indicators. They show that forecasts concerning, for example, automobile sales and unemployment can be significantly improved by including search term indices in prediction models.

Not only information from online searches but also firm website information can be used to generate economic indicators. As firm websites provide detailed information about the firm as well as its products, they appear to be suitable for measuring firm-level innovation activities (Gök et al. 2015). Kinne & Axenbeck (2020) summarise previous studies that analyse the possibility of firm website-based innovation indicators (e.g., Katz & Cothey 2006, Ackland et al. 2010, Arora et al. 2013, Gök et al. 2015, Beaudry et al. 2016, Nathan & Rosso 2022). Most studies solely focus on the hyperlink structure of websites or only conduct a simple keyword search and are limited to small amounts of firms from a particular economic sector.

Firstly applying advances in statistical learning, Kinne & Lenz (2021) attempt to predict innovation at the firm level using textual information on websites and novel machine learning tools. They use a questionnaire-based firm-level product innovation indicator (innovative/ non-innovative) from the MIP (years 2015-2017) as a target variable to train an artificial neural network classification model on website texts. The authors only consider stable product innovators in their main analysis. Firms that switch between innovation statuses, which is a phenomenon that is highly relevant in the field of innovation economics, are only observed in a secondary analysis. The average F1-score for the respective prediction is 0.68%. Moreover, Pukelis & Stanciauskas (2019) fit several machine learning models to develop a firm website-based innovation indicator, with their annotated data set being limited to 500 firms. One important characteristic of their work is the individual analysis of websites’ subpages instead of predicting the innovation status of an entire website, i.e., a firm. Additionally, their subpages are manually labelled as either innovation or non-innovation-related messages instead of using survey or patent data as target

variables. The best performance is achieved with an artificial neural network. Even though the predictive performance is very high, the authors cannot show a credible external validity of their indicator.

Furthermore, another issue of both approaches is that neural networks do not reveal any decision rule that can be easily interpreted by humans, which is why they are often called black box models. It should also be noted that both studies only consider text. Nonetheless, previous results show that there must be distinct website characteristics that relate to a firm's innovation status, but the particular website characteristics are not identified yet.

Gandin & Cozza (2019) analyse whether firms' expenditures on innovation can be predicted by means of administrative records and balance sheet data. Using a Random Forest regression approach, the authors identified firm size, sectoral affiliation, and investment in intangible assets as the most important predictors. Random Forests usually provide better predictive performance than linear methods while retaining the interpretability of feature relevance.

By applying a Random Forest approach to large-scale firm-level web data, we are able to analyse which website characteristics are linked to firms' innovation activity.

2.3 Data

Based on the Oslo Manual, we define an innovation as "a new or improved product or process (or combination thereof) that differs significantly from the unit's previous products or processes and that has been made available to potential users (product innovation) or brought into use by the unit (process innovation)" (OECD/Eurostat 2019, p. 20). Furthermore, we consider all expenditures spent for innovation purposes as innovation expenditures and summarise firm-level product or process innovation as well as innovation expenditures as innovation activity.

We use data from the MIP 2019 to classify firms as either innovative or non-innovative.⁵ The MIP is an annual survey conducted by the ZEW – Leibniz Centre for European Economic Research. The survey covers firms from manufacturing and service sectors and is conducted as a mail survey with the option to respond online.

In the MIP 2019, firms were asked whether they introduced a product or process innovation within the last three years (between 2016 and 2018) and for the total amount spent on innovation activities in the last year (2018). We consider a firm that stated it introduced a product innovation within the considered time frame as a product innovator, and a firm that stated that it introduced a process innovation within the considered time frame as a process innovator. A firm is an innovator if

⁵See [data set] ZEW – Leibniz Centre for European Economic Research (2019).

it introduced at least one of both. Every firm that spent financial resources on innovation - independent of the magnitude - is regarded as a firm with innovation expenditures. Our initial sample consists of 13,747 firms from the MIP 2019. We merge these firms with the Mannheim Enterprise Panel (MUP), which consists of more than 3.2 million economically active firms (see Bersch et al. 2014), to receive information about the firms' website addresses. The MUP serves as a sampling frame for surveys like the MIP and, e.g., contains firm-level information on turnover, number of employees, and sector affiliation. Only 54% of the firms in our sample can be linked to website addresses, as we limit ourselves to quality-assured observations. In total, we have 6,368 firms with information on the website address and at least one innovation indicator. We extract website content by applying the ARGUS web-scrapers, which allows us to collect texts as well as hyperlinks to other websites.⁶ Firm websites were first scraped in September 2018 to collect texts, then again in January 2019 for adding hyperlinks. We scraped a third time in October 2019 to add information about technical features, e.g., information on the existence of firm websites for mobile end user devices. The maximum limit of scraped subpages per website is set to 50. We consider this to be a sufficient number, as the median number of subpages in the MUP is 15 (see Kinne & Axenbeck 2020), and only 1.5% of all firms in our subsample have 50 or more subpages. Moreover, the scraping program is set to prefer subpages with shorter website addresses because we assume these subpages include more important information about the firm. Also, ARGUS is set to prefer websites in German language. Hence, when we calculate the share of different languages on a website, we expect a small bias. However, since only a few firms exceed the subpage limit, we assume this bias to be negligible. While scraping the data, especially while collecting meta information features, we received several error messages. Furthermore, we only use observations for which all features are non-missing. If, for example, a meta information feature is not available, the observation will not be used for training or testing with other feature sets. Therefore, after the entire data collection process, we have a sample of 4,487 firms for predicting product innovators and innovators, and 4,484 firms for predicting process innovators.⁷ For predicting whether a firm has innovation expenditures, the sample size is 1,893 (Table 2.1).

Additionally, a random sample of approximately 32,000 website addresses of firms not included in the MIP is drawn from the MUP and scraped with the ARGUS web-scrapers using the same settings as for the MIP sample. The sample is used for topic modelling. We train a topic model on a separate sample for two reasons.

⁶For a detailed description of the ARGUS web-scrapers, see Kinne & Axenbeck (2020) and Kinne (2018).

⁷There are three observations more for product innovators than for process innovators. Since we know that these observations have at least a product innovation, we also consider them in the innovator sample.

Table 2.1: Summary statistics for product innovators, process innovators, innovators, as well as firms with innovation expenditures.

Variable	Definition	N	Mean	SD	Min	Max
Product innovators	1: If firm is a product innovator 0: Otherwise	4,487	0.39	0.49	0	1
Process innovators	1: If firm is a process innovator 0: Otherwise	4,484	0.52	0.50	0	1
Innovators	1: If firm is a product or / and process innovator 0: Otherwise	4,487	0.61	0.49	0	1
Innovation expenditures	1: If firm innovation expenditures were reported 0: Otherwise	1,893	0.39	0.49	0	1

First, it allows to include more data points. Second, it ensures that no observation used for calibrating topics is considered for evaluating the Random Forest models. Hence, it prevents data leakage. The sample is hereinafter referred to as the LDA sample.

As we need to exclude a large share of observations due to missing values in our MIP sample, we cannot rule out a selection bias. Also, firms from certain industries and smaller firms are less likely to have a website and may, therefore, be underrepresented. In machine learning, adverse selection might lead to two issues: It could cause that our model is better fitted for groups that are overrepresented in our sample, and it could induce that the class correlated with the overrepresented group is predicted more often. To identify whether a potential selection bias exists, we analyse how the sample distribution changes with respect to the number of employees and industry sectors, when excluding observations with missing information (see A.1 and A.2). Except for “transportation and post” (sector 15), we do not see a notable change in the distribution of firms that could be linked to a severe selection bias.

To capture website characteristics, we apply several methods to generate features like a keyword search and natural language processing as well as an analysis of hyperlinks (network analysis methods). We use Python as programming language for calculating our features and for training our Random Forest models. For an overview of feature sets see Table 2.2.

Table 2.2: Features related to text, meta information, and network measures.

<i>Text-based features</i>	
1) Textual content	Term-document matrix with the 5,000 most frequent words (TF-IDF applied).
2) Emerging technologies	Dummy variable that measures whether a technology of Wikipedia's list of emerging technologies appears on a firm's website.
3) Latent patterns	Topic-document probabilities of 150 topics generated by the LDA approach.
4) Topic popularity index	The sum of LDA topic probabilities per document. Each probability is weighted with the relative frequency of its appearance in the entire LDA sample.
5) International orientation	Share of subpages in English language and the share of all other non-German subpages in all subpages.
6) Share of numbers	The share of numbers in a text of a website (measured in characters).
7) Flesch-reading-ease score	Numerical metric assessing the readability of texts.
<i>Meta information features</i>	
8) Website size	Number of subpages on a website, total amount of characters on a website.
9) Loading time	The time from sending a request (http/https) to a webserver (to get the start page of a website) until the arrival of the response (in ms).
10) Mobile version	Dummy variable that is one if a version for mobile end user devices exists and zero otherwise.
11) Domain purchase year	The year of the first entry on web.archive.org.
<i>Network features</i>	
12) Centrality	The total number of incoming and the total number of outgoing hyperlinks, as well as the PageRank centrality.
13) Social media	Number of hyperlinks to Facebook, Instagram, Twitter, YouTube, Kununu, LinkedIn, XING, GitHub, Flickr, and Vimeo.
14) Bridges	Number of bridges a firm is part of in the hyperlink network.

2.3.1 Text-based Features

Information from website texts is analysed, as it might be related to a firm's innovation status for the following reasons: Presumably, most firms are using their websites to inform customers about new products or services and might mention whether their product is new or innovative, i.e., it is likely that innovative firms use particular innovation-related words. Information about process innovations can also be detected and used if reported on the website.

Moreover, a firm might report that it uses a recently emerging technology like blockchain, 3D printing or augmented reality (for an overview of recently emerging

technologies, see Appendix A.2). Hence, an emerging technology term might appear on a firm's website and if so it is likely that the firm can be considered as innovative, as it makes use of technologies that are fairly new.

Additionally, there might be latent patterns on a website that reveal a firm's innovation status, these latent patterns can be captured by the LDA topic modelling approach as successfully shown in Bellstam et al. (2021). Furthermore, innovative firms might follow some general technological trends like the digital transformation. As these technological trends are quite general, LDA topics related to these trends might appear quite often on firm websites. To capture this, we construct a topic popularity index that indicates the distribution of popular and less popular topics on a website.

We additionally analyse the following text-based metrics: Languages that appear on a website might relate to the export status of a firm and this could provide information about a firm's innovation status because the export status is linked to firm-level innovation (e.g., Lachenmaier & Wößmann 2006, Kirbach & Schmiedeberg 2008, Cassiman & Golovko 2011). Also, we test whether the share of numbers in all string characters (text) as well as the text complexity measured by the Flesch-reading-ease score (Flesch 1948) differ between innovative and non-innovative firms.

2.3.2 Meta Information Features

Second, meta information of firm websites (see Table 2.2) might allow to distinguish innovative from non-innovative firms. For example, the website size might help to predict a firm's innovation status. Large firms are more likely to be innovative (Rammer et al. 2019). As the number of subpages of a website correlates with the number of employees of a firm (Kinne & Axenbeck 2020), the size of a website might provide information about whether a firm introduced an innovation. Also, the technological properties of a website could be relevant. Innovative firms might have a better technical knowledge and are able to apply more technologically advanced features on their websites. For example, the loading time of a website could be faster and a mobile version might be more often available when firms are more technologically advanced. However, there might be some noise because the loading time may also be short if the website is relatively simple.

Another potentially relevant feature is the age of a website, i.e., the domain purchase year, as it might relate to the actual firm age.⁸ One has to consider, however, that this relationship is unlikely to be linear. On the one hand, a website that is fairly new might indicate a start-up with an innovative idea. On the other hand, having a

⁸We approximate a website's domain purchase year by the year of the first entry on web.archive.org.

very old website means the firm has adopted this new technology very early. This could also relate to a more technologically advanced, hence, innovative firm.

2.3.3 Network Features

Third, hyperlinks between websites (see Table 2.2) might also help to identify the firm-level innovation status. Firms that have more business relationships with other firms or are more relevant according to centrality measures might be better informed and know earlier about new profitable applications. Hence, firms with more relationships to other firms could be more likely to be innovative. Moreover, innovation projects are often realised in cooperation with other firms (e.g., Becker & Dietz 2004). Thus, patterns in firm-level cooperation are expected to be of interest. A firm that connects (or bridges) different network parts is usually relevant and its removal will decompose the network. Lastly, Bertschek & Kesler (2022) show that a firm's use of the social network Facebook is linked to product innovations. Hence, the use of social media might reveal information about a firm's innovation status, as well.

Our study analyses whether the three groups of features differ in their performance when predicting a firm's innovation status. A more detailed description of the feature generation can be found in Appendix A.3.

2.4 Descriptive Analysis

The descriptive statistics for our predictor variables are presented in this section. Table 2.3 shows mean values for innovative and non-innovative firms, as well as p-values, obtained from a t-test, regarding the difference of both means for selected features.

Differences exist for most variables. Looking at 'text' features, innovative firms are more likely to mention an emerging technology term and have more subpages in English language. The share of subpages in other languages, however, does not show any significant difference between both groups. Differences are also small for the share of numbers, our topic popularity index, and the Flesch-reading-ease score, but the deviation is statistically significant for some forms of innovation activity.

The descriptive statistics for 'meta' features show that innovative firms have larger websites with respect to the number of subpages as well as with respect to the number of characters. The loading time is slightly faster for process innovators and innovators, but not for product innovators and firms with innovation expenditures. However, differences are not statistically significant. The first occurrence on web.archive.org is significantly later for non-innovative firms indicating their domain purchase year, i.e., website age, is slightly lower. Additionally, non-innovative firms have less often a version of their website for mobile end user devices. Looking

Table 2.3: Descriptive statistics for selected variables.

Feature (Variable name)	Group-specific means											
	Product innovator			Process innovator			Innovator			Innovation expend.		
	Yes	No	P-val.	Yes	No	P-val.	Yes	No	P-val.	Yes	No	P-val
Text-based features												
Emerging technology term (<i>emerging_tech</i>)	0.18	0.07	0.00	0.15	0.07	0.00	0.15	0.05	0.00	0.19	0.06	0.00
Percentage of English language (<i>english_language</i>)	0.16	0.10	0.00	0.14	0.10	0.00	0.14	0.09	0.00	0.17	0.08	0.00
Percentage of other languages (<i>other_lang</i>)	0.02	0.02	0.45	0.02	0.02	0.25	0.02	0.02	0.74	0.02	0.02	0.30
Topic popularity index (<i>pop_score</i>)	34.64	34.35	0.36	34.78	34.11	0.03	34.68	34.13	0.08	35.07	33.82	0.01
Share of numbers (<i>share_numbers</i>)	0.025	0.028	0.00	0.025	0.028	0.00	0.026	0.028	0.00	0.027	0.027	0.97
Flesch-reading-ease score (<i>flesch_score</i>)	40.09	41.22	0.01	40.54	41.03	0.26	40.47	41.26	0.09	39.28	41.28	0.01
Meta information features												
Website size: length (<i>text_length</i>)	75,269.35	56,746.84	0.00	71,629.95	55,685.73	0.00	71,193.63	52,859.37	0.00	75,334.75	52,462.63	0.00
Website size: nr. of pages (<i>nr_subpages</i>)	30.37	24.65	0.00	28.75	24.87	0.00	28.92	23.75	0.00	31.23	23.58	0.00
Loading time (<i>load_time</i>)	0.57	0.55	0.69	0.51	0.60	0.25	0.55	0.57	0.76	0.51	0.49	0.57
Mobile version (<i>mobile_version</i>)	0.76	0.70	0.00	0.76	0.68	0.00	0.75	0.67	0.00	0.73	0.69	0.06
Domain purchase year (<i>domain_purchase_year_proxy</i>)	2004.22	2004.98	0.00	2004.42	2004.96	0.00	2004.37	2005.17	0.00	2004.38	2005.01	0.01
Network features												
Outgoing hyperlinks (<i>outgoing_links</i>)	15.93	12.95	0.00	15.18	12.97	0.00	15.19	12.46	0.00	16.23	12.38	0.00
Incoming hyperlinks (<i>incoming_links</i>)	14.78	5.22	0.00	13.24	4.30	0.00	12.11	4.09	0.00	12.09	3.70	0.00
Use of social media (<i>social_media</i>)	1.62	1.02	0.00	1.51	0.98	0.00	1.47	0.92	0.00	1.62	0.91	0.00
PageRank centrality (<i>pagerank_index</i>)	$2 * 10^{-6}$	$1 * 10^{-6}$	0.00	$2 * 10^{-6}$	$1 * 10^{-6}$	0.00	$1 * 10^{-6}$	$1 * 10^{-6}$	0.00	$1 * 10^{-6}$	$1 * 10^{-6}$	0.01
Bridges (<i>bridge_index</i>)	0.43	0.26	0.01	0.38	0.28	0.05	0.37	0.27	0.04	0.31	0.27	0.35
Number of observations	4,487			4,484			4,487			1,893		

Notes: All variables were rounded to the second decimal place except PageRank centrality, which was rounded to the sixth decimal place and share of numbers which was rounded to the third decimal place.

at ‘network’ features, significant differences also exist for outgoing and incoming hyperlinks as well as for hyperlinks to social media websites. Innovative firms have on average more hyperlinks. Moreover, the difference is larger for incoming than for outgoing or social media hyperlinks. Additionally, innovative firms also are significantly more important in firm networks looking at the PageRank centrality. The statistical significance of differences regarding the bridge index is, however, limited to the form of innovation activity. In summary, Table 2.3 confirms previous assumptions. Innovative firms seem more likely to apply emerging technologies, to have more technically advanced websites, and to be better connected with each other according to most network indicators.

Figure 2.1 shows the average occurrence of different emerging technology terms on a firm website with respect to product innovators. The emerging technology terms differ strongly in their likelihood of occurrence. The emerging technology term *Internet of Things* is the most likely to occur. It appears on more than 8% of all product innovator websites and only on less than 2% of all non-product innovator websites. Also, terms relating to different machine learning applications, *biometrics*, *blockchain* technology and *mobile collaboration* appear relatively often. Moreover, for nearly every emerging technology term it is more likely to appear on a product innovator website than on a non-product innovator website. This result is the same for all innovation indicators.

Table 2.4 shows the ten most innovation-relevant LDA topics. The highest average value of Pearson correlation coefficients for all four innovation indicators and the document-topic probabilities is used to identify the most relevant LDA topics. The topics are sorted in descending order. LDA topic 98, which relates according to its keywords to research & development, has a positive and by far the strongest relationship to innovation. Also, LDA topic 35, which relates to ICT infrastructure, has a comparatively strong positive correlation with our innovation indicators. Among the top ten, the LDA topics 20 (tourism), 120 (consulting & customer support) and 23 (family business & craftsmanship) have the weakest correlation. Moreover, the correlation is negative.




Figure 2.1: Average occurrence of different emerging technology terms on firm websites with and without product innovations. Emerging technology terms not appearing on firm websites are not illustrated. The y-axis has a scale break at 2%.

Table 2.4: Content of the LDA topics with the strongest relationship to MIP-based innovation indicators.

Topic number	Content	Translated	Top words	Correlation*
LDA topic 98	Research & development	yes	'company' 'customer' 'development' 'to develop' 'department' 'employee' 'partner' 'project' 'successful'	positive (0.15)
LDA topic 35	ICT infrastructure	yes	'system' 'software' 'data centres' 'server' 'version' 'support' 'date' 'windows' 'automatic' 'document'	positive (0.10)
LDA topic 65	Construction	yes	'to build' 'project' 'new building' 'architect' 'planning' 'renovation' 'reconstruction' 'construction' 'to plan' 'architecture'	negative (-0.09)
LDA topic 134	Business software	no	'array' 'value' 'news' 'office' 'paket' 'error' 'data' 'page' 'SAP' 'search'	positive (0.08)
LDA topic 7	Product experience	no	'centro' 'company' 'best' 'use' 'experience' 'world' 'please' 'product' 'may' 'find'	positive (0.08)
LDA topic 41	Common terms	yes	'and' 'far' 'to take place' 'to put' 'frame' 'that' 'information' 'total' 'receive' 'department'	negative (-0.07)
LDA topic 5	Carpentry	yes	'to tile' 'woods' 'to lay' 'laminat' 'tile' 'to put' 'material' 'stairs' 'floor' 'to glaze'	negative (-0.07)
LDA topic 20	Tourism	yes	'region' 'city' 'to be located' 'to offer' 'museum' 'old' 'historical' 'nature' 'tour' 'landscape'	negative (-0.06)
LDA topic 120	Consulting & customer support	yes	'pleased' 'to offer' 'customer' 'to advise' 'individual' 'consulting' 'available' 'question' 'competent' 'to find'	negative (-0.06)
LDA topic 23	Family business & craftsmanship	yes	'company' 'to operate' 'visit' 'to stand' 'roofing' 'Michael' 'son' 'specialize' 'work'	negative (-0.06)

*Measured by the average of all Pearson correlation coefficients between the average topic share per document and each innovation indicator.

Figure 2.2 also relates to the ten most innovation-relevant LDA topics. It shows for every topic the average share in a document for innovative and non-innovative firms. The figure reflects the results presented in Table 2.4. The selected topics considerably differ between innovative and non-innovative firms. Also, relationships are constant, e.g., if a topic has a larger share on product innovator than on non-product innovator websites, it will also be relatively stronger represented on process innovator websites. Nonetheless, differences between innovation indicators exist. Average topic share differences diverge between indicators and are larger when considering firms' innovation expenditures than when taking product or process innovators into account.




Figure 2.2: Differences in the topic share of the top ten topics with the strongest average correlation with MIP-based innovation indicators. For instance, the LDA topic 98 has an average share of 10% in a document if a firm has innovation expenditure, compared to merely 6% if a firm does not have innovation expenditure.

2.5 Methodology

The objective of our work is the identification of website characteristics that allow for predicting firm-level innovation activities. For this purpose, we integrate the described features as predictor variables in Random Forest classification models (Breiman 2001, Friedman et al. 2001). For each of our feature sets ('text', 'meta', and 'network' features), as well as for all features jointly, a separate Random Forest model is fitted. We use the Python package *scikit-learn* for the exercise. The Random

Forest algorithm is an ensemble method used for classification or regression tasks. Like any other machine learning algorithm, as defined in Mohri et al. (2018), it uses past experience (in our case survey data) to learn how to perform predictions. The Random Forest algorithm makes its decision based on the modus or mean of a multitude of decorrelated decision trees. Each tree is built based on bootstrapped samples of training data. By splitting the data at nodes into branches that are more “pure” with respect to the target variable, the algorithm learns to improve. We chose the Random Forests algorithm because it has the advantage that it allows for the calculation of feature importances, while providing high predictive power and enabling the consideration of complex interactions.

A formal description of the “decrease in impurity” is given by Equation (2.1). $i(t)$ measures impurity at the node level, which is in our case indicated by the Gini impurity index. t is a node within one tree and s is a split at a certain value of a variable. N_x is the number of samples reaching node $x \in \{t, t_L, t_R\}$. Lastly, if t is the parent node, t_L is the left child node and t_R is the right child node for the split s at node t . The split s for node t that maximizes $\Delta i(s, t)$ is iteratively chosen.

$$\Delta i(s, t) = i(t) - N_{t_R}/N_t * i(t_R) - N_{t_L}/N_t * i(t_L) \quad (2.1)$$

Feature importance is then derived by the sum of “decreases in impurity” of a single variable divided by the sum of “decreases in impurity” of all features used to build the tree. The value is additionally averaged over all trees in the forest and again normalised so that all values sum up to one. If multiple variables will lead to similar impurity decreases at one node, only one variable is selected for splitting. Hence, (multi-)collinearity of features can bias feature importance. This issue can be illustrated by the following. In this example, the same variable is included twice in a model. When choosing a variable for splitting, the model can randomly choose between the two and the feature relevance is, thus, divided between both variables.

To evaluate the performance of collected website characteristics, we use a baseline model. A random coin toss model based on the sample distribution is chosen. A baseline model works as a benchmark to assess the performance of more complex solutions, i.e., it helps to analyse whether a trained model performs better than a random prediction.

To estimate whether we achieve considerable improvements in comparison to baseline predictions, we perform a McNemar test (McNemar 1947). Assuming a chi-squared frequency distribution, the McNemar test measures if predictions from two machine learning models significantly disagree with each other, as illustrated in Equation (2.2). RF captures the number of observations misclassified by a fitted Random Forest model, but not by the baseline model. BL captures the number of

observations misclassified by the baseline model, but not by a fitted Random Forest model.

$$\chi^2 = \frac{(RF - BL)^2}{(RF + BL)} \quad (2.2)$$

If a model including a distinct feature set significantly disagrees with baseline predictions according to the McNemar test and its evaluation metrics show superior values, we consider this feature set to be relevant for the prediction of firm-level innovation activity.

To further evaluate and compare models, we use the metrics “area under the curve” (AUC), accuracy, and improvement of accuracy in comparison to the baseline model. We also use precision, recall, and the F1-score for positive as well as negative observations (Fawcett 2006).

$$\text{false positive rate} = \frac{FP}{(FP + TN)} \quad (2.3)$$

$$\text{true positive rate (recall for the positive class)} = \frac{TP}{(TP + FN)} \quad (2.4)$$

The AUC can be explained as follows. The formulas listed in Equations (2.3) and (2.4) are based on the number of false positive predictions (FP), capturing non-innovative firms wrongly predicted as innovative; true positive predictions (TP), capturing innovative firms correctly predicted as innovative; false negative predictions (FN), capturing innovative firms wrongly predicted as non-innovative; and true negative predictions (TN), capturing non-innovative firms correctly predicted as non-innovative. The receiver operating characteristic curve (ROC) is a graphical illustration of a binary classifier performance. For different classification thresholds, the “false positive rate” is plotted against the “true positive rate” and the AUC value is an approximation of the area below the ROC. Accordingly, the AUC value is the probability that a randomly chosen innovative firm is assigned a higher probability of being innovative than a randomly chosen non-innovative firm. Usually, AUC values above 0.7 are considered as acceptable (Hosmer et al. 2013).

For the other metrics, a classification threshold has to be set. The classification threshold is also called cut-off value and refers to the transformation of the regression output to a binary classification. Different cut-off values can be chosen if for example “false negatives” are considered more costly than “false positives” or if certain metrics need to be optimised. We select 0.5 as a cut-off value for all fitted models because this value is commonly used and we do not prefer one metric or class over the other.

$$\text{precision for positive class} = \frac{TP}{(TP + FP)} \quad (2.5)$$

$$\text{precision for negative class} = \frac{TN}{(TN + FN)} \quad (2.6)$$

$$\text{true negative rate (recall for the negative class)} = \frac{TN}{(FP + TN)} \quad (2.7)$$

Formal definitions of precision for innovative and for non-innovative firms are illustrated in Equations (2.5) and (2.6). Recall for innovative and for non-innovative firms is measured by the “true positive rate” or “true negative rate” as illustrated in Equations (2.4) and (2.7). Precision measures, for instance, the share of correctly classified innovative firms in all firms classified as innovative, while recall measures the fraction of innovative firms that have been correctly identified as innovative.

$$\text{accuracy} = \frac{(TP + TN)}{(TP + TN + FP + FN)} \quad (2.8)$$

$$\text{F1-score}_{P,N} = 2 * \left(\frac{(\text{Precision}_{P,N} * \text{Recall}_{P,N})}{(\text{Precision}_{P,N} + \text{Recall}_{P,N})} \right) \quad (2.9)$$

Accuracy and F1-score are presented in Equations (2.8) and (2.9). Accuracy measures the share of correct predictions in all predictions. The F1-score captures the harmonic mean between precision and recall for positive (P) and negative (N) observations, respectively. Baseline outcomes of accuracy, F1-scores, as well as precision and recall for our different innovation activity indicators are presented in Table 2.5 in Section 2.6. The random coin toss model assumes a fixed chance of being innovative (based on the sample mean). Hence, results do not change when the threshold is varied and, therefore, the AUC value is not displayed for baseline outcomes.

To control for overfitting, we analyse model performance by using out-of-sample predictions. Accordingly, we do not evaluate the models’ performance with the observations that are already used for learning. The data is split into a training sample (for fitting models) and a test sample (for evaluating models). To be more precise, the test sample is a “hold-out” sample and, therefore, never used for model training.

The training sample consists of 75% and the test sample consists of 25% of our observations. In the supervised learning context, this is a common partitioning method. It constitutes a trade-off between the generalisation of the model and the validity of the evaluation. We also apply a grid-search to tune the hyperparameters of all our models (Friedman et al. 2001) on our training sample. We explore the hyperparameter space for the ‘number of trees’ (100, 500, 1,000, and 1,500), ‘maximum tree depth’ (50, 100, 150, and 200), and ‘minimum impurity decrease’ (0.01, 0.001). For all other hyperparameters, we use default values provided by *scikit-learn*.

This leads to 32 different hyperparameter combinations for every model. For each hyperparameter combination in our grid-search, a five-fold cross-validation

is performed. The k-fold cross-validation belongs to the non-exhaustive cross-validation methods. It is a technique to assess the generalisability of machine learning models to new data, and to detect overfitting as well as potential sample biases. The data are split into k subsets, so that $100 - (100/k)$ % of the data are used for training the model and $100/k$ % for validation. In each of the k iterations, a different training and validation data set is used.

Considering all models fitted in the cross-validated grid-search, we choose the model with the highest AUC value. The selected model is then evaluated on the test sample.⁹

2.6 Results

In this section, we present the predictions of MIP-based innovation indicators using a Random Forest classification approach. Table 2.5 shows evaluation metrics for all baseline as well as fitted models. We analyse four different innovation indicators (four target variables), which we predict based on three different subsets of features as well as their union (four different groups of features). Accordingly, we train 16 Random Forest models.

Looking at product innovators, the highest AUC score (73%) is realised with ‘all’ features. The baseline accuracy is 0.53. The largest increase can be observed for the ‘all’ feature model (17 percentage points). Text-based features alone, however, lead to an increase of 16 percentage points. Moreover, ‘network’ and ‘meta’ features have a relatively weak impact. They just lead to improvements of 13 and 11 percentage points, respectively. This indicates that a large share of predictive power results from website text. The baseline F1-score for product innovators is 0.39 and for non-product innovators it is 0.61. Hence, the sample is slightly imbalanced towards non-product innovators and the chances of randomly predicting this class correctly are higher. Furthermore, the F1-scores show a similar result to other metrics. Only the ‘text’ and the ‘all’ feature model improve F1-scores notably. When solely applying ‘meta’ or ‘network’ features, F1-scores for innovative firms are even worse than the baseline performance. Precision values do not considerably differ between innovative and non-innovative firms and are always higher than the baseline prediction. Moreover, there is a comparatively large increase in precision for innovative firms. In contrast, there is a great difference between both classes with respect to recall values. For innovative firms, recall values of fitted models are always worse than those of the baseline prediction. For non-innovative firms, the recall fluctuates between 88 and 95%.

⁹To ensure the reproducibility of our study, we fixate the random seed when necessary. The random seed influences the model performance to some extent, e.g., observations are assigned to the train or test sample based on the random seed.

Table 2.5: Results for Random Forest classification models using different feature sets and target variables. Evaluation metrics are presented for the test sample.

Feature sets				AUC	Accuracy		F1-Score		Precision		Recall		McNemar	Support
Baseline	Text	Meta	Network		Value	Δ	Positive	Negative	Positive	Negative	Positive	Negative	P-values	
Product innovators														
x				-	0.53	-	0.39	0.61	0.39	0.61	0.39	0.61	-	1,122
	x			0.72	0.69	0.16	0.47	0.78	0.69	0.69	0.35	0.90	0.00	1,122
		x		0.66	0.64	0.11	0.37	0.75	0.59	0.66	0.27	0.88	0.00	1,122
			x	0.65	0.66	0.13	0.30	0.77	0.72	0.65	0.19	0.95	0.00	1,122
	x	x	x	0.73	0.70	0.17	0.49	0.79	0.71	0.70	0.37	0.90	0.00	1,122
Process innovators														
x				-	0.50	-	0.52	0.48	0.52	0.48	0.52	0.48	-	1,121
	x			0.62	0.59	0.09	0.63	0.54	0.59	0.59	0.67	0.50	0.00	1,121
		x		0.60	0.57	0.07	0.64	0.46	0.57	0.58	0.74	0.39	0.01	1,121
			x	0.59	0.57	0.07	0.62	0.52	0.58	0.56	0.66	0.48	0.01	1,121
	x	x	x	0.63	0.60	0.10	0.64	0.55	0.60	0.60	0.68	0.52	0.00	1,121
Innovators														
x				-	0.52	-	0.60	0.40	0.60	0.40	0.60	0.40	-	1,122
	x			0.67	0.63	0.11	0.75	0.30	0.63	0.59	0.91	0.20	0.00	1,122
		x		0.64	0.62	0.10	0.74	0.33	0.64	0.56	0.88	0.23	0.00	1,122
			x	0.62	0.60	0.08	0.75	0.00	0.60	0.00	1.00	0.00	0.00	1,122
	x	x	x	0.68	0.63	0.11	0.75	0.31	0.64	0.59	0.91	0.21	0.00	1,122
Innovation expenditures														
x				-	0.54	-	0.36	0.64	0.36	0.64	0.36	0.64	-	474
	x			0.74	0.73	0.19	0.55	0.80	0.68	0.74	0.47	0.88	0.00	474
		x		0.67	0.65	0.11	0.33	0.76	0.53	0.67	0.24	0.87	0.00	474
			x	0.65	0.67	0.13	0.25	0.79	0.68	0.67	0.16	0.96	0.00	474
	x	x	x	0.75	0.72	0.18	0.55	0.80	0.67	0.74	0.47	0.87	0.00	474

Notes: Numerical values are rounded. The baseline values are calculated assuming perfect knowledge about the test sample distribution, which means that the test sample mean is used for predictions. P-values relate to the significance level at which a model disagrees with its baseline model according to the McNemar test for 10,000 baseline prediction rounds. The significance levels are based on mean values.

Our evaluation metrics for models predicting process innovators have predominantly lower values than those predicting the product innovator status. Nonetheless, fitted models show for nearly all evaluation metrics better results than the process innovator baseline model and the McNemar test also confirms a significant difference. Hence, website characteristics still improve predictions. The best performance, in terms of accuracy, is reached by our 'all' feature model and leads to a performance increase of 10 percentage points. Moreover, 'meta' and 'network' features perform slightly worse than 'text' features.

The performance for innovators is slightly better than for process innovators in terms of AUC and accuracy. As the sample is slightly imbalanced towards innovators, this performance difference, however, is also partly related to different baseline values. Furthermore, similar to product innovators, we see remarkably higher AUC values of models including 'text' features. However, considering all other evaluation metrics, 'meta' features perform very similarly to 'text' features. Looking at F1-scores, predictions for the negative class always perform worse than the baseline model. In particular, the prediction solely based on 'network' features leads to zero F1-scores. This means the model predicts for every firm a likelihood that a firm is innovative larger than 0.5, which implies that the model always predicts the majority class. This is known as zero rule prediction. For applying this rule, the information included in our baseline model is sufficient. In this regard, 'network' features do not provide information gains for innovators. Looking at precision and recall (and not considering the 'network' feature model), we find general improvements for innovative firms in comparison to the baseline model. For non-innovative firms, we only find improvements in precision. Recall values, however, are very low and worse than in the baseline model.

Even though the number of observations is the smallest, the predictive performance as well as the performance increase for firms with innovation expenditures is the highest in terms of AUC and accuracy. Looking at the 'all' feature model, firms with innovation expenditures can be predicted with an AUC value of 75% and an accuracy of 72%, which corresponds to an accuracy increase of 18 percentage points. The model solely based on 'text' features performs even slightly better than the 'all' feature model considering accuracy. Besides, values of all other evaluation metrics are always better than random for the 'text' and 'all' feature model. Both models only using 'network' or 'meta' features show also strict improvements in accuracy and precision, but F1-scores and recall are partly worse than the baseline model.

Furthermore, the McNemar test confirms that all fitted models significantly disagree with baseline predictions. The divergence is always highly significant (p-values are below 0.001), except for models that predict process innovators with either 'meta' or 'network' features, which are significant at the 0.01-level. This may be

due to the fact that both feature sets as well as models predicting process innovators perform relatively worse. Hence, the difference to baseline predictions is especially low when combining both. It is also noteworthy that even though the McNemar test is significant, it does not necessarily mean that the model is strictly better than the baseline model. Key evaluation metrics also have to show predominately superior values. We want to highlight one example here. The Random Forest model that predicts innovators using 'network' features has a large share of inferior values in comparison to baseline predictions. It uses the zero rule for its prediction. Accordingly, it significantly disagrees with the baseline model as it uses another decision rule. However, the fitted model is not strictly better, because its decision is also solely based on the sample mean and the fitted model is not learning sufficiently from the provided features, as the evaluation metrics show.

Lastly, we want to note that we do not find a particular combination of hyperparameters across innovation indicators and feature sets that is always selected by the grid-search algorithm. However, preferred 'number of trees', 'maximum tree depth', and 'minimum impurity decrease' do exist across feature sets and target variables. For the 'number of trees', 1000 and 1500 are mostly chosen. The most dominant 'maximum tree depth' values are 50 and 100. Moreover, a 'minimum impurity decrease' of 0.001 is more frequently selected than 0.01. For more details see Table A.5.

To analyse the robustness of presented results, we re-estimate the 'all' feature model for each indicator using all possible combinations of splits between the training and test sample from 0.1/0.9 to 0.9/0.1 (in steps of 0.01). The corresponding change of respective AUC values with respect to an increasing training sample is displayed in Figure A.3. We find that AUC values for product innovators, process innovators, and innovators increase until a training sample size of 0.6 and then stay roughly constant at levels pointed out in Table 2.5. Hence, AUC values seem robust with respect to the sample split if a sufficiently large training sample size is reached. Besides, values fluctuate more strongly between 0.8 and 0.9, which is presumably related to a declining test sample size.

The performance of the model predicting innovation expenditures constantly increases until a training sample size of about 0.85. It has a comparatively large drop afterwards and tends to be more volatile in general. Both can be explained by a much smaller overall sample size for this indicator. For instance, a train/test split of 0.5 implies fewer absolute observations included in the training sample. Also, the test sample is always smaller, which makes the evaluation of the performance less robust. Furthermore, the increasing trend indicates that the model would have continued to improve, if we had added more observations. AUC values based on training sample sizes between 75% and below 85% fluctuate around the AUC value pointed out in Tables 2.5.

In summary, it can be stated that the analysed website characteristics show a better performance in the prediction of product innovators and firms with innovation expenditures than of process innovators. Moreover, text-based features show a greater relative relevance.

To compare the relevance of single features across feature sets, the ten most important predictor variables measured by the MDI are displayed in Figure 2.3 for each ‘all’ feature model, respectively.

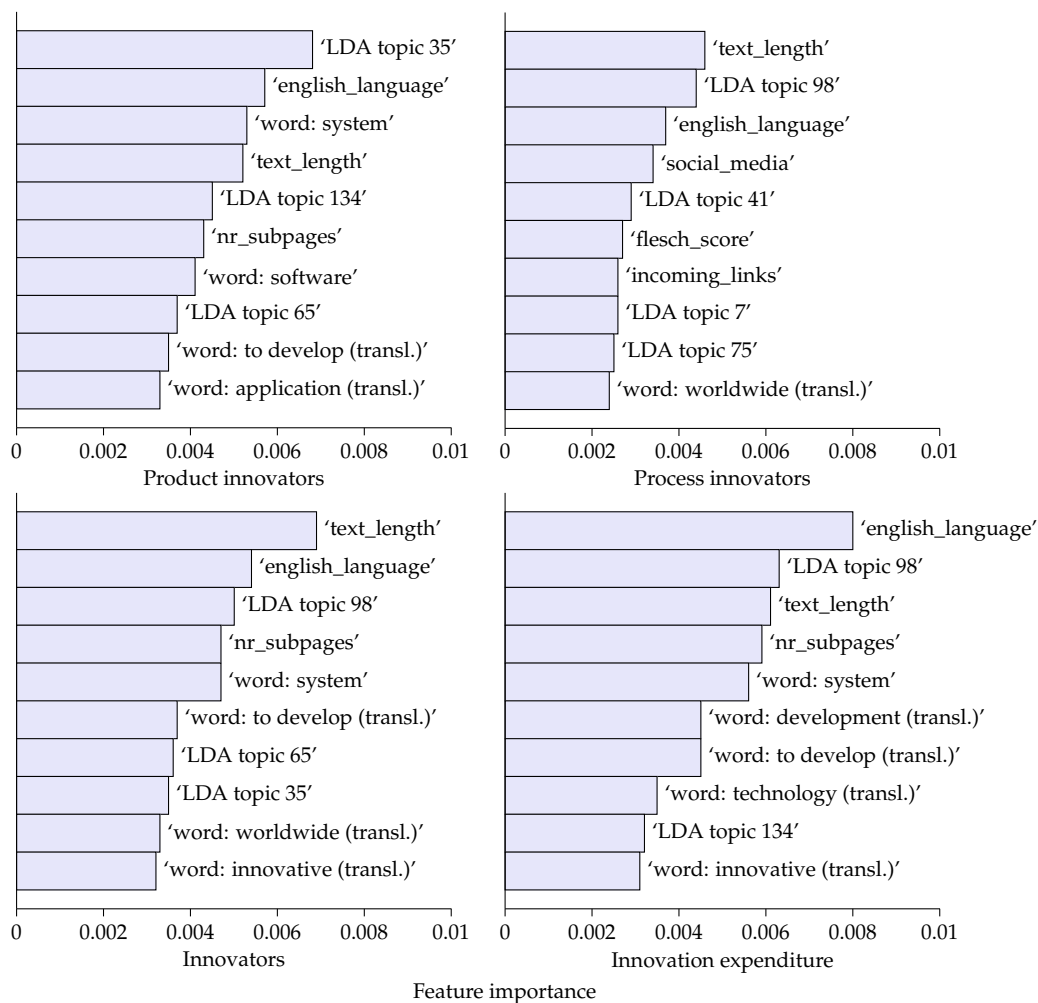


Figure 2.3: **Feature importance values for ‘all’ feature models.** For instance, a value that is two times larger implies that the mean decrease in impurity of the related feature is twice as high. Product innovators (top left), process innovators (top right), innovators (bottom left) and firms with innovation expenditures (bottom right) as target variable.

Three features exist that nearly always appear among the ten most relevant: the total number of characters (*text_length*), the number of subpages (*nr_subpages*) (this feature only appears on the twelfth position for process innovators), and the share of English language (*english_language*). A further investigation of the top 100 most

relevant features (see Appendix A.4) reveals that additional website characteristics exist with some general relevance. The words ‘worldwide’, ‘innovative’, ‘application’, ‘to develop’, ‘product’, ‘technology’ (all translated), the word ‘system’, as well as certain LDA topics, and the topic popularity index (*pop_score*), incoming (*incoming_links*), outgoing (*outgoing_links*), as well as social media hyperlinks (*social_media*), the Flesch-reading-ease score (*flesch_score*), the loading time of a website (*load_time*), and the share of numbers (*share_numbers*) are among the 100 most relevant features for every indicator. This shows that particular website characteristics exist, which have some relevance across indicators. In contrast, it is also noteworthy that features exist that show a large difference in the descriptive statistics but seem less important when predicting the innovation status. For example, the emerging technology term dummy never appears among the top ten features for any indicator and is also not frequently observed among the top 100 features. Furthermore, some features are more relevant for certain innovation indicators than for others. For instance, IT-related features seem to be highly relevant for product innovators. The IT-related LDA topics 35 (“ICT infrastructure”) and 134 (“business software”), as well as the words software and system are (only) among the top ten features for this indicator. Besides, LDA topic 7 with keywords linked to product experience and the word ‘application’ appear among the top 15 features.

On the contrary, the research & development related LDA topic 98 is more important when estimating process innovators and firms with innovation expenditure. Besides, the LDA topic 65 occurs in Figure 2.3 for product innovators and innovators, which should be related to a negative relationship to innovation activity, as the descriptive statistics show that this LDA topic is more likely to appear on websites of firms with no innovation activity. With respect to process innovators, it should be mentioned that only a single word can be found in the ten most important features and it is the only indicator that has ‘network’ features among its top ten. Furthermore, it is also interesting that the bottom left part of Figure 2.3, which relates to innovators, is at least for most features a combination of the most relevant features for product and process innovators. Last but not least, research & development-related words are highly important for predicting firms with innovation expenditures.

2.7 Discussion

Descriptive statistics as well as our fitted Random Forest models show that website characteristics are relevant predictors for firm-level innovation activity. We see a significant difference in means between innovative and non-innovative firms for most of our features. For each innovation indicator, Random Forest models using all features jointly show almost always a considerably higher performance than the baseline prediction with respect to the presented evaluation metrics. Moreover, the

McNemar test confirms a significant difference to baseline predictions for all models. Also, our results are in line with Kinne & Lenz (2021). Their statistical model has reached a similar accuracy for product innovators only observed in one MIP wave.

Our exercise also reveals – especially when predicting product innovators and firms with innovation expenditures – that ‘text’ features are relatively more important than ‘meta’ and ‘network’ features. Besides, we see a pattern regarding the most important characteristics that is independent of different target variables: Across indicators, the total number of characters, the number of subpages, and the share of English language belong always to the most relevant. It is also noteworthy that these features are more important than the word “innovative”. This finding suggests that website size and language should be considered for different types of website-based innovation indicators, which has not been done in previous studies. Meeting expectations, features that show insignificant differences in Table 2.3 almost never belong to the top ten most relevant features in Figure 2.3. An exception is the *flesch_score* in the case of process innovators. Furthermore, considering the poor performance of the ‘meta’ feature models and the result that ‘text’ is the most relevant feature set, the relevance of website size is quite counter-intuitive. One has to consider, however, that the importance of features is considered separately. The relevance of, e.g., the number of subpages is compared to the relevance of single words. If all words appearing in the term-document matrix would be considered jointly instead, their aggregated relative relevance would lie between 74 and 77%, depending on the indicator. This perspective illustrates why ‘text’ features and in particular textual content are still much more important for an accurate prediction. Nonetheless, as explained before, relative MDI importance should always be considered cautiously, as it is affected by multicollinearity. Moreover, other web-based features may exist that possess predictive power and have not been considered in our analysis. These features would most likely change the result.

Furthermore, it would also impact relative MDI importance, if this study’s web-site data would be complemented with information from other sources, for example, non-web data from the MUP. In this case, innovation activity could potentially be predicted more accurately. However, we have deliberately decided against adding non-web data to our analysis, since this study focuses on the comparison of website information, which is up-to-date and freely accessible for everyone. Nonetheless, it would certainly be interesting to investigate in a further study the effect of adding additional non-web data. For potentially relevant features, see Gandin & Cozza (2019).

Another aspect that we want to emphasise is the fact that features which are highly important for one particular indicator usually relate to its form of innovation activity. We see this as a strong indication that models use relevant information.

Especially for firms with innovation expenditures, the selected word-based features appear particularly convincing. Terms like “to develop” (transl.) and “technology” (transl.) are highly ranked and have a very strong and direct connection to research & development expenditures. Another example is that the product experience related LDA topic 7 (top 15 most important features) and the term ‘application’ have a high importance for product innovators. Additionally, the ten most relevant features of product innovators have a clear focus on information and communication (ICT) technologies, which is in line with the innovation spawning characteristic of ICT, as well as with the result of Hall et al. (2013). They find that ICT investment intensity is positively associated with innovation and is stronger linked to product than to process innovation.

Moreover, firms have a great incentive to present new products on their websites, process innovators, however, have a smaller incentive to announce innovation activity because new processes are less relevant for most website visitors. This might explain why results show a better predictive performance for product innovators than for process innovators and for innovators in general. In addition, only a single word appears among the ten most relevant features of process innovators and, even though this model differs on a higher significance level, ‘text’ features alone do only lead to slightly better predictions than ‘meta’ and ‘network’ features. This result supports the assumption that process innovations are often not mentioned explicitly.

Regarding innovators, most of the top ten features either appear in the product or process innovator ranking and the predictive performance of the ‘all’ feature model lies between both as well. This result meets our expectations as the innovator target variable is a combination of product and process innovators.

Interesting is also the fact that, contrary to our expectations, some features are not that relevant. For instance, even though the descriptive statistics show a large difference between innovative and non-innovative firms, the emerging technology dummy does not seem to be very decisive for predictions. Looking at the Pearson correlation coefficients between this and all other features reveals that the emerging technology dummy has a comparatively strong relationship with other features. Hence, their relative MDI importance is probably ranked lower due to multicollinearity. Besides, even though the descriptive statistics do not show a significant difference for every form of innovation activity, the Flesch-reading-ease score, the loading time of a website, and the share of numbers appear to be relevant for every indicator (according to the 100 most relevant features). These features, however, do not relate strongly to other features and might, therefore, provide some extra information. Hence, they are relatively relevant despite small differences.

Although we show a clear link between website characteristics and innovation

status, the predictive performance of our models leaves room for improvement as we, for example, still misclassify the existence of innovation expenditures for a considerable share of firms. Predictions might perform slightly better if neural networks were used. Our main criteria for choosing a Random Forest approach are the explainability of results and the fact that non-linear relationships can be learned. Neural networks unfortunately do not offer a direct possibility to disclose decision processes. Hence, there is a trade-off, which often occurs in practice, between performance and explainability. If explainability is not necessary, predictive performance can most likely be improved by neural networks.

Within our sample, there can be, of course, also innovative firms that do not mention their innovation activity (implicitly or explicitly) on their websites. In other words, some inaccuracy might relate to the nature of our data. In particular, product innovators, process innovators, and innovators might suffer from noise as they cover a three-year span. Websites can change a lot during this period. Comparatively good results for firms with innovation expenditures may be explained by the fact that this information is observed on an annual basis. Solving this matching problem seems to us a necessary step to improve predictions. Nonetheless, text data is always noisy and models with perfect accuracy are almost never identified.

Furthermore, it could be criticised that website-based innovation indicators can only be applied to firms that have a website. Another point of criticism would be that it could cause noise if firms falsely claim on their website that they are innovative, e.g., for marketing purposes. The MIP contains self-reported data as well, however, firms do not have the incentive to make false declarations, as answers should not affect their public image. For this reason, we expect MIP data to reveal the actual innovation status and we consider the usage of MIP-based information as target variables as a solution to the problem of false declarations of innovation activity on firm websites. Besides, patent data could have also been used as an alternative target variable. However, patent-based indicators rather measure inventions than innovations.

2.8 Conclusion

In this research article, we contribute to the discussion on whether web-based innovation indicators are a feasible alternative to survey-based innovation indicators. We conduct our analysis with data on 4,487 German firms, which reported different forms of innovation activity in a large-scale questionnaire-based survey (the MIP 2019). We extract website texts, additional website-related meta information, as well as hyperlinks of these firms and use the information to predict firm-level innovation activity reported in the MIP. The performance of our machine learning models shows that website characteristics unambiguously relate to MIP-based innovation

indicators. Furthermore, we find that website characteristics better predict product innovators and firms with innovation expenditures than process innovators. Hence, website characteristics rather appear to be suitable for measuring only certain aspects of innovation. Additionally, the importance of certain website characteristics varies between indicators. Accordingly, different features should be taken into account depending on the kind of innovation activity that is analysed. Lastly, our work and related studies show that state-of-the-art web-based predictive modelling cannot fully replace traditional surveys as error rates remain quite high. However, our models provide information about innovation activities that can be quickly updated, are on a very granular level (firm level), and are less expensive than questionnaire-based surveys.

Chapter 3

Firm Digitalisation and Mobility – Do Covid-19-Related Changes Persist?

joint work with Irene Bertschek, Patrick Breithaupt, and Daniel Erdsiek

3.1 Introduction

The adverse environmental impacts associated with transportation have continuously increased in recent decades and transportation accounted globally for about 23% of total energy sector's direct carbon emissions in 2019, with passenger cars being responsible for a large proportion thereof (IEA 2020). In 2020, however, the Covid-19 pandemic fundamentally interrupted mobility patterns. In order to prevent infections, digital technologies have been widely leveraged to avoid physical contact. For instance, remote access and virtual meetings allowed employees to work from home (WFH) and online shopping and delivery services enabled customers to purchase goods without leaving the house. Empirical evidence by Alipour et al. (2021) and Alcedo et al. (2022) indicates that these behavioural changes can be linked to lower levels of mobility around the beginning of the crisis. The observed decline sparked hope that the intensified use of digital technologies will result in persistent mobility reductions and a long-run decrease in associated carbon emissions. Comprehensive evidence on the relationship between digitalisation and mobility changes over the entire course of the pandemic, however, is missing. This study aims to close this research gap and quantifies the extent to which firm digitalisation effectively contributed to mobility reductions throughout the pandemic, covering the time frame from January 2020 until December 2022.

The long-term effects of the initial Covid-19 shock are not clear a priori. On the one hand, substantial investments in digital infrastructure and human capital, technological innovation, as well as a persistent increase in WFH arrangements give reason to believe that reductions in mobility are long-lasting (Barrero et al. 2021, Bloom

et al. 2021, Bachelet et al. 2022, Erdsiek & Rost 2022). On the other hand, after the initial Covid-19 shock, social isolation and physical inactivity while working from home may have been compensated by social meetings and an increase in mobility during leisure time. In particular, such compensatory behaviour may have been amplified by the decreasing severity of the pandemic and the lifting of restrictions in 2022. Also, working from home and improved access to services online increased the incentive to move further away from commercial districts, where rents are cheaper, potentially resulting in fewer trips taken but longer distances travelled (Marz & Şen 2022).

For our analysis, we make use of unique, web-scraped firm-level data for Germany, a large European country with an average level of digitalisation among EU countries (e.g., European Commission 2022). Moreover, we concentrate on firm digitalisation, which we approximate by estimating the extent to which firms communicate about digital topics on their websites. To this end, we apply a novel text-mining approach based on transfer learning.¹⁰ The indicator has the advantage that it contains, on the one hand, information on online and delivery services and, on the other hand, information relating to firms' potential of offering WFH, since it is likely that more tasks can be carried out remotely if firms have a high level of digital proficiency. Shortly before the start of the pandemic, we predicted the level of digitalisation for all 750,000 firms whose website addresses were available in the Mannheim Enterprise Panel (MUP).¹¹ The prediction was repeated in December 2022 based on 1,300,000 firms. For our analysis, we average firm-level predictions for German districts ("Kreise") to approximate the local economy's level of digitalisation. Furthermore, we use mobile network data to measure daily changes in mobility over the observed time frame (e.g., Persson et al. 2021).

We use an event study approach based on a difference-in-differences design to analyse how the link between mobility and firm digitalisation evolves over the course of the pandemic. Our regression results indicate a significant decrease in mobility associated with firm digitalisation for the time after the first lockdown up until the end of most Covid-19 restrictions in March 2022. After the lifting of most restrictions, however, ICT-related mobility reductions are no longer significant. The main contribution of our study is, thus, twofold. First, we show that firm digitalisation can indeed be leveraged to reduce physical travel during times of severe health threats. Secondly, however, if health threats and government restrictions ease up, the potential of digital technologies to reduce overall mobility is hardly exploited. This holds, even though factors that facilitate the substitution between physical travel and remote work or online services greatly improved during the pandemic. Hence,

¹⁰See Axenbeck & Breithaupt (2022) for a detailed description of the method.

¹¹The MUP is the most comprehensive micro database of German firms besides the official business register, which is not publicly accessible.

we cannot confirm long-term changes in mobility behaviour that might result in environmental improvements. Our results withstand an extensive set of robustness tests. We also contribute to the literature by using a novel text-based measure to approximate firm digitalisation and are able to cover a much longer time frame than previous studies analysing the effects of the Covid-19 crisis.

The remainder of this paper is structured as follows: In Section 3.2, we summarise related literature and derive our research question. In Section 3.3, we explain the data as well as the applied transfer learning approach. In Section 3.4, we present descriptive insights and in Section 3.5 econometric results. Section 3.6 includes robustness checks. In Section 3.7, we discuss our findings and conclude.

3.2 Related Literature

Whether telecommunications and travel are complements or substitutes has been a subject of intense debate for several decades (e.g., Kraemer 1982, Mokhtarian 1990). According to Kraemer & King (1982), telecommunications-transportation substitution depends on certain factors, such as transportation costs and the quality of telecommunication technologies. In contrast, Salomon (1986) puts forward the hypothesis that human beings have an intrinsic need for mobility, and thus travelling is unlikely to decline due to new technological opportunities, rather mobility patterns will change.

A strand of the empirical literature on ICT-enabled mobility reductions focuses on telecommuting. Remote work is highly relevant for the overall environmental debate, as a large proportion of the daily distance travelled is for work, and most people use their car to get there.¹² Conducting a meta-analysis of 39 empirical studies, Hook et al. (2020) find that the large majority of studies observe environmental improvements associated with telecommuting, which are mainly driven by a reduction in work-related trips. Considering also non-work travel, however, Wöhner (2022) only observes mobility savings for people that fully work remotely. People who only partly work from home completely offset saved commutes by an increase in non-work travel.

In most countries, working from home was only occasionally practised until the start of the Covid-19 pandemic, when the number of people working from home tremendously increased. For instance, the share of employees that fully work from home grew from 4 to 27% during the first lockdown in Germany (Emmler & Kohlrausch 2021). In addition, by June 2020, the share of firms using WFH increased from 48 to 74% in service industries and from 24 to 46% in the manufacturing sector

¹²For instance, between 27 and 47% of the daily distance travelled of employees in EU member states is for the purpose of work (Eurostat 2021); in EU countries, such as Germany and France, roughly two third of all workers use the car to get to their workplace (Destatis 2021, Insee 2021).

(Erdsiek 2021). Similarly, Brynjolfsson et al. (2020) find that regions with a higher share of employment in information work were more likely to shift towards WFH. Moreover, Alipour et al. (2021) confirm for the first weeks of the pandemic that a district's WFH potential can indeed be associated with overall mobility reductions, but only shortly before the first lockdown was put in place. Using the Global Survey of Working Arrangements, Aksoy et al. (2023) estimate that WFH approximately saved two hours of commuting per worker each week during the pandemic and will save one hour on average each week after the end of the pandemic. Bachelet et al. (2022) calculate for Germany that if 15% of all full-time employees in Germany will continue to work from home, 3% of carbon emissions attributed to the transport sector could be saved.

Besides WFH, online shopping may improve environmental outcomes. In theory, e-commerce can lead to mobility reductions, as orders can be consolidated and distributed more efficiently (e.g., Siikavirta et al. 2002, Durand & Gonzalez-Feliu 2012, Wiese et al. 2012).¹³ For example, Siikavirta et al. (2002) estimate that the maximum greenhouse gas emissions saving potential of e-grocery home delivery services is roughly between 0.3 and 1.3% for Finland. Using empirical data, Jaller & Pahwa (2020) find for Dallas and San Francisco that e-commerce has cut vehicle miles travelled by 7% on average in 2016, but highlight that environmental improvements depend on the modal split and are lower if people visit commercial districts by foot, bike, or public transportation instead of using the car (cf. Durand & Gonzalez-Feliu 2012, Wiese et al. 2012).

In contrast to WFH, online spending already substantially grew before the Covid-19 crisis (e.g., Alcedo et al. 2022). Nonetheless, the share of online spending in total consumer spending extremely increased at the beginning of the pandemic, jumping from roughly 17% to above 35% during the first two lockdowns. In mid-2022, however, the share of online revenue declined but remained above the pre-crisis level at roughly 24% (Alipour et al. 2022).¹⁴ Moreover, Alcedo et al. (2022) find for the first phase of the pandemic that online spending is positively linked to Google's index of residential activity at the country level, i.e., the approximated relative time spent at home, but the correlation declined until mid-2021.

In addition, empirical evidence shows that digitalisation supports firm resilience during economic crises (e.g., Bertschek et al. 2019, Reveiu et al. 2022). With respect to the Covid-19 crisis, Ben Yahmed et al. (2022) find that a region's digital capital relates to a lower level of short-time work usage at the beginning of the crisis. Also, firms in countries with a better digital infrastructure had comparatively higher revenue in

¹³Please note that we refrain from discussing further consequences for the environment of e-commerce that result, for instance, from frequent returns or additional packaging.

¹⁴Also, absolute online revenue increased during the first two years of the pandemic (bev 2022), indicating that the relative increase in online spending did not solely happen due to a decline in offline sales, but there must have been an additional shift towards online commerce.

2020 (Doerr et al. 2021). Comin et al. (2022) confirm that technological sophistication can be associated with higher sales at early stages of the crisis. Bertschek et al. (2022) show that the self-employed whose businesses were highly digitalised, benefited much more from the state aid provided by the German government during the pandemic compared to those whose businesses were less digitalised. By analysing World Bank data, Cariolle & Léon (2022) as well as Wagner (2021) show that having a website is related to firm survival during the pandemic. Moreover, Cariolle & Léon (2022) find that having a firm website is positively correlated with strategies that helped to cope with Covid-19 restrictions, such as home-delivery services, online sales, and remote work. Also, Bai et al. (2021), who observe that firm-level WFH feasibility can be associated with higher sales, net income, and stock returns during the pandemic, highlight the complementarity between digital technologies and WFH practices.

Although there are some insights into how digital strategies helped firms during the crisis, comprehensive empirical evidence is missing that accurately quantifies the extent to which firm digitalisation has effectively contributed to mobility reductions over the course of the pandemic, and, most importantly, whether changes have been sustained after most restrictions were lifted. Persistent reductions may exist because factors determining the substitutability between telecommunications and transportation greatly improved in order to cope with the pandemic (cf. Kraemer & King 1982), such as the technical quality of online communication due to large investments into digital infrastructure and a pandemic-driven surge in technological innovations (Barrero et al. 2021, Bachelet et al. 2022). For instance, the share of new patent applications that support WFH technologies more than doubled from January to September 2020 (Bloom et al. 2021). Moreover, geopolitical instability in Europe has driven gasoline prices extremely high in 2022, providing an additional incentive to replace fuel-based travel with digital solutions.¹⁵ The assumption of long-lasting mobility reductions is supported by surveys as well. These confirm that the share of people working from home did not largely decline after the pandemic became less severe in March 2022 (ifo Institute 2022, Aksoy et al. 2022).¹⁶ They also show that many customers anticipate doing more online shopping after the pandemic than before (Shaw et al. 2022).

However, in the light of an intrinsic human need for travel (cf. Salomon 1986), there is reason to believe that increased remote work and improved online access to services and products do not necessarily result in permanent mobility reductions.

¹⁵See <https://www.dashboard-deutschland.de> [online; accessed on 5 Jan. 2023].

¹⁶The Google Covid-19 Community Mobility Trends indicator suggests as well that less people visited their workplace after the pandemic than before (see <https://ourworldindata.org/covid-google-mobility-trends> [online; accessed on 5 Jan. 2023]).

For instance, as people have social and self-realisation needs, they may have enhanced social interaction after work over the course of the pandemic, especially after the crisis became less severe in March 2022 and the need to avoid physical contact declined. Furthermore, WFH increases the difficulty of managerial control (Felstead et al. 2003) and the stigma that people work less (efficiently) when they work remotely may not have diminished as much as expected during the pandemic. For instance, results from a questionnaire-based survey indicate that more than half of the firms did not noticeably change their assessment of remote work productivity (ZEW Mannheim 2022). As a consequence, employers still may prefer on-site or hybrid work arrangements because both facilitate monitoring employees as well as interaction. It could also be that a lack of supervision has the opposite effect, i.e., individuals who work from home are more efficient in order to work fewer hours. However, the resulting spare time may, in turn, also increase mobility.

In addition, working from home and increased access to online services may incentivise people to move further away from commercial districts, where rents are cheaper. This phenomenon can lower or even offset mobility reductions, as people may travel less often but longer distances (Marz & Şen 2022). Moreover, if workers are allowed to work remotely they usually can work from everywhere. Hence, long-distance trips, e.g., at weekends, become more appealing as it is possible to work while travelling. Finally, even though studies indicate that individuals buy more often products online than before the pandemic (Alipour et al. 2022, Shaw et al. 2022), people may prefer hybrid shopping modes and search for products offline and only buy them online.

Hence, whether a relationship between digitalisation and mobility reductions during different phases of the pandemic exists is a priori unclear. Therefore, this study analyses how the link between firm digitalisation and changes in mobility evolved over time in comparison to the pre-crisis level. In particular, we aim to shed light on whether changes in mobility persist after the lifting of most Covid-19 restrictions.

3.3 Data

For our analysis, we combine several data sources at the district level, which are all listed in Table B.1 in the Appendix.

3.3.1 Mobility

We use mobile network data provided by the German Federal Statistical Office (Destatis) for 400 German districts between January 2020 and December 2022 ([data

set] Destatis 2023a).¹⁷ The data stem from the Telefónica network and is processed by the company Teralytics before being forwarded to Destatis. Mobility is measured by the number of switches between mobile network cells per device within one district.¹⁸ The *change in mobility* (in %) is our variable of interest and captures the difference between mobility on a given day to the monthly average mobility in 2019 for the same weekday. For instance, switches between cells on the first Monday in September 2020 are compared to the average switches between cells on all Mondays in September 2019. We observe mobility changes for the entire day as well as for daytime and nighttime separately (6 a.m. to 10 p.m. and 10 p.m. to 6 a.m.).

The use of mobile network data has the advantage that it allows measuring mobility within precise, short time intervals. Nonetheless, we would like to acknowledge some of the shortcomings of our mobility indicator. For instance, it should be noted that the Telefónica network does not cover the entire mobile network market. As a result, we only observe changes in mobility for approximately one third of the German population, with varying market shares at the district level. To address this limitation, the data provider extrapolates the data to ensure representativeness.¹⁹ Furthermore, mobile network cells have an average size of 2.8 km to 4.8 km in rural areas and 0.7 km to 1.9 km in urban and suburban areas (Stobbe et al. 2023). Since changes in mobility can only be detected when there is a switch between mobile network cells, it is important to acknowledge that we are unable to observe a large portion of trips that are below these thresholds. Before the pandemic, however, the average distance travelled per day was 46 km, with an average distance of 12 km for a single trip (infas 2018). Therefore, we assume that we capture the majority of the daily distance travelled and consider this as a minor issue. One further limitation is that mobile network cells differ in size. The size mainly depends on population density since each cell can only handle a certain number of users. This makes it more difficult to capture changes in short distance trips in rural areas. To address this limitation, we control for the average population density and network quality in a district, as well as for whether a district is a city ("Stadt") or a countryside area ("Landkreis"). It is also worth noting that rural travel generally involves longer distances, which mitigates this limitation (infas 2018). Additionally, most of the German population lives in urban or suburban areas. As we weight our data based on population size, we further reduce the impact of this limitation.²⁰

¹⁷Please note that we consider "Wartburgkreis" and "Eisenach" as one district.

¹⁸In addition to mobile phones, also tablets, laptops, and vehicles can have SIM cards, which are removed from the analysis by approximation to avoid double counting.

¹⁹See [data set] Destatis (2023a).

²⁰It is possible that decisions for WFH and e-commerce are different in rural areas than in urban areas because of the greater distances that have to be travelled. However, due to the different mobile network cell sizes, we refrained from heterogeneity analyses that address this difference, as the described measurement problems could strongly distort the result here. However, since longer distances are travelled in rural areas, and mostly by car, mobility reductions presumably lead to greater carbon

3.3.2 Firm Digitalisation

To measure firm *digitalisation*, we take advantage of the fact that nowadays a large share of firms have a website, which usually provides insights into a firm's use of digital technologies, such as online shops, digital products, and social media. We collect these insights using a two-step transfer learning approach. The procedure is thoroughly described in the paper "Measuring the Digitalisation of Firms – A Novel Text Mining Approach" by Axenbeck & Breithaupt (2022).

Firstly, we train a text-based machine learning model that allows for automatically determining whether a text contains content on digitalisation. For this purpose, we exploit news articles, as we can easily identify whether they deal with digitalisation topics.²¹ This is because news articles appear within clearly defined sections, such as "business" and "politics". Also, news outlets can create special sections if a current topic is particularly relevant, such as "the digital transformation". We use these section titles as labels for supervised machine learning. Online articles have the additional advantage that their HTML code includes keywords for search engine optimisation (SEO), which also relate to the overarching subject of an article. Accordingly, we label all news articles appearing in a section about digitalisation or having the SEO keyword 'digital' embedded in their HTML code as digital and all other articles as non-digital.²² Based on the annotated newspaper corpus, we train a supervised machine learning model that allows for predicting whether a text is about digitalisation.²³ Secondly, we apply the fitted model to German firms, estimating the likelihood that their website text is about digitalisation. This is the transfer learning step of our text mining approach. The result is a continuous indicator measuring a firm's degree of digitalisation between zero and one. The entire procedure is illustrated in Figure 3.1.

We retrieve website addresses (URLs) from the Mannheim Enterprise Panel (Bersch et al. 2014). Using all available URLs and the ARGUS Web Scraping Tool (Kinne & Axenbeck 2020), we scraped 740,875 firm websites in January 2020 as part of the TOBI project.²⁴ Additionally, we collected information on 1,257,832 firm websites in December 2022 in order to have information covering the end of the observed time frame. In both years, we scraped up to 50 subpages of a firm website. As Kinne

emission savings here. For this reason, it would be beneficial for future research to investigate whether there are differences between urban and rural areas in terms of how digital technologies impact mobility changes.

²¹We use news articles from a large German newspaper corpus, which is described in detail in Axenbeck & Breithaupt (2022).

²²Moreover, we only consider articles before 2020, as articles related to the Covid-19 crisis might bias later firm-level predictions.

²³To this end, we use a Random Forest regression model suggested by Breiman (2001).

²⁴A research project on the potential of firm websites to measure technological progress funded by the German Federal Ministry of Education and Research (funding ID: 16IFI001).

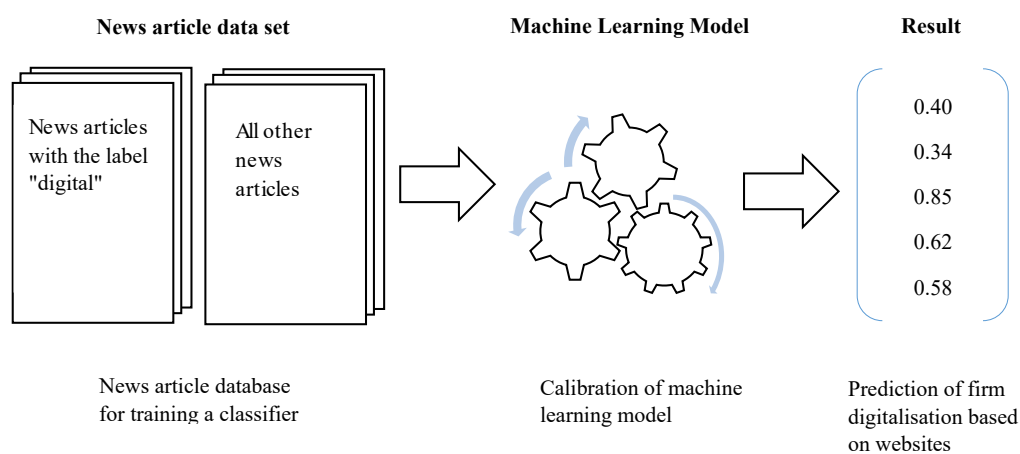


Figure 3.1: **Transfer learning approach for measuring firm digitalisation.** News article data with binary labels (left), machine learning model (middle), and continuous firm digitalisation scores based on scraped websites (right). Illustration from Axenbeck & Breithaupt (2022).

& Axenbeck (2020) show that the median number of subpages of a firm website is 15, we consider this threshold to be sufficient.

After applying the machine learning model to firm websites, we average and standardise predictions (mean zero and unitary standard deviation) for all firms in a district for both scraping periods, respectively. Regional distributions are displayed in Figure 3.2. It is apparent that firms in western Germany are more digital than those located in the eastern part of the country, which is plausible for historical reasons. Moreover, average firm digitalisation only slightly changed between both scraping periods.²⁵

Despite small changes, we consider firm digitalisation in levels and do not focus on the effect of changes in firm digitalisation over time in our main analysis. We do this for two reasons. Firstly, we face the issue that the scraping software has changed within the observed time frame, which affects the scale of the indicator and we cannot reliably compare changes. The second reason why we cannot simply compare the change in average firm digitalisation per district between both scraping periods is that the number of available firm websites notably grew.²⁶ In a robustness check in Section 3.6, we present a potential approach on how to consider the growing number of firm websites when analysing changes in firm digitalisation at the district

²⁵The Pearson correlation coefficient between both periods is 0.85.

²⁶The increase in available firm websites may have been due to more research effort being put into identifying web addresses by the data provider or because the number of firms that have a website has increased. Presumably, both factors played a part, but we assume the second reason had a larger impact, since establishing a firm website helped sustain business during the pandemic.

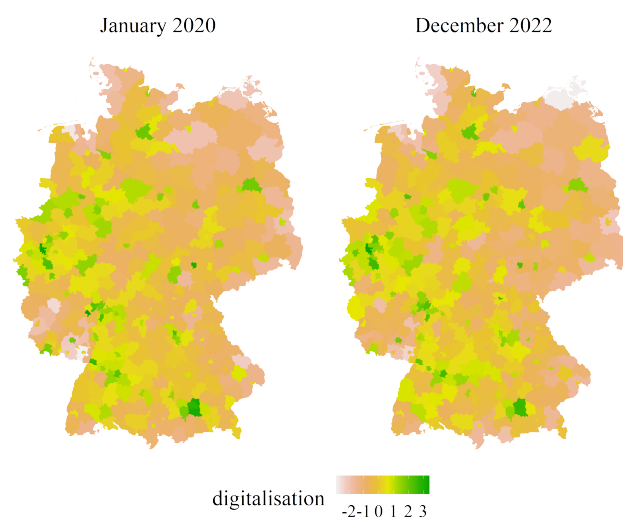


Figure 3.2: **Regional distribution in January 2020 and in December 2022 of the web-based firm digitalisation indicator.** Values are standardised (mean zero and unitary standard deviation).

level. The changing number of firms with a website also points us to the fact that we only observe the degree of digitalisation for firms that have a website. We address this issue in the same robustness check. Besides, a further minor issue is that most online services are available nationwide. However, we mostly observe small and medium-sized enterprises in our sample and it is likely that their online services are rather locally relevant, even though they are available across district borders.

3.3.3 Control Variables

Furthermore, as changes in mobility and firm digitalisation can be simultaneously affected by several factors, we add a broad variety of control variables.

3.4 Descriptive Insights

In the following, we provide descriptive insights into the relationship between firm digitalisation and changes in mobility, measured as the difference between mobility on a given day and the average monthly mobility in 2019 on the same day of the week.

In total, mobility increased by 1.02 % over the observed time frame (see Table B.2 in the Appendix, which also provides descriptive statistics of control variables).

Moreover, the increase in mobility is largely driven by a rise in short-distance travel below 30 km (see Figure B.1 in the Appendix).²⁷

In a next step, we examine how the link between firm digitalisation and mobility changes evolved over the course of the pandemic. Figure 3.3 presents a scatter plot displaying weekly changes in mobility in each district during the observation period. The dots are coloured based on average firm digitalisation in 2020. Greener dots indicate a higher average level of digitalisation.

Mobility increases in the first weeks of 2020 in comparison to 2019 for most districts. Moreover, districts that are more digital tend to be at the centre of the distribution and no clear correlation between firm digitalisation and changes in mobility is visible. We date the start of the first Covid-19 wave to March 22nd, 2020, as this is the day the first lockdown in Germany started and many businesses, such as restaurants and coffee places, had to close in order to slow down the spread of the virus. In the last week before the start of the lockdown, a large drop in mobility can be observed. The distribution also changes and districts that are more digital tend to be located at the bottom of the distribution, where the drop in mobility is the most pronounced. With the onset of the lockdown, the distribution is altered again and digital districts are gradually shifting back towards the centre of the distribution.²⁸ After the first shock, mobility slowly increases, however, districts that are more digital are moving back to the bottom end of the distribution. They remain there during the second lockdown when mobility starts decreasing again for most districts. In January 2021, as the number of Covid-19 infections did not decline, the German government implemented an additional obligation for employers to offer working from home to employees if feasible. However, according to Figure 3.3 mobility only declines for most districts at the very beginning of this first WFH obligation. Then, we observe a slightly increasing trend until the period of the first WFH obligation ends. Nonetheless, districts that are more digital still tend to show a decrease in mobility in comparison to 2019. After the end of the first WFH obligation, districts that are more digital do not show a clear reduction in mobility anymore.

Due to high infection rates, the WFH obligation as well as stricter restrictions on contacts were re-implemented during the last quarter of 2021. As shown in the corresponding part of Figure 3.3, however, mobility does not drop. The correlation between mobility changes and firm digitalisation remains but it is slightly smaller than during the first WFH obligation (perhaps because people took the second WFH

²⁷Please note that the increase in overall mobility could also be due to the expansion of the 5G network, which potentially involves smaller grid cells that could imply changes in the measurement of short-distance mobility over time. In the later analysis, we control for the network quality by considering the area in a district that is not covered or not covered by all network providers.

²⁸Alipour et al. (2021) observe a similar phenomenon with respect to a district's WFH potential and explain it by the strictness of confinement rules during the first lockdown that pushed people into short-time work when WFH was not possible. Hence, WFH feasibility may have played only a marginal role in reducing mobility during the first very strict lockdown.




Figure 3.3: **Average weekly change in mobility per district over the observed time frame.** The change in mobility measures the difference between mobility on a given day and the average monthly mobility in 2019 on the same day of the week. The dots are coloured based on the average degree of firm digitalisation in a district observed in 2020. The dots are plotted on top of each other so that districts that are more digital are more visible. The red line denotes the start of the first lockdown on March 22nd, 2020. White areas mark periods with no or few Covid-19 restrictions, grey areas mark lockdown periods and blue areas mark periods where the government-imposed WFH obligation was additionally in place. We observe gaps when missing data exist at the beginning of the week.

obligation less seriously). In spring 2022, the severity of the pandemic lessened and the second WFH obligation ended on March 20th, 2022.²⁹ Also, many other Covid-19 restrictions were lifted around that date, such as restrictions on contacts for unvaccinated people. It is apparent that after the end of these restrictions, the correlation between changes in mobility and firm digitalisation continues to decline.³⁰

Figure B.2 in the Appendix shows average changes in mobility for different phases of the pandemic with respect to quintiles of firm digitalisation observed in 2020 as well as in 2022 and a district’s WFH potential calculated by Alipour et al. (2023) and used in Alipour et al. (2021), respectively. Average mobility changes barely differ across the corresponding quintiles of a district’s WFH potential and both firm digitalisation indicators. Moreover, higher quintiles of all three indicators can unambiguously be linked to lower levels of mobility between the first open period and the end of the second WFH obligation, whereas no clear pattern is visible during the pre-pandemic phase, the first lockdown, and after the end of restrictions. Differences may exist with respect to different phases of the pandemic because firm digitalisation only represents a potential to reduce mobility that can be leveraged if needed. As the pandemic’s severity as well as the strictness of restrictions fluctuated greatly over time, the mobility-reducing potential of digital technologies may have been fully realised only during specific periods of the pandemic.

3.5 Econometric Approach and Results

We conduct an event study based on a dynamic difference-in-differences (DiD) design with two-way fixed effects and clustered standard errors at the district level to provide inferential statistical insights into how the relationship between firm digitalisation and changes in mobility evolved. In the context of the Covid-19 pandemic, similar approaches at the regional level have been conducted by Alipour et al. (2021), Ben Yahmed et al. (2022), and Alipour et al. (2022). The link between firm digitalisation and mobility changes is modelled as follows:

$$\begin{aligned} \Delta \text{mobility}_{i,t} = & \sum_{m \neq \text{Feb}'20} \beta^m (\text{digitalisation}_i \times \text{year-month}_m) \\ & + \sum_{m \neq \text{Feb}'20} \sum_{c \in C} \gamma_c^m (c_{i,t} \times \text{year-month}_m) + \text{year-month}_m \quad (3.1) \\ & + \text{district}_i + u_{i,t}. \end{aligned}$$

²⁹This day is referred to as the German Freedom Day by many media outlets.

³⁰The outlier at the bottom of the distribution during the end of the observed time frame is “Jena”, which is a small district with a digital hub that is characterised by a large level of emigration. See <https://www.zeit.de/gesellschaft/grossstaedte/jena-bevoelkerungsentwicklung-zuwanderung-abwanderung> [online; accessed on 5 Jan. 2023].




Figure 3.4: **Monthly change in mobility associated with digitalisation (estimated β^m coefficients)**. Our reference period is February 2020 (grey line). The red line denotes the start of the first lockdown on March 22nd, 2020. White areas mark periods with no or few Covid-19 restrictions, grey areas mark lockdown periods and blue areas mark periods where the government-imposed WFH obligation was additionally in place. Confidence intervals are at the 90% significance level.

Changes in mobility relative to 2019 are observed for district i on day t . β^m represents the change in mobility in month m for a given year (measured as the difference to our reference period which is February 2020) that is related to firm digitalisation. We focus on firm digitalisation observed in January 2020, as this measure is exogenous to the onset of the Covid-19 crisis.³¹ γ_c^m captures the parallel varying trend of control variable c in control variable set C for each month.³² We include year-month fixed effects to control for common shocks that affect all districts simultaneously. Moreover, we incorporate district-level fixed effects to address potential confounding factors that result from unobserved differences between districts. In addition, observations are weighed based on their population size.

Digitalisation did significantly impact changes in mobility over the course of the pandemic as shown in Panel A of Figure 3.4, which displays results only with digitalisation interacted with time dummies, as well as year-month and district-level fixed effects. Importantly, the β^m coefficient is close to zero and insignificant in January 2020, i.e., the month before the reference period. Thus, we do not observe a notable

³¹A threat to endogeneity would be, for example, if more online services have emerged as a response to the onset of the crisis in districts with a greater adherence to social distancing. This could be the case because online services were more important for reaching customers in these areas.

³²Please note that not all control variables vary over time. Moreover, we also interact the weekly incidence rate and the containment measure index with time dummies, even though they are time-varying. We do this because the sensitivity to the incidence rate as well as to Covid-19 restrictions most likely changed over time.

pre-trend.³³ After the first Covid-19 outbreak, the effect size starts to increase until it peaks in the last quarter of 2020. In October 2020, the average reduction in mobility is roughly 6.6 percentage points (pp) for every standard deviation of digitalisation. One explanation for the increase in ICT-related mobility reductions in the first months of the pandemic is that the first lockdown was very strict and many people had to stay at home anyway. Therefore, differences in firm digitalisation may have had a less visible impact on mobility reductions at the beginning of the crisis. In the subsequent summer months, incidence rates were low and people worked less due to the holiday season. Hence, the mobility-reducing potential of digital technologies might not have been fully exploited during this period. In autumn 2020, however, incidence rates increased again, but restrictions were less severe than during the first lockdown. In consequence, differences in firm digitalisation may have become more critical for changes in mobility. Moreover, the effect size may also have increased during the initial months, as digital capacities that allow for social distancing had to be built, such as online sales channels as well as VPN and fast internet connections (e.g., Barrero et al. 2021, Bloom et al. 2021). We assume that firms which already had a certain digital proficiency prior to the crisis had advantages in this regard (cf. Cariolle & Léon 2022).

In 2021, the effect size slightly decreases and levels off at around roughly -4 pp until the end of the year. Moreover, during the second WFH obligation, the effect size further declines. After lifting most restrictions in March 2022, the effect size continues to diminish and becomes insignificant for most months.

When we allow for differential time trends of control variables (displayed in Panel B of Figure 3.4), digitalisation coefficients still tend to be negative and significant for most months of the first two years of the pandemic, however, the effect size is notably smaller. Surprisingly, we find a significantly positive effect in April 2020 (the middle of the first lockdown period). Alipour et al. (2021) observe a very similar phenomenon when analysing the link between a district's WFH potential and changes in mobility at the beginning of the crisis and controlling for covariates. The authors explain the positive link by the strictness of the first lockdown, in which many employees were put on short-time work and many (rather non-digital) establishments had to close in order to avoid contagion. This reasoning is confirmed by Ben Yahmed et al. (2022), who show that regions with lower digital capital had higher short-time work usage rates at the beginning of the crisis. Thus, during this

³³As our pre-crisis time frame is very short, we cannot thoroughly verify whether the assumption of parallel trends, on which our event study is based, is fulfilled. To address this issue, we conduct a robustness check by re-estimating our model for the beginning of the crisis using weeks instead of months (results are presented in Figure B.3 in the Appendix and also discussed in Section 3.6). Also, in the specification that uses weeks, we do not observe a significant pre-trend between firm digitalisation and changes in mobility. Hence, we can assume that the parallel trends assumption holds.

time, the widespread use of short-time work and closed factories might have reduced mobility especially in low-digitalised regions. This situation at the beginning of the pandemic attenuates the association between a district's mobility reductions and its digitalisation level and can even lead to a positive coefficient. In fact, the sign of our focal coefficient is only reversed, if we include time-varying effects of socioeconomic and demographic characteristics that might partially capture the effect of a district's feasibility to work from home on mobility (see Table 3.1, discussed in the next paragraph).

In Panel B, we observe the maximum realised mobility reduction in September 2020, in which the differential decrease is 3.2 pp for every standard deviation of digitalisation. Moreover, after the end of the second WFH obligation, the effect size declines and becomes insignificant. Hence, we also find diminishing effects after the end of restrictions when we condition on differential time trends of control variables. Thus, we do not find evidence for long-lasting environmental improvements.

In a next step, we re-estimate Equation (3.1) but summarise differential time trends by the different phases of the pandemic. Table 3.1 shows that the digitalisation coefficient for the post-pandemic phase is always insignificant, independent of the considered control variables. Furthermore, the coefficient of the post-pandemic phase is always less negative than the coefficients of previous phases in which digitalisation can be unambiguously linked to a decrease in mobility. The difference to these previous phases is predominantly significant at the 10% threshold. Thus, this specification also strongly indicates that the effect size diminishes in the post-pandemic phase.

In the last step, we estimate the average change in mobility that can be linked to firm digitalisation during the two years of the pandemic (from March 22nd, 2020 to March 19th, 2022), considering all covariates. Column (1) of Table B.3 in the Appendix displays a negative and significant digitalisation coefficient, indicating that firm digitalisation can on average be associated with mobility reductions during the crisis. The effect size is -1.68 pp. Column (2) and Column (3) show results for daytime and nighttime mobility separately. We also find a significantly negative link between firm digitalisation and daytime mobility changes, but the effect size is much smaller and insignificant for nighttime mobility, which is highly plausible as most people only work and engage in commercial activities during daytime. Column (4) presents findings only for working days and Column (5) only for weekends estimated with daytime mobility changes. We find that the coefficient is slightly smaller at weekends, which is plausible as people usually do not work during these days, but commercial activities may continue. Another reason for a decrease in mobility

Table 3.1: DiD results providing insights into changes in the link between mobility reductions and firm digitalisation for different phases of the pandemic.

	dependent variable: Δ mobility						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
digitalisation (Jan '20)							
× (1) 1st lockdown	-2.079*** (-3.68)	-2.050*** (-3.87)	1.096+ (1.79)	-0.340 (-0.50)	1.289+ (1.89)	-0.315 (-0.51)	1.732** (2.76)
× (2) 1st open period	-3.920*** (-8.70)	-3.902*** (-8.97)	-2.479*** (-3.58)	-2.776*** (-4.50)	-1.851** (-2.80)	-2.843*** (-5.35)	-1.504* (-2.17)
× (3) 2nd lockdown/ 1st WFH o.	-4.895*** (-13.72)	-5.142*** (-14.17)	-2.914*** (-5.20)	-2.817*** (-5.30)	-2.214*** (-4.14)	-3.328*** (-8.10)	-1.728** (-3.05)
× (4) 2nd open period	-3.837*** (-7.34)	-4.046*** (-7.72)	-3.177*** (-3.82)	-3.056*** (-4.02)	-3.319*** (-4.17)	-3.408*** (-5.73)	-2.237** (-2.69)
× (5) 2nd WFH obligation	-3.028*** (-6.83)	-3.296*** (-7.66)	-2.352*** (-3.44)	-1.738* (-2.54)	-3.487*** (-4.92)	-3.186*** (-5.67)	-2.207** (-3.19)
× (6) post-pandemic	0.242 (0.40)	-0.0244 (-0.04)	0.502 (0.55)	0.612 (0.63)	-1.542 (-1.57)	-0.276 (-0.35)	-0.291 (-0.30)
year-month fixed effects	x	x	x	x	x	x	x
district-level fixed effects	x	x	x	x	x	x	x
pandemic controls		x					x
socioeconomic controls			x				x
infrastructure controls				x			x
demographic controls					x		x
geographic controls						x	x
observations	433999	433999	433999	433999	433999	433999	433999
R ²	0.57	0.58	0.58	0.57	0.59	0.58	0.62
$\beta^1 = \beta^6$	0.00	0.00	0.52	0.34	0.00	0.96	0.03
$\beta^2 = \beta^6$	0.00	0.00	0.00	0.00	0.67	0.00	0.08
$\beta^3 = \beta^6$	0.00	0.00	0.00	0.00	0.32	0.00	0.02
$\beta^4 = \beta^6$	0.00	0.00	0.00	0.00	0.00	0.00	0.00
$\beta^5 = \beta^6$	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Notes: β^m coefficients of Equation (3.1) estimated using OLS. t statistics in parentheses. Clustered standard errors. + $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Observations are weighted based on their population size. The time frame is split into different phases as presented in Figure 3.3. The pre-Covid-19 phase is used as a reference period. Fixed effects for every phase are additionally included. The table also includes t-tests for the equality of coefficients. (1) no control variables; (2) only controlled for the pandemic situation; (3) only controlled for socioeconomic characteristics; (4) only controlled for characteristics that relate to a district's infrastructure; (5) only controlled for demographic characteristics; (6) only controlled for geographic characteristics; (7) all control variables included.

associated with digitalisation on weekends could be that people spread their working hours over the entire week when working from home and also work at weekends (e.g., McDermott & Hansen 2021).

3.6 Robustness

Our event study relies on the assumption that no difference in mobility changes with respect to firm digitalisation between districts would have occurred if the Covid-19 outbreak did not happen. To explore this issue, we re-estimate Equation (3.1) for the period between January 7th, 2020 and May 4th, 2020, but analyse weekly instead of monthly differences. We use the week before Shrove Monday 2020 as a reference period because the first large-scale Covid-19 outbreaks occurred in Germany as part of the carnival festivities in 2020. Figure B.3 in the Appendix shows that no statistically significant difference in mobility changes with respect to firm digitalisation exists for the weeks before the first large-scale outbreaks.

As firms heavily invested in digital infrastructure in the course of the pandemic, one reason for insignificant effects in the post-pandemic phase could be that firm digitalisation changed to such an extent in the observed time frame that we do not find an effect in 2022 if we consider firm digitalisation observed in 2020. To explore this issue, we conduct the same event study but with firm digitalisation measured in 2022. Figure B.4 and Table B.4 in the Appendix show that changes in the degree of firm digitalisation during the pandemic do not appear to cause the diminishing effect size, as the results are generally very similar to our main results.

As stated above, one drawback of our web-based digitalisation indicator is that we only observe the degree of digitalisation for firms that have a website. To address this issue, we modify the way we average the degree of firm digitalisation at the district level. To this end, we assume that there are $J + K = N$ firms in a district i . Firm $j \in J$ has a website and firm $k \in K$ does not. We conjecture that firms without a website address available in the MUP do not have a website, i.e., they are part of set K . Moreover, we suppose that these firms have a lower degree of digitalisation than firms with a website and set their digitalisation score to zero. Since firm digitalisation of firms in set K is zero, we divide the sum of our web-based predictions by the total number of firms in a district to adjust our indicator for firms without a website (see Equation [3.2]):

$$\begin{aligned} \text{digitalisation}_i &= \frac{\sum_{j_i \in J_i} \text{digitalisation}_{j_i} + \sum_{k_i \in K_i} \text{digitalisation}_{k_i}}{N_i} \\ &= \frac{\sum_{j_i \in J_i} \text{digitalisation}_{j_i}}{N_i}. \end{aligned} \quad (3.2)$$

Column (1) and Column (2) of Table B.5 in the Appendix show results with respect to different phases of the pandemic, including all control variables for modified firm digitalisation observed in 2020 and in 2022, respectively. The coefficients of both modified digitalisation indicators point in the same direction as in the estimation with unmodified firm digitalisation. For the indicator observed in 2020, we find mostly insignificant effects. For 2022, however, we observe weakly significant results until the end of the second open period and no significant effect in the post-pandemic period.

Despite the fact that the scraping algorithm has changed between both scraping periods and the size of predictions at the firm level is not reliably comparable between both scraping periods, we also provide insights about the extent to which results differ if we consider the change in predictions instead of the level of firm digitalisation. To this end, we calculate the standardised difference between modified firm digitalisation observed in 2020 and in 2022. This allows us to consider changes in the degree of digitalisation for firms that have a website as well as to take the increase in the number of firms with a website into account. Column (3) in Table B.5 displays that coefficients point into familiar directions when considering changes in firm digitalisation at the district level, but only the coefficient for the first open period shows a weakly significant negative effect.

Column (4) of Table B.5 shows coefficients of household broadband availability as a proxy for the level of household digitalisation. Firm digitalisation is excluded from the estimation. In this specification, coefficients are predominately insignificant, indicating that household broadband availability is not as relevant for changes in mobility as firm digitalisation. Column (5) and Column (6) show results of unmodified firm digitalisation in 2020 and in 2022 but with districts not weighted by their population size. Digitalisation coefficients are comparable to our main results.

Moreover, in our main analysis, we conjecture that firm digitalisation leads to a decrease in mobility via remote work and e-commerce. To provide some evidence in favour of this mechanism, we analyse whether our web-based digitalisation indicator can indeed be associated with increased firm-level remote work and e-commerce at the onset of the crisis. For this purpose, we use the Mannheim Innovation Panel (MIP) in 2021,³⁴ in which German firms were asked about the percentage of employees that worked from home before the pandemic as well as during the first and second lockdown, and whether they increased e-commerce activities at the beginning of the crisis. Merging the MIP 2021 with the MUP allows us to analyse this information for 3014 firms. Results with respect to remote work are displayed in Table B.6 in the Appendix. We find that the share of employees that work from home

³⁴See [data set] ZEW – Leibniz Centre for European Economic Research (2021).

increased by 6 to 7 pp if firm digitalisation observed in 2020 is one standard deviation larger. Results with respect to e-commerce are provided in Table B.7 in the Appendix. We also observe that firms which are more digital are more likely to have expanded digital products, services, and sales channels with the onset of the crisis.

Furthermore, we analyse how firm digitalisation relates to the link between a district's WFH potential and mobility reductions. In Table B.8 in the Appendix, we replace firm digitalisation by a district's capacity to work from home according to Alipour et al. (2023).³⁵ Our results indicate that mobility reductions at daytime are also related to a district's WFH potential. However, the link becomes insignificant when we additionally include firm digitalisation, whereas the latter remains significant. This finding suggests that firm digitalisation better explains variation in changes in mobility than a district's WFH potential, which either could be because firm digitalisation is the more precisely measured variable or because firm digitalisation has further impact channels such as e-commerce.³⁶

3.7 Discussion & Conclusion

Given the climate crisis and sharply rising energy costs, discussions on the exploitation of the mobility-reducing potential of digital technologies are repeatedly part of the public debate. It is generally believed that the Covid-19 pandemic has accelerated the utilisation of this potential. We contribute to the discussion by quantifying the actual extent to which firm digitalisation can be linked to mobility reductions from January 2020 to December 2022. Using German data at the district level and considering a broad variety of control variables, we find that mobility decreased on average by 1.68 pp in comparison to 2019 for every standard deviation of firm digitalisation during the first two years of the pandemic. We observe the largest mobility reductions associated with digital technologies in the last quarter of 2020. During this period mobility decreased up to 3.2 pp for every standard deviation of firm digitalisation if we control for differential trends of covariates and up to 6.6 pp if we do not. Moreover, we observe that the effect size diminishes and becomes insignificant after most Covid-19 restrictions were lifted in March 2022, suggesting no long-lasting environmental improvements.

This result raises the question of why ICT-enabled mobility reductions declined after the end of most restrictions. In Section 3.2, we hypothesise that even though

³⁵For a description of the variable, see Appendix B.1.

³⁶Also, we looked into the robustness of results with respect to spatial correlation and estimated a spatial Durbin regression model, including spatial lags of our dependent variable and independent variables. Our results are robust in the sense that we find a diminishing effect in the post-pandemic period and no significant effect of spatially lagged firm digitalisation. Results can be retrieved from the authors upon request.

factors that promote telecommunications-transportation substitution improved during the pandemic, mobility reductions can diminish for several reasons, which we now discuss individually. A mentioned reason that explains why mobility increased again, is that managers still prefer employees to work on-site as it facilitates supervision and interaction. However, we consider this reason rather unlikely as questionnaire-based surveys confirm that a large share of employees still work remotely. The same applies to the argument that, after the pandemic, people take long-distance trips more often. We do not consider this to be a significant reason because after the pandemic long-distance travel is roughly at the same level as during the second open period, in which digitalisation can still be linked to mobility reductions (see Figure B.1 in the Appendix). An additional argument is that people move further away from their working place where rents are cheaper because they have to commute less frequently. However, if people would move to another district a notable positive correlation between firm digitalisation and changes in in-commuters should exist. In fact, we find a slightly negative correlation of -0.04 if we do not weigh by population size and only a small positive correlation of 0.08 if we do. Another argument why this reason is not substantial is that if people move away, ICT-enabled mobility reductions should decline gradually, but both, the descriptive and the econometric analysis, strongly suggest that the relationship changed at a certain point in time, namely when most restrictions were lifted. Figure B.1 in the Appendix also reveals that short-distance travel increased to a greater extent than long-distance travel. Hence, it could also be that an expansion of the 5G network caused an overall increase in mobility, as smaller grid cells allow us to observe short-distance mobility at a more granular level. However, this issue only affects the link between digitalisation and changes in mobility if the 5G expansion correlates with our measure of firm digitalisation. If this is the case, ICT-enabled mobility reductions should also tend to gradually decline, but, as stated above, the relationship changes rather abruptly.

The compensation of social and self-realisation needs during leisure time as well as a preference for hybrid shopping remain as possible reasons for diminishing ICT-enabled mobility reductions, which both point to Salomon's (1986) argument that people have an intrinsic need for mobility that prevents them from comprehensive telecommunications-transportation substitution. Since this is likely to be the case, we conclude that it is more desirable from a societal perspective to promote green, carbon neutral mobility patterns, than advocating the replacement of physical travel by digital solutions.

Chapter 4

Digital Technology Adoption and Energy Intensity in Manufacturing – Firm-Level Insights

joint work with Thomas Niebel

4.1 Introduction

The pressing need to slash global carbon emissions in combination with more recent disruptions to energy markets have spotlighted the importance of reducing the energy intensity of the manufacturing sector.³⁷ At the same time, the use of digital technologies has strongly increased in recent decades, with ICT exerting significant impacts on how firms produce and provide goods and services, not least given the ever-wider adoption of the Internet of Things and big data analytics (see, e.g., Brynjolfsson et al. 2021). Most likely, digital technologies have also changed the energy use patterns of manufacturing firms and will continue to do so in the future. However, how digital technologies influence energy intensity is ambiguous. On the one hand, ICT adoption may increase firm-level energy intensity, as digital technologies consume electrical power. On the other hand, despite their power demand, such technologies may reduce overall firm-level energy intensity due to energy efficiency improvements and dematerialisation (Berkhout & Hertin 2004). For instance, digital technologies improve the quantity and quality of information, which allows for an improved prevention of excess production and a reduction in error rates.

Despite the ambiguous impact of ICT on energy consumption, a wide range of industrialised countries have launched programs to promote smart manufacturing, such as the German “Industrie 4.0” and the US “Smart Manufacturing Leadership

³⁷In 2020, for instance, the manufacturing sector accounted for 28.5% of energy demand in Germany (German Environment Agency 2021).

Coalition” (SMLC) (Thoben et al. 2017), assuming large potentials for more sustainable production.³⁸ Similarly, with its “New Industrial Strategy for Europe”, the EU is vigorously promoting a “twin green and digital transition” (European Commission 2021a). However, digital technologies could either contribute to or impair the achievement of climate targets set forth by the Paris Agreement (see UNFCCC 2015). Therefore, it is highly important from a policy perspective to accurately assess whether digital technologies and energy are actual substitutes within production processes, and if so, whether the relationship is substantial enough to drive sustainability.

By employing aggregated data at the industry level, previous econometric studies find strong evidence for ICT-energy substitutability within economic sectors (Schulte et al. 2016, Taneja & Mandys 2022) as well as ICT-related electricity intensity improvements in manufacturing industries (Bernstein & Madlener 2010). However, the use of aggregated data has several drawbacks. For instance, it is not possible to observe whether actual substitution between inputs within firms takes place or whether energy intensity decreases due to a change in the composition of firms and associated products over time. Furthermore, as production processes differ, digital technologies may have varying energy-saving potentials depending on the industry (Bernstein & Madlener 2010). As a consequence, different effects across industries can result in a heterogeneity bias (see Imbs & Mejean 2015, Campello et al. 2019). Solow (1987) states that the question of substitutability between capital and energy can only be satisfactorily answered at the micro level, as aggregation will bias results.³⁹ This argumentation also applies to ICT capital. The aim of this paper is, therefore, to measure the substitutability between digital technologies and energy at the firm level and to answer to what extent relationships change when more granular data is considered. Additionally, we aim to identify characteristics that can explain differences between observational levels.

To the best of our knowledge, no large-scale microeconomic study exists yet that analyses the link between digital technologies and energy intensity. We contribute to the literature by filling this gap, investigating administrative panel data on 28,600 German manufacturing firms (AFiD)⁴⁰ collected between 2009 and 2017 and provided by the Research Data Centres of the Statistical Offices of the Federation and the Federal States (RDC). AFiD data are of particular high quality, as reporting to the statistical offices is obligatory and the data is thoroughly checked.⁴¹ We use

³⁸The use of sensors, computing platforms, communication technology, control and simulation methods, data-intensive modelling, and predictive engineering within production processes is subsumed under the term smart manufacturing (Kusiak 2018).

³⁹Also, see Koetse et al. (2008) and Haller & Hyland (2014).

⁴⁰Amtliche Firmendaten für Deutschland.

⁴¹See Berner et al. (2022), Kube et al. (2019), and Richter & Schiersch (2017) for highly regarded studies on environmental issues employing AFiD data.

firm-level software capital stocks as an indicator of ICT usage, which is a commonly used indicator in firm-level studies of digitalisation (see, e.g., Almeida et al. 2020, Bessen & Righi 2020, Barth et al. 2022).⁴² In comparison to other digital technologies, such as 3D printing or cloud computing, software capital has the advantage that it is a comprehensive indicator for the firm-level degree of digitalisation, as almost all hardware requires additional software. Also, respective investments are observed by German official statistics in a panel format and in monetary terms, which allows for a thorough calculation of capital stocks.

By focusing on energy intensity as the outcome of interest and applying a translog cost function approach for the econometric analysis, we enhance comparability to previous findings at the industry level. We find that a one per cent increase in software capital relates to a decrease in energy intensity of 0.003% at average relative energy costs, i.e., for the average firm. Hence, we also find evidence for ICT-energy substitutability – however, at a much smaller magnitude than previous industry-level studies suggest. Our results are robust to different sample restrictions as well as software capital stock modifications. To further analyse the robustness of firm-level results, we conduct a reduced form estimation with a selection of variables based on a CES production function. The respective result leads to the same conclusion, which is that the use of digital technologies cannot be linked to substantial energy intensity improvements at the firm level, even though substitutability is suggested.

Moreover, we contribute to the literature by showing that the link between ICT and energy intensity entails properties that can lead to biased estimates at the aggregate level. We find that effects are heterogeneous across industries and that the link is rather statistically significant in industries that are more energy-intensive. We explain this phenomenon by pointing to the larger incentives to reduce energy consumption in these industries (whether using digital or non-digital technologies). As differences are systematic in the sense that they relate to different levels of energy intensity, they can potentially bias results at the industry level. Moreover, we find larger differences between than within firms. Hence, firms with a high software capital stock are on average less energy-intensive, but when the software capital stock changes within a firm, effects have a much smaller magnitude. This phenomenon can result in a large omitted variable bias if firm characteristics are not appropriately considered. Additionally, it can contribute to an aggregation bias if one does not adequately control for changes in the composition of firms and associated products. Both issues are highly relevant for future studies on ICT-energy substitutability as they point to problems that can arise from aggregation.

⁴²We use the term digitalisation synonymously to digital transformation.

The remainder of this paper is structured as follows: Section 4.2 summarises related literature and Section 4.3 presents theoretical frameworks. Section 4.4 describes the data and provides descriptive statistics. Our results are reported in Section 4.5 and discussed in Section 4.6. Section 4.7 concludes.

4.2 Related Literature

ICT adoption may influence overall energy use and intensity in various ways.⁴³ Using aggregated data to measure energy intensity improvements within manufacturing and service industries, studies come to mixed results but tend to support the hypothesis that digital technologies are associated with a decrease in energy intensity. Using a CES production function, Collard et al. (2005) investigate the relationship between ICT and energy use in the French service sector from 1986 to 1998. The authors find that electrical energy intensity decreased with the diffusion of communication devices, while it increased with the use of computers and software. Applying the same approach, Bernstein & Madlener (2010) analyse the impact of ICT capital on electrical energy intensity in five manufacturing industries and eight EU countries from 1991 to 2005. Even though the effect seems to depend on the industry-specific production processes, the authors conclude that the diffusion of ICT is generally linked to electricity intensity improvements.

Analysing 27 industries in ten OECD countries between 1995 and 2007 and using a translog cost function approach, Schulte et al. (2016) conclude that there is strong evidence for substitutability between ICT capital and energy. Additionally, a sample split into the manufacturing and service sector shows only significant effects for the manufacturing sector. Employing a translog cost function approach on aggregated data, but analysing disaggregated ICT capital with quantile regressions, Taneja & Mandys (2022) confirm economically relevant substitution behaviour for 13 countries and 28 industries for the same time frame.⁴⁴

While an empirical link between digital technologies and energy intensity seems to exist, questions surround the nature of this relationship. Berkhout & Hertin (2004), Hilty et al. (2006), and Lange et al. (2020) develop frameworks that posit potential impact mechanisms by which ICT influences environmental outcomes. Based on these frameworks, the net effect consists of three different channels.⁴⁵

⁴³Energy intensity measures the actual amount of energy used to generate one unit of output, not necessarily considering differences in prevailing conditions, e.g., the type of product or local weather (cf. IEA 2022), whereas “energy efficiency is a generic term” and “refers to using less energy to produce the same amount [of output]” (Patterson 1996, p. 377).

⁴⁴We have to acknowledge that Taneja & Mandys (2022) find no significant effect for software capital when controlling for other ICT equipment.

⁴⁵Please note that within the framework of Lange et al. (2020), there are four different channels. Instead of behavioural and structural effects, growth effects and tertiarisation effects exist.

(I) Direct (Berkhout & Hertin 2004, Lange et al. 2020) or first-order effects (Hilty et al. 2006) relate to the energy and resource consumption during the production, usage, and disposal of ICT. Accordingly, direct effects have a negative environmental impact and increase energy and resource use.⁴⁶

(II) Indirect (Berkhout & Hertin 2004), second-order effects (Hilty et al. 2006) or energy efficiency improvements (Lange et al. 2020) refer to changes in consumption due to the application of digital technologies. Due to improvements in energy efficiency as well as substitution by dematerialised solutions, digital technologies have the potential to decrease energy intensity.⁴⁷ For example, big data and artificial intelligence allow for an improved prediction of demand and may prevent excess production. They also help to reduce error rates. Simulation methods as well as 3D printing may drastically reduce resource and energy use associated with the design and development of new products (OECD 2017a, IEA 2017). Hence, even though digital technologies consume energy, they can have a positive net effect on the firm, especially if digital systems serve as a substitute for rather than a complement to existing solutions (Berkhout & Hertin 2004).

(III) Structural and behavioural impacts (Berkhout & Hertin 2004) or third-order effects (Hilty et al. 2006) describe fundamental changes associated with the use of digital technologies. For instance, a decrease in overall energy use due to energy efficiency improvements is only possible when these are not largely offset by rebound effects. Moreover, structural and behavioural impacts have no clear direction of impact. For example, additional consumption resulting from ICT-induced economic growth may lead to increased energy and resource consumption, while shifts to less energy-intensive products and services may contribute to environmental improvements (see Lange et al. 2020).

These frameworks illustrate the complexity of the relationship between digital technologies and environmental improvements.⁴⁸ Accordingly, accurately identifying the magnitude of ICT-related energy savings within production processes is not trivial. According to Schulte et al. (2016), substitutability between ICT and energy is determined by the net effect between an energy use effect and an energy efficiency-enhancing effect (comparable to direct vs. indirect effects within the presented frameworks). The more aggregated the data, however, the more difficult it becomes to disentangle direct and indirect effects from structural or behavioural impacts. By employing industry-level data, for instance, it is not possible to determine

⁴⁶For examples of findings on the energy consumption of YouTube, see Preist et al. (2019); for the cryptocurrency Bitcoin, see Stoll et al. (2019), Corbet et al. (2021), and Jones et al. (2022); and for data centres, see Masanet et al. (2020).

⁴⁷For instance, see Zhang et al. (2019), Ghobakhloo & Fathi (2021) or Friedrich et al. (2021) for studies that qualitatively discuss ICT-enabled energy savings in manufacturing.

⁴⁸Not without reason do studies on overall trends come to different conclusions. See Table 1 in Chimbo et al. (2020).

whether actual changes in inputs are taking place or whether only the composition of firms and associated products is changing (Solow 1987).⁴⁹

Despite arguments that technological substitutability between capital and energy can only be satisfactorily measured at the micro level (Solow 1987, Koetse et al. 2008, Haller & Hyland 2014), empirical studies at the firm level analysing the link between ICT usage and energy intensity are scarce. While considering the few firms that apply industrial robots in order to assess changes in coal consumption, Huang et al. (2022) find improvements in coal intensity. Further firm-level insights are based on questionnaire-based surveys with non-technical self-assessments. For instance, in a survey conducted in 2020, 1,700 German manufacturing and service-sector firms were asked about measures in the areas of energy efficiency and digitalisation (Bertschek et al. 2020). Energy savings was the least frequently named reason for conducting ICT projects. Moreover, most manufacturing firms stated that their absolute and (relative) ICT-related energy use remained constant during the last three years. The largest study in this regard was conducted on behalf of the European Commission in 2021. For this purpose, 10,006 firms were interviewed. In this survey, firms confirmed that improving the environmental footprint is not the dominant motivation for implementing digital technologies. Nonetheless, 70% of all firms reported energy savings due to their usage.

4.3 Theoretical Frameworks

Previous industry-level studies employing production function approaches to measure the link between ICT and energy intensity apply a translog cost function approach, such as Schulte et al. (2016) or Taneja & Mandys (2022), or, alternatively, a CES production function approach, such as Collard et al. (2005) or Bernstein & Madlener (2010). In our firm-level analysis, we also focus on a translog cost function approach for two reasons. Firstly, translog cost functions have the advantage that they are more flexible than CES production functions, i.e., they make no restrictive assumptions on estimated substitution behaviour and on the optimal path of input factor adjustments induced by price changes (Christensen et al. 1973, Koetse et al. 2008, Wurlod & Noailly 2018). Secondly, previous industry-level estimates employing a translog cost function are not only limited to electricity but to overall energy. We, therefore, believe that the translog cost function approach is better suited for a direct comparison between different levels of aggregation.

However, in comparison to the CES production function formulated by Collard et al. (2005), the employed translog cost function has the disadvantage that energy intensity is not used directly as a dependent variable. To provide confidence that we

⁴⁹Moreover, further aggregation issues can occur (see , e.g., Imbs & Mejean 2015).

are measuring actual changes in energy intensity, we also estimate a reduced form CES production function in the later analysis to show that results derived at the firm level are robust to the two theoretical approaches. We will explain both – the translog cost function and the CES production function approach – in the following.

4.3.1 Translog Cost Function

The applied dual translog cost function approach is based on the seminal work of Christensen et al. (1973), Berndt & Wood (1975), Brown & Christensen (1980), and Berndt & Hesse (1986). Schulte et al. (2016) were the first to adjust a translog function to measure substitution behaviour between ICT and energy. In the spirit of Shephard (1953), Diewert (1971), and Berndt & Wood (1975), Schulte et al. (2016) apply the duality theorem. Accordingly, a cost function instead of a production function is estimated. The cost function is twice differentiable, linearly homogeneous, concave in factor prices, and corresponds to a given level of output Y . Different forms of capital are considered as quasi-fixed factors and materials as weakly separable, which results in a restricted variable cost (VC) function that depends on the following parameters:

$$VC = f(P_E, P_L, K_{ICT}, K_N, Y, t). \quad (4.1)$$

E indicates energy, L labour, and P respective prices. K_{ICT} relates to ICT capital and K_N to tangible (non-ICT) capital. Disembodied technological change is captured by time t . Variable costs consist of energy and labour costs ($VC = P_E E + P_L L$).

Equation (4.1) can be approximated by a translog specification (Schulte et al. 2016):

$$\begin{aligned} \ln VC = & \beta_0 + \beta_Y \ln Y + \frac{1}{2} \beta_{YY} \ln(Y)^2 + \beta_T t + \frac{1}{2} \beta_{TT} t^2 + \sum_k \beta_k \ln P_k \\ & + \sum_m \beta_{K_m} \ln K_m + \frac{1}{2} \sum_k \sum_l \beta_{kl} \ln P_k \ln P_l + \frac{1}{2} \sum_m \sum_n \beta_{K_m K_n} \ln K_m \ln K_n \\ & + \sum_k \beta_{kY} \ln P_k \ln Y + \sum_m \beta_{K_m Y} \ln K_m \ln Y + \sum_k \sum_m \beta_{kK_m} \ln P_k \ln K_m \\ & + \sum_k \delta_{kT} \ln P_k t + \sum_m \delta_{K_m T} \ln K_m t + \delta_{YT} \ln Y t, \end{aligned} \quad (4.2)$$

with $k, l \in \{E, L\}$ and $m, n \in \{ICT, N\}$. Applying Shephard's lemma, assuming symmetry ($\beta_{EL} = \beta_{LE}$), and homogeneity of degree one ($\beta_{EL} = -\beta_{EE}$) allows estimating the following equation (see Christensen et al. [1973] and Berndt & Wood [1975]), in which the share of energy costs in variable costs, S_E , is a function of the

energy price relative to the labour price, time, output, as well as ICT and non-ICT capital intensity:⁵⁰

$$\begin{aligned} \frac{\partial \ln VC}{\partial \ln P_E} = \frac{P_E E}{VC} = S_E = & \beta_E + \beta_{EE} \ln \left(\frac{P_E}{P_L} \right) + \beta_{EK_{ICT}} \ln \left(\frac{K_{ICT}}{Y} \right) \\ & + \beta_{EK_N} \ln \left(\frac{K_N}{Y} \right) + \beta_{EY}^* \ln Y + \delta_{ET} t. \end{aligned} \quad (4.3)$$

Under the zero profit condition, total costs equal the production price, P_Y , times output. The elasticity of energy intensity with respect to ICT capital can be obtained by approximating total costs by variable costs ($VC \approx P_Y Y$). Then, multiplying S_E by P_Y/P_E allows observing that $E/Y = (P_Y/P_E) S_E$ (cf. Welsch & Ochsens 2005, Ma et al. 2008, Wurlod & Noailly 2018). Using this property enables the calculation of the elasticity for energy intensity, assuming no effect of software investments on prices:

$$\begin{aligned} \epsilon_{E/Y, K_{ICT}} &= \frac{\partial \ln(E/Y)}{\partial \ln K_{ICT}} = \frac{\partial (E/Y) Y}{\partial \ln K_{ICT} E} = \frac{\partial ((P_Y/P_E) S_E) Y}{\partial \ln K_{ICT} E} \\ &= \frac{\partial S_E}{\partial \ln K_{ICT}} \frac{Y P_Y}{E P_E} = \beta_{EK_{ICT}} \frac{P_Y Y}{P_E E} = \beta_{EK_{ICT}} \frac{1}{S_E} = \frac{\beta_{EK_{ICT}}}{S_E}. \end{aligned} \quad (4.4)$$

4.3.2 CES Production Function

To test the robustness of results derived by the translog model, we estimate a reduced form of a CES production function in the spirit of Collard et al. (2005) and Bernstein & Madlener (2010). This approach has the advantage that energy intensity is directly considered as a dependent variable. It is a nested approach with 3-inputs (KL; E) and constant returns to scale, in which output is generated by:

$$Y = \left[\omega \{AE\}^{\frac{\sigma-1}{\sigma}} + (1-\omega) \{F\}^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \quad (4.5)$$

with $\omega \in [0, 1]$. Tangible capital and labour are combined to form a composite input F . $\sigma > 0$ captures the elasticity of substitution between energy and the capital/labour nest. A denotes the energy-related level of technology, which evolves as function of ICT capital relative to tangible capital and disembodied technological change:

$$\ln A = \theta_0 + \theta_{ICT} \ln \left(\frac{K_{ICT}}{K_N} \right) + \theta_t t. \quad (4.6)$$

⁵⁰ $\beta_{EY}^* = \beta_{EY} + \beta_{EK_N} + \beta_{EK_{ICT}}$; Schulte et al. (2016) scale capital by output to be consistent with literature that measures effects of ICT on labour and output. Consequently, β_{EY} has to be modified to β_{EY}^* .

By assuming perfect competition and price-taking behaviour of firms, which implies that the unit of the cost function gives the price of output, we obtain the following equation after applying Shephard's lemma, taking logarithms, rearranging, and plugging Equation (4.6) into Equation (4.5):⁵¹

$$\ln\left(\frac{E}{Y}\right) = \sigma \ln(\omega) - \sigma \ln\left(\frac{P_E}{P_Y}\right) + (\sigma - 1)(\theta_0 + \theta_{ICT} \ln\left(\frac{K_{ICT}}{K_N}\right) + \theta_{it}). \quad (4.7)$$

Accordingly, energy intensity can be expressed as a function of the energy-related level of technology and the energy price relative to the production price. The equation is later estimated in reduced form.

4.4 Data

Our analysis focuses on firm-level data on the German manufacturing sector (AFiD) collected between 2009 and 2017 and provided by the RDC. Within our data, firms are assigned to the manufacturing sector if they have the highest value added in associated industries.

4.4.1 Data Sources

We combine two AFiD data sets merged by internal identifiers from the RDC:

- (A) The “AFiD Panel Industrial Units”, which contains two sub-data sets that are relevant for our analysis:⁵² The Census on Investment includes information about investments in tangible and intangible assets. It is a full census covering all German firms in the manufacturing sector with 20 employees or more. From this survey, we retrieve our indicator for the firm-level degree of digitalisation, which is software capital. Information on software investments is available from 2009 onward. We include information on investments in property, plant, and equipment from 2003 onward. This allows considering investments in tangible assets before the observation period and improves the calculation of respective capital stocks.⁵³ The second applied sub-data set is the Cost Structure Survey. It contains comprehensive annual information at the firm level about produced output as well as inputs, such as energy costs, labour costs, and the number of employees. The Cost Structure Survey is a stratified, partly

⁵¹Please note that we follow Van der Werf (2008) and Lagomarsino (2020) by considering the logarithmised ratio between the energy and output price.

⁵²See [data set] Research Data Centres of the Statistical Offices of the Federation and the Federal States (2019a).

⁵³Software investments have a very high depreciation rate. Therefore, not observing such investments before the observation period is not a substantive issue.

rotating panel. Firms with 500 employees or more are fully covered in the survey, whereas firms with fewer employees are generally observed for at least four consecutive years if they are surveyed.⁵⁴

- (B) The “AFiD Module Use of Energy” contains detailed information about the use of different energy sources at the plant level.⁵⁵ The data set is also a full census including all German manufacturing plants with 20 employees or more. For information on firm-level energy use, we aggregate plant-level information for each firm.⁵⁶

Additionally, we add information from several data sources. We combine AFiD with gross value added deflators from Eurostat at the two-digit industry level (NACE Rev. 2 classification) to calculate real output. Annual software deflators are also taken from Eurostat. This allows us to consider real software investments and thus to take into account quality improvements in software. EU KLEMS data is added (also at the two-digit industry level) to receive information about capital growth rates, depreciation rates, as well as tangible capital deflators. The data are also supplemented by the yearly producer price index maintained by the German Federal Statistical Office (Destatis) as well as information on the prices of different energy carriers. For a detailed overview of supplementary data, see Table C.1 in the Appendix.

4.4.2 Variable Description

Based on the raw data described in Section 4.4.1, we conduct the following additional calculations. We define overall firm-level energy use (E) as the sum of the energetic use of different energy carriers plus electricity use (in kWh).⁵⁷ The descriptive statistics in Table 4.1 show that mean energy use is above 30 GWh and that the median fluctuates around 2 GWh. Hence, the distribution of energy use is highly skewed; some firms consume far more energy than the large body of firms. Energy costs ($P_E E$) can be directly retrieved from the Cost Structure Survey. Furthermore,

⁵⁴Accordingly, our entire observation period can be divided into three sequences (2009-2011, 2012-2015, 2016-2017).

⁵⁵See [data set] Research Data Centres of the Statistical Offices of the Federation and the Federal States (2019b).

⁵⁶One minor drawback is that we do not observe the firm units that are assigned to the service sector. Hence, when we observe software investments, it may be that they were made in a service sector unit and we cannot observe corresponding changes in energy use in that unit. However, service sector sites consume a much smaller fraction of energy compared to plants in the manufacturing sector. We also do not expect to see large differences in the degree of digitalisation within firms, as digitalisation projects are most likely implemented for entire companies.

⁵⁷We consider the following energy carriers: biomass, natural gas, coal, heating oil, district heat, liquid gas, and the category “other energy sources”. Additionally, we subtract self-generated electricity by means of the listed energy carriers from electricity use to avoid double counting when calculating overall energy use.

the analysis requires information on energy prices, which are not directly available in AFiD. Following Haller & Hyland (2014), we divide energy costs by energy use (E) to receive information on the energy price for each firm (P_E ; in €/kWh). The energy price borne by most firms is between 0.02 and 0.20 €/kWh, which seems plausible considering industry prices for different energy sources.⁵⁸

Gross wages and salaries, statutory contributions, and other social insurance costs are summarised to receive information on labour costs ($P_L L$). The amount of full-time equivalents (L) is measured by the total number of employees adjusted for part-time employees. In the analysed time frame, firms employ slightly more than 270 full-time equivalents on average. The yearly wage is derived by dividing labour costs by full-time equivalents. For hourly wages, we adjust values by the average yearly hours worked in 2016 in German manufacturing.⁵⁹ The average hourly labour price (P_L) is €29.⁶⁰

Variable costs (VC) are calculated based on the sum of energy and labour costs. S_E measures the share of energy costs in variable costs; S_L the share of labour costs. The average share of energy costs in variable costs is around 0.09, which is comparable to the average industry-level share derived by Schulte et al. (2016). Output (Y) is measured by real value added based on information specified in the Cost Structure Survey and deflated using Eurostat data at the two-digit industry level.⁶¹

Software capital (K_{SW}) approximates the firm-level degree of digitalisation and tangible capital, i.e., property, plant, and equipment, represents the non-software capital stock (K_N).⁶² Please note that we only account for purchased software capital and firms may also use software that is free of charge. We deflate capital stocks based on Eurostat (software) and EU KLEMS (non-software) data. Furthermore, the perpetual inventory method (PIM) is applied to estimate capital stocks (Griliches 1980, Berlemann & Wesselhöft 2014, Lutz 2016, Löschel et al. 2019, Dhyne et al. 2021a). If calculated correctly, PIM allows us to measure the total productivity-relevant capital by considering previous investments and depreciation rates alongside current investments.⁶³ Moreover, PIM requires assumptions about initial capital stocks, which

⁵⁸To control for outliers, we exclude the highest and lowest percentile with respect to the energy price from our analysis. The resulting price distribution is displayed in Figure C.1 in the Appendix. See Figure C.2 in the Appendix for a comparison with the average energy price calculated using prices for different energy sources (if available) from official statistics.

⁵⁹See <https://iab.de/en/daten/iab-working-time-measurement-concept/> [Online; accessed on 11 Apr. 2023].

⁶⁰The value is slightly higher in statistics adjusted for the overall population (<https://www-genesis.destatis.de/genesis/online?language=en&sequenz=tabelleErgebnis&selectionname=62431-0001>) [Online; accessed on 11 Apr. 2023].

⁶¹We do not subtract energy costs to calculate value added, as capital, energy, and labour (KLE) are part of the optimisation problem in the later analysis. Materials are considered as weakly separable.

⁶²Leasing capital is excluded.

⁶³The depreciation rate of software capital in our preferred specification is 31.5%, as in EU KLEMS. We also calculate an average depreciation rate for non-software capital based on EU KLEMS.

are calculated based on average investments in the first three observation periods as well as depreciation and capital growth rates. Consequently, we only consider observations that are observed at least three years in a row.⁶⁴ Our calculated capital stocks confirm the findings of Kaus et al. (2020), who analyse tangible and intangible capital within the German manufacturing sector. Software capital (as a form of intangible capital) is growing faster in comparison to tangible capital. Furthermore, both distributions of respective investments are heavily skewed and lumpy, but software investments show these characteristics to a greater extent. For instance, we find approximately 25% of firms without any software investments in the analysed period. Accordingly, we add €1 to every software capital stock, as this allows us to take the logarithm when software capital stocks are zero.⁶⁵ To evaluate whether the estimated software capital stocks are a sufficient proxy for the firm-level degree of digitalisation, we compare our results with the Survey on the Use of Information and Communication Technologies in Companies (ICT survey, 2012 – 2017), which is a stratified random sample and contains more detailed information on ICT usage.⁶⁶ Figure 4.1 shows mean software capital intensity, i.e., the amount of software capital used to generate one unit of output, for firms in which at least 20% of employees use a personal computer (PC) and for firms in which less than 20% use a PC. Firms have a much higher software capital intensity when at least every fifth employee uses a PC. Figure 4.2 illustrates software capital intensity by the firm-level maximum data transmission rate, i.e., the internet speed. The figure shows that the higher the Mbit/s range, the higher the mean software capital intensity. Consequently, a clear relationship between software capital and the use of other digital technologies exists.⁶⁷

Additionally, the following control variables are included in the analysis. We add federal state dummies, industry dummies at the two-digit level, and dummies capturing different size classes, measured by the number of employees. We consider six different size classes: (1) 20 to 49 employees, (2) 50 to 99 employees, (3) 100 to 249 employees, (4) 250 to 499 employees, (5) 500 to 999 employees, and (6) 1,000 or more employees. Moreover, we generate a dummy indicating whether a firm has multiple units. By means of the electric energy consumption and the ratio of electric energy costs to value added, we approximate whether firms receive a full or a partial

⁶⁴For a detailed description of PIM see C.2.

⁶⁵We consider this issue in various robustness checks presented in C.6.1.

⁶⁶The ICT survey is additionally provided by the [data set] Research Data Centres of the Statistical Offices of the Federation and the Federal States (2019c). We are able to match 16,813 observations from our sample with the ICT survey. Unfortunately, different questions are asked every year and there is a large share of missing values, so the number of observations is much lower for each survey item.

⁶⁷See C.3.4 for an analysis of whether industry-level and regional differences with respect to software usage are plausible.

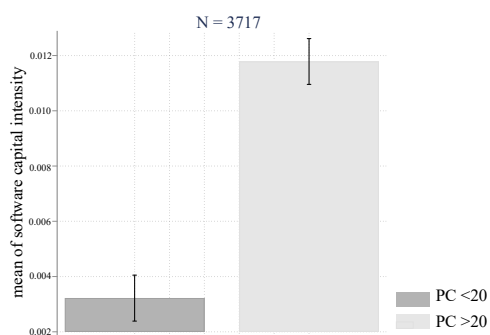


Figure 4.1: Software capital intensity by firms' PC usage. The brighter grey relates to firms in which 20% of all employees or more use a computer. Error bars relate to confidence intervals at the 95% level.




Figure 4.2: Software capital intensity by maximum data transmission rate. Error bars relate to confidence intervals at the 95% level.

exemption from the EEG levy.⁶⁸ Accordingly, we include two dummies relating either to a full or partial exemption. Also, a dummy that controls for whether a firm produces energy is included, as this may affect energy costs as well. Last but not least, we include a dummy that is set to one if a firm is trading commodities.

Although AFiD data are the cornerstone of many official German governmental statistics and several plausibility checks are conducted by Destatis, we find small shares of implausibly small or high values. To address this, we trim our sample by the labour and energy price at the 1th and 99th percentile, and winsorise all growth rates at the 0.1th and the 99.9th percentile. We also exclude firms with zero labour, energy or non-software capital use, as well as firms with a negative output. Additionally, we exploit the panel structure to identify outliers and exclude firms for which the standard deviation relative to the median of input-output ratios as well as labour and energy prices is higher than 100.

4.4.3 Additional Descriptive Statistics

After the described preprocessing steps our sample includes 123,362 observations, 28,600 firms in total, and on average about 13,700 firms per year (see Table C.2 in the Appendix). Around 13% of these firms have multiple units. Moreover, we apply the first-difference estimator in the subsequent statistical analysis. This reduces our main estimation sample to 89,653 observations.

⁶⁸A special surcharge used to support the expansion of renewable energy. The government grants exemptions to some energy-intensive firms in order to avoid harming their international competitiveness.

Table 4.1: Summary statistics of selected variables.

variable description	(1)			(2)		
	sample statistics in levels			$100 \times \Delta \ln^*$		
	mean	median	sd	mean	median	sd
E energy use in kWh	33,331,227.84	2,021,331.63	405,383,675.97	2.23	1.76	25.28
$P_E E$ energy costs in €	1,848,518.09	245,000.00	14,004,206.33	3.81	2.93	29.79
L full-time employees	273.15	88.50	1,957.65	1.28	0.86	11.88
$P_L L$ labour costs in €	16,269,208.85	3,778,850.00	158,115,569.30	3.80	3.75	11.46
P_E energy price per kWh in €	0.13	0.11	0.09	1.57	1.32	35.98
P_L hourly labour price in €	28.74	27.99	9.17	2.52	2.46	11.24
K_{SW} software capital in €	258,458.48	13,389.11	2,703,507.41	18.05	0.00	133.26
K_N tangible capital in €	20,110,799.14	3,011,534.63	204,487,549.41	2.11	-1.50	18.06
Y output in €	22,929,312.43	5,060,662.00	213,216,121.19	3.82	3.31	27.48
Y/L output per employee in €	65,881.09	57,474.03	42,847.85	2.54	1.94	29.16
E/Y energy intensity in kWh/€	1.07	0.38	3.86	-1.59	-1.70	36.24
K_{SW}/Y software capital intensity	0.01	0.0024	0.06	14.37	-9.56	136.68
K_N/Y tangible capital intensity	0.93	0.54	4.50	-1.64	-3.33	33.02
VC variable costs in €	18,117,726.95	4,190,796.00	165,630,044.11	3.71	3.62	11.35
S_L labour cost share	0.91	0.94	0.10	0.05	0.02	2.24
S_E energy cost share	0.09	0.06	0.10	-0.05	-0.02	2.14
Observations		123,362			89,653	

*Please note that all change rates are in per cent except those of S_E and S_L , which are not logarithmised. Respective change rates are, therefore, denoted in percentage points.

An overview of the mean, median, and standard deviation for selected variables can be found in Table 4.1. Values are also presented for annual change rates ($100 \times \Delta \ln$).

Column group (1) displays the sample statistics in levels. Mean software capital intensity is 0.01 and median software capital intensity is 0.0024. In comparison, mean tangible capital intensity is 0.93 and median tangible capital intensity is 0.54.⁶⁹

Column group (2) shows descriptive statistics for growth rates. Mean growth rates for energy use, labour use, tangible capital, and software capital are positive.⁷⁰ Hence, the absolute use of input factors grows over time at the firm level. Moreover, the growth rate for output is also positive. Looking at intensities, we see that labour, energy, and capital intensity decrease over time. In contrast, software capital intensity strongly increases. It has an average growth rate of 14.37%. At 18.05%, unscaled software capital rises even more sharply. As a clear relationship between software usage and the use other digital technologies exists, we can assume that overall ICT capital also grew strongly within the analysed time frame. Furthermore, the descriptive statistics of growth rates point to an issue: Median software capital growth is zero and median growth rates for software capital intensity, tangible capital, and tangible capital intensity are negative. Negative median growth rates can be explained by the fact that we generally observe a highly skewed distribution of investments. In addition, we measure a zero median software capital stock growth rate because we allow for firms with no software investments at all. Related software capital stocks remain constant at one (obligatory) euro. Thus, they cannot shrink and their growth rate is zero. These observations are potentially problematic for the econometric analysis. Therefore, a considerable share of our robustness checks address this issue (see C.6.1).

4.5 Econometric Analysis

4.5.1 Translog Cost Function

First, we estimate the translog cost function. A potential omitted variable bias with respect to unobserved firm characteristics is a common problem in empirical studies. To address this issue, we remove time-invariant firm-specific fixed effects from the estimation by taking first differences from Equation (4.3). Accordingly, $\Delta u_{i,t}$ captures the time-specific deviation of firm i . Moreover, we add a dummy variable for every year to capture disembodied technological change at time t . To be accurate,

⁶⁹Ratios are comparable to aggregated EU KLEMS data.

⁷⁰It should also be noted here that standard deviations for all logarithmic growth rates are larger than those of aggregated industry-level data.

we replace ICT capital (K_{ICT}) by software capital (K_{SW}). Equation (4.8) denotes the empirical specification of the translog model:⁷¹

$$\begin{aligned} \Delta S_{Ei,t} = & \hat{\beta}_{EE} \Delta \ln \left(\frac{P_E}{P_L} \right)_{i,t} + \hat{\beta}_{EK_{SW}} \Delta \ln \left(\frac{K_{SW}}{Y} \right)_{i,t} + \hat{\beta}_{EK_N} \Delta \ln \left(\frac{K_N}{Y} \right)_{i,t} \\ & + \hat{\beta}_{EY}^* \Delta \ln Y_{i,t} + \sum_{t=2010}^T \hat{\delta}_{Et} t_{i,t} + \Delta u_{i,t}. \end{aligned} \quad (4.8)$$

Table 4.2: First-difference estimation results of Equation (4.8).

	dependent variable: ΔS_E				
	(1)	(2)	(3)	(4)	(5)
$\Delta \ln \left(\frac{P_E}{P_L} \right)$	0.0285*** (62.02)	0.0284*** (169.56)	0.0295*** (51.02)	0.0251*** (35.08)	0.0288*** (53.34)
$\Delta \ln \left(\frac{K_{SW}}{Y} \right)$	-0.000245*** (-5.20)	-0.000238*** (-5.19)	-0.000206*** (-4.47)	-0.000214*** (-4.11)	-0.000220*** (-3.85)
$\Delta \ln \left(\frac{K_N}{Y} \right)$	-0.0013*** (-3.43)	-0.0015*** (-4.42)	-0.0013** (-3.27)	-0.0011 (-1.31)	-0.00181** (-3.28)
$\Delta \ln(Y)$	0.0017** (3.21)	0.0013*** (3.32)	0.0010 (1.48)	0.0014 (1.43)	0.00151* (2.15)
$\bar{\epsilon}_{E/Y, K_{SW}}$	-0.0069	-0.0067	-0.0059	-0.0062	-0.0062
$\epsilon_{E/Y, K_{SW}}$ at \bar{S}_E	-0.0026	-0.0026	-0.0022	-0.0024	-0.0024
Year	x	x	x	x	x
Industry		x	x	x	x
Multi-unit		x	x	x	x
Federal state		x	x	x	x
Size class		x	x	x	x
EEG exemption		x	x	x	x
Producer		x	x	x	x
Trading		x	x	x	x
Firm					x
Observations	89,653	89,653	59,405	25,715	89,653
Adjusted R^2	0.267	0.271	0.290	0.250	0.267

Notes: Column (1): Basic specification. Column (2): Preferred specification including control variables. Column (3): Only changes in software capital stocks after observed in their third period or later. Column (4): Only increasing software capital stocks. Column (5): Additional firm-level fixed effects. t statistics in parentheses. First-difference estimation. Clustered standard errors. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. $\bar{\epsilon}_{EK_{SW}}$ displays the average elasticity for energy intensity. $\epsilon_{EK_{SW}}$ at \bar{S}_E displays the elasticity for energy intensity at the average of relative energy costs.

Column (1) of Table 4.2 presents results for the baseline specification, not including any additional control variables. Column (2) includes control variables as described in Section 4.4.2. It is our preferred specification. Both columns show very similar results.⁷² The coefficient for software capital intensity (hereafter “software

⁷¹We allow for clustering of observations at the firm level when calculating the standard errors of estimates.

⁷²If Schulte et al.’s (2016) and Taneja & Mandys’s (2022) estimated coefficients point in the same direction, this studies estimated coefficients do so as well.

coefficient”) is negative and significant at a high threshold, but its effect size is much smaller than in previous industry-level estimates. Calculating the energy intensity elasticity by Equation (4.4) leads to the following result: An increase in software capital of one per cent is associated with a 0.007% decrease in energy intensity on average. The energy intensity elasticity for the average firm, i.e., at average relative energy costs, is -0.003% .⁷³ Consequently, the relationship is highly inelastic in this assessment of microeconomic data.

Comparing these results to previous industry estimates reveals large differences between both levels of aggregation. The results derived by Schulte et al. (2016) allow us to compute the industry-level energy intensity elasticity at average relative energy costs, which is -0.173 ($\epsilon_{E/Y, K_{SW}}$ at \bar{S}_E).⁷⁴ We consider this a fundamental difference from the firm-level value of -0.003 .

One reason why effects at the firm level have a smaller magnitude could be a potential measurement error. For instance, initial capital stocks may be biased and investments need to be considered for a couple of periods to calculate reliable capital stocks. To shed light on whether the effects are smaller due to this reason, we estimate the translog model only with firms observed in their third period or later. Column (3) shows that, when exclusively considering more reliable software capital stocks, the magnitude of the software coefficient is consistent with our preferred specification.⁷⁵

Additionally, we may observe a misleading correlation. In general, if firms do not invest, their capital stock is depreciated. Hence, it decreases automatically. If especially those firms that do not invest, increase their relative energy use, we would also measure negative capital intensity coefficients. However, this result would be deceptive, as we cannot relate energy intensity improvements to investments. To analyse whether this is an issue with respect to software usage, we re-estimate Equation (4.8) and only consider observations for which the software capital stock is increasing. Column (4) shows that if only increasing software capital stocks are considered, the effect size is comparable to our preferred specification and the software coefficient is significantly negative at a high threshold. In Column (5), we additionally control for firm-level fixed effects. The software coefficient remains highly significant at a comparable magnitude to our preferred specification.⁷⁶

The fact that the energy intensity elasticity is small does not necessarily mean that it is not relevant, as software capital grew strongly in our sample in the observed

⁷³The difference between the average effect and the effect at the average can be explained by the skewed distribution of relative energy costs. See Table C.3 in the Appendix.

⁷⁴The coefficient for ICT capital derived by Schulte et al. (2016) is -0.016 . It is divided by the average energy cost share at the industry level, which is 0.092 .

⁷⁵For further robustness checks with respect to endogeneity issues caused by a potential measurement error of our main variable of interest see C.6.1.

⁷⁶Besides, including firm-level fixed effects marginally downsizes the R-squared. That is why, we consider the specification without fixed effects as our preferred specification.

time frame. Hence, software capital may still relate to considerable energy intensity improvements due to its large growth rate. To discuss this, we perform a back-of-the-envelope calculation. We multiply the average elasticity by the average annual growth rate of software capital,⁷⁷ which is 18.05%. This translates into an annual decrease in energy intensity of 0.12%, using the average elasticity of our preferred specification. Over ten years this would result in energy intensity improvements of roughly 1.5%, assuming that software capital would continue to grow at such a high rate and the relationship between software and energy remains stable. This shows that software investments do relate to energy intensity improvements and cost savings to some extent, but are not a key driver for achieving sustainability targets.

A further question is whether it would be economically rational to invest in software to save energy. To shed light on this economic consideration, we perform a second back-of-the-envelope calculation and approximate average energy costs savings per euro invested in software in the year of investment. We use again the average energy intensity elasticity based on our preferred specification.⁷⁸ We measure that €1 invested in software saved approximately €0.02 in energy costs on average in the analysed time frame. This calculation illustrates that investing in software to save energy (costs) generally does not appear to be economical from a firm's perspective. This confirms a rationale already observed in questionnaire-based surveys, in which firms were asked for non-technical self-assessments: Savings in energy consumption due to the use of digital technologies are more accurately described as a welcome side effect and do not appear to be large enough to be the main motivation for conducting digitalisation projects (Bertschek et al. 2020, European Commission 2021b).

To sum up, an increase in software capital is associated with a decrease in relative energy use, but the relationship has a much smaller magnitude than previous industry-level estimates suggest. Effect sizes at the firm level are robust with respect to various econometric specifications of the translog model.⁷⁹

4.5.2 Reduced Form CES Production Function

In this section, we analyse the robustness of firm-level results with respect to the theoretical approach. To this end, we additionally estimate a reduced form CES-based approach, following Collard et al. (2005) and Bernstein & Madlener (2010). To derive our empirical specification, we take first differences of a reduced form

⁷⁷Please note that we could also have used here mean software capital intensity. However, as it does barely influence the coefficient whether we consider software as an intensity or not (see Column (1) of Table C.8 in the Appendix), we selected the software capital growth rate because it is larger. This allows us getting an upper bound estimate of energy intensity savings.

⁷⁸In C.4, our approach is described in detail.

⁷⁹For further robustness checks see C.6.2.

of Equation (4.7). Additionally, we measure the general input price level by the producer price index, which we retrieve at a two-digit industry level from Destatis.⁸⁰ Similar to the translog model, we capture disembodied technological progress by time dummies and replace ICT capital (K_{ICT}) by software capital (K_{SW}). Accordingly, the empirical specification of the CES-based approach is denoted as follows:

$$\Delta \ln \left(\frac{E}{Y} \right)_{i,t} = \hat{\beta}_{\frac{P_E}{P_{PPI}}} \Delta \ln \left(\frac{P_E}{P_{PPI}} \right)_{i,t} + \hat{\beta}_{\frac{K_{SW}}{K_N}} \Delta \ln \left(\frac{K_{SW}}{K_N} \right)_{i,t} + \sum_{t=2010}^T \hat{\delta}_{Et} t_{i,t} + \Delta u_{i,t}. \quad (4.9)$$

We estimate the same five specifications that we present for the translog model in Table 4.2. Table 4.3 shows that the energy intensity elasticity with respect to the software-tangible capital ratio is robust across different specifications.

Table 4.3: First-difference results of Equation (4.9).

	dependent variable: $\Delta \ln E/Y$				
	(1)	(2)	(3)	(4)	(5)
$\Delta \ln \left(\frac{P_E}{P_{PPI}} \right)$	-0.449*** (-58.72)	-0.446*** (-58.43)	-0.416*** (-44.54)	-0.481*** (-34.27)	-0.455*** (-51.64)
$\Delta \ln \left(\frac{K_{SW}}{K_N} \right)$	-0.00278*** (-3.57)	-0.00289*** (-3.71)	-0.00278*** (-3.54)	-0.00239** (-2.68)	-0.00240* (-2.52)
Year	x	x	x	x	x
Industry	x	x	x	x	x
Multi-unit	x	x	x	x	x
Federal state	x	x	x	x	x
Size class	x	x	x	x	x
EEG exemption	x	x	x	x	x
Producer	x	x	x	x	x
Trading	x	x	x	x	x
Firm					x
Observations	89,267	89,267	59,405	25,609	89,267
Adjusted R^2	0.224	0.228	0.208	0.252	0.249

Notes: Column (1): Basic specification. Column (2): Preferred specification including control variables. Column (3): Only changes in software capital stocks after observed in their third period or later. Column (4): Only increasing software capital stocks. Column (5): Additional firm-level fixed effects. t statistics in parentheses. First-difference estimation. Clustered standard errors. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

However, it has to be acknowledged that results presented in the last two columns are significant at lower thresholds. The software/tangible capital coefficient varies between -0.002 and -0.003 . Thus, the CES-based elasticity is slightly smaller than the average elasticity derived with the translog approach ($\bar{\epsilon}_{E/Y, K_{SW}}$). However, elasticities overlap when the translog model is considered at average relative energy costs ($\epsilon_{E/Y, K_{SW}}$ at \bar{S}_E). Hence, both approaches are consistent at mean

⁸⁰We lose a small fraction of observations as the producer price index is not available for the repair and installation industry (Division 33) for 2009.

values. A possible explanation for this phenomenon is the fact that we assume a strictly linear relationship in the CES-based approach. Accordingly, the CES-based model is less flexible and could be biased at large divergences from mean values. Nonetheless, the qualitative interpretation of the CES-based approach is comparable to the translog approach in that sense that the link between ICT and energy intensity is only small and reductions in energy intensity related to the use of digital technologies cannot be associated with substantial energy intensity improvements – independent of the theoretical approach.⁸¹

4.5.3 Properties of the Link Between ICT and Energy Relating to Aggregation Issues

It is a common issue in environmental economics that results diverge between different levels of aggregation (see Fezzi & Bateman 2015, Heerink et al. 2001), especially when analysing the substitutability between production function inputs (e.g., Haller & Hyland 2014). For instance, in a meta-analysis focusing on empirical results based on translog cost functions and one-, two-, or four-digit industry-level data, Koetse et al. (2008) find that studies using more aggregate data tend to find larger substitution elasticities between tangible capital and energy. An aggregation bias occurs when estimates derived from aggregated data do not accurately reflect individual behaviour. According to Koetse et al. (2008), the fact that substitution is a microeconomic phenomenon results in an overestimation of actual substitution behaviour between tangible capital and energy at aggregate levels.⁸²

Given that this study's results also differ from previous industry-level studies, the question arises as to which level of aggregation correctly reflects the link between software capital and energy intensity. In the following, we demonstrate that the relationship at the firm level indeed entails certain properties that can cause an overestimation of individual substitutability when aggregated data is employed.

Differences Across Industries

The manufacturing sector produces a variety of goods that require different production processes, diverging in their energy intensity. This being the case, the energy-saving potential of digital technologies may be different across industries, such as the paper industry or the cement industry (see, e.g., Bernstein & Madlener 2010, Zhang et al. 2018, and Ateş et al. 2021). However, diverging effects, i.e., slope heterogeneity, can result in a heterogeneity bias if slopes correlate with the variance of

⁸¹Besides, in the CES-based approach we specify the use of digital technologies by the software/tangible capital ratio and not by software capital intensity. We examine whether the difference between both variables can explain the small divergence in elasticities between both theoretical approaches. However, we find that elasticities are robust to different specifications of software use.

⁸²See also Solow (1987) and Haller & Hyland (2014) for similar arguments.

the variable of interest (Imbs & Mejean 2015, Campello et al. 2019), even if we control for industry-specific fixed effects.⁸³

To analyse differences across industries, we split our sample based on industry affiliations and individually fit the translog model for industries at the two-digit NACE level.⁸⁴ Estimation coefficients for software capital intensity by industry are displayed in Figure 4.3. If estimated independently, the software coefficient is negative but insignificant for most industries. However, it shows significant negative effects at the 90% threshold⁸⁵ for manufacturers of paper and paper products (Division 17); chemicals and chemical products (Division 20); other non-metallic mineral products – including the cement industry – (Division 23); basic metals – including the iron and steel industry – (Division 24); electrical equipment (Division 27); as well as for the repair and installation industry (Division 33). Most of these industries are considered energy-intensive.⁸⁶ Consequently, a reduction in energy costs as a share of variable costs, as well as related energy intensity improvements appear to be driven by industries that consume relatively more energy, such as the non-metallic mineral products and the basic metals industries. One economic explanation for this phenomenon may be that the incentives to promote energy efficiency improvements (through ICT) are larger in industries consuming higher amounts of energy.⁸⁷

Further, we analyse whether differences between industries are statistically significant. Equal coefficients would imply that no bias is involved in simple linear aggregation (Zellner 1962). To test for coefficient heterogeneity, we follow Zellner (1962) and Baltagi (1981) and apply the F-test: We interact industry dummies with software capital intensity growth rates and test the null that coefficients of interaction terms are all equal, as illustrated by Equation (4.10):

$$H_0 : \hat{\beta}_{EK_{ICT}N10} = \hat{\beta}_{EK_{ICT}N11} = \dots = \hat{\beta}_{EK_{ICT}N33}. \quad (4.10)$$

The resulting F-statistic is 4.20 with 22 degrees of freedom⁸⁸ and a p-value of 0.00. As a consequence, we reject the null, as differences between industries are statistically significant and assume that slope heterogeneity exist.

As results suggest slope heterogeneity, we may face a heterogeneity bias at different levels of aggregation. This issue exists if observed industry-specific slopes

⁸³For a formal illustration of the problem, see Appendix C.5.

⁸⁴We only consider firms that do not switch between industries in the analysis (N = 85220).

⁸⁵We lower the threshold as there are relatively few observations available for certain industries.

⁸⁶See https://www.destatis.de/EN/Themes/Economic-Sectors-Enterprises/Energy/_Graphic/_Interactive/energy-consumption-industry.html [Online; accessed on 11 Apr. 2023].

⁸⁷Besides, only particular (energy-intensive) industries, such as cement and steel, are part of the European Union Emissions Trading System (EU ETS). Hence, the EU ETS potentially generates extra cost pressure for these industries to save energy-related carbon emissions. See https://ec.europa.eu/clima/eu-action/eu-emissions-trading-system-eu-ets_en [Online; accessed on 11 Apr. 2023].

⁸⁸Note that the tobacco industry is excluded.




Figure 4.3: **Industry-specific estimations by Equation (4.8).** Each colour relates to an industry at the two-digit NACE level. The dots mark respective estimation coefficients and the corresponding lines represent confidence intervals at the 90% level. We change the threshold as much fewer observations are available if we consider each industry separately. Additionally, the tobacco industry is excluded because of few observations and a very low R-squared. Regression results including all variables can be found in Table C.10 in the Appendix.

correlate with the variance of the software capital growth rate (hereafter referred to as the “variance of software”). Moreover, an aggregation bias is present when this correlation occurs only at aggregate levels. It is not possible to test for this issue directly, as most software coefficients are not statistically significant. Therefore, we have to take a different approach to provide insights into this issue. To construct an argumentative bridge, we use our claim that the magnitude of ICT-enabled energy efficiency improvements within one industry relates to the level of energy intensity. In a first step, we show that the more energy-intensive industries and firms are, the larger the magnitude of ICT-related energy intensity improvements. Then, we demonstrate that energy intensity correlates with variance of software at the industry level but not at the firm level.

Firstly, we illustrate that a higher average energy intensity relates to greater ICT-related reductions in the energy cost share. For this purpose, we sort our sample into quartiles based on average energy intensity at the two-digit NACE level and at the firm level, respectively. After splitting the sample, we estimate Equation (4.8) for every quartile separately. The left panel of Figure 4.4 shows that for firms being in the two quartiles corresponding to industries with a low average energy intensity, the point estimates have a much lower magnitude than the average effect for the pooled sample (displayed in our main results), even though improvements are significant in the lowest quartile. For the two quartiles corresponding to firms in industries with higher average energy intensity, the magnitude of the coefficient increases and the coefficient of the highest quartile is much larger than the average effect for the pooled sample. We observe a similar phenomenon when we sort by average firm-level energy intensity (right panel). Here, the software coefficient is also much more pronounced for the most energy-intensive quartile. We conclude that the higher the average energy intensity of an industry or firm, the more evident the ICT-related reductions in energy costs as a share of variable costs.




Figure 4.4: Differences with respect to the average level of energy intensity estimated by Equation (4.8). Quartiles sorted by average energy intensity of the associated industry (left panel) and at the firm level (right panel). 1st quartile: Lowest level of average energy intensity, respectively. The dots mark respective estimation coefficients and the corresponding lines represent confidence intervals at the 90% level.

Secondly, we analyse to what extent average energy intensity correlates with the variance of software at different levels of aggregation.⁸⁹ Table 4.4 shows Pearson correlation coefficients between average energy intensity and the variance of software at the industry (two-digit NACE level) and at the firm level. The correlation coefficient is 0.55 at the industry level and 0.02 at the firm level. Accordingly, a higher average energy intensity at the industry level can be strongly associated with a higher variance of software. In contrast, a higher average energy intensity is to a much lower degree linked to a higher variance of software at the firm level. We conclude that due to a more pronounced correlation at the industry level, systematic differences across industries predominantly relate to problematic properties when aggregating the data and, therefore, potentially involve an aggregation bias. As results indicate that the variance of software at the industry level is larger when ICT-related energy intensity improvements are higher, industries with steeper slopes are given more weight in econometric estimations.⁹⁰ We, therefore, assume an overestimation of substitution behaviour at the industry level. In addition, we have to acknowledge that a small but significant correlation also exists at the firm level. Hence, we also cannot completely rule out a slight heterogeneity bias at this observational level as well.⁹¹

Table 4.4: Correlation between average energy intensity at the industry (two-digit NACE level) as well as at the firm level and the variance of the software capital growth rate, respectively. *p*-values in parentheses.

variance of software	industry level	firm level
energy intensity	0.55 (0.00)	0.02 (0.00)

Differences Between and Within Firms

Furthermore, there might be an overestimation of substitution behaviour at the aggregate level, as it remains unobserved whether ICT-related energy intensity improvements take place because firms substitute between ICT and energy or because the composition of firms and associated products changes in the market (see Solow 1987). This issue is particularly relevant when examining the substitutability between ICT capital and energy within production processes, as newly established

⁸⁹Please note that we do not consider software capital as an intensity here, as then both variables would be influenced by output.

⁹⁰An unambiguous direction of the heterogeneity bias can be identified if one controls for appropriate fixed effects. For a technical explanation see C.5.

⁹¹To analyse whether a potential heterogeneity bias exists at the firm level, it is possible to compare coefficients derived by the mean group estimator (every industry is estimated separately and then averages are taken) with a standard panel estimator (Pesaran et al. 1996). Here, we do not find any significant statistical difference. Hence, we cannot reject the null that no difference between both estimators exists. However, as most industries have very large confidence intervals this result should be taken with a grain of salt. Results can be retrieved from the authors upon request.

firms with innovative products are likely to be more ICT-intensive and energy-efficient, making it challenging to accurately assess the extent of substitutability between these two factors.

Our main results demonstrate that if software intensity increases within a firm, the magnitude of reductions in energy intensity is only small. However, if average energy intensity and software usage relate to each other, we might observe a large difference *between* firms, indicating an omitted variable bias when not controlling for (unobserved) firm characteristics, such as management ability or diverging products. Firm-level data allows for the disentanglement of differences in effects *between* and *within* firms.

In order to analyse to what extent omitted firm characteristics can affect results at the aggregate level, we re-estimate the translog model, but apply a pooled OLS and a hybrid Mundlak model instead of taking first differences (e.g., Mundlak 1978, Allison 2009). A comparison of pooled OLS with more sophisticated panel estimators allows us to approximate the potential omitted variable bias that occurs when firm characteristics are not taken appropriately into account. The hybrid Mundlak estimator additionally enables splitting the effect size into a within effect, reflecting substitution behaviour within firms, and a between effect, indicating the effect of firm-level characteristics at their averages. In technical terms, it is a random effects estimator in which variables are decomposed into firm-level means (between effect) as well as their distance to the firm-level mean (within effect).⁹² The within-effect coefficients of the Mundlak estimator have to be consistent with coefficients of the fixed effects (FE) estimator. To illustrate this consistency, we additionally present results for the translog model estimated with fixed effects (but not taking first differences).

Column (1) of Table 4.5 displays results for the pooled OLS estimator. The software coefficient is now nearly seven times larger than in our preferred specification (see Column [2] of Table 4.2). Column (2) provides results for the fixed effects estimator. The software coefficient of the fixed effects model is comparable to the coefficient derived by the first difference estimator. Both point in the same direction and have more or less the same magnitude. Column (3) shows results for the hybrid Mundlak model, in which the overall effect is decomposed into a between and within effect. A clear finding is that the within-effects software coefficient is nearly the same as the one obtained with the fixed effects estimator. Moreover, the between effect is comparable to the software coefficient of the pooled OLS estimator. Hence, firms that have a higher software capital intensity tend to have on average lower relative energy costs.

⁹²Including group means allows to relax assumptions of the random-effects estimator. In the case of one independent variable, a hybrid Mundlak model would be $y_{i,t} = \beta_0 + \beta_W(x_{i,t} - \bar{x}_i) + \beta_B(\bar{x}_i) + \varepsilon_{i,t}$.

These results confirm that differences between firms are associated with a much larger software coefficient than changes within a firm. However, this property most likely relates to structural characteristics, rather than to technological substitution within firms. This phenomenon may have consequences at the aggregate level. If changes in the composition of firms and products are not appropriately considered, they can lead to an overestimation of substitution elasticities.⁹³

Table 4.5: Comparison of software coefficients for Pooled OLS, FE, and Mundlak

	dependent variable: ΔS_E		
	(1)	(2)	(3)
$\ln(\frac{K_{SW}}{Y})$	-0.00165*** (-32.95)		
$\ln(\frac{\overline{K_{SW}}}{\overline{Y}})$			-0.00151*** (-16.13)
$\Delta \ln(\frac{K_{SW}}{Y})$		-0.000213*** (-3.48)	-0.000214*** (-3.48)
Observations	123362	123362	123362
R^2	0.577	0.268	0.268

Notes: Column (1): Pooled OLS specification. Column (2): Fixed effects specification. Column (3): Mundlak specification. The following additional control variables are included: Year, industry, multi-unit, federal state, size class, EEG exemption, producer, and trading. For coefficients of all model variables see Table C.9 in the Appendix. t statistics in parentheses. Robust standard errors in Column (1). Clustered standard errors in Column (2) and Column (3). * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

4.6 Discussion

Previous studies employing industry-level data indicate substantial synergies between the ongoing digital transformation and improving energy intensity. Using software capital intensity as a proxy for the firm-level degree of digitalisation, we also observe that an increase in digitalisation is associated with a decrease in relative energy use, but of a much smaller magnitude than suggested by previous industry-level estimates.

It is not unusual that effects are smaller when more granular data is employed. In a meta-analysis on the relationship between tangible capital and energy demand,

⁹³The finding that a higher level of relative energy costs is on average associated with lower software capital intensity can also be related to a larger variance of software. If the software capital stock is relatively small, even small amounts of investment can lead to a relatively large growth rate. The issue that growth rates are more volatile when levels are lower is, for instance, a well-known phenomenon when comparing developing and industrialised countries (Pritchett 2000).

Koetse et al. (2008) observe a similar phenomenon. Nevertheless, we would like to acknowledge that the link between software capital and energy intensity improvements could also be less pronounced at the firm level due to endogeneity issues related to a measurement error in software capital. However, the software coefficient is robust to several sample restrictions and different modifications of software capital stocks (see C.6.1). This consistency may provide some confidence.

A potential aggregation bias explains diverging outcomes between observational levels. Looking at individual industries, we find mostly insignificant effects, but software-related reductions in relative energy costs tend to be significant for exactly those industries that are energy-intensive. Hence, the heterogeneity between industries appears to be systematic. Moreover, results indicate a large omitted variable bias if firm-specific characteristics are not appropriately considered. Both issues can result in an aggregation bias. This insight is highly valuable for future research because it points to potential problems that need to be considered when analysing the link between ICT and energy. Nevertheless, we want to emphasise that these characteristics cannot explain the entire divergence between observational levels. The potential omitted variable bias as well as slopes of highly energy-intensive industries, which determine the heterogeneity bias, are still small in comparison to results derived at the industry level, for instance, by Schulte et al. (2016). In other words, there have to be further reasons for the discrepancy between observational levels. Differences in study design, such as a divergent selection of countries and years, as well as further aggregation issues, may be additional arguments that explain the entire discrepancy.⁹⁴

Furthermore, it could also be that effects are small because software capital is insufficient to approximate the firm-level degree of digitalisation. Considering all possible indicators, we believe that for the purpose of this study, software capital is the most suitable indicator. Unlike other digitalisation indicators, e.g., the number of employees working with a computer, software capital has the advantage that it is measured in monetary values. Another advantage in using software capital is that it is very general in comparison to other digital technologies, such as cloud computing or 3D printing. Almost all hardware requires software and we show a clear relationship between software usage and the use of other digital technologies in Section 4.4.2. Especially technologies that optimise production by analysing large amounts of data, and, thus, potentially improve energy efficiency, rely heavily on software. Nonetheless, further analyses looking at different types of digital technologies may be useful, as heterogeneous effects in this respect could exist as well. Additionally, we have to acknowledge that we do not consider the use of software that is free of charge. However, as we only look at relative percentage changes and it is likely that

⁹⁴It may also be worth looking into diverging effects with respect to short-run and long-run elasticities in future studies.

for most firms both, the use of free of charge and paid software, are proportional to each other, we assume that this does not have a large effect on our results.

One further issue that could explain why energy intensity improvements are small is that energy efficiency improvements may be accompanied by rebound effects, which can also take place at the firm level (Amjadi et al. 2018). For example, potential energy savings may not be fully realised because improvements in energy efficiency increase the attractiveness of using energy as an input factor. How digital technologies relate to this issue could be worth analysing in further research.

Last but not least, the question arises as to what our results imply for the net impact of ICT on total energy consumption. By means of the translog model, we analyse the relationship between software capital intensity and the ratio between energy and labour costs. By estimating a reduced form of a CES production function, we consider the relationship between software usage and the ratio between energy use and output. Many economic studies show a clear link between labour and ICT (e.g., Van Reenen 2011, Michaels et al. 2014, and Atasoy et al. 2016) as well as productivity and ICT (Stiroh 2005, Cardona et al. 2013). In other words, the observed relationship may be driven by the positive effects of software capital on labour and output as well.⁹⁵ Therefore, we refrain from coming to conclusions on absolute energy consumption.⁹⁶

4.7 Conclusion and Policy Implications

Climate change and the digital transformation are hugely influential megatrends. Consequently, analysing their interrelationships is of major importance. Previous studies that employ aggregated data indicate that ICT adoption is associated with a substantial decrease in energy intensity, especially in manufacturing industries.

This is the first large-scale empirical study to analyse the relationship between the use of digital technologies and energy intensity improvements at the firm level. For this purpose, we employ administrative panel data on 28,600 firms in the German manufacturing sector collected between 2009 and 2017. Furthermore, we use software capital intensity as an indicator for the firm-level degree of digitalisation and apply a translog cost function approach for our main analysis. Our results show

⁹⁵We want to emphasise that even if output or labour increase due to software usage and energy consumption remains constant or grows to a lower extent, energy intensity improvements still occur, as energy is used relatively less.

⁹⁶Moreover, we solely measure energy intensity improvements inside firms, i.e., we cannot draw conclusions about additional energy that is consumed in external data centres due to an increase in the use of cloud computing. However, cloud computing has not been used very frequently in the observed time frame and its use has only picked up in more recent years. See https://digital-agenda-data.eu/charts/desi-see-the-evolution-of-an-indicator-and-compare-breakdowns/embedded#chart={%22indicator%22:%22desi_idt_cloud%22,%22breakdown-group%22:%22total%22,%22unit-measure%22:%22pc_ent%22,%22ref-area%22:%22DE%22} [Online; accessed on 11 Apr. 2023].

a statistically significant negative link between software capital and energy costs as a share of variable costs, but the effect size is much smaller than expected. Our findings are robust to several econometric specifications. Thus, we conclude that an increase in the firm-level software capital stock cannot be associated with substantially lower levels of energy intensity within firms.

Moreover, we find that the link between ICT and energy intensity entails properties that can result in an overestimation of substitution behaviour at more aggregate levels. This insight is highly valuable for further studies as it clearly points to issues that arise from aggregation, which need to be considered when analysing the substitutability between digital technologies and energy.

In the light of the current energy crisis, our results may be especially relevant for policy makers, consultants, and managers who aim to mitigate escalating energy costs within firms, yet who overestimate the ability of digital technologies to induce energy intensity improvements. Our results also have policy implications for strategies limiting global climate change and fostering sustainability. As the adoption of digital technologies in manufacturing, at least at its current stage, does not appear to significantly contribute to energy intensity improvements, our findings underpin the importance to combine digital efficiency improvements with environmental strategies (Digitalization for Sustainability (D4S) 2022) and “[...] to include specific holistic sustainability and resilience targets within Europe’s [and individual countries’] digital road map[s] [...]”, which was recently requested by the European Commission (p.12, 2022) in its Industry 5.0 vision.⁹⁷

⁹⁷See https://research-and-innovation.ec.europa.eu/research-area/industry/industry-50_en [Online; accessed on 11 Apr. 2023].

Chapter 5

What Drives the Relationship Between Digitalisation and Industrial Energy Demand? Exploring Firm-Level Heterogeneity

joint work with Anne Berner and Thomas Kneib

5.1 Introduction

The growing number of applications and the rapidly evolving performance of information and communication technologies (ICT) have raised hopes of increasing productivity while simultaneously reducing greenhouse gas emissions and energy use (Kander et al. 2015, IEA 2019). Digital technologies such as smart sensors and advanced data analytic tools offer the opportunity to make energy use more efficient and help to save resources. As a result, current European environmental policies consider digitalisation as a key element in lowering environmental burdens (European Commission 2019, 2020, 2021*a*). However, digital technologies also consume energy and resources, and the negative environmental impacts of producing, using, and disposing of digital devices are becoming increasingly apparent (Williams 2011, Andrae & Edler 2015, Belkhir & Elmeligi 2018, Lange et al. 2020). Therefore, it is a priori unclear whether the ongoing digital transformation will bring synergies or trade-offs between technological progress and environmental benefits. Moreover, the impact of digital technologies on energy consumption could also be heterogeneous and vary across firm- and market-specific characteristics. For instance, the amount of energy used in a digitalised production process may depend on the industry association, such as the chemical and the automotive industry. Additionally,

market concentration and input prices can influence how digital technologies affect energy use, as low competition and input prices may reduce firms' incentives to save on energy costs.

Given the continued growth in the deployment of digital technologies, it is crucial to better understand which factors shape the impact of digital technologies on energy consumption. In the current study, we examine drivers of the relationship in manufacturing. We focus on the manufacturing sector, as it is responsible for a large share of global carbon emissions.⁹⁸ Previous empirical studies on the link between ICT and environmental impacts mainly attempt to prove a homogeneous and directional link between ICT and CO₂ emissions (Zhang & Liu 2015, Chen et al. 2019, Kopp & Lange 2019) or energy outcomes (Collard et al. 2005, Bernstein & Madlener 2010, Schulte et al. 2016, Huang et al. 2022).⁹⁹ However, as the relationship may be heterogeneous and, thus, non-linear (Ben Lahouel et al. 2021, Taneja & Mandys 2022, Xu et al. 2022), standard regression models fall short of fully uncovering the complexity of the relationship. To fill this research gap, we aim to reveal effect heterogeneity by applying a non-parametric, flexible tree-based algorithm, which is called the Generalised Random Forest (GRF) algorithm (Athey et al. 2019). By allowing for heterogeneous effects of observables, this method enables the identification of specific firm-level and external characteristics that influence energy demand. Moreover, the algorithm has the advantage that no assumptions have to be made in advance about the relationships among variables that may cause differences in effects. Instead, relationships are identified in an exploratory manner.

Previous studies have already demonstrated the usefulness of tree-based algorithms for analysing heterogeneity (Davis & Heller 2017, Johnson et al. 2020, Knaus et al. 2021) and apply them to evaluate environmental outcomes (e.g., Valente 2023, O'Neill & Weeks 2019, Prest 2020, Miller 2020, Knittel & Stolper 2021). We contribute to this literature by analysing an extensive administrative panel data set on German manufacturing firms (AFiD)¹⁰⁰ for the years 2009 to 2017. Besides, previous microeconomic studies on the relationship between digitalisation and energy use tend to focus on energy intensity or specific energy carriers. We also extend previous literature by firstly analysing ICT-related changes in overall energy demand at the firm level. To account for sources of self-selection and to considerably reduce potential endogeneity issues, we combine the GRF algorithm with R-learning (Nie & Wager 2021) and apply a difference-in-difference approach to leverage the panel structure of our data. Firm digitalisation is measured by a binary variable that takes a value of one if a firm experiences an increase in software capital and is zero otherwise.

⁹⁸For example, manufacturing industries accounted for 26 % of global CO₂ emissions and for 38 % of global energy use in 2020 (IEA 2021).

⁹⁹This also includes the results from Chapter 4. Examples for energy outcomes are energy use, energy efficiency, and energy intensity.

¹⁰⁰Amtliche Firmendaten für Deutschland.

Our results confirm heterogeneity but indicate a trade-off between an increase in the use of digital technologies and absolute energy savings for the majority of firms. We identify multiple characteristics that explain heterogeneity. For instance, the relative increase in energy consumption is smaller in energy-intensive industries, an increase in market concentration is associated with a higher rise in energy use, and digital firms appear to be less sensitive to electricity price changes (and price policies).

On average, we find that an increase in the firm-level degree of digitalisation relates to a simultaneous rise in energy use of 1.03%. Analysing electricity use and non-electric fossil fuel use separately reveals that the magnitude of the effect is even larger for electricity use (1.34%), yet we do not find a significant effect for fossil fuel use, and the respective point estimate is close to zero. Thus, the results suggest that the overall increase is driven by an intensified use of electricity, which is intuitive as digital technologies mostly consume electric power. In the context of policy objectives, however, our results contradict the expectation that digital technologies will lead to a significant reduction in energy consumption. Nonetheless, as electricity is potentially renewable, our analysis suggests that digital technologies may facilitate an increase in the use of sustainable energy sources, thereby enabling the decarbonisation of energy production.

Currently, a considerable share of European industrial digitalisation policies involves funding for small and medium-sized enterprises (SMEs) as well as for regions that are considered structurally weak. Selective targeting of digitalisation that relates to lower levels of energy demand may allow for greater progress toward climate targets. To evaluate the current synchronisation of industrial digitalisation and climate policies, we examine whether present funding criteria are associated with a smaller increase in energy consumption. In a subgroup analysis, we reveal that smaller firms in structurally weak regions show higher average growth in energy use than larger firms in regions that are considered economically strong. Therefore, the results also indicate a policy trade-off between lowering energy use and supporting technological progress in firms with a need for economic assistance.

The remainder of this paper is organised as follows: The next section deciphers the link between energy use and digitalisation in the light of the current literature (Section 5.2). Section 5.3 explains our empirical strategy with a focus on the Generalised Random Forest methodology to measure heterogeneous relationships. Our empirical analysis relies on an extensive administrative firm-level panel data set that will be described in Section 5.4. Section 5.5 presents and discusses the main results, while Section 5.6 discusses the robustness of our results. Section 5.7 concludes.

5.2 Digitalisation and Energy Use in Manufacturing

In economic literature, the introduction of digital technologies is usually linked to changes in productivity, for example, due to increased process efficiency and the optimisation of work practices (Brynjolfsson & Hitt 2000, Brynjolfsson & McAfee 2011, Cardona et al. 2013). We already know from this strand of literature that digital technologies can have very different effects on productivity improvements at the firm level (Bresnahan et al. 2002, Gal et al. 2019, Dhyne et al. 2021b, Cirillo et al. 2023).

Additionally, more and more studies focus on the environmental impacts connected to digitalised production processes, in particular, on the effect on energy consumption. In this context, the literature identifies four impact channels that drive or mitigate the overall effect on energy demand. At the economy-wide level, these transmission channels can be characterised by the following keywords: (1) direct effects, (2) economic growth, (3) energy efficiency, and (4) sectoral change (Lange et al. 2020).¹⁰¹

Direct effects comprise the energy that is embodied in the production, usage, and disposal of ICT and lead to an increase in energy demand (Williams 2011). The same holds for the second channel, which subsumes that digital technologies can act as a multiplier for *economic growth*. Subsequently, the resulting enhanced consumption of products and services can increase energy use indirectly (Belkhir & Elmeligi 2018, Lange et al. 2020). The third channel implies that *energy efficiency* improvements may lower energy intensity. Especially, grey literature assigns high climate protection potentials to the application of ICT. For instance, GeSI & Accenture (2015) state that digital technologies could abate 2.7Gt of CO₂ emissions by 2030 in manufacturing industries.¹⁰² This is asserted because, for example, industrial control systems allow for an improved fault detection, which potentially reduces per-unit energy and resource consumption as well as wastage (Berkhout & Hertin 2004, Baer et al. 2002). Also, simulation methods and 3D printing can considerably decrease the environmental footprint during product design and engineering processes (OECD 2017b). More generally, Berkhout & Hertin (2004) identify five areas in which ICT can lower relative energy use: a) simulation of production processes, b) intelligent design and operation of products and services, c) intelligent distribution and logistics, e.g., supply chain efficiency or alternative distribution structures, d) changing seller-buyer relationships, e.g., mass customisation, and e) work organisation, e.g., teleworking.

¹⁰¹Please note that other frameworks, such as those proposed by Berkhout & Hertin (2004) and Hilty et al. (2006), have only three transmission channels. Economic growth and sectoral change are combined as third-order effects.

¹⁰²It should be noted here that the study is financially related to telecommunication companies.

However, Lange et al. (2020) point out that the desired effects of energy efficiency improvements on energy demand can be mitigated by rebound effects. These describe the energy-increasing consequences that might be triggered by energy efficiency improvements and lead to a situation where potential savings will not be fully realised (cf. Khazzoom 1980, Gillingham et al. 2016). Last but not least, *sectoral change*, i.e., tertiarisation, relates to a shift to a more service-oriented economy. For instance, software-based solutions do not need to be physically manufactured and thus potentially require less energy and capital.

In a nutshell, ICT directly consume energy and stimulate economic growth, which can increase energy use indirectly, but digital technologies can also foster energy-efficient manufacturing as well as the dematerialisation of goods. Consequently, their usage may have simultaneous positive and negative impacts on energy use, and the respective net environmental impact is a priori ambiguous from a theoretical perspective.

The wide range of mechanisms may explain why it is still under debate whether digital technologies increase or decrease energy use. Studies that find synergies between energy savings and ICT highlight that the energy mix, sector association, production factors, and regional characteristics may influence empirical results: Analysing ten OECD countries, Schulte et al. (2016) conduct a parametric econometric analysis at the sectoral level and confirm that reductions in relative energy demand can be linked to ICT usage. They highlight that relative demand decreases, in particular, for non-electric energy, while relative demand for electric energy is not significantly affected. Accordingly, the relationship may depend on the energy source. Bernstein & Madlener (2010) find mixed results with respect to the effect of computers and software on relative electricity demand for European manufacturing industries. They state that the sign of the effect depends heavily on the involved sector-specific production processes. Applying quantile regression, Taneja & Mandys (2022) find a reduction in relative energy demand, but the magnitude of the reduction varies depending on the level of energy intensity.¹⁰³ Focusing on industrial robots, as well as considering 38 countries and 17 manufacturing industries, Wang et al. (2022) find energy intensity improvements due to robot usage. A closer look at the mechanism reveals that the level of energy use is barely affected, while output increases in response to the intensified use of robots. Thus, the authors do not find absolute environmental improvements. In addition, their results indicate effect heterogeneity with respect to labour and capital intensity. Using a compound index to measure digitalisation, Xu et al. (2022) find reductions in absolute energy use and improvements in the share of renewable energy in total energy at the country level. They also show that effects are mediated by technological innovation and

¹⁰³Energy intensity denotes the ratio between energy demand and output.

are more pronounced in low-income countries. Therefore, they postulate heterogeneous effects with respect to regional characteristics. Majeed (2018) confirms diverging effects of ICT on CO₂ emissions between developed and developing countries. Moreover, applying a non-linear model Ben Lahouel et al. (2021) find that ICT have increased carbon efficiency in Tunisia within the last decades.

In contrast to these rather optimistic findings, other studies indicate a trade-off between environmental outcomes and technological progress. Ren et al. (2021) find that internet development can be linked to an increase in energy use per capita in China. Sadorsky (2012) measures that digital technologies are positively linked to an increase in electricity consumption in emerging economies. Covering 93 countries over the period 1995–2016, Alataş (2021) confirms that ICT increase CO₂ emissions at the country level.

Econometric evidence at the firm level is scarce. To the best of our knowledge, no econometric study to date examines absolute energy use in the manufacturing sector, yet empirical evidence exists with respect to changes in energy intensity: In Chapter 4, we observe only marginal average energy intensity improvements related to software usage. Besides, we find that, even though overall effects are small, relative savings are more pronounced in energy-intensive industries, which indicates effect heterogeneity with respect to different production processes. Applying propensity score matching and focusing on the effect of industrial robots on coal consumption, Huang et al. (2022) find improvements in coal intensity. However, as described in Wang et al. (2022) above, the origin of the improvements is mainly an increase in output. A study conducted by Wen et al. (2021) focuses on environmental pollution measured by chemical oxygen demand (COD) and sulphur dioxide (SO₂). The authors find that an increase in ICT investments and services at the provincial-city level relates to a significant firm-level reduction of pollutants. On the contrary, a study conducted by Brozzi et al. (2020) states that firms seldom consider digital improvements (summarised under the term "Industry 4.0") beneficial for environmental targets but pursue predominately economic opportunities in this regard. A questionnaire-based survey with 1,700 German firms indicates diverging effects. According to non-technical self-assessments, 65% of all surveyed manufacturing firms said that their ICT-related energy use remained constant during the last three years, 22% stated it decreased, and 13% mentioned an increase (Bertschek et al. 2020).

To sum up, previous studies on the relationship between digital technologies and energy use show ambiguous results. One reason for different study outcomes could be that parallel impact channels might lead to diverging effects of ICT on the environment (cf. Lange et al. 2020). While the described channels relate to entire economies, analogous mechanisms can also exist within firms. For instance, digital firms may experience more rapid growth. Also, it is possible that digitalisation

can induce a shift in a firm's product offerings towards an increased provision of services, which tend to be less energy intensive. As already indicated in previous studies but not yet comprehensively investigated, various firm- and market-specific characteristics could influence the magnitude of these individual effects and, in turn, moderate the net effect. In this vein, Berkhout & Hertin (2004, p. 903) argue for moving "beyond the dichotomy between pessimism and optimism" to recognise that the relationship between ICT and energy is "complex, interdependent, [and] deeply uncertain". It is an urgent political task to create the conditions for placing digitalisation at the service of sustainable development. To optimally use the potential of digital technologies for climate protection, Lange et al. (2020) argue that fields of application with a positive environmental impact should be promoted without favouring effects that have negative environmental impacts. Horner et al. (2016, p.16) also conclude from their review study that a "focus on identification of important parameters driving the energy use in ICT-infused systems" is important in future research studies.

Focusing on manufacturing firms, our contribution is not only to measure the impact of ICT on energy use and but also to identify characteristics that moderate the net effect in an exploratory manner. More precisely, we are interested in the following three questions:

1. What role do digital technologies generally play for energy consumption in manufacturing firms?
2. Which firm-level and external characteristics relate to heterogeneity?
3. To what extent does current targeting of industrial digitalisation policies influence energy use?

5.3 Methodology

The literature review shows that identifying the role and importance of ICT for energy use is a complex endeavour. Accordingly, the identification of characteristics that moderate energy consumption in digitalised production processes by applying a linear OLS model would quickly result in estimating too many interaction coefficients. Interpreting all of them would get soon out of hand and hardly be useful from a scientific perspective (Prest 2020, Gulen et al. 2021). As a consequence, we apply a flexible tree-based algorithm, which is suitable to measure complex non-linear relationships. Our estimation approach builds on the Generalised Random Forest (GRF) algorithm (Athey et al. 2019), which is a non-parametric modelling approach that allows us revealing heterogeneity and uncovering subgroup differences by applying the potential outcome framework (Rubin 1974).

5.3.1 Measuring Heterogeneous Relationships

In order to capture the effect that digital technologies may have on energy use, we compare a sample of $i = 1, \dots, n$ firms F over a time period of $t = 1, \dots, T$ years. For each firm, we define a binary variable $W_{i,t} = \mathbb{1} \{ \Delta D_{i,t} > 0 \}$ that indicates whether the firm i increases its use of digital technologies D in period t or not. As we follow a method that has its origin in the causal inference literature, we consider firms for which $W = 1$ as “treated” and firms for which $W = 0$ as “untreated” or “control group”.

Our variable of interest is energy consumption $Y_{i,t}$. We denote the potential energy consumption of a firm that increases its use of digital technologies in period t as $Y_{i,t}(W_{i,t} = 1)$ and the corresponding energy consumption that we would have observed if the firm had not increased its use of digital technologies as $Y_{i,t}(W_{i,t} = 0)$. We define the expected difference between the two potential energy outcomes as the average treatment effect (ATE) τ . If we additionally condition on different covariates $X_{i,t} = x$, we receive the conditional average treatment effect (CATE), which is formally defined as (Athey & Wager 2019):

$$\tau(x) = \mathbb{E} [Y_{i,t}(W_{i,t} = 1) - Y_{i,t}(W_{i,t} = 0) \mid X_{i,t} = x]. \quad (5.1)$$

5.3.2 Generalised Random Forests

A promising method to reveal these heterogeneous treatment effects from observational data is the Causal Forest algorithm (Wager & Athey 2018, Knaus et al. 2021). While the name promises to automatically determine causal relationships, in fact it allows the measurement of high-dimensional interaction. The Causal Forest is a special case of the GRF approach introduced by Athey et al. (2019). This approach builds on the recursive partitioning, sampling, and split selection of the Random Forest algorithm (Breiman 2001), an aggregation method applied to decision trees, i.e., classification and regression trees (CART). The goal of this algorithm is to predict an outcome \hat{y} using a non-parametric function of splitting variables, for instance, various covariates. Within one decision tree, the sample is recursively split into subgroups, optimising the accuracy of the prediction. If a further split does not result in accuracy improvements, we call the subgroup at this node a final “leaf” of the tree.

Variation, and hence, decorrelation between decision trees is achieved, on the one hand, by basing each tree on a subsample S_b of the entire data set (bagging), and on the other hand, by choosing a random subset of all possible covariates to build each tree. This procedure also allows for out-of-bag predictions. Hence, we only consider trees where $i \notin S_b$ to determine relationships and predict $\hat{y}^{-i}(X_{i,t})$ (Athey & Wager 2019). This encounters problems, when working with panel data, as a firm constitutes a cluster of observations. This means that we have to exclude

trees containing the same observation, i.e., firm i at period t , and trees including the same firm i at period $t + s$ to avoid information leakage.

To account for the clustered structure of our data when drawing subsamples for each decision tree, we manipulate the sampling of observations as follows (Athey et al. 2020): Instead of directly drawing \mathcal{S}_b , we first sample clusters J_b from $\{F_1, \dots, F_n\}$. Based on each sampled J_b , we then draw k observations to build each tree.

The ensemble method applied to single trees can be described as a data-adaptive kernel method and formulated by the following, when considering clusters:

$$\hat{y}(x) = \sum_{i=1}^n \sum_{t=1}^T \alpha_{i,t}(x) Y_{i,t}, \quad \alpha_{i,t}(x) = \frac{1}{B} \sum_{b=1}^B \frac{\mathbb{1}(\{X_{i,t} \in L_b(x), F_i \notin J_b\})}{|\{i : X_{i,t} \in L_b(x), F_i \notin J_b\}|}, \quad (5.2)$$

where B indicates the number of “grown” trees, indexed by $b = 1, \dots, B$. $L_b(x)$ is the leaf of the b -th tree containing test point x . Accordingly, $\alpha_{i,t}(x)$ indicates how often an observation falls in the identical leaf as x and it can be used to calculate a weighted average of $Y_{i,t}$ based on the forest-based adaptive neighbourhood of x .

The weighting procedure is one of the main building blocks of the “Generalised Random Forest” framework (Athey et al. 2019). It is implemented in the `grf` package in R, on which we base our analysis.

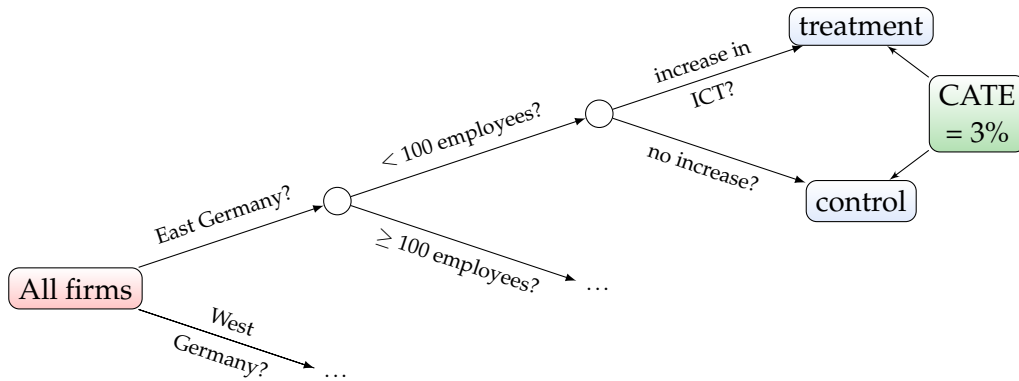


Figure 5.1: **Illustration of Causal Forest partitioning.** The conditional average treatment effect (CATE) is calculated by comparing the effect of an increase of digital technologies between firms within groups of similar firms.

The Causal Forest algorithm aims to predict treatment effects $\hat{\tau}$, which denote the difference between treated and untreated observations within leaves. Accordingly, splits are conducted by maximising treatment effect heterogeneity. Nevertheless, the work horse of the algorithm remains a decision tree. See Figure 5.1 for a graphical illustration of a respective causal tree. The sample is split at each node recursively into two child nodes according to the covariates that maximise the discrepancy between the subgroup ATE. Unequal child node sizes are penalised. Final nodes report the

estimated ATE conditional on the covariates that were responsible for the splitting, which is also known as CATE (Athey et al. 2019).

The size of our database allows us to follow an “honest” estimation procedure, which means that we split the firm panel into two groups: With the first half of the sample, we build the tree structure to calculate weights. Based on these weights, we use the second half of the training sample to estimate CATEs. This procedure prevents overstating the goodness of fit (Athey & Imbens 2019). For the analysis, we grow a forest of 10,000 trees.¹⁰⁴

Identifying Assumptions

Since it is not possible to observe both, firm i increasing its use of digital technologies and not increasing its use in period t , we need the following additional assumptions to accurately estimate Equation (5.1):¹⁰⁵

A.1 Common support: $0 < \mathbb{P}[W_{i,t} = 1 \mid X_{i,t} = x] < 1$, for all x in the support of $X_{i,t}$.

A.2 Unconfoundedness: $\{Y_{i,t}(1), Y_{i,t}(0)\} \perp W_{i,t} \mid X_{i,t}$.

A.3 Exogeneity of covariates: $X_{i,t}^1 = X_{i,t}^0$.

A.4 Stable Unit Treatment Value Assumption (SUTVA): $Y_{i,t} = W_{i,t}Y_{i,t}^1 + (1 - W_{i,t})Y_{i,t}^0$.

The first assumption requires that no subgroup of firms defined by the covariates $X_i = x$ is located in either the treatment or the control group only, which implies that the (inverse) treatment probability must be bounded away from zero and one. The second assumption ensures that potential outcomes are independent of the treatment status, conditional on the covariates. The third assumption imposes that covariates are not affected by the treatment. The fourth assumption requires that there is no interference or no spillover between treated and untreated observations.

Throughout our rigorous analysis, we acknowledge the possibility that all underlying assumptions, whether implicit or explicit, could be scrutinised and called into question. For instance, selection effects may occur, as investments in digital technologies could correlate with specific firm characteristics (Athey & Wager 2019, Gulen et al. 2021). As an illustration, firms that generate more output might consume more energy and have a higher probability to invest in digital technologies. This phenomenon may result in confounding effects and also increase the difficulty to identify counterfactual observations for these firms.

¹⁰⁴In addition to the size of the sample and the covariates used, the forest estimation is also influenced by the maximum split imbalance (between treatment and control group in the child-node) and the minimum node size (minimum number of observations in a final leaf). We tune all parameters by using cross-validation. See Athey & Imbens (2019) for details.

¹⁰⁵We refer here to Knaus et al. (2021) for an extended explanation.

To ensure a substantial degree of overlap (or common support), we trim our sample and only use observations which have propensity scores that match the counterfactual group (Dehejia & Wahba 1999, 2002). For instance, we drop all observations in the group that does not increase its use of digital technologies with an estimated propensity score lower than the smallest estimated propensity score in the group that increases its use of digital technologies, and vice versa for observations with a rising level of digitalisation.

Second, we have to ensure unconfoundedness. In Section 5.3.2 and 5.3.2, we describe how we improve robustness to confounding by employing orthogonalisation and exploiting the panel structure of our data.

The assumption of exogenous covariates might also be violated, since the use of digital technologies can, in addition to energy use, influence other production function inputs, such as tangible capital, labour, as well as output. To solve this issue, we refrain from including critical variables measured concurrently in the same period as the treatment status. Instead, we incorporate them in lagged levels. This procedure allows for the consideration of these variables without risking that the assumption of homogeneity of covariates is violated.¹⁰⁶

We cannot assume with certainty that the fourth assumption of Stable Unit Treatment Values (SUTVA) is fulfilled a priori. For instance, digital technologies can improve the efficiency of entire supply chains and alternate distribution structures. In particular, improved coordination can enable energy and resource savings across decision-making units. Potential changes in energy consumption are, thus, not only a function of a firm's own level of digitalisation, but may also depend on the use of digital technologies by other firms. We assume that such effects are most pronounced between subsidiaries within a firm. Since we consider companies and not plants as the unit of observation, we are able to integrate these kinds of effects into the analysis. However, we would like to acknowledge that taking into account energy efficiency improvements due to enhanced coordination between companies is beyond the scope of our analysis.¹⁰⁷

¹⁰⁶We do not include the variables in lagged growth rates because this would require an additional year to be considered with potentially missing values and we wanted to maintain the maximum amount of observations available in our data set.

¹⁰⁷In addition, a possible violation of the SUTVA may exist due to changes in the market price resulting from ICT-related shifts in energy demand. However, we assume that the potential violation of the SUTVA from this factor is of negligible size.

Orthogonalisation

The assumption of independent assignment of treatment conditional on firm characteristics X is important for unbiased estimates (Assumption A.2). To fulfil this assumption, we account for variables that determine selection into treatment. Previous empirical studies reveal that firm characteristics, including firm size, R&D expenditure, export-intensity, and industry association can drive ICT adoption. Also, external characteristics, such as competition intensity and firm location, as well as the policy situation may play a part (e.g., Giunta & Trivieri 2007, Haller & Siedschlag 2011, Guerrieri et al. 2011, Kinkel et al. 2022, Cho et al. 2023). Accordingly, we add information related to these characteristics to our control variable set.¹⁰⁸

To ensure unconfoundedness by the mentioned characteristics, we orthogonalise treatment and outcome variables by regressing X on Y and W and then subtracting predictions (Robinson 1988, Nie & Wager 2021).¹⁰⁹ This procedure allows for differencing out the variation in outcome and treatment variables attributed to covariates. To this end, we train separate Random Forests to compute estimates of propensity scores $e(x) = \mathbb{P}[W_{i,t} | X_{i,t} = x]$ and expected outcomes $m(x) = \mathbb{P}[Y_{i,t} | X_{i,t} = x]$. This approach is also known as R-learning or local centring.

The $(-i)$ -superscript in this case stands for leave-one-out estimates, indicating that the i -th observation was not used to compute, e.g., $\hat{m}^{(-i)}(X_{i,t})$. The resulting residualised outcome ($Y - m(x)$) and treatment ($W - e(x)$) variables, as well as the weights are combined in the estimation. Hence, treatment effects are estimated by solving the following equation:

$$\hat{\tau} = \frac{\sum_{i=1}^n \sum_{t=1}^T \alpha_{i,t}(x) \left(Y_{i,t} - \hat{m}^{(-i)}(X_{i,t}) \right) \left(W_{i,t} - \hat{e}^{(-i)}(X_{i,t}) \right)}{\sum_{i=1}^n \sum_{t=1}^T \alpha_{i,t}(x) \left(W_{i,t} - \hat{e}^{(-i)}(X_{i,t}) \right)^2}. \quad (5.3)$$

Table 5.1 summarises the main steps of the Causal Forest algorithm including orthogonalisation and honesty.

Table 5.1: Summary of the steps of the Causal Forest algorithm with orthogonalisation and honesty.

1.	Regress $W_{i,t}$ on $X_{i,t}$ to obtain a prediction model for $\hat{e}^{(-i)}(X_{i,t})$.
2.	Regress $Y_{i,t}$ on $X_{i,t}$ to obtain a prediction model for $\hat{m}^{(-i)}(X_{i,t})$.
3.	With the first half of the sample generate Causal Trees but replace $W_{i,t}$ and $Y_{i,t}$ with $W_{i,t} - \hat{e}^{(-i)}(X_{i,t})$ and $Y_{i,t} - \hat{m}^{(-i)}(X_{i,t})$. Then calculate $\alpha_{i,t}$ as in Equation (5.2).
4.	Use the second half of the sample and weights obtained in Step 3 to calculate $\hat{\tau}(x)$ by solving Equation (5.3).

¹⁰⁸For a detailed description of all control variables, see Section 5.4.2.

¹⁰⁹This step can be compared to the consideration of main effects in an OLS regression.

Panel Structure

To reduce confounding due to unobservable characteristics, which can either be time-invariant or time-varying, we exploit the panel structure of our data. Firstly, similar to Athey et al. (2020) and Knittel & Stolper (2021), we take first differences from our outcome variable as well as from control variables to remove individual fixed effects.¹¹⁰ This enables the elimination of a potential time-invariant omitted variable bias. Secondly, in the spirit of Prest (2020), Knittel & Stolper (2021), and Valente (2023), we additionally include a lagged outcome variable, to reduce possible time-varying confounding due to unobservables. This can help to reduce potential confounding, as the outcome from the previous period may be influenced by the same unobservables as current firm characteristics (Lechner 2015). In other words, conditioning on pre-treatment outcomes allows controlling for previous behaviour that might motivate investment in ICT.

5.4 Data

5.4.1 Microdata on the German Manufacturing Sector

Our analysis builds on firm-level data on the German manufacturing sector (AFiD) collected by the Research Data Centres of the Statistical Offices of the Federation and the Federal States (RDC) between 2009 and 2017 ([data set] Research Data Centres of the Statistical Offices of the Federation and the Federal States 2019 a,b). We combine two different AFiD data sources: (1) The AFiD-Panel Industrial Units and (2) the AFiD-Module Use of Energy with additional information such as energy prices and deflators.¹¹¹

Our final panel contains annual information on German manufacturing firms with at least 20 employees at the firm level (yielding around 90,000 observations in total). The longitudinal data set covers basic information about production value, employees, wages, as well as details on production function inputs (e.g., machines and resources). Most importantly, it contains information about energy use, the related energy sources, and software investments.

Even though our data set covers an extensive set of firms, it is a rolling window survey (most firms are observed for four or more consecutive years), which means that not every firm is participating in the survey every year. This makes it difficult to assess whether firms are exiting or entering the market, and, therefore, aggregated effects at the sectoral level cannot be assessed properly. However, as our analysis concentrates on the firm level, this is only a minor limitation.

¹¹⁰Note here that our treatment is also dichotomised based on the growth rate of ICT usage. Hence, also for our variable of interest first differences are taken before it is converted into a binary indicator.

¹¹¹The data set is also used and described in detail in Chapter 4.

Table 5.2: Variable overview.

variable	description	variation	transformation
Outcome			
Y	energy use	firm, year	$\Delta \ln$
	electricity use	firm, year	$\Delta \ln$
	(non-electric) fossil fuel use	firm, year	$\Delta \ln$
Treatment			
W	binary indicator for an increase in digitalisation	firm, year	$\mathbb{1}\{\Delta D > 0\}$
Covariates			
X	output (Q)	firm, year	$\ln t - 1$
	tangible capital (K)	firm, year	$\ln t - 1$
	number of employees (L)	firm, year	$\ln t - 1$
	producer price index (p_M)	year, sector	$\Delta \ln$
	energy price (p_E)	year, sector/district	$\Delta \ln$
	prices for electricity and gas	year, consumption level	$\Delta \ln$
	prices for other energy sources	year	$\Delta \ln$
	lagged outcome (Y_{t-1})	firm, year	$\ln t - 1$
	share of energy source (e.g., natural gas/energy use)	firm, year	$t - 1$
	R&D intensity (R&D divided by Q)	firm, year	Δ
	tax intensity	firm, year	Δ
	subsidy intensity	firm, year	Δ
	trading intensity	firm, year	Δ
	HHI	year, sector	Δ
	relative use of self-produced fossil-based energy	firm, year	Δ
	relative use of self-produced renewable energy	firm, year	Δ
	proxy for renewable levy (EEG) exemption	firm, year	one-hot
	multi/single unit	firm, year	one-hot
	main industrial grouping	firm, year	one-hot
	structurally weak region	district	one-hot
	sector association	sector	LASSO vector
	location	federal state	LASSO vector
	time or disembodied technological change (t)	year	LASSO vector

5.4.2 Variable Description

In this section, we briefly characterise the variables included in the analysis. Unless stated explicitly, first differences are taken. Please find an overview of all employed variables in Table 5.2 and a detailed description of the variables in Appendix D.1. We provide descriptive statistics in Appendix D.2.

We look at three different outcomes of interest (denoted by Y): energy use, electricity use, and non-electric energetic fossil fuel use (hereafter abbreviated by fossil fuel use). Energy use represents the sum of consumed energy sources (renewable and fossil, e.g., natural gas or biomass) plus electricity consumption. All variables are measured in kWh and are log-transformed.¹¹²

The degree of firm-level digitalisation D is approximated via a software capital stock. We consider software capital to be a suitable indicator for firm-level ICT usage, as it is a precursor to almost all digital hardware. In particular in manufacturing, technologies that optimise production processes usually require additional software. The monetary measurement of the software capital stock makes it easy to

¹¹²Note here that electricity consumption and fossil fuel use do not sum up to energy use, since non-electric non-fossil energy, such as biomass, cannot be accounted to either of the two.

compare the proxy across different sectors and provides a certain generality in contrast to investments in single technologies, such as Cloud Computing or robotics.¹¹³ Not without reason, it is a commonly used indicator at the firm level (cf. Almeida et al. 2020, Bessen & Righi 2020, Barth et al. 2022). We also integrate a tangible capital stock K . We calculate both capital stocks by applying the perpetual inventory method (PIM), which allows for generating a productivity-relevant capital stock (cf. Griliches 1980, Lutz et al. 2017). For this purpose, we use deflated investments. Moreover, we base the calculation of software capital on information on software investments, while tangible capital is approximated using information on investments in property, plants, and equipment. We include tangible capital in logarithmised lagged levels in the estimation.

Furthermore, we take first differences of the software capital stock. Based on this transformation, we define a binary treatment indicator W that approximates an increase in the use of digital technologies. Accordingly, the indicator is one if firm i shows an increase in software capital in the year t and zero otherwise.¹¹⁴ Although the Generalised Random Forest (GRF) algorithm allows for the consideration of a continuous treatment indicator, we decided to dichotomise our variable of interest. We do this for two reasons. Firstly, changes in software capital are often accompanied by hardware investments and the use of open source software. We, therefore, rather see the rise in software capital as an indicator for a digital event that takes place inside the firm. Secondly, a continuous treatment indicator would result in the estimation of linear treatment effects, which we believe is a rather unrealistic assumption.¹¹⁵ In summary, we observe for approximately 30% of firms an increase in software capital.

We additionally include numerous covariates for each firm in the analysis. These can serve two purposes. Firstly, they can moderate the impact of ICT on energy consumption. Secondly, they may influence selection into treatment, as well as energy use in general and we, therefore, have to control for them via orthogonalisation. We group these covariates in five categories: Production function in- and outputs, external factors, firm structure, policy situation, and energy mix. We provide a brief

¹¹³For a detailed description of the capital stock approximation, for a descriptive analysis of the suitability of software capital as an indicator for firm digitalisation, as well as for robustness checks with respect to the depreciation rate, we refer to Chapter 4.

¹¹⁴We are aware of the fact this approach generates an unconventional composition of the control group, comprising companies that do not have any software capital as well as those whose software capital stock remains constant or declines. To analyse whether this may be an issue, we compared descriptive statistics between both subgroups. If weighted by their propensity scores, these descriptive statistics revealed no significant difference between the two groups. Results can be retrieved from the authors upon request.

¹¹⁵Still, we would like to acknowledge that by dichotomising our variable of interest we are discarding information and accepting a possible measurement error. Using a continuous variable may provide additional insights, but since we only have limited access to the administrative data, we refrain from conducting this analysis as an additional robustness check at this stage of our work.

description of the variables here, but refer to the overview in Appendix D.1 for a detailed description and the data sources.

We select *production function in- and outputs* based on a simple energy demand model,¹¹⁶ in which energy and materials are treated as flexible inputs. Hence, in addition to tangible capital K , we consider labour use L , which is approximated by the number of employees, the price for materials p_M , and the energy price p_E , as well as the firm-level production value Q in the analysis. The number of employees and the firm-level production value are integrated in lagged levels to ensure that the exogeneity of the covariates is fulfilled (Assumption A.3 in Section 5.3.2). We approximate the price for materials by the producer price index. For the energy price, we use the location-specific industry average of firm-level expenditure for one kilowatt-hour of energy. We additionally add prices for different energy carriers from external data sources: We merge electricity, natural gas, coal, heating oil, district heat, biomass, and liquid gas prices. Also, we log-transform all price variables.

Information on *external factors* covers variables, such as location (federal state), year of observation t , which approximates disembodied technological change, and industry association. In a standard OLS regression, all three characteristics would typically be included as one-hot-encoded fixed effects. However, trees-based algorithms have difficulties with large one-hot-encoded matrices. Therefore, we follow Jens et al. (2021) and modify them in a two-step procedure. First, we estimate the effect of each variable, coded as fixed effects dummies in a LASSO regression, on Y . For instance, we estimate the effect of each manufacturing industry, such as the automotive industry, on energy use. Second, we create a vector of the respective estimation coefficients for each variable and include this vector as a feature in the GRF estimation instead of a one-hot-encoded matrix. Jens et al. (2021) show the effectiveness of this approach in Monte Carlo simulations. Further external factors that are integrated in the estimation are the competitive situation in each industry approximated by the Herfindahl–Hirschman Index (HHI) and a dummy indicating whether the firm is situated in a region considered “structurally weak” due to its limited economic productivity.

Additionally, we include information on the *firm structure*, such as information on the number of plants, industrial grouping (intermediate goods, capital goods, durable consumer goods, non-durable consumer goods, and energy producer), and the volume of traded commodities, as well as investment in research and development (R&D) relative to output. Except for the last two variables, which are continuous, we integrate all information in levels and one-hot-encoded. Here, we refrain from applying LASSO-based fixed effects vectors, as the number of categories is small.

¹¹⁶ $Y^* = Y(Q, p_E, p_M, K, L, \theta)$, with θ relating to all other parameters that may impact energy use.

The *policy situation* of the firm is characterised by paid taxes and received subsidies. The information is considered proportional to output. Information whether the firm is potentially fully or partly exempt from the EEG levy,¹¹⁷ which is the case for various energy-intensive firms, is also included as a categorical variable.

Last but not least, we include covariates that describe the *energy mix* of the firm, starting with the share of different energy sources used in the production process. All shares are integrated as lagged levels, as current changes may be strongly correlated with the outcome variable. Moreover, we add the share of self-generated energy (fossil and renewable energies).

After this preprocessing our data set contains $p = 78$ covariates and our sample includes 92,315 observations based on 28,734 firms.¹¹⁸

5.5 Results

We structure our results according to the three research questions posed in Section 5.2. Hence, we first discuss the general role of digital technologies for energy consumption in the manufacturing sector. Then, we turn to heterogeneity-driving characteristics before presenting results on selective targeting of current industrial digitalisation policies and their influence on energy use.

5.5.1 Conditional Average Treatment Effects

We start by estimating the conditional average treatment effects of an increase in the use of digital technologies, approximated by a binary indicator, on total energy use, electricity use, and non-electric fossil fuel use. We use each outcome in a separate analysis, i.e., we estimate three separate Causal Forest models.

Figure 5.2 depicts the distributions of the treatment effects predicted by the Causal Forest for the three different outcomes. All panels show out-of-bag (OOB) predictions, which are average predictions for each observation, using only trees that do not include the respective observation (James et al. 2021).¹¹⁹ We find for total energy use that the ICT-related increase in energy consumption ranges roughly from -3% to 6% and has its mean at 1.03% . When electricity use is our dependent variable, the ICT-related increase in electricity consumption is slightly higher and at 1.34% . The contrary holds for fossil fuels, where the ATE decreases to 0.23% and becomes insignificant. It has to be acknowledged here that the range of the CATE

¹¹⁷A levy paid in Germany for electricity consumption to promote renewable energies.

¹¹⁸Please note that the number of observations in Chapter 4 and Chapter 5 marginally diverges because the variable preprocessing is slightly different.

¹¹⁹We excluded a small test sample of 2% of our observations from the training procedure of the Causal Forest model to analyse the external validity. Figure D.7 in the Appendix shows the CATE predictions for this test sample. The similarity between the distributions indicates that the model is well calibrated.

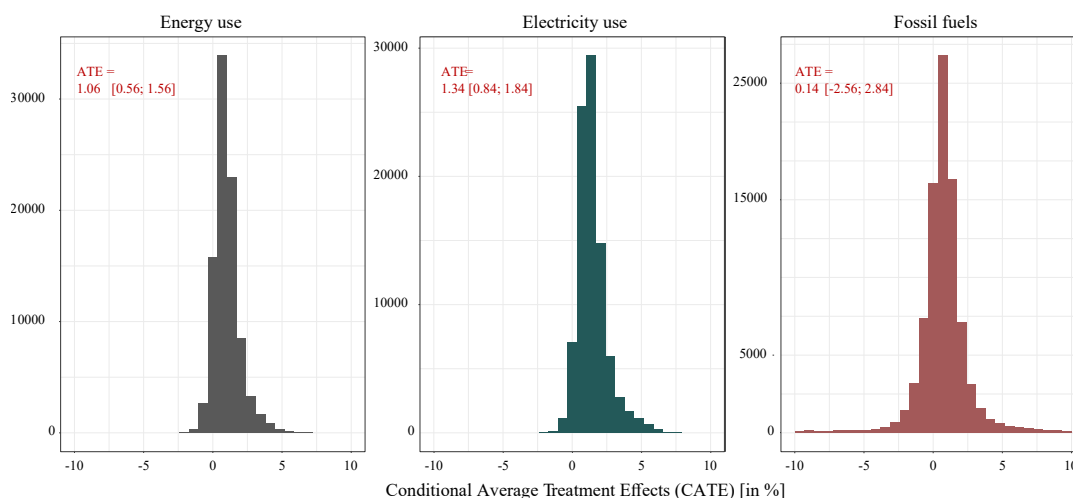


Figure 5.2: Distribution of the conditional average treatment effect (CATE) for the three different outcome variables: energy use, electricity use, and non-electric fossil fuel use.

distribution is now much broader and spans roughly from -30% to 30% .¹²⁰ Overall, the average treatment effect indicates that an increase in the firm-level degree of digitalisation is significantly related to higher levels of energy use. However, there is a small share of firms for which the potential outcome declines. Thus, we conclude that for some firms we can observe both, energy savings and an increase in digital technologies, i.e., potential synergies. Nonetheless, firms for which an increase in the software capital stock relates to growing energy use are far more frequent. The positive relationship seems to be particularly pronounced for the electricity use of a firm, while we cannot determine an unambiguous direction of ICT-related changes in energy consumption for fossil fuel use. Accordingly, results suggest that the change in overall energy use is driven by an increase in electricity use. This is in line with the reasoning that ICT consume mainly electric energy.

At first sight, this finding contradicts previous results from Schulte et al. (2016), who observe that ICT relate to a reduction in non-electric energy, but do not significantly affect the demand for electric energy. However, comparing both studies reveals that Schulte et al. (2016) use different outcome variables. For instance, instead of considering absolute electricity use, they use the share of electricity costs in variable costs as a dependent variable. This divergence may explain the differences between the two studies.

¹²⁰Please note that if not stated otherwise average treatment effects are estimated doubly robust.

We evaluate the Causal Forest fit by applying the Best Linear Prediction Test (Chernozhukov et al. 2018). The test uses the OOB predictions of ICT-related changes in energy consumption to predict actual changes and thereby evaluates the quality of estimates with the following linear model:¹²¹

$$(Y_{i,t} - Y_{i,t-1}) - \hat{m}^{(-i)}(X_{i,t}) = \beta_{\text{ATE}} \bar{\tau} \left(W_{i,t} - \hat{e}^{(-i)}(X_{i,t}) \right) + \beta_{\text{CATE}} \left(\hat{\tau}^{(-i)}(X_{i,t}) - \bar{\tau} \right) \left(W_{i,t} - \hat{e}^{(-i)}(X_{i,t}) \right) + \epsilon_{i,t}. \quad (5.4)$$

The results for the two β -coefficients are reported in Table 5.3 with respect to overall energy use. Since β_{ATE} is close to 1, the model captures the average ICT-related changes in energy consumption well. We also find evidence that the covariates adequately capture the underlying heterogeneity, as the second coefficient (β_{CATE}) is also close to 1 and significant. The results of the other two outcomes are reported in Table D.3.¹²² Although the results for the electricity model are comparable to those of the overall energy model, the fossil fuel model does not appear to adequately predict ICT-related changes in fossil fuel consumption. Thus, for fossil fuels, we cannot reject the null that no heterogeneity exists.

Table 5.3: **Best Linear Predictor Test for the forest with total energy use as outcome.**

	Estimate	SE	t-stat	p-value
β_{ATE}	0.998	0.235	4.245	$1.09e - 05^{***}$
β_{CATE}	1.261	0.366	3.448	0.0003^{***}

Notes: Results of the best linear predictor test for model calibration and heterogeneity that seeks to fit the estimated CATE as a linear function of the out-of-bag predictions (see Equation 5.4).

As our results confirm an increase in energy use at the firm level. It is intuitive to ask how this result affects the overall energy consumption of the manufacturing sector. However, we have to face a limitation in this regard, as even though our data set covers an extensive set of firms, it is a rolling window survey, which means that not every firm of the manufacturing sector has to answer the survey every year. This makes it difficult to assess whether firms exit or enter the market and, therefore, aggregated effects at the sectoral level cannot be assessed properly. Thus, we refrain from conclusions with respect to changes in aggregated energy consumption.

¹²¹The model is calibrated well if β_{ATE} and β_{CATE} are close to one.

¹²²Table D.3 in the Appendix also contains a t-test that approximately examines based on rounded coefficients and standard errors, as well as 10,000 degrees of freedom whether the coefficients of the Best Linear Prediction Test are significantly different from one. For energy use and electricity use, we cannot reject the null that coefficients are equal to one. Hence, we do not find statistical evidence that the model over- or underestimates changes in total energy consumption and electricity consumption related to ICT usage.

5.5.2 Analysing Effect Heterogeneity

While the CATE distributions indicate that the relationship is heterogeneous for total energy and electricity use, it does not clarify how the observed covariates are associated with ICT-related changes in energy consumption.

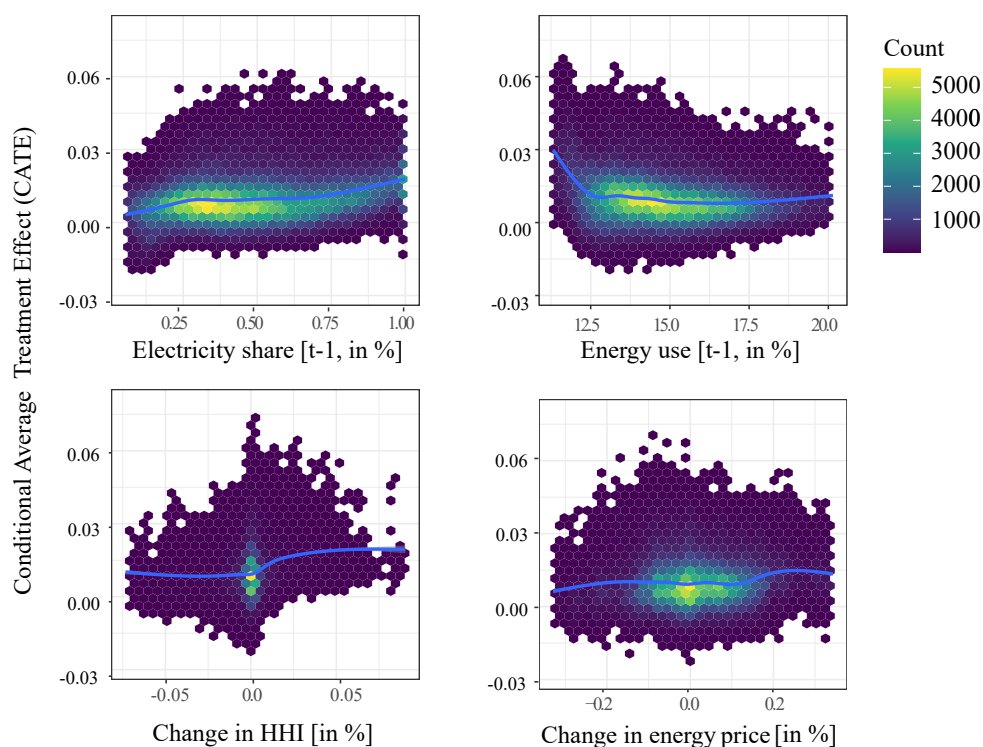


Figure 5.3: **Bivariate distributions and smoothed regression lines for ICT-related changes in energy consumption and selected variables (total energy use).** The colour of the hexagons symbolises the density of the observations and each hexagon comprises at least 5 individual observations. Individual observations cannot be presented due to anonymity constraints.

Figure 5.3 shows bivariate distributions and smoothed regression lines for predicted ICT-related changes in overall energy consumption with respect to the following variables: energy use and relative electricity consumption in the previous period, changes in market concentration, and changes in the overall energy price. The four variables were chosen according to the variable importance in the splitting algorithm for overall energy consumption (see Figure D.5 in the Appendix).¹²³

Previous level of energy use and share of electricity. The upper right panel of Figure 5.3 indicates that firms which used relatively little energy in the previous period are associated with a greater increase in ICT-related energy use. This may imply that smaller firms increase their energy use to a greater extent when investing in ICT, which can be explained by the phenomenon that digital technologies spark economic growth. In addition, the joint distribution of ICT-related changes in overall

¹²³The importance of prices is considered jointly.

energy consumption and the electricity share in the previous period indicates a positive relationship (upper left panel). This result potentially confirms that digitalisation more strongly affects electricity-using firms.

HHI. The HHI is positively correlated with predicted CATEs (lower left). This might imply that digital firms which face less competition use relatively more energy than digital firms in less concentrated markets. Accordingly, fierce competition may provide larger incentives to save costs and mitigate the additional energy consumed by digital technologies.

Energy prices. The lower right panel of Figure 5.3 relates to the overall energy price. It suggests that the association between ICT-related changes in energy consumption and changes in the energy price is positive. Assuming a negative “baseline” effect for energy prices (Labandeira et al. 2017), i.e., a negative own price elasticity, this result indicates that the sensitivity to the energy price decreases for firms that increase their use of digital technologies, since the slope of their energy demand curve potentially becomes less steep compared to firms that do not increase their use of digital technologies.

The energy price only reflects the average price of the energy sources consumed, weighted by their usage. However, in fact, the effects for different energy outcomes may diverge with respect to prices for different energy sources. We assume this because different energy sources can be used as substitutes, and digital technologies may influence own and cross-price elasticities differently. As digital technologies consume electricity, we conjecture that firms that increase their use of digital technologies become more dependent on electricity. Thus, on the one hand, their sensitivity to an increase in the electricity price may decline. On the other hand, if firms increase their use of digital technologies, they may also respond differently to changes in fossil fuel prices. We assume this because they can potentially substitute fossil fuels more easily with electricity and, therefore, may become more sensitive to fossil fuel prices. In summary, we hypothesise that own and cross-price sensitivity for different energy sources is affected if a firm increases its use of digital technologies.

To analyse this claim, we compare the difference between the prices of different energy sources between the 20% of firms (Q5) with the highest predicted increase in ICT-related energy consumption and the 20% of firms (Q1) with the lowest predicted increase. The first panel of Figure 5.4 depicts results for overall energy use as dependent variable. Each bar represents the price difference of an energy source. We see that the electricity price per kWh is higher in Q5 than in Q1. Hence, the firms for which the ICT-related difference in energy consumption is the largest face higher electricity prices. For natural gas, district heat, and coal, we do not observe any notable price differences between Q1 and Q5. For heating oil and liquid petroleum gas




Figure 5.4: Difference between energy prices with respect to the 20% of firms with the highest predicted ICT-related change in energy, electricity, or fossil fuel consumption and the 20% with the lowest predicted ICT-related change. We calculate $Q5 - Q1$.

(LPG), we find a negative divergence. Hence, we observe lower respective prices where the difference in energy consumption between ICT-increasing and not ICT-increasing firms is the largest.

The second panel of Figure 5.4 shows price differences for changes in electricity consumption. It is straightforward to see that electricity and fossil fuel prices are higher where the difference in electricity consumption between ICT-increasing and not ICT-increasing firms is the largest. As explained above, two different mechanisms that work in parallel may explain this difference. On the one hand, sensitivity to electricity prices declines for digital firms. On the other hand, digital firms can more easily switch to electricity if prices of fossil fuels increase; hence, the sensitivity to other prices may increase.

The third panel of Figure 5.4, shows for fossil fuel use that a higher respective positive divergence between ICT-increasing and not ICT-increasing firms can be associated with lower fossil fuel prices. Furthermore, there is no difference between both quintiles with respect to the electricity price. This result is in line with our assumption that price sensitivity increases for fossil fuel prices. However, since the Best Linear Prediction Test does not confirm heterogeneity for fossil fuels, results for fossil fuel use should be interpreted with caution.

In summary, we find that when the electricity difference between ICT-increasing and not ICT-increasing firms is larger than electricity prices are also higher. Furthermore, we find that a smaller increase in energy consumption is more frequently linked to higher fossil fuel prices. Policymakers should be aware that this result suggests that digital firms may be less responsive to an electricity price policy, such as

a levy to promote renewable energies, but may be more responsive to a fossil fuel price policy (targeting non-electric energy consumption).

5.5.3 Group Differences in the Light of Current Policies

In the following, we look at the differences between subgroups with respect to current digitalisation policies. So far, German and also European digitalisation policies,¹²⁴ involve subsidies and funding for small and medium-sized enterprises (SMEs) and for regions that are considered structurally weak. To analyse the interplay of this strategy with climate targets, we conduct a subgroup analysis investigating whether and how the estimated ICT-related increase in energy consumption varies along firm size and regional structure.

We use group average treatment effects (GATEs) for the analysis. GATEs refer to the average of individual treatment effects over pre-defined, low-dimensional characteristics (Knaus et al. 2021). Therefore, they are more granular than the overall ATE but are easier to interpret than the previously described firm-level effects. In the spirit of Athey et al. (2020), we split the sample into quintiles for the exercise.¹²⁵

We estimate GATEs for firms located in regions that are considered either structurally weak or strong along three continuous variables that indicate firm size: the number of employees, the tangible capital stock, and output. We consider these variables all from the previous period, as decision-makers usually observe firm characteristics, and funding decisions are subsequently made. Figure 5.5 shows the GATEs for each of the three “size” variables separately. The horizontal axes depict quintiles for “size” variables. The vertical axes show the estimated ICT-related increase in energy consumption. Note that we calculate the quintiles before grouping the data by region. Green lines relate to firms in structurally weak regions and purple lines to firms in structurally strong regions.

All three panels indicate that the ICT-related increase in energy consumption declines with firm size in both, structurally weak and strong regions. The effects in structurally weak regions vary between 1.45% for firms in the lowest quintiles of labour and output and 0.9% for firms in respective highest quintiles. Furthermore, the effect size is smaller for structurally strong regions, for which the effect range is between 1.25% – 0.85% for quintiles of labour and output. Effect differences for quintiles of tangible capital are slightly less pronounced. Besides, the difference between the energy use of firms with increasing software capital and those without is,

¹²⁴For instance “go digital”, <https://www.innovation-beratung-foerderung.de/INNO/Navigation/DE/go-digital/Foerdermodell/foerdermodell.html> [Online; accessed 17 Mar. 2023] and “digital jetzt”, see <https://www.foerderdatenbank.de/FDB/Content/DE/Foerderprogramm/Bund/BMWi/digital-jetzt-investitionsfoerderung-kmu.html> [Online; accessed 17 Mar. 2023]

¹²⁵Note that we do not estimate GATEs doubly robust, as AIPW-scores tend to not perform well on smaller samples and the common support assumption may not be fulfilled anymore (Glynn & Quinn 2010).

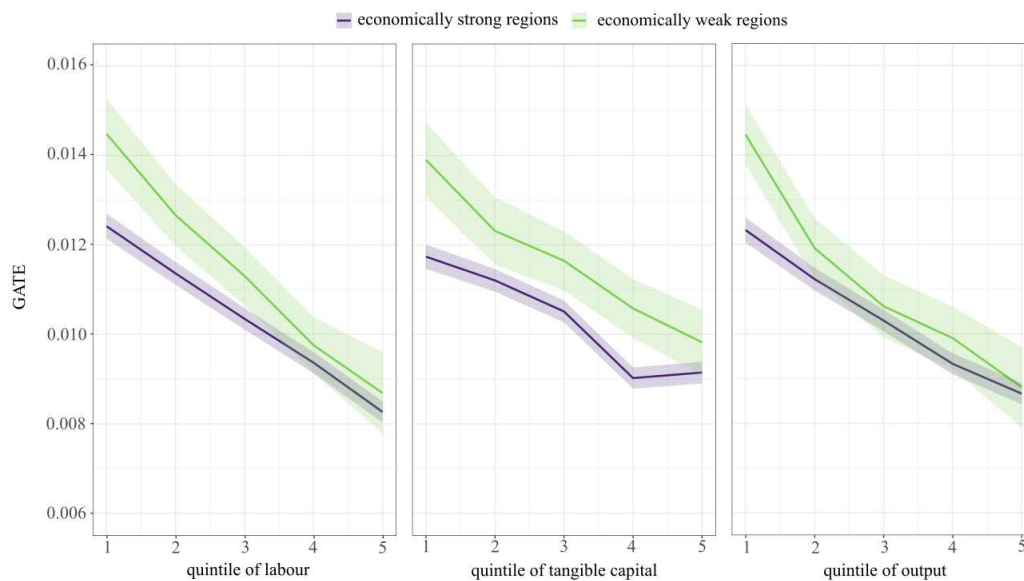


Figure 5.5: Group average treatment effects (GATE) grouped by the economic strength of the corresponding region for different quintiles of labour (number of employees; L), tangible capital (K) and output (Q). The two lines relate to the economic strengths of the region, shaded areas denote 90% confidence intervals

in particular, strong for small firms in structurally weak regions, while the difference between regions is partly insignificant for higher quintiles for each “size” variable and never significant for the highest quintile.

One potential explanation for the higher increase in energy consumption in small firms in structurally weak regions is the fact that digital technologies are a catalyst for economic growth by improving productivity, especially for laggard firms (Borowiecki et al. 2021). Related efficiency improvements exist for economic reasons, such as the generation of scale and scope economies and the reduction of transaction costs (Brynjolfsson & Hitt 2000). Since larger firms in industrialised regions potentially have advantages in economies of scale and scope and fewer transaction costs, digital technologies may spark here productivity improvements and economic growth to a lower magnitude. This phenomenon may explain why we observe a larger increase in energy consumption for smaller firms in structurally weak regions. We conclude that a policy trade-off between the goal of saving energy and economic assistance by increasing the use of digital technologies may be especially pronounced for those firms.

In the next step, we analyse group differences with respect to energy-intensive and other industries. We already put forward the hypothesis that relationships may diverge between industries as production processes vary and, hence, can be differently affected by digitalisation. Considering that a large share of manufacturing’s

total energy consumption is driven by a few industries, differences between industries are policy relevant and may be crucial for achieving climate targets.




Figure 5.6: **Group average treatment effects (GATE) grouped by energy-intensive (Divisions: 10-12, 17,19, 20, 23, 24) and remaining industries for different quintiles of labour (number of employees; L), tangible capital (K) and output (Q).** The two lines relate to the economic strengths of the region, shaded areas denote 90% confidence intervals

Similarly to differences between structurally strong and weak regions, we calculate GATEs for energy-intensive and other industries. We consider the following industries as energy intensive, as they jointly account for more than 80% of the total energy consumption in manufacturing: “food, beverages, tobacco products” (Division 10–12, 5.8%), “paper & paper products” (Division 17, 5.7%), “coke, refined petroleum products” (Division 19, 14.4%), “chemicals & chemical products” (Division 20, 32.9%) “non-metallic products” (Division 23, 7.4%), “basic metals” (Division 24, 16.9%).¹²⁶ We also calculate sector GATEs with respect to quintiles of different “size” variables, as the previous analysis shows large differences in this regard.

Figure 5.6 shows that the increase in energy consumption is less for firms in energy-intensive industries. For the lowest quintile of labour, for example, the increase in energy consumption for energy-intensive industries is only 1.1%, whereas it is 1.35% for other industries. However, the differences decrease for the higher quintiles of “size” variables and are only significant for the highest quintile of tangible capital.

¹²⁶See German Environmental Agency; www.umweltbundesamt.de/daten/umwelt-wirtschaft/industrie/branchenabhaengiger-energieverbrauch-des#primarenergienutzung-des-verarbeiten-den-gewerbes [Online; accessed 9 Apr. 2023].

As assumed, the results suggest that industry differences affect ICT-related changes in energy consumption.¹²⁷ Energy-intensive industries are part of the European Union Emissions Trading System (EU ETS). This system generates an additional incentive to save carbon emissions. Hence, an increase in energy consumption may be attenuated for energy-intensive firms by an increasing pressure to save energy-related carbon emissions.¹²⁸

5.6 Robustness

The role of output, labour use, and tangible capital for the relationship between ICT and energy use is ambiguous in our analysis. On the one hand, these variables might be influenced by digital technologies. Therefore, they are a potential source of biased results, as this would violate Assumption A.3, and we cannot consider contemporaneous changes in variables. However, on the other hand, production function in- and outputs may also be potential confounders. For instance, an increase in tangible capital may correlate with the use of digital technologies. This might lead to the rise in tangible capital being the reason for a higher energy consumption, while digital technologies had actually no impact on energy use. Consequently, by integrating only lagged levels, we cannot fully control for confounding due to simultaneous changes in production function in- and outputs.

To control for respective simultaneous changes, we re-estimate our model and replace lagged output, tangible capital, and labour use by logarithmised growth rates (see Appendix D.5). Since the results are comparable to those of our main model, we conclude that the results are robust and contemporaneous changes only play a minor role.

Moreover, we conduct a second robustness check in which we constrain our definition of an increase in digitalisation and only consider firms as digital for which the software capital stock per employee increases additionally. In this specification, the ATE is now 0.006%, but with a p-value of 0.12 (see Appendix D.5). The Best Linear Prediction Test shows significant results at the 95%-level. Hence, considering software capital per employee, our results also confirm firm-level heterogeneity. However, it should be acknowledged that this result is significant at a much lower level. An explanation for an attenuated statistical power may be, on the one hand, that we now control more strictly for firm growth (also for the ICT-induced one). On

¹²⁷This result does not contradict findings of Chapter 4, in which we show that energy intensity improvements are rather statistically significant in energy-intensive industries. If output increases parallel to ICT investments than energy intensity can decrease. It will probably improve to a greater extent where the rise in absolute energy consumption is smaller.

¹²⁸We have to acknowledge that prices of emission allowances were rather low between 2011 and 2017.

the other hand, we observe firms now in the control group which have been previously considered as digital, i.e., firms with an increasing software capital stock but with a decreasing software capital stock per employee. Observing these firms in the control group may decrease measured ICT-related changes in energy consumption. A further study should shed more light on whether electricity consumption rises because digital firms grow faster or whether digital technologies spark electricity use independent of economic growth.¹²⁹

5.7 Summary and Conclusion

On the one hand, the ongoing digital transformation has raised hopes of climate protection potentials in the energy-intensive manufacturing sector. On the other hand, digital technologies may actually contribute further to environmental damage because they themselves consume energy and resources. However, there is little evidence in the literature that identifies key parameters that determine this relationship.

The main contribution of the article is to disentangle the heterogeneity at the firm level regarding the relationship between ICT and energy use in manufacturing. For this purpose, we apply the Generalised Random Forest algorithm proposed by Athey et al. (2019) to a large administrative panel data set. We harness the panel structure of the data to reduce confounding and mitigate endogeneity issues.

We find that for most firms with an increase in ICT capital, energy use increases relative to firms that do not or barely invest in ICT. Comparing electricity and non-electric fossil fuel use, we additionally show that the relationship differs with respect to different energy sources. We find no significant changes in the use of non-electric fossil fuels, but an average increase in electricity use of 1.34%. Contrary to political hopes, digital technologies seem to increase energy use at the firm level. However, the increase is particularly related to electricity consumption, for which decarbonisation can be realised by renewable energy sources. Furthermore, there is a small share of firms for which energy use declines. Looking closer at the external and firm-level characteristics that may explain heterogeneity, our analysis confirms anticipated rationales. Most interestingly, we observe a growing ICT-related increase in energy consumption with respect to the electricity price, which indicates that the sensitivity to the electricity price declines for digital firms.

Analysing current policy rationales to target SMEs and firms in regions that are considered structurally weak, the analysis reveals that digitalisation policies might not mitigate energy use, while simultaneously fostering technological progress.

¹²⁹It may also be beneficial for future research to examine whether effects of digital technologies on energy consumption in the periods following the investment still persist and if effects differ from the initial period.

However, since our study is the first to shed light on characteristics that determine a change in firms' energy consumption as a response to the ongoing digital transformation, there is a strong need for further research. As digital technologies become even more important in the next few years, so will the question of how to actively shape this process into a direction that supports sustainability goals. To be able to systematically align both policies that support technological progress and instruments that reduce energy use, a better understanding of drivers and moderators, i.e., of firm-level heterogeneity, is essential.

Chapter 6

Concluding Remarks

The goal of this thesis is to shed new light on measuring technological change and its environmental impacts, in order to better understand to what extent current technological progress enables the achievement of sustainability targets. To this end, I make use of new opportunities for empirical economists arising from the widespread adoption of digital technologies. My contributions to this field of research include (1) insights into the usefulness of web-based information to measure technological change, and (2) a novel understanding of the link between digitalisation and energy use patterns, as well as changes in mobility, based on large-scale firm-level data.

In detail, jointly with my co-authors, I show in Chapter 2 that firm websites in combination with emerging text and data mining techniques serve particularly as a useful complement to the survey-based measurement of firm-level product innovation activities as well as innovation expenditure. In Chapter 3, we show that this approach also provides plausible estimates of the degree of digitalisation for a local economy. Moreover, one specific advantage of web data is that information can be accessed in real-time and at a large scale, which is additionally leveraged in Chapter 3 in the context of the Covid-19 pandemic.

Furthermore, by employing large-scale administrative panel data, my co-authors and I find in Chapter 4 for the years 2009 to 2017 that the use of digital technologies, approximated by software capital, in German manufacturing can only be linked to a marginal decline in energy intensity at the firm level. In Chapter 5, we additionally observe an increase in absolute energy consumption, using the same data source. It has to be noted, however, that the observed increase in absolute energy consumption differs between firms and depends, among other factors, on energy prices and on the utilised energy sources. In addition, we find in Chapter 3 that firm digitalisation can only be associated with mobility reductions during the Covid-19 pandemic. An ICT-enabled decline in mobility as well as (most likely) in related carbon emissions did not sustain after most Covid-19 restrictions were lifted. This result holds despite the fact that factors that facilitate telecommunications-transportation substitution greatly improved in the course of the pandemic. Thus, the findings of my thesis with respect to environmental outcomes strongly suggest that it is not given that the

digital and green transitions will promote each other naturally in all areas of our economy. This result is especially policy-relevant because it clearly indicates that digitalisation, at least for this thesis' research subjects, does not sufficiently support the achievement of climate targets and more action is needed to unleash (potential) synergies between both transitions.

My results raise several specific questions for future research that are thoroughly discussed in the respective chapters.¹³⁰ Additionally, I would like to draw attention to two broader issues that pertain to more than one of the presented essays and merit, in my opinion, consideration by future research, as they allow for further insights into the link between technological change and environmental outcomes:

1. Merging web-based and administrative firm-level data.

This thesis uses, on the one hand, data generated in the digital world and, on the other hand, large-scale administrative panel data from German statistical offices. However, both data sources are not analysed jointly at the firm level. They are not merged because very strict data protection regulations of the research data centres of the German statistical offices do not allow the two data sources to be merged at a granular level (with manageable effort). Yet, combining firm websites with large firm-level administrative panel data would enable generating more in-depth insights into the impact of digital technologies and environmental outcomes. For instance, Chapter 4 and Chapter 5, in which digitalisation is approximated by software capital, would have profited from additional robustness checks with software capital being replaced by the web-based digitalisation indicator from Chapter 3, as the latter aims to measure firm digitalisation in a more general way.¹³¹

Moreover, combining information on firm websites with the employed administrative panel data would have provided the opportunity to consider additional variables. Latent characteristics such as a firm's management style and attitude towards environmentally-friendly business practices may influence differences in ICT-related changes in energy consumption. Hence, Chapter 5, which focuses on heterogeneity in the link between digital technologies and energy consumption, would have particularly profited from this additional information. However, such insights are generally not available in German administrative panel data. In contrast, firm websites frequently contain information on key management decisions and environmental strategies. Thus, merging both data sources would allow for taking such aspects into account.

¹³⁰For example, Chapter 3 addresses the problem that technological progress can only be measured by firms that have a website. Chapter 4 discusses whether software capital is a suitable indicator to measure firm digitalisation.

¹³¹I want to acknowledge here that such robustness checks are only possible for firms that have a web presence.

Furthermore, the joint analysis of web data and administrative data can provide a variety of insights which are relevant for guiding the green transition, beyond analysing the impact of digital technologies. For instance, greenwashing practices of firms are generally seen as a threat to the achievement of climate targets. A comparison of self-reported environmental claims on firm websites with actual environmental outcomes would enable a more precise understanding of the extent to which firms are genuinely striving for environmental progress and to what degree such claims are merely used to gain competitive advantages. In addition, factors such as competition intensity or industry affiliation could be identified that may determine the extent to which firms engage in greenwashing. Moreover, apart from environmental questions, it may be fruitful for economic research to link self-reported information on firm websites to information on a firm's cost structure, as researchers could then analyse which claims on firm websites can be linked to an increase in wages and productivity, as well as firm growth. This may allow for finding new answers to questions related to economic prosperity. Hence, the possibility of combining both data sources can pave the way for future research to provide new insights into a broad range of pressing questions. Thus, researchers and decision-makers should strive to facilitate merging both, e.g., through less restrictive data protection regulations.

2. Energy intensity improvements but increased energy consumption.

At first sight, it seems contradictory that the results of Chapter 4 indicate that digital technologies (marginally) improve energy intensity, whereas the results of Chapter 5 suggest that ICT increase on average total energy consumption. One plausible explanation for this phenomenon is that if firms invest in digital technologies, a concurrent increase in output is taking place. More precisely, both, energy consumption and output, rise with ICT adoption, but output grows to a larger extent, allowing for a simultaneous improvement in energy intensity.

I look into the plausibility of this reasoning by conducting an additional reduced-form regression analysis in Appendix E. Here, I estimate the impact of an increase in software capital on changes in absolute energy consumption, output, and energy intensity, respectively (see Equation [E.1]). The results of

Table E.1 in the Appendix confirm that, when firms increase their software capital, output grows more strongly than absolute energy consumption, parallel to observed improvements in energy intensity. Therefore, a concurrent increase in output resolves the apparent inconsistencies between both chapters.¹³²

Nonetheless, I want to highlight that the findings of the Chapters 4 and 5, as well as of Appendix E do not allow for a thorough understanding of why the adoption of digital technologies does not result in absolute energy savings, despite promised potentials to improve energy efficiency.¹³³ It is difficult to obtain comprehensive insights as the results do not indicate clear causal relationships between digitalisation, energy use patterns, and changes in output.

The following three reasons may explain why the results do not indicate a notable decrease in energy use. Firstly, it may be that the energy-saving potential of digital technologies is not as significant as expected.¹³⁴ Secondly, it is possible that specific digital solutions exist that improve energy efficiency, but these were not largely adopted in the observed time frame. Instead, investments were made in other digital technologies that may not necessarily enhance energy efficiency. Thirdly, it could be that digital technologies which improve energy efficiency are available and being used, but adjustments in production factors in response to lower energy service costs (direct rebound effects) sparked additional energy consumption (cf. Amjadi et al. 2018).

To better understand the extent to which digital technologies are being used to improve energy efficiency and whether rebound effects are occurring, it would be helpful to have more detailed information on the current deployment of digital technologies that promise energy efficiency improvements, such as artificial intelligence.¹³⁵ To gain more insights, for instance, a text-based analysis of firm websites or business reports could be conducted to measure the extent to which firms are mentioning that they employ environmentally-beneficial digital technologies.

Furthermore, digitalisation could lead to increased energy consumption through different chains of events. For instance, digitalisation may improve energy efficiency. This can give firms a cost advantage, leading to increased demand and output. As a result of this increased output, energy consumption may also rise (cf. Berner et al. 2022). Another possible mechanism would be

¹³²It is important to note here that a matching is performed in Chapter 5, but not in Chapter 4 and in Appendix E. Thus, respective results also have slightly different implications.

¹³³See Chapter 4 for an explanation of the difference between energy intensity and energy efficiency.

¹³⁴For instance, see Friedrich et al. (2021), Ghobakhloo & Fathi (2021), IEA (2017), OECD (2017a) or Zhang et al. (2019) for examples that promise potentials to improve energy efficiency.

¹³⁵Additionally, it would be interesting to analyse whether digital technologies with energy-saving potentials are being increasingly used since the start of the current energy price crisis.

that firms solely invest in digital technologies to gain a competitive advantage, which, in turn, may allow firms to increase prices as well as sales.¹³⁶ As the employed digital technologies also consume energy, total energy consumption would increase at the firm level as well. Moreover, energy intensity would improve as long as sales grow to a greater extent than energy use.

The lack of clarity regarding these impact channels makes it difficult to develop effective digitalisation policies that allow for lowering energy consumption. Therefore, future research on ICT-related changes in energy consumption should put more effort into identifying causal relationships between digitalisation, energy use patterns, and changes in output. One possible path for future research could be to analyse the impact of an external digitalisation shock, such as subsidies for digitalisation. Then, by analysing temporal variations in changes in output, energy consumption, and energy intensity improvements, it may be possible to determine which of these factors is affected first and whether changes in one variable have an impact on the other.

In addition, the derived results on the link between digital technologies and firm-level energy consumption do not necessarily allow conclusions on changes in aggregated energy consumption. Two potential scenarios introduce ambiguity. Either it could be that aggregated output remains roughly the same, but more digital and energy-efficient firms displace those that are less digital and energy-efficient (displacement effect). Or it could be that the increase in output of digital firms leads to overall output growth, allowing the larger, more digital firms and less digital firms to coexist (growth effect). Hence, aggregated energy consumption could either decrease or increase, depending on whether displacement or growth effects dominate. As descriptive results of Chapter 5 suggest that energy consumption and output generally increase regardless of whether firms invest in software capital, it is possible that the growth effect predominates and digitalisation can be associated with an increase in aggregated energy consumption (see Table D.2). However, this argument needs further investigation. Hence, in addition to identifying directions of causality in the nexus of digitalisation, energy consumption, and output, future research should also focus on the question of how changes at the firm level impact dynamics at the aggregated level.

¹³⁶In this paragraph, the term 'output' refers to the quantity of produced goods or services, while the term 'sales' refers to the revenue generated by selling output.

As the adoption of digital technologies continues to grow, actively guiding this process in a direction that allows for reaching climate targets becomes increasingly important. This thesis shows that the use of novel data and methods can provide a new understanding of how digitalisation is linked to environmental outcomes. However, my concluding remarks also illustrate that while the results obtained in this thesis are important, deeper insights still need to be gained. The continuation of the utilisation of detailed information at the firm level, combining different data sources and observational levels, as well as a stronger focus on causal relationships between digital technologies and environmentally-relevant outcomes are ways to obtain these insights.

Bibliography

- Ackland, R., Gibson, R., Lusoli, W. & Ward, S. (2010), 'Engaging with the public? Assessing the online presence and communication practices of the nanotechnology industry', *Social Science Computer Review* **28**(4), 443–465.
- Aksoy, C. G., Barrero, J. M., Bloom, N., Davis, S. J., Dolls, M. & Zarate, P. (2022), Working from home around the world, Working Paper 30446, National Bureau of Economic Research, Cambridge, MA.
- Aksoy, C. G., Barrero, J. M., Bloom, N., Davis, S. J., Dolls, M. & Zarate, P. (2023), Time savings when working from home, Working Paper 30866, National Bureau of Economic Research, Cambridge, MA.
- Alataş, S. (2021), 'The role of information and communication technologies for environmental sustainability: Evidence from a large panel data analysis', *Journal of Environmental Management* **293**, 112889.
- Alcedo, J., Cavallo, A., Dwyer, B., Mishra, P. & Spilimbergo, A. (2022), E-commerce during Covid: Stylized facts from 47 economies, Working Paper No. 2022/019, International Monetary Fund, Washington, D.C.
- Alipour, J.-V., Fadinger, H. & Schymik, J. (2021), 'My home is my castle – The benefits of working from home during a pandemic crisis', *Journal of Public Economics* **196**, 104373.
- Alipour, J.-V., Falck, O., Krause, S., Krolage, C. & Wichert, S. (2022), The future of work and consumption in cities after the pandemic: Evidence from Germany, CESifo Working Paper No. 10000, CESifo, München.
- Alipour, J.-V., Falck, O. & Schüller, S. (2023), 'Germany's capacity to work from home', *European Economic Review* **151**, 104354.
- Allison, P. D. (2009), *Fixed Effects Regression Models*, Sage, London.
- Almeida, R. K., Fernandes, A. M. & Viollaz, M. (2020), 'Software adoption, employment composition, and the skill content of occupations in Chilean firms', *Journal of Development Studies* **56**(1), 169–185.

BIBLIOGRAPHY

- Amjadi, G., Lundgren, T. & Persson, L. (2018), 'The rebound effect in Swedish heavy industry', *Energy Economics* **71**, 140–148.
- Andrae, A. S. & Edler, T. (2015), 'On global electricity usage of communication technology: Trends to 2030', *Challenges* **6**(1), 117–157.
- Archibugi, D. & Planta, M. (1996), 'Measuring technological change through patents and innovation surveys', *Technovation* **16**(9), 451–468, 519.
- Arora, S. K., Youtie, J., Shapira, P., Gao, L. & Ma, T. (2013), 'Entry strategies in an emerging technology: A pilot web-based study of graphene firms', *Scientometrics* **95**(3), 1189–1207.
- Arundel, A. & Kabla, I. (1998), 'What percentage of innovations are patented? Empirical estimates for European firms', *Research Policy* **27**(2), 127–141.
- Atasoy, H., Banker, R. D. & Pavlou, P. A. (2016), 'On the longitudinal effects of IT use on firm-level employment', *Information Systems Research* **27**(1), 6–26.
- Ateş, K. T., Şahin, C., Kuvvetli, Y., Küren, B. A. & Uysal, A. (2021), 'Sustainable production in cement via artificial intelligence based decision support system: Case study', *Case Studies in Construction Materials* **15**, e00628.
- Athey, S., Friedberg, R., Mühlbach, N. S., Steimer, H. & Wager, S. (2020), Between work, public programs, and retirement: Heterogeneous responses to a retirement reform, in N. Søndergaard, ed., 'Essays in Applied Econometrics and Causal Machine Learning', chapter 2, pp. 41–140.
- Athey, S. & Imbens, G. W. (2019), 'Machine learning methods that economists should know about', *Annual Review of Economics* **11**(1), 685–725.
- Athey, S., Tibshirani, J. & Wager, S. (2019), 'Generalized Random Forests', *Annals of Statistics* **47**(2), 1148–1178.
- Athey, S. & Wager, S. (2019), 'Estimating treatment effects with Causal Forests: An application', *Observational Studies* **5**(2), 37–51.
- Axenbeck, J. & Breithaupt, P. (2022), Measuring firm-level digitalisation – A novel text-mining approach, ZEW Discussion Paper No. 22-065, ZEW – Leibniz Centre for European Economic Research, Mannheim.
- Bachelet, M., Kalkuhl, M. & Koch, N. (2022), What if working from home will stick?, CEPA Discussion Papers No. 41, Center for Economic Policy Analysis, Potsdam.
- Baer, W. S., Hassell, S. & Vollaard, B. A. (2002), *Electricity Requirements for a Digital Society*, RAND Corporation, Santa Monica, CA.

BIBLIOGRAPHY

- Baeza-Yates, R. & Ribeiro-Neto, B. (1999), *Modern Information Retrieval*, ACM press, New York, NJ.
- Bai, J., Brynjolfsson, E., Jin, W., Steffen, S. & Wan, C. (2021), Digital resilience: How work-from-home feasibility affects firm performance, Working Paper 28588, National Bureau of Economic Research, Cambridge, MA.
- Baltagi, B. H. (1981), 'Pooling: An experimental study of alternative testing and estimation procedures in a two-way error component model', *Journal of Econometrics* **17**(1), 21–49.
- Barrero, J. M., Bloom, N. & Davis, S. J. (2021), Why working from home will stick, Working Paper 28731, National Bureau of Economic Research, Cambridge, MA.
- Barth, E., Davis, J. C., Freeman, R. B. & McElheran, K. (2022), 'Twisting the demand curve: Digitalization and the older workforce', *Journal of Econometrics* (forthcoming).
- Beaudry, C., Héroux-Vaillancourt, M. & Rietsch, C. (2016), Validation of a web mining technique to measure innovation in high technology canadian industries, in 'CARMA 2016: 1st International Conference on Advanced Research Methods in Analytics', pp. 100–115.
- Becker, W. & Dietz, J. (2004), 'R&D cooperation and innovation activities of firms—evidence for the German manufacturing industry', *Research Policy* **33**(2), 209–223.
- Belderbos, R., Carree, M. & Lokshin, B. (2004), 'Cooperative R&D and firm performance', *Research Policy* **33**(10), 1477–1492.
- Belkhir, L. & Elmeligi, A. (2018), 'Assessing ICT global emissions footprint: Trends to 2040 & recommendations', *Journal of Cleaner Production* **177**, 448–463.
- Bellstam, G., Bhagat, S. & Cookson, J. A. (2021), 'A text-based analysis of corporate innovation', *Management Science* **67**(7), 4004–4031.
- Ben Lahouel, B., Taleb, L., Ben Zaied, Y. & Managi, S. (2021), 'Does ICT change the relationship between total factor productivity and CO2 emissions? Evidence based on a nonlinear model', *Energy Economics* **101**, 105406.
- Ben Yahmed, S., Berlingieri, F. & Brüll, E. (2022), Adjustments of local labour markets to the Covid-19 crisis: The role of digitalisation and working-from-home, ZEW Discussion Paper No. 22-031, ZEW – Leibniz Centre for European Economic Research, Mannheim.

BIBLIOGRAPHY

- Berkhout, F. & Hertin, J. (2001), Impacts of information and communication technologies on environmental sustainability: Speculations and evidence, Report to the OECD, University of Sussex, Brighton.
- Berkhout, F. & Hertin, J. (2004), 'De-materialising and re-materialising: Digital technologies and the environment', *Futures* **36**(8), 903–920.
- Berlemann, M. & Wesselhöft, J.-E. (2014), 'Estimating aggregate capital stocks using the perpetual inventory method', *Review of Economics* **65**(1), 1–34.
- Berndt, E. R. & Hesse, D. M. (1986), 'Measuring and assessing capacity utilization in the manufacturing sectors of nine OECD countries', *European Economic Review* **30**(5), 961–989.
- Berndt, E. & Wood, D. (1975), 'Technology, prices, and the derived demand for energy', *The Review of Economics and Statistics* **57**(3), 259–268.
- Berner, A., Lange, S. & Silbersdorff, A. (2022), 'Firm-level energy rebound effects and relative efficiency in the German manufacturing sector', *Energy Economics* **109**, 105903.
- Bernstein, R. & Madlener, R. (2010), 'Impact of disaggregated ICT capital on electricity intensity in European manufacturing', *Applied Economics Letters* **17**(17), 1691–1695.
- Bersch, J., Gottschalk, S., Müller, B. & Niefert, M. (2014), The Mannheim Enterprise Panel (MUP) and firm statistics for Germany, ZEW Discussion Paper No. 14-104, ZEW - Leibniz Centre for European Economic Research, Mannheim.
- Bertschek, I., Block, J., Kritikos, A. & Stiel, C. (2022), German financial state aid during Covid-19 pandemic: Higher impact among digitalized self-employed, ZEW Discussion Paper No. 22 - 045, ZEW – Leibniz Centre for European Economic Research, Mannheim.
- Bertschek, I., Erdsiek, D., Niebel, T., Schuck, B., Seifried, M., Ewald, J., Lang, T., Hicking, J., Wenger, L. & Walter, T. (2020), Schwerpunktstudie Digitalisierung und Energieeffizienz - Erkenntnisse aus Forschung und Praxis, Technical report, Bundesministerium für Wirtschaft und Energie (BMWi), Berlin. https://www.bmwk.de/Redaktion/DE/Publikationen/Digitale-Welt/schwerpunktstudie-digitalisierung-energieeffizienz.pdf?__blob=publicationFile&v=12 [Online; accessed 9 Apr. 2023].
- Bertschek, I. & Kesler, R. (2022), 'Let the user speak: Is feedback on Facebook a source of firms' innovation?', *Information Economics and Policy* **60**, 100991.

BIBLIOGRAPHY

- Bertschek, I., Polder, M. & Schulte, P. (2019), 'ICT and resilience in times of crisis: Evidence from cross-country micro moments data', *Economics of Innovation and New Technology* **28**(8), 759–774.
- Bessen, J. E. & Righi, C. (2020), Information technology and firm employment, Law and Economics Research Paper No. 19-6 (2019), Boston University School of Law, Boston, MA.
- bevh (2022), 'Interaktiver Handel in Deutschland - Ergebnisse 2020'. https://www.bevh.org/fileadmin/content/05_presse/Auszuege_Studien_Interaktiver_Handel/Demo_bevh_Gesamtbericht_Interaktiver_Handel_in_Deutschland_2022.pdf [Online; accessed 17 Mar. 2023].
- Blei, D. M., Ng, A. Y. & Jordan, M. I. (2003), 'Latent Dirichlet allocation', *Journal of Machine Learning Research* **3**(Jan), 993–1022.
- Bloom, N., Davis, S. J. & Zhestkova, Y. (2021), 'Covid-19 shifted patent applications toward technologies that support working from home', *AEA Papers and Proceedings* **111**, 263–66.
- Borowiecki, M., Pareliussen, J., Glocker, D., Kim, E. J., Polder, M. & Rud, I. (2021), The impact of digitalisation on productivity: Firm-level evidence from the Netherlands, OECD Economics Department Working Papers 1680, OECD, Paris.
- Breiman, L. (2001), 'Random Forests', *Machine Learning* **45**(1), 5–32.
- Bresnahan, T. F., Brynjolfsson, E. & Hitt, L. M. (2002), 'Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence', *Quarterly Journal of Economics* **117**(1), 339–376.
- Bresnahan, T. F. & Trajtenberg, M. (1995), 'General purpose technologies "Engines of growth"?', *Journal of Econometrics* **65**(1), 83–108.
- Brown, R. S. & Christensen, L. R. (1980), Estimating elasticities of substitution in a model of partial static equilibrium: An application to U.S. agriculture, 1947-1974, SSRI Workshop Series 292581, University of Wisconsin-Madison, Social Systems Research Institute.
- Brozzi, R., Forti, D., Rauch, E. & Matt, D. T. (2020), 'The advantages of Industry 4.0 applications for sustainability: Results from a sample of manufacturing companies', *Sustainability* **12**(9), 3647.
- Brynjolfsson, E. & Hitt, L. M. (2000), 'Beyond computation: Information technology, organizational transformation and business performance', *Journal of Economic Perspectives* **14**(4), 23–48.

BIBLIOGRAPHY

- Brynjolfsson, E., Horton, J. J., Ozimek, A., Rock, D., Sharma, G. & TuYe, H.-Y. (2020), Covid-19 and remote work: An early look at US data, Working Paper 27344, National Bureau of Economic Research, Cambridge, MA.
- Brynjolfsson, E., Jin, W. & McElheran, K. (2021), 'The power of prediction: Predictive analytics, workplace complements, and business performance', *Business Economics* 56(4), 217–239.
- Brynjolfsson, E. & McAfee, A. (2011), *Race Against the Machine: How the Digital revolution is Accelerating Innovation, Driving Productivity, and Irreversibly Transforming Employment and the Economy*, Digital Frontier Press, Lexington, MA.
- Campello, M., Galvao, A. F. & Juhl, T. (2019), 'Testing for slope heterogeneity bias in panel data models', *Journal of Business & Economic Statistics* 37(4), 749–760.
- Cardona, M., Kretschmer, T. & Strobel, T. (2013), 'ICT and productivity: Conclusions from the empirical literature', *Information Economics and Policy* 25(3), 109–125.
- Cariolle, J. & Léon, F. (2022), How internet helped firms to cope with Covid-19, FERDI document de travail P300, FERDI, Clermont-Ferrand.
- Cassiman, B. & Golovko, E. (2011), 'Innovation and internationalization through exports', *Journal of International Business Studies* 42(1), 56–75.
- Chen, X., Gong, X., Li, D. & Zhang, J. (2019), 'Can information and communication technology reduce CO2 emission? A quantile regression analysis', *Environmental Science and Pollution Research* 26(32), 32977–32992.
- Chernozhukov, V., Demirer, M., Duflo, E. & Fernández-Val, I. (2018), Generic machine learning inference on heterogeneous treatment effects in randomized experiments, with an application to immunization in India, Working Paper 24678, National Bureau of Economic Research, Cambridge, MA.
- Chimbo, B. et al. (2020), 'Information and communication technology and electricity consumption in transitional economies', *International Journal of Energy Economics and Policy* 10(3), 296–302.
- Cho, J., DeStefano, T., Kim, H., Kim, I. & Paik, J. H. (2023), 'What's driving the diffusion of next-generation digital technologies?', *Technovation* 119, 102477.
- Choi, H. & Varian, H. (2012), 'Predicting the present with Google Trends', *Economic Record* 88, 2–9.
- Christensen, L. R., Jorgenson, D. W. & Lau, L. J. (1973), 'Transcendental logarithmic production frontiers', *The Review of Economics and Statistics* 55(1), 28–45.

BIBLIOGRAPHY

- Cirillo, V., Fanti, L., Mina, A. & Ricci, A. (2023), 'New digital technologies and firm performance in the Italian economy', *Industry and Innovation* **30**(1), 159–188.
- Collard, F., Fève, P. & Portier, F. (2005), 'Electricity Consumption and ICT in the French Service Sector', *Energy Economics* **27**(3), 541–550.
- Comin, D. A., Cruz, M., Cirera, X., Lee, K. M. & Torres, J. (2022), Technology and resilience, Working Paper 29644, National Bureau of Economic Research, Cambridge, MA.
- Corbet, S., Lucey, B. & Yarovaya, L. (2021), 'Bitcoin-energy markets interrelationships - New evidence', *Resources Policy* **70**, 101916.
- Crepon, B., Duguet, E. & Mairessec, J. (1998), 'Research, innovation and productivity: An econometric analysis at the firm level', *Economics of Innovation and New Technology* **7**(2), 115–158.
- [data set] atene KOM GmbH (2021). Breitbandatlas des Bundes - Release 2/2021 (retrieved on 09/04/2021), for further information see <https://atenekom.eu/project/breitbandatlas> [Online; accessed 2 Jan 2023].
- [data set] Bundesagentur für Arbeit (2022). Pendlerverflechtungen der sozialversicherungspflichtig Beschäftigten nach Kreisen - Deutschland (retrieved on 01/12/2022), for further information see <https://statistik.arbeitsagentur.de> [Online; accessed 9 Apr. 2023].
- [data set] Bundesnetzagentur (2022). Mobilfunkmonitoring (retrieved on 22/12/2022), for further information see <https://www.breitband-monitor.de/mobilfunkmonitoring/download> [Online; accessed 22 Feb. 2023].
- [data set] Destatis (2023a). Mobile Network Data (retrieved on 03/01/2023), for further information see <https://www.destatis.de/EN/Service/EXDAT/Datensatze/mobility-indicators-mobilephone.htm> [Online; accessed 7 Feb 2023].
- [data set] Destatis (2023b). Corona-Daten Deutschland (retrieved as stated in Table B.1), for further information see <https://www.corona-daten-deutschland.de/dataset> [Online; accessed 12 Apr. 2023].
- [data set] Research Data Centres of the Statistical Offices of the Federation and the Federal States (2019a). AFiD Panel Industrial Units (doi: 10.21242/42221.2017.00.01.1.1.0), for the years 2009-2017, for further information see <https://www.forschungsdatenzentrum.de/en> [Online; accessed 9 Apr. 2023].
- [data set] Research Data Centres of the Statistical Offices of the Federation and the Federal States (2019b). AFiD Module Use of Energy (doi:

BIBLIOGRAPHY

- 10.21242/43531.2017.00.03.1.1.0), for the years 2009-2017, for further information see <https://www.forschungsdatenzentrum.de/en> [Online; accessed 9 Apr. 2023].
- [data set] Research Data Centres of the Statistical Offices of the Federation and the Federal States (2019c). Survey on the Use of Information and Communication Technologies in Companies (doi: 10.21242/52911.2012.00.00.1.1.0 - 10.21242/52911.2017.00.00.1.1.0), for the years 2012-2017, for further information see <https://www.forschungsdatenzentrum.de/en> [Online; accessed 9 Apr. 2023].
- [data set] ZEW – Leibniz Centre for European Economic Research (2019). Mannheim Innovation Panel 2019, for further information see <https://www.zew.de/forschung/mannheimer-innovationspanel-innovationsaktivitaeten-der-deutschen-wirtschaft> [Online; accessed 30 Mar. 2023].
- [data set] ZEW – Leibniz Centre for European Economic Research (2021). Mannheim Innovation Panel 2021, for further information see <https://www.zew.de/forschung/mannheimer-innovationspanel-innovationsaktivitaeten-der-deutschen-wirtschaft> [Online; accessed 30 Mar. 2023].
- [data set] ZEW – Leibniz Centre for European Economic Research (2022). Mannheim Enterprise Panel 2020 & 2022, for further information see <https://www.zew.de/publikationen/the-mannheim-enterprise-panel-mup-and-firm-statistics-for-germany> [Online; accessed 30 Mar. 2023].
- Davis, J. M. & Heller, S. B. (2017), 'Using Causal Forests to predict treatment heterogeneity: An application to summer jobs', *American Economic Review* **107**(5), 546–50.
- Dehejia, R. H. & Wahba, S. (1999), 'Causal effects in nonexperimental studies: Reevaluating the evaluation of training programs', *Journal of the American Statistical Association* **94**(448), 1053–1062.
- Dehejia, R. H. & Wahba, S. (2002), 'Propensity score-matching methods for nonexperimental causal studies', *The Review of Economics and Statistics* **84**(1), 151–161.
- Destatis (2021), 'Press – 68% of the persons in employment went to work by car in 2020'. https://www.destatis.de/EN/Press/2021/09/PE21_N054_13.html [Online; accessed 9 Apr. 2023].
- Destatis (2023), 'Experimentelle Daten – Mobilitätsindikatoren auf Basis von Mobilfunkdaten'. <https://www.destatis.de/DE/Service/EXDAT/Datensaetze/mobilitaetsindikatoren-mobilfunkdaten.html> [Online; accessed 12 Apr. 2023].
- Dhyne, E., Konings, J., Van den Bosch, J. & Vanormelingen, S. (2021a), 'The return on information technology: Who benefits most?', *Information Systems Research* **32**(1), 194–211.

BIBLIOGRAPHY

- Dhyne, E., Konings, J., Van den Bosch, J. & Vanormelingen, S. (2021b), 'The return on information technology: Who benefits most?', *Information Systems Research* **32**(1), 194–211.
- Diewert, W. E. (1971), 'An application of the Shephard duality theorem: A generalized Leontief production function', *Journal of Political Economy* **79**(3), 481–507.
- Digitalization for Sustainability (D4S) (2022), Digital reset. Redirecting technologies for the deep sustainability transformation, Technical report, Technical University of Berlin, Berlin. <https://doi.org/10.14279/depositonce-16187> [Online; accessed 30. Mar. 2023].
- Doerr, S., Erdem, M., Franco, G., Gambacorta, L. & Illes, A. (2021), 'Technological capacity and firms' recovery from Covid-19', *Economics Letters* **209**, 110102.
- Durand, B. & Gonzalez-Feliu, J. (2012), 'Urban logistics and e-grocery: Have proximity delivery services a positive impact on shopping trips?', *Procedia - Social and Behavioral Sciences* **39**, 510–520.
- Einav, L. & Levin, J. (2014), 'The data revolution and economic analysis', *Innovation Policy and the Economy* **14**(1), 1–24.
- Emmler, H. & Kohlrausch, B. (2021), Homeoffice: Potenziale und Nutzung. Aktuelle Zahlen aus der HBS-Erwerbspersonenbefragung, Welle 1-4, WSI Policy Brief 52, Hans-Böckler-Stiftung, Düsseldorf.
- Erdsiek, D. (2021), Working from home during Covid-19 and beyond: Survey evidence from employers, ZEW Discussion Paper No. 21-051, ZEW – Leibniz Centre for European Economic Research, Mannheim.
- Erdsiek, D. & Rost, V. (2022), Working from home after Covid-19: Firms expect a persistent and intensive shift, ZEW-Kurzexpertise Nr. 22-06, ZEW – Leibniz Centre for European Economic Research, Mannheim.
- European Commission (2019), 'Communication from the Commission to the European Parliament, the European Council, the Council, the European Economic and Social Committee and the Committee of the Regions: The European Green Deal'. COM(2019) 640 final. https://eur-lex.europa.eu/resource.html?uri=cellar:b828d165-1c22-11ea-8c1f-01aa75ed71a1.0002.02/DOC_1&format=PDF [Online; accessed 9 Apr. 2023].
- European Commission (2020), 'Communication from the Commission to the European Parliament, the European Council, the Council, the European Economic and Social Committee and the Committee of the Regions: Stepping up Europe's 2030

BIBLIOGRAPHY

- climate ambition – Investing in a climate-neutral future for the benefit of our people’. COM(2020) 562 final. <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52020DC0562&from=EN> [Online; accessed 9 Apr. 2023].
- European Commission (2021a), ‘Communication from the Commission to the European Parliament, the European Council, the Council, the European Economic and Social Committee and the Committee of the Regions: Updating the 2020 New Industrial Strategy: Building a stronger single market for Europe’s recovery’. COM(2021) 350 final. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=COM:2021:350:FIN> [Online; accessed 9 Apr. 2023].
- European Commission (2021b), ‘Survey on the contribution of ICT to the environmental sustainability actions of EU enterprises’. <https://digital-strategy.ec.europa.eu/en/library/survey-contribution-ict-environmental-sustainability-actions-eu-enterprises> [Online; accessed 9 Apr. 2023].
- European Commission (2022), ‘The Digital Economy and Society Index (DESI)’. <https://digital-strategy.ec.europa.eu/de/policies/desi> [Online; accessed 9 Apr. 2023].
- Eurostat (2021), ‘Passenger mobility statistics’. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Passenger_mobility_statistics [Online; accessed 9 Apr. 2023].
- Fawcett, T. (2006), ‘An introduction to ROC analysis’, *Pattern Recognition Letters* **27**(8), 861–874.
- Felstead, A., Jewson, N. & Walters, S. (2003), ‘Managerial control of employees working at home’, *British Journal of Industrial Relations* **41**(2), 241–264.
- Fezzi, C. & Bateman, I. (2015), ‘The impact of climate change on agriculture: Nonlinear effects and aggregation bias in Ricardian models of farmland values’, *Journal of the Association of Environmental and Resource Economists* **2**(1), 57–92.
- Flesch, R. (1948), ‘A new readability yardstick.’, *Journal of Applied Psychology* **32**(3), 221.
- Freitag, C., Berners-Lee, M., Widdicks, K., Knowles, B., Blair, G. S. & Friday, A. (2021), ‘The real climate and transformative impact of ICT: A critique of estimates, trends, and regulations’, *Patterns* **2**(9), 100340.
- Frenz, M. & Ietto-Gillies, G. (2009), ‘The impact on innovation performance of different sources of knowledge: Evidence from the UK Community Innovation Survey’, *Research Policy* **38**(7), 1125–1135.

BIBLIOGRAPHY

- Friedman, J., Hastie, T. & Tibshirani, R. (2001), *The Elements of Statistical Learning*, 1 edn, Springer Series in Statistics, New York, NJ.
- Friedrich, R., Ploner, F., Schäfer, C. T., Disselhoff, T., Petkau, A., Hennemann, C., Moecke, J., Wätzig, T., Zimmert, O., Waltersmann, L., Kiemel, S., Mieke, R. & Sauer, A. (2021), Potenziale der schwachen künstlichen Intelligenz für die betriebliche Ressourceneffizienz, Technical report, VDI – Zentrum Ressourceneffizienz, Berlin. <https://www.ressource-deutschland.de/publikationen/studien/> [Online; accessed 9 Apr. 2023].
- Frisch, R. & Waugh, F. V. (1933), 'Partial time regressions as compared with individual trends', *Econometrica* **1**(4), 387–401.
- Gal, P., Nicoletti, G., Renault, T., Sorbe, S. & Timiliotis, C. (2019), Digitalisation and productivity: In search of the holy grail – firm-level empirical evidence from eu countries, Technical report 1533, Paris.
- Gandin, I. & Cozza, C. (2019), 'Can we predict firms' innovativeness? The identification of innovation performers in an Italian region through a supervised learning approach', *PLoS One* **14**(6), e0218175.
- Gentzkow, M., Kelly, B. & Taddy, M. (2019), 'Text as data', *Journal of Economic Literature* **57**(3), 535–74.
- German Environment Agency (2021), 'Energieverbrauch nach Energieträgern und Sektoren'. <https://www.umweltbundesamt.de/daten/energie/energieverbrauch-nach-energietraegern-sektoren> [Online; accessed 9 Apr. 2023].
- GeSI & Accenture (2015), Smarter 2030 - ICT solutions for 21st century challenges, Technical report, Global e-Sustainability Initiative (GeSI), Brussels. http://smarter2030.gesi.org/downloads/Full_report.pdf [Online; accessed 9 Apr. 2023].
- Ghobakhloo, M. & Fathi, M. (2021), 'Industry 4.0 and opportunities for energy sustainability', *Journal of Cleaner Production* **295**, 126427.
- Gillingham, K., Rapson, D. & Wagner, G. (2016), 'The rebound effect and energy efficiency policy', *Review of Environmental Economics and Policy* **10**(1), 68–88.
- Ginsberg, J., Mohebbi, M. H., Patel, R. S., Brammer, L., Smolinski, M. S. & Brilliant, L. (2009), 'Detecting influenza epidemics using search engine query data', *Nature* **457**(7232), 1012–1014.
- Giunta, A. & Trivieri, F. (2007), 'Understanding the determinants of information technology adoption: Evidence from Italian manufacturing firms', *Applied Economics* **39**(10), 1325–1334.

BIBLIOGRAPHY

- Glynn, A. N. & Quinn, K. M. (2010), 'An introduction to the augmented inverse propensity weighted estimator', *Political Analysis* **18**(1), 36–56.
- Gök, A., Waterworth, A. & Shapira, P. (2015), 'Use of web mining in studying innovation', *Scientometrics* **102**(1), 653–671.
- Griffith, R., Huergo, E., Mairesse, J. & Peters, B. (2006), 'Innovation and productivity across four European countries', *Oxford Review of Economic Policy* **22**(4), 483–498.
- Griliches, Z. (1980), 'R&D and the productivity slowdown', *American Economic Review* **70**(2), 343–348.
- Guerrieri, P., Luciani, M. & Meliciani, V. (2011), 'The determinants of investment in information and communication technologies', *Economics of Innovation and New Technology* **20**(4), 387–403.
- Gulen, H., Jens, C. E. & Page, T. B. (2021), 'The heterogeneous effects of default on investment: An application of Causal Forest in corporate finance'. <https://mays.tamu.edu/departments-of-finance/wp-content/uploads/sites/2/2021/03/jens.pdf> [Online; accessed 9 Apr. 2023].
- Hall, B. H., Jaffe, A. & Trajtenberg, M. (2005), 'Market value and patent citations', *RAND Journal of Economics* **36**(1), 16–38.
- Hall, B. H., Lotti, F. & Mairesse, J. (2013), 'Evidence on the impact of R&D and ICT investments on innovation and productivity in Italian firms', *Economics of Innovation and New Technology* **22**(3), 300–328.
- Haller, S. A. & Hyland, M. (2014), 'Capital-energy substitution: Evidence from a panel of Irish manufacturing firms', *Energy Economics* **45**, 501–510.
- Haller, S. A. & Siedschlag, I. (2011), 'Determinants of ICT adoption: Evidence from firm-level data', *Applied Economics* **43**(26), 3775–3788.
- Haque, N. U., Pesaran, M. H. & Sharma, S. (1999), Neglected heterogeneity and dynamics in cross-country savings regressions, IMF working papers 99/128, International Monetary Fund, Washington, DC.
- Heerink, N., Mulatu, A. & Bulte, E. (2001), 'Income inequality and the environment: Aggregation bias in environmental Kuznets curves', *Ecological Economics* **38**(3), 359–367.
- Hilty, L. M., Arnfalk, P., Erdmann, L., Goodman, J., Lehmann, M. & Wäger, P. A. (2006), 'The relevance of information and communication technologies for environmental sustainability – A prospective simulation study', *Environmental Modelling & Software* **21**(11), 1618–1629.

BIBLIOGRAPHY

- Hook, A., Court, V., Sovacool, B. K. & Sorrell, S. (2020), 'A systematic review of the energy and climate impacts of teleworking', *Environmental Research Letters* **15**(9), 093003.
- Horner, N. C., Shehabi, A. & Azevedo, I. L. (2016), 'Known unknowns: Indirect energy effects of information and communication technology', *Environmental Research Letters* **11**(10), 103001.
- Hosmer, Jr., D. W., Lemeshow, S. & Sturdivant, R. X. (2013), *Applied Logistic Regression*, 3 edn, John Wiley & Sons, Hoboken, NJ.
- Huang, G., He, L.-Y. & Lin, X. (2022), 'Robot adoption and energy performance: Evidence from Chinese industrial firms', *Energy Economics* **107**, 105837.
- IEA (2017), *Digitalization and Energy*, IEA, Paris. <https://www.iea.org/reports/digitalisation-and-energy> [Online; accessed 9 Apr. 2023].
- IEA (2019), *World Energy Outlook 2019*, IEA, Paris. <https://www.iea.org/reports/world-energy-outlook-2019> [Online; accessed 9 Apr. 2023].
- IEA (2020), *Energy Technology Perspectives 2020*, IEA, Paris. <https://www.iea.org/reports/energy-technology-perspectives-2020/technology-needs-in-long-distance-transport> [Online; accessed 9 Apr. 2023].
- IEA (2021), *Tracking Industry 2021*, IEA, Paris. <https://www.iea.org/reports/tracking-industry-2021> [Online; accessed 9 Apr. 2023].
- IEA (2022), *Energy Efficiency Indicators Data Explorer*, IEA, Paris. <https://www.iea.org/data-and-statistics/data-tools/energy-efficiency-indicators-data-explorer> [Online; accessed 9 Apr. 2023].
- ifo Institute (2022), 'Working from home declines only slightly in Germany'. <https://www.ifo.de/en/press-release/2022-09-05/working-home-declines-only-slightly-germany> [Online; accessed 9 Apr. 2023].
- Imbs, J. & Mejean, I. (2015), 'Elasticity optimism', *American Economic Journal: Macroeconomics* **7**(3), 43–83.
- infas (2018), *Mobilität in Deutschland – Ergebnisbericht*, Technical report, Bundesministerium für Verkehr und digitale Infrastruktur, Bonn. https://bmdv.bund.de/SharedDocs/DE/Anlage/G/mid-ergebnisbericht.pdf?__blob=publicationFile [Online; accessed 24 Mar. 2023].
- Insee (2021), 'The car remains the main mode of transport to go to work, even for short distances'. <https://www.insee.fr/en/statistiques/5019715> [Online; accessed 9 Apr. 2023].

BIBLIOGRAPHY

- Jaffe, A. B., Newell, R. G. & Stavins, R. N. (2003), Technological change and the environment, *in* K.-G. Mäler & J. R. Vincent, eds, 'Handbook of Environmental Economics', Vol. 1, Elsevier, Amsterdam, chapter 11, pp. 461–516.
- Jaller, M. & Pahwa, A. (2020), 'Evaluating the environmental impacts of online shopping: A behavioral and transportation approach', *Transportation Research Part D: Transport and Environment* **80**, 102223.
- James, G., Witten, D., Hastie, T. & Tibshirani, R. (2021), *An Introduction to Statistical Learning*, Springer, New York, NJ.
- Jens, C., Page, T. B. & Reeder III, J. (2021), 'Controlling for group-level heterogeneity in Causal Forest'. <https://ssrn.com/abstract=3907601> [Online; accessed 9 Apr. 2023].
- Johnson, M. S., Levine, D. I. & Toffel, M. W. (2020), Improving regulatory effectiveness through better targeting: Evidence from OSHA, Harvard Business School Technology & Operations Management Unit Working Paper 20-019, Harvard Business School, Boston, MA.
- Jones, B. A., Goodkind, A. L. & Berrens, R. P. (2022), 'Economic estimation of Bitcoin mining's climate damages demonstrates closer resemblance to digital crude than digital gold', *Scientific Reports* **12**(1), 14512.
- Jovanovic, B. & Rousseau, P. L. (2005), General purpose technologies, Working Paper 11093, National Bureau of Economic Research, Cambridge, MA.
- Kander, A., Malanima, P. & Warde, P. (2015), *Power to the people: Energy in Europe over the last five centuries*, Princeton University Press, Princeton, NJ.
- Katz, J. S. & Cothey, V. (2006), 'Web indicators for complex innovation systems', *Research Evaluation* **15**(2), 85–95.
- Kaus, W., Slavtchev, V. & Zimmermann, M. (2020), Intangible capital and productivity: Firm-level evidence from German manufacturing, IWH Discussion Papers 1/2020, Leibniz-Institut für Wirtschaftsforschung Halle (IWH), Halle (Saale).
- Kelly, B., Papanikolaou, D., Seru, A. & Taddy, M. (2021), 'Measuring technological innovation over the long run', *American Economic Review: Insights* **3**(3), 303–20.
- Khazzoom, J. D. (1980), 'Economic implications of mandated efficiency in standards for household appliances', *The Energy Journal* **1**(4), 21–40.
- Kinkel, S., Baumgartner, M. & Cherubini, E. (2022), 'Prerequisites for the adoption of AI technologies in manufacturing – Evidence from a worldwide sample of manufacturing companies', *Technovation* **110**, 102375.

BIBLIOGRAPHY

- Kinne, J. (2018), 'ARGUS - An automated robot for generic universal scraping'. <https://github.com/datawizard1337/ARGUS> [Online; accessed 6 Apr. 2023].
- Kinne, J. & Axenbeck, J. (2020), 'Web mining for innovation ecosystem mapping: A framework and a large-scale pilot study', *Scientometrics* **125**, 2011–2041.
- Kinne, J. & Lenz, D. (2021), 'Predicting innovative firms using web mining and deep learning', *PLoS One* **16**(4), e0249071.
- Kirbach, M. & Schmiedeberg, C. (2008), 'Innovation and export performance: Adjustment and remaining differences in East and West German manufacturing', *Economics of Innovation and New Technology* **17**(5), 435–457.
- Klomp, L. & Van Leeuwen, G. (2001), 'Linking innovation and firm performance: A new approach', *International Journal of the Economics of Business* **8**(3), 343–364.
- Knaus, M. C., Lechner, M. & Strittmatter, A. (2021), 'Machine learning estimation of heterogeneous causal effects: Empirical Monte Carlo evidence', *The Econometrics Journal* **24**(1), 134–161.
- Knittel, C. R. & Stolper, S. (2021), 'Machine learning about treatment effect heterogeneity: The case of household energy use', *AEA Papers and Proceedings* **111**, 440–444.
- Koetse, M. J., de Groot, H. L. F. & Florax, R. J. G. M. (2008), 'Capital-energy substitution and shifts in factor demand: A meta-analysis', *Energy Economics* **30**(5), 2236–2251.
- Kogan, L., Papanikolaou, D., Seru, A. & Stoffman, N. (2017), 'Technological innovation, resource allocation, and growth', *Quarterly Journal of Economics* **132**(2), 665–712.
- Kopp, T. & Lange, S. (2019), The climate effect of digitalization in production and consumption in OECD countries, in 'Proceedings of the 6th International Conference on ICT for Sustainability'.
- Kraemer, K. L. (1982), 'Telecommunications/ transportation substitution and energy conservation: Part 1', *Telecommunications Policy* **6**(1), 39–59.
- Kraemer, K. L. & King, J. L. (1982), 'Telecommunications/ transportation substitution and energy conservation Part 2', *Telecommunications Policy* **6**(2), 87–99.
- Kube, R., von Graevenitz, K., Löschel, A. & Massier, P. (2019), 'Do voluntary environmental programs reduce emissions? EMAS in the German manufacturing sector', *Energy Economics* **84**, 104558.

BIBLIOGRAPHY

- Kusiak, A. (2018), 'Smart manufacturing', *International Journal of Production Research* **56**(1-2), 508–517.
- Labandeira, X., Labeaga, J. M. & López-Otero, X. (2017), 'A meta-analysis on the price elasticity of energy demand', *Energy Policy* **102**, 549–568.
- Lachenmaier, S. & Wößmann, L. (2006), 'Does innovation cause exports? Evidence from exogenous innovation impulses and obstacles using German micro data', *Oxford Economic Papers* **58**(2), 317–350.
- Lagomarsino, E. (2020), 'Estimating elasticities of substitution with nested CES production functions: Where do we stand?', *Energy Economics* **88**, 104752.
- Lange, S., Pohl, J. & Santarius, T. (2020), 'Digitalization and energy consumption. Does ICT reduce energy demand?', *Ecological Economics* **176**, 106760.
- Lechner, M. (2015), Treatment effects and panel data, in B. H. Baltagi, ed., 'The Oxford Handbook of Panel Data', Oxford University Press, Oxford, chapter 9, p. 257–284.
- Lenz, D. & Winker, P. (2020), 'Measuring the diffusion of innovations with paragraph vector topic models', *PLoS One* **15**(1), e0226685.
- Löschel, A., Lutz, B. J. & Managi, S. (2019), 'The impacts of the EU ETS on efficiency and economic performance – An empirical analyses for German manufacturing firms', *Resource and Energy Economics* **56**, 71–95.
- Lutz, B. J. (2016), Emissions trading and productivity: Firm-level evidence from German manufacturing, ZEW Discussion Paper No. 16-067, ZEW - Leibniz Centre for European Economic Research, Mannheim.
- Lutz, B. J., Massier, P., Sommerfeld, K. & Löschel, A. (2017), Drivers of energy efficiency in German manufacturing: A firm-level stochastic frontier analysis, ZEW Discussion Paper No. 17-068, ZEW - Leibniz Centre for European Economic Research, Mannheim.
- Ma, H., Oxley, L., Gibson, J. & Kim, B. (2008), 'China's energy economy: Technical change, factor demand and interfactor/interfuel substitution', *Energy Economics* **30**(5), 2167–2183.
- Mairesse, J. & Mohnen, P. (2010), Using innovation surveys for econometric analysis, in B. H. Hall & N. Rosenberg, eds, 'Handbook of the Economics of Innovation', Vol. 2, North-Holland, Amsterdam, chapter 26, pp. 1129–1155.

BIBLIOGRAPHY

- Majeed, M. T. (2018), 'Information and communication technology (ICT) and environmental sustainability in developed and developing countries', *Pakistan Journal of Commerce and Social Sciences* **12**(3), 758–783.
- Marz, W. & Şen, S. (2022), 'Does telecommuting reduce commuting emissions?', *Journal of Environmental Economics and Management* **116**, 102746.
- Masanet, E., Shehabi, A., Lei, N., Smith, S. & Koomey, J. (2020), 'Recalibrating global data center energy-use estimates', *Science* **367**(6481), 984–986.
- McDermott, G. R. & Hansen, B. (2021), Labor reallocation and remote work during Covid-19: Real-time evidence from GitHub, Working Paper 29598, National Bureau of Economic Research, Cambridge, MA.
- McNemar, Q. (1947), 'Note on the sampling error of the difference between correlated proportions or percentages', *Psychometrika* **12**(2), 153–157.
- Michaels, G., Natraj, A. & Van Reenen, J. (2014), 'Has ICT polarized skill demand? Evidence from eleven countries over twenty-five years', *The Review of Economics and Statistics* **96**(1), 60–77.
- Miller, S. (2020), 'Causal Forest estimation of heterogeneous and time-varying environmental policy effects', *Journal of Environmental Economics and Management* **103**, 102337.
- Mohri, M., Rostamizadeh, A. & Talwalkar, A. (2018), *Foundations of Machine Learning*, 2 edn, MIT Press, Cambridge, MA.
- Mokhtarian, P. L. (1990), 'A typology of relationships between telecommunications and transportation', *Transportation Research Part A: General* **24**(3), 231–242.
- Mrozik, W., Rajaeifar, M. A., Heidrich, O. & Christensen, P. (2021), 'Environmental impacts, pollution sources and pathways of spent lithium-ion batteries', *Energy & Environmental Science* **14**, 6099–6121.
- Muench, S., Stoermer, E., Jensen, K., Asikainen, T., Salvi, M. & Scapolo, F. (2022), Towards a green & digital future, EUR 31075 EN, Publications Office of the European Union, Luxembourg.
- Mullainathan, S. & Spiess, J. (2017), 'Machine learning: An applied econometric approach', *Journal of Economic Perspectives* **31**(2), 87–106.
- Mundlak, Y. (1978), 'On the Pooling of Time Series and Cross Section Data', *Econometrica* **46**(1), 69–85.

BIBLIOGRAPHY

- Nathan, M. & Rosso, A. (2022), 'Innovative events: Product launches, innovation and firm performance', *Research Policy* **51**(1), 104373.
- Nie, X. & Wager, S. (2021), 'Quasi-oracle estimation of heterogeneous treatment effects', *Biometrika* **108**(2), 299–319.
- OECD (2017a), *The Next Production Revolution: Implications for Governments and Business*, OECD Publishing, Paris. <https://doi.org/10.1787/9789264271036-en> [Online; accessed 9 Apr. 2023].
- OECD (2017b), *OECD Digital Economy Outlook 2017*, OECD Publishing, Paris. <https://www.oecd-ilibrary.org/content/publication/9789264276284-en> [Online; accessed 9 Apr. 2023].
- OECD/Eurostat (2019), *Oslo Manual 2018: Guidelines for Collecting, Reporting and Using Data on Innovation, 4th Edition*, OECD publishing, Paris & Luxembourg. <https://www.oecd-ilibrary.org/content/publication/9789264304604-en> [Online; accessed 9 Apr. 2023].
- O'Neill, E. & Weeks, M. (2019), 'Causal Tree estimation of heterogeneous household response to time-of-use electricity pricing schemes'. <https://arxiv.org/abs/1810.09179> [Online; accessed 9 Apr. 2022].
- Patterson, M. G. (1996), 'What is energy efficiency? Concepts, indicators and methodological issues', *Energy Policy* **24**(5), 377–390.
- Persson, J., Parie, J. F. & Feuerriegel, S. (2021), 'Monitoring the Covid-19 epidemic with nationwide telecommunication data', *Proceedings of the National Academy of Sciences* **118**(26), e2100664118.
- Pesaran, H., Smith, R. & Im, K. S. (1996), Dynamic linear models for heterogenous panels, in L. Mátyás & P. Sevestre, eds, 'The Econometrics of Panel Data: A Handbook of the Theory with Applications', Vol. 33, Springer Netherlands, Dordrecht, chapter 8, pp. 145–195.
- Peters, B. & Rammer, C. (2013), Innovation panel surveys in Germany, in Fred Gault, ed., 'Handbook of Innovation Indicators and Measurement', Edward Elgar Publishing, Cheltenham, chapter 6, pp. 135–177.
- Preist, C., Schien, D. & Shabajee, P. (2019), Evaluating sustainable interaction design of digital services: The case of YouTube, in 'Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems', pp. 1–12.
- Prest, B. C. (2020), 'Peaking interest: How awareness drives the effectiveness of time-of-use electricity pricing', *Journal of the Association of Environmental and Resource Economists* **7**(1), 103–143.

BIBLIOGRAPHY

- Pritchett, L. (2000), 'Understanding patterns of economic growth: Searching for hills among plateaus, mountains, and plains', *World Bank Economic Review* **14**(2), 221–250.
- Pukelis, L. & Stanciauskas, V. (2019), 'Using internet data to compliment traditional innovation indicators'. <https://www.ippapublicpolicy.org/file/paper/5d073ea805eb6.pdf> [Online; accessed 9 Apr. 2023].
- Rammer, C., Behrens, V., Doherr, T., Hud, M., Köhler, M., Krieger, B., Peters, B., Schubert, T., Trunschke, M. & von der Burg, J. (2019), *Innovationen in der deutschen Wirtschaft: Indikatorenbericht zur Innovationserhebung 2018*, Technical report, ZEW - Leibniz Centre for European Economic Research, Mannheim.
- Ren, S., Hao, Y., Xu, L., Wu, H. & Ba, N. (2021), 'Digitalization and energy: How does internet development affect China's energy consumption?', *Energy Economics* **98**, 105220.
- Reveiu, A., Vasilescu, M. D. & Banica, A. (2022), 'Digital divide across the European Union and labour market resilience'. *Regional Studies*, 1–15.
- Rhoades, S. A. (1993), 'The Herfindahl-Hirschman Index', *Federal Reserve Bulletin* **Mar**, 188–189.
- Richter, P. M. & Schiersch, A. (2017), 'CO2 emission intensity and exporting: Evidence from firm-level data', *European Economic Review* **98**, 373–391.
- Robinson, P. M. (1988), 'Root-n-consistent semiparametric regression', *Econometrica* **56**(4), 931–954.
- Robinson, W. S. (1950), 'Ecological correlations and the behavior of individuals', *American Sociological Review* **15**(3), 351–357.
- Rubin, D. B. (1974), 'Estimating causal effects of treatments in randomized and non-randomized studies.', *Journal of Educational Psychology* **66**(5), 688.
- Sadorsky, P. (2012), 'Information communication technology and electricity consumption in emerging economies', *Energy Policy* **48**, 130–136.
- Salomon, I. (1986), 'Telecommunications and travel relationships: A review', *Transportation Research Part A: General* **20**(3), 223–238.
- Schulte, P., Welsch, H. & Rexhäuser, S. (2016), 'ICT and the demand for energy: Evidence from OECD countries', *Environmental and Resource Economics* **63**(1), 119–146.
- Shaw, N., Eschenbrenner, B. & Baier, D. (2022), 'Online shopping continuance after Covid-19: A comparison of Canada, Germany and the United States', *Journal of Retailing and Consumer Services* **69**, 103100.

BIBLIOGRAPHY

- Shephard, R. W. (1953), *Cost and Production Functions*, Princeton University Press, Princeton, NJ.
- Siikavirta, H., Punakivi, M., Kärkkäinen, M. & Linnanen, L. (2002), 'Effects of e-commerce on greenhouse gas emissions: A case study of grocery home delivery in Finland', *Journal of Industrial Ecology* 6(2), 83–97.
- Solow, J. L. (1987), 'The capital-energy complementarity debate revisited', *The American Economic Review* 77(4), 605–614.
- Solow, R. M. (1956), 'A contribution to the theory of economic growth', *The Quarterly Journal of Economics* 70(1), 65–94.
- Solow, R. M. (1957), 'Technical change and the aggregate production function', *The review of Economics and Statistics* pp. 312–320.
- Stiroh, K. J. (2005), 'Reassessing the impact of IT in the production function: A meta-analysis and sensitivity tests', *Annales d'Économie et de Statistique* (79/80), 529–561.
- Stobbe, L., Richter, N., Quaack, M., Knüfermann, K., Druschke, J., Fahland, M., Höller, V. W., Wahry, N., Zedel, H., Kaiser, M., Hoffmann, S., Töpfer, M. & Nissen, N. (2023), Umweltbezogene Technikfolgenabschätzung Mobilfunk in Deutschland, Texte 26/2023, Umweltbundesamt, Dessau-Roßlau.
- Stoll, C., Klaaßen, L. & Gallersdörfer, U. (2019), 'The carbon footprint of Bitcoin', *Joule* 3(7), 1647–1661.
- Tacchella, A., Napoletano, A. & Pietronero, L. (2020), 'The language of innovation', *PLoS One* 15(4), 1–20.
- Taneja, S. & Mandys, F. (2022), 'The effect of disaggregated information and communication technologies on industrial energy demand', *Renewable Sustainable Energy Reviews* 164, 112518.
- Theil, H. (1971), *Principles of Econometrics*, John Wiley & Sons, Hoboken, NJ.
- Thoben, K.-D., Wiesner, S. A. & Wuest, T. (2017), "'Industry 4.0" and smart manufacturing – A review of research issues and application examples', *International Journal of Automation Technology* 11(1), 4–19.
- UNFCCC (2015), 'Paris Agreement', United Nations Treaty Collection. <https://unfccc.int/resource/docs/2015/cop21/eng/l09r01.pdf> [Online; accessed 9 Apr. 2023].
- Valente, M. (2023), 'Policy evaluation of waste pricing programs using heterogeneous causal effect estimation', *Journal of Environmental Economics and Management* 117, 102755.

BIBLIOGRAPHY

- Van der Werf, E. (2008), 'Production functions for climate policy modeling: An empirical analysis', *Energy Economics* **30**(6), 2964–2979.
- Van Reenen, J. (2011), 'Wage inequality, technology and trade: 21st century evidence', *Labour Economics* **18**(6), 730–741.
- Wager, S. & Athey, S. (2018), 'Estimation and inference of heterogeneous treatment effects using random forests', *Journal of the American Statistical Association* **113**(523), 1228–1242.
- Wagner, J. (2021), With a little help from my website: Firm survival and web presence in times of Covid-19 – Evidence from 10 European countries, Working Paper Series in Economics No. 399, Leuphana Universität Lüneburg, Institut für Volkswirtschaftslehre, Lüneburg.
- Wang, E.-Z., Lee, C.-C. & Li, Y. (2022), 'Assessing the impact of industrial robots on manufacturing energy intensity in 38 countries', *Energy Economics* **105**, 105748.
- Weber, T., Bertschek, I., Ohnemus, J. & Ebert, M. (2018), 'Monitoring-Report Wirtschaft DIGITAL 2018'. <https://www.bmwk.de/Redaktion/DE/Publikationen/Digitale-Welt/monitoring-report-wirtschaft-digital-2018-langfassung.pdf> [Online; accessed 21 Mar. 2023].
- Welsch, H. & Oxen, C. (2005), 'The determinants of aggregate energy use in West Germany: Factor substitution, technological change, and trade', *Energy Economics* **27**(1), 93–111.
- Wen, H., Lee, C.-C. & Song, Z. (2021), 'Digitalization and environment: How does ICT affect enterprise environmental performance?', *Environmental Science and Pollution Research* **28**(39), 54826–54841.
- Wiese, A., Toporowski, W. & Zielke, S. (2012), 'Transport-related CO2 effects of online and brick-and-mortar shopping: A comparison and sensitivity analysis of clothing retailing', *Transportation Research Part D: Transport and Environment* **17**(6), 473–477.
- Williams, E. (2011), 'Environmental effects of information and communications technologies', *Nature* **479**(7373), 354–358.
- Wooldridge, J. M. (2002), *Econometric Analysis of Cross Section and Panel Data*, MIT press, Cambridge, MA.
- Wurlod, J.-D. & Noailly, J. (2018), 'The impact of green innovation on energy intensity: An empirical analysis for 14 industrial sectors in OECD countries', *Energy Economics* **71**, 47–61.

BIBLIOGRAPHY

- Wöhner, F. (2022), 'Work flexibly, travel less? The impact of telework and flextime on mobility behavior in Switzerland', *Journal of Transport Geography* **102**, 103390.
- Xu, Q., Zhong, M. & Li, X. (2022), 'How does digitalization affect energy? International evidence', *Energy Economics* **107**, 105879.
- Zellner, A. (1962), 'An efficient method of estimating seemingly unrelated regressions and tests for aggregation bias', *Journal of the American Statistical Association* **57**(298), 348–368.
- ZEW Mannheim (2022), 'One in three firms views remote work more positively than before the pandemic'. <https://www.zew.de/en/press/latest-press-releases/one-in-three-firms-views-remote-work-more-positively-than-before-the-pandemic> [Online; accessed 17 Feb. 2023].
- Zhang, C. & Liu, C. (2015), 'The impact of ICT industry on CO2 emissions: A regional analysis in China', *Renewable and Sustainable Energy Reviews* **44**, 12–19.
- Zhang, W., Gu, F. & Guo, J.-F. (2019), 'Can smart factories bring environmental benefits to their products? A case study of household refrigerators', *Journal of Industrial Ecology* **23**(6), 1381–1395.
- Zhang, X. & Wei, C. (2022), 'The economic and environmental impacts of information and communication technology: A state-of-the-art review and prospects', *Resources, Conservation and Recycling* **185**, 106477.
- Zhang, Y., Ma, S., Yang, H., Lv, J. & Liu, Y. (2018), 'A big data driven analytical framework for energy-intensive manufacturing industries', *Journal of Cleaner Production* **197**, 57–72.

Appendices

Appendix A

Innovation Indicators Based on Firm Websites – Which Website Characteristics Predict Firm-Level Innovation Activity?

A.1 Comparison of the Distribution Between the MIP and the Applied Subsample




Figure A.1: Firm distribution based on the number of employees.

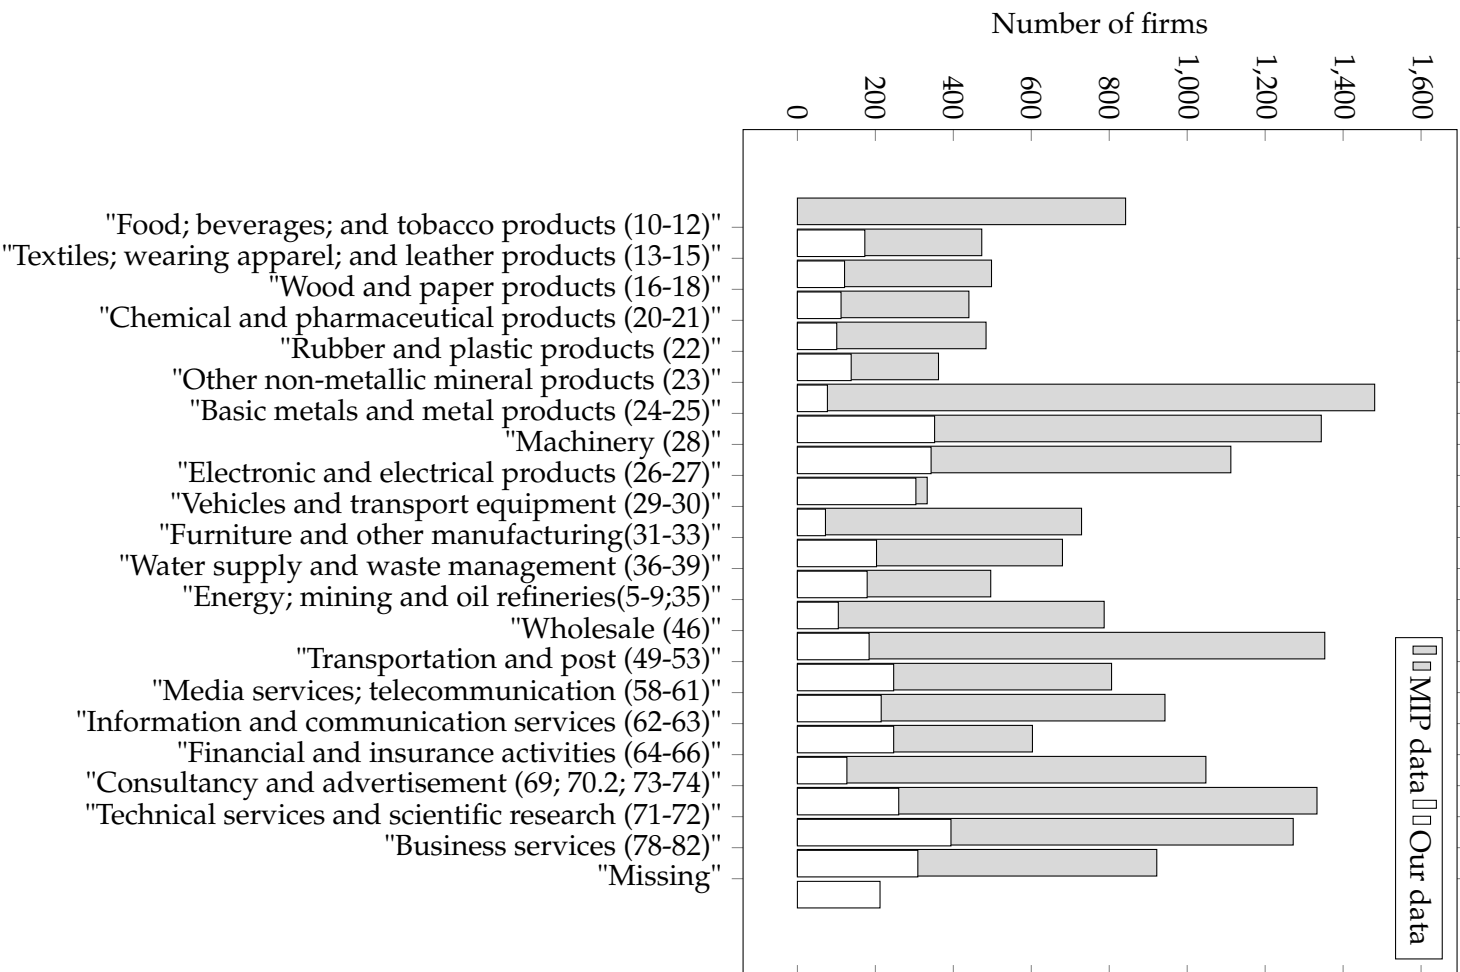


Figure A.2: Firm distribution for economic sectors based on 2 digit NACE codes.

A.2 List of Emerging Technology Terms Used in the Conducted Keyword Search

English terms: agricultural robot, closed ecological systems, cultured meat, precision agriculture, vertical farming, micro air vehicle, neural-sensing headset, four-dimensional printing, arcology, aerogel, bioplastic, conductive polymers, cryogenic treatment, fullerene, graphene, lab-on-a-chip, magnetorheological fluid, metamaterials, metal foam, multi-function structures, nanomaterials, carbon nanotube, quantum dots, superalloy, synthetic diamond, translucent concrete, 3D displays, ferroelectric liquid crystal display, holography, interferometric modulator display, laser video displays, OLED displays, micro LED displays, telescopic pixel display, time-multiplexed optical shutter, volumetric display, biometrics, digital scent technology, electronic nose, e-textiles, flexible electronics, memristor, molecular electronics, nano electro mechanical systems, spintronics, thermal copper pillar bump, three-dimensional integrated circuit, concentrated solar power, electric double-layer capacitor, flywheel energy storage, grid energy storage, home fuel cell, lithium iron phosphor battery, lithium-sulfur battery, magnesium battery, nanowire battery, ocean thermal energy conversion, smart grid, vortex engine, wireless energy transfer, zero-energy building, computer-generated imagery, virtual reality, ultra-high-definition television, 5G cellular communications, artificial general intelligence, augmented reality, blockchain, carbon nanotube field-effect transistor, civic technology, cryptocurrency, exascale computing, gesture recognition, internet of things, emerging memory technologies, emerging magnetic data storage technologies, fourth generation optical discs, holographic data storage, general purpose computing on graphics processing units, exocortex, machine translation, machine vision, mobile collaboration, nano radio, optical computing, quantum computing, quantum cryptography, radio-frequency identification, semantic web, smart speaker, software-defined radio, speech recognition, subvocal recognition, hybrid forensics, body implants, prosthesis, cryonics, de-extinction, genetic engineering of organisms and viruses, suspended animation, artificial hibernation, immunotherapy/oncology, nano medicines, nano sensors, oncolytic viruses, personalized medicine, whole genome sequencing, robotic surgery, stem cell treatments, synthetic biology, synthetic genomics, tissue engineering, tricorder, brain-computer interface, neuro informatics, electro encephalography, neuro prosthetics, caseless ammunition, directed energy weapon, electro laser, electromagnetic weapons, electrothermal-chemical technology, green bullet, laser weapon, particle beam weapon, sonic weapon, stealth technology, vortex ring gun, wireless long-range electric shock weapon, artificial gravity, stasis chamber, inflatable space habitat, miniaturized satellite, android, gynoid, nanorobotics, powered exoskeleton,

self-reconfiguring modular robot, unmanned vehicle, airless tire, alternative fuel vehicle, electro hydrodynamic propulsion, flying car, fusion rocket, hoverbike, jetpack, backpack helicopter, maglev train, vactrain, magnetic levitation, mass driver, personal rapid transit, physical internet, scooter-sharing system, propellant depot, reusable launch system, space elevator, spaceplane, supersonic transport, vehicular communication systems.

German terms: Agrarroboter, geschlossenes ökologisches System, Zuchtfleisch, Präzisionslandwirtschaft, vertikale Landwirtschaft, Mikro-Luftfahrzeug, neuronales Headset, vierdimensionales Drucken, Arkologie, Aerogel, Bio-Kunststoff, leitfähige Polymere, kryogene Behandlung, Fulleren, Graphen, Labor auf einem Chip, magnetorheologische Flüssigkeit, Metamaterialien, Metallschaum, Multifunktionsstrukturen, Nanomaterialien, Kohlenstoffnanoröhre, Quantenpunkte, Superlegierung, synthetischer Diamant, durchsichtiger Beton, 3D-Display, ferroelektrische Flüssigkristallanzeige, Holographie, interferometrische Modulatoranzeige, Laser-Video-Display, OLED Display, Mikro-LED Display, Teleskop-Pixelanzeige, zeitgemultiplexer optischer Verschluss, volumetrische Anzeige, Biometrie, digitale Dufttechnologie, elektronische Nase, E-Textil, flexible Elektronik, Memoristor, molekulare Elektronik, nanoelektromechanisches System, Spintronik, Thermo-Kupfer-Säulen-Stoß, dreidimensionale integrierte Schaltung, konzentrierte Solarenergie, elektrischer Doppelschicht-Kondensator, Schwungradspeicherung, Speicherung von Netzenergie, Heim-Brennstoffzelle, Lithium-Eisen-Phosphor-Batterie, Lithium-Schwefel-Batterie, Magnesium-Batterie, Nanodraht-Batterie, Ozean-Thermische Energieumwandlung, intelligentes Netz, Vortex-Motor, drahtlose Energie-Übertragung, Nullenergiehaus, computergeneriertes Bild, virtuelle Realität, hochauflösendes Fernsehen, 5G zellulare Kommunikation, künstliche Intelligenz, erweiterte Realität, Blockchain, Kohlenstoffnanoröhren-Feldeffekttransistor, zivile Technik, Kryptowährung, Exascale-Computing, Gestenerkennung, Internet der Dinge, neue Speichertechnologie, neue magnetische Speichertechnologie, optische Platten der vierten Generation, holografischer Speicher, allgemeines Rechnen auf Grafikprozessoren, Exokortex, maschinelle Übersetzung, maschinelles Sehen, mobile Zusammenarbeit, Nano-Funk, optische Datenverarbeitung, Quantencomputer, Quantenkryptographie, Radiofrequenz-Identifikation, semantisches Web, intelligenter Lautsprecher, Software-definiertes Radio, Spracherkennung, subvokale Erkennung, Hybrid-Forensik, Körperimplantat, Kryonik, Wiederbelebung ausgestorbener Tierarten, Gentechnik, verzögerte Reanimation, künstlicher Winterschlaf, Immuntherapie/-onkologie, Nanomedizin, Nanosensoren, onkolytische Viren, individualisierte Medizin, whole

genome sequencing, Roboterchirurgie, Stammzellentherapie, synthetische Biologie, synthetische Genomik, Gewebezüchtung, Tricorder, Gehirn-Computer-Schnittstelle, Neuroinformatik, Elektroenzephalographie, Neuroprothetik, hülsenlose Munition, gerichtete Energiewaffe, Elektro-Laser, elektromagnetische Waffen, elektrothermisch-chemische Technologie, grünes Geschoss, Laser-Waffe, Strahlenwaffe, Schallwaffe, Tarntechnologie, Wirbelringkanone, Elektroschockwaffe, künstliche Schwerkraft, Stasiskammer, aufblasbares Weltraum-Habitat, Miniatursatellit, Android, Nanorobotik, Exoskelett, selbstkonfigurierender Roboter, unbemanntes Fahrzeug, luftlose Reifen, Fahrzeug mit alternativen Kraftstoffen, Elektrodynamischer Antrieb, Fluidik, Fusionsrakete, Schwebefahrrad, Jetpack, Rucksackhelikopter, Magnetschwebbahn, Vactrain, magnetische Schwebetechnik, Massenantrieb, Personal Rapid Transit, physisches Internet, Roller-Sharing-System, fliegendes Treibstofflager, wiederverwendbares Startsystem, Raumaufzug, Raumflugzeug, Überschalltransport, Fahrzeugkommunikationssystem.

A.3 Detailed Information on the Calculation of Features

Text-Based Features

1) **Texts** – To identify the most relevant terms when predicting a firm's innovation status, we transform the scraped texts into a format that allows us to do mathematical operations: We convert the website texts into a term-document matrix, e.g., see Baeza-Yates & Ribeiro-Neto (1999), Blei et al. (2003), which is a matrix that counts the frequency of terms that occur in a collection of documents (websites in this particular case). Every column represents a document and a row represents a word from a predefined vocabulary space. Accordingly, every cell counts how often a particular word appears in a particular document. We define our vocabulary space as the 5,000 most frequent words in our entire training text corpus. Before we calculate the term-document matrix, we conduct the following preprocessing steps. First, we merge all scraped subpages related to a single firm and delete irrelevant subpages (imprints, information about cookies or texts that are prescribed by law) by using the gold standard approach based on a supervised machine learning regression model (see Kinne & Lenz 2021). Also, every word is converted into lower case and lemmatised by means of the Python package *spacy*. We exclude punctuation as well as English and German stop words (word lists are derived from the Python package *nltk*). Additionally, we manipulate the term-frequency counts by the TF-IDF scheme (Baeza-Yates & Ribeiro-Neto 1999), as it usually improves predictions. Therefore, each document is tokenised and the term-document frequency is calculated by means of the *TfidfVectorizer* algorithm from *scikit-learn*.

2) **Emerging technology terms** – To capture firms that mention emerging technologies, we conduct a keyword search in which we calculate whether a technology from Wikipedia’s list of emerging technologies appears on a firm’s website using all subpages and the entire vocabulary as well as the Python package *regex*.¹³⁷ We only search for a selection of technologies that are in a research, development, diffusion or commercialisation stage, as it is a criterion for an innovation to be brought into use. A detailed list of all used keywords is provided in Appendix A.2. The feature *emerging_tech* is a dummy variable that captures whether an emerging technology term appears on a firm website.

3) **Latent patterns** – Latent patterns on a website, which might reveal a firm’s innovation status, are captured by the latent Dirichlet allocation model (LDA) (see Blei et al. 2003). The LDA algorithm assumes that a document consists of a set of topics, while every topic is a distribution of words. By linking each word in a document to a topic and iteratively improving assignments, the algorithm learns the distribution of topics in the text corpus as well as the distribution of words related to each topic. Moreover, after applying the LDA algorithm, the topic-document matrix shows how much every topic contributes to a document (website). We do not want our topic model to be exclusively valid for our sample. Hence, we calibrate our topics on a separate sample, which consists of 32,276 websites of firms observed in the MUP 2019 but not in the MIP 2019. We apply the same text preprocessing to it as to our MIP sample, but with two differences. First, we use a larger vocabulary space (15,000 most frequent words). Second, we do not manipulate word counts by means of the TF-IDF formula, but generate a TF-IDF stop word dictionary excluding words with a lower sum of TF-IDF scores than three within the LDA corpus. The latter is applied to ensure that rather words that are characteristic for particular websites are considered. Also, to improve our model performance, we delete all words that appear less than 50 times and in more than 90 per cent of all documents in the LDA corpus. We use the *TfidfVectorizer* to calculate the stop word dictionary. This dictionary as well as the *CountVectorizer* from *scikit-learn* is applied to generate a term-document matrix for our LDA sample. A term-document matrix for the MIP sample is calculated in the same manner. The Python package *scikit-learn* is used to train the LDA model. In the standard LDA approach, the number of topics needs to be defined. To solve this issue, we apply the grid-search technique to optimize the number of topics. For this, we use the *GridSearchCV* algorithm from *scikit-learn*. It is evaluated which model parameter combination leads to the best result according to the log likelihood. We conduct a grid-search over different values for the ‘number of topics’-parameter as well as the document-topic prior. We try 200, 180, 150, 250 topics and values of 0.05, 0.1 for the document-topic prior. The optimal number of

¹³⁷See https://en.wikipedia.org/wiki/List_of_emerging_technologies [Online; accessed 16 Aug. 2018].

topics is 150, the highest log-likelihood is achieved with a document-topic prior of 0.1. After fitting the LDA model with the separate sample, the topic distribution for each website in our MIP sample is predicted (*LDA topic*) and used in our Random Forest models, i.e., the predicted topic share in a document for each topic is used as a feature.

4) **Topic popularity index** – The topic popularity index is the sum of document-topic probabilities weighted by the relative frequency each topic appears in the entire text corpus (*pop_score*). A topic is considered to appear in a document if the document-topic probability is larger than 2%.

5) **Language classification** – The export orientation of a website might provide information about a firm’s innovation status. English is worldwide the most widely spoken language by the total number of speakers. Therefore, it is quite likely that firms with international customers describe their products in English. We measure the share of subpages in English language, as well as all other languages except German to approximate the export orientation of a firm (*english_language, other_lang*). For the language classification of subpages, we apply the Python package *langdetect*.

6) **Share of numbers** – We also test whether the share of numbers in the total text length per document relates to the innovation status. The share was calculated by the ratio of digits within a string (document). For example, the text ‘This book costs 500 dollars.’ has a ratio of 3/28, i.e., 10.7 per cent. The corresponding variable is named *share_numbers*.

7) **Flesch-reading-ease score** – The Flesch-Reading-Ease score is a metric used to assess the complexity of texts. The main idea for the index is that short words and short sentences are easier for readers to understand. The Python package *ReadabilityCalculator* was used to calculate the score.¹³⁸ The full definition can be found in Flesch (1948) and the corresponding variable is named *flesch_score*.

Meta Information Features

8) **Website size** – Approximating firm size might help to predict a firm’s innovation status. For example, Kinne & Axenbeck (2020) show that the number of subpages correlates with firm size and larger firms tend to be more likely to implement an innovation. Hence, we use the number of subpages as a feature to predict a firm’s innovation status (*nr_pages*). One problem related to this feature is that it is truncated at 50 subpages due to the scraping limit of the web-scrapers. However, as only 1.5 per cent of our observations exceed the scraping limit, we do not see a severe problem here. Moreover, we use a Random Forest model that selects cut-off points for splitting. Hence, it can cope with truncated features. We additionally analyse

¹³⁸Retrieved from <https://pypi.org/project/ReadabilityCalculator/> [Online; accessed 15 Apr. 2023].

to what extent the number of characters per website (*text_length*), which might also relate to firm size, informs about the firm's innovation status.

9) **Loading time** – This feature serves as a proxy for a firm's hardware structure. A website's loading time (*load_time*) is determined by a http or https request. The time from sending the request until the arrival of the response is measured. Servers which are far away or which only process the requests slowly (e.g., due to bad hardware or an overload) have a higher loading time (in milliseconds). However, it should be noted that the IT infrastructure can also be outsourced to professional hosting firms. We retrieved the loading time by means of the the Python packages *requests* and *time*. The latter is a standard Python library.

10) **Mobile version** – For each website, it is retrieved whether a version for mobile end user devices exists. A Google API is used to extract this information from the websites. The data is delivered as JSON object. Within the delivered data, the binary variable "score" within the data structure "usability" is used (*mobile_version*). It indicates Google's mobile version passing score. The Python packages *json*, *mechanize*, *socket* and *urllib* are used for this exercise.

11) **Website age** – To determine the website age, we use web.archive.org. The website includes an Internet archive that allows to look at websites at earlier stages. We wrote a small program that automatically goes to web.archive.org and searches for the first entry of a particular website. This characteristic serves as a proxy for the digital age of a firm (*domain_purchase_year_proxy*). Our program uses the Python package *urllib*.

Network Features

12) **Centrality** – Relationships with other firms might also link to a firm's innovation status. If a firm is related to another firm, it is likely that the firm will refer on its website to it. Hence, to capture relationships with other firms, the sum of outgoing (*outgoing_links*) and incoming (*incoming_links*) hyperlinks to other firms is observed. Outgoing hyperlinks are measured by the number of external links on a firm website. We measure incoming hyperlinks by counting how often firms which are listed in the entire MUP refer to a particular firm. Additionally, a directed graph is constructed. Here, a vertex represents a firm and an edge a hyperlink from one firm to another. The Pagerank centrality measure is calculated with the Python package *igraph*¹³⁹ and the function "pagerank". The default parameters are used and the resulting variable is called *pagerank_index*.

13) **Social media** – The use of social media could also be correlated with the firm's innovation status. Therefore, the sum of hyperlinks to the websites Facebook, Instagram, Twitter, YouTube, Kununu, LinkedIn, XING, GitHub, Flickr, and Vimeo

¹³⁹See <https://igraph.org> [Online; accessed 15 Apr. 2023].

is counted and used as another feature (*social_media*). This is calculated by means of *regex* again.

14) **Bridges** – An undirected graph is constructed, as well. A bridge is an edge of a graph whose removal increases the number of connected components. For each vertex, we count the number of times it is part of a bridge. The Python package *networkx* (<https://networkx.github.io> [Online; accessed 15 Apr. 2023]) and the function "bridges" is used to calculate the bridges and the described measure. The resulting variable is named *bridge_index*.

A.4 Most Relevant Features for Each ‘All’ Feature Model¹⁴⁰

Table A.1: Most relevant features for product innovators.

Model	Top 100 most relevant features
Product innovators	‘LDA topic 35’, ‘english_language’, ‘word: system’, ‘text_length’, ‘LDA topic 134’, ‘nr_subpages’, ‘word: software’, ‘LDA topic 65’, ‘word: to develop (transl.)’, ‘word: application (transl.)’, ‘LDA topic 105’, ‘word: test’, ‘LDA topic 7’, ‘word: product (transl.)’, ‘incoming_links’, ‘word: worldwide (transl.)’, ‘word: innovative (transl.)’, ‘LDA topic 98’, ‘domain_purchase_year_proxy’, ‘word: version’, ‘word: innovative’, ‘LDA topic 41’, ‘LDA topic 20’, ‘word: technology (transl.)’, ‘share_numbers’, ‘word: sensor’, ‘LDA topic 127’, ‘social_media’, ‘flesch_score’, ‘word: development (transl.)’, ‘emerging_tech’, ‘LDA topic 34’, ‘word: technology’, ‘LDA topic 38’, ‘LDA topic 96’, ‘LDA topic 75’, ‘LDA topic 46’, ‘pop_score’, ‘LDA topic 39’, ‘word: automatic (transl.)’, ‘LDA topic 101’, ‘LDA topic 70’, ‘LDA topic 78’, ‘LDA topic 84’, ‘LDA topic 128’, ‘outgoing_links’, ‘LDA topic 148’, ‘LDA topic 97’, ‘word: to optimize (transl.)’, ‘word: software development (transl.)’, ‘word: application (transl.)’, ‘LDA topic 119’, ‘LDA topic 36’, ‘word: component (transl.)’, ‘LDA topic 69’, ‘load_time’, ‘LDA topic 52’, ‘LDA topic 56’, ‘LDA topic 60’, ‘LDA topic 143’, ‘word: digital’, ‘LDA topic 8’, ‘LDA topic 113’, ‘LDA topic 120’, ‘word: complex (transl.)’, ‘LDA topic 53’, ‘LDA topic 138’, ‘LDA topic 144’, ‘LDA topic 51’, ‘LDA topic 15’, ‘LDA topic 19’, ‘word: support’, ‘LDA topic 103’, ‘LDA topic 106’, ‘word: user (transl.)’, ‘LDA topic 57’, ‘LDA topic 107’, ‘LDA topic 49’, ‘LDA topic 104’, ‘word: deployment (transl.)’, ‘LDA topic 5’, ‘LDA topic 111’, ‘word: interfaces (transl.)’, ‘LDA topic 85’, ‘LDA topic 61’, ‘LDA topic 114’, ‘LDA topic 43’, ‘LDA topic 45’, ‘LDA topic 26’, ‘LDA topic 132’, ‘LDA topic 16’, ‘word: production (transl.)’, ‘LDA topic 125’, ‘LDA topic 146’, ‘word: year (transl.)’, ‘LDA topic 140’, ‘LDA topic 91’, ‘word: integrate (transl.)’, ‘LDA topic 79’, ‘word: special (transl.)’

Table A.2: Most relevant features for process innovators.

Model	Top 100 most relevant features
Process innovators	‘text_length’, ‘LDA topic 98’, ‘english_language’, ‘social_media’, ‘LDA topic 41’, ‘flesch_score’, ‘incoming_links’, ‘LDA topic 7’, ‘LDA topic 75’, ‘word: worldwide (transl.)’, ‘outgoing_links’, ‘nr_subpages’, ‘LDA topic 84’, ‘word: product (transl.)’, ‘word: system’, ‘LDA topic 65’, ‘LDA topic 20’, ‘LDA topic 57’, ‘LDA topic 53’, ‘share_numbers’, ‘LDA topic 106’, ‘LDA topic 148’, ‘LDA topic 104’, ‘load_time’, ‘LDA topic 99’, ‘LDA topic 122’, ‘LDA topic 140’, ‘word: technology (transl.)’, ‘pop_score’, ‘word: to develop (transl.)’, ‘LDA topic 35’, ‘LDA topic 31’, ‘LDA topic 127’, ‘LDA topic 12’, ‘word: ISO’, ‘LDA topic 39’, ‘LDA topic 121’, ‘LDA topic 32’, ‘LDA topic 36’, ‘word: innovative (transl.)’, ‘LDA topic 2’, ‘LDA topic 100’, ‘LDA topic 6’, ‘LDA topic 13’, ‘LDA topic 120’, ‘word: standard’, ‘word: successful (transl.)’, ‘LDA topic 43’, ‘LDA topic 103’, ‘LDA topic 60’, ‘LDA topic 64’, ‘LDA topic 96’, ‘LDA topic 23’, ‘LDA topic 133’, ‘LDA topic 93’, ‘LDA topic 78’, ‘LDA topic 40’, ‘LDA topic 146’, ‘LDA topic 74’, ‘LDA topic 101’, ‘LDA topic 97’, ‘word: to start (transl.)’, ‘word: international’, ‘LDA topic 147’, ‘LDA topic 86’, ‘LDA topic 73’, ‘LDA topic 144’, ‘LDA topic 14’, ‘LDA topic 46’, ‘word: partner’, ‘LDA topic 19’, ‘LDA topic 68’, ‘word: team’, ‘LDA topic 30’, ‘LDA topic 141’, ‘LDA topic 123’, ‘LDA topic 111’, ‘LDA topic 34’, ‘LDA topic 134’, ‘word: application (transl.)’, ‘LDA topic 22’, ‘word: as well as (transl.)’, ‘LDA topic 0’, ‘LDA topic 24’, ‘LDA topic 113’, ‘LDA topic 88’, ‘LDA topic 105’, ‘LDA topic 8’, ‘LDA topic 94’, ‘LDA topic 44’, ‘LDA topic 79’, ‘LDA topic 114’, ‘LDA topic 5’, ‘LDA topic 126’, ‘LDA topic 83’, ‘LDA topic 45’, ‘LDA topic 129’, ‘LDA topic 56’, ‘LDA topic 117’, ‘LDA topic 145’

¹⁴⁰transl.: translated from German to English language.

Appendix A. Innovation Indicators Based on Firm Websites – Which Website Characteristics Predict Firm-Level Innovation Activity?

Table A.3: Most relevant features for innovators.

Model	Top 100 most relevant features
Innovators	'text_length', 'english_language', 'LDA topic 98', 'nr_subpages', 'word: system', 'word: to develop (transl.)', 'LDA topic 65', 'LDA topic 35', 'word: worldwide (transl.)', 'word: innovative (transl.)', 'LDA topic 84', 'LDA topic 134', 'LDA topic 41', 'LDA topic 20', 'word: product(transl.)', 'LDA topic 7', 'LDA topic 31', 'social_media', 'flesch_score', 'domain_purchase_year_proxy', 'word: development (transl.)', 'word: application (transl.)', 'incoming_links', 'LDA topic 78', 'outgoing_links', 'LDA topic 96', 'LDA topic 75', 'word: successful (transl.)', 'LDA topic 103', 'word: complex (transl.)', 'LDA topic 101', 'LDA topic 100', 'LDA topic 140', 'share_numbers', 'LDA topic 5', 'LDA topic 105', 'LDA topic 122', 'LDA topic 0', 'LDA topic 56', 'LDA topic 114', 'load_time', 'LDA topic 127', 'LDA topic 50', 'LDA topic 6', 'LDA topic 53', 'LDA topic 69', 'LDA topic 94', 'LDA topic 51', 'LDA topic 46', 'LDA topic 120', 'pop_score', 'LDA topic 102', 'LDA topic 90', 'LDA topic 113', 'word: to offer (transl.)', 'LDA topic 121', 'LDA topic 36', 'LDA topic 52', 'LDA topic 32', 'LDA topic 19', 'LDA topic 89', 'word: experience (transl.)', 'LDA topic 2', 'LDA topic 60', 'LDA topic 142', 'word: innovative', 'LDA topic 43', 'LDA topic 23', 'LDA topic 87', 'LDA topic 28', 'LDA topic 39', 'LDA topic 148', 'LDA topic 133', 'LDA topic 106', 'LDA topic 11', 'LDA topic 34', 'LDA topic 82', 'LDA topic 37', 'LDA topic 13', 'LDA topic 86', 'word: as well as (transl.)', 'LDA topic 61', 'LDA topic 33', 'LDA topic 12', 'LDA topic 126', 'word: high (transl.)', 'LDA topic 22', 'LDA topic 71', 'LDA topic 85', 'LDA topic 138', 'LDA topic 144', 'LDA topic 117', 'LDA topic 83', 'LDA topic 16', 'word: deployment (transl.)', 'LDA topic 136', 'LDA topic 147', 'LDA topic 123', 'LDA topic 64', 'LDA topic 68'

Table A.4: Most relevant features for innovation expenditure.

Model	Top 100 most relevant features
Innovation expend.	'english_language', 'LDA topic 98', 'text_length', 'nr_subpages', 'word: system', 'word: development (transl.)', 'word: to develop (transl.)', 'word: technology', 'LDA topic 134', 'word: innovative (transl.)', 'word: innovation', 'incoming_links', 'word: international', 'LDA topic 148', 'LDA topic 105', 'word: product (transl.)', 'word: application (transl.)', 'word: research (transl.)', 'word: worldwide (transl.)', 'LDA topic 84', 'LDA topic 7', 'domain_purchase_year_proxy', 'LDA topic 36', 'LDA topic 106', 'outgoing_links', 'LDA topic 35', 'flesch_score', 'LDA topic 28', 'LDA topic 5', 'LDA topic 20', 'LDA topic 65', 'load_time', 'LDA topic 100', 'word: innovative', 'LDA topic 39', 'LDA topic 125', 'share_numbers', 'LDA topic 41', 'LDA topic 120', 'LDA topic 73', 'LDA topic 1', 'integration', 'pop_score', 'LDA topic 82', 'LDA topic 13', 'social_media', 'emerging_tech', 'LDA topic 104', 'LDA topic 57', 'LDA topic 6', 'LDA topic 53', 'LDA topic 109', 'LDA topic 26', 'LDA topic 75', 'word: high', 'LDA topic 34', 'LDA topic 32', 'LDA topic 89', 'LDA topic 49', 'LDA topic 140', 'LDA topic 81', 'word: workshop', 'LDA topic 83', 'LDA topic 113', 'word: management', 'LDA topic 22', 'LDA topic 59', 'LDA topic 56', 'LDA topic 31', 'LDA topic 67', 'LDA topic 24', 'LDA topic 0', 'LDA topic 79', 'LDA topic 68', 'LDA topic 102', 'LDA topic 61', 'LDA topic 3', 'LDA topic 138', 'LDA topic 44', 'LDA topic 40', 'LDA topic 128', 'LDA topic 146', 'LDA topic 141', 'word: to optimize', 'LDA topic 70', 'LDA topic 78', 'LDA topic 132', 'LDA topic 95', 'word: process (transl.)', 'LDA topic 80', 'LDA topic 127', 'LDA topic 60', 'LDA topic 93', 'LDA topic 133', 'LDA topic 114', 'LDA topic 46', 'word: high', 'word: as well as', 'LDA topic 96', 'LDA topic 8'

A.5 Learned Hyperparameters for Random Forest Models Using Different Feature Sets and Target Variables

Table A.5: Learned hyperparameters for Random Forest models using different feature sets and target variables.

Feature sets			Number of trees	Max. depth	Min. impurity decrease
Text	Meta	Network			
Product innovators					
x			1000	50	0.001
	x		1000	50	0.001
		x	1500	50	0.001
x	x	x	1500	100	0.001
Process innovators					
x			1000	50	0.001
	x		1500	50	0.01
		x	1000	50	0.001
x	x	x	1500	50	0.001
Innovators					
x			1500	50	0.001
	x		1000	50	0.001
		x	500	50	0.001
x	x	x	1000	50	0.001
Innovation expenditures					
x			1500	100	0.001
	x		1000	50	0.01
		x	1000	50	0.01
x	x	x	1000	50	0.001

A.6 AUC Values for Different Splits Between Training and Test Sample




Figure A.3: AUC values for different splits between training and test sample. Line plot that illustrates for each indicator how AUC values of the 'all' feature model increase if the train/test split changes from (0.1/0.9) to (0.9/0.1) in steps of 0.01.

Appendix B

Firm Digitalisation and Mobility – Do Covid-19-Related Changes Persist?

B.1 Variable Description

Table B.1: Description of variables.

variable	description	source
Main variables		
<i>mobility</i>	The change in switches between phone cells per mobile device at a given day relative to monthly pre-crisis averages in 2019 for the same weekday (in %, 01/01/'20 - 31/12/'22) at the district level. The change in mobility is reported for daytime and nighttime separately (daytime mobility: 6 a.m. to 10 p.m.; nighttime mobility: 10 p.m. to 6 p.m.). Overall mobility is a weighted mean of daytime and nighttime mobility. Mobility data is being processed and provided by the Teralytics AG. Please note that we observe missing values over the entire time period.	[data set] Destatis (2023a). Mobile Network Data. (https://www.destatis.de/EN/Service/EXDAT/Datensaetze/mobility-indicators-mobilephone.htm [retrieved on 03/01/2023]). Data was obtained upon request.
<i>digitalisation</i>	We train a Random Forest regression model on a large German newspaper corpus. The fitted model allows for predicting the likelihood that a firm's website content relates to digitalisation. These predictions are used as a continuous indicator for firm digitalisation. For this purpose, we scraped 750,000 firm websites in January 2020 and 1,300,000 firm websites in December 2022. For more details see Axenbeck & Breithaupt (2022). Validity checks with external data show a clear relationship with already established digitalisation indicators at the firm, sectoral, and regional level. We average predictions at the district level.	[data set] ZEW – Leibniz Centre for European Economic Research (2022). Web addresses are retrieved from the Mannheim Enterprise Panel (MUP), which comprises a large set of German firms (Bersch et al. 2014). The MUP is fed by data from Creditreform.

Appendix B. Firm Digitalisation and Mobility – Do Covid-19-Related Changes Persist?

Table B.1: Description of variables.

variable		description	source
Control variables			
Pandemic characteristics			
<i>weekly cases</i>		Sum of confirmed Covid-19 cases in the last 7 days per 100,000 inhabitants (01/03/'20 - 31/12/'22) at the district level considered as a rolling window. We include weekly cases and not daily cases as a control variable, because weekly cases better reflect the number of people that are infectious and it is also the number, which was mostly communicated in the media. We, therefore, assume that people made their mobility decisions on weekly cases. We set values before 01/03/'20 to zero.	[data set] Destatis (2023b) (https://www.corona-daten-deutschland.de/dataset/infektionen_kreise [retrieved on 04/01/2021]).
<i>containment measures</i>	<i>measures</i>	Index that captures the severity of containment measures at the district level (01/03/'20 - 30/11/'22). The index was calculated by infas 360. We set values that are before the observed time frame to zero. We use the last observed value to replace missing values after the observed time frame.	[data set] Destatis (2023b) (https://www.corona-daten-deutschland.de/dataset/massnahmenindex_kreise_pro_tag [retrieved on 04/01/2023]).
socioeconomic characteristics			
<i>share of academics</i>		Number of persons aged 15 or older with a Bachelors, Masters, Ph.D., or comparable university degree divided by the number of inhabitants being 15 or older in 2019 at the district level.	[data set] Destatis (2023b) (number of academics, https://www.corona-daten-deutschland.de/dataset/bildungsniveau [retrieved on 02/06/2022], number of inhabitants: https://www.corona-daten-deutschland.de/dataset/bevoelkerung [retrieved on 02/06/2022]).
<i>GDP per inhabitant</i>		Gross domestic product in €1,000 per person in 2020 at the district level.	[data set] Destatis (2023b) (https://www.corona-daten-deutschland.de/dataset/volkswirtschaftliche_gesamtrechnung [retrieved on 2/06/2022]).
<i>low-income households</i>	<i>households</i>	The number of low-income households (\leq €1,000 per month) in 2019 at the district level per 1,000 inhabitants.	[data set] Destatis (2023b) (number of low-income households: https://www.corona-daten-deutschland.de/dataset/private_finanzen [retrieved on 12/11/2022], number of inhabitants: https://www.corona-daten-deutschland.de/dataset/bevoelkerung [retrieved on 02/06/2022]).
<i>people on social benefits</i>		The number of recipients of benefits under SGB II and the number of recipients of benefits under SGB XII per 1,000 inhabitants in 2017.	[data set] Destatis (2023b) (https://www.corona-daten-deutschland.de/dataset/sozialindikatoren [retrieved on 12/11/2022]).
<i>share of workers in the service sector</i>		The number of people that work in the service sector divided by all workers observed at the end of 2019.	[data set] Destatis (2023b) (https://www.corona-daten-deutschland.de/dataset/arbeitsmarktstruktur [retrieved on 12/11/2022]).
infrastructure			

Appendix B. Firm Digitalisation and Mobility – Do Covid-19-Related Changes Persist?

Table B.1: Description of variables.

variable	description	source
<i>cars per person</i>	Number of cars per 1,000 inhabitants in 2021.	[data set] Destatis (2023b) (https://www.corona-daten-deutschland.de/dataset/verkehr [retrieved on 12/11/2022]).
≥ 50 Mbit/s	Share of households with a broadband availability of at least 50 Mbit/s in each district in 2020.	[data set] atene KOM GmbH (2021). Breitbandatlas des Bundes (German Broadband Atlas) - Release 2/2021. Data is restricted in usage. Access can be requested at atene KOM GmbH (https://atenekom.eu/project/breitbandatlas/ [retrieved on 09/04/2021]).
<i>not covered by all</i>	Area that is not covered by 4G, 5G or 5G DSS by every network provider in % for the year 2022. Please note that information is only available for 388 districts. Missing vs are set to zero and an additional dummy variable is added that controls whether the information is available or not.	[data set] Bundesnetzagentur (2022). Mobilfunkmonitoring (https://www.breitband-monitor.de/mobilfunkmonitoring/download [retrieved on 22/12/2022]).
<i>not covered</i>	Area that is not covered by 4G, 5G, or 5G DSS by any network provider in % for the year 2022. Please note that information is only available for 388 districts. Missing values are set to zero and an additional dummy variable is added that controls whether the information is available or not.	[data set] Bundesnetzagentur (2022). Mobilfunkmonitoring (https://www.breitband-monitor.de/mobilfunkmonitoring/download [retrieved on 22/12/2022]).
demographic characteristics		
<i>share of men</i>	The number of male inhabitants divided by all inhabitants in 2020.	[data set] Destatis (2023b) (number of male inhabitants & number of inhabitants, https://www.corona-daten-deutschland.de/dataset/bevoelkerung [retrieved on 02/06/2022]).
<i>not of working age</i>	The percentage of the population in a district that is either under 15 years old or over 65 years old.	[data set] Destatis (2023b) (number of people younger than 15 years or older than 65 years & number of inhabitants, https://www.corona-daten-deutschland.de/dataset/bevoelkerung [retrieved on 02/06/2022]).
<i>number of inhabitants</i>	The number of people that are registered in a district (divided by 1,000).	[data set] Destatis (2023b) (https://www.corona-daten-deutschland.de/dataset/bevoelkerung [retrieved on 02/06/2022]).
<i>changes in population</i>	Changes in % of the number of people that are registered in a district. We consider changes in the population between 2019 and 2020 for the year 2020 and changes between 2019 and 2021 for the years 2021 and 2022 (as information for 2022 was not available when the analysis was conducted).	GENESIS-ONLINE: Table 12411-0015 (https://www-genesis.destatis.de/genesis/online [retrieved on 02/06/2022]).
<i>population density</i>	Inhabitants per square kilometre in 2019 at the district level.	[data set] Destatis (2023b) (https://www.corona-daten-deutschland.de/dataset/besiedlung [retrieved on 2/06/2022]).

Appendix B. Firm Digitalisation and Mobility – Do Covid-19-Related Changes Persist?

Table B.1: Description of variables.

variable		description	source
<i>in-commuters</i>		Changes in the number of employees that work in a district but live elsewhere in %. We consider changes in in-commuters between 2019 and 2020 for the year 2020 and changes between 2019 and 2021 for the years 2021 and 2022 (as information for 2022 was not available when the analysis was conducted).	[data set] Bundesagentur für Arbeit (2022). Pendlerverflechtungen der sozialversicherungspflichtig Beschäftigten nach Kreisen - Deutschland (https://statistik.arbeitsagentur.de [retrieved on 01/12/2022]).
<i>out-commuter</i>		Changes in the number of employees that live in a district but work elsewhere in %. We consider changes in out-commuters between 2019 and 2020 for the year 2020 and changes between 2019 and 2021 for the years 2021 and 2022 (as information for 2022 was not available when the analysis was conducted).	[data set] Bundesagentur für Arbeit (2022). Pendlerverflechtungen der sozialversicherungspflichtig Beschäftigten nach Kreisen - Deutschland (https://statistik.arbeitsagentur.de [retrieved on 01/12/2022]).
<i>one-person households</i>	<i>house-</i>	Number of people per 1,000 inhabitants that live in a one-person household in a district for the year 2019.	[data set] Destatis (2023b) (number of one-person households: https://www.corona-daten-deutschland.de/dataset/haushalte [retrieved on 01/11/2022], number of inhabitants: https://www.corona-daten-deutschland.de/dataset/bevoelkerung [retrieved on 02/06/2022]).
<i>living space per household</i>	<i>per</i>	Average living space per household in a district for the year 2019.	[data set] Destatis (2023b) (https://www.corona-daten-deutschland.de/dataset/wohnsituation [retrieved on 01/11/2022]).

Appendix B. Firm Digitalisation and Mobility – Do Covid-19-Related Changes Persist?

Table B.1: Description of variables.

variable	description	source
geographic characteristics		
<i>city</i>	A dummy variable that is one if a district is a “Stadtkreis” or a “Kreisfreie Stadt” (city) and zero if a district is a “Landkreis” or “Kreis” (countryside area).	[data set] atene KOM GmbH (2021). Breitbandatlas des Bundes (German Broadband Atlas) - Release 2/2021. Data is restricted in usage. Access can be requested at atene KOM GmbH (https://atenekom.eu/project/breitbandatlas/) [retrieved on 09/04/2021]).
<i>West Germany</i>	A dummy that is one if a district is in the former Federal Republic of Germany (West Germany) and that is zero if a district is in the former German Democratic Republic (East Germany).	[data set] Destatis (2023b) (https://www.corona-daten-deutschland.de/dataset/raumordnung) [retrieved on 01/11/2022]).
other variables		
<i>WFH potential</i>	The percentage of employees who can potentially work from home according to their self-assessment and considered at the location of their workplace. The calculation is described in detail in Alipour et al. (2023).	[data set] Destatis (2023b) (https://www.corona-daten-deutschland.de/dataset/arbeitsmarktstruktur) [retrieved on 01/11/2022]).
<i>number of firms</i>	Number of firms in a district in March 2020.	[data set] Destatis (2023b) (https://www.corona-daten-deutschland.de/dataset/firmeninformati) [retrieved on 05/1/2023]).

B.2 Additional Descriptive Statistics

Table B.2: Overview of descriptive statistics.

	N	mean	sd	p10	p90
Δ mobility	433999 (daily)	1.02	17.87	-21.33	22.67
Δ mobility daytime	433999 (daily)	5.04	18.24	-17.00	27.00
Δ mobility nighttime	433999 (daily)	-7.02	21.69	-34.00	18.00
digitalisation (Jan '20)	400	0.00	1.00	-1.16	1.33
digitalisation (Dec '22)	400	0.00	1.00	-1.22	1.33
digitalisation modified (Jan '20)	400	0.00	1.00	-1.35	1.39
digitalisation modified (Dec '22)	400	0.00	1.00	-1.19	1.34
Δ digitalisation modified (Dec '22)	400	0.00	1.00	-1.12	1.35
weekly cases	433999 (daily)	289.68	464.14	0.80	852.60
containment measures	433999 (daily)	31.58	20.03	8.01	58.76
share of academics	400	0.16	0.05	0.11	0.22
GDP per inhabitant in €1000	400	37.08	16.05	24.80	53.80
low-income households per 1,000 inhabitants	400	69.71	43.48	23.26	133.49
people on social benefits per 1,000 inhabitants	400	8.99	4.14	4.20	15.20
share of workers in the service sector	400	0.49	0.07	0.41	0.58
cars per 1,000 person	400	549.82	77.11	434.00	630.00
≥ 50 mbit/s	400	91.94	7.23	81.50	98.80
not covered by all	400	15.22	10.21	1.39	29.21
not covered	400	2.45	3.16	0.00	6.55
share of men	400	0.49	0.01	0.49	0.50
share not of working age	400	0.36	0.03	0.33	0.40
number of inhabitants divided by 1,000	400	207.77	244.81	72.04	346.97
change in population	800 ('19-'20, '19-'21)	0.00	0.00	-0.01	0.01
population density per square kilometre	400	536.46	709.62	83.00	1484.00
change in out-commuters	800 ('19-'20, '19-'21)	0.02	0.02	-0.01	0.05
change in in-commuters	800 ('19-'20, '19-'21)	0.02	0.03	-0.02	0.06
one-person households per 1,000 inhabitants	400	198.38	59.32	135.33	292.64
living space per household	400	115.32	18.40	89.00	138.00
city	400	0.26	0.44	0.00	1.00
West Germany	400	0.81	0.39	0.00	1.00




Figure B.1: **Change in mobility compared to 2019 by distance in % (7-day average).** The red line denotes the start of the pandemic on March 22nd, 2020. Grey areas mark lockdown periods and blue areas mark periods where a government-imposed WFH obligation was additionally in place. Data Source: Destatis (2023).

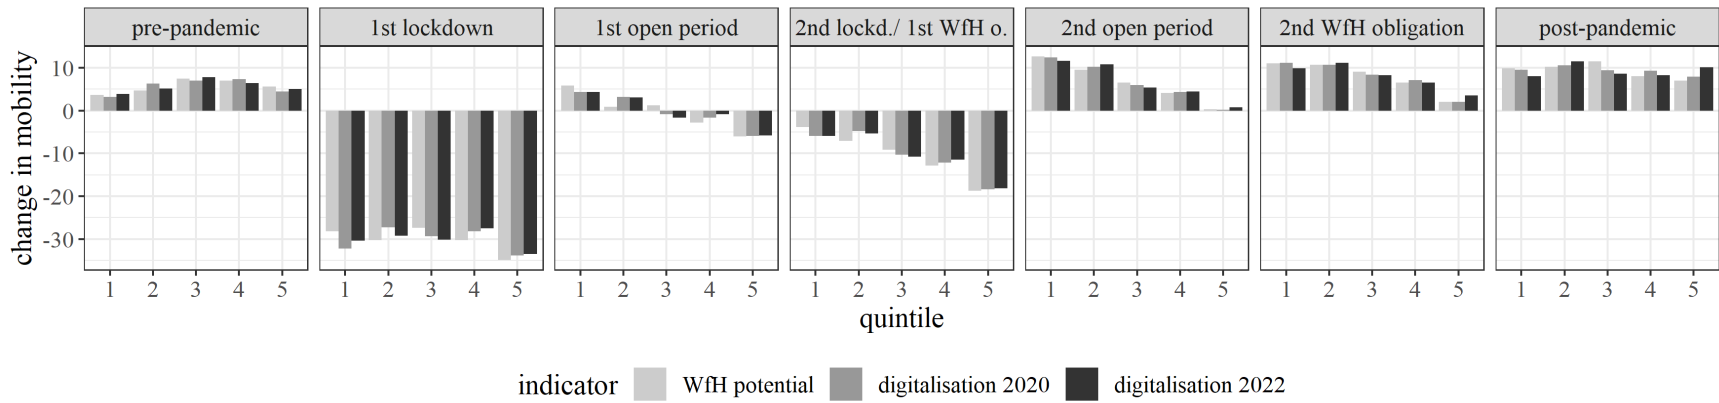


Figure B.2: **Comparison between quintiles.** Comparison between quintiles of the WFH potential derived by Alipour et al. (2023) and the web-based digitalisation indicator with respect to differences in average mobility.

B.3 Results With Respect to the Average Effect

To measure the average effect during the two years of the pandemic (from March 22nd, 2020 to March 19th, 2022), we estimate the following linear model using year-month fixed effects and clustered standard errors at the district level:

$$\Delta \text{mobility}_{i,t}^h = \alpha + \beta \text{digitalisation}_i + \sum_{c \in C} \gamma_c c_{i,t} + u_{i,t}. \quad (\text{B.1})$$

h refers to the period (daytime, nighttime, or the entire day) for which mobility is observed in district i at day t . C denotes our set of control variables c . Observations are weighted based on their population size.

Table B.3: Average decrease in mobility associated with digitalisation considering mobility changes over the entire day, daytime mobility changes, nighttime mobility changes, as well as differences between working days and weekends during the two pandemic years. Firm digitalisation is observed in 2020.

	dependent variable: Δ mobility				
	(1)	(2)	(3)	(4)	(5)
digitalisation (Jan '20)	-1.681*** (-3.44)	-2.374*** (-4.60)	-0.297 (-0.59)	-2.565*** (-5.01)	-1.900*** (-3.50)
year-month fixed effects	x	x	x	x	x
pandemic controls	x	x	x	x	x
socioeconomic controls	x	x	x	x	x
infrastructure controls	x	x	x	x	x
demographic controls	x	x	x	x	x
geographic controls	x	x	x	x	x
observations	288399	288399	288399	205599	82800
R^2	0.566	0.532	0.474	0.565	0.501

Notes: Equation B.1 estimated using OLS and all control variables. t statistics in parentheses. Clustered standard errors. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Observations are weighted based on their population size. (1) mobility changes over the entire day; (2) mobility changes during daytime; (3) mobility changes during nighttime; (4) mobility changes on working days during daytime; (5) mobility changes on weekends during daytime.

B.4 Results of Robustness Checks

B.4.1 Parallel Trends




Figure B.3: **Analysis of parallel trends before the Covid-19 pandemic (estimated β^m coefficients).** Equation (3.1) is estimated at the weekly level. Digitalisation and control variables are interacted with time dummies. We estimate the time frame between January 7th, 2020, and May 4th, 2020. The latter date denotes the end of the first lockdown. We exclude the first week in January because it is a holiday period with irregular mobility patterns. We use the week before Shrove Monday 2020 as the reference period, as the first large-scale Covid-19 outbreaks occurred in Germany as part of carnival festivities (grey line). Digitalisation is observed in 2020. Confidence intervals are at the 90% significance level.

B.4.2 Digitalisation Observed in December 2022




Figure B.4: **Monthly change in mobility associated with digitalisation observed in December 2022 (estimated β^m coefficients)**. Our reference period is February 2020 (grey line). The red line denotes the start of the first lockdown on March 22nd, 2020. White areas mark periods with no or few Covid-19 restrictions, grey areas mark lockdown periods and blue areas mark periods where the government-imposed WFH obligation was additionally in place. Confidence intervals are at the 90% significance level.

Appendix B. Firm Digitalisation and Mobility – Do Covid-19-Related Changes Persist?

Table B.4: DiD results providing insights into changes in the link between mobility reductions and firm digitalisation with respect to different phases of the pandemic using digitalisation observed in 2022.

	dependent variable: Δ mobility						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
digitalisation (Dec '22)							
× (1) 1st lockdown	-1.816** (-3.10)	-1.900*** (-3.53)	1.439* (2.31)	0.254 (0.39)	2.069** (3.20)	0.411 (0.62)	2.168*** (3.83)
× (2) 1st open period	-3.769*** (-8.15)	-3.777*** (-8.54)	-2.352*** (-3.38)	-2.303*** (-3.92)	-1.531* (-2.33)	-2.431*** (-4.20)	-1.250+ (-1.90)
× (3) 2nd lockdown/ 1st WFH o.	-4.654*** (-13.07)	-4.959*** (-13.88)	-2.529*** (-4.72)	-2.287*** (-4.51)	-1.709** (-3.22)	-2.761*** (-6.16)	-1.306* (-2.41)
× (4) 2nd open period	-3.561*** (-7.02)	-3.813*** (-7.57)	-2.811*** (-3.47)	-2.358** (-3.28)	-2.758*** (-3.55)	-2.897*** (-4.62)	-2.004* (-2.50)
× (5) 2nd WFH obligation	-2.436*** (-5.25)	-2.748*** (-6.20)	-1.492* (-2.27)	-0.496 (-0.76)	-2.236** (-3.25)	-2.466*** (-4.29)	-1.735** (-2.60)
× (6) post-pandemic	0.893 (1.50)	0.591 (1.00)	1.384 (1.52)	1.929* (2.18)	-0.0779 (-0.09)	0.618 (0.80)	0.359 (0.39)
year-month fixed effects	x	x	x	x	x	x	x
district-level fixed effects	x	x	x	x	x	x	x
pandemic controls		x					x
socioeconomic controls			x				x
infrastructure controls				x			x
demographic controls					x		x
geographic controls						x	x
observations	433999	433999	433999	433999	433999	433999	433999
R ²	0.57	0.59	0.58	0.58	0.59	0.58	0.62
$\beta^1 = \beta^6$	0.00	0.00	0.95	0.07	0.02	0.81	0.04
$\beta^2 = \beta^6$	0.00	0.00	0.00	0.00	0.03	0.00	0.02
$\beta^3 = \beta^6$	0.00	0.00	0.00	0.00	0.01	0.00	0.01
$\beta^4 = \beta^6$	0.00	0.00	0.00	0.00	0.00	0.00	0.00
$\beta^5 = \beta^6$	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Notes: Equation 3.1 estimated using OLS. *t* statistics in parentheses. Clustered standard errors. + $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Observations are weighted based on their population size. The time frame is split into different phases as presented in Figure 3.3. The pre-Covid-19 phase is used as a reference period. Fixed effects for every phase are additionally included. The table also includes *t*-tests for the equality of coefficients. (1) no control variables; (2) only controlled for the pandemic situation; (3) only controlled for socioeconomic characteristics; (4) only controlled for characteristics that relate to a district's infrastructure; (5) only controlled for demographic characteristics; (6) only controlled for geographic characteristics; (7) all control variables included.

B.4.3 Further Robustness Checks

Table B.5: Further robustness checks.

	Δ mobility '20 modified	Δ mobility '22 modified	Δ mobility '22 – '20 modified	Δ mobility ≥ 50 Mbit/s	Δ mobility '20 no weights	Δ mobility '22 no weights
	(1)	(2)	(3)	(4)	(5)	(6)
digitalisation						
× (1) 1st lockdown	1.900*** (3.55)	0.838* (2.20)	0.0639 (0.23)	0.193* (2.46)	2.120** (3.24)	2.032*** (3.56)
× (2) 1st open period	-0.210 (-0.33)	-0.831 ⁺ (-1.68)	-0.633 ⁺ (-1.72)	-0.0456 (-0.47)	-1.539* (-2.00)	-1.503* (-2.00)
× (3) 2nd lockdown/ 1st WFH o.	-0.711 (-1.26)	-0.922* (-2.05)	-0.544 (-1.60)	0.0149 (0.19)	-1.471* (-2.31)	-1.249 ⁺ (-1.93)
× (4) 2nd open period	-0.359 (-0.44)	-1.051 ⁺ (-1.65)	-0.773 (-1.61)	-0.133 (-1.16)	-2.259* (-2.48)	-2.132* (-2.28)
× (5) 2nd WFH obligation	-0.821 (-1.14)	-0.711 (-1.28)	-0.324 (-0.74)	-0.171 ⁺ (-1.89)	-2.203** (-3.03)	-1.685* (-2.25)
× (6) post-pandemic	0.700 (0.65)	0.266 (0.37)	-0.00608 (-0.01)	0.0445 (0.41)	-0.425 (-0.47)	0.177 (0.19)
year-month fixed effects	x	x	x	x	x	x
district-level fixed effects	x	x	x	x	x	x
pandemic controls	x	x	x	x	x	x
socioeconomic controls	x	x	x	x	x	x
infrastructure controls	x	x	x	x	x	x
demographic controls	x	x	x	x	x	x
geographic controls	x	x	x	x	x	x
observations	433999	433999	433999	433999	433999	433999
R^2	0.619	0.619	0.619	0.619	0.558	0.558

Notes: Equation 3.1 estimated using OLS. t statistics in parentheses. Clustered standard errors. ⁺ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Observations weighted by population size if not stated otherwise. The time frame is split into different phases as presented in Figure 3.3. The pre-Covid-19 phase is used as a reference period. Models are estimated with all control variables. Fixed effects for every phase are additionally included. (1) All firms are considered when calculating the average degree of firm digitalisation in a district, i.e., firms with no available website are set to zero; the indicator is calculated for digitalisation in 2020; (2) same modification but the indicator is calculated for digitalisation in 2022; (3) change in a district's degree of digitalisation is calculated by subtracting the modified indicators for 2020 from the modified indicator for 2022; (4) firm digitalisation is excluded and coefficients for broadband availability are used as a proxy for effects of household digitalisation over time and displayed; (5) - (6) firm digitalisation is calculated for 2020 and for 2022 as in the main analysis, but observations are not weighted by population size.

B.4.4 Firm-Level Link

To analyse the link between firm digitalisation and the share of employees that work from home, we estimate the following model using MIP 2021 data:¹⁴¹

$$\begin{aligned} \text{WFH share}_{j,t} = & \alpha + \beta_d \text{digitalisation}_j^{2020} + \beta_{ld} \text{lockdown}_t \\ & + \beta_{dld} \text{digitalisation}_j^{2020} * \text{lockdown}_t^p + u_{j,t}. \end{aligned} \quad (\text{B.2})$$

Digitalisation is considered for firm j at time t , which can either be the time before the first lockdown (January/February 2020) or a lockdown period. p denotes the considered lockdown, which can either be the first or second one. Results are displayed in Table B.6.

Table B.6: **Link between firm digitalisation and WFH at the firm level.**

	dependent variable: WFH share			
	(1)	(2)	(3)	(4)
digitalisation (Jan '20)	0.729 (1.45)	0.634 (1.30)	0.624 (1.33)	0.522 (1.16)
1st lockdown	13.47* (2.33)	13.47* (2.22)		
2nd lockdown			15.72* (2.57)	15.71* (2.51)
digitalisation (Jan '20) × 1st lockdown	6.364* (2.30)	6.361* (2.19)		
digitalisation (Jan '20) × 2nd lockdown			7.022* (2.53)	7.017* (2.47)
constant	-0.999 (-0.15)	1.196 (0.20)	-2.121 (-0.31)	-0.268 (-0.04)
ln(sigma)	2.990*** (7.00)	2.977*** (6.63)	3.012*** (7.75)	3.000*** (7.54)
industry	x	x	x	x
federal state		x		x
log-likelihood	-18466.7	-18392.3	-18297.7	-18225.1
observations	6028	6028	6028	6028

Notes: Equation B.2 estimated using an interval-censored regression model. Coefficients can be directly interpreted (see Wooldridge 2002). Sigma is comparable to the standard error of an OLS estimate. t statistics in parentheses. Clustered-standard errors at the firm level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Digitalisation is standardised with respect to the firm sample. (1) first lockdown only controlling for different industries; (2) first lockdown controlling for different industries and federal-state fixed effects; (3) second lockdown only controlling for different industries; (4) second lockdown controlling for different industries and federal-state fixed effects.

¹⁴¹See [data set] ZEW – Leibniz Centre for European Economic Research (2021).

Appendix B. Firm Digitalisation and Mobility – Do Covid-19-Related Changes Persist?

Furthermore, we analyse whether our web-based firm digitalisation indicator can be associated with a greater likelihood that a firm increased its digital products, services, and sales channels with the onset of the Covid-19 crisis based on the MIP 2021. To this end, we fit the following linear model:

$$\text{e-commerce}_j = \alpha + \beta_d \text{digitalisation}_j^{2020} + u_j. \quad (\text{B.3})$$

The increase in digital business activities is denoted by the binary variable *e-commerce*. Results are displayed in Table B.7:

Table B.7: Link between firm digitalisation and increased e-commerce activity at the firm level.

	dependent variable: increased e-commerce	
	(1)	(2)
digitalisation (Jan '20)	0.0723*** (7.54)	0.0717*** (7.42)
constant	0.665* (2.11)	0.702* (2.25)
industry	x	x
federal state		x
R-squared	0.0661	0.0760
observations	3014	3014

Notes: Equation B.3 estimated using OLS. *t* statistics in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Robust standard errors. Digitalisation is standardised with respect to the firm sample. (1) only controlled for different industries; (2) controlled for different industries and federal-state fixed effects.

B.4.5 Link to a District’s WFH Potential

Table B.8: Equation B.1 with digitalisation replaced by a district’s WFH potential. Only daytime mobility changes are considered.

	dependent variable: Δ mobility	
	(1)	(2)
WFH potential	-0.506** (-2.93)	-0.271 (-1.50)
digitalisation (Jan '20)		-1.593** (-2.78)
constant	27.41 (0.57)	14.41 (0.30)
year-month fixed effects	x	x
socioeconomic controls	x	x
infrastructure controls	x	x
demographic controls	x	x
geographic controls	x	x
observations	433999	433999
adjusted R^2	0.452	0.454

Notes: Equation B.1 estimated using OLS and all control variables. t statistics in parentheses. Clustered standard errors. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Observations are weighted based on their population size. (1) Equation B.1 estimated for daytime mobility with digitalisation replaced by a district’s WFH potential; (2) Equation B.1 estimated for daytime mobility with digitalisation and a district’s WFH potential included.

Appendix C

Digital Technology Adoption and Energy Intensity in Manufacturing

C.1 Additional Data

For our analysis, we add, inter alia, information on prices of different energy sources, gross value added deflators to calculate real value added, and growth and depreciation rates, as well as investment deflators to calculate capital stocks. All data sources are listed in Table C.1. The identifier denotes the variable that is used to merge the dataset with AFiD.

Table C.1: Description of additional data sources.

Information	Data source	Comments	Identifier
Price for energy source (electricity, natural gas, heating oil, coal)	Gesamtausgabe der Energiedaten, Federal Ministry for Economic Affairs and Energy (BMWi); Status: 31.03.2020, https://www.bmwi.de/Redaktion/DE/Artikel/Energie/energiedaten-gesamtausgabe.html [Online; accessed on 15 Apr. 2023] (Retrieved on: 01.04.2020)	Prices for hard coal (import prices), heavy heating oil (industry prices, VAT excluded), light heating oil (light, industry prices, VAT excluded), electricity, and natural gas prices independent from the consumption level are retrieved. The respective units have all been converted to €/kWh.	Year
Price for energy source (district heat)	Fernwärme – Preisübersicht, AGFW Der Energieeffizienzverband für Wärme, Kälte und KWK e. V.; Status: 01.10.2017, https://www.agfw.de/energiewirtschaft-recht-politik/wirtschaft-und-markt/markt-preise/preisanpassung/ [Online; accessed on 15 Apr. 2023] (Retrieved on: 14.08.2019)	Absolute price development from 2009-2017 for the connected loads of 160 kW (p.8) are used. Values are converted from €/MWh to €/kWh. Prices are retrieved without VAT.	Year

Appendix C. Digital Technology Adoption and Energy Intensity in Manufacturing

Table C.1: Description of additional data sources.

Information	Data source	Comments	Identifier
Price for energy source (biomass)	Brennstoffkostenentwicklung von Gas, Öl und Pellets, Deutsches Pelletinstitut GmbH (DEPI); Status: 2019, https://depi.de/de/pelletpreis-wirtschaftlichkeit#dau2v [Online; accessed on 15 Apr. 2023] (Retrieved on: 13.09.2019)	Pellet price for 2015 is taken, value converted from cent/kWh to €/kWh (VAT excluded).	Year
Price for energy source (biomass)	Index der Erzeugerpreise gewerblicher Produkte (5.10 Holzprodukte - GP09-1629 14 908 Pellets, Briketts, Scheiten o.ä. Formen aus Sägespänen u.a. Sägenebenprodukt), from: Daten zur Energiepreisentwicklung - Lange Reihen von Januar 2005 bis Mai 2020, Statistisches Bundesamt (Destatis); Status: 26.06.2020, https://www.destatis.de/DE/Themen/Wirtschaft/Preise/Publikationen/Energiepreise/energiepreisentwicklung-pdf-5619001.pdf?__blob=publicationFile [Online; accessed on 15 Apr. 2023] (Retrieved on: 16.07.2020)	The base year of the Destatis index is 2015. Therefore, the DEPI-price is taken from the year 2015 and multiplied by the index for each year to receive information about the change in the price for biomass.	Year
Price for energy source (liquid gas)	IEA Energy Prices and Taxes Statistics, International Energy Agency; Status: 1.Quarter 2019, https://www.oecd-ilibrary.org/energy/data/iea-energy-prices-and-taxes-statistics_enep-ice-data-en [Online; accessed on 15 Apr. 2023] (Retrieved on: 04.09.2019)	Prices (VAT excluded) from 2009-2017 for liquid gas are retrieved. Values are converted from €/1 to €/kWh.	Year
Producer price index (PPI)	Index der Erzeugerpreise gewerblicher Produkte (Inlandsabsatz) nach dem Güterverzeichnis für Produktionsstatistiken Ausgabe 2009 (GP 2009) - Lange Reihen der Fachserie 17, Reihe 2 von Januar 2005 bis September 2020, Statistisches Bundesamt (Destatis), Status: 20.10.2020, https://www.destatis.de/DE/Themen/Wirtschaft/Preise/Erzeugerpreisindex-gewerbliche-Produkte/Publikationen/Downloads-Erzeugerpreise/erzeugerpreise-lange-reihen-pdf-5612401.html [Online; accessed on 15 Apr. 2023] (Retrieved on: 12.11.2020)	Index on the yearly average change is retrieved.	Year, industries (two-digit NACE code)

Appendix C. Digital Technology Adoption and Energy Intensity in Manufacturing

Table C.1: Description of additional data sources.

Information	Data source	Comments	Identifier
Gross value added deflators	National accounts aggregates by industry, Eurostat, Status: 24.03.2020, https://ec.europa.eu/eurostat/de/ [Online; accessed on 15 Apr. 2023] (Retrieved on: 01.04.2020)	Price index (implicit deflator), base year 2010, national currency.	Year
Capital stock	Cross-classification of gross fixed capital formation by industry and by asset (flows) - Computer software and databases (gross), Eurostat, Status: 30.03.2020, https://ec.europa.eu/eurostat/de/ [Online; accessed on 15 Apr. 2023] (Retrieved on: 01.04.2020)	Table PD10_NAC, price index (implicit deflator), base year 2010, national currency. Software deflators are retrieved. See C.2 for detailed information on how we calculate software as well as non-software capital stocks.	Year
Capital stock	EU KLEMS database - 2019 release, Germany capital input data, see Stehrer, R., A. Bykova, K. Jäger, O. Reiter and M. Schwarzhappel (2019): Industry level growth and productivity data with special focus on intangible assets, wiiw Statistical Report No. 8. https://euklems.eu/excel/DE_Capital_SDB_2019.xlsx [Online; accessed on 15 Apr. 2023] (Retrieved on: 18.04.2020)	Real gross fixed capital formation (in prices from 2010) to calculate growth rates, depreciation rates as well as investment deflators (except software deflators) are taken from the EU KLEMS database for the years 2003-2017. See C.2 for detailed information on how we calculate software as well as non-software capital stocks	Year, industries (two-digit NACE code)
Household broadband availability	Breitbandatlas des Bundes (German Broadband Atlas) - Release 2/2021. Data is restricted in usage. Access can be requested at ateneKOM GmbH, https://atenekom.eu/project/breitbandatlas/ [Online; accessed on 15 Apr. 2023] (Retrieved on: 9.04.2021)	Not integrated in the analysis	municipality level (AGS)

C.2 Perpetual Inventory Method (PIM)

In the spirit of Griliches (1980), Berlemann & Wesselhöft (2014), Lutz (2016), Löschel et al. (2019), and Dhyne et al. (2021a) capital stocks are calculated for software capital and non-software capital by means of the perpetual inventory method (PIM).

Given geometric constant depreciation, the capital stock K_t at period t can be written as a function of previous period's capital stock K_{t-1} , gross investments I_t , and the consumption of fixed capital at rate δ . Hence, capital stocks except initial ones can be calculated by the following equation:

$$K_t = (1 - \delta)K_{t-1} + I_t. \quad (\text{C.1})$$

To calculate initial capital stocks, one can express annual percentage increase in capital as the amount of investments minus the capital depreciated in the previous period:

$$\frac{K_t - K_{t-1}}{K_{t-1}} = \frac{I_t}{K_{t-1}} - \delta. \quad (\text{C.2})$$

Assuming that capital grows at a constant rate ($g_K = (K_t - K_{t-1})/K_{t-1}$), one can obtain the following expression:

$$K_{t-1} = \frac{I_t}{g_K + \delta}. \quad (\text{C.3})$$

Setting $t = 1$ allows to calculate the initial capital stock:

$$K_0 = \frac{I_1}{g_K + \delta}. \quad (\text{C.4})$$

For the calculation of firm-level initial capital stocks, it is recommended to use average investments of the first three years within the observation period because investments highly fluctuate over time:¹⁴²

$$\hat{I}_1 = \frac{\sum_{t=1}^3 I_t}{n}. \quad (\text{C.5})$$

Accordingly, in this study we calculate initial capital stocks by applying Equation (C.4) and (C.5), subsequent capital stocks are calculated by Equation (C.1).

PIM requires information on capital growth rates. These are estimated by calculating the compound annual growth rate at the industry level using real gross fixed capital formation at prices from 2010. Information on the gross fixed capital formation volume of software and total capital is retrieved from the EU KLEMS database.

¹⁴²Please note here that we do robustness checks with respect to different period lengths to calculate initial capital stocks.

Appendix C. Digital Technology Adoption and Energy Intensity in Manufacturing

Depreciation rates and deflators for non-software capital are also taken from the EU KLEMS database. Software capital deflators are retrieved from Eurostat (see Table C.1).

C.3 Additional Descriptive Statistics

C.3.1 Number of Observations per Year

Table C.2: Number of observations per year. We point out that the last panel sequence includes slightly fewer observations than the first two.

Panel sequence Year	Year									
	1			2				3		Total
	2009	2010	2011	2012	2013	2014	2015	2016	2017	
% multi-unit firms	13.5%	13.8%	11.3%	11.3%	13.5%	13.8%	13.8%	13.9%	13.9%	13.2%
Observations	13,886	14,196	13,671	13,672	14,139	13,931	13,581	13,306	12,980	123,362

C.3.2 Distribution of Energy Prices

Energy prices may be endogenous as they depend, for instance, on the chosen quantity. To solve this issue, we calculate a second price variable using external energy prices (P_E [external]). We use prices of different energy sources (if available) from official statistics and weight them by the firm-level use of the respective energy source (see Table C.1). The distribution of external prices is displayed in Figure C.2 and is similar to internal energy prices (Figure C.1), but the distribution is less skewed to the right. We use the external energy price variable in a later robustness check (see Column (3) and Column (4) of Table C.8).




Figure C.1: Distribution of P_E .




Figure C.2: Distribution of P_E [external].

C.3.3 Details on the Distribution of S_E

Table C.3: Detailed descriptive statistics on the distribution of S_E .

	mean	sd	p5	p50	p95
S_E	0.090	0.102	0.012	0.055	0.292
Observations	123,362				

C.3.4 Plausibility of Differences in Software Usage with Respect to Industry and Regional Characteristics

In the following, we analyse whether industry-level and regional differences with respect to software usage are plausible. Figure C.3 shows the average software capital intensity for different industries. Manufacturers of wearing apparel (Division 14) and basic pharmaceutical products (Division 21) show the highest average software capital intensity. The pharmaceutical industry (combined with the chemical industry) was the most digital German manufacturing industry in 2018 according to Weber et al. (2018). The high software capital intensity of the wearing apparel industry can be explained by the fact that it is a market with highly interconnected supply chains and fast-changing trends. In addition, digitalisation allows for an increased individualisation of products, which is especially important for this industry. Furthermore, it is also intuitive that the computer industry (Division 26) uses more software than most other industries. Manufacturers of other transport equipment (Division 30) may have a comparatively high software capital intensity because related industries, such as aircraft and spacecraft construction, are highly innovative.




Figure C.3: Average software capital intensity by industry between 2009 and 2017. Each bar relates to an industry at the two-digit NACE level. The tobacco industry is excluded because of too few observations.

The geographic distribution of software capital intensity is displayed in Figure C.4. The darker the blue colour of the respective area, the higher the average software capital intensity. The white area in between marks regions for which we either

Appendix C. Digital Technology Adoption and Energy Intensity in Manufacturing

observe no or fewer than three enterprises.¹⁴³ We find that areas with a very high software capital intensity coincide with major German cities. For example, Berlin, Munich, Dresden, Stuttgart, and Hanover show very high values. As digital enterprises usually concentrate in larger cities, we consider this as a further indicator that software capital is suitable for measuring the firm-level degree of digitalisation.




Figure C.4: Average software capital intensity by region between 2009 and 2017. The dark blue regions represent those with the highest average software capital intensity. Regions with fewer than three observations per year or with no observations are not displayed.

¹⁴³As the RDC is not allowed to provide information at this granular level due to German data protection laws.

C.4 Calculation of Energy Cost Savings per Software Investment in the Year of Investment

By Equation (C.6), we initially calculate relative improvements in energy use in the year of investment related to the increase in software capital (relative savings). To do this, we multiply the energy intensity elasticity ($\epsilon_{E/Y, K_{SW}}$) by the relative change in software capital ($\Delta \ln K_{ICT}$) for each firm i in year t :

$$\text{relative savings}_{i,t} = \epsilon_{E/Y, K_{SW_{i,t}}} \times \Delta \ln K_{SW_{i,t}}. \quad (\text{C.6})$$

To calculate savings in energy consumption, we assume that output is constant and calculate how much energy consumed in the previous period has been saved in the current period with respect to changes in the software capital stock. Savings in energy costs are then approximated by multiplying savings in energy consumption by the firm-specific energy price:

$$\text{cost savings}_{i,t} = \text{relative savings}_{i,t} \times E_{i,t-1} \times P_E. \quad (\text{C.7})$$

In order to estimate the average savings in energy costs per euro invested in software, we sum up firm-level energy cost savings in the year of investment over all periods for which we have information and divide them by the sum of all software investments that have taken place in the same time period:

$$\text{savings per investment} = \frac{\sum_{t=2010}^{2017} \sum_{i=1}^N \text{cost savings}_{i,t}}{\sum_{t=2010}^{2017} \sum_{i=1}^N \text{software investments}_{i,t}}. \quad (\text{C.8})$$

C.5 Technological Illustration of the Heterogeneity Bias

For illustrating the issue, we consider a uni-variate model in which changes in the energy cost share are regressed on the ICT intensity growth rate at the aggregation level a :¹⁴⁴

$$\Delta S_{E_{a,t}} = \beta_{EK_{ICT_h}} \Delta \ln \left(\frac{K_{ICT}}{Y} \right)_{a,t} + \Delta e_{a,t}. \quad (C.9)$$

Assume that $\beta_{EK_{ICT_h}}$ captures the group-specific link between growth in ICT intensity in each industry and changes in the energy cost share of group h .¹⁴⁵ In the spirit of Imbs & Mejean (2015), this link differs by the amount of η_h with $E[\eta_h] = 0$ between industries:

$$\beta_{EK_{ICT_h}} = \beta_{EK_{ICT}} + \eta_h. \quad (C.10)$$

Previous industry-level studies constrain the link between ICT and energy to homogeneity and push heterogeneity into the residual ($\Delta u_{a,t} = \Delta e_{a,t} + \eta_h \Delta \ln (K_{ICT}/Y)_{a,t}$) by estimating:

$$\Delta S_{E_{a,t}} = \beta_{EK_{ICT}} \Delta \ln \left(\frac{K_{ICT}}{Y} \right)_{a,t} + \Delta u_{a,t}. \quad (C.11)$$

According to Imbs & Mejean (2015), the respective point estimate can than be written as:

$$\hat{\beta}_{EK_{ICT}} = \beta_{EK_{ICT}} + \frac{\text{cov}(\Delta \ln K_{ICT}/Y_{a,t}, \Delta u_{a,t})}{\text{var}(\Delta \ln K_{ICT}/Y_{a,t})}, \quad (C.12)$$

where $\text{cov}(\cdot)$ [$\text{var}(\cdot)$] denotes the covariance (variance) operator. For simplicity, $\Delta \ln K_{ICT}/Y$ is now expressed as z and the notation EK_{ICT} is omitted from $\hat{\beta}$ and β . After rearranging, plugging in Equation (C.10), and assuming no endogeneity bias, the expression can be written as:

$$\hat{\beta} - \beta = \frac{\text{cov}(z_{a,t}, \eta_h z_{a,t})}{\text{var}(z_{a,t})}. \quad (C.13)$$

Equation (C.13) is further rearranged. $E(\cdot)$ denotes the expectation operator. The expectation of β equals the expectation of β_h . Moreover, since β is a constant, it can be placed either inside or outside the expectation operator. If not controlled for appropriate fixed effects, then its bias can be expressed as:

¹⁴⁴Please note here that an aggregation bias due to systematic slope differences can also occur when aggregate control variables correlate with the variable of interest at the firm level as suggested by Haque et al. (1999) and Theil (1971).

¹⁴⁵Note that a can be equal to h . h may be the level of aggregation at which heterogeneous effects occur.

$$\begin{aligned}
 \hat{\beta} - \beta &= \frac{\text{cov}(z_{a,t}, \eta_h z_{a,t})}{\text{var}(z_{a,t})} \\
 &= \frac{E(z_{a,t} \cdot \eta_h z_{a,t}) - E(z_{a,t}) \cdot E(\eta_h z_{a,t})}{\text{var}(z_{a,t})} \\
 &= \frac{E(z_{a,t}^2(\beta_h - \beta)) - E(z_{a,t})E((\beta_h - \beta)z_{a,t})}{\text{var}(z_{a,t})} \\
 &= \frac{E(z_{a,t}^2\beta_h) - \beta E(z_{a,t}^2) - E(z_{a,t})E(\beta_h z_{a,t} - \beta z_{a,t})}{\text{var}(z_{a,t})} \\
 &= \frac{\text{cov}(z_{a,t}^2, \beta_h) - E(z_{a,t})E(\beta_h z_{a,t} - \beta z_{a,t})}{\text{var}(z_{a,t})}.
 \end{aligned} \tag{C.14}$$

However, most studies that employ industry-level panel data control for average effects at the industry level, such as Schulte et al. (2016). Hence, when applying fixed effects at the level of a , Equation (C.15) determines the bias:

$$\begin{aligned}
 \hat{\beta}_{FE_a} - \beta_{FE_a} &= \frac{\text{cov}((z_{a,t} - \bar{z}_a)^2, \beta_h) - E(z_{a,t} - \bar{z}_a)E(\beta_h(z_{a,t} - \bar{z}_a) - \beta(z_{a,t} - \bar{z}_a))}{\text{var}(z_{a,t} - \bar{z}_a)} \\
 &= \frac{\text{cov}((z_{a,t} - \bar{z}_a)^2, \beta_h)}{\text{var}(z_{a,t} - \bar{z}_a)}.
 \end{aligned} \tag{C.15}$$

Now, it can be seen that a heterogeneity bias can evolve if the variance of a demeaned variable of interest within one industry correlates with group-specific slopes, even if controlled for specific fixed effects.

In our case, the bias becomes zero if all industry-specific ICT intensity growth rates share the same variance, if the variance is uncorrelated with β_h or if slopes are homogeneous. However, the bias becomes negative (positive), i.e., b will be biased downwards (upwards) if industry-specific ICT intensity growth rates that face a larger variance correlate with a more (less) negative $\beta_{EK_{ICT}_h}$.

C.6 Additional Econometric Estimations

C.6.1 Robustness Checks with Respect to a Potential Measurement Error

Calculation of Software Capital Stocks

In the following, we provide robustness checks with respect to different modifications in the calculation of software capital stocks and respective growth rates. For instance, we analyse how results change if we calculate software capital stocks, assuming depreciation rates of 25, 33, and 50%. Also, different period lengths are employed to calculate initial capital stocks: We estimate initial software capital stocks based on two, four, and six observation periods if available. Table C.4 shows that changes in the depreciation rate of the software capital stock only lead to marginal differences between coefficients. Hence, the results appear to be robust in this regard. The results for different maximum lengths of observation periods considered for the initial capital stock calculation are displayed in Table C.5. We find slight differences for initial software capital stocks that include two as well as up to six periods. For initial stocks based on two periods, we find effects that are marginally smaller. For initial stocks based on up to six periods, the effect size is slightly larger and the software coefficient becomes -0.0003 . However, we do not consider this deviation to be large enough to have an effect on the economic interpretation of results.

Table C.4: Equation (4.8) with software capital stocks modified by different depreciation rates.

	dependent variable: ΔS_E		
	(1)	(2)	(3)
$\Delta \ln(\frac{P_E}{P_L})$	0.0284*** (61.74)	0.0284*** (61.74)	0.0284*** (61.74)
$\Delta \ln(\frac{K_{SW}}{Y})$	-0.000219*** (-5.13)	-0.000237*** (-5.18)	-0.000242*** (-5.16)
$\Delta \ln(\frac{K_N}{Y})$	-0.0015*** (-3.76)	-0.0015*** (-3.77)	-0.0015*** (-3.78)
$\Delta \ln(Y)$	0.0014* (2.56)	0.0013* (2.51)	0.0013* (2.49)
Year	x	x	x
Industry	x	x	x
Multi-unit	x	x	x
Federal state	x	x	x
Size class	x	x	x
EEG exemption	x	x	x
Producer	x	x	x
Trading	x	x	x
Observations	89,653	89,653	89,653
Adjusted R^2	0.271	0.271	0.271

Notes: Column (1): Depreciation rate is 25%. Column (2): Depreciation rate is 33%. Column (3): Depreciation rate is 50%. t statistics in parentheses. First-difference estimation. Clustered standard errors. ⁺ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table C.5: Equation (4.8) with software capital stocks modified by different lengths of periods considered for the initial capital stock calculation.

	dependent variable: ΔS_E		
	(1)	(2)	(3)
$\Delta \ln(\frac{P_E}{P_L})$	0.0284*** (61.74)	0.0284*** (61.74)	0.0284*** (61.74)
$\Delta \ln(\frac{K_{SW}}{Y})$	-0.000167*** (-4.68)	-0.000212*** (-4.25)	-0.000325*** (-4.69)
$\Delta \ln(\frac{K_N}{Y})$	-0.0015*** (-3.83)	-0.0015*** (-3.82)	-0.0015*** (-3.82)
$\Delta \ln(Y)$	0.0014** (2.61)	0.0013* (2.52)	0.0012* (2.29)
Year	x	x	x
Industry	x	x	x
Multi-unit	x	x	x
Federal state	x	x	x
Size class	x	x	x
EEG exemption	x	x	x
Producer	x	x	x
Trading	x	x	x
Observations	89,653	89,653	89,653
Adjusted R^2	0.271	0.271	0.271

Notes: Column (1): 2 periods maximal included for initial software capital stock calculation. Column (2): 4 periods maximal included for initial software capital stock calculation. Column (3): 6 periods maximal included for initial software capital stock calculation. t statistics in parentheses. First-difference estimation. Clustered standard errors. + $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

“Zero” Software Capital Stocks

One reason why software capital stocks may be imprecisely estimated may be that a large share of firms do not report any software investments. As we impute these capital stocks in every period by an obligatory euro, they can neither rise nor shrink. However, it is not clear whether this is accurate, since firms could very well have invested in software before the observation period and their software capital stock would actually decrease in the observed time frame. Another problem may occur for firms that do not invest in periods used to calculate initial capital stocks but start to invest afterwards. We then observe huge percentage increases in software capital stocks, since change rates from “zero” to large natural numbers are large by construction.¹⁴⁶

We conduct the following robustness checks to analyse issues with respect to “zero” software capital stocks. Firstly, we exclude all observations that have “zero” software capital stocks as well as those observations that have a software capital stock that increases from “zero” and re-estimate our model. This approach allows us to measure to what extent results differ when potentially problematic observations are excluded. Secondly, we look closer at observations that have “zero” software capital stocks. Hence, their software capital growth rate is zero. We do not know whether firms actually have acquired no software capital or whether they invested before the observation period and their software capital stock decreases due to depreciation. To analyse whether this makes a difference, we impute “zero” growth rates. We replace them by the logarithmic change rate that we would have observed in a firm that has software capital but does not invest in the current period. Hence, in a further re-estimation, software capital decreases by the depreciation rate for all observations that do not invest. Thirdly, we deal with the issue that some software capital stocks increase from “zero”. This may result in implausibly large growth rates. Therefore, we censor very large values that increase from “zero”. We consider increases more than 5-fold as implausible, limit them at this threshold, and re-estimate the model.

Results with respect to issues related to “zero” software capital stocks can be found in Table C.6 in the Appendix. In the first column, we exclude observations that have “zero” software capital stocks as well as observations that have a software capital stock that increases from “zero”. The effect size of the software coefficient is comparable to our preferred specification; however, it is only significant at the 10% level. Consequently, even though the exclusion of potentially problematic observations does only marginally alter the coefficient size, we have to acknowledge that the relationship is now significant at a lower threshold. The second column displays results for imputed depreciation rates for “zero” software capital stocks that would

¹⁴⁶In fact, they actually rise from €1 as zero values are imputed.

Appendix C. Digital Technology Adoption and Energy Intensity in Manufacturing

occur if a firm had invested in previous periods. We find that this modification does not notably affect results and the coefficient of interest is comparable to baseline results. Hence, it does not make a difference whether we depreciate “zero” software capital stocks or not. The last two columns relate to estimates in which increases from “zero” software capital stocks are limited to a threshold. In Column (3), “zero” software capital stocks are additionally excluded and in Column (4) included. The results in both columns are comparable and the coefficient slightly increases but not substantially. To sum up, different treatments of growth rates related to “zero” software capital stocks only marginally affect our results, i.e. they can hardly be the reason why we find smaller effects in comparison to aggregated estimates.

Table C.6: **Robustness checks with respect to “zero” software capital stocks.**

	dependent variable: ΔS_E			
	(1)	(2)	(3)	(4)
$\Delta \ln(\frac{P_E}{P_L})$	0.0270*** (50.15)	0.0284*** (61.74)	0.0270*** (51.13)	0.0284*** (61.74)
$\Delta \ln(\frac{K_{SW}}{Y})$	-0.000225+ (-1.87)	-0.000238*** (-5.14)	-0.000356*** (-4.86)	-0.000350*** (-4.77)
$\Delta \ln(\frac{K_N}{Y})$	-0.0017*** (-3.42)	-0.0015*** (-3.78)	-0.0015** (-3.17)	-0.0015*** (-3.76)
$\Delta \ln(Y)$	0.0011+ (1.70)	0.0013* (2.52)	0.0011+ (1.69)	0.0012* (2.32)
Year	x	x	x	x
Industry	x	x	x	x
Multi-unit	x	x	x	x
Federal state	x	x	x	x
Size class	x	x	x	x
EEG exemption	x	x	x	x
Producer	x	x	x	x
Trading	x	x	x	x
Observations	65,226	89,653	66,841	89,653
adj. R^2	0.258	0.271	0.260	0.271

Notes: Column (1): Potentially problematic observations excluded. Columns (2): “Zero” growth rates imputed. Column (3): Growth rates starting from “zero” limited and “zero” capital stocks excluded. Column (4): Growth rates starting from “zero” limited. t statistics in parentheses. First-difference estimation. Clustered standard errors. + $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Additional Robustness Checks

Moreover, we conduct a robustness check with respect to the economic crisis and exclude observations before 2011. We also estimate our model only with single-unit firms to analyse to what extent our results may be biased due to inaccurately matched information in multi-unit firms. Additionally, we test whether the inclusion of tangible capital may lead to multicollinearity issues, as software investments are often complementary. Further, we estimate a specification in which we consider software capital intensity from the previous period to examine whether there is a time lag in effects. For this purpose, we apply fixed effects instead of first difference, as this allows us to consider all observations that are included in the main analysis. In an additional specification, we include industry-year fixed effects.

Table C.7 shows effects for single-unit firms (Column [1]) and estimation results, in which only observations after 2011 are considered (Column [2]). The restricted estimates are consistent with our baseline results. Both software coefficients point in a negative direction and are significant, but software coefficients are slightly smaller for both restricted samples. Moreover, our results are also robust with respect to the exclusion of tangible capital (Column [3]). We additionally estimate the influence of lagged software capital and find a negative but insignificant effect for lagged software capital (Column [4]). Including industry-year-level fixed effects does not affect baseline results notably (Column [5]).¹⁴⁷

¹⁴⁷In addition, we also performed an estimation in which we replaced software capital intensity with software investments and measured a comparable relationship. The results are available from the authors upon request.

Table C.7: Further robustness checks.

	dependent variable:				
	ΔS_E			S_E	ΔS_E
	(1)	(2)	(3)	(4)	(5)
$\Delta \ln(\frac{E}{L})$	0.0292*** (52.32)	0.0290*** (59.68)	0.0284*** (61.74)	0.0328*** (43.27)	0.0284*** (61.80)
$\Delta \ln(\frac{K_{SW}}{Y})$	-0.000209*** (-4.52)	-0.000177*** (-3.78)	-0.000249*** (-5.37)	-0.000205** (-3.10)	-0.000228*** (-4.95)
$\Delta \ln(\frac{K_M}{Y})$	-0.0013*** (-3.30)	-0.0015*** (-3.81)		-0.000120 (-3.13)	-0.0014*** (-3.60)
$\Delta \ln(Y)$	0.00087 (1.39)	0.0015* (2.71)	0.0028*** (7.43)	0.00318 (4.09)	0.0018*** (3.37)
$\Delta \ln(\frac{K_{SW}}{Y})_{t-1}$				-0.0000400 (-0.56)	
Year	x	x	x	x	x
Industry	x	x	x	x	x
Multi-unit	x	x	x	x	x
Federal state	x	x	x	x	x
Size class	x	x	x	x	x
EEG exemption	x	x	x	x	x
Producer	x	x	x	x	x
Trading	x	x	x	x	x
Economic industry \times Year					x
Observations	62,821	77,029	89,653	89,653	89,653
adj. R^2	0.285	0.284	0.271	0.277	0.281

Notes: After 2011. Column (2): Single-unit firms, i.e., firms with only one business location. Column (3): Not controlled for tangible capital. Column (4): Additional lagged K_{SW} and estimated in fixed effects (to ensure the same amount of observations). Column (5): Additional industry-year fixed effects. t statistics in parentheses. First-difference estimation except Column (4), which uses fixed effects. Clustered standard errors. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

C.6.2 IV Estimates

Endogenous control variables do not lead to biased coefficients when uncorrelated with the variable of interest, but they do if they relate to each other (Frisch & Waugh 1933). In particular, the relationship between output and software capital could potentially bias results, because both variables may highly depend on each other. A similar problem may exist with respect to the energy-labour-price ratio, as the use of software usually requires skills that are in high demand. To test whether these issues affect results, we conduct the following IV estimates as further robustness checks. The results are displayed in Table C.8.

For the analysis of endogeneity problems related to our output indicator, a problem arises because the variable is included in the model and at the same time both capital stocks are scaled by output. However, from an econometric perspective it is not necessary to scale capital stocks by output since we already control for it.¹⁴⁸ Accordingly, we rearrange Equation (4.3) and do not scale both capital stocks by output anymore. Hence, we now estimate $\hat{\beta}_{EY}$ instead of $\hat{\beta}_{EY}^*$ (see Section 4.3). The translog model is re-estimated and displayed in Column (1). It is straightforward to see that this modification barely affects software and tangible capital coefficients. In a second step, we instrument output by a firm's market share in terms of sales. Market shares are calculated using four-digit and two-digit industry levels and employing the Census on Investments; accordingly, two different instruments are used. Additionally, we calculate market shares only if at least five observations per industry are available. Hence, we exclude a small share of observations from the estimation. The results are displayed in Column (2). The output coefficient gently increases, but the software coefficient is barely affected. In a last step, we instrument the energy-labour-price ratio by an exogenous energy price variable. To calculate the exogenous energy price, we use prices of different energy sources (if available) from official statistics and weight them by the individual use of the respective energy source.¹⁴⁹ The results are displayed in Column (3); in Column (4) output is instrumented as well. The effect size of the price coefficient decreases and it is now significant at a lower threshold, but the software coefficient is not affected and is comparable to the baseline specification. Furthermore, tests for underidentification and weak identification, as well as the Sargan–Hansen test indicate that the exogenous energy price and market shares are appropriate instruments.

¹⁴⁸We scale software capital by output in our preferred specification to be consistent with Schulte et al. (2016).

¹⁴⁹See Table C.1 in the Appendix for more details on data sources. The distribution of the exogenous price variable and its relationship to the potentially endogenous energy price variable is displayed in C.3.2.

Table C.8: IV estimates.

	dependent variable: ΔS_E			
	(1)	(2)	(3)	(4)
$\Delta \ln(\frac{P_E}{P_L})$	0.0284*** (0.000)	0.0283*** (0.000)	0.0023* (0.019)	0.0018+ (0.068)
$\Delta \ln(K_{SW})$	-0.000242*** (0.000)	-0.000260*** (0.000)	-0.000245*** (0.000)	-0.000265*** (0.000)
$\Delta \ln(K_N)$	-0.0016*** (0.000)	-0.0019*** (0.000)	-0.0012** (0.006)	-0.0015** (0.002)
$\Delta \ln(Y)$	0.0031*** (0.000)	0.011*** (0.001)	0.0036*** (0.000)	0.0108** (0.001)
Year	x	x	x	x
Industry	x	x	x	x
Multi-unit	x	x	x	x
Federal state	x	x	x	x
Size class	x	x	x	x
EEG exemption	x	x	x	x
Producer	x	x	x	x
Trading	x	x	x	x
Observations	89,653	89,017	89,653	89,017
Underidentification		64.16	816.0	63.83
Weak identification		42.61	1139.2	27.93
P-value Hansen J statistic		0.864		0.909

Notes: Column (1): Capital not scaled by output. Column (2): Y instrumented. Column (3): P_E/P_L instrumented. Column (4): Y and P_E/P_L instrumented. p -values in parentheses. First-difference estimation. Clustered standard errors. + $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. The underidentification test displays the Kleibergen-Paap LM statistic and the weak identification test displays the Kleibergen-Paap Wald F-statistic.

C.6.3 Pooled OLS, FE, and Mundlak

Table C.9: Pooled OLS, FE, and Mundlak (all model coefficients).

	dependent variable: S_E		
	(1)	(2)	(3)
$\ln(\frac{P_E}{P_L})$	-0.00376*** (-8.73)		
$\ln(\overline{\frac{P_E}{P_L}})$			-0.00494*** (-5.83)
$\Delta \ln(\frac{P_E}{P_L})$		0.0316*** (57.23)	0.0316*** (57.21)
$\ln(\frac{K_{SW}}{Y})$	-0.00165*** (-32.95)		
$\ln(\overline{\frac{K_{SW}}{Y}})$			-0.00151*** (-16.13)
$\Delta \ln(\frac{K_{SW}}{Y})$		-0.000213*** (-3.48)	-0.000214*** (-3.48)
$\ln(\frac{K_N}{Y})$	0.0061*** (30.85)		
$\ln(\overline{\frac{K_N}{Y}})$			0.00499*** (14.62)
$\Delta \ln(\frac{K_N}{Y})$		-0.000116 (-0.26)	-0.000116 (-0.26)
$\ln(Y)$	0.0397*** (69.70)		
$\ln(\overline{Y})$			0.0356*** (29.68)
$\Delta \ln(Y)$		0.00396*** (6.18)	0.00396*** (6.18)
Year	x	x	x
Industry	x	x	x
Multi-unit	x	x	x
Federal state	x	x	x
Size class	x	x	x
EEG exemption	x	x	x
Producer	x	x	x
Trading	x	x	x
Observations	123,362	123,362	123,362
R^2	0.577	0.268	0.268

Notes: Column (1): Pooled OLS specification. Column (2): Fixed effects specification. Column (3): Mundlak specification. t statistics in parentheses. Robust standard errors in Column (1). Clustered standard errors in Column (2) and Column (3). * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

C.6.4 Industry-Specific Effects – Regression Table

Table C.10: Differences across industries – Regression results.

dependent variable: ΔS_E												
	(1) N10	(2) N11	(3) N12	(4) N14	(5) N15	(6) N16	(7) N17	(8) N18	(9) N19	(10) N20	(11) N21	(12) N22
$\Delta \ln(\frac{P_E}{P_L})$	0.0489*** (26.23)	0.0586*** (12.01)	0.0405*** (10.53)	0.0199*** (5.83)	0.0165*** (5.75)	0.0235*** (9.41)	0.0419*** (10.18)	0.0289*** (8.43)	0.0536*** (6.19)	0.0365*** (13.51)	0.0264*** (7.11)	0.0322*** (15.69)
$\Delta \ln(\frac{K_{SW}}{Y})$	-0.000220 (-1.45)	-0.000331 (-1.05)	-0.000167 (-0.72)	0.000284 (1.31)	0.000509 (1.00)	-0.000245 (-0.59)	-0.000516+ (-1.88)	-0.0000985 (-0.40)	0.000292 (0.35)	-0.000426* (-2.01)	0.000158 (0.51)	-0.000209 (-1.07)
$\Delta \ln(\frac{K_N}{Y})$	-0.00488* (-2.24)	-0.00308 (-0.71)	-0.000797 (-0.49)	-0.000863 (-0.67)	-0.00689 (-1.51)	0.00309 (0.86)	0.000420 (0.16)	-0.000133 (-0.08)	-0.000993 (-0.84)	-0.00116 (-0.36)	-0.00176 (-0.96)	-0.00172 (-1.10)
$\Delta \ln(Y)$	0.000123 (0.05)	-0.00398 (-0.86)	0.000311 (0.91)	0.000459 (0.30)	-0.00546 (-1.08)	0.0138** (2.96)	0.00475 (1.04)	-0.00349 (-1.14)	0.000910 (0.47)	0.00996* (2.42)	-0.000809 (-0.25)	0.00214 (0.89)
$\bar{\epsilon}_{E/Y, K_{SW}}$	(-)	(-)	(-)	(-)	(-)	(-)	-0.0068	(-)	(-)	-0.0083	(-)	(-)
$\epsilon_{E/Y, K_{SW}}$ at $\bar{S}_{K_{SW}}$	(-)	(-)	(-)	(-)	(-)	(-)	-0.0031	(-)	(-)	-0.0029	(-)	(-)
Observations	11,043	1,483	1,975	1,079	566	1,838	2,522	1,698	217	4,965	1,090	4,825
Adjusted R^2	0.355	0.439	0.318	0.316	0.193	0.204	0.293	0.343	0.422	0.263	0.330	0.304

dependent variable: ΔS_E											
	(13) N23	(14) N24	(15) N25	(16) N26	(17) N27	(18) N28	(19) N29	(20) N30	(21) N31	(22) N32	(23) N33
$\Delta \ln(\frac{P_E}{P_L})$	0.0451*** (14.01)	0.0448*** (10.53)	0.0311*** (26.78)	0.0144*** (10.33)	0.0169*** (16.34)	0.0180*** (23.08)	0.0224*** (11.54)	0.0238*** (8.86)	0.0246*** (9.48)	0.0145*** (18.01)	0.0158*** (13.18)
$\Delta \ln(\frac{K_{SW}}{Y})$	-0.000790* (-2.42)	-0.000338+ (-1.74)	-0.0000753 (-0.63)	-0.000163 (-1.06)	-0.000228+ (-1.92)	-0.0000248 (-0.38)	-0.0000923 (-0.52)	0.000293 (1.42)	0.0000621 (0.33)	0.0000947 (0.75)	-0.000322+ (-1.91)
$\Delta \ln(\frac{K_N}{Y})$	-0.00246 (-1.26)	-0.00216 (-0.99)	-0.00117 (-1.40)	-0.0000138 (-0.01)	0.000422 (0.52)	-0.000687 (-1.34)	-0.00127 (-1.07)	-0.00171 (-1.03)	0.0000232 (0.01)	0.00113 (1.38)	-0.00148 (-1.63)
$\Delta \ln(Y)$	0.00749** (2.60)	0.00858** (2.94)	-0.000107 (-0.09)	0.000141 (0.93)	0.000617 (0.61)	-0.000665 (-1.06)	-0.000833 (-0.61)	-0.00198 (-0.90)	-0.000207 (-0.09)	-0.00259+ (-1.89)	-0.00162 (-1.19)
$\bar{\epsilon}_{E/Y, K_{SW}}$	-0.0093	-0.0038	(-)	(-)	-0.110	(-)	(-)	(-)	(-)	(-)	-0.0272
$\epsilon_{E/Y, K_{SW}}$ at $\bar{S}_{K_{SW}}$	-0.0043	-0.0020	(-)	(-)	-0.0058	(-)	(-)	(-)	(-)	(-)	-0.105
Observations	4,314	4,021	10,942	3,737	5,157	12,079	3,173	1,079	1,797	3,093	2,402
Adjusted R^2	0.333	0.310	0.385	0.260	0.324	0.342	0.278	0.382	0.361	0.337	0.348

Notes: N10: Food products; N11: Beverages; N13: Textiles; N14: Wearing apparel; N15: Leather & related products; N16: Wood, wood & cork products; N17: Paper & paper products; N18: Printing, recorded media; N19: Coke, refined petroleum products; N20: Chemicals & chemical products; N21: Basic pharmaceutical products; N22: Rubber & plastic products; N23: Other non-metallic mineral products; N24: Basic metals; N25: Fabricated metal products; N26: Computer, electro, optical products; N27: Electrical equipment; N28: Machinery and equipment n.e.c.; N29: Motor vehicles, (semi-)trailers; N30: Other transport equipment; N31: Furniture; N32: Other manufacturing; N33: Repair and installation. The following additional control variables are included: Year, multi-unit, federal state, size class, EEG exemption, producer, and trading. Firms that switch between industries are excluded. The tobacco industry is excluded because of too few observations and a very low R-squared. *t* statistics in parentheses. First-difference estimation. Clustered standard errors. + $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Appendix D

What Drives the Relationship Between Digitalisation and Industrial Energy Demand?

D.1 Description of Variables

Table D.1: Description of variables.

Main variables	
total energy use	Overall firm-level energy use, i.e. the sum of energetic use of different energy carriers plus electricity use (in kWh) observed in the AFiD-Module Use of Energy.
electricity use	Total electricity consumption (in kWh) observed in the AFiD-Module Use of Energy.
fossil fuel use	The sum of firm-level use of natural gas, coal, heating oil, district heat, and liquid petroleum gas (in kWh) observed in the AFiD-Module Use of Energy.
treatment	Software capital approximates the degree of firm-level digitalisation (in €). We calculate firm-level software capital stocks as in Chapter 4 and base them on software investments reported in the AFiD-Panel Industrial Units. Firstly, we generate real software investments using software deflators from Eurostat. Secondly, we apply the perpetual inventory method (PIM) to estimate capital stocks (Griliches 1980, Lutz et al. 2017). We consider a depreciation rate of 31.5%. The value is retrieved from the EU KLEMS database. ¹⁵⁰ Based on these software capital stocks, we calculate software capital growth rates and dichotomise them to generate treatment W . Accordingly, W is one if the software capital stock of firm i increases in period t and zero otherwise. It has to be acknowledged that we only account for purchased software capital and firms may also use open source software. We refer to Appendix C.5 for a detailed description of the calculation of software capital stocks and to Section 4.4.2 for a discussion of their representativeness for the firm-level degree of digitalisation. Moreover, Section C.6.1 contains robustness checks with different depreciation rates for the translog model applied in Chapter 4. The link between software capital intensity and energy intensity appear to robust to different depreciation values (25%, 33%, and 50%).

¹⁵⁰ See EU KLEMS database - 2019 release, Germany capital input data, see Stehrer, R., A. Bykova, K. Jäger, O. Reiter and M. Schwarzhappel (2019): Industry level growth and productivity data with special focus on intangible assets, wiiw Statistical Report No. 8. https://euklems.eu/excel/DE_Capital_SDB_2019.xlsx [Online; accessed on 11 Apr. 2023] (Retrieved on: 18.04.2020).

Appendix D. What Drives the Relationship Between Digitalisation and Industrial Energy Demand?

Table D.1: Description of variables.

Covariates	
lagged outcome	We include the lagged outcome in log-levels in the estimation. If we integrated change rates from the previous period, we would need to consider $t - 2$ as well. This would imply that we lose a large share of observations, as our panel is imbalanced.
output	We take the gross production value and subtract turnover from trade and other activities to calculate output (in €). All variables are observed in the AFiD-Panel Industrial Units.
tangible capital	Tangible capital is calculated using real investments in property, plant, and equipment (AFiD-Panel Industrial Units, in €) and applying the PIM. Deflators and depreciation rates are taken from the EU KLEMS data. We refer also to Chapter 4 for a detailed description of the calculation of tangible capital stocks.
labour use	Labour use is measured by the number of employees observed in the AFiD-Panel Industrial Units. We convert part-time employees to full-time employees and adjust the number of employees in this regard.
producer price index	Average material prices are approximated by the index of producer prices of industrial products (domestic sales) retrieved from Destatis. https://www.destatis.de/DE/Themen/Wirtschaft/Preise/Erzeugerpreisindex-gewerbliche-Produkte/Publikationen/Downloads-Erzeugerpreise/erzeugerpreise-1-ange-reihen-pdf-5612401.html [Online; accessed on 11 Apr. 2023] (retrieved on: 12.11.2020).
price energy	We calculate overall firm-level energy prices in a two-step procedure. Firstly, we divide firm-level energy costs by firm-level energy use to approximate the firm-level energy price (in €/kWh). However, this approach may be endogenous and prone to issues resulting from misreporting. Consequently, we calculate in a second step, based on the firm-level energy price, the average energy price within an industry (four-digit NACE level) in one region (five-digit AGS level [Kreisbene]). We then approximate the firm-level energy price by the regional industry average. This allows considering a more robust energy price. Moreover, if we observe less than five firms in a region within one industry, we approximate the firm-level energy price by the federal-state average at the two-digit NACE level.
price electricity	Electricity prices are retrieved from Eurostat (status: 08.04.2019, in €/kWh). We consider prices for non-household consumers, which is bi-annual data and we, therefore, take the yearly average. Moreover, prices depend on the consumption level and we exclude VAT and other recoverable taxes and levies. As firms switch their consumption level over time, we consider a firm's consumption level of the first period that we observe to match prices. This allows for not considering price variations due to changes in the consumption level. https://ec.europa.eu/eurostat/de/ [Online; accessed on 11 Apr. 2023] (retrieved on: 15.07.2020).
price natural gas	Natural gas prices are also retrieved from Eurostat (status: 10.02.2020, in €/GJ). We consider bi-annual natural gas prices (average price per year is calculated) for non-household consumers. Natural gas prices depend on the consumption level. Accordingly, we consider a firm's consumption level of the first period that we observe to match prices. Prices are retrieved excluding VAT and other recoverable taxes and levies. Natural gas prices are converted from GJ to kWh. https://ec.europa.eu/eurostat/de/ [Online; accessed on 11 Apr. 2023] (retrieved on: 15.07.2020).
prices of other energy carriers	Other energy prices are retrieved from the IEA (liquid petroleum gas, retrieved on: 04.09.2019), Destatis & DEPI (biomass, retrieved on: 16.07.2020 [Destatis], and retrieved on: 13.09.2019 [DEPI]), AGFW (district heat, retrieved on: 14.08.2019), and BMWK former BMWi (heating oil, retrieved on: 01.04.2020). For a more detailed description on sources for energy prices see Appendix C.1.

Appendix D. What Drives the Relationship Between Digitalisation and Industrial Energy Demand?

Table D.1: Description of variables.

share of energy sources	To consider the energy mix, we divide the use of electricity, natural gas, coal, heating oil, district heat, liquid petroleum gas, and biomass by overall energy consumption and consider each share as a variable in the Causal Forest model. All variables are observed in the AFiD-Module Use of Energy. We include each energy share in lagged levels in our estimation.
R&D intensity	We divide the total expenditure on research & development observed in the AFiD-Panel Industrial Units by output.
tax intensity	The amount of taxes (e.g. property tax, motor vehicle tax, excise duties; excluding income and corporation tax, equalization levies on burdens and VAT) observed in the AFiD-Panel Industrial Units is divided by output.
subsidy intensity	The amount of subsidies received for current production in the business year observed in the AFiD-Panel Industrial Units is divided by output.
trading intensity	The total turnover of trading goods during the business year observed in the AFiD-Panel Industrial Units is divided by output. Trading goods are considered to be goods of foreign origin that are generally resold unprocessed and without a production-related connection to own products.
Herfindahl–Hirschman Index	The HHI captures the competitive situation that a firm has to face. It is calculated using yearly revenue-based market shares at the four-digit NACE level observed in the AFiD-Panel Industrial Units. For a detailed description of the HHI, see Rhoades (1993). We exclude industries for which we observe less than five firms per year.
share renewable production	Own electricity generation from renewable power observed in the AFiD-Module Use of Energy is divided by overall energy use.
share fossil production	Own electricity generation from fossil sources observed in the AFiD-Module Use of Energy is divided by overall energy use.
weak region	We include a dummy indicating whether a firm is situated in a region (five-digit AGS level) that is considered as “structurally weak” (0) due to its limited economic productivity or “structurally strong” (1). An overview map of structurally weak regions can be found at: https://www.bmwi.de/Redaktion/DE/Dossier/Digital-Jetzt/digital-jetzt-infografik-strukturschwache-regionen.html [Online; accessed on 11 Apr. 2023] (retrieved on: 14.03.2022)
EEG exemption	A one-hot encoded variable is generated that indicates whether a firm is not (1), partly (2), or fully (3) exempted from charges under the law on renewable energies (EEG). This is calculated by means of the approximated ratio between electricity costs and value added as well as electricity use. For this purpose, we combine information from the AFiD-Panel Industrial Units and the AFiD-Module Use of Energy.
energy intensive industry	We define an energy-intensive industry as an industry or a group of industries at the two-digit NACE level that accounts for more than 5 % of total energy consumption of the manufacturing sector (Divisions: 10-12, 17,19, 20, 23, 24). The information is retrieved from the German Environmental Agency (https://www.umweltbundesamt.de/daten/umwelt-wirtschaft/industrie/branchenabhaengiger-energieverbrauch-des#primarenergienutzung-des-verarbeitenden-gewerbes) [Online; accessed on 11 Apr. 2023]).
main industrial grouping	We add a one-hot encoded variable that indicates the industrial main group of the firm: intermediate goods producer (1), capital goods producer (2), durable goods producer (3), consumer goods producer (4), and energy producer (5).
single- / multi-unit firm	A one-hot encoded variable is included that indicates whether a firm is a single-unit firm (1), a multi-unit firm in one federal-state (2), or a multi-unit firm in several federal states (3).
industry association	A LASSO-based fixed effects vector controlling for the industry assignment is calculated based on two-digit NACE codes. For a detailed description, see Jens et al. (2021).
year	A LASSO-based fixed effects vector controlling for the observation period (year) is generated based on Jens et al. (2021).

Appendix D. What Drives the Relationship Between Digitalisation and Industrial Energy Demand?

Table D.1: Description of variables.

federal states

A LASSO-based fixed effects vector controlling for the federal state of the firm's registered office is calculated based on Jens et al. (2021).

D.2 Descriptive Statistics

Table D.2: Averages and standard errors of firm characteristics for treated and untreated firms (2010 to 2017).

variable	mean control	s.d. control	mean treated	s.d. treated
total energy use (in GWh)	29.64	414.65	45.14	396.57
total energy use ln Δ	0.02	0.27	0.03	0.27
total energy use ln $t - 1$	14.72	1.9	15.2	1.92
electricity use (in GWh)	9.61	74.67	16.18	132.21
electricity use ln Δ	0.02	0.27	0.03	0.28
electricity use ln $t - 1$	13.88	1.9	14.42	1.9
fossil fuel use (in GWh)	17.71	281.98	25.89	262.55
fossil fuel use ln Δ	0.05	1.38	0.06	1.34
fossil fuel use ln $t - 1$	13.23	3.55	13.84	3.29
treatment (W)	0.0	0.0	1.0	0.0
output (in million €)	62.15	572.48	122.86	1250.08
output ln Δ	0.04	0.2	0.06	0.21
output ln $t - 1$	16.53	1.44	17.11	1.44
tangible capital (in million €)	16.55	152.71	33.75	326.58
tangible capital ln Δ	0.02	0.24	0.05	0.22
tangible capital ln $t - 1$	14.78	1.93	15.51	1.76
number of employees	241.07	1699.02	431.28	2836.52
number of employees ln Δ	0.01	0.13	0.03	0.13
number of employees ln $t - 1$	4.59	1.09	5.04	1.16
producer-price index	99.5	3.63	99.57	3.74
producer-price index ln Δ	0.01	0.03	0.01	0.03
price energy	0.13	0.03	0.13	0.03
price energy ln Δ	0.01	0.12	0.01	0.13
price biomass	0.04	0.0	0.04	0.0
price biomass ln Δ	0.02	0.08	0.02	0.07
price coal	0.01	0.0	0.01	0.0
price coal ln Δ	0.02	0.19	-0.0	0.21
price district heat	0.07	0.0	0.07	0.0
price district heat ln Δ	0.01	0.05	0.01	0.05
price electricity	0.14	0.03	0.14	0.02
price electricity ln Δ	0.04	0.06	0.05	0.06
price heating oil	0.06	0.01	0.06	0.01
price heating oil ln Δ	0.0	0.22	-0.01	0.24
price liquid petroleum gas	0.06	0.01	0.06	0.01
price liquid petroleum gas ln Δ	-0.02	0.16	-0.03	0.18
price natural gas	0.04	0.01	0.04	0.01
price natural gas ln Δ	-0.01	0.09	-0.02	0.1
biomass share [in %]	0.02	0.1	0.02	0.09
biomass share [in %] $t - 1$	0.02	0.1	0.01	0.09
coal share [in %]	0.01	0.06	0.01	0.06
coal share [in %] $t - 1$	0.01	0.05	0.01	0.06
district heat share [in %]	0.03	0.12	0.04	0.13
district heat share [in %] $t - 1$	0.03	0.12	0.03	0.13
electricity share [in %]	0.49	0.25	0.51	0.24
electricity share [in %] $t - 1$	0.49	0.25	0.51	0.24
natural gas share [in %]	0.32	0.29	0.33	0.28
natural gas share [in %] $t - 1$	0.31	0.29	0.32	0.28
N	64,933		27,382	

Appendix D. What Drives the Relationship Between Digitalisation and Industrial Energy Demand?

Table D.2: Averages and standard errors of firm characteristics for treated and untreated firms (2010 to 2017).

variable	mean control	s.d. control	mean treated	s.d. treated
heating oil share [in %]	0.11	0.22	0.09	0.19
heating oil share [in %] $t - 1$	0.12	0.22	0.1	0.2
liquid petroleum gas share [in %]	0.01	0.06	0.01	0.05
liquid petroleum gas share [in %] $t - 1$	0.01	0.06	0.01	0.05
R&D intensity [in %]	0.01	0.03	0.02	0.04
R&D intensity [in %] Δ	-0.0	0.02	0.0	0.02
tax intensity [in %]	0.01	0.02	0.01	0.03
tax intensity [in %] Δ	-0.0	0.01	-0.0	0.01
subsidy intensity [in %]	0.0	0.01	0.0	0.01
subsidy intensity [in %] Δ	-0.0	0.01	0.0	0.01
trading intensity [in %]	0.11	1.06	0.12	0.5
trading intensity [in %] Δ	0.01	1.0	-0.0	1.54
share self-produced fossil-based energy [in %]	0.0	0.02	0.01	0.03
share self-produced fossil-based energy [in %] Δ	0.0	0.01	0.0	0.02
share self-produced renewable energy [in %]	0.01	0.04	0.01	0.04
share self-produced renewable energy [in %] Δ	0.0	0.03	0.0	0.03
HHI	0.07	0.09	0.07	0.09
HHI Δ	0.0	0.03	0.0	0.03
no EEG exemption [in %]	0.92	0.28	0.92	0.28
partial EEG exemption [in %]	0.08	0.27	0.07	0.26
full EEG exemption [in %]	0.01	0.09	0.01	0.1
single-unit firm	0.78	0.41	0.75	0.43
multi-unit firm in one federal state [in %]	0.08	0.27	0.07	0.26
multi-unit firm in several federal states [in %]	0.14	0.35	0.18	0.38
intermediate goods producer [in %]	0.44	0.5	0.44	0.5
capital goods producer [in %]	0.29	0.45	0.34	0.47
durable goods producer [in %]	0.04	0.19	0.04	0.2
consumer goods producer [in %]	0.23	0.42	0.18	0.39
energy producer [in %]	0.0	0.02	0.0	0.03
year 2010 [in %]	0.14	0.35	0.13	0.34
year 2011 [in %]	0.14	0.34	0.15	0.36
year 2012 [in %]	0.06	0.23	0.07	0.25
year 2013 [in %]	0.14	0.35	0.13	0.34
year 2014 [in %]	0.15	0.36	0.13	0.34
year 2015 [in %]	0.15	0.35	0.13	0.34
year 2016 [in %]	0.08	0.28	0.13	0.33
year 2017 [in %]	0.14	0.35	0.12	0.33
Industry: Food products [in %]	0.14	0.35	0.09	0.29
Industry: Beverages [in %]	0.02	0.13	0.02	0.12
Industry: Tobacco products [in %]	0.0	0.0	0.0	0.0
Industry: Textiles [in %]	0.02	0.15	0.02	0.14
Industry: Wearing apparel [in %]	0.01	0.11	0.01	0.1
Industry: Leather & related products [in %]	0.01	0.08	0.0	0.07
Industry: Wood, wood & cork products [in %]	0.02	0.15	0.02	0.13
Industry: Paper & paper products [in %]	0.03	0.17	0.03	0.16
Industry: Printing, recorded media [in %]	0.02	0.14	0.02	0.13
Industry: Coke, refined petroleum products [in %]	0.0	0.02	0.0	0.03
Industry: Chemicals & chemical products [in %]	0.06	0.23	0.06	0.24
Industry: Basic pharmaceutical products [in %]	0.01	0.1	0.02	0.12
Industry: Rubber & plastic products [in %]	0.06	0.23	0.06	0.24
Industry: Other non-metallic mineral prod. [in %]	0.05	0.22	0.05	0.21
Industry: Basic metals [in %]	0.05	0.21	0.05	0.21
N	64,933		27,382	

Appendix D. What Drives the Relationship Between Digitalisation and Industrial Energy Demand?

Table D.2: Averages and standard errors of firm characteristics for treated and untreated firms (2010 to 2017).

variable	mean control	s.d. control	mean treated	s.d. treated
Industry: Fabricated metal products [in %]	0.13	0.34	0.12	0.33
Industry: Computer, electro, optical prod. [in %]	0.04	0.2	0.05	0.23
Industry: Electrical equipment [in %]	0.06	0.24	0.07	0.25
Industry: Machinery and equipment n.e.c. [in %]	0.13	0.34	0.18	0.38
Industry: Motor vehicles, (semi-)trailers [in %]	0.04	0.19	0.05	0.21
Industry: Other transport equipment [in %]	0.01	0.12	0.02	0.12
Industry: Furniture [in %]	0.02	0.14	0.02	0.14
Industry: Other manufacturing [in %]	0.04	0.19	0.04	0.19
Industry: Repair and installation [in %]	0.03	0.17	0.02	0.15
N	64,933		27,382	

D.3 Evaluating Assumptions and the Fit of the Causal Forest

To assess the overlap assumption, we plot the propensity scores that indicate the probability of treatment for each observation in Figure D.1. Note that we lose four observations due to trimming. The histograms for the (trimmed) treated and not treated firms overlap in a way that makes it impossible to deterministically decide on the treatment status of a firm, as the scores are bounded away from zero and one. Hence, the overlap assumption is fulfilled.




Figure D.1: Distribution of propensity scores between treatment and control group.

Furthermore, the internal validity of the Causal Forest approach is based on the idea that the treatment is random. Therefore, we assess the balance of the covariates between firms that increase software capital and firms that do not increase software capital. Figure D.2 depicts these differences in the distributions of the treated and untreated samples after re-weighting the covariates with the inverse propensity score (as we apply AIPW weights). Except for some very rare outliers, the distributions do not show any notable differences between both groups. Thus, our model is able to appropriately balance covariates.

Additionally, we test the model calibration by comparing OOB predictions to actual changes in energy consumption (see Equation 5.4). The results for all three model outcomes are presented in Table D.3. According to the test results, the model for electricity use seems to be calibrated well and the performance is comparable to the model with energy use as an outcome variable. In contrast, the model using fossil fuels as an outcome variable fails in predicting an average treatment effect that is different from zero and does not seem to capture the underlying heterogeneity adequately. Model results of this outcome should therefore be interpreted cautiously.

Appendix D. What Drives the Relationship Between Digitalisation and Industrial Energy Demand?

Table D.3: **Best Linear Predictor Test for the forest with all outcomes.**

Outcome variable	Coefficient	Estimate	SE	difference from zero		difference from one	
				<i>t</i> -stat	<i>p</i> -value	<i>t</i> -stat	<i>p</i> -value
Energy use	β_{ATE}	0.998	0.235	4.245	1.09e05***	-0.008	0.993
	β_{CATE}	1.261	0.366	3.448	0.0003***	0.713	0.475
Electricity use	β_{ATE}	0.980	0.172	5.695	1.96e-09***	-0.116	0.907
	β_{CATE}	0.914	0.316	2.897	0.002***	-0.272	0.785
Fossil fuel use	β_{ATE}	1.442	3.870	0.373	0.355	0.1142	0.909
	β_{CATE}	-0.833	0.784	-1.061	0.856	-2.338	0.019*

Notes: Results of the Best Linear Predictor Test for model calibration and heterogeneity that seeks to fit the estimated CATE as a linear function of the out-of-bag predictions (see Equation 5.4). Difference from zero: one-tailed t-test that tests whether estimated coefficients are significantly larger than zero; difference from one: two-tailed t-test that tests whether estimated coefficients significantly differ from one. Please note that the test results of the “difference from one” test are approximated by the presented rounded β -coefficients and standard errors, as well as 10,000 degrees of freedom.

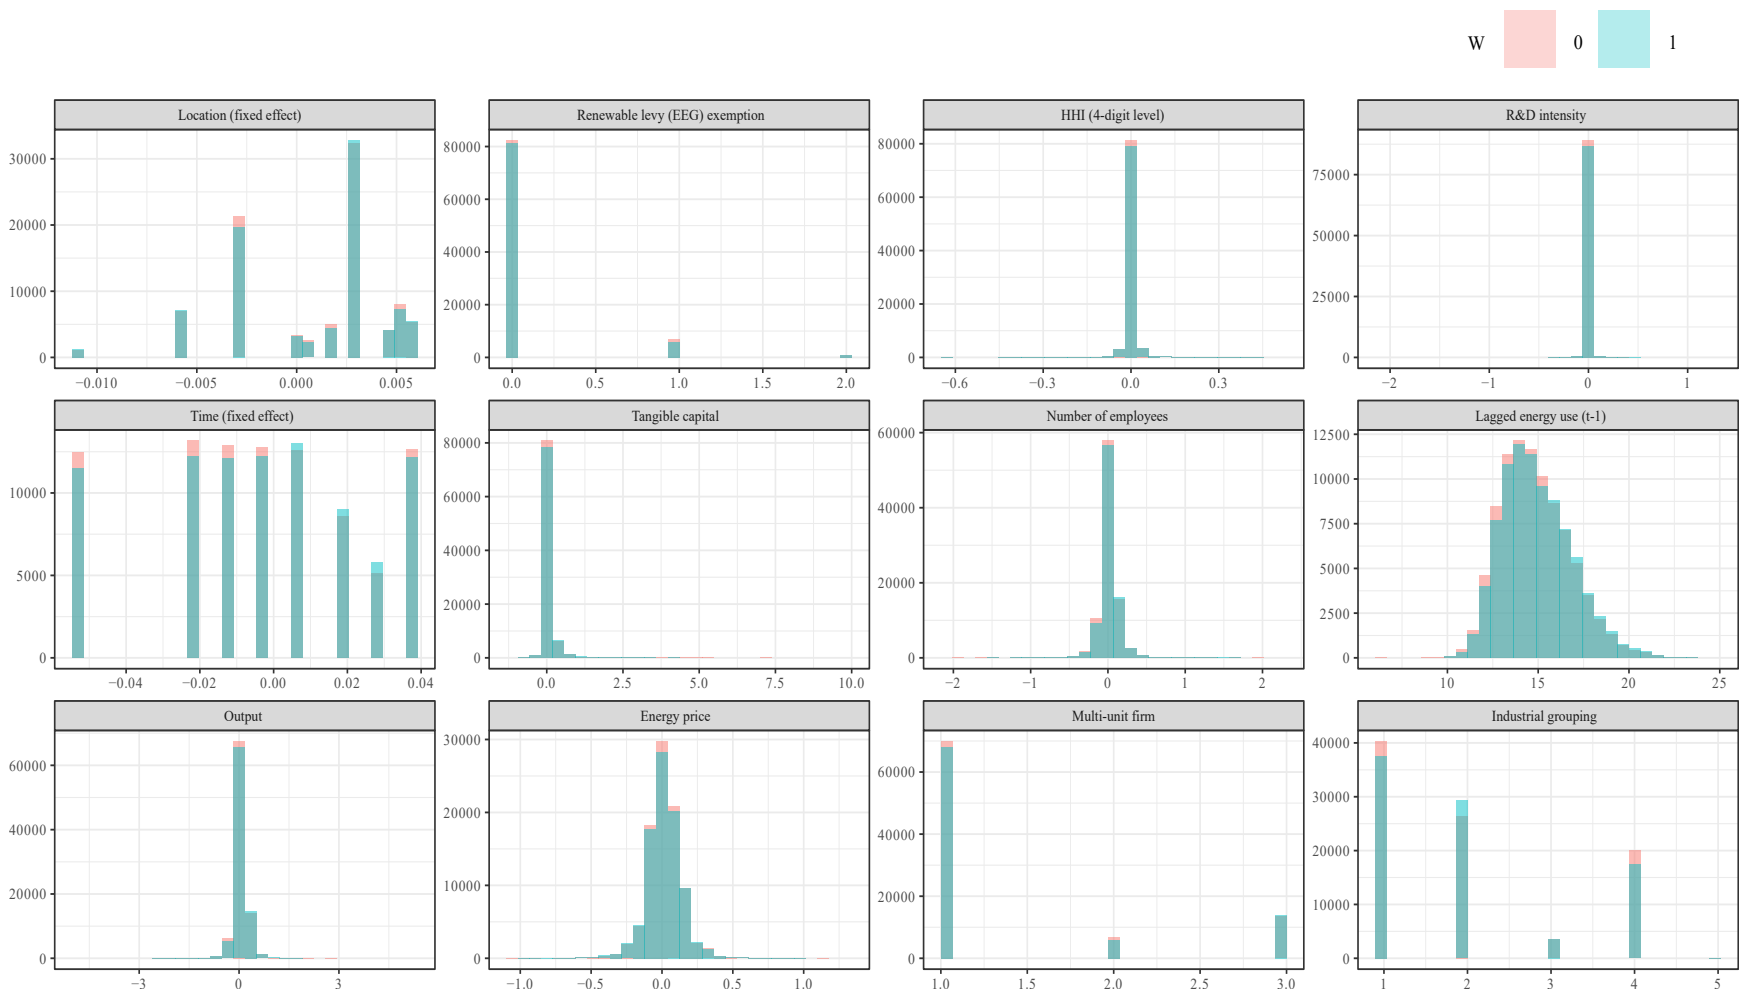


Figure D.2: Inverse-propensity weighted histograms for treated and untreated observations (Part I).

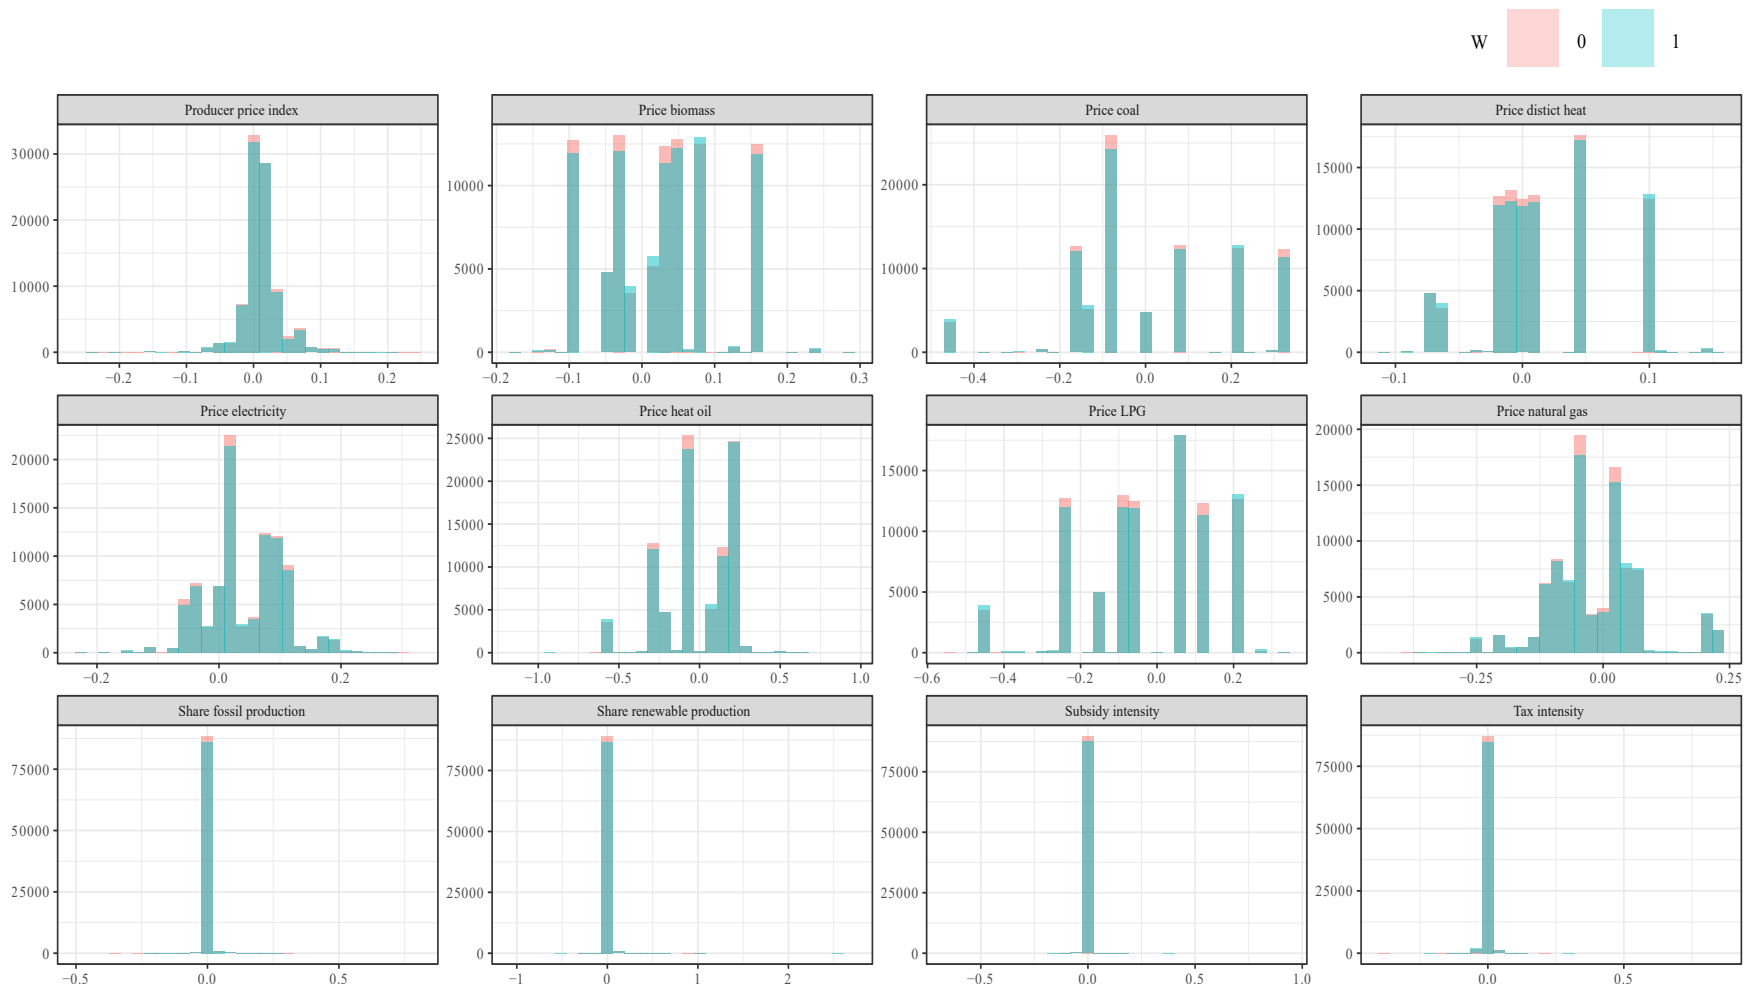


Figure D.3: Inverse-propensity weighted histograms for treated and untreated observations (Part II).

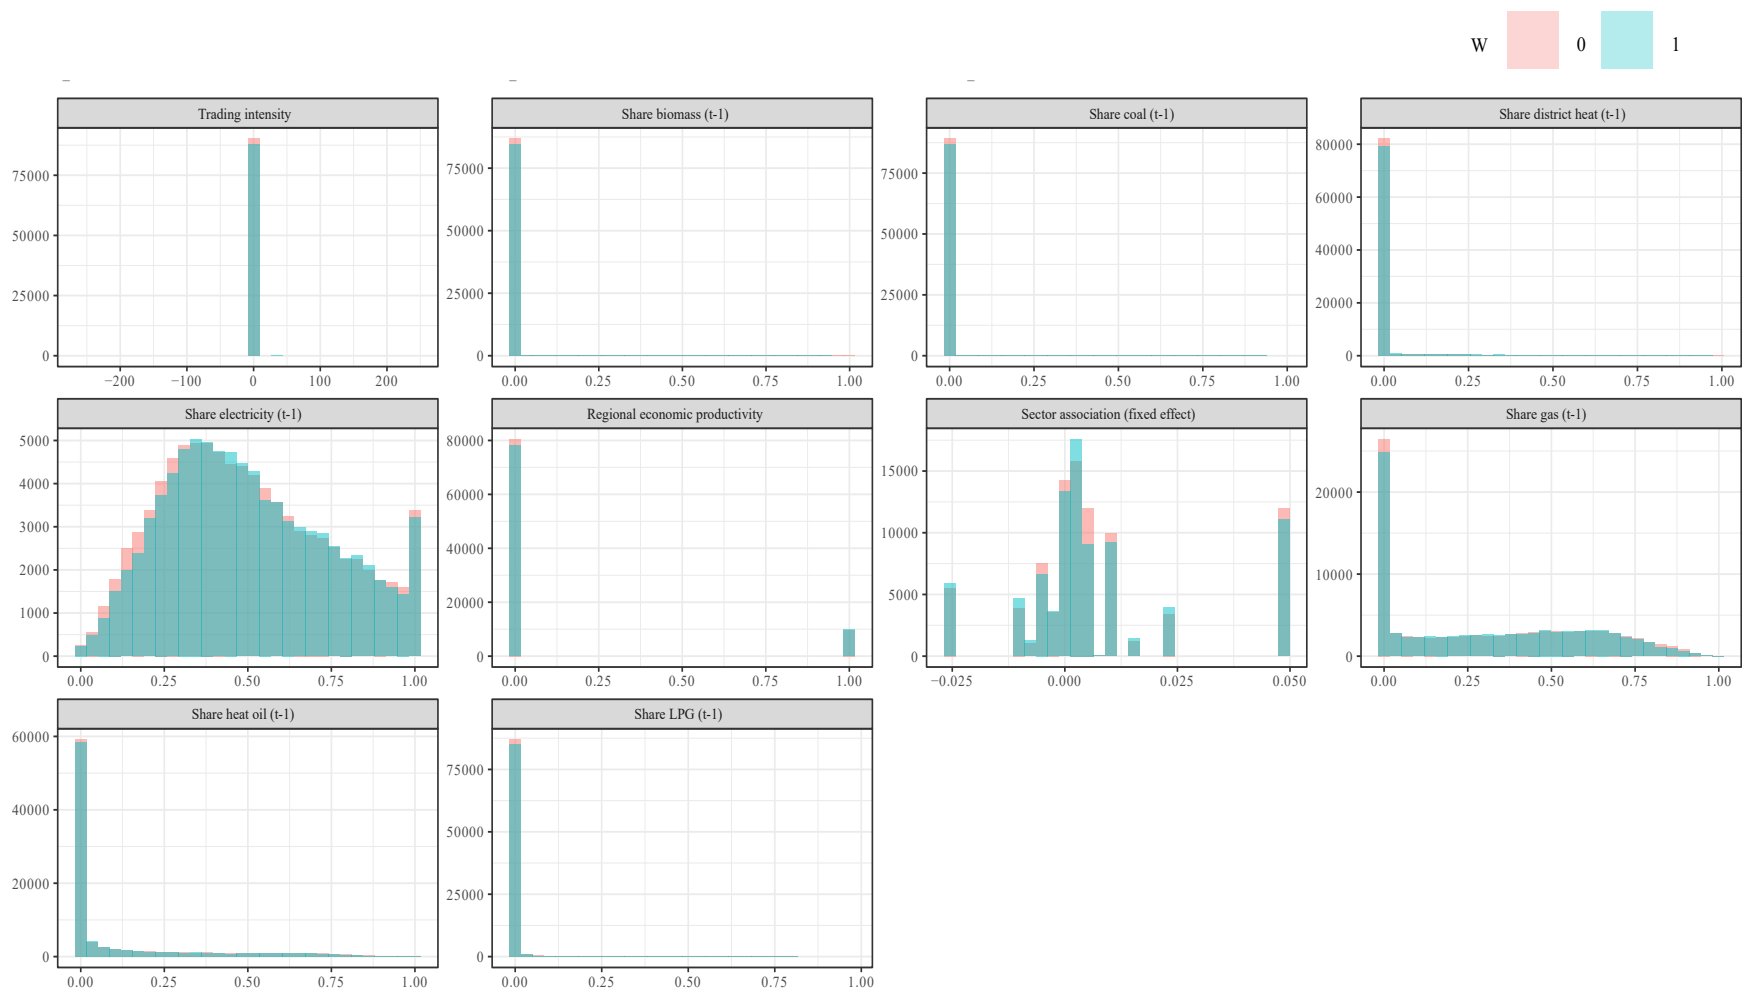


Figure D.4: Inverse-propensity weighted histograms for treated and untreated observations (Part III).

D.4 Variable Importance

To understand the main drivers of treatment heterogeneity, we analyse the firm characteristics that were used as splitting variables in the forest. Variable importance (VI) measures how many times a covariate was used for splitting at level l across all trees t , where relative split frequency (RSF) denotes the split frequency (SF) of variable m divided by all splits at level l . Additionally, weights ($w_l = l^{-2}$) are used that exponentially favour higher tree levels.

$$VI_m = \frac{\sum_{l=1}^L RSF_{ml} * w_l}{\sum_{l=1}^L w_l} \quad (D.1)$$

$$RSF_{ml} = \frac{SF_{ml}}{\sum_{m=1}^M SF_{ml}} \quad (D.2)$$




Figure D.5: Variable Importance for the three Causal Forests with the outcome variables total energy use, electricity use, and fossil fuels.

Figure D.5 lists the 15 most important variables for splitting the sample into groups. Combined, energy prices are by far the most important variable if we sum up the importance values of the energy prices (total energy price, price of heat oil, price of LPG, price of electricity, and price of coal). Furthermore, the share of electricity use (of the previous period) is important for the splitting procedure of all three outcomes.

D.5 Robustness Analysis

We conduct the following robustness analysis:

Growth rates: We repeat our analysis and replace output, tangible capital, and labour use in lagged levels by logarithmic growth rates as one alternative specification.

D/L: In a further robustness test, we modify our treatment variable and consider only firms as treated if their software capital per employee increases. We repeat the analysis and replace the treatment dummy D with a relative dummy that represents the change in the capital stock (ΔK_{ICT}) relative to the number of employees (L).

Table D.4 shows estimated ATEs as well as the performance of the Best Linear Prediction Test for both robustness checks. In the first specification (growth rates), the ATE is now at 0.01 and significant. The mean and differential forest prediction indicate that treatment effects are well calibrated. Hence, different formulations of production function in- and outputs do not alter main results. However, it is noteworthy that the variable importance of these critical variables increases when they are included in logarithmic growth rates (not displayed). In the second robustness test (D/L), the ATE is now at 0.006, which is slightly smaller than in our main specification and the p-value is at 0.12. The Best Linear Prediction Test confirms that the model is well calibrated. Hence, we can also confirm heterogeneity by our modified digitalisation indicator.

Table D.4: **Robustness tests.**

Robustness type	Outcome variable	Variable	Estimate	SE	t -stat	p -value
Growth rates	Energy use	ATE	0.010	0.003	3.330	0.0004***
		β_{ATE}	1.052	0.308	3.413	0.0003***
		β_{CATE}	1.117	0.420	2.660	0.0004**
D/L	Energy use	ATE	0.006	0.005	1.180	0.119
		β_{ATE}	1.011	0.459	2.202	0.012*
		β_{CATE}	0.815	0.418	1.950	0.026*

Notes: Results for the different robustness models.

Appendix D. What Drives the Relationship Between Digitalisation and Industrial Energy Demand?

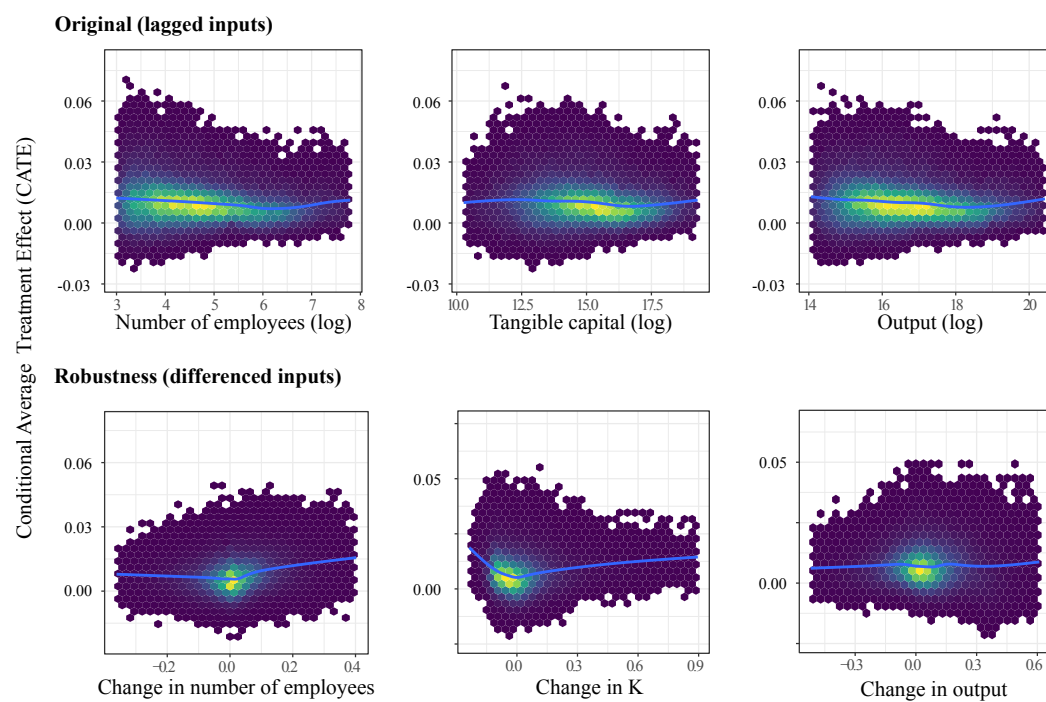


Figure D.6: **Bivariate distributions of the predicted treatment effect and production factors.** Comparison of the original model (lagged input variables) and the robustness check (logarithmic growth rates).

Appendix D. What Drives the Relationship Between Digitalisation and Industrial Energy Demand?

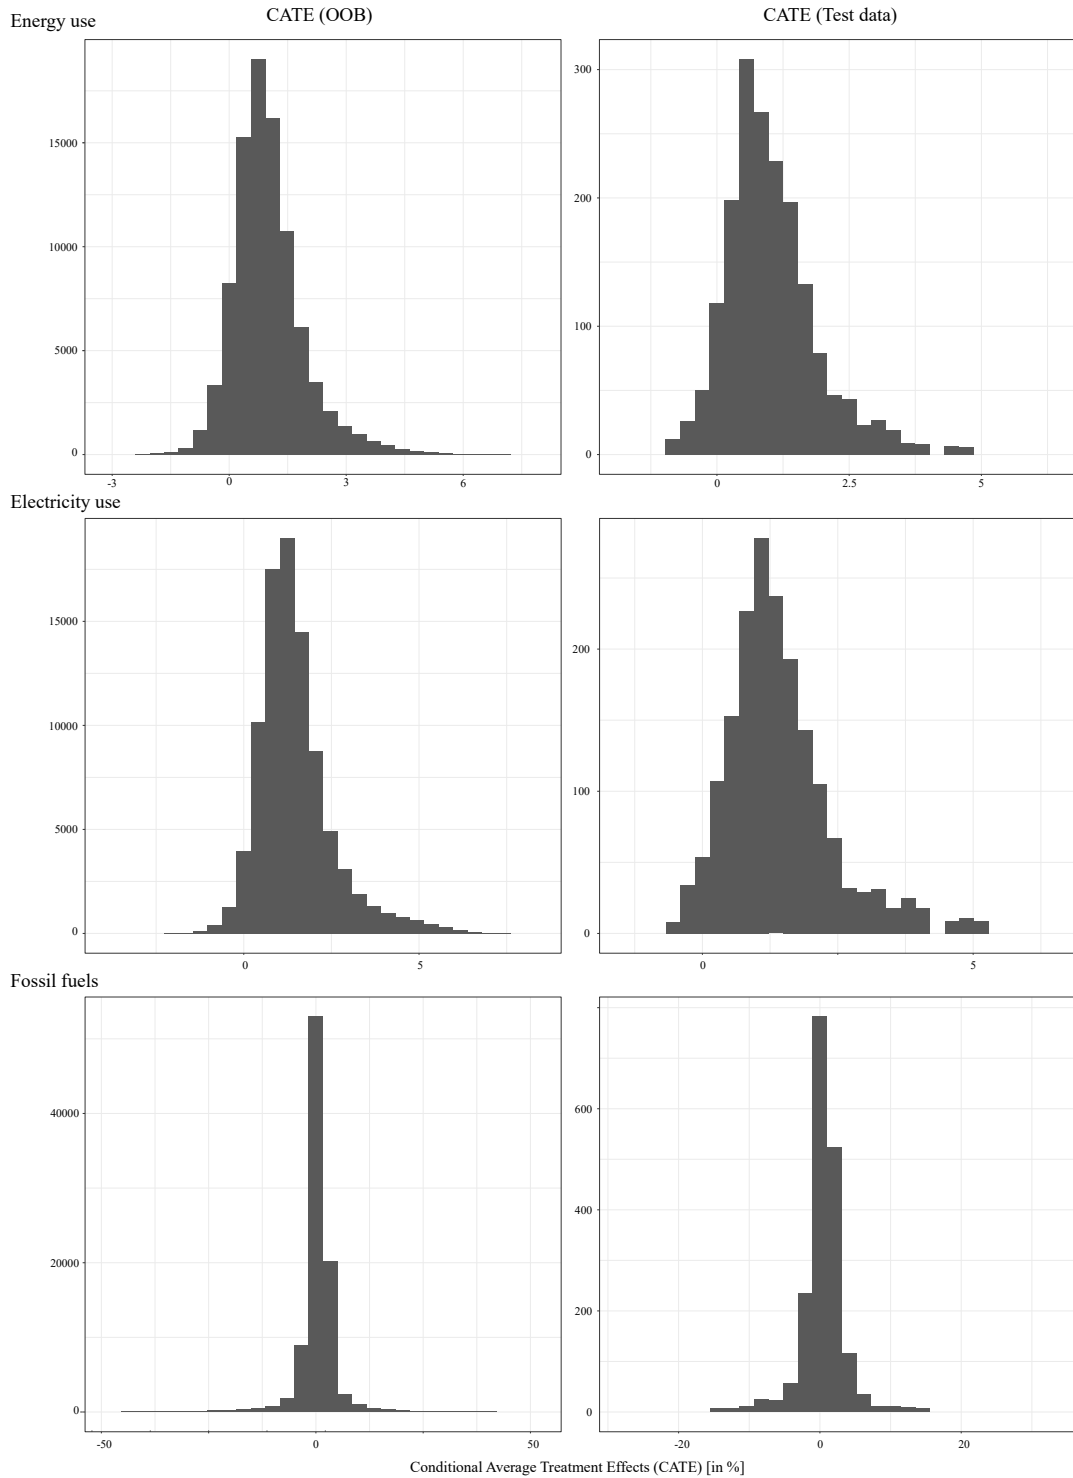


Figure D.7: Comparison between OOB predictions and predictions of the test sample.

Appendix E

Link Between ICT Adoption and Changes in Total Energy Consumption, Output, as well as Energy Intensity

To analyse how an increase in ICT adoption relates to changes in total energy consumption, output, and energy intensity, I estimate the following linear model:

$$\Delta \ln Z_{i,t}^j = \alpha + \beta \frac{P_E}{P_{PPI}} \Delta \ln \left(\frac{P_E}{P_{PPI}} \right)_{i,t} + \beta_W W_{i,t} + \sum_{c \in C} \gamma_c c_{i,t} + \Delta u_{i,t}. \quad (\text{E.1})$$

Z captures changes in total energy consumption (E), output (Y), and energy intensity (E/Y), respectively (with $j \in E, Y, E/Y$). W captures a dummy that is one if firm i has an increase in software capital in period t and zero otherwise. I additionally control for changes in the energy price relative to changes in the producer price index and include the same control variables as in Chapter 4 (yearly fixed effects, industry dummies, a dummy indicating whether a firm has multiple units, federal state dummies, two dummies relating either to a full or partial exemption of the EEG levy, a dummy that controls for whether a firm produces energy, and a dummy that is set to one if a firm is trading commodities). Table E.1 shows that if firms increase their software capital, energy consumption approximately increases by 1% and output grows by 1.95%. Simultaneously, energy intensity roughly improves by 0.9%.¹⁵¹ Hence, output increases to a greater extent than total energy consumption when firms invest in digital technologies. This phenomenon resolves the contradiction that we observe energy intensity improvements in Chapter 4, but an average increase in energy consumption in Chapter 5.

¹⁵¹ Assuming that roughly 30% of all firms increase their software capital stock every year, the average yearly improvement in energy intensity is a bit larger but qualitatively comparable with the value derived in Chapter 4.

Appendix E. Link Between ICT Adoption and Changes in Total Energy Consumption, Output, as well as Energy Intensity

Please be aware that the estimation strategy in Chapter 5 includes a matching and the results of Table E.1 include a comparison of firms that may be very different from each other. Hence, the results of both approaches involve slightly different implications.

Table E.1: Link between ICT adoption and changes in total energy consumption, output, and energy intensity.

	(1)	(2)	(3)
	$\Delta \ln E$	$\Delta \ln Y$	$\Delta \ln(E/Y)$
$\Delta \ln(\frac{P_E}{P_{PPI}})$	-0.406*** (-55.22)	0.0532*** (16.62)	-0.459*** (-61.02)
W	0.0100*** (6.86)	0.0195*** (10.05)	-0.00940*** (-4.18)
Year	x	x	x
Industry	x	x	x
Multi-unit	x	x	x
Federal state	x	x	x
Size class	x	x	x
EEG exemption	x	x	x
Producer	x	x	x
Trading	x	x	x
Observations	89267	89267	89267
Adjusted R^2	0.332	0.006	0.206

Notes: OLS estimation of Equation (E.1). t statistics in parentheses. First-difference estimation. Clustered standard errors. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Column (1): Changes in energy consumption as the dependent variable. Column (2): Changes in output as the dependent variable. Column (3): Changes in energy intensity as the dependent variable.

Declaration of Authorship

Hiermit erkläre ich, dass ich die vorgelegten Aufsätze selbstständig und nur mit den Hilfen angefertigt habe, die für den jeweiligen Aufsatz angegeben sind. In der Zusammenarbeit mit den angeführten Koautoren war ich wie angegeben anteilig beteiligt. Bei den von mir durchgeführten und in den Aufsätzen erwähnten Untersuchungen habe ich die Grundsätze guter wissenschaftlicher Praxis eingehalten, wie sie in der Satzung der Justus-Liebig-Universität Gießen zur Sicherung guter wissenschaftlicher Praxis niedergelegt sind.

Ort, Datum

Unterschrift