NORTHWESTERN UNIVERSITY

Essays in Technology, Finance, and Labor

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Managerial Economics and Strategy

By

Michał Zator

EVANSTON, ILLINOIS

June 2020

ABSTRACT

Essays in Technology, Finance, and Labor

Michał Zator

This dissertation is a wide-range study of the relationships between the three central elements of the production function: technology, capital and its financing, and labor.

Chapter 1 analyzes the relationship between labor and recent wave of automation and digitization technologies, showing that while they typically substitute for workers, in several service industries the complementarity effect dominates. This is shown in two ways: (1) labor scarcity, instrumented with population aging, increases investment in technology on average but impedes it in selected industries; (2) technology typically reduces employment but increases it in selected industries. Additional results show that financial constraints impede technology adoption and that the new technology is skill-biased. Overall, the study unwinds the heterogeneous link between new technologies and labor, highlighting the importance of analyzing a broad set of technologies and studying patterns of their adoption.

Chapter 2 studies the link between household debt and labor supply. Using income tax data from Poland and exploiting variation in floating-rate mortgage payments driven by inter-bank rates fluctuations, I show that households work and earn more when their mortgage payments

are higher. Higher income covers around 35% of the increase in the payment. The effect is stronger for households with higher payment-to-income ratio and for more flexible income sources. The increase in labor supply is accompanied by a decrease in consumption and savings and is driven by several mechanisms, including spousal labor supply, change of job, and additional income from after-hours contracts. The analysis shows that interests rates can affect labor supply of mortgage holders, with implications for monetary policy and debt relief policies.

Chapter 3 studies the effect of debt on Danish exporters' response to a negative demand shock: the 2005 boycott of Danish products in Muslim countries after publication of Muhammad caricatures. Combining balance-sheet data with firms' sales by product-destination in a triple-difference design, we find that only low-leverage firms recoup lost demand by increasing investment, introducing new products and entering new markets. In contrast, high-leverage firms reduce sales and employment, turning to outsourcing to reduce operating risk. These results highlight important flexibility costs of debt, consistent with declarations of practitioners.

Acknowledgements

I am grateful to my advisors, Efraim Benmelech, Benjamin Friedrich, Benjamin Jones, and David Matsa, for their guidance and support during the duration of my PhD studies. I thank many faculty members at Northwestern, especially Daniel Barron, Anthony DeFusco, and Nicola Persico, and fellow PhD students, in particular Roman Acosta, Edwin Munoz, Sanket Patil, and Ruozhou Yan, for numerous invaluable and inspring discussions and comments. I am thankful for continuous support and motivation coming from my beloved family: my wife Patrycja, children Filip and Teresa, and parents Janina and Ryszard. My road to PhD would not be possible without them.

Table of Contents

List of Tables

List of Figures

CHAPTER 1

Digitization and Automation: Firm Investment and Labor Outcomes

I am grateful to my committee, Efraim Benmelech, Benjamin Friedrich, Benjamin Jones and David Matsa, numerous colleagues and visitors, especially Roman Acosta, Daniel Barron, Ricardo Dahis, Joao Guereiro, Riccardo Marchingiglio, Sarit Markovich, Sara Moreira, Edwin Munoz, Jane Olmstead-Rumsey, Sanket Patil, Nicola Persico, Michael Powell, Elena Prager, Luis Rayo, Amanda Starc, Jeroen Swinkels, Miguel Talamas, Benjamin Vatter, and Ruozhou Yan; and participants of NBER Doctoral Workshop on Economics of AI for helpful comments and guidance. I thank Center for Research on Technology and Innovation at Kellogg, in particular James Conley, for financially supporting the project.

1.1. Introduction

Recent advances in artificial intelligence, robotics and other digital technologies suggest that the global economy is facing important technological changes (Brynjolfsson and McAfee, 2014; Schwab, 2017). In the past, steam engines, electricity, and computers vastly improved productivity and standards of living, but at the same time caused widespread reorganization of economic activity. Today, new technologies again raise hopes and fears, demonstrating impressive capabilities and attracting broad business and public interest.[1] What will be the impact of these technologies on the future of work across different industries, areas and types of workers?

A small existing empirical literature that studies recent technological change focuses on industrial robots and shows that robots either reduce or do not affect employment (Graetz and Michaels, 2018; Acemoglu and Restrepo, 2019; Dauth et al., 2018). Yet, robots are a specific technology class, highly concentrated in manufacturing. Their impact need not be fully informative about the effects of advances in other technology classes, such as data processing, networks or cloud computing. To have an economic impact, these technologies must first be adopted by firms, which does not happen at random. Hence, to fully understand the effects of technological change, we need to consider a broad set of technologies and study patterns of their adoption.

This paper studies the adoption and impact of the most recent wave of digitization and automation technologies. I use firm-level administrative data from Germany that cover all industries and contain measures of several new technologies. The main component of the

[1]In 2015 ImageNet Competition, algorithms' accuracy in image recognition exceeded the typical levels recorded for humans. Numerous companies have made large investment in digital technologies (e.g. in 2018 Samsung announced $22 bln investment in AI and 5G). Many countries, including Canada, China, and Germany, have created national strategies for AI.

paper studies the relationship between technology and labor and consists of two parts. The first part shows that an exogenous increase in labor scarcity faced by a firm typically results in increased investment in technology. The second part documents that adoption of technology typically leads to lower employment. Both results reveal that new technology substitutes for workers on average across the whole economy. These average effects, however, hide the important main result - significant heterogeneity across industries. The aggregate substitution pattern is driven by industries such as manufacturing or trade. In contrast, in finance or IT the complementarity effect dominates, creating the opposite pattern of adoption and employment changes. The remaining parts of the paper contain additional results, revealing significant heterogeneity by worker's skill and occupation, demonstrating that firm's financial constraints impede technology adoption, and showing how technology affects the number and size of firms as well as productivity.

The setting of the analysis is Germany – a large country at the frontier of modern technology.[2] The data contain establishment-level measures of digitization and automation coming from IAB Establishment Panel: a combination of social security records and an annual survey conducted by the German Employment Agency. The measures reflect adoption of a wide set of new technologies, including automation and modern digital technologies related to communication and data processing (e.g. robots, Internet of Things, cloud computing). I combine these measures with a rich set of variables related to firm personnel, investment, finances, and output from 1993 to 2017. I also use additional industry-level measures of technology: count of robots and capital stock of software and databases.

[2]For example, robot density in German manufacturing is now 32.2 robots per 1000 workers, with the world average being 8.5 and the comparable figure for US and France being 20 and 13.7 (based on International Federation of Robotics data).

I begin by describing the data and presenting several facts about technology usage. I then introduce a theoretical framework that guides the analysis and illustrates the interplay of the three elements of a firm's production function – technology, labor, and capital. The firm considers the adoption of a new technology in a partial equilibrium model with a task-based production function (Acemoglu and Autor, 2011). The technology can be characterized by the extent to which it substitutes or complements labor. The model shows that if the substitution effect dominates, technology adoption reduces employment and, at the same time, adoption is higher for firms facing high labor cost (or, more generally, high labor scarcity). Conversely, if the complementarity effect dominates, technology adoption increases employment and labor cost decreases adoption. Motivated by this if-and-only-if relationship implied by the model, I empirically analyze both how labor scarcity affects technology adoption and how technology affects employment.[3]

The first empirical part of the paper studies how labor scarcity affects firms' investment in digitization and automation. I measure difficulties in finding workers in four different ways that combine firms' survey declarations with actual hiring decisions. The technology adoption measure is based on firms' responses to questions about intensity of digitization and automation usage and on firms' investment expenditures. The results show that firms have higher adoption

[3]We can view these two approaches as two alternative strategies for identification of the unknown technology characteristics. At the same time, both bring a unique perspective. They are in a different position on a spectrum between partial and general equilibrium analysis and shed light not only on the characteristics of the technology but also on the determinants of corporate investment.

of these technologies when they face difficulties in finding suitable workers.[4] This association is robust to various controls, specifications, and variable definitions.

The relationship between labor scarcity and technology adoption is subject to endogeneity concerns. In particular, firms adopting the technology may be different in an unobservable way, e.g., more successful. In an initial attempt to alleviate these concerns, I show that technology adoption is accelerated by labor scarcity measured on the district level, and hence not directly related to firm characteristics. This relationship also holds for exporters, thus easing concerns that local labor scarcity proxies for local product market demand. To ease remaining concerns, I use plausibly exogenous variation in labor scarcity driven by aging patterns. A higher share of older workers in the local labor market – which is driven mostly by fertility decisions made couple of decades ago, unlikely to be related to today's technology adoption decisions – significantly increases firms' difficulties in finding workers. 2SLS regression with the 10-year change in the share of older workers as an instrument confirms the positive impact of labor scarcity on technology adoption. Additional tests use predicted aging and analyze the effect on exporters to alleviate remaining endogeneity concerns about migration and changes in product markets.

The aggregate positive effect of labor scarcity hides significant heterogeneity across industries, worker types and technology classes. Substitution – defined as a positive relationship between labor scarcity and technology adoption – dominates not only in manufacturing but also in retail and hospitality industries, among others. At the same time, several service industries, e.g., IT,

[4]Labor scarcity can be manifested either through higher price of labor or through difficulties in finding suitable workers. In a perfectly competitive labor market, the price of labor should adjust. However, because of many labor market rigidities (e.g. industry-wide and firm-wide wage agreements), German labor market is not perfectly competitive and labor scarcity is often manifested by firms being unable to find suitable workers. Economic intuition remains the same independently on whether scarcity affects prices or ability to find workers, since we can interpret the latter as high labor cost as well (e.g. high search cost).

health and education, and finance are characterized by complementarity or no clear relationship. This suggests that while the substitution effect typically dominates, the complementarity effect is strong enough to produce employment growth in some industries.[5] Looking into narrower definitions of the technology reveals that robots display the substitution pattern in manufacturing and mining while digitization generates both substitution and complementarity in various service industries. Scarcity of low skilled workers shows strong positive relationship with technology adoption, while the effect of high skilled workers scarcity is insignificant. This pattern suggests that the substitution effect dominates over the complementarity effect for low skilled workers but the two forces balance each other for high skilled workers. Further tests show that the substitution effect is higher for firms that employ many low skilled workers and many non-administrative workers. I show that this last pattern is in contrast to computers in the early 2000s, which mostly affected administrative workers.

The second part of the paper studies employment effects of technology directly. The heterogeneous relationship between labor scarcity and technology adoption implies that the employment effects should also be heterogenous. To test that, I utilize another data set, with employment records for over 2 million establishments, and use other measures of technology. Unlike other papers that rely only on industry-level variation, I study the effects of robots and digitization using within-industry variation. I analyze 10-year changes in employment in a difference in differences framework, taking the first difference across industries and the second difference across local

[5]An example of such phenomenon is the introduction of ATM in the US that did not lead to the decrease of bank clerks employment (even though ATMs clearly substituted for some workers) because banks reacted to increased productivity by opening more branches (Bessen, 2016). Other technological advances in banking, such as improvement in communication technologies, could have also contributed to the increase in productivity and employment.

areas. Global intensity of the technological change – changes in robot density and the per-worker value of software and databases investment – is a treatment that continously varies at the industry level; treatment group is firms in places where technology adoption is high; control group is firms in places with low adoption. The identifying assumption is that the difference in employment change between high- and low-adoption areas would not have systematically varied across industries but for differences driven by technological change. This specification is similar to that used in Rajan and Zingales (1998).

I find negative employment effects of automation and insignificant positive effects of digitization on average. Comparing employment change in high- and low-adoption areas reveals that technology reduces employment in industries such as manufacturing or trade, but increases it in finance, IT or education and health. This pattern exactly mimics the heterogeneous effect of labor scarcity on adoption, consistent with my theoretical model.

I complement the analysis of the relationship between technology and labor by studying another determinant of technology adoption: financial constraints. Many new technologies are explicitly designed to limit capital investment (e.g. cloud computing) and it is not clear whether financial constraints are still an important impediment to investment. I show that firms that report difficulties in obtaining credit are less likely to adopt digitization and automation. To ease endogeneity concerns, I instrument financial constraints with area's exposure to Commerzbank, which significantly cut its lending after the financial crisis (Huber, 2018). The result, which also holds outside of manufacturing, confirms the importance of access to finance even though many new technologies appear to be less capital-intensive.

The remaining part of the paper analyzes the effect of technology on skill structure of the workforce, number of firms, firm size, and labor productivity. Both digitization and automation increase the share of high-skill workers but the pattern for digitization is not always significant. Digitization increases the number of establishments and reduces the average size of the establishment, but automation appears to have the opposite effects. The association of new technology adoption and labor productivity is positive and driven mostly by robotization.

My results show that even though the last wave of new technologies typically substitutes for workers, this average effect hides significant heterogeneity across industries, worker types, and technology classes. The substitution effect dominates in manufacturing, retail or hospitality, but technology mostly complements workers in IT, finance or education and health. This empirical evidence complements recent theoretical work on digitization and automation (Acemoglu and Restrepo, 2018b; Agrawal et al., 2019; Moll et al., 2019), improving upon a small empirical literature in this area (Graetz and Michaels, 2018; Acemoglu and Restrepo, 2019; Dauth et al., 2018) in several ways: I analyze a broader set of technologies and full set of industries, study determinants of technology adoption, use firm-level data,[6] study the most recent period, and improve identification in the employment effects analysis. Previous literature focuses on a specific automation technology – industrial robots – and finds some evidence that it reduces employment. This paper shows that the effect of robots gives only a partial view of the effects of technological change overall. Other technologies that play an important role in many industries have opposite net effects and lead to employment growth. My findings suggest that while the decline in the number of available jobs may be a concern, facilitation of workers' reallocation between industries is likely to be a more pressing challenge associated with the current technological

[6]Recently, some contemporaneous papers also analyze firm-level data (Cheng et al., 2019; Koch et al., 2019; Bessen et al., 2019) but their focus is on automation and industrial robots. I analyze a broader set of technologies, ask partially different set of questions and use various empirical strategies to go beyond descriptive analysis.

change. This challenge is particularly severe because, as several of my results show, the substitution effect is concentrated among low-skill workers.

The paper also highlights the importance of studying patterns of firms' investment in technology adoption. Adoption is higher in places where labor is scarce. Hence, while in many industries technology is indeed associated with lower employment, it does not necessarily lead to a displacement of workers. Instead, it may allow firms to grow even if they have difficulties filling vacancies with labor. Moreover, even if the displacement happens, it is more likely to happen in the areas where jobs are abundant. These findings are important in the context of society aging (Abeliansky and Prettner, 2017; Acemoglu and Restrepo, 2018a) and have implications for designing place-based and industry-based policies.

The results also improve our understanding of determinants of firm investment. The negative effect of financial constraints is well known (Fazzari et al., 1988; Han Kim et al., 2019) but my findings highlight the importance of labor scarcity (Mao and Wang, 2018). Contrary to financial constraints, labor scarcity does not necessarily decrease investment. Depending on the characteristics of a particular investment project, it may encourage it or impede it (Xu, 2018). The paper also demonstrates how the analysis of corporate investment can be used to infer the effects of technological change. It is related to the broader literature on firm innovation and labor (Acharya et al., 2013; Babina and Howell, 2019; Bena et al., 2018) and on new technologies and finance (Chaboud et al., 2014; Buchak et al., 2018; Zhang, 2019; D'Acunto et al., 2019).

The paper is also related to a large literature that studies information and communication technologies at the end of the 20th century (Brynjolfsson and Hitt, 2000; Bresnahan et al., 2002; Autor et al., 2003; Autor, 2019). It documents that ICT leads to job and wage polarization and influences firm organization and productivity. Several other papers analyze older technologies (Doms et al., 1997; Lewis, 2011; Clemens et al., 2018), studying their relationship to workers' skills and the link between technology adoption and immigration.

1.2. Data and Measures of Technology

The main data used in this paper comes from the German Employment Agency and is a combination of administrative personnel records and establishment survey conducted by the Agency every year. The establishment data contains firm-level measures of digitization and automation adoption. The data is supplemented with industry-level measures of digitization and automation coming from independent sources. This section provides a brief description of the data and a descriptive analysis of the technology measures. Additional details about the data sources are presented in the Online Appendix.

1.2.1. Data Sources

I use the data from the Institute for Employment Research (IAB) of the German Employment Agency which administers several data sets based on social security records and other complementary data collection efforts. The two main administrative data sets used in this paper are IAB Establishment Panel (IAB-EP) and Establishment History Panel (BHP).[7]

IAB Establishment Panel is a yearly survey of stratified random sample of German establishments. It covers years 1993-2017 and, as of now, over 16 thousand establishments from all industries.

[7]More precisely, the study uses following datasets: weakly anonymous Establishment History Panel 1975-2016, DOI: 10.5164/IAB.BHP7516.de.en.v1; IAB Establishment Panel (years 1993-2017), DOI: 10.5164/IAB.IABBP9317.de.en.v1; and Sample of Integrated Labour Market Biographies (years 1975-2014). Data access was provided via on-site use at the Research Data Centre (FDZ) of the German Federal Employment Agency (BA) at the Institute for Employment Research (IAB) and remote data access. The documentation for the data is contained in Ellguth and Kohaut (2014) and Schmucker et al. (2018).

It contains rich information about firms' personnel, investment, business policies, R&D and other areas of firm operations. While some variables are available in every year, others are only present for a subset of years. In particular, the digitization and automation measures used in this paper were included in 2016 and 2017 waves of the panel. The survey is merged with administrative personnel records, containing information about firms' workforce size and structure.

Establishment History Panel is a large firm-level data set containing information for a 50% random sample of all German establishments. It covers over 2 million establishments over the years 1976-2016. It contains yearly snapshots of firms' personnel structure and wage information, based on Social Security records. The data contains establishment location and industry code, but no firm-level data on technology adoption.

I supplement the IAB data with two independent country-industry-year data sources. The first is robot data from International Federation of Robotics, available for 1993-2016. The second is EU KLEMS database (Jäger, 2016), which contains information on employment and the stock of various types of capital, including that related to digital technologies: software & databases and ICT equipment for 1995-2015.

1.2.2. Measures of Digitization and Automation

I measure digitization and automation using several firm-level and industry-level measures. Various ways of measuring technology usage have advantages and disadvantages (for a related discussion, see McElheran, 2018). The benefits of those based on hard information, such as stock of robots, include the objective nature, natural quantitative interpretation and comparability across time and space. The drawbacks are focus on single technology class (i.e., in the robots example, inability to capture related automation technologies that do not meet the definition of robot) and quality and availability of measures.[8] The alternative method is based on surveys - an approach that is dominant in the existing literature on the adoption and impact of technology (Doms et al., 1997; Bresnahan et al., 2002; Lewis, 2011; Bloom et al., 2016). Surveys allow to overcome data availability issues and provide a measurement of technologies for which constructing an objective measure of usage would be hard (e.g. big data analytics), or which are hard to define and can take different form in different firms (e.g. Internet of Things). Main drawbacks are related to concerns about informativeness of the survey and limited ability to perform international and intertemporal comparisons.

In this paper I combine various measures of technology, trying to balance their strengths and weaknesses and provide more convincing and comprehensive results. In section 1.4 I analyze firm-level responses based on the survey data. In section 1.5 I use industry-level data. I apply these two types of measures in two very different specifications and show that both paint a similar picture of the technology.

Firm-Level Measures. Firm-level measures come from 2 waves of IAB-EP and combine subjective measures of intensity of technology adoption with binary indicators of technology usage. In 2016 the Panel contained an extra set of questions asking firms about "Automation and Digitization" technologies. The interviewer specified that these technologies include "autonomous robotics, smart factories, Internet of Things, big data analytics, cloud computing, online platforms, among other technologies". Respondents were asked to asses familiarity, potential and current

[8]Counting robots is not an ideal way to aggregate robots of different types and sizes. Data is also not widely available: Raj and Seamans (2019) discuss the lack of firm-level data on the usage of automation and AI. Recently, however, there are attempts to use existing customs data to measure robot usage. Fort et al. (2018) and ongoing project of Kwon and Zator (2019) use firm-level customs data on direct imports of robots, which has several advantages but also many limitations.

adoption of the technologies on a scale from 1 to 10 (with an option "Difficult to say"): A) how intensively has the establishment dealt with this topic so far? B) what potential is there for application of such technologies in the establishment? C) how well is the establishment equipped with these technologies compared to other establishments in the sector? My main measure of the technology is the answer to part C of the question, which measures the adoption. Importantly, the adoption measure is relative to other firms in the sector and hence should not capture differences in technologies across industries. Asking for relative assessment also makes it easier for respondents to give a meaningful answer by providing some reference point. Upper panel of Table 1.1 demonstrates that firms do not overshoot their assessment of adoption - the median response is 6 and the average response is 5.7. Figure 1.11 shows that this remains true across most industries. The median response is 6 in 8 out of 10 broad industry groups. In ICT and Finance the median is higher (this can be partially explained by IT firms identifying themselves with other sectors, e.g. IT firm providing solutions to car manufacturers may compare themselves to other automotive firms) and hence in my analysis I appropriately take into account industry fixed effects and focus on exploiting within-industry variation. To compare survey responses to industry-level technology measures I use answer to part A of the question, which is highly correlated with adoption but includes industry-wide differences in technology.

In 2017 the Panel did not contain the same questions but it did contain additional measures of digitization and automation. In particular, firms were asked whether they use different classes of technology. These technologies included 1) program controlled means of production requiring indirect handling by humans (e.g. industrial robots, CNC machines); 2) Software, algorithms and web interfaces for IT-based optimization of business processes (e.g. big data analytics, cloud computing); 3) networking and data exchange between facilities, processes or products (e.g. Smart Factory, drones, Internet of Things). Firms were also asked about how important each technology class is for them. I use 3 binary indicators of usage and technology being important (i.e. I define the measure as 1 if firm uses the technology and reports its importance to be at least 3 on a scale from 1 to 5) - robotics, data and networks - to measure the three respective technology classes. I will use these measures to analyze heterogeneous effect of different technologies.

Summary statistics of all measures are presented in Table 1.1. The exact wording of the questions is presented in the Online Appendix. The limitation of the data is that it is only available for a cross-section of firms, instead of being observed over time. These measures are used as the dependent variable in section 1.4 and are used to construct local adoption intensity that is one of the independent variables in section 1.5. Notice that in the survey firms were also asked about usage of other technologies that will not be the focus of this paper (e.g. usage of computers or mobile phones, which is positive for almost all firms).

Industry-Level Measures. Industry-level measures are stock of robots and stock of software and databases capital. They are available in the time-series and are used to construct technology measures that are independent variables in section 1.5. Robot density, i.e. number of robots per 1000 employees, is calculated by combining the count of robots from International Federation of Robotics with employment data from EU KLEMS. For Germany, robot density is available since 1993, the longest period among all countries. The Federation computes the robot count based on the information from robot suppliers that are responsible for over 90% of world industrial robot production. Robots are highly concentrated in manufacturing, although small count of robots is also reported for mining, construction and some other industries.

Figure 1.1. Firm-Level and Industry-Level Measures of Technology
The left panel presents the relationship between robot density at the industry level and the intensity of dealings with digitization and automation based on firm-level data. Each dot represents one of 2-digit industries (see Online Appendix for the list). Robot density is defined as the standardized logarithm of the count of robots per 1000 workers from IFR data. Robots are concentrated in manufacturing and only 14 industries have separately reported positive number of robots. For remaining industries robot density is calculated using "Other" category and is close to zero. Automation and digitization is the average of firms' responses to part A of the automation and digitization question from the IAB Establishment Panel (how intensively have you dealt with it so far, scale 1-10). The right panel presents the relationship between software & databases capital stock at the industry level and the intensity of dealings with digitization and automation based on firm-level data. Software and databases capital stock is the standardized natural logarithm of per-worker software & database capital stock from EU KLEMS.

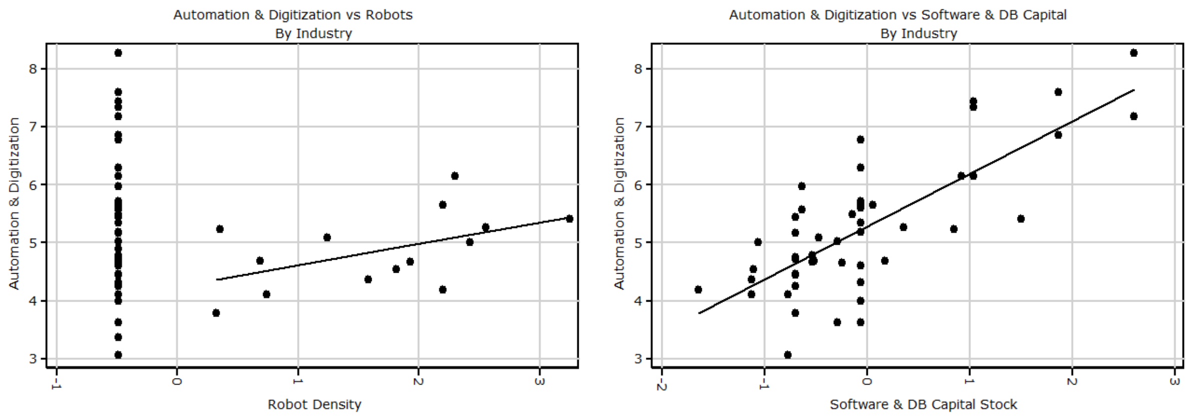

Stock of software and databases capital comes from EU KLEMS database. I combine the capital stock with employment levels, also from EU KLEMS, and calculate stock of software and databases capital per worker, expressed in thousands of Euros. I use the real value expressed in 2010 prices. Because data collection and publication practice changed for Germany in 2014, the capital series experiences an unexpected shift. To deal with this problem I use changes of capital stock between 2004 and 2014, instead of changes between 2005 and 2015 in section 1.5.

Validation of Measures. To cross-validate technology measures, Figure 1.1 presents the relationship of industry-level measures and survey-based measure of familiarity with automation and digitization (Online Appendix confirms these results in a tabular form). There is a positive and significant correlation between survey responses and industry-level measures of technology, which confirms that the survey does capture the information about technology in a meaningful way. Interestingly, while the correlation of survey measures with digitization is visible across all industries, the relationship with robots is driven only by industries within manufacturing. This is simply because almost no robots are being used outside of manufacturing.[9]

Bottom panel of Table 1.1 presents the relationship of the survey-based measure of automation and digitization adoption to other firm-level variables, including binary indicators of technology usage from 2017. There is a strong and positive correlations between overall technology adoption and probability of using each particular technology. Moreover, firms reporting higher adoption have higher investment (share in sales). They also report lower age of their equipment and the share of R&D workers in their personnel records is higher. The fact that self-reported measure of adoption is strongly correlated with hard information on investment and personnel structure again suggests that the adoption measure captures the real differences in technology across firms, as opposed to incorrect perceptions.

Descriptive Analysis of Digitization and Automation. Figure 1.3 shows how usage of various technologies varies across industries, based both on IAB Establishment Panel firm-level measures and on industry-level measures from International Federation of Robotics and EU KLEMS. Clearly, robots are highly concentrated in manufacturing. The strict and narrow definition of a robot used in IFR data implies that there is almost no robots in other sectors. The looser definition used in IAB-EP causes firms in other sectors to report some usage of

[9]Recent survey conducted by World Economic Forum (WEF, 2018) shows that relatively small share of technology-adopting firms expect to be using robots. Among 19 technology classes included in the survey, different types of robots occupy 4 out of 6 bottom spots when technologies are ranked according to the probability of adoption in near future. The list is opened by big data analytics, followed by app- and web-enabled markets, internet of things, machine learning and cloud computing. This primacy of digital technologies highlights the importance of looking at broader of set of technologies rather than focusing on industrial robots.

Table 1.1. Digitization and Automation: Summary Statistics and Relation to Other Variables

Panel A: Summary Statistics

Variable		Mean	Std Dev	P25	Median	P75	Num Obs
2016 Digitization	A (familiarity)	4.89	3.02	2	5	8	14036
and Automation	C (adoption)	5.72	2.68	4	6	8	10255
2017 Robots		0.154	0.361	0	0	0	11577
2017 Digitization	(Data)	0.525	0.499	0	1	1	11577
2017 Digitization	(Networks)	0.127	0.332	0	0	0	11577

Panel B: Relation to Other Variables

	Y = Adoption of Digitization and Automation					
	(1)	(2)	(3)	(4)	(5)	(6)
Robots (2017)	0.86***					
	(0.089)					
Digitization: Data (2017)		1.16***				
		(0.057)				
Digitization: Networks (2017)			1.15***			
			(0.084)			
Investment (% sales)				2.39***		
				(0.383)		
Age of Equipment					-1.08***	
					(0.032)	
Share of R&D Workers						1.28***
						(0.419)
N	8407	8407	8407	10255	10255	10255
Industry FE	✓	✓	✓	✓	✓	✓

Top panel shows summary statistics for the firm-level measures of technology. Summary statistics for other variables are presented in Table 1.12. In the bottom panel regressions, the dependent variable is adoption of digitization and automation from the IAB Establishment Panel (wave 2016, part C). Independent variables are binary indicators of usage of different technology classes coming from 2017 wave of the IAB-EP; share of gross investment in sales; firm assessment of their equipment age; and share of R&D workers in total employment. Industry fixed effects are included as a control variable. (*) denotes significance at 10% level, (**) at 5% level and (***) at 1% level.

robots[10], but their prevalence is low. Based on IFR data we can see that usage of robots in Germany is highly correlated with usage of robots in other countries, although in some

[10]Robots in IAB-EP are broadly defined as "program controlled means of production requiring indirect human intervention"; establishments outside of manufacturing reporting some usage of robots defined in such a way include e.g. airport services firms.

important industries - such as automotive or chemical manufacturing - Germany seems to have higher robotization rates. Digital technologies are more popular and more homogeneously distributed across industries. Technologies related to IT enhancements of business processes, such as big data analytics or cloud computing, are prevalent across many industries and their rates of usage is high. Technologies related to networks and communications between machines (e.g. Internet of Things) are less prevalent and tend to be concentrated in ICT and manufacturing sectors, although the contrast is not as stark as in the case of robots. In the EU KLEMS data software and databases capital is highest in sectors like ICT and finance. The stock of capital is lower in Germany than in other countries, but this is partially due to different reporting methods causing a level shift – when we examine changes in the capital stock between 2005 and 2015 (which are used in the empirical analysis) the values are more similar (see Figure 1.13 in the Appendix).

Usage of technology varies by geographical area, even after controlling for the industry structure. Figure 1.4 shows the intensity of digitization and automation adoption across Germany. The measure is an average of firm assessments of adoption relative to other firms in the sector and hence it is not driven by industry composition, but rather by other factors affecting technology adoption, such as proximity of R&D centers, technological spillovers and labor market conditions.

Usage of technologies also varies by firm characteristics, see Fig 1.5. Large firms are more likely to report using robots and digital technologies. Firms with high levels of adoption are also more productive, which can be partially explained by the fact that their workforce is more skilled. Interestingly, high-adopters have also been growing faster in the last 5 years. Hence, the naive firm-level regression does not support the hypothesis that technology reduces

Figure 1.2. Digitization and Automation Usage by Industry Group - Robots and Digitization based on IAB Establishment Panel

The graph shows the frequency with which technologies are used, based on firm-level responses from 2017 IAB Establishment Panel. On the left, the share of firms declaring use of "means of production requiring indirect human intervention" (robots, CNC machines) is shown. On the right, the graph presents the share of firms that use digital technologies related to IT-based optimization of business processes (big data, cloud computing) and networking and data exchange between facilities or processes. For data confidentiality reasons the exact value of robot usage for ICT and Finance was censored - the value is below 0.1 and the Figure shows it as 0.05.

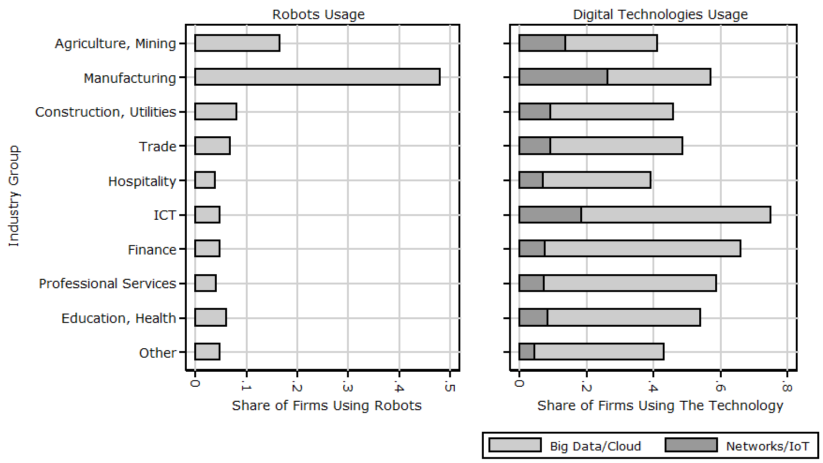

employment. Clearly, however, such a conclusion is premature because technology adoption is endogenous and faster growing firms may be more likely to adopt new technologies.

Figures 1.14 and 1.15 in the Appendix document the relationship of technology adoption with wages, innovation, being an exporter and other firm characteristics. High-adopters pay higher wages, are more innovative, are more likely to be exporters and more likely to be a part of multi-establishment group. However, there is no relationship between technology adoption and foreign ownership.

Figure 1.3. Digitization and Automation Usage by Industry Group - Robots and Digitization based on IFR and EU KLEMS data

The top graph shows the number of robots per 1000 workers in Germany and 6 other European countries by industry, based on data from International Federation of Robotics. The bottom graph shows the stock of software and databases capital, in thousands of Euros per worker, based on EU KLEMS data. For both measures, the data comes from 2015.

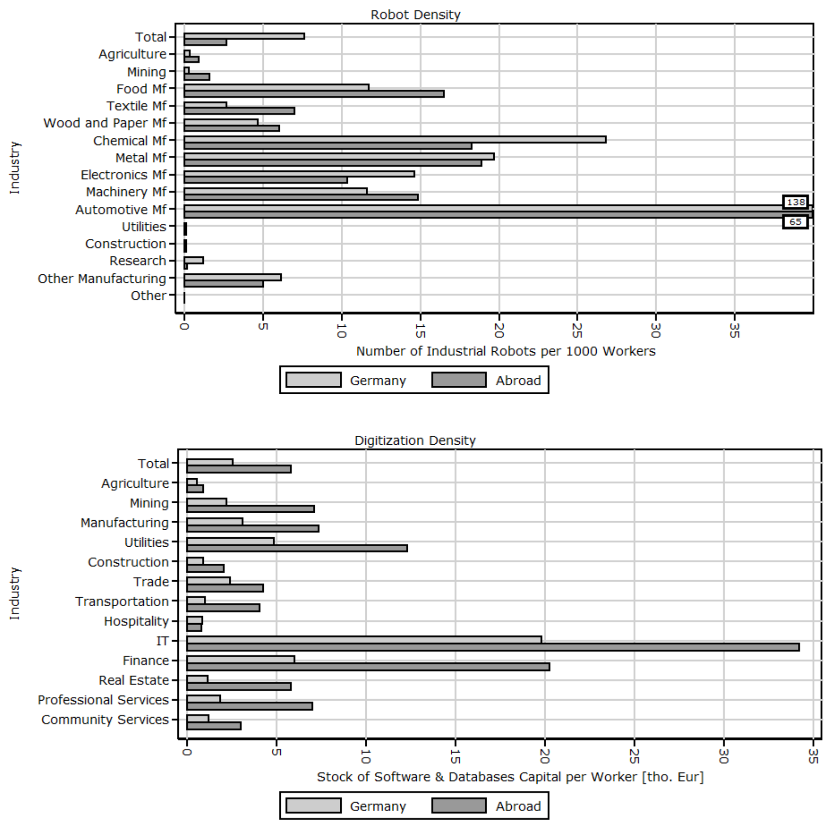

1.2.3. Measures of Labor Scarcity and Employment

Both part 1.4 and part 1.5 of the paper analyze the relationship of technology and labor. This section provides a description and basic analysis of the variables used to measure labor market conditions and employment evolution.

The first part of the analysis, section 1.4, studies the impact of labor scarcity on the investment in digitization and automation. I use 4 labor scarcity measures that come from IAB Establishment Panel. Unlike the technology measures, they are available in multiple, although not all, waves

Figure 1.4. Geographic Distribution of Digitization and Automation Adoption
The map presents values of digitization and automation adoption index from 2016 IAB Establishment Panel. The original index is computed at the district level but for data confidentiality reasons presented values were computed on the spatial planning regions (RORs) level – each ROR contains ca. 4 districts. Moreover, some values cannot be shown due to data provider restrictions. The index is the average of firms' responses to a question about the intensity of digitization and automation adoption from 2016 wave of IAB Establishment Panel, but industry-level averages are subtracted and economy-wide average is added (so that the local index does not depend on industry composition). The response are on scale (1,10) and ROR-level averages vary between 3.41 and 7.04.

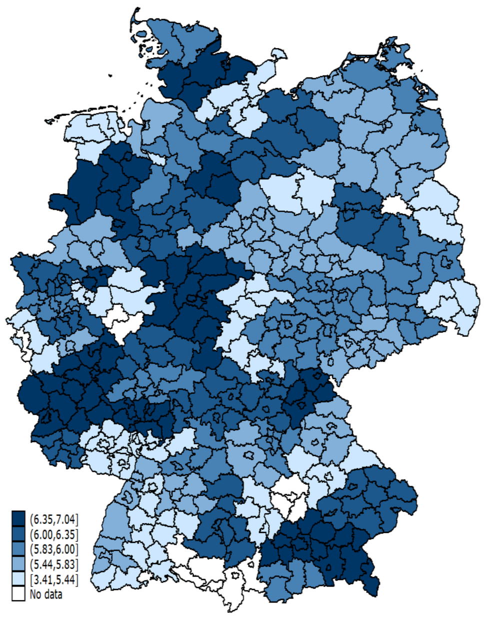

Figure 1.5. Digitization and Automation Usage and Firm Characteristics
The left panel shows how firm size (measured by count of employees) and firm growth (measured by 5-year relative change in number of employees) varies with different levels of digitization and automation adoption. High-adopting firms are larger and typically have been growing faster. Right panel shows how labor productivity (measured by sales per worker) and skill level of the workforce (measured by share of skilled workers, defined based on skilled/unskilled assignment of 12 occupational groups) varies with adoption intensity. High-adoption firms are more productive and have higher share of skilled workers.

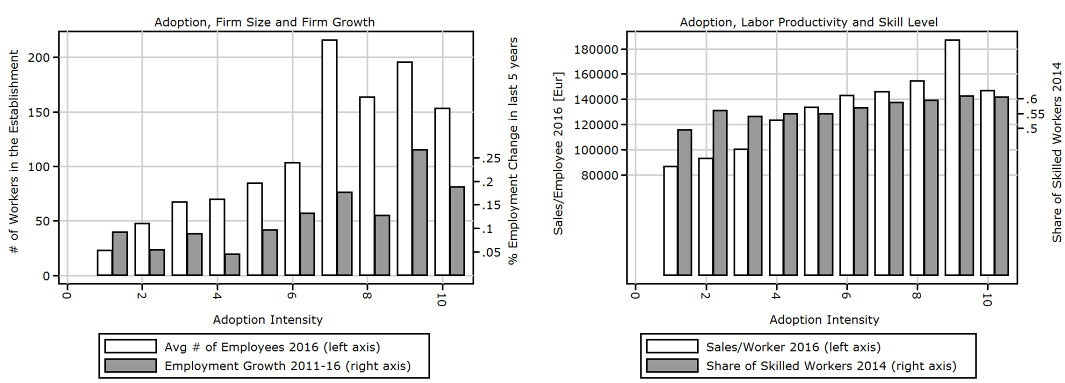

of the panel. My main measure is firm's binary declaration that they have "difficulties in finding required workers on the labor market", which is a part of "Staffing Problems" module of the Panel. The second measure is a binary indicator that the firm would like to recruit additional staff, on top of the staff they actually recruited, and is based on the "Recruitment" module. The third measure captures labor-driven capacity constraints, i.e. firm's declaration that they are unable to increase the production without hiring new staff. The fourth measure, only available for a subset of firms, is a binary indicator of abandoning a project because of personnel shortage. All firms are asked whether they had an innovative project that they planned to carry out but did not. If the answer is positive, firms are asked about different possible reasons for abandoning the project, including personnel shortage.

Top panel of Table 1.3 shows summary statistics of labor scarcity measures. In the basic specification, I use labor scarcity measure from 2014 but additional tests exploit values from years 2008-2016. In 2014, around 40% of German firms declared that they have difficulties in finding workers. This average hides significant variation across industries and regions. Figure 1.6 shows that the share of firms with labor scarcity problem is lower than 35% in the bottom quintile of district and higher than 50% in the top quintile. Bottom panel of Table 1.3 shows that labor scarcity is strongly and negatively related to district's unemployment rate. Labor scarcity perception of a firm is also positively related to scarcity faced by other firms in the same industry, as well as to the firm's productivity, employment growth and wages.

The second part of the paper, section 1.5, studies the impact of technology on employment change in the last 10 years. Employment change is computed at the area-industyr level by aggregating employment records of 50% random sample of all German firms. As bottom rows of Table 1.3, Panel A, illustrate, average cell has over 3,000 workers. The size of German

Table 1.2. Labor Scarcity and Employment Measures: Summary Statistics

VARIABLE	MEAN	STD DEV	P25	MEDIAN	P75	NUM OBS
LABOR SCARCITY MEASURES						
Hard to Find Workers	0.40	0.49	0	0	1	10391
Would Like to Hire More Workers	0.19	0.39	0	0	0	10365
Can't Produce More w/o Hiring Extra Labor	0.41	0.49	0	0	1	8777
Project Abandoned – Can't Find Workers	0.21	0.39	0	0	0	1812
Hard to Find Workers - Agricul., Manufacturing	0.46	0.50	0	0	1	2619
Hard to Find Workers - Construction, Trade	0.37	0.48	0	0	1	3461
Hard to Find Workers - Prof. Services	0.41	0.49	0	0	1	3776
Hard to Find Workers (Area Index)	0.40	0.11	0.35	0.39	0.47	14202
EMPLOYMENT MEASURES						
Employment	3154	5615	283	1276	3503	5703
Δ%Employment (2005-2015)	33.7	189.8	-6.2	14.2	43.4	5592

workforce increased between 2005 and 2015, with median area-industry cell increasing employment by 14%, but there is substantial variation in employment changes across industries and local areas. Figure 1.16 in the Appendix shows employment changes by industry for Germany and other European countries, based on EU KLEMS data.

1.3. Theoretical Framework

To guide the empirical analysis, I introduce a model of firm's decisions regarding technology adoption and inputs choice in the face of labor and capital costs/constraints. I model firm's adoption decision in the spirit of Davies (1979): the firm decides whether or not to adopt the new technology by comparing benefits and costs of adoption. The benefits depend on the cost of labor and capital, but the shape of this dependence is different for different types of technology. The firm production function I employ is similar to task-based model of Acemoglu and Autor (2011) and related to pioneering work of Zeira (1998). Similar approach was recently employed by Acemoglu and Restrepo (2018b) to model automation. This modeling choice is common for me and them but the focus is different – while they propose an endogenous growth model and

Table 1.3. Labor Scarcity and Employment Measures: Relation to Other Variables

	Y = Hard to Find Workers					
	(1)	(2)	(3)	(4)	(5)	(6)
Local Unemployment	-0.0058***					
	(0.0021)					
Hard to Find Workers		0.912***				
(Industry Index)		(0.020)				
Labor Productivity			0.0001***			
			(0.00003)			
Total Employment				0.0062		
				(0.0049)		
Employment Growth (5Y)					0.034***	
					(0.008)	
Avg Wage						0.0013***
						(0.0002)
N	9616	10389	6771	10390	6808	7968
Industry FE	✓		✓	✓	✓	✓
Area FE		✓				

Top panel shows summary statistics for the firm-level measures of labor scarcity and area-industry-level measures of employment changes. Labor scarcity measures come from 2014 wave of the IAB-EP, except for the last measure – Project Abandonded Because of Lack of Suitable Personnel – which is an average from waves 2009, 2011, 2013 and 2015. “Agricul., Manufacturing” refers to all industries in agriculture, mining and manufacturing, i.e. 1-31 industry codes. “Construction, Trade” refers to industry codes 35-57, which includes construction, utilities, trade, hospitality, transport. “Prof. Services” refers to industries with codes (58-90), which includes IT, finance, professional services, health, and education. Employment and its percentage change is computed for industry-area cells based on the Establishment History Panel data. In the bottom panel regressions, the dependent variable is firm-level measure of difficulties in finding workers. Local unemployment rate comes from German Statistical Office and varies by district. Industry index is an average of labor scarcity declarations of all firms in the same industry, excluding firm’s own declaration. Labor productivity is sales per worker in 2015 measured in EUR/worker. Standard errors are clusterd on the area level in column 1 and industry level in column 2. (***) denotes significance at 1% level.

analyze automation in general equilibrium, I consider a firm decision problem where prices are treated as given and highlight not only the automation process, but also the part of technological progress that complements labor.

A firm produces output by combining a continuum of tasks in the unit interval [0,1] in the Cobb-Douglas form:[11]

[11] Cobb-Douglas form is chosen to simplify the exposition. Allowing for different elasticity of substitution between tasks does not change the main intuition of the model, formalized in Proposition 1.

Figure 1.6. Geographic Distribution of Labor Scarcity
The map presents values of labor scarcity index. The original index is computed at the district level but for data confidentiality reasons presented values were computed on the spatial planning regions (RORs) level – each ROR contains ca. 4 districts. Moreover, some values cannot be shown due to data provider restrictions. The index is the average of firms' responses to a question about the difficulties in finding workers from 2014 wave of IAB Establishment Panel. The response are binary and district-level fraction of firms that respond positively (i.e. who say that they have difficulties finding workers) varies between 0.15 and 0.65.

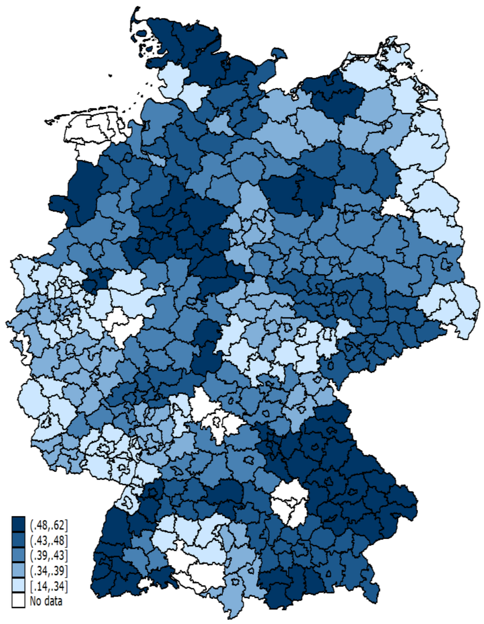

$$Y = A \cdot exp\left[\int_0^1 a \cdot lny(i)di\right] \tag{1.1}$$

where $a < 1$, i.e. the firm faces decreasing returns to scale and sells its product on a competitive market with price of output equal to 1. Alternatively, Y can be interpreted as firm's revenue and a as a way of capturing the downward-sloping demand function for firm's product. A is a firm-specific productivity parameter. Different tasks represent different parts of the production process. For example, car manufacturer needs to perform welding, painting, design, marketing etc. to produce and sell a car.

Each task i can be produced by capital/machines or by labor and has the following production function:

$$y(i) = \alpha_K(i)k(i) + \alpha_L(i)l(i) \tag{1.2}$$

Vector $\alpha(i) = [\alpha_K(i), \alpha_L(i)]$ represents technology. For each i it determines the productivity of capital and labor at task i. Terms $k(i)$ and $l(i)$ represent the amounts of capital and labor assigned to the production of task i (chosen by the firm). For simplicity, I consider only one type of labor but the Online Appendix contains an extension with two labor types.

The firm can hire labor paying cost w and rent capital at rate r, which are exogenously given. For simplicity, input markets are perfectly competitive in the model. High labor cost represents a tight local labor market – which in reality may manifests itself not only through high wages but also through high search costs and difficulties in finding suitable workers. High cost of capital in the model may represent also financial constraints, i.e. difficulties in obtaining credit.

The firm maximizes profits $Y - rK - wL$ by choosing inputs K, L and assigning them to tasks (for each i choosing $k(i), l(i)$). Because capital and labor are perfect substitutes within a single task, whenever $\alpha_K(i)/r$ (cost-adjusted capital productivity) is higher than $\alpha_L(i)/w$, task i will be performed by capital (automated). Figure 1.7 demonstrates the determination of the automation threshold R under an arbitrary new technology - tasks in the interval $[0, R]$ are performed by capital, while remaining tasks are performed by labor.[12] The Figure shows that

[12]For simplicity I assume that tasks are ordered in such a way that α_K is decreasing in i and α_L is weakly increasing; this assumption is inconsequential since the symmetry of all tasks implies that only the mass of automated tasks matters, not their index.

Figure 1.7. Productivity of Labor and Capital Across Tasks Under New and Old Technology
The graph illustrates mechanisms behind the production function in the theoretical model. It shows cost-adjusted productivity of labor and capital for different tasks for an example of technology parameters before and after technological change. For simplicity, it is assumed that the initial technology has machines that are not productive ($\alpha_K(i) \equiv 0$) and labor is uniformly productive in all tasks ($\alpha_L = 1$ and, also for simplicity, $w = 1$, and hence $\frac{\alpha_L(i)}{w} \equiv 1$). For new technology, labor productivity schedule is assumed to be weakly increasing in the task index i and capital productivity is decreasing in the task index. Under old technology, all tasks are performed with labor and the quantity of labor demanded is determined by productivity that is equal to 1. Under new technology, point R marks the limit of capital-labor substitution. For tasks in the interval $[0,R]$, $\alpha_K(i)/r$ (cost-adjusted capital productivity) is higher than $\alpha_L(i)/w$ and hence they are performed by capital. For the remaining tasks, $\alpha_K(i)/r$ is lower than $\alpha_L(i)/w$ and hence the tasks are performed by labor. Quantity of capital and labor is determined by the productivity parameter θ which is an average of the upper-envelope of new productivity curves.

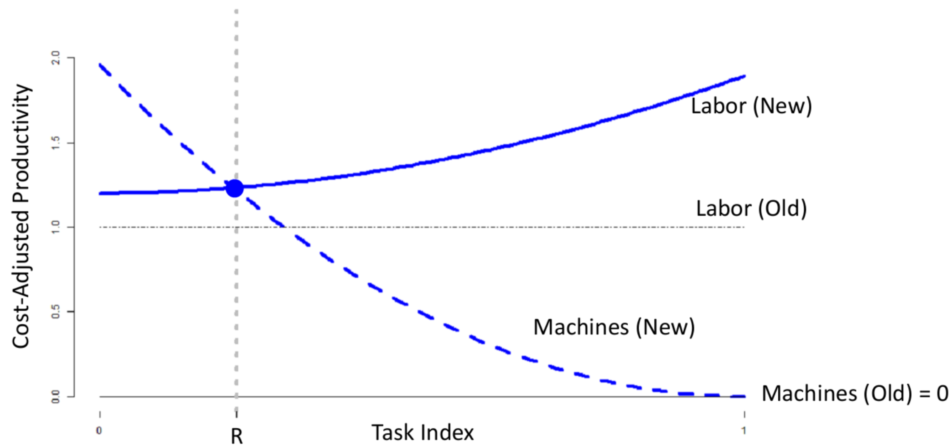

the decision whether capital or labor performs a given task depends on the interplay of costs and productivities of labor and capital.

What is the total labor demand of the firm? Using features of Cobb-Douglas production function one can derive[13] the expression for firm's profit:

---

[13]The firm maximizes profits:

$$\max_{l,k,R} Aexp\left\{\int_0^R a \cdot ln[\alpha_K(i)k(i)]di + \int_R^1 a \cdot ln[\alpha_L(i)l(i)]di\right\} - w\int_R^1 l(i)di - r\int_0^R k(i)di$$

Cobb-Douglas production function with symmetric tasks is characterized by equal expenditures on each task: $p(i)y(i) = p(i')y(i')$, where $p(i)$ denotes the price of one unit of task i (e.g. one widget). For two tasks produced by labor we therefore have $p(i)\alpha_L(i)l(i) = p(i')\alpha_L(i')l(i')$. At the same time, the (hourly) wage of workers employed in each tasks is the same, and hence $p(i) = w/\alpha_L(i)$. This implies that $l(i) = l(i')$ for each 2 tasks performed by labor and similarly $k(i) = k(i')$ for each 2 tasks performed by capital. Moreover, for each task performed by labor and each task performed by capital we have $l(i) = k(i')\frac{r}{w}$. We can therefore write $L = l(1-R)$

$$\Pi = A\left(\frac{L}{1-R}\right)^{a} exp\left\{\int_0^R ln[\alpha_K(i)\frac{w}{r}]di + \int_R^1 ln[\alpha_L(i)]di\right\}^{a} - w\frac{L}{1-R} \tag{1.3}$$

Let us denote $\theta = exp(\int_0^R ln[\alpha_K(i)\frac{w}{r}]di + \int_R^1 ln[\alpha_L(i)]di)$ (notice that θ captures overall productivity). Profit maximization implies that total labor demand equals:

$$L = (1-R)\left(\frac{Aa\theta^{a}}{w}\right)^{\frac{1}{1-a}} \tag{1.4}$$

Different Effects of Technological Change. Improvements in technology can be viewed as changes to the productivity schedule $\alpha(i) = [\alpha_K(i), \alpha_L(i)]$. These changes, combined with prices of capital and labor, determine the value of R and θ, which govern the substitution and productivity effects of technology. Increases in α_K generally increase both R and θ. Intuitively, an increase in productivity of machines increases the set of tasks that are automated and at the same time rises overall productivity. Notice, however, that the increase in both parameters does not need to be strict. If α_K increases for tasks $i < R$, that is those which are already automated, there is no change in the automation threshold R, but there is an increase in the productivity θ. If α_K increases for tasks $i \gg R$ but the increase is small, so that $\frac{\alpha_L}{w} > \frac{\alpha_K}{r}$ still holds for all $i > R$, there is no impact on R nor on θ.

Increases in α_L generally reduce automation threshold R but increase productivity θ. While empirically less relevant, an increase in labor productivity for $i < R$ may lead to deautomatization of some tasks, i.e. decrease in R. More plausibly, however, α_L increases for tasks $i > R$, rising θ. This effect can be thought of labor-augmenting technological change.

Overall, technological change is a combination of: 1) automation, i.e. changes in R, coming from changes in the relative productivity of capital, α_K, and labor, α_L; 2) labor-augmenting technological change, coming from increases in α_L for $i > R$; 3) productivity effect of capital improvements, i.e. increases in α_K that increase θ (possibly without influencing R). The intensity of these 3 channels can vary across industries and technology classes and can lead to different technology adoption patterns and different employment effects.

Adoption and Labor Scarcity. Consider now the firm's decision whether or not to adopt a new technology. For simplicity, let us assume that current technology is such that machines are not productive ($\alpha_K(i) \equiv 0$) and productivity of labor does not vary across tasks (productivity is normalized to 1, i.e. $\alpha_L(i) \equiv 1$; notice that Eq. 1.4 then implies that labor demand under current technology equals $\left(\frac{Aa}{w}\right)^{\frac{1}{1-a}}$). The firm is contemplating adopting new technology with arbitrary characteristics $\alpha' = [\alpha'_L, \alpha'_K]$. While this is not necessary in the model, in empirically relevant scenario $\alpha' > \alpha$, i.e. the new technology is unambiguously better. However, the adoption has a fixed cost $C(r)$. The firm decides whether or not they want to pay it and produce using α' or stick with the old technology and produce using α. They will adopt the new technology if and only if the cost of adoption $C(r)$ is lower than an increase in profits that can be attained, i.e. will maximize:

$$\Delta\Pi = \rho[\Pi_1 - \Pi_0 - C(r)] \tag{1.5}$$

where $\rho \in \{0, 1\}$ represents the decision to adopt. How does labor scarcity, represented by high w, affect the adoption? The adoption would be increasing in labor scarcity if and only if the above function was supermodular in adoption and wages. Under appropriate differentiability assumptions, supermodularity corresponds to:

$$(1.6) \qquad \frac{d^2(\Delta\Pi)}{d\rho dw} \geq 0$$

It can be shown[14] that this is the case if and only if:

$$(1.7) \qquad (1-R)\cdot\theta^{\frac{a}{1-a}} - 1 < 0$$

Intuitively, the adoption of the new technology is increasing in labor scarcity if the technology mostly substitutes for labor. If the new technology reduces the demand for labor, firms that face highest labor cost (which could possibly come from high search costs) are most likely to adopt. Conversely, if the technology mostly complements labor, i.e. labor productivity increases and it is optimal to hire more workers after the technology is adopted, firms that face the lowest labor cost are most likely to adopt.

Employment Change. How would the quantity of labor demanded change if the firm adopted the new technology? We can simply compare labor demand (given by Eq. 1.4) under new technology with arbitrary parameters $\alpha' = [\alpha'_L, \alpha'_K]$ to labor demand under the current technology

---

14 $\frac{d^2(\Delta\Pi)}{d\theta dw} \geq 0$ is true if and only if $\frac{d(\Pi_1-\Pi_0)}{dw} \geq 0$. Substituting expressions for optimal labor demand from Eq. 1.4, $\Pi_1 - \Pi_0$ can be expressed as:

$$\Pi_1 - \Pi_0 = A\left(\frac{Aa\theta^a}{w}\right)^{\frac{a}{1-a}}\theta^a - w\left(\frac{Aa\theta^a}{w}\right)^{\frac{1}{1-a}} - A\left(\frac{Aa}{w}\right)^{\frac{a}{1-a}} + w\left(\frac{Aa}{w}\right)^{\frac{1}{1-a}} = w^{\frac{-a}{1-a}}\underbrace{A^{\frac{1}{1-a}}}_{h_1>0}\underbrace{[a^{\frac{a}{1-a}} - a^{\frac{1}{1-a}}]}_{h_2>0}[\theta^{\frac{a}{1-a}} - 1]$$

The derivative of this expression with respect to wage equals (notice that $\frac{d\theta}{dw} = \theta\frac{R}{w}$):

$$\frac{d(\Pi_1-\Pi_0)}{dw} = h_1h_2[-\frac{a}{1-a}w^{\frac{-1}{1-a}}(\theta^{\frac{a}{1-a}} - 1) + w^{\frac{-a}{1-a}}\frac{a}{1-a}\theta^{\frac{2a-1}{1-a}}\frac{R\theta}{w}] = -H\left[(1-R)\theta^{\frac{a}{1-a}} - 1\right]$$

where $H = -h_1h_2\frac{-a}{1-a}w^{\frac{-1}{1-a}} > 0$. Therefore, technology adoption is increasing in labor scarcity if and only if $(1-R)\cdot\theta^{\frac{a}{1-a}} - 1 < 0$

with parameters $\alpha_K(i) \equiv 0$ and $\alpha_L(i) \equiv 1$ (continuing to assume it for simplicity). The change in employment is given by:

$$\Delta\%L = (1-R) \cdot \theta^{\frac{a}{1-a}} - 1 \qquad (1.8)$$

The adjustment of employment after the adoption can be decomposed into two parts. First, when productivity of machines (α'_K) is large relatively to productivity of labor (α'_L), some tasks are more efficiently performed by machines and thus get automated (R increases). This is the substitution effect, which reduces the demand for labor. Second, when the productivity of machines or labor increases, the firm experiences increase in the average productivity θ and increases its production and the demand for inputs. This is the productivity effect, which increases the demand for labor. Which of these two effects dominates is an empirical question.

Connecting Employment Change and Labor Scarcity Effect. Suppose that we are interested in learning about unknown characteristics of a new technology and in particular about the extent to which it substitutes labor versus complements it by increasing the productivity. Suppose also that we observe both firms' adoption decisions, as well as information about their employment change and about labor costs they face. Then because both condition 1.7 and equation 1.8 contain the same expression, we can learn about the characteristics of the new technology in two ways: by analyzing how labor scarcity affects the adoption and by studying how the adoption affects employment change.

Proposition 1. *Technology adoption is higher when firm faces labor scarcity if and only if total employment decreases after technology is adopted. This is the case when substitution effect dominates over productivity effect, i.e. when:*

$$\underbrace{(1-R)}_{Substitution\ Effect} \cdot \underbrace{exp\{\int_0^R ln[\alpha_K(i)\frac{w}{r}]di + \int_R^1 ln[\alpha_L(i)]di\}^{\frac{a}{1-a}}}_{Productivity\ Effect} - 1 < 0$$

This proposition provides a motivation for the empirical analysis in which I analyze both how labor scarcity affects technology adoption as well as how technology affects employment. Both parts shed light on the same, unknown technology characteristics that generate substitution and productivity effects.

Online Appendix presents additional theoretical results. I formalize an intuitive result that financial constraints impede the adoption of technology if and only if there is any technological progress embodied in machines or if adoption itself requires capital expenditures. I also present an extension of the model in which I consider different worker groups, corresponding e.g. to different skill groups. I show that the change in the share of workers of each type depends both on the changes in labor productivity as well as the degree to which tasks performed by different workers get automated.

1.4. Labor Scarcity and Technology Adoption

This section analyzes how availability of labor influences firm's investment in digitization and automation. As formalized in Proposition 1, the shape of this relationship depends on the features of the technology. If the technology purely substitutes existing workers, then scarcity of labor should increase the investment. If technology is purely complementing existing workers, then scarcity of labor should have the opposite effect.[15] I regress technology adoption on labor scarcity and demonstrate that the average effect is positive, consistent with substitution

[15] Both effects can clearly coexist even within a single industry. In retail, for example, introduction of self-checkout is a clear example of labor-substituting technology. Smartphone applications collecting customer shopping patterns, on the other hand, complement the work provided by data analysts and marketing specialists.

effect. I use various approaches, including instrumenting labor scarcity with aging, to address endogeneity concerns. Finally, I perform heterogeneity analysis which reveals that effect of labor scarcity significantly varies by industry, worker type and technology class.

1.4.1. Basic Results

The analysis in this section uses a cross-section of firms from IAB Establishment Panel. Basic specification is the OLS regression:

$$Technology_i = \beta \cdot Labor\,Scarcity_i + \gamma \cdot I_j + \phi \cdot Z_i + e_i \tag{1.9}$$

The dependent variable is a measure of digitization and automation that comes from the IAB Establishment Panel. The main measure is the firm assessment of intensity of adoption: a continuous variable varying from 1 to 10. I also employ alternative measures, such as binary indicator of adoption being above median and interaction of this indicator with investment share in sales being above median. I_j denotes set of industry fixed effects that correspond to 2-digit classification based on NACE Rev. 2 (see Online Appendix for a complete list of industries). Z_i contains set of firm-level controls. In the main specification I control for firm size (measured as total employment, but robust to using total sales). Several additional controls which do not substantially influence the main coefficients of interest, such as profitability, establishment age, past employment growth, type of management, international ownership, being part of a group or being a public firm are included in the robustness checks. Due to limited data availability, including additional controls significantly reduces sample size. I choose the main specification to be parsimonious but present results that demonstrate that including additional controls does not meaningfully change the magnitudes of the coefficients. Another potentially important control variable is the area fixed effect, since part of the variation in labor scarcity is common for all firms in the area. I present the results both with and without area fixed effects. Standard errors in the main specification are clustered on the industry level.

The main independent variable, $LaborScarcity_i$, is defined in four different ways described in section 1.2.3. The main measure is firm's declaration that they have difficulties finding workers. It is defined based on answers to 2014 survey that predates the technology-adoption measure by 2 years. Lagging the independent variable is the first attempt to circumvent the reverse causality problem but the results would remain similar if I used measures from 2016 (see Table 1.14).

Table 1.4 presents the estimates of equation 1.9. Columns 1-7 present the results with adoption measure being a continuous assessment of adoption from the survey; columns 8-12 present the results for adoption measure that interacts above-median adoption assessment with above-median capital expenditures (share in total sales). Each measure points to a clear positive relationship between technology adoption and labor scarcity: the harder it is to find workers, the higher is the level of technology adoption. This result suggests that on average across firms, marginal worker is substituted, rather than complemented by the new technology. The magnitudes suggested by different measures are similar: changing labor constraints measure from 0 to 1 increases technology adoption by 10-15% of standard deviation. Inclusion of area fixed effects, in columns 4 and 12, slightly decreases the magnitude of the effect, consistent with part of the variation in labor scarcity being driven by area-level characteristics. However, even when those characteristics are purged off, significant variation across firms remains and is positively related to technology adoption.

Table 1.4. OLS Regression of Digitization and Automation Adoption on Labor Scarcity

	Basic Specification				Alternative Measures of Labor Scarcity			Alternative Measure of Technology				
	Y = Digitization and Automation Adoption							*Y = High Adoption X High Investment*				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Hard to Find Workers	0.307***	0.285***	0.372***	0.260***				0.066***				0.065***
	(0.072)	(0.076)	(0.105)	(0.070)				(0.011)				(0.013)
Demand for Hiring					0.180**				0.085***			
> Hired					(0.083)				(0.015)			
Can't Increase Sales						0.361***				0.059***		
without New Staff						(0.083)				(0.016)		
Investment Prevented							0.381**				0.077**	
By Lack of Personnel							(0.159)				(0.034)	
N	7469	5855	3479	7469	7449	6260	1431	5781	5771	5615	1191	5781
Industry FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Size	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Profits & Growth		✓	✓									
Firm Type			✓									
Area FE				✓								✓

Dependent variable in columns 1-7 is the digitization and automation adoption measure from IAB Establishment Panel. Dependent variable in columns 8-12 is the binary indicator for the adoption measure being above median and the investment to sales ratio (average from 2011-2016) being above industry-wide median. Independent variables are various measures of labor scarcity from IAB Establishment Panel: "Hard to Find Workers" is a binary variable that takes value 1 when establishment confirmed that they face this staffing problem; "Demand for Hiring > Hired" is the dummy variable taking value 1 if establishment declared that they would like to hire more workers than they did hire; "Can't Increase Sales Without New Staff" takes value 1 when establishment declares that they are capacity constrained and would have to hire new staff to increase production; "Investment Prevented by Lack of Personnel" is defined only for a subset of firms that declare that they have abandoned an investment in product or process innovation (and thus smaller sample size). It takes value 1 if the establishment declares "lack of qualified personnel" as a reason for abandoning the project. Control variables include 5 binary variables for different size quintiles (by employment) and industry fixed effects (2-digit). Additional controls under "Profits & Growth" include categorical measure of firm's profitability, employment growth in the last 3 years and establishment age. Further controls under "Firm type" include categorical measures of type of management, being part of a group, establishment age, public ownership, foreign ownership and being a startup. Standard errors are clustered on the industry level. (**) denotes significance at 5% level and (***) at 1% level.

Table 1.13 in the Appendix shows that the estimate remains similar after including additional controls such as profitability, past employment growth, establishment age and management and ownership characteristics. Basic labor scarcity measures are defined based on data from 2014, but table 1.14 shows the results for measures from different periods. Overall, the relationship seems to be somewhat persistent, but it is no longer significant if the measures comes from years earlier than 2010. Table 1.16 presents the results for other staffing problem variables, which may be thought of as placebo checks. There is no relationship between technology adoption and firms declarations about problems with worker motivation or with having too many employees. Interestingly, there is also no relationship with an indicator that takes value 1 if a firm reports high labor costs as one of their staffing problems. This might be because of labor market institutions in Germany make wages rigid. There exists a positive relationship between adoption and the demand for further training. While there are many ways of interpreting this relationship, one possibility is that firms that have difficulties finding suitable workers are also forced to hire employees that need to be intensely trained.

One weakness of my digitization and automation measure is that it is only available in one 2-year time period (I treat 2016 and 2017 jointly and do not analyze the time-series variation between these 2 years given the short period and technology measures not being directly comparable) and hence it measures the stock of technology as opposed to its change. While I do not observe the digitization and automation adoption variable in the past, and thus constructing the natural measure of change is not possible, Table 1.15 presents the results of a specification that attempts to approximate the analysis of technology changes. Using the information about computer equipment from 2001-07 I compute firms' technological sophistication in the past. I then use it as a control variable and to create a synthetic measure of changes in technology.

Both approaches confirm that digitization and automation adoption in 2016/17 is accelerated by labor scarcity.

1.4.2. Addressing Endogeneity Concerns

Table 1.4 demonstrates that the positive relationship between labor scarcity and technology adoption is robust to using various specifications but the results are subject to important endogeneity concerns. First, there is a concern about reverse causality: a firm that adopted new, sophisticated technology may have troubles finding workers because skills required to operate the technology are scarce. This concern is to some extent mitigated by using lagged measures of labor scarcity but it nonetheless remains valid. Second, there is a concern about omitted variables: firms that adopt new technology more intensely may be different in a way that is unobserved. Typically we would expect such firms to be more productive and successful than other firms. If such firms are more attractive to workers and hence have less difficulties recruiting, OLS coefficient may be downward biased. But those firms may also have higher demand for their products, which can be accommodate both by hiring more workers (and thus having more problems finding them) and technology adoption, introducing upward bias to OLS coefficients.

I take several steps to address these concerns. I start by using labor scarcity measure that is not specific to the firm but captures labor market conditions in the firm's local area. Because each firm is small compared to their local labor market, firm-specific factors do not influence local labor scarcity. However, local productivity shocks may affect both the labor market and the output market: when local economy is booming, the demand for goods sold locally is high and it is hard to find workers, because unemployment is low. High demand, in turn, may lead to higher technology adoption. To deal with this scenario, I limit sample of my firms to those who export significant share of their production, and hence are unlikely to be sensitive to the local economic conditions in their district differently than to conditions in the rest of Germany.

Nonetheless, some threats to the identification remain. For example, some areas may have better access to information and expertise about technology because they have more universities or other research institutes. This can cause bias if presence of these institutions is correlated with labor scarcity. To alleviate this type of concerns, I use a specific part of variation in labor scarcity that comes from aging of the workforce. When older workers retire and there are few young workers in the local area, it is harder to fill vacancies. The main factor driving differences in aging are fertility decisions made many years ago that are unlikely to be strongly related to today's technology adoption. While this approach alleviates previous worries, some concerns may still remain. In particular, linking back to previous example, the coefficient of labor scarcity may be upward biased if places with e.g. more universities are also places that have aged the most. Certain threats to validity of the instrument are related to migration, even though it is likely too small to drive the results and would bias my estimates towards zero (see Online Appendix). To alleviate this concern I use predicted aging based on the age distribution of population in 2004. Another concern could be related to product market effects of aging, i.e. firms can adopt new technologies because the age structure of their customers makes it more attractive. Again, this concern probably biases my estimates downward, because we typically expect younger customers to be more technology-savvy and thus aging should discourage firms from adopting new technologies. Nonetheless, I alleviate this concern by estimating my 2SLS specification in the subsample of exporters.

Labor Scarcity in the Local Area. Following the discussion of endogeneity concerns, I substitute firm-level measures of labor scarcity with measures defined on the local area level.

The equation I estimate is:

$$Technology_i = \tilde{\beta} \cdot LaborScarcity_a + \tilde{\gamma} \cdot I_j + \tilde{\phi} \cdot X_i + e_i \tag{1.10}$$

Labor scarcity in the local area is the share of firms in the area that report difficulties in finding workers.[16] The geographical variation in labor scarcity index is presented in Figure 1.6. The analysis is performed with ca. 400 districts, but it is robust to using ca. 100 larger areas (Raumordnungsregionen) instead.

Columns 1-2 of Table 1.5 present the results. Being located in an area where many other firms declare that they have troubles finding workers is associated with higher levels of digitization and automation adoption. Moving local labor constraints index from 10th percentile to 90th percentile increases adoption by around 10% of its standard deviation. Local area conditions affect firms through local labor market, but they may also affect it through the output market. To mitigate this concern, I estimate equation (1.10) using only the sample of firms that export at least 20% of their production. The result presented in column 2 confirms the positive relationship between labor scarcity and adoption of digitization and automation, suggesting that the demand channel likely does not explain my findings.

Instrumenting Labor Scarcity with Aging. While local area labor market conditions are exogenous to the firm (ignoring endogenous firm location decision), it still might be the case that unobservable characteristics of firms differ by area in a way which is correlated with labor scarcity, even after controlling for industry. To further alleviate the endogeneity concerns, I instrument local labor scarcity with aging patterns in the local area. German society is aging

[16] I employ leave-one-out procedure: for each firm in the sample I exclude its own declaration when calculating local averages. Therefore, denoting the variable with subscript *a* slightly abuses the notation

Table 1.5. Digitization and Automation and Labor Scarcity: Addressing Endogeneity Concerns

	Labor Scarcity in Local Area		Instrumenting Scarcity with Aging				
	Full Cross-Section	Exporters	Panel	Full Cross-Section			Exporters
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
			Y = Hard to Find Workers (Firm Level)				
% Workers > 55			0.196*				
(in local market)			(0.118)	First Stage:			
Δ % Workers >55				0.259*	0.373**	0.684**	0.489
(in local market)				(0.141)	(0.173)	(0.323)	(0.691)
IV				Aging	Aging	Pred. Aging	Aging
F-Stat				5.61	6.08	6.38	2.58
	Y=Digitization and Automation Adoption						
Labor Constraints	0.078*	0.170**					
Index	(0.047)	(0.081)					
				Second Stage:			
Hard To Find				3.882*	4.234**	5.245**	3.402
Workers				(2.127)	(2.125)	(2.090)	(2.473)
N	10250	1006	122694	8770	8746	8805	925
Industry FE	✓	✓	✓	✓	✓	✓	✓
Size	✓	✓	✓	✓	✓	✓	✓
Area Controls					✓		
Firm FE			✓				
Year FE			✓				

Dependent variable in the top panel is firm's declaration of difficulties in finding workers. Dependent variable in the bottom panel is the digitization and automation adoption measure from IAB Establishment Panel. Columns 1 and 2 present OLS regressions with a z-score of local area measures of labor scarcity, computed using leave-one-out procedure (excluding firm's own declaration). Column 3 presents panel regression of labor scarcity on aging, which can be interpreted as a conceptual first stage for 2SLS regressions presented in columns 4-7. In the top panel, columns 4-7 present first stage regressions with the dependent variable being firm's declaration about difficulties in finding workers. In the bottom panel, second stage regressions are presented and the dependent variable is firm's technology adoption. Independent variable in row 1 is the district-level share of workers above 55 in the workforce, while row 2 contains change in this share between 2004 and 2014. This change is used as an instrument together with 2004 level of labor constraints in the area, which proxies for unobserved fixed characteristics of the area. Column 6, instead of using actual aging, uses predicted aging based on the age distribution in the district in 2004. In column 7, the sample is limited to firms that in the last year (2016) exported at least 20% of their production. All columns include industry (2-digit) and firm size quintiles fixed effects. Area controls in column 5 refers to area-level measures of economic conditions: average wage and average sales per worker. Standard errors are clustered by area and industry. (*) denotes significance at 10% level, (**) at 5% level and (***) at 1% level.

rapidly and its fertility rate, around 1.4 in the last decade, is one of the lowest in the world. The total population was decreasing since 2003, although recent influx of immigrants have actually reversed this trend. With the median age of 47, 3rd highest in the world, aging is commonly considered to cause shortages of workers (Börsch-Supan, 2003). Figure 1.17 in the Appendix presents the evolution of the share of older workers over time and shows that it is accompanied by an upward trend in labor constraints declarations.

I use the differences in the intensity of aging across local areas to isolate part of variation in labor scarcity driven by demographics. I define aging as the change in the share of workers above 55 in the entire workforce and compute it using the information on the age composition from Establishment History Panel. This data set contains personnel records of half of all the German establishments and allows to calculate the share of workers above 55 per area with large precision. The retirement age in Germany is now slightly above 65 years but early retirement schemes allow many workers to retire earlier, depending on their work experience (e.g. at 60). The instrument is supposed to capture the fact that when there are relatively fewer young workers and many workers nearing retirement, finding a new worker is difficult. Figure 1.8 presents the map of the intensity of aging across Germany. On average, aging is higher in East Germany but there is substantial heterogeneity within the two parts of the country: central parts of the West are also aging fast, while aging is less intense in the Southern parts of the East.

Column 3 of Table 1.5 uses of the fact that I observe age composition and declarations about staffing problems for several years and shows that aging is related to more difficulties in finding workers even after controlling for firm and year fixed effects. Interestingly, a simple cross-sectional correlation between labor constraints and share of older workers has negative sign. Of course places with many older people are very different from places with many younger

Figure 1.8. Geographic Distribution of Aging
The map presents 2005-2015 changes in the share of workers above 55. The original changes are computed at the district level but for data confidentiality reasons presented values were computed on the spatial planning regions (RORs) level – each ROR contains ca. 4 districts. The values are computed based on the data from Establishment History Panel and are ratios of total number of workers who are 55 or more to total employment.

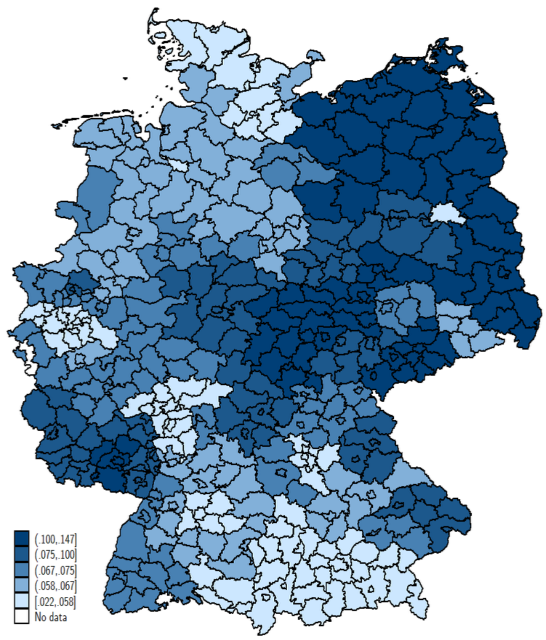

people and hence not taking into consideration these fixed differences obscures the real impact of aging.

Column 3 may be interpreted as a conceptual first stage, but it is not the same as the actual first stage presented in the upper part of column 4. Because my measures of digitization and automation are only available in 2016/17, the 2SLS regression has to be cross-sectional. This precludes the usage of firm or area fixed effects and requires a different specification. I instrument labor constraints today with the change in share of workers above 55 between 2004 and 2014 and with labor constraints index in 2004. The results of the first stage regression are presented in column 2: the local-area change in the share of older workers in the last decade positively and significantly predicts labor constraints today. The F-statistic is 5.8 which indicates that the relationship between instruments and labor scarcity measure is significant, although the value of the statistic is below the common rule-of-thumb for weak instrument assessment.

The results of the second stage regression are presented in the lower part of column 4. There is positive and significant relationship between labor scarcity, instrumented with aging, and adoption of digitization and automation, which confirms the findings from previous specifications and suggests that the relationship is causal. As discussed in section 1.4.2, OLS coefficient may suffer from both downward and upward bias. While the IV coefficient is larger than OLS, suggesting that the downward bias may be more pronounced, the confidence interval is large enough to accommodate bias of different direction and hence I do not draw strong conclusions regarding the true sign of the bias. Column 5 includes additional controls for area economic characteristics. The results are very similar when average wage and average labor productivity in the district are included in the regression. Column 6 presents the coefficients from IV regression where predicted aging is used instead of observed aging. I use age distribution by local labor market region in 2004, obtained from German Statistical office, and construct predicted change in the share of workers above 55. I assume that individuals who were between 55 and 65 years old in 2004 are no longer in the workforce in 2014, while individuals who were between 15 and 25 in 2004 are in the workforce in 2014. Predicted aging is strongly correlated with actual aging but eliminates the part of aging that comes from migration. Using predicted aging as an instrument confirms the results. Column 7 presents the coefficients from IV regression estimated on the subsample of exporters. While the small sample prevents me from obtaining significant estimates, the coefficient remains positive and of similar magnitude as in the main sample, suggesting the demand considerations related to aging are unlikely to drive the results.

1.4.3. Heterogeneity of Labor Scarcity Effect

The results presented so far suggest that the substitution effect dominates over the productivity effect on average, but the average result may mask heterogeneous impact of different technologies in different industries. Figure 1.9 shows the results of specification from Eq. 1.9 modified by adding an interaction of labor scarcity (main measure, difficulties in finding workers) with 10 broad industry groups (see Online Appendix for details on industry group definition). Panel A reveals that the aggregate positive relationship hides significant heterogeneity across industries. Industries like manufacturing or utilities display significant, positive relationship, consistent with the idea that recent technological progress there is mostly labor-substituting. At the same time, industries such as finance or professional services see either significant negative relationship between labor scarcity and technology adoption, or an effect close to zero, consistent with the view that recent technological progress there is complementing labor on net. Importantly, industries such as retail and wholesale trade or hospitality – which are responsible for large part of employment – display clear positive relationship, consistent with the substitution effect.

Panel B of Figure 1.9 provides further details about the heterogeneity, showing the effects on two alternative dependent variables: digitization and automation defined separately (using the measures from 2017 IAB-EP interacted with adoption intensity from 2016). Robots drive the substitution pattern in manufacturing and agriculture but play no significant role in other industries. Digitization is responsible for substitution in industries such as trade or hospitality and for complementarity in industries such as finance or education and health.

These results highlight the importance of considering a broad set of technologies and allowing for heterogeneous relationship of technology and labor across industries. Recent literature (Acemoglu and Restrepo, 2019; Dauth et al., 2018) analyzes industrial robots in manufacturing and shows that they can reduce employment. However, the current wave of technological change involves many other technologies which may affect different industries differently. In most industries, the substitution effect of technology dominates – and is often driven by technologies other than robots – and hence the overall employment problem is still a potential concern. At the same time, the complementarity effect of digital technologies seems to dominate in selected industries.

Another dimension of heterogeneity is related to different types of jobs or different skill levels of workers. Table 1.6 presents the results of Eq. 1.9 modified by adding interactions of labor scarcity with the share of administrative and non-administrative and skilled and unskilled workers in the firm. Columns 1 and 2 demonstrate that labor scarcity accelerates the adoption of technology especially when a firm employs many non-administrative workers. Columns 5 and 6 present similar analysis for ICT technologies in years 2001-02, where the opposite pattern is visible. This suggests that recent technology is affecting workers outside of office jobs, contrary to ICT technology in the early 2000s, which mainly affected administrative workers. Columns 3 and 4 show that labor scarcity accelerates technology adoption particularly when a firm employs many unskilled workers. This suggests that the substitution effect of new technologies is most pronounced for unskilled workers and hence the technologies are skill-biased.

Skill bias of the technology can also be analyzed using skill-specific measures of labor scarcity. To do that, I construct local area indices of labor scarcity for unskilled and skilled workers. I compute them by averaging labor scarcity declarations of all firms in the area whose

Figure 1.9. Labor Scarcity Effect by Industry

Panel A. Digitization and Automation Adoption

The Figure presents coefficients from regression of digitization and automation adoption (main measure of the adoption intensity from IAB-EP 2016) on labor scarcity interacted with indicators for 10 industry groups, controlling for industry fixed effects and firm size. Positive coefficients indicate that labor scarcity increases the adoption of the technology (and hence the substitution effect dominates), while negative coefficients indicate that labor scarcity impedes the adoption (and hence the productivity effect dominates).

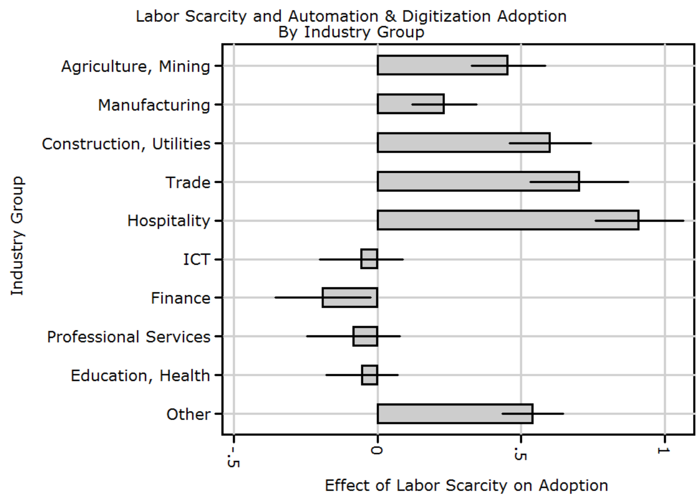

Panel B. Digitization and Automation Separately

The Figure presents coefficients from regression of two technology measures – automatin and digitization defined separately – on labor scarcity interacted with indicators for 10 industry groups, controlling for industry fixed effects and firm size. Automation is defined as the interaction of the main measure of the intensity of adoption (from IAB-EP 2016) with an indicator for using robots (from IAB-EP 2017) and considering them at least somewhat important (>=3 on 1-5 scale). Digitization is defined analogously, but with measures of digital technologies (data and networks).

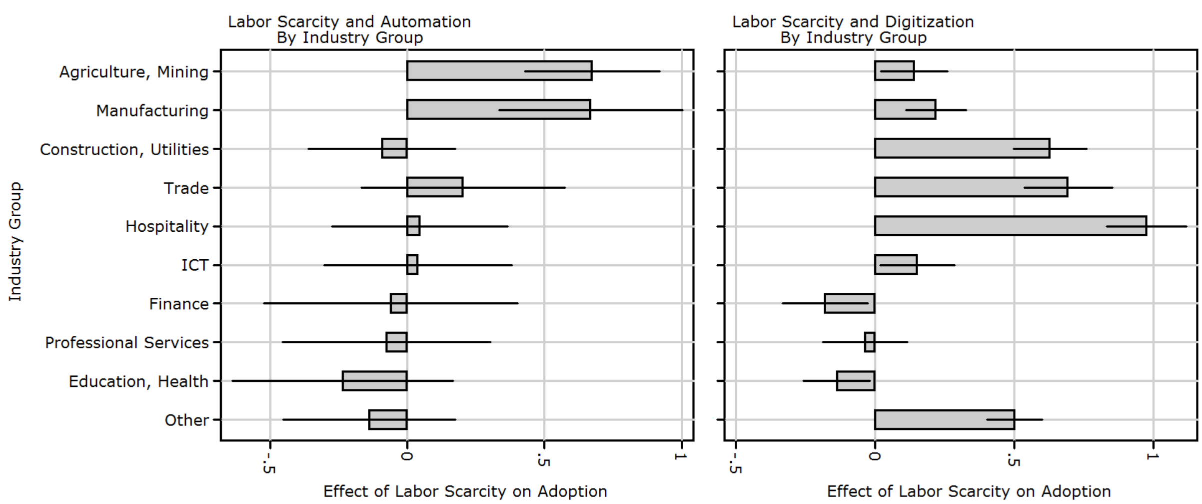

share of unskilled workers is above and below median, respectively. Column 5 of Table 1.6 show that there is no significant relationship between technology adoption and scarcity index for skilled workers which suggests that substitution and complementarity effects of the technology are roughly balanced. Column 6, on the other hand, shows strong and significant positive relationship that again suggests that the substitution effect dominates for low-skill workers.

1.5. Impact of Technology on Employment

This section studies how digitization and automation affect firms' employment. I regress 10-year changes in total employment on measures of technological change, defined on industry-area level, controlling for industry and area fixed effects. I find that automation significantly reduces employment, while digitization insignificantly increases it on average. However, aggregate results hide significant heterogeneity across industries: employment decreases in high- vs low-adoption areas in industries such as manufacturing, retail or hospitality, but increases in industries such as IT, finance or education and health. Consistent with Proposition 1, these findings paint the same picture of the technology as the analysis of the adoption patterns in section 1.4.

The two sections can be therefore viewed as alternative approaches to identifying unknown characteristics of the technology. At the same time, both bring a unique value. The adoption analysis informs us about the determinants of firm investment but is of partial equilibrium type and cannot show how technology affects patterns of firms entry or exit. The employment analysis, performed at the industry-area level, can take these patterns into account and moves one step towards general equilibrium analysis. More precisely, it captures any employment change at the industry-area cell level, although it still cannot capture the employment changes

Table 1.6. Technology and Labor Scarcity: Worker Type Heterogeneity

	Y=Digitization & Automation (2016)						Y=ICT (2001-02)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Hard to Find	0.296***	0.140	0.018	0.474***			0.094	1.069***	0.458*	0.459
Workers	(0.089)	(0.109)	(0.095)	(0.194)			(0.233)	(0.330)	(0.255)	(0.298)
% Admin	0.038						1.301**			
X HFW	(0.205)						(0.646)			
% Non-Admin		0.337*						-1.212**		
X HFW		(0.179)						(0.524)		
% Unskilled			0.785***						-0.193	
X HFW			(0.189)						(0.573)	
% Skilled				-0.423**						-0.185
X HFW				(0.194)						(0.567)
Skilled Scarcity					0.198					
Index					(0.220)					
Unskilled Scarcity						0.865***				
Index						(0.271)				
N	6626	6626	6626	6626	10169	10233	1047	1047	1047	1047
Industry FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Size	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Dependent variable in columns 1-6 is the adoption of digitization and automation from IAB Establishment Panel. Dependent variable in columns 7-10 is the average investment in Information and Communication technologies in 2001-02. Independent variable in rows 2-5 is firm-level indicator of difficulties in finding workers, interacted with the employment share of administrative and non-administrative workers (rows 2 and 3) and of unskilled and skilled workers (rows 4 and 5). The two classifications and based on occupational structure of the firm reported in 12 occupational groups. Since some occupational groups are difficult to classify, shares of the two groups (admin and non-admin or skilled and unskilled) do not sum up to 1. In rows 6-7 independent variables are indices of local labor scarcity for skilled and unskilled workers defined separately. Difficulties in finding workers are measured in 2014 for digitization and automation adoption and in 1999 for ICT investment. Each specification includes firm size fixed effects and industry fixed effects. Standard errors are clustered on the industry level, except for columns 5 and 6, where they are clustered on area and industry level. (*) denotes significance at 10% level, (**) at 5% level and (***) at 1% level.

in the whole industry, whole local area or in the whole economy. While I can perform analysis at the industry-level, the ability to cleanly identify the effects of technology is then limited.

Methodology. The naive approach to studying the effects of automation and digitization would regress the employment change in the last few years on the measure of digitization and automation from IAB Establishment Panel. This approach, however, suffers from several concerns that might lead to unwarranted conclusions. For example, the decision to adopt technology is endogenous and related to many characteristics of the firm. In particular, firms that are growing are more likely to adopt new technology and hence regressing e.g. employment change on technology adoption will not reveal the true causal effect of technology (see section 1.2.2 for evidence). To produce more reliable estimates, I employ an alternative specification and use other measures of digitization and automation. For digitization, I use stock of software and databases capital per worker from EU KLEMS database. For automation, I use number of robots per 1000 workers, based on robot shipments data from International Federation of Robotics. Both measures are at the country-industry-year level, The variation in these measures across industries can be used to analyze the employment effects of technology, as demonstrated by Graetz and Michaels (2018) and Acemoglu and Restrepo (2019). However, if all the identifying variation is at the industry level, it is difficult to disentangle the effects of technology from other industry-level changes correlated with technology adoption. I combine the variation in the intensity of technological change across industries with the variation in technology adoption levels across local areas. Doing this allows me to look within 2-digit industry and identify the effect of technology by comparing firms in high-adoption areas to those in low-adoption areas. The geographic variation in adoption is captured by local area[17] measures of intensity

[17]Area in this section is defined using spatial planning regions (ROR, Raumordnungsregionen), constructed by German Federal Authority of Construction and Regional Planning (BBR) taking into account the commuting

of digitization and automation adoption, computed as average of firm-level measures from IAB Establishment Panel.

I use data from the Establishment History Panel (BHP) – an administrative data set with information on 50% of all German establishments – and aggregate employment information to industry X area level, which is a natural choice given the desire to properly take into account firm entry and exit and that the variation in the independent variables is at the area-industry level. My main empirical specification is:

(1.11)

$$\Delta Y_{a,j} = \beta_R \cdot \left(\Delta Robots_j \cdot Adoption_a\right) + \beta_D \cdot \left(\Delta Digitization_j \cdot Adoption_a\right) + \phi I_j + \xi A_a + \varepsilon_{a,j}$$

This is a long differences specification with all changes in the above equation, denoted by Δ, corresponding to 10-year change between 2005 and 2015 (in the main model; other periods are considered in alternative specifications). Subscripts a and j denote area and 2-digit industry, respectively. $\Delta Robots_j$ is the change in number of robots per 1000 workers used in a given industry, coming from International Federation of Robotics data. $\Delta Digitization_j$ is the change in software and databases capital stock per worker in a given industry, coming from EU KLEMS data. Because of changes in reporting that happened in 2014/15, I use 2004-2014 change in digitization. $Adoption_a$ is a measure of digitization and automation adoption in area a. In the basic specification, this is an indicator of the intensity of adoption being above median. The intensity of adoption is the average of firms' declarations from the IAB Establishment Panel. Because in the declarations the firm compares itself to other firms in the same sector, the

patterns of workers. Germany is divided into 96 ROR regions. While RORs are good proxies for local labor markets, a possible alternative definition would use districts, which are smaller. However, for data confidentiality reasons, performing the analysis on the level of districts is not possible.

measure is not driven by industry composition of the area. The independent variables include vectors of industry fixed effects I_j and area fixed effects A_a. In the basic specification, I weight all observations with 2005 level of employment.

Interpreting the Empirical Specification. There are two ways to interpret the empirical specification. The first one is to consider it to be a difference in differences estimator in which the differences are taken across industries and areas (not across time, as in traditional DiD settings; difference across time is included in the dependent variable, which is the change in employment). The treatment is the change in digitization and automation intensity at the industry level and the treated group are firms in high-adoption areas, while control group are firms in low-adoption areas. The identifying assumption is that absent technological change, the difference between change in outcomes of the treatment and control group would not be systematically different across industries.

The second way to interpret the specification is in terms of the propensity to adopt technology. Two independent forces are pushing for the adoption: large technological change in the industry and being located in a high-adoption area. Looking at firms in industries with large technological change which are located in high-adoption areas and partialling out the effect of industry alone and area alone should therefore allow to isolate the effect of technology.

To intuitively understand the specification, consider an example of two industries - car manufacturing and paper manufacturing - with two firms in each of them. Let us assume that there is a large technological change in car manufacturing, but negligible technological change in paper manufacturing. In each industry, one firm is located in a high-adoption area, the other in an area with negligible adoption level. We are interested in learning how technology affects employment. To calculate this effect we need to compare how change in employment differs between high- and low-adoption-area firms in car manufacturing. The observed difference is a combination of the "technology effect" and of the "location effect". Comparing high- and low-adoption-area firms within paper industry – that has a negligible technological change and therefore "technology effect" is negligible – allows us to compute "location effect". Assuming that this effect does not systematically differ across industries, this allows us to back out the "technology effect" in the car industry.

Endogeneity. Changes in robots density and digitization intensity in Germany may be endogenous. For example, when German firms in a given industry face large demand, they may be adopting more robots and digital technologies and at the same time increasing employment – in which case my estimates of employment effect would be biased upwards. While in my specification the technology coefficient is identified using within-industry variation and thus the concern is less severe, the degree of technological change still influences the estimates (intuitively, the coefficient of technology is a weighted average of differences between high and low adoption areas across industries, with weights equal to intensity of technological change in the industry). To better isolate exogenous variation in technology, I follow the approach of Autor et al. (2013) and Acemoglu and Restrepo (2019) and use changes in robot density and software and databases capital in a group of other European countries.[18] I present both the reduced form estimates with technology abroad as the independent variable and IV estimates in which I use technology abroad to instrument the domestic technological change.

Differences in the adoption across local areas are not random and can be correlated with various other factors affecting employment. I do not assume that high- and low-adoption areas

[18]Both for robots and digitization I use 6 other countries but the group is different because of data availability. For robots, it includes France, Italy, Denmark, Netherlands, Sweden and United Kingdom. For software and databases capital, the group includes France, Italy, Belgium, Netherlands, Finland and Austria.

are similar except for the levels of technology adoption. Instead, I include area fixed effects with the goal of capturing all area-specific factors other than technology. The key identifying assumption is that the effect of these factors does not systematically vary across industries in a way correlated with the technology. For example, I allow for the presence of many universities in the area to affect employment, but I assume that it will affect employment in each industry in a way which is uncorrelated with the intensity of technological change.

Assuming that the differences between high- and low-adoption areas are similar across industries, except for the effect of technology, seems more plausible than assuming that all industries are similar, except for the effect of the technology. Nonetheless, the former assumption can still be violated. Most plausible concern is related to agglomeration effects differentially affecting different industries. High-skill industries may prefer to be located in selected business hubs more than manufacturing firms. The existence of these preferences alone does not pose a problem to my strategy. However, if these preferences are becoming more and more prevalent and if business hubs also have higher levels of technology adoption, we may see that employment in high-skill services (which have high digitization) increases, while employment in manufacturing (which has high robotization) decreases in high-adoption areas.

While this is a possible concern, it is unlikely to explain the whole set of results presented in this paper. In particular, employment effects also hold when controlling for past employment changes, and hence are unlikely to be driven by differential employment trends across industry-area pairs. Moreover, section 1.4 shows that technology adoption overall is higher in areas with more labor scarcity. If these characteristics identify "business hubs" that are attracting most productive firms in high-skill services industries, we should see positive relationship between technology adoption and labor scarcity for high-skill services – but we see the opposite.

In addition, the analysis of adoption patterns across Germany (Fig. 1.4) reveals that many areas with high adoption (e.g. northeastern Bavaria or western Lower Saxony) are not the typical business centers. Finally, the differential importance of agglomeration effects can be to large extent driven by technology and hence it may be viewed as a mechanism through which technology affects employment, rather than as an alternative explanation.

Results. Table 1.7 presents the results of employment effects analysis. I estimate Eq. 1.11 with dependent variable being the percentage change in employment between 2005 and 2015. The results for the main specification, presented in column 3, show that robotization has negative and significant effect on employment. One additional robot per 1000 workers reduces employment in high adoption areas by 0.36% in the 10-year period, compared to firms in the same industry in low adoption areas. The effect of digitization is positive, but insignificant in the main specification. When a continuous measure of adoption is used instead (column 4), effect of robots is still negative and significant, while effect of digitization remains positive and becomes significant. Interestingly, naive approach on regressing employment change on area- or industry-level measures of technological change (columns 1 and 2) shows very different results and highlights the necessity to properly take into account other industry-level changes.

The intensity of automation and digitization in Germany may be endogenous. To deal with this concern, I estimate an alternative specification that, instead of using domestic change in technology, uses change in the technology abroad (column 5) or instrument domestic changes with the changes abroad (column 6). Both results confirm the negative and significant effect of automation and positive but insignificant effect of digitization. Consistent with expectations, the coefficients of employment effect are lower, suggesting that endogeneity concerns can indeed to some extent bias the coefficients upwards.

Table 1.7. Employment Effects of Technology

	Y=%ΔEmployment									
	(2005-2015)						(2005-2010)		(2010-2015)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Adoption	1.894									
	(0.0134)									
Robots		-0.022								
		(0.328)								
Digitization		-1.674								
		(1.204)								
Robots X			-0.357*			-0.544*	-0.185		-0.768**	
Adoption>P(50)			(0.194)			(0.298)	(0.157)		(0.379)	
Digitization X			0.756			0.700	0.746**		1.054	
Adoption>P(50)			(0.464)			(0.542)	(0.380)		(0.882)	
Robots X				-0.278**						
Adoption (cont.)				(0.113)						
Digitization X				0.962*						
Adoption (cont.)				(0.545)						
Robots Abroad X					-0.635**			-0.325		-0.970***
Adoption>P(50)					(0.308)			(0.285)		(0.324)
Digitization Abroad X					0.179			0.271**		0.053
Adoption>P(50)					(0.174)			(0.113)		(0.231)
N	5275	5275	5275	5275	5275	5275	5202	5202	5202	5202
Area FE		✓	✓	✓	✓	✓	✓	✓	✓	✓
Industry FE	✓		✓	✓	✓	✓	✓	✓	✓	✓

Dependent variable is the relative change in employment for a given industry-area cell expressed in percentage points. Independent variables are robotization – change in number of robots per 1000 workers on industry level in Germany – and digitization – investment in Software and Databases capital per worker (in thousands of Euro) in Germany – and their interactions with indicators of a firm being located in high technology adoption area. The analysis is conducted on the industry-area level (2-digit industry; RORs/commuting zones). High/above-median and continuous local indicator for adoption defined based adoption measures from 2016 IAB Panel. Robots abroad and digitization abroad are defined analogously to German measures, except they are averages for several other European countries. Column 5 presents their reduced form relationship to employment change, while column 6 presents IV results. The first stage of the IV specification is presented in the Appendix. All regressions, except column 1 and 2, include industry and area fixed effects and are weighted using employment levels from 2005. Standard errors, reported in parentheses, are two-way clustered by area and industry. (*) denotes significance at 10% level, (**) at 5% level and (***) at 1% level.

Figure 1.10. Employment Effects of Technology by Industry
The Figure presents the difference in 2005-2015 employment change between high- and low-adoption areas for different industry groups. High-adoption area is defined as an area in the 4th quartiles of the technology adoption index, while low-adoption area is defined as an area in the 1st quartile of the adoption index. The adoption index is an area-level average of firm-level declarations of digitization and automation adoption from IAB-EP. The estimates are obtained by regressing industry-level differences between high- and low-adoption areas on indicators for 10 industry groups, weighting the observation by number of areas and controlling for average difference between high and low-adoption areas. Details of assignment to industry groups are presented in the Online Appendix. Whiskers represent 5% confidence intervals for the coefficients.

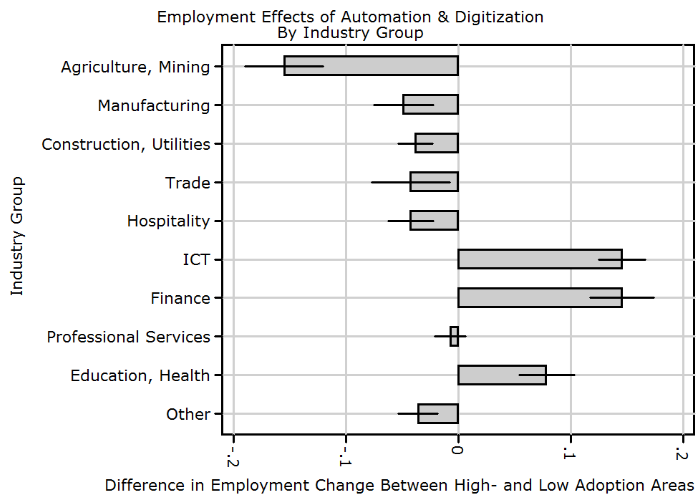

Columns 7-10 analyze alternative periods. The signs of the coefficients and main conclusion remain unchanged, although in the recent period negative effect of robotization is more evident, while positive effect of digitization is significant between 2005 and 2010.

Table 1.17 presents the results which confirm robustness of the main findings and contain additional details. The main result is robust to alternative measures of adoption, excluding automotive industry, assigning equal weights to each observation (as opposed to weighting by employment in 2005) or adding controls for past employment changes. In addition, the table presents first stage of the IV regressions.

Heterogeneity. The results presented in Table 1.7 show average effects of technology but Figure 1.9 suggests that they may be hiding important heterogeneity across industries. To shed more light on the across industry heterogeneity, I compute differences in employment changes between high- and low-adoption areas across industries. These differences, demeaned and aggregated to 10 industry groups in a regression with industry-group fixed effects, are presented in Figure 1.10. The Figure presents differences between areas in the 4th and 1st quartiles of adoption. The differences between above and below median areas show similar pattern, but the estimates are less precise. The difference in employment effects between high- and low-adoption areas is consistently negative for industries in which robots are present, i.e. mining, manufacturing and construction and utilities. However, the difference is of mixed sign in industries in which digitization is playing some role. While in some industries, including IT or finance, the employment change is higher in high-adoption areas, in others, such as trade and hospitality, the employment change is higher in low-adoption areas. This result confirms the findings presented in Figure 1.9. Consistent with Proposition 1, in industries in which labor scarcity increases the adoption of technology, the employment effect of the technology is negative, while in industries in which labor scarcity decreases the adoption, the employment effect is positive.

1.6. Additional Results

1.6.1. Financial Constraints and Technology Adoption

This section completes the analysis of determinants of technology adoption by studying the role of financial constraints. As shown in Table 1.1, technology adoption is associated with higher investment. This investment can be impeded if a firm faces financial constraints, understood as difficulties in accessing the capital. How important is this mechanism for the investment in digitization and automation, given that many new technologies – cloud computing, software-as-a-service type of programs – may not require any sizable capital investment?

Measures. My measures of financial constraints come from the IAB Establishment Panel. Because majority of firms in this data are private, and because the data is collected for different purposes, many traditional accounting variables (e.g. measures of liabilities) are not available. Instead, in selected years firms are explicitly asked about financial constraints, which is an interesting advantage of this data set. The survey asks firms if in the last year they had difficulties in obtaining credit. If the answer is positive, the firm is asked to give more details (credit application rejected, credit volume decreased, credit costs increased). I use a binary variable coded as 1 if the firm reports difficulties in obtaining credit as my first measure of financial constraints. In addition, firms report size of their investment in a given year, together with sources of its financing. In selected years, each firm that reports non-zero investment is asked what share of expenditures was financed by equity and debt. I define my leverage measure to be the share of debt in total investment. Unfortunately, these financial variables are part of Additional Modules of the IAB Establishment Panel and are not available after 2010. I use lagged measures which may introduce downward bias into my analysis because they may have limited power in explaining the cumulative level of investment in 2016-17. I complement these measures with another variable that is available every 2 years (last time in 2015) but only for a subset of firms. The variable is defined for a subset of firms that declare abandoning a planned project related to product or process innovation, and takes value of 1 if among reasons for abandonment the firm lists “lack of financing sources” (this is different than another reason that can be listed, “costs too high”), and 0 otherwise.

Table 1.8. Digitization and Automation Adoption and Financial Constraints

	Y= Digitization and Automation Adoption								
	All Industries					Manufacturing		Non-Manufacturing	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Difficulties	-0.640*			-0.676**		-0.782*		-0.578	
in Obtaining Credit	(0.327)			(0.295)		(0.392)		(0.434)	
Investment Prevented		-0.549***					-0.360		-0.576**
Can't Obtain Financing		(0.211)					(0.321)		(0.266)
Leverage			-0.041**	-0.031**					
			(0.017)	(0.014)					
Financial Constraints					-3.054*				
(area-level average)					(1.831)				
N	2217	1431	2129	1965	8763	604	489	1613	948
Industry FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Size	✓	✓	✓	✓	✓	✓	✓	✓	✓
IV					CB				
F-Stat					4.23				

Dependent variable in all columns is the digitization and automation adoption measure from IAB Establishment Panel. Independent variables are firms' reported difficulties in obtaining credit; declarations of abandoning innovative project because of inability to access financing; share of debt in the financing of capital expenditures; and local area share of firms abandoning the innovative project because of inability to access financing. The last independent variable is instrumented in column 5 with quartile of area-level dependence on Commerzbank (based on Huber (2018)). Each specification includes firm size quintiles fixed effects and industry fixed effects. Standard errors are clustered on the industry level, except for column 5, where they are clustered on area and industry level. (*) denotes significance at 10% level, (**) at 5% level and (***) at 1% level. Columns 1-5 show the results for the full sample. Columns 6-9 show the results for manufacturing and non-manufacturing industries separately.

Table 1.8 presents the results. Column 1 shows that there is a significant, negative relationship between financial constraints and adoption of digitization and automation. Column 2 confirms this negative relationship using an alternative measure of financial constraints based on reasons for abandoning investment. The magnitudes of the coefficients are similar, even though the measure from column 1 is based on declarations from 2008. It is consistent with the view that financial constraints are persistent and that the current level of technology adoption stems from investment decisions made in the last several years. Columns 3 and 4 use share of debt in total investment to proxy for financial constraints and again confirm the negative relationship.

Comparison of columns 2 and column 7 in Table 1.4 unveils an interesting contrast between financial and labor constraints. The sample in both cases consists of the same subset of firms that declare abandoning an investment in an innovative project. When the reason for abandoning the investment is the lack of access to finance, the firm has lower levels of digitization and automation (which may directly result from abandoning the innovative project, possibly involving new technologies). Yet, when the reason for abandoning the investment is the lack of qualified personnel, adoption of automation and digitization is higher.

Endogeneity. The evidence presented in columns 1-4 of Table 1.8 can suffer from endogeneity concerns, similar to those discussed in section 1.4.2. To some extent, these concerns are less severe, e.g. the reverse causality is rather implausible in the case of financial constraints. Nonetheless, to alleviate other concerns, I follow Huber (2018) and use lending cut of a large commercial bank in Germany – Commerzbank – as an exogenous shock to the availability of credit. For historical reasons, some areas in Germany have a larger share of firms with relationship to Commerzbank than others. When in the course of financial crisis Commerzbank significantly limited lending because of losses they suffered in their international trading activities, the ability to obtain credit decreased in areas more exposed to the bank. In my data, being located in an area more exposed to Commerzbank is indeed related to higher probability of firms reporting difficulties in obtaining credit in years 2009-2015, but is uncorrelated with financial constraints in 2008. Using area-level exposure to Commerzbank as an instrument for area-level average difficulties in obtaining financing (2009-15) confirms the negative effect of financial constraints on technology adoption. In addition, Table 1.18 in the Appendix shows the estimates of the main specification after including additional controls, which also alleviates the concern that financial constraints are only proxying for other firm characteristics.

Columns 6-9 of Table 1.8 confirm that financial constraints impede technology adoption not only in manufacturing, where capital-intensive robots are ubiquitous, but also in other industries, where many new technologies, e.g. cloud computing, are explicitly designed to limit capital investment. Even though smaller sample size prevents me from obtaining significant estimates in all specifications, coefficients are consistently negative and significant results are obtained both within and outside of manufacturing.

1.6.2. Technology and Changes in the Workforce Skill Structure

New technology may affect not only firm's total employment level but also the skill level of the workforce. In theory, the existence and direction of a skill bias is ambiguous, as formalized in Proposition 3 (Online Appendix). On the one hand, new technologies are commonly thought to substitute low skilled workers and complement high skilled workers. This assumption is explicitly built into the model of Acemoglu and Restrepo (2018b) and is consistent with (Graetz and Michaels, 2018) who show that higher usage of robots is associated with lower low skill employment. An extensive literature on ICT revolution points to the polarization effect – share of low and high skilled workers increases at the expense of middle skilled workers. At the same time, artificial intelligence is often thought to threaten also high skilled workers – performing tasks such as credit application approval – for which labor input can be strongly decreased thanks to advances in big data analytics and machine learning.

I define educational structure of the workforce using skill level information from Establishment History Panel (BHP) administrative records. For each firm, number of workers in low, medium and high education group is reported every year. The groups are based on German educational system and are defined as follows: low skilled include workers without vocational qualifications; medium skilled include workers with vocational education but no higher degree; high skilled workers include employees with university degree or applied university degree (Fachhochschule). In the data, around 12% of workers are low skilled, 73% medium skilled and remaining 15% high skilled.

Table 1.9 presents the results of educational structure analysis. I estimate Eq. 1.11 and use 10-year change in the share of low-, medium- and high-skill workers as my dependent variables. The table shows that both digitization and automation are associated with skill upgrading, even though the significance of coefficients appears in different columns for the two technologies and for digitization the results are only significant when technology abroad is used as a proxy for technological change. These results are consistent with the findings from Table 1.6, which suggest that the substitution effect of new technologies affects mostly unskilled workers (in Table 1.6 workers are classified based on whether they performed simple or complex tasks, as opposed to level of education used in Table 1.9).

If digitization and automation require new skills, not possessed by firm's existing workers, firms can adapt not only by hiring better educated workers but also by training their existing

Table 1.9. Effects of Technology on Skill Structure and Training

Panel A: Skill Structure

	Y=Δ% Low-Skilled		Y=Δ% Medium-Skilled		Y=Δ%High-Skilled	
	(1)	(2)	(3)	(4)	(5)	(6)
Robots X Adoption	-0.029		-0.027		0.057***	
	(0.035)		(0.033)		(0.016)	
Digitization X Adoption	-0.113		0.045		0.068	
	(0.097)		(0.210)		(0.114)	
Robots Abroad X		-0.002		-0.068		0.070**
Adoption>P(50)		(0.061)		(0.051)		(0.030)
Digitization Abroad X		-0.050*		0.018		0.032
Adoption>P(50)		(0.027)		(0.072)		(0.045)
N	5268	5268	5268	5268	5268	5268
Area FE	✓	✓	✓	✓	✓	✓
Industry FE	✓	✓	✓	✓	✓	✓

Panel B: Training

	Y=% Workers Trained			Y=# Training Measures		
	(1)	(2)	(3)	(4)	(5)	(6)
Adoption	1.352***			0.933***		
	(0.316)			(0.227)		
Robots X Adoption		0.053	-0.033		0.065	0.064
		(0.061)	(0.078)		(0.043)	(0.060)
Digitization X Adoption		0.335***	0.847**		0.226**	0.589
		(0.119)	(0.351)		(0.108)	(0.369)
N	983	8746	983	986	8780	986
Area FE	✓	✓	✓	✓	✓	✓
Industry FE	✓	✓	✓	✓	✓	✓
Firm Controls	✓	✓	✓	✓	✓	✓
Lagged Dep. Var.	✓		✓	✓		✓

Dependent variables in the top panel are 2005-2015 changes in the share of low-, medium- and high-skill workers expressed in percentage points. The analysis is conducted on the industry-area level (2-digit industry; RORs/commuting zones). Independent variables are interactions of the average industry-level robotization – measured as 2005-2015 change in number of robots per 1000 workers – and digitization – measured as 2004-2014 investment in Software and Databases capital per worker (in thousands of Euro) – with indicators of a firm being located in an area with high technology adoption. The local indicator for adoption is defined based on area-level average of responses to Digitization and Automation adoption question from IAB Establishment Panel (measured in 2016) and is defined as having the adoption indicator above median. All regressions include industry and area fixed effects and are weighted using employment levels from 2005. The analysis in the bottom panel is conducted on the establishment level. Dependent variable in columns 1-3 is the average share of workers undergoing any training between 2005 and 2015 in percentage points; dependent variable in column 4-6 is the average number of training methods in use reported by the firm in 2005-2015. The independent variable is the digitization and automation adoption measure from IAB Establishment Panel and its interaction with 2005-2015 changes in the number of robots per 1000 workers and in the 2004-2014 investment in software and databases capital per worker on the industry level. Columns 1, 3, 4 and 6 also include values of the dependent variable in the earlier period, i.e. between 1995 and 2000. All regressions include industry and area fixed effects. (*) denotes significance at 10% level, (**) at 5% level and (***) at 1% level. Standard errors, reported in parentheses, are two-way clustered by area and industry (upper panel) or clustered by industry (lower panel).

employees. While the second solution may not always be feasible (certain skills might be very hard to acquire or require many years of education), its advantage is that it allows the firm to retain existing workers that have valuable experience and firm-specific human capital. Also, if particular technology requires skills that are specific to the firm (e.g. because a digital system installed in the firm is unique), internal training might be the only way to acquire necessary skills.

Panel B of Table 1.9 analyzes measures of training intensity coming from the IAB Establishment Panel. Each firm reports how many workers took part in training activities in the last year and declares what type of training methods were used (external or internal courses, symposia, on the job training, etc.). Based on this information, I construct the share of workers who underwent training and the number of training methods used. I calculate averages of these two variables between 2005 and 2015. Using them as dependent variables, I estimate an equation analogous to Eq. 1.11, but at the firm-level and with firm-level measure of adoption:

$$\Delta Y_i = \beta_R' \cdot \left(\Delta Robots_j \cdot Adoption_i\right) + \beta_D' \cdot \left(\Delta Digitization_j \cdot Adoption_i\right) + \phi' I_j + \xi' A_a + \varepsilon_i \tag{1.12}$$

In this specification, instead of interacting industry-level changes in technology with area-level adoption propensity, I directly interact technological change with firm adoption measure. This allows me to measure the technology with less noise but requires availability of firm-level adoption measure (and thus can be used only in a subsample of firms from IAB Establishment Panel; however training intensity is also observed only for those firms) and may suffer from endogeneity concerns, requiring more caution when interpreting the results.

There is a positive and significant association of adoption and training intensity. Firms that adopt the technology are training more workers and use more training methods. This increase in training is especially pronounced for industries that have high levels of digitization and is not significantly higher for industries with high levels of automation. These results suggest that new technologies do require new skills and training existing workers plays a significant role in the adaptation to new technologies. The training, however, seems to be used mostly in case of digital technologies.

1.6.3. Technology and Number and Size of Firms

Employment effects of technology may not affect all firms equally. Instead, they may mask both employment changes in existing firms as well as firm creation and destruction. Table 1.10 analyzes the effects of technology on the number of firms and average firm size in the area-industry cell. Robotization decreases the number of firms and insignificantly increases firm size. Digitization has the opposite effect – it increases the number of firms and decreases average firm size. This is consistent with the fact that robots are most useful for firms with large scale of production. At the same time, it is consistent with modern digital technologies, such as cloud computing, being available also to small firms and reducing barriers to entry in some sectors. It is worth remembering that, compared to other publicly available data sets (e.g. Compustat), the data used in this paper contains many small, private establishments. Therefore even though some technological forces may lead to increasing concentration among the very large firms (as evidenced by the example of Google and other similar firms), among smaller establishments the effect of digitization appears to be the opposite.

Table 1.10. Effects of Technology on the Number and Size of Establishments

	Y=ΔLOG(NUMBER OF FIRMS)		Y=ΔAVG FIRM SIZE	
	(1)	(2)	(3)	(4)
Robots X	-0.0017***		0.381	
Adoption>P(50)	(0.0006)		(0.712)	
Digitization X	0.0065**		-0.985	
Adoption>P(50)	(0.0031)		(0.609)	
Robots Abroad X		-0.0033***		0.542
Adoption>P(50)		(0.0006)		(0.785)
Digitization Abroad X		0.0023**		-0.473***
Adoption>P(50)		(0.0011)		(0.131)
N	5269	5269	5269	5269
Area FE	✓	✓	✓	✓
Industry FE	✓	✓	✓	✓

Dependent variable in columns 1-2 is the change in the logarithm of the number of firms in the industry-area cell. Dependent variable in columns 3-4 is the change in average number of employees per firm in the industry-area cell between 2005 and 2015. The analysis is conducted on the industry X area level (2-digit industry; RORs/commuting zones). Independent variables are robotization – measured as change in number of robots per 1000 workers on industry level in Germany – and digitization – measured as stock of Software and Databases capital per worker (in thousands of Euro) in Germany – and their interactions with indicators of a firm being located in an area with high technology adoption. The local indicator for adoption is defined based on area-level average of responses to digitization and automation adoption question from IAB Establishment Panel (measured in 2016). High adoption area is defined as having the adoption indicator above median. Robots abroad and digitization abroad are defined analogously to German measures, except they are averages for several other European countries. All regressions include industry and area fixed effects and are weighted using employment levels from 2005. Standard errors, reported in parentheses, are two-way clustered by area and industry. (*) denotes significance at 10% level, (**) at 5% level and (***) at 1% level.

1.6.4. Technology and Labor Productivity

This section studies how technology adoption affects labor productivity. While in the basic theoretical framework technology typically leads to productivity improvements[19], in reality productivity gains from technology adoption are not always evident. In 1987 Robert Solow famously said that "You can see the computer age everywhere but in the productivity statistics", which succinctly captures the concept of productivity paradox – an observed slowdown in productivity growth in the 1980s despite rapid adoption of IT technologies. The possible return

[19] Because otherwise it would not be adopted. However, in a richer dynamic model it is possible to observe that technology adoption decreases initial productivity, but increases it in the future.

of this paradox, i.e. lack of productivity gains from IT investment in recent decade, is discussed by Acemoglu et al. (2014), while Brynjolfsson et al. (2019) discuss a similar paradox in the context of artificial intelligence.

In my data, productivity – defined as sales or value added per worker – can only be observed for establishments surveyed in the IAB Establishment Panel. While Eq. 1.11 can be estimated using only these firms, the smaller sample size combined with the fact that proxying for technology usage with local area adoption necessarily introduces noise, makes it difficult to obtain precise estimates. To deal with this problem, I employ an analogous specification that makes direct use of firm-level technology adoption declarations, see Eq. 1.12.

Table 1.11 presents the results. Higher adoption of robots at the firm-level is associated with an increase in labor productivity (column 2) but the effect of digital technologies is insignificant. Using industry-level measures of technological change confirms the positive effect of robots, both when using variation across industries (column 4) or when relying on within-industry variation across firms with different levels of adoption (column 5). The findings are the same also when weighting observations by initial employment (column 6) and when using value added per worker instead of sales (column 7). One extra robot per 1000 workers is associated with 3% higher labor productivity in high-adoption firms compared to low-adoption firms. Interestingly, the effect of digitization is significant in column 5, but not in column 6, which is weighted by employment. This may suggest that digitization increases productivity to a limited extent and the gains are concentrated in small establishments.

Table 1.11. Effects of Technology on Labor Productivity

	Y=%ΔSALES PER WORKER						Y=ΔVA P.WRK
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Adoption	1.325						
	(1.013)						
Adoption X		3.328*					
Has Robots		(1.769)					
Adoption X			0.534				
Has Digit Tech			(0.882)				
Robots				0.818***			
				(0.245)			
Digitization				0.974			
				(1.083)			
ΔRobots X					2.318*	2.442**	6.142***
Adoption>P(50)					(1.319)	(1.248)	(1.835)
ΔDigitization X					8.984***	0.278	3.311
Adoption>P(50)					(3.021)	(7.344)	(6.484)
N	1364	1364	1364	2223	1364	1364	1398
Ind FE	✓	✓	✓		✓	✓	✓
Weights						Emp	

The analysis is conducted on the establishment level. Dependent variable in columns 1-6 is the relative change in labor productivity (sales per worker) between 2005 and 2015 in percentage points. Dependent variable in column 7 is the relative change in value added per worker. Independent variable "Adoption" is the firm-level measure of the intensity of digitization and automation adoption from IAB Establishment Panel (wave 2016). Adoption X Has Robots and Adoption X Has Digit Tech is the adoption measure indicator interacted with binary indicators for using robots and using digital technologies (from 2017 wave of the IAB Establishment Panel). Robots and Digitization, respectively, denote 2005-2015 changes in the number of robots per 1000 worker and in the stock of software and databases capital per worker on industry level. They are interacted with high adoption variable – binary indicator of firm-level adoption being above median. All regressions include industry fixed effects. Standard errors, reported in parentheses, are clustered by industry. All observations are weighted equally, except for column 6, where firm-level employment in 2005 is used as weights. (*) denotes significance at 10% level, (**) at 5% level and (***) at 1% level.

1.7. Conclusions

The main contribution of this paper is to inform the debate on automation, AI and other digital technologies and their impact on the future of work. Numerous countries and organizations are devoting significant attention to these new technologies, but the debate often remains on a superficial level and is based on anecdotes and futuristic visions. While some voice concerns that automation will destroy majority of jobs and impoverish large parts of the society, others are enthusiastic about the benefits of the technology and expect it to virtually end the problem of scarcity. This paper attempts to confront these predictions by taking more systematic approach and providing comprehensive empirical evidence based on rich firm-level data for a broad set of technologies and a complete set of industries.

The results show that the adoption of the most recent wave of new technologies – digitization and automation – is typically increased by labor scarcity, suggesting that these technologies substitute for workers on average. Consistent with that, the new technologies typically reduce employment. These effects, however, vary significantly across industries, worker types and technology classes. Average effects are driven by industries such as manufacturing, retail and hospitality, but in industries such as finance and education and health, technology seems to complement workers and leads to increased employment.

The fact that technology adoption is driven by labor scarcity means that machines are not necessarily stealing workers' jobs, even in sectors where substitution dominates. Instead, their adoption could be a response to the lack of workers. In addition, to the extent that they do displace some workers, they do so in places where jobs are most abundant. This endogeneity of adoption should be taken into account in the design of policies aiming to help workers and firms affected by technological change, especially those targeted to specific local areas.

Most importantly, the heterogeneity of the technology-labor relationship highlights that technology leads some industries to shrink, but others to grow. This pattern suggests that a key challenge associated with the current technological change is the facilitation of workers' transition between different sectors. This is not a new challenge, since a similar transition from agriculture to manufacturing accompanied previous waves of technological change in

XIX and XX century. Past transitions, however, were often slow and costly for large parts of the population. Exerting effort to make the current transition smoother and more equitable is the major way in which economics can help the society face the challenges associated with the technological change.

Appendix

Figure 1.11. Digitization and Automation Adoption: Summary Statistics by Industry Group

The Figure presents summary statistics for the intensity of digitization and automation adoption from the IAB Establishment Panel (part C - intensity of adoption on the scale from 1 to 10) by industry group. Bold line inside the box represents the median of firms declarations. Box limits represent one standard deviation below and above the mean declaration (and hence the center of the box represents the mean). The whiskers represent 10th and 90th percentile of the declarations. Minimum and maximum for each group, not depicted, equals 1 and 10 respectively. 10 broad industry groups are defined based on grouping consecutive 2-digit NACE Rev. 2 codes – the details are reported in the Online Appendix.

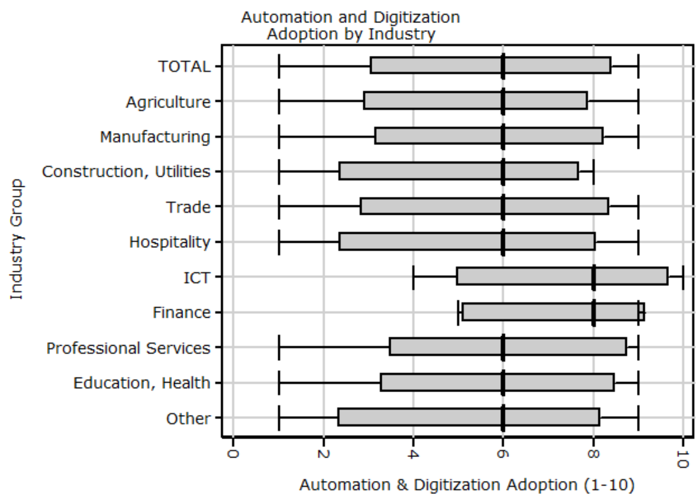

Figure 1.12. Evolution of Robotization and Digitization for Germany and Other Countries

The figure shows the evolution of robot density (number of robots per 1000 workers, based on IFR data) and of digitization (stock of software and databases capital per worker, in tho. Euro, based on EU KLEMS data) in Germany and other European countries. Both for robots and digitization I use 6 other countries but the group is different because of data availability. For robots, it includes France, Italy, Denmark, Netherlands, Sweden and United Kingdom. For software and databases capital, the group includes France, Italy, Belgium, Netherlands, Finland and Austria.

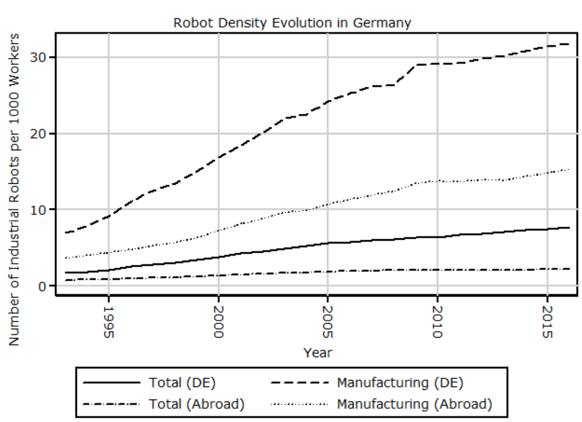

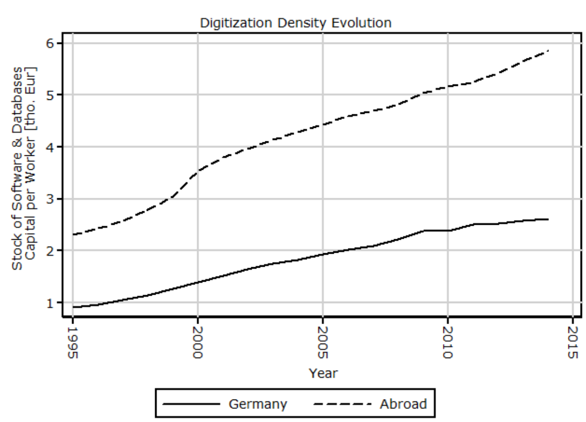

Table 1.12. Summary Statistics of Selected Variables

		Mean	Std Dev	P25	Median	P75	Num Obs
Investment	(% sales)	6.81	24.80	0.5	2.37	6.55	10302
High Adoption		0.33	0.47	0	0	1	7512
Unemployment Rate	(Area)	7.56	3.00	5.2	7.3	9.8	13109
Share of Workers >55	(Area)	0.25	0.03	0.23	0.25	0.27	12761
Financial Constraints		0.04	0.18	0	0	0	3010
Debt/Other Sources		0.47	1.68	0	0	0	2894
Number of Employees		106.8	853.3	4	14	59	14202
Share Unskilled		0.36	0.35	0.03	0.25	0.67	12141
Share Admin		0.29	0.33	0.02	0.14	0.47	12141
Share Workers Trained		0.31	0.29	0.06	0.24	0.50	12160
Sales	(mln Euro)	25.7	640	0.25	1.03	5.4	7874
Sales per Employee	(tho Eur)	131	230	42	75	140	7874
Δ%Sales Per Employee	(2005-15)	28.9	97.4	-11.2	13.0	44.6	2223
Robots (Ind)	(2015)	4.1	18.7	0	0	0	5511
Digitization (Ind)	(2015)	4.2	7.6	1.0	1.5	2.2	5640
ΔRobots (Ind)	(2005-15)	0.84	4.79	0	0	0	5402
ΔDigitization (Ind)	(2005-15)	1.56	3.41	0.17	0.45	0.50	5535
Adoption (Area)	2016	5.78	0.65	5.48	5.88	6.25	5642
% Low Skill	(2015)	11.9	8.5	6.7	10.5	15.2	5655
% Medium Skill	(2015)	72.8	13.9	66.7	75.8	81.8	5655
Δ% Low Skill	(2005-15)	-3.3	7.6	-5.6	-2.9	-0.6	5557
Δ% Medium Skill	(2005-15)	-1.1	9.7	-4.5	-0.6	0.3	5557

Summary statistics for technology, labor scarcity and employment measures are presented in Tables 1.1 and 1.3. Investment is the average value of investment in 2011-2016, expressed as the share of firm's sales. Variable is missing if a firm has not reported any positive investment in that period. High adoption is a binary measure that combines survey declaration about automation and digitization adoption (part C) with information about firm investment: it equals 1 if both adoption and investment are above industry-wide median. District-level unemployment rate and share of workers above 55 are from 2014. Share of unskilled and administrative workers comes from BHP extension to IAB-EP and represents 2014 value for the share of workers performing unskilled and administrative tasks, based on 12-group Blossfeld Occupational Classification used in Social Security records. Share of workers trained is based on average of firms' declarations in in 2005-2015 waves of IAB Establishment Panel. Financial constraints is firms' declaration that they had troubles getting credit (from 2008). Leverage is the ratio of debt to other sources (equity and subsidies) of investment financing in 2008. Sales are in thousands of Euro and are from 2015. Change in sales per worker is in relative terms and only available for a subset of firms for whom both 2005 and 2015 IAB Establishment Panel responses are observed. Robots and their change are expressed as number of robots per 1000 workers and come from International Federation of Robotics data (employment comes from EU KLEMS database). Digitization is the stock of software and databases capital in thousands of Euro per worker, coming from EU KLEMS database. Adoption is the Raumordnungsregion (ROR/commuting zone) average of firm declarations about intensity of digitization and automation adoption from 2016 IAB Establishment Panel. Shares of low- and medium- workers are based on workers' three educational groups reported in the BHP data. High-skill workers are the remaining group, omitted for brevity.

Figure 1.13. Changes in Robotization and Digitization in 2005-2014/15

The Figure shows the 2005-2015 change of robot density (number of robots per 1000 workers, based on IFR data) and 2004-2014 change of digitization (stock of software and databases capital per worker, in tho. Euro, based on EU KLEMS data) in Germany and other European countries. Both for robots and digitization I use 6 other countries but the group is different because of data availability. For robots, it includes France, Italy, Denmark, Netherlands, Sweden and United Kingdom. For software and databases capital, the group includes France, Italy, Belgium, Netherlands, Finland and Austria.

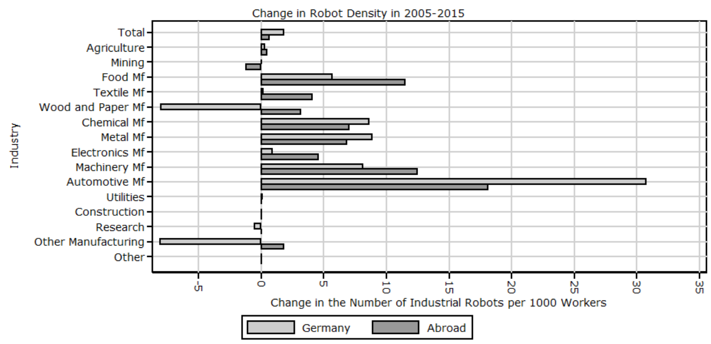

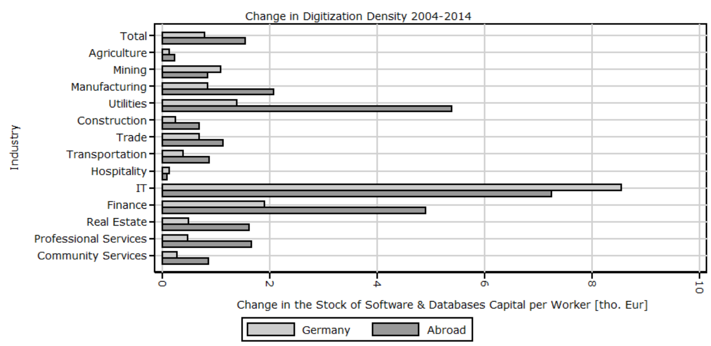

Figure 1.14. Digitization and Automation Usage and Other Firm Characteristics

The Figure shows the relationship between Digitization and Automation adoption and various firm characteristics: introducing product innovation in the last year, establishment being part of a multi-establishment firm, establishment having foreign owner, and being part of public firm. All variables come from the most recent wave of the IAB Establishment Panel in which a variable is available.

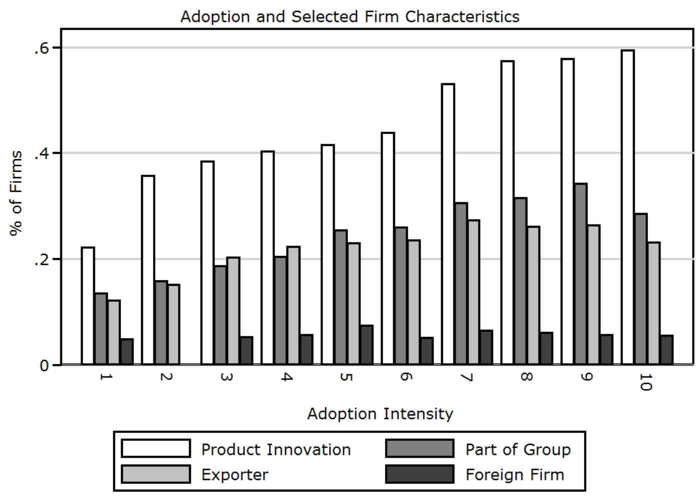

Figure 1.15. Digitization and Automation Usage and Wages

The graph shows the relationship between Digitization and Automation adoption and average wage in the establishment. Both adoption and wages data come from IAB Establishment Panel.

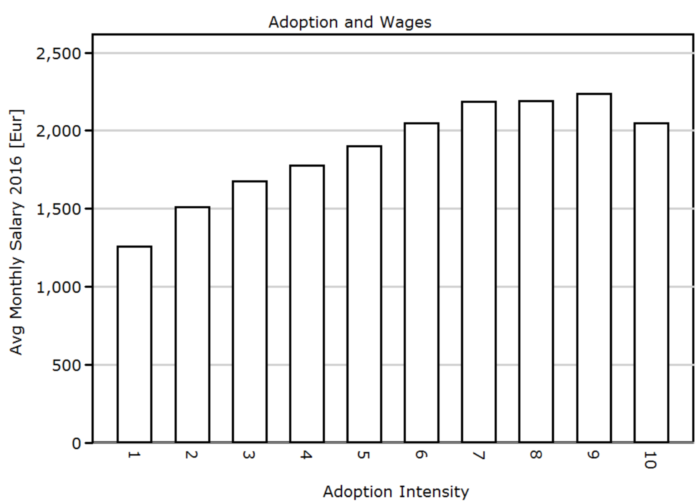

Figure 1.16. Changes in Employment by Industry in 2005-2015 for Germany and Other Countries

Based on EU KLEMS data. Foreign countries include Austria, Belgium, France, Finland, Italy and Netherlands.

Figure 1.17. Aging and Labor Scarcity in Germany

The graph presents the evolution of workforce aging and labor scarcity index. Aging is measured using the share of workers above 55 years among all workers. Labor scarcity index is the average of firms' declarations of having difficulties infinding workers in different waves of IAB Establishment Panel (the question is not asked every year and hence no continuous series can be plotted; instead, linear fit is shown on the graph together with values for each available year). Because of changes in reporting in 1999, the values of share of workers above 55 before 1999 were adjusted to remove discontinuity (increased by 0.04).

Table 1.13. OLS Regression of Digitization and Automation Adoption on Labor Scarcity with Additional Control Variables

	Y = Digitization and Automation Adoption (2016)								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Hard to Find Workers	0.294***	0.301***	0.308***	0.299***	0.312***	0.347***	0.334***	0.396***	0.372***
	(0.067)	(0.068)	(0.064)	(0.068)	(0.066)	(0.070)	(0.073)	(0.094)	(0.105)
N	7469	7346	7401	7469	7434	6590	6232	4419	3479
Industry FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Size	✓	✓	✓	✓	✓	✓	✓	✓	✓
Profitability	✓								✓
Part of Group		✓							✓
Establishment Age			✓						✓
Employment Growth				✓					✓
Public Firm					✓				✓
Foreign Owner						✓			✓
Professional Management							✓		✓
Startup (not Spin-out)								✓	✓

All columns present specification analogous to column 1 from Table 1.4, but with additional controls. The controls include dummies for profitability assessment, being part of multi-establishment group; dummies for establishment age, the speed of employment growth in last 3 years, being a public firm, having a foreign owner, being managed by a professional manager and being a startup (i.e. the establishment was started as startup, as opposed to being spun off from other existing establishment). Because of missing values in additional control variables the sample size varies between columns.

Table 1.14. Persistence of the Labor Scarcity Effect

	Y = Digitization and Automation Adoption (2016)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Hard to Find Workers (2016)	0.272***						
	(0.057)						
Hard to Find Workers (2014)		0.308***					
		(0.068)					
Hard to Find Workers (2012)			0.286***				
			(0.087)				
Hard to Find Workers (2010)				0.225***			
				(0.076)			
Hard to Find Workers (2008)					0.142		
					(0.087)		
Hard to Find Workers (2006)						0.089	
						(0.116)	
Hard to Find Workers (2004)							0.132
							(0.128)
N	10196	7469	5666	4498	3604	2832	2262
Industry FE	✓	✓	✓	✓	✓	✓	✓
Size	✓	✓	✓	✓	✓	✓	✓

All columns present specification analogous to column 1 from Table 1.4, but with labor scarcity measures coming from different waves of the IAB Establishment Panel.

Table 1.15. Semi Difference-in-Difference Approach: Including ICT Adoption in the Early 2000s

	Y= Digitization & Automation Adoption (2016)		Y=ΔTechnology (A&D - Comp 01)	Y=ΔTechnology (A&D - ICT01-07)
	(1)	(2)	(3)	(4)
Hard to Find	0.315**	0.195***	0.317	0.227*
Workers (2014)	(0.132)	(0.073)	(0.285)	(0.124)
Computers (2001)	0.072**			
	(0.029)			
ICT Investment		0.158***		
(2001-07)		(0.026)		
N	1351	2840	1351	2840
Industry FE	✓	✓	✓	✓
Size	✓	✓	✓	✓

In columns 1 and 2, specification analogous to column 1 from Table 1.4 is presented, but additional independent variables - decile of computer usage in 2001 and decile of ICT investment in 2001-07 period - are included. Using these variables reduces sample size because only selected firms were interviewed in past waves of the IAB Establishment Panel. Columns 3 and 4 present specification analogous to column 1 from Table 1.4, but with the dependent variable being the difference between intensity of 2016 digitization and automation adoption and the decile of computer usage in 2001 (column 3) or decile of ICT investment in 2001-07 (column 4).

Table 1.16. Other Staffing Problems

	Y = Digitization and Automation Adoption (2016)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Too Many Employees	-0.017							
	(0.136)							
High Labor Costs		-0.043						
		(0.068)						
Aging Population			-0.114					
			(0.066)					
High Labor Turnover				0.164				
				(0.111)				
Demand For					0.343***			
Further Training					(0.127)			
Lacking Motivation						0.018		
						(0.120)		
Many Absences							-0.093	
							(0.010)	
Staff Shortage								0.125
								(0.078)
N	7469	7469	7469	7469	7469	7469	7469	7469
Industry FE	✓	✓	✓	✓	✓	✓	✓	✓
Size	✓	✓	✓	✓	✓	✓	✓	✓

All columns present specification analogous to column 1 from Table 1.4, but with the main independent variable being an indicator for different types of labor problems. All indicators are defined based on firm response to the same module ("Staffing problems") of the 2014 IAB Establishment Panel. Standard errors are clustered on the industry level. (*) denotes significance at 10% level, (**) at 5% level and (***) at 1% level.

Table 1.17. Robustness Checks of Employment Changes Regression and Additional Results

	Y=%ΔEmployment (2005-2015)										Y=ΔRobotization X High Adoption	Y=ΔDigitization X High Adoption
	Basic Specification		Quartiles Dummies		Not Weighted By Employment		Control for Past Empl. Changes		Exclude Automotive		First Stage Regressions	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Robots X	-0.357*				-1.470***		-0.352*		-0.616			
Adoption>P(50)	(0.194)				(0.440)		(0.195)		(0.390)			
Digitization X	0.756				-0.303		0.701		0.611			
Adoption>P(50)	(0.464)				(0.919)		(0.461)		(0.489)			
Robots Abroad X		-0.635**				-2.446***		-0.630**		-0.801**	1.155***	-0.009
Adoption>P(50)		(0.308)				(0.662)		(0.309)		(0.392)	(0.271)	(0.0256)
Digitization Abroad X		0.179				-0.330		0.162		0.140	0.097	0.330***
Adoption>P(50)		(0.174)				(0.377)		(0.173)		(0.184)	(0.074)	(0.026)
Robots X			-0.635***									
Adoption>P(75)			(0.134)									
Digitization X			1.155*									
Adoption>P(75)			(0.702)									
Robots Abroad X				-0.816**								
Adoption>P(75)				(0.324)								
Digitization Abroad X				0.276								
Adoption>P(75)				(0.243)								
N	5275	5275	5275	5275	5275	5275	5275	5275	5275	5275	5275	5275
Area FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Industry FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
F-Stat											30.26	163.57

The Table present robustness checks for the main specification presented in column 3 of Table 1.7. Columns 1 and 2 present the basic specification from Table 1.7. Columns 3 and 4 use quartiles of adoption instead of above-median indicator (all quartile dummies are included, but only 4th quartile is presented, the value is relative to the first quartile). Columns 5 and 6 present basic specification with equal weights for every industry-area cell (as opposed to weighting by initial employment). Columns 7 and 8 include change in employment between 1995 and 2000 as a control. Columns 9 and 10 exclude automotive industry, which has the highest robot density. Columns 11 and 12 present first stage regressions for 2SLS specification (second stage is presented in column 6 of Table 1.7). All regressions are weighted using employment levels from 2005 (except for columns 7 and 8). Standard errors, reported in parentheses, are two-way clustered by area and industry. (*) denotes significance at 10% level, (**) at 5% level and (***) at 1% level.

Table 1.18. Effect of Financial Constraints - Additional Controls

	Y = Digitization and Automation Adoption								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Difficulties in	-0.619**	0.634**	-0.627**	-0.644**	-0.663**	-0.779**	-0.660**	-0.713+	-0.818+
Obtaining Credit	(0.298)	(0.294)	(0.302)	(0.285)	(0.289)	(0.330)	(0.294)	(0.436)	(0.502)
N	2217	2185	2199	2186	2210	1892	1784	1114	837
Industry FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Size	✓	✓	✓	✓	✓	✓	✓	✓	✓
Profitability	✓								✓
Part of Group		✓							✓
Establishment Age			✓						✓
Employment Growth				✓					✓
Public Firm					✓				✓
Foreign Owner						✓			✓
Professional Management							✓		✓
Startup (not Spinout)								✓	✓

All columns present specification analogous to column 1 from Table 1.8, but with additional controls. The controls include dummies for profitability assessment, being part of multi-establishment group; dummies for establishment age, the speed of employment growth in last 3 years, being a public firm, having a foreign owner, being managed by professional manager and being a startup (i.e. the establishment was started as startup, as opposed to being spun off from other existing establishment). Because of missing values in additional control variables the sample size varies between columns.

CHAPTER 2

Working More to Pay the Mortgage: Interest Rates and Labor Supply

I thank Bence Bardoczy, Efraim Benmelech, Anthony DeFusco, Benjamin Friedrich, David Matsa, Riccardo Marchingiglio, Brittany Lewis, Nicola Persico and seminar participants at Northwestern University for helpful comments on this project. I thank current and former staff of the Ministry of Entrepreneurship and Technology, in particular Paweł Maryniak, for their help in obtaining the data.

2.1. Introduction

There is an extensive evidence that household balance sheets play a significant role in the amplification of macroeconomic shocks (Mian et al., 2013; Mian and Sufi, 2014) and in the transmission of monetary policy (Bernanke and Gertler, 1995; Di Maggio et al., 2017). The literature studies several mechanisms which link financial positions of households, interest rates, and the real economy, but these mechanisms are usually related to households consumption behavior (Iacoviello, 2005; Calza et al., 2013; Garriga et al., 2017; Jappelli and Scognamiglio, 2018; Cloyne et al., 2020; La Cava et al., 2016; Hedlund et al., 2017; Flodén et al., 2017). This paper documents that changes in interest rates can affect the real economy also through the impact on labor supply of mortgage holders.

Using the data from Poland, where almost all mortgages are floating-rate, I show that an increase in mortgage payment driven by fluctuations in the reference rate (LIBOR/WIBOR), leads households to work and earn more. The magnitude of this effect is substantial: around 35% of the increase in payment is covered with the increase in income. The effect is higher for households with higher payment-to-income ratio, more pronounced for more flexible sources of income and driven by several mechanisms, including spousal labor supply, change of job or additional income from after-hours work.

The labor supply response can be interpreted through the consumption commitment model as in Chetty and Szeidl (2007). An increase in mortgage payment increases the price of housing consumption. Such change, which is likely to persist for some time, decreases household wealth and creates demand for additional liquid funds today. Housing consumption is costly to adjust – especially when financed with a mortgage – and hence, ignoring bankruptcy considerations, household needs to secure additional funds in some way. The household will typically decrease other consumption (consistent with extensive evidence, e.g. Di Maggio et al., 2017) and try to increase income. My results show that the income reaction can be sizable.

The mechanism I am identifying captures an intuitive idea that people work harder when their obligations are higher[1] and is consistent with multiple stylized facts from various contexts: more indebted households work more hours per week, delay retirement, are less likely to quit a job and are faster to find a new job when unemployed (see Appendix Figures 2.6, 2.7 and 2.8). It is worth noting, however, that the relationship between level of debt and household labor supply may be non-linear. While I document the positive effect of debt on labor supply in the sample of borrowers with low bankruptcy rates and recourse loans, debt overhang effect may suppress labor supply for borrowers who are close to bankruptcy, have negative equity and non-recourse loans (Bernstein (2018) and Donaldson et al. (2019)).

I conduct the analysis in the context of mortgage market in Poland, where 99.8% of mortgages are floating-rate loans. Mortgage interest rate is the sum of reference rate (typically 3-Month WIBOR or LIBOR - *Warsaw/London Inter-bank Offered Rate*) and a fixed markup. Every 3-6 months the mortgage rate is updated to reflect the current level of the reference rate. The variation in payments driven by changes in reference rate could be large: in my data interests payments change by up to 50% in the period of 2 years.

I analyze the evolution of mortgage payments and income in the period of 2005-2015 using the income tax data for the universe of Polish population. Using the mortgage interests tax deduction, I identify a near-universe of mortgage holders with loans originated between 2002

[1]The mechanism I am identifying can also be interpreted as a household finance analogy of Jensen (1986) free cash flows result. Jensen's mechanism operates in corporate finance world through reduced agency problem, while the mechanism in this paper operates through the increase in effort, but both suggest that debt leads to higher income or profits.

and 2006[2] and thus largely preexisting in the time period I analyze. I study how changes in the size of mortgage payment affect households labor income and other labor behavior. The within-household variation in mortgage payment is driven predominantly by changes in the reference rate, and I use the level of reference rate interacted with mortgage size to instrument for the actual payments.

I start by documenting the strong relationship between the level of the reference rate (which is an average of WIBOR and LIBOR) and the size of interests payments in my data. 1 percentage point change in the reference rate changes the typical yearly payment by around 700 PLN, which corresponds to 16% of the average interest payment. The time evolution of the average amount of interests paid closely follows the evolution of the reference rate, illustrating the strong relationship between reference rate and interest payments. The strength of this relationship is confirmed by the F-statistic of the first stage regression, which is of the order of 10^8.

The main result of the paper is the positive effect of debt payments on labor income, which I interpret as the the effect on labor supply. This effect is substantial - around 35% of the increase in mortgage payment is covered with the increased labor income. My basic specification regresses household's income on the size of the mortgage payment, controlling for household and year fixed effects as well as for fixed effects of age–previous-year-income bins, and age–income specific time trends. The results are similar when using only intensive variation in the mortgage size and when using both intensive and extensive margin variation with no-mortgage households being a control group. This basic specification is further enriched by using Instrumental Variable approach, which explicitly uses only reference-rate-driven variation in payments.

[2]But I observe interests payments made by these households in the entire analyzed period, i.e. 2005-2015.

I complement the main findings by exploring the heterogeneity of the income effect. The strength of the labor supply effect may depend on the size of mortgage and on the relative magnitude of adjustment costs for consumption and labor income. I document a clear pattern of heterogeneity with respect to payment-to-income ratio (PTI). The increase in income after the increase in mortgage payment is monotonic in PTI and varies from essentially zero for households with low PTI to over twice the size of average effect for those with large PTI. This heterogeneity pattern suggests that adjustment of labor supply is more important for relatively large mortgages, for which responding to higher payments through reduced consumption may be very costly or not feasible. The size of the effect varies also by income type. The relative increase of income from self-employment is 35% higher than the increase in wages, which likely reflects greater ability of the self-employed to benefit from adjusting labor supply. Nonetheless, the effect among wage-earners is also significant and sizable and I identify several mechanisms which shed light on the ways in which this adjustment takes place. The change in pension income is not significant, which can be thought of as a placebo check for the validity of the specification.

What mechanisms are responsible for the income increase? I construct several proxies for additional labor market activities and show that they contribute to the observed effect. First, households whose mortgage payments increase are more likely to change jobs and their new position is more likely to be in a different town than where they live. This suggests that when individuals face higher mortgage payments, they are more likely to take a higher-paying job which they might have previously rejected because of negative compensating differentials (e.g. long commute). Second, higher interests increase the probability that households receive a supplementary income from additional income-bearing gigs. Third, exploiting the fact that my data allows me to observe couples who file taxes jointly, I document a significant effect of spousal labor supply channel: households are more likely to become dual-earner household after their mortgage payment increases. This set of mechanisms is not exhaustive. While I do not observe effort and hours worked, I expect that their increase also contributes to the observed income effect, in particular for households with more flexible sources of income (self-employment, piece rate compensation).

The analysis includes several attempts to address the endogeneity concerns which arise in my setup. Because variation in mortgage payment is partially endogenous, due to possible prepayment and refinancing, I instrument payment size with the reference rate level interacted with a measure of mortgage size. The IV specification shows an effect which is similar to OLS estimate. My identifying assumption is that, conditional on individual's age and previous year household income (I control for the interaction of fixed effects and trends specific to age- and previous income bins), the effect of interest rates on household income is not systematically related to mortgage size except through the size of the payment. The key threat to validity of this assumption is that mortgage holders may have incomes which are differently sensitive to macroeconomic fluctuations and merely happen to be higher in the years with high interest rates. But if that were the case, we would expect high interest rates also leading to higher consumption and savings for mortgage holders. I show that the opposite is true. Using information on tax deductions for selected types of consumption and savings (charitable donations, private pension contributions, expenses on broadband internet), I show that they decrease following the increase in mortgage payment. While these measures do not represent overall household consumption, the results are consistent with the fact that consumption and savings adjustment is another way in which households respond to higher debt payments.

Related Literature. My findings contribute to household finance literature and the literature on the link between finance and labor. Several papers analyze the link between debt, interest rates, and consumption (Gross and Souleles, 2002; Agarwal et al., 2017; Kartashova and Zhou, 2020). Among them, Di Maggio et al. (2017) analyzes the effect of the mortgage payment decline due to ARM mortgages reset and shows that it leads to increased consumption of cars and to voluntary deleveraging. Under several assumptions, their estimates suggest that around 80% of the decrease in payment may be consumed. I complement their findings by showing that changes in the size of debt payments can also affect labor supply, at the same time also providing evidence for the consumption response. In my setting, labor income changes by around 35% of the change in the mortgage payment.

Some existing studies show that household debt can depress labor supply through debt overhang effect (Bernstein, 2018; Donaldson et al., 2019). My findings suggest that debt can have the opposite effect and increasing the value of debt obligations may lead to higher labor supply. In reality, both positive and negative impact of debt on labor supply can coexist and their relative importance depends on the presence of negative equity, strength of recourse laws, labor market conditions[3] and types of jobs under consideration.[4] Several other papers show evidence consistent with my findings. Fortin (1995) and Del Boca and Lusardi (2003) show that women are more likely to work when their household has a mortgage, while the study of Bednarzik et al. (2017) shows that indebted individuals return to work faster after job displacement. Rothstein and Rouse (2011) analyze student loans and show that higher debt leads students to choose

[3]In my setting, in the case of default the borrower remains liable for the portion of debt remaining after seizure and sale of the house. Moreover, throughout the analyzed period Polish economy was continuously growing and the labor market was healthy.

[4]Brown and Matsa (2016b) show that indebted households apply for more local jobs, but fewer positions outside of their commuting zone

higher-salary jobs, while Fos et al. (2017) show that student debt decreases probability of enrollment in graduate school.

Documenting that interest rates can affect labor supply of mortgage holders contributes to the literature on the transmission of monetary policy (Bernanke and Gertler, 1995; Kashyap and Stein, 2000). The labor supply channel is novel and is likely to be quantitatively important, especially in countries with high reliance on floating-rate or adjustable-rate mortgages. It counters the typical effects of monetary policy: when interests rates increase, the contractionary impulse transmitted through other channels is mitigated by the increase in labor supply. My findings also have implications for designing programs aimed at helping distressed borrowers, suggesting that optimal policy could involve directing resources to households with limited ability to increase income or structuring programs in ways which provide incentives to increase labor supply. The labor supply incentives should also be taken into consideration when designing bankruptcy laws or rules regarding the recourse, not only for mortgages but also e.g. for student loans. In addition, my findings have implications for the methods of risk assessment which banks and other institutions use when issuing the loan. The fact that labor income reaction is an important method of adjustment to tightened budget constraint suggests that ability to adjust income should be an important factor determining credit-worthiness of a potential borrower.

This paper contributes also to the literature on the relationship between consumption and labor income, especially in the presence of consumption commitments. While it is generally recognized that the link between consumption and labor supply can go in both directions (Heckman, 1974), the existing literature focuses on analyzing how income shocks affect consumption adjustment (Jappelli and Pistaferri, 2010).[5] I show the causal effect in the other direction:

[5] Chetty and Szeidl (2007) discuss how consumption commitments can explain the added worker effect, i.e. the labor supply response of the spouse to the loss of job by the primary earner.

shocks to the consumption prices[6] can affect labor supply and thus income. The income response is a mechanism through which households can smooth their consumption, contributing to widely documented "excess smoothness" (Blundell et al., 2016). This effect is most evident when analyzing change in prices for a category of expenditures which is large and has high adjustment costs. While mortgage payment is a prime example of such category[7], many other expenditures can have these characteristics (e.g. child care, medical bills). The approach taken in this paper can also be interpreted as studying elasticity of labor supply with respect to consumption prices. If the consumption truly cannot be changed, this is similar to studying the elasticity of labor supply with respect to wealth or unearned income (Imbens et al., 2001; Deshpande, 2016; Cesarini et al., 2017). One difference, however, is related to the persistence of the shock and its perception by households.

2.2. Data and Institutions

2.2.1. Data and Summary Statistics

I use a panel dataset with 2005-2015 income tax records for the universe of Polish population. For each individual that have filed a tax declaration in a given year, I observe their income from various sources (e.g. salary, pensions, self-employment); a set of characteristics such as sex, age or place of residence; and the value of claimed tax deductions. Filing the tax declaration is mandatory and the process is comparable to that in other countries; additional details are discussed in the Appendix. The data allows me to follow individuals over time and match

[6] I simplify the exposition by referring to debt payments as consumption. While technically debt payments are not consumption, the debt-financed purchases are. We can therefore think about increases in mortgage payments as increases in the cost of housing.

[7] Debt-financed consumption is likely to involve commitment by the very nature of debt, i.e. the fact that it is a way to pay for consumption which already took place. In the most stark example, student loan is a way of paying for consumption which cannot be adjusted, since an individual cannot go back in time and change its education.

married couples who are filing taxes jointly. The dataset was obtained from Polish Ministry of Entrepreneurship and Technology; according to my knowledge, this is the first paper which uses the entire population of this dataset.[8] The data is confidential and has been anonymized so that it is impossible to identify any single person; the person identifiers are synthetic and monetary values were modified by adding a small random noise component to mask the exact values.

My key variable of interest, which allows me to identify mortgage holders and observe their mortgage interest payments, is mortgage tax deduction. The deduction was introduced in 2002 and abolished in 2007,[9] but households who started deducting interests during that period keep the right to deduct them until the end of their mortgage contract (usually 25-30 years). Therefore, if a household originated a mortgage and started deducting interests e.g. in 2003, I am able to observe the amount deducted in the whole period of my data (and hence I identify them as mortgage holders). However, if a household originated a mortgage in 2007, they are not allowed to use the deduction and in my data I do not identify them as mortgage holders.

The group of mortgage holders which I analyze is, therefore, a near-universe of households who initiated a mortgage between 2002 and 2006 (while households did not have to use the deduction, there were no incentive not to do so); the remaining part of the universe of taxpayers is a control group. The control group contains households without a mortgage and those with a mortgage originated after 2006 or before 2002. In practice, the number of mortgages originated before 2002 is very limited (membership in building societies was more common way to finance real estate purchases in these earlier years) and popularity of renting is low. The majority of

[8]Kopczuk (2012), who analyzes the effects of business tax reform on income and tax revenues, is another paper which uses micro data from the same source. However, he only analyzes a sub-sample of all taxpayers.

[9]The official reason for abolishing the deduction is related to the incompatibility of the law with the rules of European Union (which Poland joined in 2004). However, the fact that the law was abolished instead of being just slightly modified suggests that budgetary reasons were an important motivation.

the control group are owners without mortgage, but a non-negligible part are households with mortgages originated after 2006. My estimates will be therefore biased towards zero because part of the control group consists of households who also have mortgages and are subject to the treatment I analyze. In practice, however, this problem should not be very severe since mortgage holders form less than 20% of the control group (see the statistics for the entire mortgage market in Poland in AMRON, 2015). At the same time, focusing on the subset of mortgages originated between 2002 and 2006 is convenient because at the time of large interest rate changes – which are the shocks which I want to exploit – all the mortgages are preexisting. I can therefore abstract from the problem of mortgage origination endogenously responding to the level of interest rate.

The deduction allows households to deduct all interests paid on their mortgage, irrespective of the level of the interest rate, if their initial mortgage size is below a threshold stipulated by the tax code. If the mortgage size is above the threshold, the household can deduct amount of interests paid multiplied by the ratio of the threshold to their mortgage size. The threshold varies over time, depending on the time of the first deduction, but the majority of mortgages do not exceed it. My data contains only the amount of interests deducted (I do not observe the mortgage size) and hence I cannot exactly determine whether a given mortgage exceeds the threshold. Based on auxiliary sources, however, I estimate that on average interests observed in my data correspond to around 90% of the true interests paid. The details of this estimate are discussed later, when I take this differences into account when interpreting the magnitudes of the effects.

I limit the sample to individuals who are observed in the entire analyzed period. To better tailor the control to the treatment group, I drop all people born before 1946 or after 1986, who were either too old or too young to be a potential mortgage holder in 2002-2006. My final data set is strongly balanced panel with 9.9 million individuals and over 100 million observations. There are over 160 thousands of individuals identified as mortgage holders but I drop those whose interests deductions time series seems incomplete, i.e. interests bounce back and forth between zero and positive value. This has minor effect on the data: there are 156 thousands of mortgage holders in the final sample.

Summary statistics for the main variables are presented in Table 2.1. Mortgage holders have two times higher income and are on average 6 years younger than the control group. Related to the age difference, they are less likely to receive pensions and more likely to be self-employed. Average household with a mortgage deducts 4376 PLN of interests per year. Not knowing the amount of principal paid every month, I am unable to compute the size of the total payment, but a reasonable estimate would imply that on average mortgage payment constitutes around 10% of household income. This average is relatively low, partially due to the fact that while nominal and real incomes were constantly growing after 2006, interests rates were lower in the second half of my sample.

2.2.2. Institutions: Mortgage and Labor Markets in Poland

A crucial feature of the Polish mortgage market is that the vast majority of mortgages - 99.8% as of 2016 – are floating-rate.[10] In a typical mortgage contract in Poland, the interest rate is defined as a reference rate – usually 3-month Warsaw Inter-bank Offer Rate, WIBOR, or 3-month LIBOR – plus a fixed markup. There is no initial period during which the rate is fixed.

[10]Strong dominance of floating-rate mortgages is not unique for Poland. Other countries where floating-rate mortgages are strongly dominant include Spain, Australia or Ireland and most countries in Europe have a large share of floating-rate mortgages (e.g. about half in the United Kingdom).

Table 2.1. Summary Statistics by Mortgage Status.

Statistics are calculated for the main sample used in the analysis. The sample contains all individuals born between 1946 and 1986 who have tax records for the entire 2005-2015 period. I drop individuals with seemingly incomplete mortgage interests information, i.e. those with more than one hole in the series of interests deductions (drops around 5% of mortgage holders). Family income is calculated as the sum of incomes of two individuals who file taxes jointly in a given year. Number of observations in each row is the same and given in the last row of the table, except for rows which condition on positive value (e.g. of wages or donations), where total number of observations is given next to variable name. For internet expenses top 1% of outliers was dropped, because of unrealistically high values most likely reflecting data error. For interests, only positive values were included (to exclude zero values for mortgage holders who started paying mortgage later than in 2005).

	Mortgage = 0		Mortgage = 1	
Variable	Mean	SD	Mean	SD
Gross Household Income	54004	62622	104271	88729
Wages	43459	55220	91414	86205
Business Profits	2724	30233	4783	20187
Pension	4924	12803	2962	11803
Share Self-Employed	15.5%	-	20.8%	-
Share Receiving Pension	20.6%	-	6.7%	-
Wages (I >0) (N = 84.8 ml)	51559	56568	97378	85647
Business Profits (I>0) (N = 10.0 ml)	27483	92406	32901	43340
Pension (I>0) (N = 20.5 ml)	23879	18498	30611	24361
Interests Paid	-	-	4376	3823
Donations	33.4	1096	91.2	2450
Expenses - Private Pension	6.37	191	16.5	310
Expenses - Internet	147	268	220	317
Donations (I >0) (N = 2.6 ml)	1284	6681	2407	12368
Expenses - Private Pension (I >0) (N = 0.2 ml)	3280	2840	3998	2718
Expenses - Internet (I >0) (N = 25.7 ml)	581	177	606	209
Year Born	1966.8	10.5	1971.7	8.3
Number of Individuals (tho.)	8 998 141	-	156 229	-
Number of Observations	98 644 025	-	1 713 939	-

While the reference rate changes every day, each mortgage contract specifies the frequency with which the interests rate is updated, usually once in 3-6 months. In addition, some banks may not change the rate if the reference rate changed only slightly. In general, however, the variability in reference rate leads to changes in monthly mortgage payments.

There are few other characteristics of Polish mortgage market which are important for interpreting my results. Mortgage's length can vary but most of the borrowers have 25-30 year contracts. Refinancing is rare because the main motivation to refinance – to benefit from a decrease in interest rates – is not relevant, as mortgage payments automatically incorporate changes in the interest rate. Only around 3-4% of mortgages are refinanced, usually when the situation of the borrower significantly changes. While there is no exact data on prepayment, anecdotally it is also a rare event. All mortgages in Poland are recourse loans which means that borrower still has to pay back the rest of the debt when the house is foreclosed and revenue from its sale is not enough to cover the total liability. Therefore there are no strategic bankruptcies and consumer bankruptcy, while possible, is rare in general. Around 2% of mortgages have delays in payments of more than 30 days, substantially less than 3.7% delinquent loans in the US.

Another characteristic of Polish mortgage market is its currency composition. Large fraction (the exact data for this time period is not available, extrapolation from later years suggests that the fraction is 25-50%) of mortgages are denominated in foreign currency, mostly Swiss Franc or Euro and use LIBOR as their reference rate. I do not observe currency in my data and hence I will treat all mortgages in the same way and use synthetic reference rate (an average of 3M WIBOR and 3M LIBOR CHF) to isolate the effects of changing interest rates. As shown in the Appendix Figure 2.4, the evolution of WIBOR and LIBOR is closely related and hence using their average yields similar results as using any of them individually.

Labor market institutions in Poland are similar to other European countries. The dominant type of contract is permanent employment which usually features 40-hour work week. The personal income tax rates have been 18% and 32% throughout most of the analyzed period, on top of social security contributions. In total, employees typically take home around 70% of their gross salary and 60% of the total cost to the employer. Unemployment rate throughout most of the analyzed period was between 9% and 12%. The unemployment was to large extent driven by rural areas (where mortgages are significantly less popular) and hence unemployment faced by my treatment group was lower. Importantly, the entire analyzed period was a period of economic growth and relatively healthy labor market. Poland was the only member of European Union which did not experience recession as an aftermath of the financial crisis. Due to healthy financial system, no construction boom in the previous years and large demand for infrastructural investments, every quarter in the analyzed period had positive GDP growth. While unemployment hit the lows in 2008 and it increased slightly afterwards, the change was small (from 9 to 10-11%) compared to other European countries (e.g. in Spain unemployment went up from around 10% to 20-25%). As a result, in the entire analyzed period both nominal and real incomes were growing, as illustrated in Figure 2.3 in the Appendix.

2.3. Research Design

My strategy exploits the within-household variation in the size of mortgage payment driven by interest rate fluctuations to analyze the impact of the size of payment on household's labor income and other outcomes. The main specification is as follows:

$$Y_{i,t} = \alpha \cdot (Interests_{i,t} = WIBOR/LIBOR_t X Exposure_i) + \sum_{t=2005}^{2015} Year_t + \mu_i + \beta X_{i,t} + \varepsilon_{i,t}$$

The main explanatory variable is the amount of interests paid by household i in year t. While interests are only one part of total mortgage payment, they capture the majority of non-deterministic variation in the payment (since capital payments are set in advance, except for foreign-currency denominated loans where principal payments are subject to exchange rate movements) and constitute often more than a half of the entire payment in the initial few years. In the basic panel regression, I directly include interests which I observe in the data. The variation in the size of mortgage payment is driven mostly by fluctuations in WIBOR or LIBOR, which are the reference rates for most mortgages. To isolate only that part of variation, and to disregard other more endogenous mechanisms such as prepayment, I instrument the size of interests paid with the level of reference rate – an average of WIBOR and LIBOR – multiplied by a an estimate of the mortgage size. My specification is therefore an instrumental variable panel estimation with fixed effects, year fixed effects, and additional time-varying controls. Conceptually, it studies how households with large mortgages react to changes in the interest rate, compared to household with smaller or no mortgage.

The basic specification controls for year fixed effects $Year_t$ and individual fixed effects. Preferred specification adds fixed effects for previous-year income bin, age, and the interaction of these two factors, as well as time trend specific to age-income group. Main outcome variable $Y_{i,t}$ is a measure of household income: the default is gross income but I also use additional measures such as wages, pension or business income. The economic mechanism I am trying to analyze suggest that appropriate specification involves variables in levels, not logarithms: I expect absolute income increase to be proportional to interest increase (because the extra income is supposed to cover an increase in interests) as opposed to relative change in income being proportional to change in interests. However, in robustness analysis I also include specification in logs. Income and interests payments are measured at family level since couples file taxes jointly and they claim only one tax deduction. The sample contains individual-level observations and my main results are obtained with weights of 0.5 for 2-person household observations. Standard errors are clustered on the household level.

I am trying to capture the following mechanism: in years in which household is paying high interests, I expect their income to be higher because household increases labor supply to cover additional expenses. In a perfectly flexible world we would expect that one additional unit of interests increases income by a fraction of unit. In practice, however, it is possible that due to labor market inflexibility, households must increase income by more than the increase in interests payments (e.g. to meet mortgage payments individual needs to work in a second job that does not have flexible hours). The effect may therefore be larger than the increase in the payment. On the other hand, it is entirely possible that households cover the entire increase in the mortgage payment via reduction of consumption or savings (or by additional borrowing), which means that the effect on income is zero.

2.3.1. Impact of Interest Rates on Mortgage Payments

Because all regressions include individual fixed effects, the entire identifying variation comes from time-variation in interests paid by the household. Conditional on paying interests in a given year, this variation reflects mostly fluctuations in the reference rate, usually WIBOR 3M (Warsaw Inter-bank Offer Rate for 3 months) or LIBOR 3M CHF. Reference rates are influenced by macroeconomic conditions such as National Bank of Poland interest rates, foreign exchange rates, and international money market situation. Of course these factors cannot be affected by individual household and hence from household perspective reference rate change is an exogenous shock to the size of their mortgage payment. While National Bank of Poland sets interest rates taking into account macroeconomic situation, Poland is a small open economy and the interests rates are to large extent driven by international conditions (see Appendix Figure 2.4 for comparison of payments for WIBOR-based mortgage in Polish zloty and LIBOR-based mortgage in Swiss franc). Overall, the changes in reference rates are unlikely to be directly related to the situation in Polish mortgage market; undoubtedly, however, they are related to overall economic conditions worldwide, which affects also labor market situation in Poland.

Figure 2.1 shows that there is a strong relationship between average level of interests in the data and the reference rate, which is an average of WIBOR and LIBOR rates. The relationship is not perfect for several reasons: interest payments consists of reference rate and fixed markup, banks adjust contract rate with some delays or do not adjust them at all if the changes are too small, some mortgages may use different reference rates and the exact split between WIBOR and LIBOR is unknown, etc. Nevertheless, the graph shows that the reference rate is an important driver of interests payments. Moreover, the magnitude of changes is large. Between 2006 and 2008 interests increased by over 40% and they went down again by almost 50% between 2008 and 2010. Table 2.2 confirms the relationship using regression analysis. There is a significant relationship between reference rate and the amount of interests payments, especially when we interact reference rate with the proxy for the size of the mortgage. Column 4 illustrates that the F-statistic in a first stage regression which uses reference rate to instrument for the size of mortgage payment is of the order of 10^8. Distribution of interests payments in 2008 and 2015 is shown in Appendix Figure 2.5, illustrating both the large changes in interest payments across years as well as substantial cross-sectional variation in exposure.

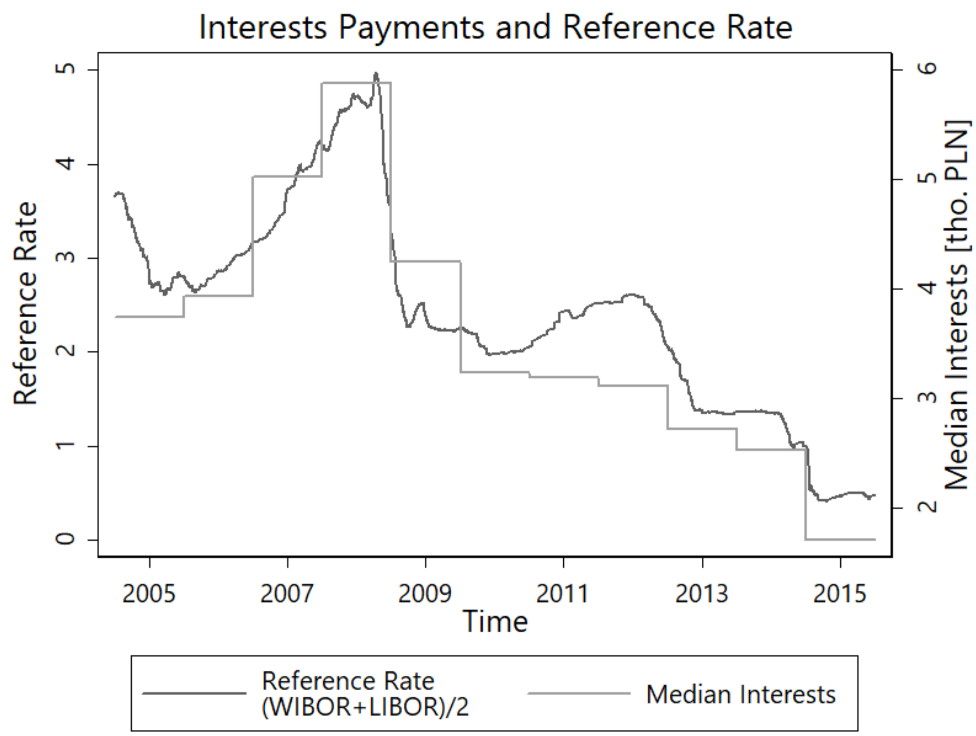

Figure 2.1. Median Interests Payments and Reference Rate.
The dark line is the median value of interests deducted by all mortgage holders using the deduction in a given year. Light-colored line is the average of 3-month WIBOR (Warsaw Inter-bank Offer Rate) and 3-month LIBOR CHF, which are the typical reference rates for mortgages in Poland. Absent borrower-level information on the reference rate used, I assume that equal share of loans are indexd to WIBOR and LIBOR.

Endogeneity of Time Variation in Interests Payments. While most of the variation in interests paid is caused by the movements of reference rate, it is possible that some variation is related to household decisions. For example, when a household member loses his job, the bank may sometimes allow for suspending debt payments for a couple of months. Alternatively, when the household increases its income e.g. because of getting better paid and more stable job, it may decide to refinance the mortgage and receive a lower rate thanks to being now less risky borrower. While all these factors seem to be rarely playing a role and there is no clear indication

Table 2.2. Interests Payments and Inter-bank Lending Rates

The dependent variable is the amount of interests deducted in a given year by the household. The main independent variables are the reference rate, which is an average of 3-month WIBOR and LIBOR CHF rates, and its interaction with a proxy for mortgage size, which is the average amount of interests paid by a given household in all years. All columns use strongly balanced panel with individuals fixed effects with observations weighted by the inverse of number of people in the household (1 or 2). In column 3, linear time trend is added. Column 4 presents the first stage of the 2SLS regression of income on interests payments (the main specification); it controls for individual age, previous income, year fixed effects and triple-interaction of these three factors. Standard errors are clustered on a household level. (***) denotes significance at 0.001 level.

	Interests Paid			
	(1)	(2)	(3)	(4)
Reference Rate	0.0135***			
	(0.0001)			
Reference Rate X Mortgage Size		0.335***	0.335***	0.331***
		(<0.001)	(0.001)	(0.001)
N	100 314 612	100 314 612	100 314 612	91 164 643
Individual FE	✓	✓	✓	✓
Trend			✓	
Controls				✓
F-Stat				$1.3 \cdot 10^8$

in which direction they would bias the result, they do cast some doubt on the exogeneity of the amount of interests.

To alleviate these concerns and explicitly use only the variation from the reference rate I use Instrumental Variable strategy, exploiting only variation driven by the changes in reference rate and fixed differences in exposure. The identifying assumption in that specification is that the influence of macroeconomic situation captured by evolution of WIBOR and LIBOR – after controlling for age, previous year's income, their interaction and time trends specific to these interactions – does not depend on the presence and size of the mortgage except through the size of mortgage payment. This assumption, while reasonable, is clearly not innocuous. What it precludes is that mortgage holders are more reactive to interest rate changes and that, within mortgage holders group, people with larger mortgages are more reactive than those with smaller mortgages. The rationale for the assumption is that there is no direct link between mortgage status and individual's job. Mortgage holders are hardly a very special group - it is common to have a mortgage and while there are demographic and income-level differences between mortgage holders and overall population (since purchasing a house is a decision usually made by young adults who are relatively well off), after controlling for these effects mortgage holders should not stand out as unique.

Nonetheless, to support the identifying assumption, I present results which suggests that macroeconomic considerations are not driving my results. I look at the patterns of consumption proxies and show that they move in the opposite direction than income. If the effects were due to different exposure to macroeconomic conditions, we would expect that income and consumption go in the same direction. However, if the observed effect is a response to higher mortgage payments, we should expect income and other consumption to go in the opposite directions, which indeed is the case (Table 2.6). I also look at the heterogeneity of the effect with respect to payment-to-income ratio and observe a clear monotone pattern (Figure 2.2), which would not to be expected if the effect was due to macroeconomic factors.

Endogeneity of Exposure. It is clear that households make a rational choice about the size of a mortgage and whether or not to get mortgage at all. When making this decision, the expectation about household's income stream is one of the key variables taken into account. I now discuss the implications of this potential endogeneity of the size of the exposure.

There are two elements of this potential concern. The first one is understanding whether observed effect can be generalized to the whole population. Suppose that we observe that mortgage holders increase income when their debt payment increases. If we randomly allocated debt in the entire population, should we observe the same reaction to the increased debt payments? Not necessarily, because people self-select into having a mortgage and may do so on the basis of their earnings upside potential. For example, a household considering taking a mortgage may rationally expect that payment may increase in the future, if the interest rates go up. If household members have limited possibilities of increasing earnings in case of payment increase, they may decide not to get the mortgage. This self-selection issue means that the results may not generalize to the whole population, even if we ignore demographic differences between mortgage holders and the rest of the population. The results obtained in this paper should be understood as effects observed in the population of mortgage holders. However, mortgage holders are large and important group and sheer size of mortgage balance sheets makes it important for the overall economy.

The second concern is related to the causal interpretation of the results. If the choice of mortgage and its size is driven by some unobserved characteristics, such as expectations of high earnings growth, is not mortgage just a proxy for these characteristics? While this concern can be to some extent valid, its severity is greatly reduced by features of my analysis. All mortgages in my data are preexisting at the time of shock: a typical scenario is a household who decides to originate a mortgage in 2004, makes the payments for 4 years and in 2008 sees large increase in their payment. Of course it is theoretically possible that when choosing their mortgage 4 years earlier, the household had expected the increase in payment in 4 years and increase in their income in 4 years. This simple expectation story, however, does not seem plausible and it is very unlikely that the timing of the increase in payments and incomes would be the same. It is more likely that a household had a sophisticated belief about their earnings are sensitive and to macroeconomic conditions. While I am not able to fully rule out this possibility, it does not seem very plausible given that the shock in interest rates in a small open economy does not always have clear relationship to conditions relevant for labor market. In my data, a large part of movement in interest rates is related to international financial markets but while most of Europe saw large GDP and employment drops during the last recession, Poland had no single quarter during which GDP decreased.

2.4. Results - Labor Supply

2.4.1. Main Results

The main result of the paper is presented in Table 2.3, which shows that an increase in interests paid by the household is associated with higher income. Columns 1 and 2 include only mortgage holders and use intensive-margin variation in the size of the mortgage payment. Column 3 includes uses the full sample, including individuals without a mortgage. Columns 2 and 3 include additional controls: interactions of age- and income-group fixed effects and trends specific to age-previous income bins. I control for individual's age using fixed effects with 1 year accuracy; in a similar way I control also for previous year family income, including indicators for 10 thousand zloty bins in the regression. I use 20 bins and incomes above 200 thousand zlotys are grouped in the last one. Because of computational considerations, in the interaction terms previous income is included through decile fixed effects, while age is included as 10-year age group. Dependent variable in all regressions is gross household income which is the sum of incomes of both spouses. Similarly, interests is the total amount of interests deducted by the household (since spouses are filing jointly, there is only one deduction; in fact even if they decided to file separately, they would be allowed to deduct interests only once). Standard errors are clustered on the household level.

Basic specification in column 1 shows that 1 zloty increase in interests deducted leads to income higher by 0.13 zloty. In columns 2 and 3, which are the preferred OLS specifications, 1 zloty increase in the size of payment is associated with income higher by 0.3-0.35 zloty.

In columns 4 and 5 instrumental variable approach is implemented. Interests are instrumented with reference rate multiplied by the proxy for the size of exposure, i.e. average amount of interests paid in the entire period. Column 4 includes only individuals with a mortgage, while column 5 includes the whole sample. Both specifications produce positive and significant estimates of the impact of interests payments on labor income. Estimate from column 4 is almost identical to the corresponding OLS estimate (column 2); estimate in column 5 is larger, but still similar to OLS estimate from column 3. The first stage regressions for IV specification (column 5) is presented in column 4 of Table 2.2.

Interpretation of Magnitudes – Deduction Cap and Income Taxes. To interpret the magnitudes, it is important to discuss the cap on the mortgage size, out of which interests can be deducted, as well as highlight that the analysis is performed for gross as opposed to net income. These factors have the opposite consequences for the relationship of the estimated coefficient to the true effect of debt on labor supply and their joint effect is likely to be small.

Consider a household with mortgage size of 400 thousand zlotys and an applicable cap of 200 thousand zlotys. Suppose that the yearly interests for such household amount to 20 thousand zlotys. When these interests go up by 10%, it is a 2 thousand zlotys increase in the real amount of interest paid. However, in my data, I only see half of this increase which will lead to overestimating the true effect of debt on labor supply. Unfortunately, the data is not rich enough to directly compute by how much the true effect is overestimated, but additional data sources allow me to perform back-of-envelope calculations to estimate the importance of this channel.

The applicable cap on the size of mortgage is approximately 191 thousand zlotys.[11] Based on Amron Sarfin Reports (AMRON, 2009), the average size of newly originated mortgage varied between 76 and 139 thousand zlotys between 2002 and 2006, with the volume-weighted average being equal to 107 thousand. I do not have more information about the distribution of mortgage size between 2002 and 2006, but I use the information for 2008 and 2009 to estimate the share of mortgages above the 191 thousand and their average value.[12] This allows me to estimate that on average, the interests reported in my data represent around 90% of interests paid and hence to obtain the true magnitudes we need to divide the coefficients by roughly 1.11.

To understand the effect of debt on labor supply one also needs to take into account that the dependent variable in my analysis is the gross income. By construction, part of the increase in interest payment is also automatically countered by the increased tax deduction. Each 1 zloty of extra interests decreases taxable income and hence the net increase in mortgage expenditures is typically 0.82 zloty.

In the end the cap on the mortgage size means that I overestimate the true effect while the tax shield means that I underestimate the effect on the net income. Back of envelope calculations suggest that the original coefficients should be divided by 0.82 x 1.11, which implies that the

[11] The cap varied between years and changed from 189 thousand zloty between 2002 and 2007 to over 326 thousand zloty in 2013 and later. To determine the size of the cap, the household first determines in which year their investment was completed (e.g. the house was built) and then uses the applicable value of the mortgage size limit. That is, even though the mortgage origination moment needs to be between 2002 and 2006 to use the deduction, the applicable limit depends not on the mortgage origination moment, but rather on the moment when the investment was completed. I assume that each investment is completed within 3 years of mortgage origination with uniform probability.

[12] For mortgages denominated in Polish zlotys, the growth in average mortgage size between 2006 and 2008 was around 60% and hence a mortgage of size 191 tho. in 2002-2006 would represent mortgage of size 305 tho. in 2008. In the first quarter of 2008 almost 22% of all mortgages were above 300 thousand zlotys with the average mortgage size in this tail being 602 thousand. I am going to assume that between 2002-2006 there were 22% of mortgages above the cap with the average value of the mortgage being 383 thousand ($191 \cdot \frac{602}{300}$). This implies that the true value of interests paid for 22% of mortgage holders was higher than interests deducted and on average they have paid 2 times more than they deducted.

Table 2.3. Interests Payments and Income

The dependent variable is the total gross household income in Polish zlotys. The main independent variable is the amount of mortgage interests deducted from the taxable income. The amount deducted may be smaller than the true amount paid for some borrowers which means the true effect on the gross income is lower - see the main text for the discussion of magnitudes. All columns use strongly balanced panel with individuals fixed effects with observations weighted by the inverse of number of people in the household (1 or 2). The data includes 9.1 ml individuals in 2005-2015 period. In column 2, age fixed effects represent individuals age in years and previous year income fixed effects represent previous year household income rounded to nearest 10 000 zl. Columns 3 and 4 present the specification from column 2 in which interests variable is instrumented with the interaction of WIBOR rate (the usual reference rate for the mortgage) with either an indicator for paying interests (column 3) or with this indicator multiplied by the proxy for mortgage size (average value of interests paid in the entire period, column 4). Standard errors are clustered on the household level. (***) denotes significance at 0.001 level.

	Gross Family Income				
	(1)	(2)	(3)	(4)	(5)
Interests	0.128***	0.347***	0.309***	0.351***	0.461***
	(0.030)	(0.028)	(0.027)	(0.052)	(0.028)
N	1 557 718	1 557 718	91 164 643	1 557 718	91 164 643
Sample	Mortgage Holders		All	Mrtg Hold	All
Individual FE	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓
Age X Prev. Income FE		✓	✓	✓	✓
Age X Prev. Inc Trends		✓	✓	✓	✓
IV				Ref Rate X Mrtg Size	Ref Rate X Mrtg Size
F-Stat				$4.3 \cdot 10^5$	$1.3 \cdot 10^8$

estimated increase in labor income lies between 14.3% and 50.5%, with the preferred estimates being 34%-38%.

2.4.2. Effects by the Type of Income

Tax declaration requires individuals to report their total income divided into different categories. While some of the income types are sparsely populated or rather obscure (e.g. other income), others can be used to get additional insights about the observed effect and test its plausibility. In Table 2.4 I analyze the effect of mortgage interests on wages, pensions and income from self-employment. Employment contracts may often be rigid: salary may not be directly related to effort and there may be no possibility for increasing the amount of hours worked (although, as demonstrated in Section 2.5, there are other ways in which employed persons can adjust their income). I expect that self-employed individuals have more opportunities to increase their labor supply and therefore the effect for business profits should be larger. Pensions, on the other hand, cannot be adjusted in the short term and hence provide a useful placebo check. In fact, the characteristics of pension system may lead to an effect of the opposite sign: if an individual collects pension and receives income from employment which exceeds certain threshold, his pension can be reduced (the reduction is not proportional and hence from the individual's perspective it still makes sense to work). This mechanism may be especially important in my sample (in which only persons younger than 65 are included) because most persons receiving pensions are in working-age and are part of some early-retirement scheme. Such recipients are, compared to a typical retiree, more likely to work and have their pension reduced.

To analyze these predictions I employ two approaches. First, I analyze the absolute change in income from different sources in the full sample. This approach, however, does not take into account the baseline values of income from different sources and conflates both intensive and extensive margin response. For example, while higher interests can encourage extra effort of self-employed individuals and thus increase business profits, they may also discourage households from taking the risk of starting a new business. Second approach, therefore, analyzes the relative changes in each income class and limits the sample to households with positive value of income from given source in the last year. Table 2.4 presents the results. While in absolute terms (columns 1-3) increase in wages is the dominant mechanism of the increase in income, this follows from the fact that wages are by far the most common source of income (see Table 2.1).

Table 2.4. Regressions of Interests Payments and Income from Different Sources.

All columns include individuals fixed effects as well as year, age, total previous income bin fixed effects (10 tho. bins) as well as binary variables for the previous year value of the level of dependent variable (wages, rounded to nearest 10 tho.; pensions and profits; rounded to nearest 5 tho.). Dependent variable is family gross wages or log wages (columns 1 and 4), pensions and log pensions (columns 2 and 5) and business profits and log profits (column 3 and 6). Main independent variable is the value of interests paid by the family in a given year expressed in thousands of zl (i.e., the value of interests divided by 1000, for readability of coefficients). Columns 1, 2 and 3 include the entire sample (and thus analyze both intensive and extensive margin effect), while columns 4, 5 and 6 only include those individuals whose family had non-zero income from given source in the previous year (focusing on intensive margin effect). (***) denotes significance at 0.001 level. Standard errors clustered on the household level are displayed in parentheses.

	Wages	Pensions	Profits	Log(Wages)	Log(Pensions)	Log(Profits)
	(1)	(2)	(3)	(4)	(5)	(6)
Interests	239.2***	-10.3***	30.4***	0.0078***	0.0007	0.0105***
/1000	(17.1)	(1.8)	(5.5)	(0.0002)	(0.0006)	(0.0011)
N	91 164 643	91 164 643	91 164 643	75 126 227	16 886 003	8 176 337
Controls	Individual FE, Year FE, Age FE, Previous Income FE, Previous Income from Given Source FE					

In relative terms (columns 4-6) the effect for business profits is 35% higher than the effect for wages. Nonetheless, we still see a positive and quite sizable response of the employment income. Coefficient for pensions is negative in absolute terms and positive, but insignificant in relative terms. The lack of intensive margin response confirms the lack of effect in the category of income which does not directly depend on individual's choice. The negative coefficient in column 2 suggests that extensive margin negative response, e.g. individuals delaying retirement or not applying for disability benefits, may contribute to lowering the pension income. Notice, however, that this decrease likely does not lower total income, since it is provoked by receiving high income from employment.

2.4.3. Heterogeneity with Respect to Debt-to-Income Ratio

Income response to mortgage payment changes may not be uniform in the whole population. In fact, one could intuitively expect that income adjustment should be coming from those households who have tighter budget constraint and hence it is more difficult for them to cover payment increase by reducing other consumption. A useful proxy for tightness of the budget constraint could be the relative size of the mortgage. Figure 2.2 presents coefficients from the regression in which interests variable is interacted with average interests-to-income ratio (the specification is otherwise the same as the basic specification in column 2 of Table 2.3). Notice that while the measure of mortgage size underestimates the true size of the mortgage (because it is based only on interests paid and does not include capital payment), it nonetheless allows for correct ordering of households from lowest to highest PTI ratio.

The coefficients show a clear montone pattern. Relationship of income and mortgage payment is stronger for those households, for whom mortgage payment constitutes larger share of their income. In fact, the effect for households with share of interests in income of less than 5% (which are around 25% of the sample) is not significantly different from zero. The effect is much higher for households with very large mortgages but, importantly, it is also positive and significant for a large part of the population with interests-to-income ratio around 10-15%, which includes majority of the sample.

The monotone pattern supports the proposed mechanism: households increase their labor supply because their budget constraint is tight and they need to cover increased mortgage payment to avoid costly bankruptcy. Intuitively, an increase in mortgage payment should matter less for high-income household with a small mortgage, compared to a household with large mortgage and relatively low income. For the first household the change in payments could

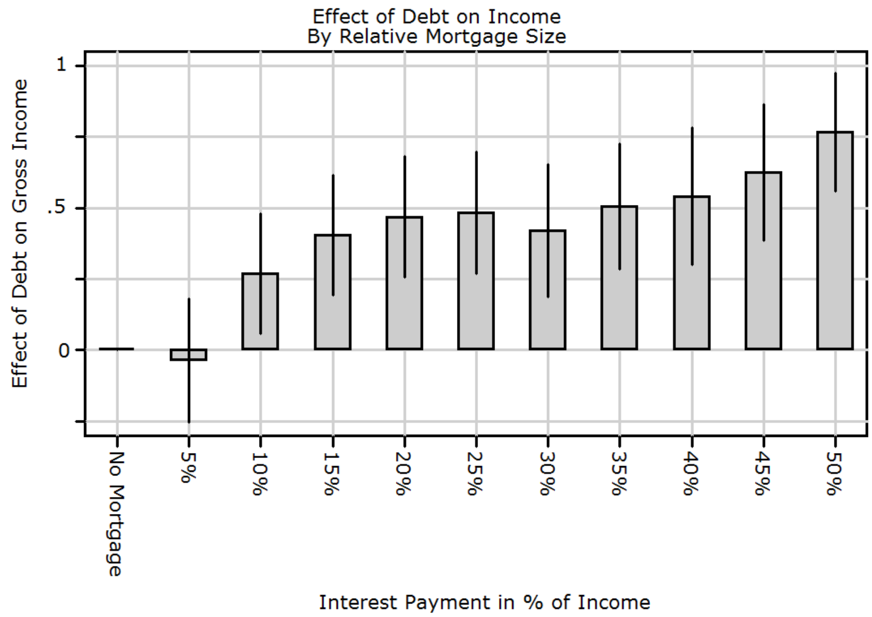

Figure 2.2. Heterogeneity of the Effect of Interest By Payment-to-Income Ratio.
The bars represent the coefficients from income-interests regression (analogous to the main specification) in which interests paid are interacted with the average size of interests paid divided by initial income ("mortgage size"; notice that the true mortgage size is higher because it also includes capital payment, although in the initial period of mortgage repayment, which is dominant in the data, interests consitute majority of the payment). The bars represent 95% confidence intervals. Only observations with mortgage size in the interval (0.05,0.5) are included; average mortgage size in this sample is 13% and the median is 10%.

be negligible and it can adjust other consumption more easily. The second household might have more troubles adjusting the consumption: given their low income and tight budget, their consumption may consists of higher share of necessities and be much harder to adjust.

Notice that while observed pattern of heterogeneity is not inconsistent with an alternative explanation based on different sensitivity to macroeconomic conditions, it is much less expected in this scenario. If we believe that mortgage holders are richer and therefore more sensitive to interest rates (e.g. high-level managers may have bonuses which vary greatly with business cycle, while factory-floor workers compensation is less volatile), we should expect to see an opposite pattern - relationship between income and payment changes should be increasing in income level and thus decreasing in PTI ratio (while richer households have also larger mortgages, the effect of higher income dominates and their PTI ratio is typically lower).

2.5. Mechanisms

What is the mechanism for observed income increase? Given that vast majority of gross income in my data are wages, part of the effect is likely driven by an increase in hours worked in the current job. Unfortunately, the tax data does not have information on hours worked, but the fact that income from self-employment increases more than earnings does lend indirect support to this mechanism. Nonetheless, on top of increased hours, there are several alternative channels which I am going to investigate: spousal labor supply, supplementary income, changing jobs and working in multiple jobs. Not all of these channels can be directly observed and hence I am unable to quantitatively asses their contribution to the income response. Instead, I provide qualitative evidence of their importance, which helps to understand the composition of the main effect. Some other channels, which I am not able to directly investigate, can also be playing a role. For example, an increase in interests can push individuals to take a tougher stance in bargaining with their employees and receive a wage increase.

2.5.1. Spousal Labor Supply

If only one spouse in a family works, the household can increase its income by having the other spouse enter the employment. Existing literature (Fortin, 1995; Del Boca and Lusardi, 2003) shows that women are more likely to work if their household has a mortgage. In a similar way, one could expect that when mortgage payment increases, non-working spouse (who usually is a woman in my data) enters the labor force. But women labor participation in Poland is relatively

Table 2.5. Mechanisms of the Increase in Income

All columns use panel data with individuals fixed effects, yearly time dummies and age dummies. Dependent variables are indicators for: single-earner household (columns 1), getting supplementary income (column 2, personal income on top of wages), changing job (column 3) and for working in the same town (column 4) or on multiple contracts (column 5). Main independent variable is value of interests paid by the family in given year expressed in thousands of zl. Dependent variables are expressed in percentage points. Sample in column 1 is limited to all married couples filling declarations jointly through the entire time period; sample in columns 3-5 is limited to individuals with employment contracts (for whom decreases in taxable income can be used to proxy for job change). (***) denotes significance at 0.001 level. Standard errors clustered on the household level are displayed in parentheses. Mean of dependent variables are presented in the bottom of the table.

	Single Earner (%)	Supplemental Income (%)	Change Job (%)	Local Worker (%)	More Contracts (%)
	(1)	(2)	(3)	(4)	(5)
Interests / 1000	-0.0888***	0.0178***	0.2300***	-0.0495***	-0.0075
	(0.0102)	(0.0042)	(0.0130)	(0.0137)	(0.0053)
N	42 296 262	91 164 643	48 034 080	48 034 080	48 034 080
Mean Dep. Var	31%	1.1%	12.6%	68%	2.7%
Controls	Individual FE, Year FE, Age FE				

high and, consistent with the mentioned literature, is even higher in households with mortgage. The margin for this adjustment can therefore be very limited. I investigate this channel in column 1 of Table 2.5. I limit the sample to couples who are filing jointly in the entire period (which allows me to disregard phenomena such as divorce etc.; on the other hand this may bias my estimates downwards because it is possible that when one spouse does not work, the other files taxes individually and hence this household is dropped from the sample). I define single earner household as a household where only one of two persons declares positive income from wages or personal income. This is to exclude categories in which income may be arbitrarily assigned to both spouses (for example, income from financial assets) and may not reflect true labor force participation.

The results indeed suggest that spousal labor supply adjustment is a significant contributor to the observed income effect. 1000 PLN increase in interest makes household 0.1 pp less likely to be single earner household.

2.5.2. Supplementary Income

Another possible way of adjustment is to perform additional, income-bearing gigs. An individual who works on a typical full time contract may decide to take extra after-hours jobs to supplement his main source of income. I am able to identify such activity to some extent by looking at the "personal activity" income category in the tax form. This category includes income obtained from activity performed personally but not subject to a formal employment contract (subject to labor regulations). An example of income which should be reported in this category could be a consulting fee which a professor - who receives a salary from a university - gets from an outside firm. The indicator does not perfectly capture additional gigs - this category of income can also include standard employment income (for tax purposes declared as subcontracting) and the extra gig can also occasionally be reported as normal employment. Nonetheless, it is a very useful proxy for capturing this type of adjustment.

I define supplemental income indicator as 1 when individual receives salary income and declares non-zero personal activity income in a given year; and 0 otherwise. The result are presented in column 2 of Table 2.5. The coefficient of interests is positive and significant. Its magnitude is small in absolute terms but quite large relative to the base levels: it indicates about 16% increase in the probability of receiving supplemental income (the baseline rate is 1%). The result shows that even individuals with rigid employment contract may increase their income by performing income-bearing activities outside of their normal workplace.

2.5.3. New and Another Job

Employed individual can increase his income by changing the job to a better-paid position or deciding to take a second job. The latter is self-explanatory, but the former may sound a bit surprising: of course it is true that better-paid job increases income, but in the context of adjusting to higher mortgage payments, why should we expect that an individual takes better-paid job only when the mortgage payment increases? If the better-paid job was available, why not to take it earlier? It is possible that jobs which pay more are worse in terms of non-monetary benefits: working conditions, required effort or location (Sorkin, 2017). An individual may be therefore hesitant to accept the better-paid position in normal times, but when extra money is badly needed, he can reconsider his choice and decide to take the job, sacrificing some non-pecuniary benefits for the sake of higher income.

I am using proxies for changing job and an indicator for working in multiple jobs based on the tax deduction available to every employee. Every worker with an employment contract can decrease its taxable income by an amount determined in the tax code. This amount is almost fixed and depends only on two factors: 1) whether an individual works in the same town in which he lives or not (higher deduction if needs to commute to different town); 2) whether individual works in a single job or in multiple jobs (higher deduction for multiple jobs). Based on this deduction I can define indicators for working locally and for working in multiple jobs. I define "change job" variable as any change in the local worker status: the variable takes value 1 if worker who worked in different town in year $t-1$ works locally in year t and when worker who worked locally in year $t-1$ works in different town in year t. In addition, I define a binary indicator for working locally. Change in these variables will capture either change of job or change of the place of residence but given the fact that I focus on mortgage holders, the second seems less likely.

The results are presented in columns 3, 4 and 5 of Table 2.5. When interests payments are high, workers are more likely to change jobs. Moreover, their new job is more likely to be outside of town of their residence. This is consistent with the compensating differential explanation if working locally is better because it reduces time needed for commute. Workers may therefore take more distant - and presumably better paid - jobs when their mortgage payment increases, sacrificing their short commute time to obtain higher income. There is no significant change in the probability of working on multiple contracts, possibly because this arrangement is rarely observed in general.

2.6. Consumption Response

An increase in mortgage payment can be covered in two ways: increasing income or decreasing other expenses. The main results confirm the importance of the first channel, but the magnitudes of the effect leave a lot of room for consumption adjustment. Unfortunately, the tax data does not contain a good measure of consumption and hence it is impossible to perform a comprehensive analysis of consumption response in this data set. Nevertheless, because the tax code allows for several deductions, I am able to create some consumption proxies. While they are very imperfect and by no means can be treated as good measures of consumption, they can provide some evidence for the existence of consumption reduction response.

Proxies include charitable donations, contributions to private pension funds and expenses on internet access. Deduction of internet access expenses is very popular, but every household can only use it twice and has to do it in 2 consecutive years. For that reason, I limit my sample to households which used the deduction in the previous and current year, since for them the decrease in reported expenses indeed captures the reduced expenditures, as opposed to starting or stopping deducting them. Donations and contributions to private pension funds are less popular (Poland has a public pension system and private pensions are in infancy) but can be deducted each year. The results for all three proxies are presented in Table 2.6. Notice that some other deductions are available but are very sparsely populated in the data, e.g. deduction for purchasing disability-related equipment such as wheelchairs. The results for these proxies were never significant and are not reported.

The results confirm the negative consumption response. When interests payments are high, households reduce charitable donations, private pension contributions and their expenses on the internet access (notice that all the coefficients are relative to other households; in absolute terms consumption may be increasing but at a slower pace). This is consistent with findings of Di Maggio et al. (2017) and highlights the importance of performing additional analysis, perhaps with consumption survey data, which would allow to analyze the consumption response.

The consumption response supports the labor supply adjustment explanation for the main results of the paper, as opposed to an alternative explanation based on differential sensitivity of mortgage holders to changes in interest rates. If we believe that the results are driven by unobservable characteristics of mortgage holders which make their income more sensitive to interest rate changes, we should expect positive response in consumption. Normally, when income increases, consumption also increases and hence if mortgage holders merely happen to earn more in years with higher interest rates, we should also see that they consume more. If households respond to the increase in payment, we should see the opposite: households need to

Table 2.6. Regressions of Interests Payments and Consumption Proxies.

All columns include individuals fixed effects as well as year, age and previous income bin fixed effects. Dependent variable is log of charitable donations (column 1), log of contributions to private pension account (IKZE, column 2) and log of expenses on internet access (column 3). The sample in column 3 is limited to households which have used internet deduction in the previous year and have non-zero deduction in the current year (because of 2 consecutive years limit for using the deduction). Main independent variable is value of interests paid by the household in a given year expressed in thousands of zl. (***) denotes significance at 0.001 level and (*) at 0.05 level. Standard errors clustered on the household level are displayed in parentheses.

	Log(Donations)	Log(Expenses - Private Pensions)	Log(Expenses - Internet)
	(1)	(2)	(3)
Interests/1000	-0.0013***	-0.0028***	-0.0012***
	(0.0003)	(0.0002)	(0.0003)
N	91 164 643	91 164 643	18 730 425
Controls	Individual FE, Year FE, Age FE, Previous Income FE		

cover higher payment and they do so by both increasing their income and decreasing their other consumption.

2.7. Discussion

In this paper I present an evidence that households cover increases in their mortgage payments with increased labor income. The effect is quantitatively sizable, as around 35% of the increase in payment is covered in that way, and hence might have important implications for understanding of the relationship between consumer debt and labor markets, as well as for monetary policy and debt relief policies.

The results I find are consistent with evidence in some of the existing papers. For example, Brown and Matsa (2016b) use border-discontinuity design to analyze job search behavior of mortgage holders depending on whether their loan is a recourse loan or not. While they focus on ability to relocate, their results also show that households living in a state with recourse are in general more active in their job search. Stronger commitment to pay back the debt causes households to increase their labor supply and search effort, consistent with my results.

More generally, various data sources show positive association between debt and labor market outcomes. While causality is difficult to establish, separate analyses of Survey of Consumer Finance, Current Population Survey and Health and Retirement Study all suggest that debt is associated with working harder. Figures 2.6, 2.7 and 2.8 in the Appendix contain some graphical evidence from these analyses. It is evident that the positive relationship between debt and supply of labor is widespread and can have widespread implications.

On the other hand, it should be noted that in the period of my analysis (2005-2015), Polish economy was constantly growing and labor market conditions were relatively good, contrary to other countries in the same period. The positive effect which I document may be therefore large because it was relatively easy for households to increase income. During the recession the change in income may be more modest.

There are several practical implications of my findings. First, they are of interests to those responsible for monetary policy. I establish that the contractionary effects of interest rate increase, which happen through several traditional channels, are accompanied by an increase in labor supply of mortgage holders. This effect is most sizable for countries with a high share of mortgages with floating or adjustable rates. Nonetheless, it may be manifested also for households with a fixed-rate mortgage through refinancing channel.

Second, policymakers are often interested in helping households with high debt burden. US Government's HAMP program is a recent example of such an action. My research shows that when designing such policies, it is worth taking into account household's potential to increase labor supply. Optimal policy would probably give higher benefits to households which cannot easily increase labor supply, e.g. because they are located in areas with high unemployment. It might be also optimal to introduce some built-in incentives which encourage labor supply adjustment.

More generally, my results have implications for design of bankruptcy and recourse laws. Debt has important implications for labor supply and hence personal bankruptcy law can have important implications for labor market. The more lenient is the law in terms of bankruptcy and recourse, the lower is the motivation for households to increase labor supply. On the one hand this can be treated as moral hazard costs. On the other hand, if increasing labor supply leads to misallocation of resources (e.g. because highly-educated people are forced to accept jobs below their qualifications), more lenient laws can improve allocative efficiency.

Finally, my results have implications for assessing credit worthiness of an individual. The potential to increase labor supply seems to be an important predictor of whether a household is a good or bad borrower. While this potential is not directly observable, it can be proxied with education, sector, labor market situation in local area and industry and current workload (i.e. persons with low current workload have more room for increasing labor supply).

Appendix

Additional Institutional Details

The data is based on the income tax declarations filled by individuals every year, usually in March-April of the following year. Employers send tax forms with income information to their employees and to the Tax Administration. Employees use the forms to fill tax declarations in which they include their total income, amount of taxes already withheld and deductions they would like to apply. These declarations are later send to the Tax Administration who processes them, making returns or requesting payments. Because the tax is normally withheld at source, a typical taxpayer receives a modest return. Even though taxpayers declare their incomes themselves, the Tax Administration has employer records to validate the declarations and hence the measure of income is highly reliable. The amount of tax evasion in Poland, as proxied by the size of shadow economy (see Medina and Schneider, 2017), is fairly similar to countries like Spain or Norway (share of shadow economy in GDP around 20%) and slightly larger compared to Germany or France (around 15%).

For the mortgage tax deduction, the amount of interest deducted is declared by the household based on the documentation received from the bank. While this documentation is not sent to the Tax Office with the tax declaration, it should be archived for at least 5 years for the purposes of potential tax audit. Only the amount of interests deducted is entered into tax declaration.

Appendix Figures and Tables

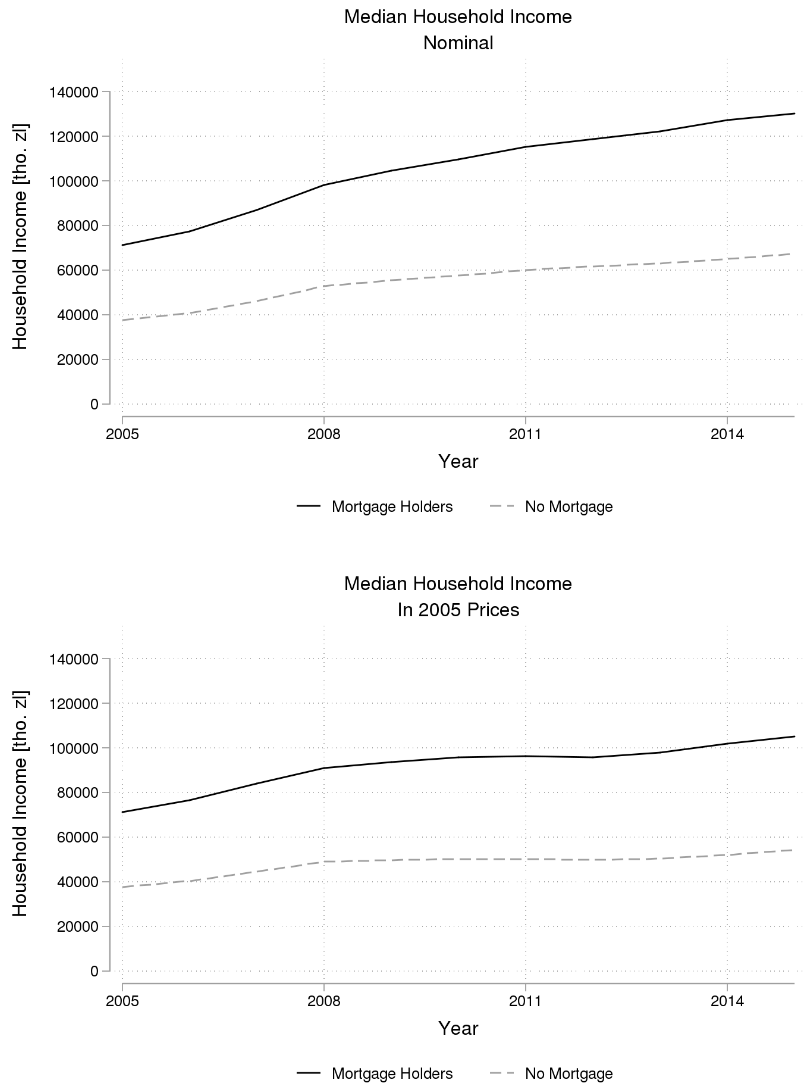

Figure 2.3. Median Household Income By Mortgage Status
The Figure shows the evolution o median household income by mortgage status. The upper panel shows nominal income, while the lower panel shows real income in 2005 prices (CPI used as the deflator). Both nominal and real incomes were consistently growing for the entire analyzed time period. Throughout the whole period mortgage holders were typically richer than households without a mortgage.

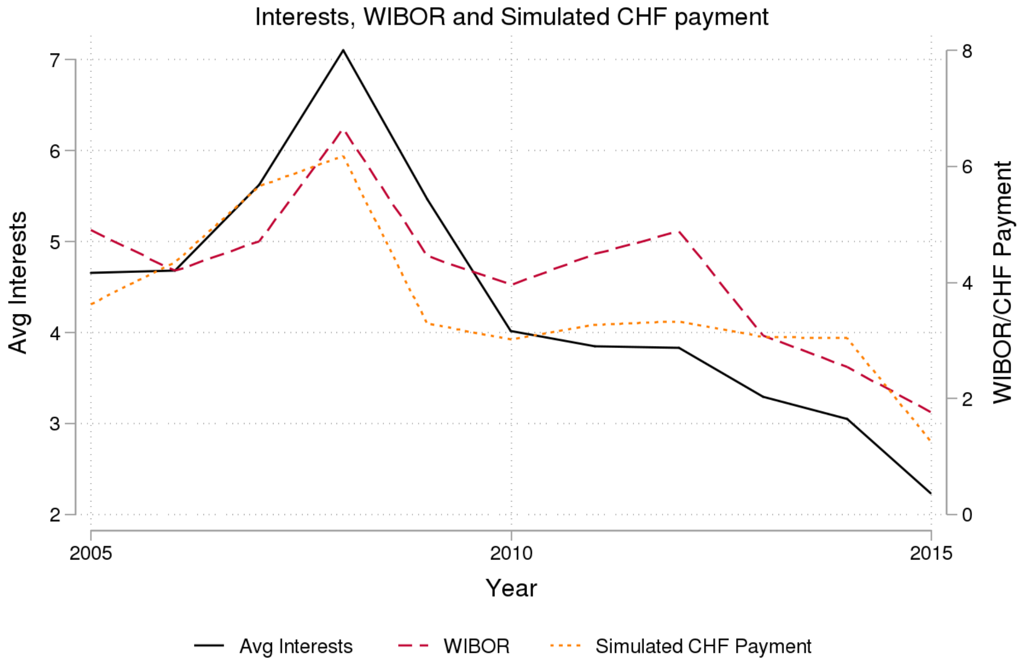

Figure 2.4. Average Mortgage Interests in the Data, WIBOR Rate and simulated LIBOR-based CHF payment.

Average interests rates are corrected by payment-schedule factor: since in typical mortgage with fixed monthly installment interests decrease in every month, I correct interests observed in the data by the scheduled factor calculated for mortgage originated for 25 years in 2004 with 3% markup and WIBOR 3M reference rate. The size of correction varies from 1.5% to 3%, depending on the year (notice that this mechanism leads to lower interests in subsequent periods but not to lower payment - interests part of the payment decreases but capital part increases accordingly). Estimated CHF-payment shows re-scaled hypothetical Swiss Franc denominated mortgage payments with interests based on LIBOR 3M. It incorporates both changes in LIBOR and in the CHF/PLN exchange rate.

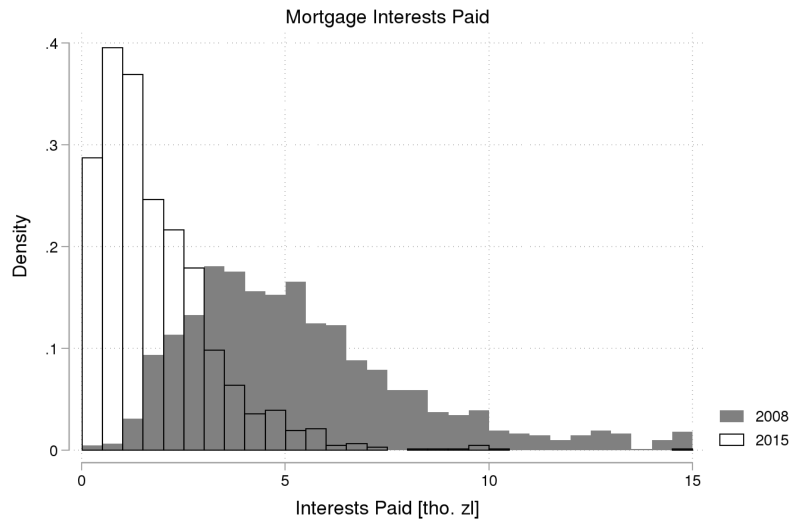

Figure 2.5. Distribution of Interests Paid in 2008 and 2015
The Figure shows the distribution of interests paid in 2008 (gray bars, year with highest interest rates) and 2015 (white bars, year with lower interest rates). The graph illustrates both the cross-sectional dispersion in the amount of interests paid, as well as the extent of changes in the size of interest payments over time. The value of interests is censored at 15 thousand zloty.

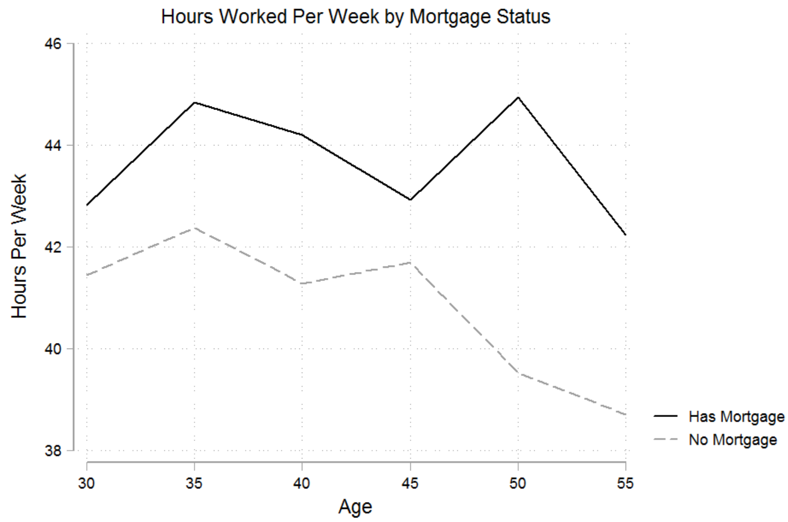

Figure 2.6. Hours Worked by Age and Mortgage Status (Survey of Consumer Finance 2016)

The data comes from Survey of Consumer Finance 2016 of United States Federal Reserve. Mortgage status is based on variable X723. Respondents with answer different than yes (1) and no (5) were dropped. Hours worked use variable X4110. Age, defined as the difference between 2016 and the year of birth, is rounded to the nearest multiple of five.

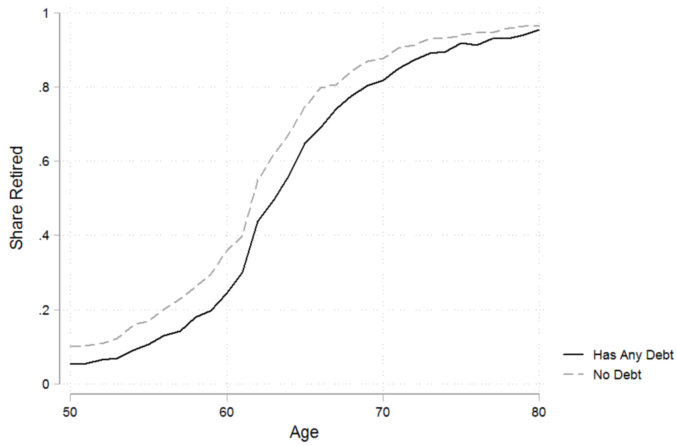

Figure 2.7. Share of Retirees by Age for Mortgage Holders and Non-Holders (Health and Retirement Study)

The data comes from RAND 2014 Health and Retirement Study longitudinal file. The sample includes all respondents between 50 and 80 years old. Retirement status is a binary indicator based on respondent's declaration if considers himself retired. If partly retired, the variable takes value 0.5. Respondents with answer "question irrelevant" were dropped. Has any debt is defined as total debt being above zero, where total debt is the sum of mortgage, other home loans, other debt and 2nd home mortgage. Age is rounded value of respondent age (agey_m variable).

Figure 2.8. Quitting and Unemployment Duration by Housing Status (Current Population Survey)

The data comes from IPUMS CPS data (Sarah Flood, Miriam King, Renae Rodgers, Steven Ruggles, and J. Robert Warren. Integrated Public Use Microdata Series, Current Population Survey: Version 6.0 [dataset]. Minneapolis, MN: IPUMS, 2018. https://doi.org/10.18128/D030.V6.0). The sample contains observations from years between 2010-2017. Individuals are included in the sample if in any of these years their mortgage status is not missing (variable spmmort, which also defines the three groups presented on the graph). Quitting and firing identified based on declared reason for being unemployed. Duration of unemployment is measured using variable durunemp. The bars show the share of all respondents who are unemployed and declare that they are other job loser (excludes layoffs). The line shows the average value of unemployment duration (in weeks) for respondents who are currently unemployed and were fired.

CHAPTER 3

Flexibility Costs of Debt: Danish Cartoons During the Cartoon Crisis

(joint with Benjamin Friedrich)

Coauthored with Benjamin Friedrich (Northwestern University, Kellogg School of Management, benjamin.friedrich@kellogg.northwestern.edu.). We thank Efraim Benmelech, David Matsa, Filippo Mezzanotti, and seminar participants at Aarhus University, Northwestern University, University of Tuebingen, University of Cologne, DIEW 2018, and GEA 2019 and MFA 2020for their comments and suggestions. We thank the Labor Market Dynamics Group (LMDG), Department of Economics and Business, Aarhus University, and, in particular, Henning Bunzel for invaluable support and making the data available. LMDG is a Dale T. Mortensen Visiting Niels Bohr professorship project sponsored by the Danish National Research Foundation.

3.1. Introduction

Holding a large amount of debt may limit firms' ability to pursue new projects as they come along and to flexibly adjust to changes in their environment. According to CFOs surveyed by Graham and Harvey (2001), these flexibility costs of debt are the most important factor shaping firms' capital structure decisions.[1] Similarly, Modigliani and Miller (1963) recognize "the need for preserving flexibility" as the main reason to limit firm's debt holdings. Flexibility is particularly valuable when the environment is uncertain and changing abruptly, which has been increasingly the case in recent years (Baker, Bloom, and Davis, 2016; Alfaro, Bloom, and Lin, 2018). Yet, obtaining comprehensive evidence on the importance of financial flexibility and on the mechanisms through which it affects firm operations is challenging.

The empirical challenges are largely due to measurement and identification problems. First, flexibility is to a large extent latent. Studying it empirically requires observing a situation in which a firm adjusts its operations to new opportunities or challenges.[2] Moreover, in order to credibly identify the role of flexibility in this adjustment, it is important to analyze an unexpected, well-defined event rather than gradual and predictable change. Finally, flexibility can be exercised through adjustments along multiple margins, such as investment, employment, innovation, and product market decisions. While some existing papers analyze the link between capital structure and one of these adjustment margins, e.g. labor hoarding in Sharpe (1994)

[1]Among CFOs, 59.4% say that financial flexibility is important or very important in affecting how they choose the appropriate amount of debt for their firms. This is the highest value among all 14 factors listed, including a firm's credit rating (57%), tax advantage (45%), bankruptcy costs (21%) or customers/suppliers comfort (19%). Flexibility is defined as "being able to pursue new projects as they come along".

[2]One can define measures of operational flexibility that can be computed ex-ante, e.g., operating leverage. Existing research shows that there is a link between a firm's financial and operational flexibility, both in theory (Dotan and Ravid, 1985) and in the data (Mandelker and Rhee, 1984; Peterson, 1994). These measures of flexibility, however, are often noisy and may be too narrow to show how well a firm can adjust to a new environment.

and Giroud and Mueller (2016), data limitations make it difficult to paint a complete picture of flexibility.

This paper aims to overcome these difficulties and provide comprehensive evidence on the link between capital structure and firms' flexibility to adapt to changes in their environment. We take advantage of a natural experiment in Denmark that led to a sudden and unexpected reduction of foreign demand for a small set of firms in an otherwise growing economy. In September 2005, a Danish newspaper published caricatures of the prophet Muhammad, which subsequently led to a widespread boycott of Danish products in Muslim countries. Danish firms' sales to these markets significantly declined and remained lower for at least one year. We study how the exposed firms reacted to the shock and how their reaction differed depending on firm's capital structure. We use a triple-difference design – taking a difference across time, exposure, and leverage – and analyze firms' decisions across several margins: product mix, export strategies, investment, employment, and outsourcing.

Our main findings show that low- and high-leverage firms adapted to the shock in very different ways. Low-leverage firms avoided a decrease in sales by introducing new products and entering new export markets, which was accompanied by an increase in investment. As a result, they did not reduce employment. In contrast, high-leverage firms innovated significantly less and did not redirect their products to new markets. Instead, compared to low leverage firms, they reduced their sales, employment, and investment. In addition, they reduced their operational risk by substituting employment with outsourcing. These results illustrate the flexibility costs of debt. In a changing environment, debt may limit firm's ability to undertake costly and risky projects, even if they are profitable. Financial obligations may also decrease the ability to take a wait-and-see approach and require making costly adjustments such as a reduction in operating leverage.

We start by establishing basic facts about firms' responses to the crisis. Low- and high-leverage firms had similar exposure to the boycott and they both significantly decreased their exports to Muslim countries. But low-leverage firms redirected their exports to other markets, significantly increasing their exports to non-Muslim destinations. For high-leverage firms, this increase was much smaller and not significantly different from zero. Consistent with the heterogeneity in export response, high-leverage firms experienced a significant decrease in total sales growth compared to low-leverage firms. In fact, sales did not significantly decline for low leverage firms, which suggests that domestic sales increased enough to cover losses caused by the boycott. While high-leverage firms would have certainly liked to increase their domestic sales and redirect exports, it is likely that financial constraints or risk considerations did not allow them to make costly investments in developing new markets and attracting new clients. Consistent with this interpretation, low-leverage firms in the treatment group increased their debt holdings, but high-leverage firms did not. The difference between the two groups is driven mainly by increasing short-term debt to suppliers and new long-term debt among low-leverage firms.

To provide more details on the differential response in exports and sales, we employ product-level data on exports and - for a subset of firms - product-specific sales. Doing so reveals that low-leverage firms were able to partially compensate for the boycott-induced losses by introducing new products into non-Muslim markets. These products were not only slight modifications of the existing product portfolio: we analyze the response using detailed product-destination data on export flows, and show that low-leverage firms increased the number of 6-, 4- and even

2-digit product categories in their exports. All these effects were close to zero for high-leverage firms. We then use annual sales data by product code for a sample of manufacturing firms to show that at least some of these newly exported products constituted product innovation and were not just products that were already sold domestically before. Again, we find no increase in product innovation for high-leverage firms. Redirecting sales and introducing new products likely required costly investment. Consistent with that, we observe a significant increase in investment among low-leverage firms, but no change for firms with high leverage. These findings suggest that financial obligations may limit firms' ability to innovate and invest in developing new products and export markets or, more generally, to flexibly adapt their product offerings to new market conditions.

On top of impeding the ability to actively adapt to new conditions, high leverage may also be an obstacle to taking a wait-and-see approach. Given the uncertainty about the length and severity of the boycott, firms may prefer to keep their scale unchanged and hoard labor. Financial obligations, however, may make this approach infeasible or too risky. We document that employment of high-leverage firms declined after the boycott, contrary to low-leverage firms. Consistent with the desire to increase operational flexibility, we show that high-leverage firms increased their use of outsourcing, which is more flexible than employment and hence reduces operating leverage. Using a subsample of manufacturing firms for which we have more detailed data on outsourcing activities, we show that firms started outsourcing high-skill tasks unrelated to their core activity, e.g., IT or marketing. Combined with the decreasing share of skilled workers, which we also document, this suggests that firms substituted employment of expensive workers performing tasks unrelated to the core activities with more flexible outsourcing arrangements. This strategy may provide cost savings in the short run, but may also imply loss of talent and firm-specific human capital, and therefore higher future recruiting and training costs.

Taken together, our results paint a complex picture of the relationship between capital structure and flexibility, documenting various mechanisms which link financial leverage and adaptation to shocks. The results suggest that high leverage limits firms' flexibility and forces them to take actions that trade off long-term optimality for improved liquidity today.[3] A potential concern with this interpretation is that leverage might proxy for other differences between firms that affect their responses. Our strategy cannot entirely rule out this concern but we present several results which support our interpretation. We first show that leverage is not simply a proxy for firm size, product variety, management quality, or differences across industries. Explicitly controlling for these differences across firms yields quantitatively very similar results across high- and low-leverage firms. Next, we analyze the role of an alternative channel related to liquidity: cash holdings. A higher cash buffer has qualitatively similar effects to having low leverage, suggesting that liquidity considerations are the driving force behind our results. However, given the estimated coefficients, cash holdings can explain only a small fraction of the observed variation, suggesting that financial leverage is the main driver of our results. Finally, we employ a "maturing-debt" approach similar to Almeida et al. (2011). We identify firms that will likely face a high share of debt maturing soon after the boycott because they acquired a large share of their long-term debt more than one year before the crisis. Our results show that these firms account for a large share of the employment decline and increase

[3] While we cannot explicitly rule out that some of the differences in adaptation are actually favorable to high-leverage firms (consistent with a disciplining role of debt as argued by Jensen, 1986), we find it unlikely that lack of product innovation, etc., is beneficial in the long term. Overall, our results suggest that reduced operational flexibility stemming from upcoming debt payments may be an important downside of high leverage, which is consistent with the significance of this factor for practitioners (Graham and Harvey, 2001).

in outsourcing, while firms with less debt maturing after the boycott drive the increase in investment, product innovation and foreign market entry. We also present a series of robustness checks using different assumptions and definitions of our key variables and find the same strong differences between high and low leverage firms.

Our key contribution is providing comprehensive evidence on the relationship between capital structure and firms' operational flexibility. While several existing papers document a link between leverage and particular dimensions of flexibility, for example, related to production technology (Reinartz and Schmid, 2016), pricing (D'Acunto, Liu, Pflueger, and Weber, 2018), or employment contracts (Simintzi, Vig, and Volpin, 2014; Serfling, 2016; Kuzmina, 2018), our setting allows us to enrich this evidence by analyzing simultaneous adjustments along several margins. Using a quasi-experimental setup, we directly measure firms' adjustments introduced in response to a sudden and unexpected shock, complementing other studies analyzing operating leverage (Novy-Marx, 2011; Chen, Harford, and Kamara, 2019) or other related proxies for flexibility (MacKay, 2003; Gu, Hackbarth, and Li, 2018).[4]

More broadly, we contribute to various strands of the literature on the real effects of financial decisions and constraints.[5] Our analysis of new products and exporting decisions links us to the literature on capital structure and the product market (Chevalier 1995b, Busse 2002, Fresard 2010). While Chevalier (1995a) shows that the financial situation of competitors may lead to strategic expansion in the supermarket industry, we show that a firm's own financial situation may influence the ability to expand in response to a negative demand shock. We show that part

[4]We also contribute to the literature on operating leverage by confirming the trade-off hypothesis (Van Horne, 1977) using a clean quasi-experimental setting. When expected value of future cash flows decreases (and hence financial leverage increases), high-leverage firms reduce their operating leverage by increasing reliance on outsourcing.

[5]A large body of work studies the relationship between firms' finances and a diverse set of real outcomes, including investment (Chava and Roberts, 2008), workplace safety (Cohn and Wardlaw, 2016), and product failures (Kini, Shenoy, and Subramaniam, 2017).

of the product-market response is due to increased innovative activity, which also links us to the literature on finance and innovation (see Kerr and Nanda (2015) for a review). While most of this literature measures innovation with the overall number of patents,[6] we contribute to a recent strand that focuses on product innovation (Krieger, Li, and Papanikolaou, 2018; Granja and Moreira, 2020).

Several recent papers analyze the complex relationship between firms' finances and labor. Chodorow-Reich (2014) shows that limited access to credit may adversely affect employment; other papers show that firm finances may affect various personnel considerations, including the pool of available employees (Brown and Matsa, 2016a); worker turnover (Baghai, Silva, Thell, and Vig, 2020); and firing decisions (Caggese, Cuñat, and Metzger, 2019). Giroud and Mueller (2016) show that debt may limit firms' ability to hoard labor. We use an empirical strategy similar to theirs, comparing high and low leverage firms subject to the same economic shock to disentangle financial and economic distress (which is a traditional challenge in the literature, see Asquith, Gertner, and Scharfstein 1994; Andrade and Kaplan 1998). However, while their goal is to highlight the role of firms' leverage in the employment decline during the financial crisis, we aim to study various mechanisms through which financial flexibility affects firms' adjustment to a changing environment. To this end, we analyze a highly targeted and unexpected shock that did not significantly affect the situation of the financial sector and overall economic conditions. Our results enrich the evidence on the impact of debt on employment by showing that labor hoarding among low leverage firms may be facilitated by innovative activity and product market responses, while the employment decline for high leverage firms may be

[6]See, e.g., work on the relationship between patenting and bank regulation (Amore et al. 2013, Chava et al. 2013, Cornaggia et al. 2015), M&A activity (Bena and Li 2014), lending relationships (Hombert and Matray 2017), and public listing (Acharya and Xu 2017) .

accompanied by an increase in outsourcing. This finding also contributes to a small literature on capital structure and outsourcing decisions. Empirical work by Moon and Phillips (2020) and the theoretical model of Kanatas and Qi (2016) argue that firms that outsource have lower financial leverage to protect relationship-specific investments. Our paper suggests a second mechanism that works in the opposite direction: firms with high leverage choose to outsource because they want to increase operational flexibility.

Finally, our analysis of a foreign demand shock also relates to the new and growing literature on the relationship between international trade and corporate finance (see Foley and Manova (2015) for a review). Access to financial capital affects both the extensive margin of entry into foreign destination markets and the intensive margin of the export value to each market (Manova, 2013). While internal capital markets can help multinationals circumvent financial constraints (Desai, Foley, and Forbes, 2008), availability of trade credit from banks substantially influences exporting activity (Amiti and Weinstein, 2011). Paravisini, Rappoport, Schnabl, and Wolfenzon (2015) find that bank credit shocks in particular affect the intensive margin of exports within product-destination. Our analysis shows that capital structure also affects the flexibility to enter new markets with new products in response to an adverse demand shock in a healthy banking environment.

The remainder of this paper is organized as follows. Section 3.2 provides details about the "Cartoon Crisis" and exposed firms. In Section 3.3 we describe our econometric approach and the data. We present our empirical findings in Section 3.4. In Section 3.5 we discuss mechanisms underlying the main results, while section 3.6 provides robustness analysis. . Section 3.7 concludes.

3.2. The Cartoon Crisis

In this section, we describe the timeline of events that led to the Cartoon Crisis and the consequences for Danish exporters across different industries. In particular, we analyze the duration of the boycott, extent of export reduction, and persistence of adverse effects after the end of the official boycott.

3.2.1. Timeline of Events

Denmark's largest newspaper, *Jyllands-Posten,* published 12 cartoons of the prophet Muhammad on September 30, 2005. According to the newspaper's accompanying article, the cartoons were a statement in favor of freedom of expression, in response to the self-censorship of Danish artists regarding illustrations in a recently published book about the life of Muhammad.

The cartoons first led to public protests among Danish Muslims that received no formal response from *Jyllands-Posten* or the Danish government. As a consequence, a group of Danish Muslims contacted ambassadors of several Muslim countries to Denmark for help in disseminating information about the cartoons in the Muslim world. The group was successful in placing the cartoons on the agenda of the Organization of Islamic Countries conference in Mecca in December 2005. This event set in motion widespread media coverage and political debate in Muslim countries. By the end of January 2006, Saudi Arabia and Kuwait were the first countries to declare an official boycott of Danish products. These announcements were followed by more violent protests at Danish embassies in Syria, Iran, Pakistan, and other countries. At the same time, the boycott quickly spread to Muslim countries around the world.[7]

[7] For a detailed time line of events, see Jensen (2008).

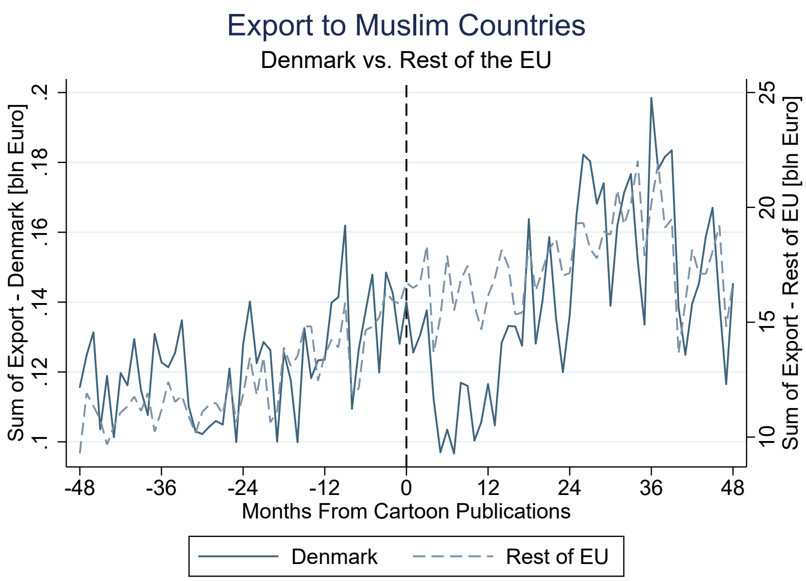

Figure 3.1. Exports to Muslim Countries for Denmark and the Rest of the European Union

The figure shows total monthly exports from Denmark and from the rest of the European union to countries with at least 50% Muslim population. The X axis defines time (in months) relative to September 2005, the time of cartoons publication. The data comes from Eurostat.

3.2.2. Consequences of the Boycott for Danish Exporters

Ex ante, firms had to build expectations about the duration of the shock to make the necessary adjustments. To illustrate the actual timing and duration of the boycott, Figure 3.1 compares total monthly exports from Denmark and from the rest of the EU to countries with at least 50% Muslim population.[8] The horizontal axis defines time (in months) relative to September 2005

[8]Consistent with the findings by Michaels and Zhi (2010), a deterioration in attitudes towards Danish products led to a substantial reduction in Danish exports even to Muslim countries that did not declare an official boycott. Muslim population shares follow a report by the Pew Research Center (2009) based on national census data from the years 2000-2006 and the World Religion Database using Muslim population estimates for the year 2005. The countries are United Arab Emirates, Afghanistan, Albania, Algeria, Azerbaijan, Bangladesh, Burkina Faso, Bahrain, Brunei, Djibouti, Egypt, Western Sahara, Gambia, Guinea, Indonesia, Iraq, Iran, Jordan, Kyrgyzstan, Comoros, Kuwait, Kazakhstan, Lebanon, Libya, Morocco, Mali, Mauritania, Maldives, Malaysia, Niger, Nigeria, Oman, Pakistan, Palestine, Qatar, Saudi Arabia, Sudan, Sierra Leone, Senegal, Somalia, Syria, Chad, Tajikistan, Tunisia, Turkey, Uzbekistan, Yemen, and Mayotte. In Table 3.11 we employ an alternative definition of the treatment group that only includes Arab countries.

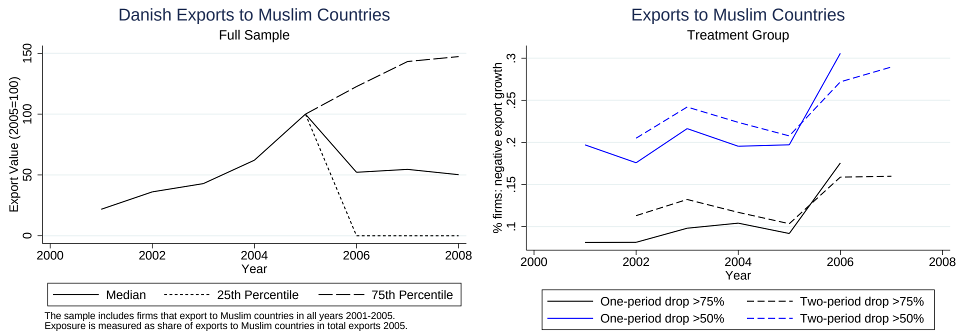

Figure 3.2. Exports to Muslim Countries for Danish Firms
The figure shows firm-level evolution of exports to Muslim countries over time. In the left panel, the value of firm's exports in 2005 is normalized to 100 and the evolution of 25th, 50th and 75th percentile of firms based on export value in 2007 is presented. In the right panel, solid lines show share of firms for whom export in a given year is 50% (blue line) or 75% (black line) lower than last year. The dashed lines show shares of firms for whom export in a given year is 50% or 75% lower than export 2 years ago.

when the cartoons were first published. As the time line illustrates, the boycott started with a delay of several months because the dissemination in the Muslim world took a considerable amount of time. The figure shows that while exports from other European countries continued to grow over this period, Danish exports experienced a sudden drop and remained at 25% below their previous level for more than one year. Danish exports largely recovered in 2007 at the aggregate level.

The aggregate time series suggests a temporary shock with a full recovery by mid-2007 compared to the export volume from other EU countries. But this aggregate time series hides important heterogeneity in the persistence of the shock across Danish exporters. The left panel of Figure 3.2 illustrates the time series of export volume to Muslim destinations among firms that exported to these markets before the boycott. We normalize their export value to Muslim countries in 2005 to 100 to illustrate the average drop in 2006 and the dispersion in outcomes over the post-boycott period. Specifically, the figure shows the interquartile range of export values across exposed firms after 2005. The median firm was unable to reach its 2005 export volume to Muslim countries again in 2007, and the bottom quartile of exposed firms remained at zero export volume to Muslim countries in 2007. The persistent drop is observable even if we restrict attention to firms who had very stable and successful business activities in Muslim countries over 2001-2005 (see Figure 3.11 in the Appendix).

Because trade volumes at the firm level can be quite volatile even without a boycott, the right panel of Figure 3.2 illustrates the share of firms with negative export growth over time. The sample in each year comprises firms who previously exported to Muslim countries. The solid blue line shows that 30% of previous exporters to Muslim countries experienced a drop by more than half to these destinations in 2006. This share is 50% (10 percentage points) larger than in the years before the crisis. More importantly, the share of firms experiencing a persistent drop in exports over a two-year period also increased by a similar margin and remained high in 2007 (dashed blue line). This substantial increase in the share of firms experiencing export reductions compared to normal times is even stronger when considering larger declines of 75% or more (black solid and dashed lines).

These results emphasize that firms experienced a large reduction in exports to Muslim countries on average, that the recovery was very heterogeneous across firms, and that the shock persisted for a substantial share of exporters.[9] The aggregate recovery in Figure 3.1 is driven by high-growth firms and by new entrants into these markets.

The left panel of Figure 3.3 illustrates the share of firms by industry that exported to Muslim countries before the boycott. Exposed firms constitute a substantial fraction of exporters in a large set of different industries, ranging from consumer products such as textiles, food, and furniture, to heavy machinery and equipment. Moreover, the right panel of Figure 3.3 reports

[9]Appendix Figure 3.11 provides additional details on firm export responses.

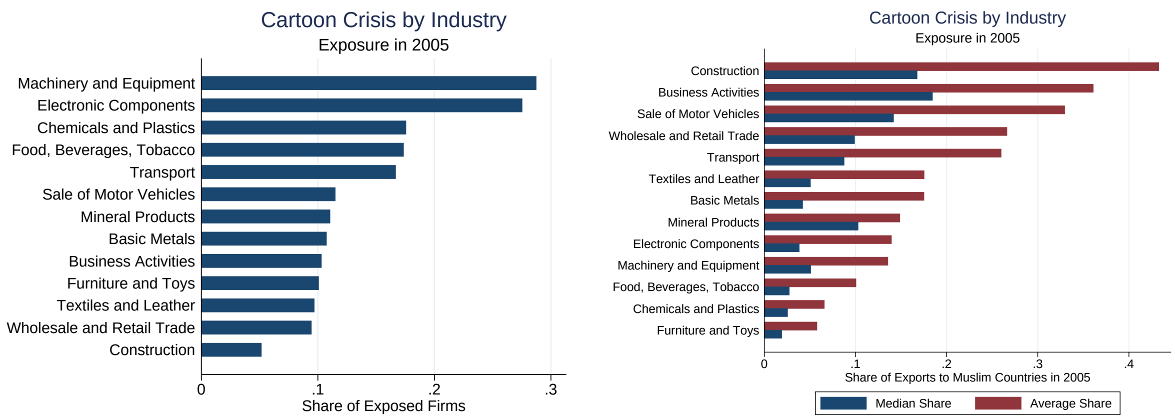

Figure 3.3. Boycott across Industries
The left panel shows the extensive margin of exposure: share of firms who were exposed to Muslim countries by industry. Exposure is defined as having at least 0.5% of total exports sold to Muslim countries. The right panel shows intensive margin of exposure: mean and median share of exports to Muslim countries among exposed firms by industry.

the average and median share of exports in these destination markets among exposed firms. There was considerable variation in the importance of Muslim destinations not only across sectors but also across firms within industries. The difference between average and median shares by industry indicates a small share of firms with high exposure in each industry, which we will characterize further below.[10]

3.3. Data and Empirical Strategy

In this section, we discuss our empirical strategy and introduce the data and main definitions.

[10]We provide additional details on the change in log exports to Muslim countries for 2005-2006 by industry in Appendix Figure 3.10. Construction and consumer products experienced the largest drop in exports. All but two industries saw a substantial reduction in exports to Muslim countries.

3.3.1. Data

An important advantage of our empirical context is that we can combine a variety of administrative data sources for the universe of firms and workers in Denmark, including firm-level trade data, financial statements of firms, and employer–employee matched data over the period 2001–2006. In addition, for a sample of firms, we are able to add detailed information on outsourcing and input expenditure, as well as product-level sales. We discuss these data sources in detail below.

The first data source is the Danish Foreign Trade Statistics Register (UHDI), which provides annual firm-level trade value by product-destination pairs at the CN-8-digit product level. Importantly, any trade flows with countries outside the EU are precisely measured by the customs authority (Extrastat). For trade with EU member states, Danish firms only have to declare exports above a threshold of approximately $250,000 per year (Intrastat). Thus, we do not include firms selling small quantities only to destinations within the EU in the sample of exporters.

Second, we add financial statements of firms from the Accounting Register (FIRE) and additional information on founding date, sales, employment, industry, and firm exit from the Danish Business Register (FIRM). The accounting data provides balance sheets and profit and loss statements with information about short-term and long-term liabilities, assets, investment, and input costs, in particular for labor services. A smaller subset of manufacturing firms also provides more detailed responses on purchases of goods and services, with a specific section of the survey listing expenditures on outsourcing across different tasks, such as transportation, accounting, consulting, catering, IT, and marketing. Moreover, a similar dataset collects information on sales by product CN-8-digit code for all manufacturing firms with at least 10 employees.[11] These data will be valuable to identify product innovation among these firms.

Finally, we use firm identifiers from the Firm-Integrated Database for Labor Market Research (FIDA) to match firm-level data with worker-level information from the Danish Integrated Database for Labor Market Research (IDA). IDA covers the universe of firms and workers in Denmark over 1980–2011. The data contain information about primary employment in November each year, including plant and firm identifiers, and detailed worker characteristics such as gender, age, education, experience, tenure, hourly wages, and annual earnings.

3.3.2. Definitions and Sample Descriptives

The main sample uses data on all private Danish firms that were exporters during 2003-2005. This yields a panel of about 15,000 firms. Among those, almost 2,000 firms were exporting to Muslim countries and hence form our treatment group.

We capture the differences in firms' capital structure by calculating the book leverage, defined as the ratio of total liabilities to total assets. Its distribution in treatment and control groups is presented in Figure 3.4. The two groups have similar distributions of leverage and, if anything, leverage is slightly higher for the control group. Figure 3.8 in the Appendix presents distributions for firms within the treatment group with low and high exposure to the boycott, and they track each other very closely. Our main measure of leverage is an indicator variable for high-leverage firms, which is defined as 1 if a firm's leverage is above the median. While we used country-wide median in the basic specification, Section 3.6 demonstrates that the results are robust to defining the median by industry and to using the continuous measure of leverage

[11] Smeets and Warzynski (2013) provide more details about this dataset and use it to measure productivity of multi-product firms.

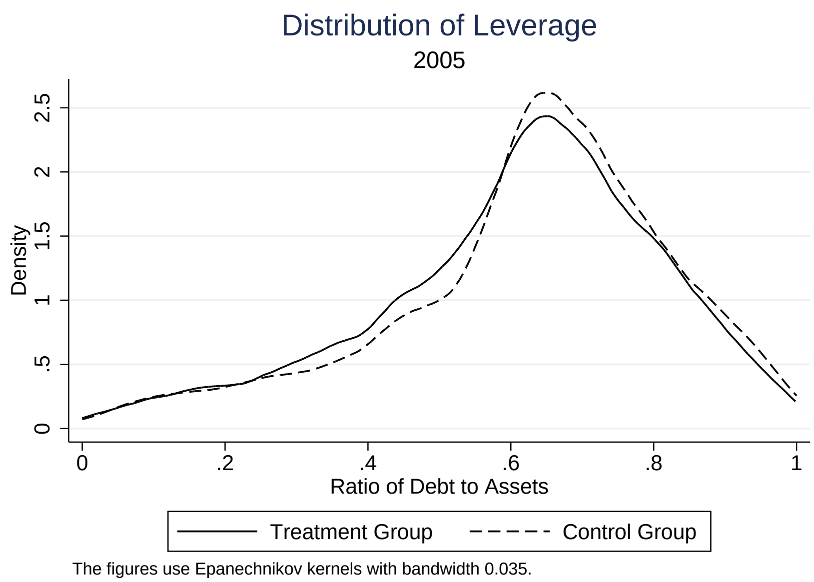

Figure 3.4. Distribution of Leverage for Firms Exposed and Not Exposed to the Boycott
The figure shows the distribution of total leverage (defined as the ratio of total liabilities to total assets), separately for firms exposed to the boycott (treatment group, firms with at least 0.5% of exports sold to Muslim countries) and those not exposed (control group). The distribution is smoothed using Epanechnikov kernel with bandwidth 0.035. Observations with leverage above 1 are censored.

As documented in Figure 3.4, the typical level of book leverage in our sample is relatively high, which several factors can explain. First, firms in Denmark - similar to other countries in continental Europe - are more bank-dependent than U.S. firms and hence have higher leverage. Second, we use only firms involved in international trade, whose levels of debt are higher. And third, our debt measure includes all liabilities and hence is significantly higher than the ratio excluding debt to suppliers. In Section 3.6 we provide alternative results, which are consistent with results from our main specification, using measures of financial leverage that exclude debt to suppliers.

We measure exposure to the Cartoon Crisis based on exports to Muslim countries in 2005, before the boycott. Given the widespread rejection of Danish products even without formal declaration of a boycott in many countries, we chose a broad definition of Muslim countries, including all countries with at least 50% Muslim population in 2005. We consider any firm with at least 0.5% of its exports in these markets as exposed to the shock.[12] The control group consists of all firms in the sample without any exports to Muslim countries in 2005.

Figure 3.5 shows the distribution of exposure to the boycott, defined as share of exports to Muslim countries before the start of the boycott, for low- and high-leverage firms. A large group of firms has low exposure that does not exceed 10% of exports. Yet, among both low- and high-leverage firms there is a sizable group almost exclusively focused on Muslim markets. Low- and high-leverage firms have very similar exposure; there is no statistical difference in the degree of exposure between the two groups.[13] Therefore, in our main specification we directly compare the reaction of all firms using a simple indicator of exposure. In Section 3.6, however, we present results for alternative specifications that explicitly take into account the cross-sectional differences in exposure. In Appendix Figure 3.9 we present the distribution of an alternative measure of exposure, which is defined as share of exports to Muslim countries in total sales.

Table 3.1 provides separate descriptive statistics for low- and high-leverage firms in control and treatment groups. We report all results for the pre-boycott period 2001-2005. One important difference between control and treatment groups is size. Not surprisingly, firms exporting to Muslim countries (which are relatively distant and exotic markets) are larger on average. There is also a size difference between low- and high-leverage firms within the treatment group, as

[12] This restriction aims to reduce false categorization of treated firms and to focus on firms with a relevant share of business in Muslim countries. We provide additional robustness checks for the definitions of treatment and exposure in Section 3.6.

[13] The Kolmogorov-Smirnov test cannot reject the null of equal distributions with p=0.27.

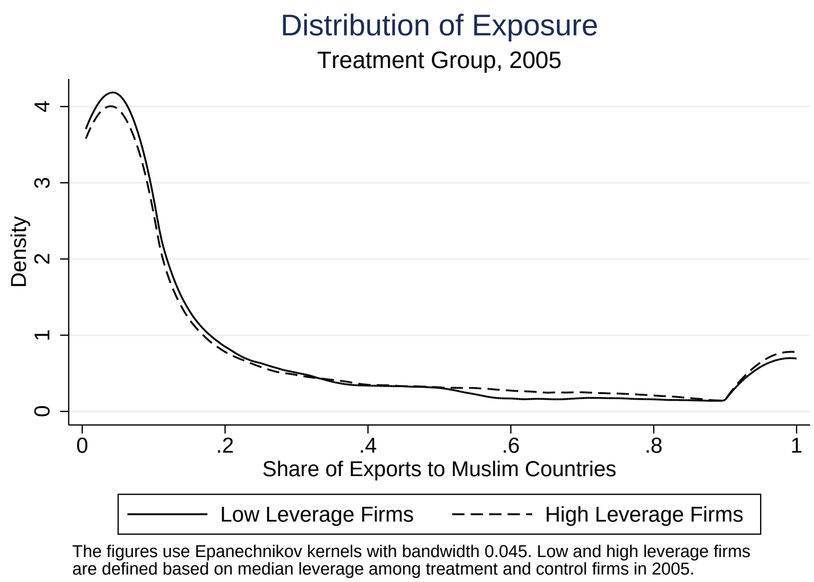

Figure 3.5. Distribution of Exposure (Share of Exports to Muslim Countries) for High- and Low-Leverage Firms

The figure shows the share of exports sold to Muslim countries for high- and low-leverage firms. High leverage is defined as share of total liabilities in total assets being above median and is computed based on values from 2005.

evidenced by sales, exports and employment. While firm-fixed effects will absorb any time-invariant effect of size, it is possible that size also influences the response to the boycott. We explicitly consider this possibility in Section 3.5, showing that the effect of leverage does not capture the differential size of firms. The number of new export destinations and newly exported products measure the degree of innovation in exports. For a subsample of manufacturing firms, our data also contains product-level data on total sales that are used to create overall measures of product innovation (number of new products). The difference in the number of export destinations and number of new products between low- and high-leverage firms are roughly proportional to the differences in their size. Section 3.5 explicitly shows that the size of product portfolio does not explain the effect of leverage. The last rows of Table 3.1 show indicator variables for outsourcing defined based on costs from the profit and loss statement (outside

Table 3.1. Summary Statistics

For sales, employment, exports, number of products, and leverage, we report the sample median. For other variables we report the sample means. We report the number of observations for the full dataset. Some variables, e.g., number of new products, are only available for a subsample of firms.

	(1)	(2)	(3)	(4)
	Control Group		Treatment Group	
	Low Lev	High Lev	Low Lev	High Lev
Firms	6,548	6,816	1,000	844
Observations	33,278	33,828	5,407	4,313
Sales	1.7m	1.7m	5.6m	3.6m
Employment	8.5	8.4	27.8	17.6
Exports	0.05m	0.04m	1.42m	0.68m
Leverage	0.586	0.782	0.558	0.766
Num Export Destinations	2	2	17	11
Num New Export Destinations	1	1	2	2
Num Export Products (2 dg)	2	2	4	4
Num Export Products (4 dg)	3	3	9	7
Num Export Products (6 dg)	4	3	11	9
Num New Export Products (2 dg)	0.038	0.036	0.081	0.043
Num New Export Products (4 dg)	0.080	0.071	0.197	0.096
Num New Export Products (6 dg)	0.139	0.131	0.345	0.163
Num New Products (2 dg)	0.104	0.126	0.155	0.108
Num New Products (4 dg)	0.162	0.184	0.328	0.185
Num New Products (6 dg)	0.238	0.257	0.544	0.293
External Labor	0.761	0.751	0.674	0.689
Temp Workers	0.532	0.523	0.493	0.491
Subcontracting	0.627	0.621	0.504	0.530

labor, separated into temporary agency workers and subcontractors). In general, treated firms are slightly less likely to outsource tasks, but there are no significant differences by leverage within the treatment and control group.

3.3.3. Empirical Strategy

Our main specification to estimate the role of leverage for adaptation to the boycott is:

$$Y_{it} = \alpha \cdot Exposed_i \cdot Post_t + \beta \cdot HighLev_i \cdot Exposed_i \cdot Post_t + d_{ht} + d_{st} + \mu_i + \varepsilon_{i,t}. \tag{3.1}$$

The dependent variable Y_{it} measures firm outcomes such as sales, number of new products, investment, employment, or outsourcing. The main coefficients of interest are α and β - the response of low-leverage firms exposed to the boycott and the differential response of high-leverage exposed firms in the period during and after the crisis, $(Post = 1\{\text{year} \geq 2006\})$. As we show in Figure 3.5, exposure to the shock is very similar for high- and low-leverage firms in the treatment group, which means the interaction with high leverage isolates the role of financial distress conditional on economic distress. This is a key advantage of our research design.

Throughout all regressions, we control for annual differences between high- and low-leverage firms using high-leverage-by-year fixed effects, $d_{ht} = 1\{\text{year} = t\}\,1\{\text{high leverage} = 1\}$. The identifying assumption of the interaction effect β is that absent the shock, and conditional on other controls, the difference between high- and low-leverage firms in the treatment group would have followed the same path as difference between high and low leverage firms in the control group. Importantly, we include industry-year fixed effects d_{st} as well as firm-fixed effects μ_i in all specifications to account for industry trends and idiosyncratic time-invariant determinants of firm outcomes. These characteristics may include location, technology, management practices, and firm size, among others. In the main specification, we measure some outcomes (exports, sales, employment, products, debt) in growth rates to allow for firm-specific trends. In additional robustness checks, we provide further results using alternative specification in levels with group-specific trends, and including time-varying firm characteristics such as size, product mix, and management quality. All results cluster standard errors at the industry level (for 53 distinct industries) because firms' adjustment may interact with other firms in their industry and we want to allow for arbitrary correlation of these responses.[14]

3.3.4. Identifying Variation

In this section, we discuss to what extent our research design addresses the two main empirical challenges common to the literature on consequences of financial leverage: (i) distinguishing financial distress from economic distress and (ii) endogeneity of capital structure choice.

Financial distress is usually accompanied or triggered by economic distress. Firms with high leverage have trouble meeting their debt obligations when their economic situation deteriorates. Oftentimes this is related to important developments inside a firm: losing key employees or important customers, facing a new aggressive competitor, or lagging behind in technological innovation. It is very challenging to distinguish the consequences of these economic factors from consequences of financial distress. For example, if we observe that a financially distressed firm decreases investment, is it because of its capital structure or because the demand for its product decreased, which caused both the financial distress and the reduction of investment?

To tackle this problem, we use a difference-in-differences strategy and explicitly control for the consequences of economic distress. Our approach is conceptually similar to that taken by Giroud and Mueller (2016), who compare the reaction of high- and low-leverage firms to local demand shocks driven by the financial crisis. In our setting, the firm-level trade data before the boycott allow us to measure the size of the economic shock precisely for each firm. This means

[14] We use the specification from Eq. 3.1 in regressions presented in all our tables. The figures present the evolution of various outcomes for low- and high-leverage firms and use the following specification:

$$Y_{it} = \delta_{lt} d_{lte} + \delta_{ht} d_{hte} + d_{ht} + d_{st} + \mu_i + \varepsilon_{i,t}. \tag{3.2}$$

where d_{lte} and d_{hte} are indicators for low- and high-leverage firms, respectively, interacted with year and exposure fixed effects. The omitted category for year fixed effects is 2005, the year before the boycott.

we can compare the differential response of firms with similar exposure to the boycott, but with different leverage before the crisis. As a result, we attribute any differences between high- and low-leverage firms to the impact of differences in financial distress. The assumption behind this procedure is that high-leverage firms would exhibit a similar reaction to low-leverage firms if their leverage ratio was lower. Given that their exposure to the shock is the same (see Table 3.5), this is a reasonable assumption. However, to explicitly account for potential differences in exposure and in firm size (since high-leverage firms are on average smaller than low-leverage firms, see Table 3.1), Sections 3.5 and 3.6 present additional evidence supporting our main results.

Endogeneity of leverage choice is another issue plaguing the empirical corporate finance literature. Because firms choose their leverage having expectations about the future, it is hard to rule out reverse causality and the influence of omitted variables. For example, when we observe that a firm has more flexible labor contracts after increasing leverage, is this the causal effect of leverage or did the anticipation of more flexible contracts make this firm choose higher leverage? Or perhaps something else happened at the firm, e.g., new management was introduced, which led to both the increase in leverage and a change in contract flexibility.

Even though we do not have exogenous variation in capital structure, our research design is well-suited to rule out reverse causality concerns. The unexpected nature of our shock did not allow firms to adjust their capital structure in anticipation of the boycott. Concerns about omitted variables are alleviated by our triple-difference design. Because we explicitly control for the difference between high- and low-leverage firms not subject to the shock, our results cannot be attributed to an omitted factor which would lead to diverging outcomes of high- and low-leverage firms even in the absence of the boycott. Some omitted factor may, however, influence firms' reaction to the boycott. To address this possibility, we explicitly discuss alternative explanations in Section 3.5. We first rule out that leverage is only a proxy for other firm characteristics such as size, product mix, or management quality that affect the response to shocks. Then we show that results based on cash holdings and a debt-maturity approach are consistent with our main findings and lend further support to the mechanism of financial constraints.

Finally, we want to emphasize that our results estimate the average treatment effect on the treated high and low leverage firms. We do not claim that the effects we find can be generalized to the entire population of firms. It is true that some firms pre-select into having large debt and we document ex-post consequences of debt for a group of these firms that were exposed to the boycott. Whether these consequences would be the same if we randomly allocated more debt to other firms is a question that goes beyond our analysis. On the one hand, traditional theory of trading off financial and economic risk suggests that firms with low operating risk – e.g. firms for whom it is easier to redirect their sales – may be more willing to take on more financial risk, and hence have high leverage. If that were the case, our results would underestimate flexibility costs of debt for the general population of firms. On the other hand, firms choosing high leverage may have lower cost of operational disruption, i.e. labor hoarding may be less important for them. In that case, our results would overestimate the value of flexibility to wait. In any case, from the perspective of policy and firm-decision making, it is important to understand the effects of debt on firms who actually do have or consider having a significant amount of debt. Hence, the treatment effect on the treated is the relevant object of analysis.

3.4. Results

In this section, we provide graphical and regression-based evidence on the differential response to the boycott by Danish exporters with high or low leverage. We first present the differences in several main outcomes - exports, sales, debt - that suggest low-leverage firms are able to adjust to the shock without losing much of their business, but high-leverage firms are not. We then show how this adjustment happens, analyzing number of export products, product innovation, and investment. Low-leverage firms undertake new projects and grow their businesses in other markets after the shock, but high-leverage firms might lack flexibility to do that. Limited flexibility also leads high-leverage firms to make decisions that increase their reliance on outsourcing and are meant to reduce operating risk.

3.4.1. Exports, Sales, and Debt

The boycott significantly reduces affected firms' exports to Muslim countries and hence has a direct negative impact on their sales. However, production capacity freed by lower demand from Muslim countries could be used for producing goods sold in other markets. To boost sales elsewhere, firms may need to create new products, increase marketing expenses, or offer more attractive prices to their customers. High-leverage firms may be unable to afford these actions or may find them too risky and, as a result, the decline of their sales could be larger, even though the initial exposure to the boycott is the same.

We present the effects of the boycott on exports, sales, and bankruptcy in Table 3.2. Both low-leverage and high-leverage firms reduce exports as a consequence of the boycott (column 1). While the difference between the two groups is not significant, the point estimate for high-leverage firms is negative, suggesting that they might experience a larger export drop.

To shed more light on potential differences by destination markets, we restrict the sample to firms exposed to Muslim countries in columns 2 and 3 and separate their changes in exports to Muslim countries and other destinations. As expected, the decline is driven by Muslim countries (column 2). However, the decline for low-leverage firms is accompanied by an increase in exports to other countries (column 3) and an increase in domestic sales. We analyze the drivers of the increase in domestic sales and exports to other countries more closely in the next section. In sum, total sales growth for low leverage firms does not decrease (column 4). High-leverage firms, however, do not significantly increase their exports elsewhere. The sum of the coefficients in column 3 is not significantly different from zero, suggesting that high-leverage firms are unable to redirect their sales to other markets. As a result, their total sales growth decreases.[15]

The shock caused by the boycott is not large enough to drive affected firms out of business. Column 5 presents the results of a cross-sectional regression for all exporters in the sample in 2005, with the dependent variable being an indicator for firm exit in 2006. There is no significant effect of the boycott for either low- or high-leverage firms. Firm survival is therefore unlikely to be an important element of the analysis. This is an interesting feature of our setting, that allows us to study potential effects of leverage on firms' operations far away from bankruptcy threats.

At the same time, the shock may still threaten firms' liquidity and require additional funds to accommodate its consequences. Consistent with this, Table 3.3 shows a significant increase in borrowing by low-leverage firms. We find a 3.9% increase in total debt (column 1), which is driven by a substantial increase in short-term debt to suppliers (column 2), as well as an increase in long-term debt for some firms. While the increase in the intensive margin of long-term debt

[15]This result emphasizes the importance of distinguishing high- and low-leverage firms. In a case study of dairy producers, Hiller et al. (2014) do not find reallocation of products across markets on average for exposed dairy producers after the boycott.

Table 3.2. Response of the Amount of Business

All regressions, except column 5, include firm fixed effects, industry-year fixed effects, and binary variables for each year interacted with indicators for high leverage. Columns 2 and 3 contain only firms exposed to the boycott (because only for them it is meaningful to analyze exports to Muslim countries around the boycott). The main independent variable is triple interaction of exposure to the boycott (treatment), having high leverage, and post-boycott period (year 2006). In columns 2 and 3 all firms are exposed. Dependent variables in columns 1-4 are log-differences in total exports, total exports to Muslim countries, total exports to non-Muslim countries, and total sales. Column 5 is a cross-sectional regression of an indicator for firm death in 2006. The bottom row presents the mean of dependent variables in the pre-boycott period. In all regressions, standard errors are clustered at the industry level (53 industries). Joint p-val row presents a p-value from the F-test for significance of the high-leverage firms' response (the sum of baseline coefficient for low-leverage firms and differential effect for high-leverage firms).

	(1)	(2)	(3)	(4)	(5)
	Δln(Export)	Δln(Exp Muslim)	Δln(Exp Other)	Δln(Sales)	Firm Exit 2006
Treatment X 2006	-0.2115***	-0.6247***	0.3476*	0.0289	-0.0021
	(0.069)	(0.081)	(0.172)	(0.022)	(0.001)
Treatment X High X 2006	-0.0664	0.0938	-0.1910	-0.0510*	0.0037
	(0.076)	(0.107)	(0.228)	(0.027)	(0.003)
Obs	53,910	6,077	9,157	76,826	12,626
R-squared	0.016	0.027	0.006	0.007	0.006
Firms	13,307	1,590	1,799	15,208	12,626
Joint p-val	3.16e-07	4.19e-06	0.404	0.314	0.507
Sample Avg 01-05	0.0298	0.163	-0.003	0.0576	0.004

*** p<0.01, ** p<0.05, * p<0.1

in column 4 is not significant, we find a large increase in long-term liabilities when including firms that did not hold long-term debt before the boycott (column 5). In contrast, the increase in debt holdings is significantly smaller for high-leverage firms, and their net change of total debt is negative.[16] Firms with high leverage may be unable to increase borrowing because lenders

[16]Note that these changes in liabilities are relative to high-leverage and low-leverage firms in the control group, thereby accounting for mean reversion among high and low leverage firms during normal times. In addition, we conduct a Placebo test by defining high leverage based on debt holdings at the end of 2002 and analyzing changes in debt in 2004. This exercise aims to capture any differential regression to the mean among firms in the treatment group during normal times. Appendix Table 3.15 shows that, if anything, treated low leverage firms reduce their

Table 3.3. Response of Liabilities

All regressions include firm fixed effects, industry-year fixed effects, and binary variables for each year interacted with indicators for high leverage. The main independent variable is triple interaction of exposure to the boycott (treatment), having high leverage, and post-boycott period (year 2006). Dependent variables are the log-changes in total debt (column 1), short-term debt to suppliers (column 2), other short-term debt (column 3), and long-term debt (column 4-5). All columns consider the log of debt, except column 5, which measures log(1 + long-term debt). The bottom row presents the mean level of the dependent variables in the pre-boycott period (in millions of DKK). In all regressions, standard errors are clustered at the industry level (53 industries). Joint p-val row presents a p-value from the F-test for significance of the high-leverage firms' response (the sum of baseline coefficient for low-leverage firms and differential effect for high-leverage firms).

	(1)	(2)	(3)	(4)	(5)
	Δln(Debt)	Δln(Short-Term Debt)		Δln(Long	Δln(1+Long
		To Suppliers	To Other	-Term Debt)	-Term Debt)
Treatment	0.0388**	0.0816**	-0.0005	0.0652	0.3534***
X 2006	(0.019)	(0.033)	(0.029)	(0.057)	(0.117)
Treatment	-0.0544*	-0.0930***	-0.0672	-0.0904*	-0.2220
X High	(0.029)	(0.032)	(0.041)	(0.050)	(0.152)
X 2006					
Obs	76,790	76,110	76,263	50,204	76,810
R-squared	0.013	0.020	0.015	0.028	0.039
Firms	15,207	15,189	15,201	13,716	15,208
Joint p-val	0.389	0.716	0.0548	0.709	0.352
Sample	38.8	7.3	20.3	17.6	11.2
Avg 01-05					

*** p<0.01, ** p<0.05, * p<0.1

consider them too risky given their already high debt. Similarly, suppliers may be willing to improve financing conditions for low-leverage firms, but may consider it too risky to do the same for high-leverage firms.

3.4.2. Flexibility To Grow: Redirecting Sales and Innovation

In Section 3.3.2 we showed that the pre-boycott export exposure to Muslim countries was similar for low- and high-leverage firms. Table 3.2, however, shows that sales of low-leverage

liabilities, while high leverage firms increase them compared to firms in the control group. These differences would lead us to understate the effects in Table 3.3.

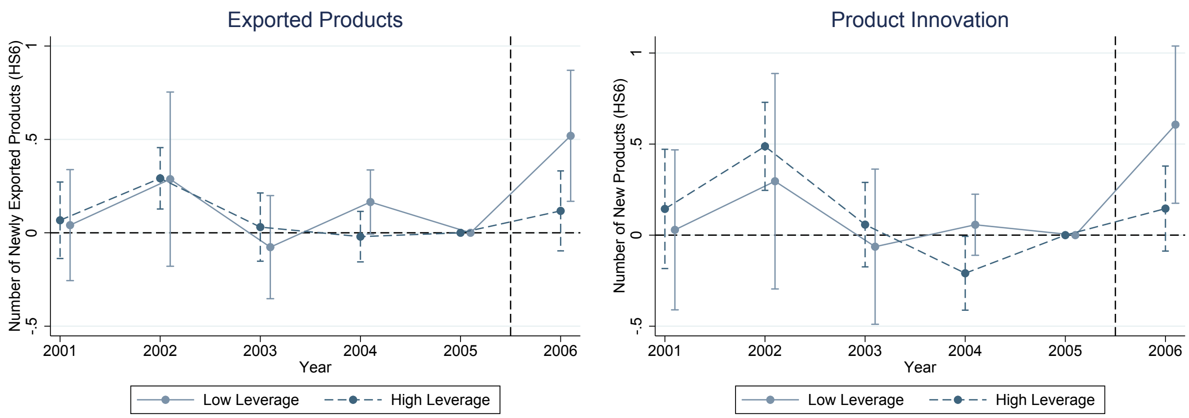

Figure 3.6. Product Innovation Response

The dependent variables are the change in number of new export products and the change in number of new products (including domestic sales) at the 6-digit HS level. The bars in the figure show coefficients from a regression analogous to the main specification in which high-leverage X treatment X pre-boycott term was split into several terms for each year separately (normalizing the year 2005). Coefficients for low- (high-) leverage firms are for firms exposed to the boycott with leverage below (above) the median.

firms did not decrease, contrary to high-leverage firms. These patterns suggest that low-leverage firms were able to pursue new business opportunities at home or abroad and redirect their sales elsewhere. In this subsection we investigate the product market response in more detail. We take advantage of the detailed product-destination-level data on export flows and total sales, and analyze entry into new export markets, product innovation, and investment by low- and high-leverage firms.

Figure 3.6 presents graphical evidence from an extended regression with annual coefficients for low- and high-leverage firms in the treatment group; see equation (3.2). The solid lines show that low-leverage firms responded to the boycott by increasing the number of exported products (left panel) and the number of total products sold (right panel). For high-leverage firms (dashed lines), the response along these margins is close to zero and insignificant. Tables 3.4 and 3.5 provide a more detailed analysis of export margins and product innovation.

Column 1 of Table 3.4 shows that as a result of the boycott, low-leverage firms stopped exporting to around 15% of their export destinations. We further estimate that the decrease was 4% higher for high-leverage firms.[17] We note, however, that Appendix Table 3.12 presents a level-based specification in which the effect on export destinations for high-leverage firms is insignificant. Since the log-based specification puts more weight on firms with a lower number of export destinations, we interpret these differences as suggestive evidence that the negative effect of high leverage on the extensive margin response is driven by firms serving few markets before the boycott.

In columns 2-4, we find substantial differences in adjustment of the exported product mix. The dependent variable is the log change in the number of products for which we observe non-zero export flows for a given firm. Column 2 defines product category based on the 6-digit classification from the harmonized system (HS), while columns 3 and 4 use the 4- and 2-digit classification, respectively. An example of a 2-digit category is "Coffee, Tee, Mate, and Spices." Within this group, 4-digit categories include "Coffee" and "Tea," while "Black Tea" or "Green Tea" are examples of 6-digit products. Although introducing new 6-digit products, as evidenced in column 2, may be viewed as slightly modifying a firm's range of export products, introducing a new 2-digit category (column 4) is presumably more complicated and requires additional investment. Our results therefore confirm that low-leverage firms actively innovate and look for new ways to use their existing capacity. This response is muted for high-leverage firms: the sum of coefficients in each column is insignificant and very close to zero. As with entry into new countries, additional analysis using level specifications suggests that the ability

[17] Part of the exit from Muslim markets could be compensated by entrepot trade. Although Danish producers may try to use nearby non-Muslim countries as a gateway to reach Muslim markets, the boycott targets all products associated with Denmark, not only those produced in Denmark. Therefore a Danish brand will be boycotted even if the product is sold by a foreign firm or even if it is produced in a different country.

Table 3.4. Redirecting Sales: New Export Markets and Products

All regressions include firm fixed effects, industry-year fixed effects, and binary variables for each year interacted with indicators for high leverage. The main independent variable is triple interaction of being exposed to the boycott (treatment), having high leverage and post-boycott period (year 2006). Dependent variables are changes of number of export destinations (column 1) and number of exported products defined as non-zero flows in 6-, 4-, and 2-digits HS product category (columns 2-4). The bottom row presents mean levels of dependent variables in the pre-boycott period. In all regressions standard, errors are clustered at the industry level. Joint p-val row presents a p-value from the F-test for significance of the high-leverage firms' response (the sum of baseline coefficient for low-leverage firms and differential effect for high-leverage firms).

	(1)	(2)	(3)	(4)
	ΔLog(Export Destinations)	ΔLog(Num Export Products) 6 digits	4 digits	2 digits
Treated X 2006	-0.1473***	0.0732***	0.0886***	0.0672***
	(0.022)	(0.026)	(0.022)	(0.016)
Treated X High X 2006	-0.0397**	-0.0883***	-0.0923***	-0.0642**
	(0.016)	(0.029)	(0.028)	(0.024)
Observations	53,910	53,910	53,910	53,910
R-squared	0.032	0.032	0.030	0.023
Firms	13,307	13,307	13,307	13,307
Joint p-val	0.000	0.393	0.809	0.871
Sample Avg 01-05	7.810	11	7.878	4.149

*** p<0.01, ** p<0.05, * p<0.1

to introduce new export products is most limited for firms with high leverage and fewer exported products (see Appendix Table 3.12).

By looking only at exports we are unable to tell if newly exported products are indeed product innovations or existing products which were sold domestically before. To answer this question, we use sales data at the product level for all manufacturing firms with at least 10 employees to define new products for each firm over time.

We report the results in Table 3.5. As suggested by the product dynamics in export markets, we find evidence that low-leverage firms exposed to the boycott introduce differentially more

Table 3.5. Product Innovation

All regressions include firm fixed effects, industry-year fixed effects, and binary variables for each year interacted with indicators for high leverage. The main independent variable is triple interaction of exposure to the boycott (treatment), having high leverage, and post-boycott period (year 2006). Dependent variables are the change in the total number of new products and indicators for any new products sold (columns 1-6) or exported (columns 7-9), defined as first-time non-zero revenue in a 6-, 4-, or 2-digits product category in the HS system. The bottom row presents the mean of dependent variables in the pre-boycott period. In all regressions, standard errors are clustered at the industry level. Joint p-val row presents a p-value from the F-test for significance of the high-leverage firms' response (the sum of baseline coefficient for low-leverage firms and differential effect for high-leverage firms).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	ΔNumber of New Products			Any New Products			ΔNum New Export Products		
	(6 digit)	(4 digit)	(2 digit)	(6 digit)	(4 digit)	(2 digit)	(6 digit)	(4 digit)	(2 digit)
Treated X 2006	0.5478**	0.3338**	0.1693**	0.0307*	0.0372**	0.0239	0.4400**	0.2855**	0.1089*
	(0.230)	(0.155)	(0.081)	(0.015)	(0.015)	(0.016)	(0.184)	(0.118)	(0.056)
Treated X High X 2006	-0.4707**	-0.3103**	-0.2081**	-0.0416**	-0.0216	-0.0416**	-0.3838**	-0.2702**	-0.0948
	(0.192)	(0.146)	(0.091)	(0.029)	(0.026)	(0.020)	(0.140)	(0.105)	(0.058)
Observations	10,258	10,258	10,258	10,258	10,258	10,258	10,258	10,258	10,258
R-squared	0.027	0.024	0.028	0.024	0.019	0.018	0.029	0.022	0.021
Firms	2,369	2,369	2,369	2,369	2,369	2,369	2,369	2,369	2,369
Joint p-val	0.568	0.791	0.436	0.560	0.382	0.083	0.538	0.818	0.681
Sample Avg 01-05	0.272	0.148	0.0604	0.0952	0.0690	0.0364	0.176	0.0872	0.0333

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

new products in 2006 compared to other firms and to their previous innovative behavior. Product innovation is less pronounced among exposed firms with high leverage and the total effect for this group is close to zero. The magnitude of changes is considerable, with low-leverage exposed firms adding on average 0.3 new products in a 4-digit category. This corresponds to a 200% increase relative to the annual average rate of innovation across firms over the period 2001-2005. Finally, we use the product-level sales data to refine the analysis on product mix in export markets. Specifically, we identify product innovation in export markets as new products immediately sold abroad. Columns 7-9 of Table 3.5 show that for the sample of manufacturing firms with information on product-specific sales, product innovation drives between 65%-85% of the increase in newly exported products. Again, the total effect for high-leverage firms is close to zero.

What drives the innovative response that we observe? Given that we analyze one year after the boycott, and that technological innovation typically requires a long R&D process, the observed effects may be surprising. However, other types of innovation may require much less time to be completed. In our data, we observe the strongest response from firms that export food, textiles, and machinery. We inspect the export data to better understand the type of innovation that takes place. Firms often start exporting products that may serve a different purpose but are technologically closely related (e.g. a producer of men's underwear may start selling men's pajamas). They also sometimes start exporting complementary products (e.g., a firm exporting meat products adding sauces or spices), probably making use of their idle sales force and trying to bundle products to make them more attractive for their existing customers. For machinery, a significant share of new products are parts, possibly for maintenance and upgrades of products that the firm sold in the past.

The development of new products and redirecting sales to new markets presumably requires significant investment. Table 3.6 analyze the investment response to the boycott for the full sample in columns 1-3 and for the subset of firms with data on product innovation in columns 4-6. We report two alternative measures, investment as a share of lagged assets and investment flows (at the intensive margin and also accounting for zeros using one plus investment). Consistent with previous findings, low-leverage firms increase investment, whereas high-leverage firms do not. The results for the smaller innovation sample are less precisely estimated, but column 6 reveals that the main investment increase occurs for low-leverage firms who have been actively investing in the past. High-leverage firms may not have enough financial flexibility to undertake additional investment, which may explain the lack of innovative response. Lack of financial flexibility among high leverage firms may also explain an additional factor contributing to the differential product market and sales response: change in receivables (Appendix Table 3.14). The level of receivables increases for low-leverage firms, suggesting that low-leverage firms may be improving financing conditions offered to their customers and thus encouraging an increase in sales.

3.4.3. Flexibility to Wait: Employment and Outsourcing

As demonstrated by the previous subsection, high-leverage firms may lack flexibility to pursue new business opportunities. But high leverage may also impede firms' ability to withstand the negative shock without any action and force them to make costly adjustments. When the boycott hits, revenues of the firm decrease but many cost categories remain fixed. As a result, profitability decreases and operating leverage increases, and the firm may need to reduce its risk to safely continue its operations. While financial leverage may be difficult to adjust, the firm

Table 3.6. Investment Response

All regressions include firm fixed effects, industry-year fixed effects, and binary variables for each year interacted with indicators for high leverage. The main independent variable is triple interaction of exposure to the boycott (treatment), having high leverage, and post-boycott period (year 2006). Dependent variables are the value of investment in % of total assets (columns 1, 4), log of 1 + total investment (columns 2, 4), and log of total investment (columns 3, 6). The bottom row presents mean of dependent variables in the pre-boycott period. In all regressions, standard errors are clustered at the industry level. Joint p-val row presents a p-value from the F-test for significance of the high-leverage firms' response (the sum of baseline coefficient for low-leverage firms and differential effect for high-leverage firms).

	(1)	(2)	(3)	(4)	(5)	(6)
	Investment (% Assets)	ln(1+Inv)	ln(Inv)	Investment (% Assets)	ln(1+Inv)	ln(Inv)
	Full Sample			Innovation Sample		
Treatment X 2006	0.0063*	0.2892***	0.2335***	0.0220	0.3158	0.2761***
	(0.004)	(0.092)	(0.062)	(0.019)	(0.161)	(0.096)
Treat X High X 2006	-0.0070*	-0.2381	-0.1931**	-0.0182	-0.3252	-0.1034
	(0.004)	(0.158)	(0.158)	(0.022)	(0.300)	(0.143)
Observations	76,826	76,826	45,821	10,258	10,258	8,674
R-squared	0.001	0.047	0.060	0.006	0.060	0.040
Firms	15,208	15,208	13,437	2,369	2,369	2,242
Joint p-val	0.674	0.627	0.604	0.583	0.969	0.419
Sample Avg 01-05	0.014	1520	2448	0.041	4212	4948

*** p<0.01, ** p<0.05, * p<0.1

may try reducing operating leverage, in line with the trade-off hypothesis (Van Horne, 1977; Mandelker and Rhee, 1984). In this subsection, we analyze whether firms exposed to the boycott can hoard labor or whether they try to gain additional flexibility by turning to outsourcing as a more flexible input in production.

Table 3.6 and the left panel of Figure 3.7 illustrate that high-leverage firms reduce employment after the boycott. This remains true across three measures of employment: total headcount, full-time equivalent workers, and total wage bill. While the decline is consistent with general downsizing of the firm, one may expect that employment does not decrease after a temporary

Table 3.7. Employment Response

All regressions include firm fixed effects, industry-year fixed effects, and binary variables for each year interacted with indicators for high leverage. The main independent variable is triple interaction of exposure to the boycott (treatment), having high leverage, and post-boycott period (year 2006). Dependent variables are the log-differences in count of workers employed (column 1), full-time equivalence employment (measure of total hours worked) in column 2, total wage bill (column 3), and the change in the share of workers with college degree or at least college (columns 4-5). The bottom row presents mean of dependent variables in the pre-boycott period. In all regressions, standard errors are clustered at the industry level. Joint p-val row presents a p-value from the F-test for significance of the high-leverage firms' response (the sum of baseline coefficient for low-leverage firms and differential effect for high-leverage firms).

	(1) Δln(Emp) (headcount)	(2) Δln(Emp) (FTE)	(3) Δln(Wages) (Total)	(4) Δ% College	(5) Δ% College+
Treatment	0.0043	0.0112	-0.0023	0.0021	0.0026
X 2006	(0.014)	(0.014)	(0.015)	(0.003)	(0.003)
Treat X High	-0.0347*	-0.0585***	-0.0426**	-0.0068**	-0.0080*
X 2006	(0.018)	(0.020)	(0.018)	(0.003)	(0.005)
Observations	76,826	76,826	74,826	76,469	76,469
R-squared	0.009	0.023	0.015	0.002	0.002
Firms	15,208	15,208	15,208	15,146	15,146
Joint p-val	0.0495	0.0067	0.0031	0.150	0.184
Sample Avg 01-05	0.0012	0.0358	0.0503	0.0027	0.0039

*** p<0.01, ** p<0.05, * p<0.1

reduction in sales. Hiring and firing is costly, and employees build firm-specific human capital that contributes to the stickiness of employment after temporary negative shocks. This stickiness, known as labor hoarding (Okun, 1963), is especially likely in our setting as firms face significant uncertainty on the length of the boycott. Hoarding labor, however, is costly. It requires paying workers' salaries today, even though they are not productive, and the benefits can only be recouped in the future in the form of reduced hiring and training costs. If a firm is financially constrained, it may be unable to engage in labor hoarding. Instead, it may choose to reduce employment more than firms with a lot of financial slack, even though it may be sub-optimal in the long run.[18] Our employment results are consistent with this hypothesis: while low-leverage firms do not reduce employment, possibly because they engage in labor hoarding, high-leverage firms do.

Whom do high-leverage firms fire? Columns 4 and 5 of Table 3.6 shows that the share of high-skilled workers with at least a college degree decreases differentially by 0.8 pp, which corresponds to 4% of the base value and is twice the magnitude of average yearly changes. This lends further support to high-leverage firms' inability to hoard labor: high-skilled workers are most expensive to hoard (because of their high salaries) but the benefits of hoarding them are arguably highest because of expensive recruitment and their extensive firm-specific knowledge. It is also consistent with firing employees who focus on non-core activities, such as accountants or marketing specialists. These workers are more likely to be highly skilled than production workers.

Interestingly, the decline in employment for high-leverage firms is accompanied by an increase in the use of outsourcing. Hence, the observed employment drop represents not only a decline in the total labor input, but also changes in the nature of labor contracts: high-leverage firms reduce their use of more rigid employment contracts and increase their flexibility by outsourcing. Table 3.8 presents the results of our outsourcing analysis. Column 1 shows that high-leverage firms are over 5 pp. more likely to report any outsourcing after the boycott. This is a large effect given the fact that 75% of firms before the boycott already declared using some outsourcing, and hence 5 pp. corresponds to 20% of firms that have not used any outsourcing before. Column 2 shows that the share of total spending on outsourcing in sales significantly increases for high-leverage firms. The right panel of Figure 3.7 presents

[18] This is related to the inability of high-leverage firms to credibly enter relational employment contracts, see e.g. Fahn et al. (2017).

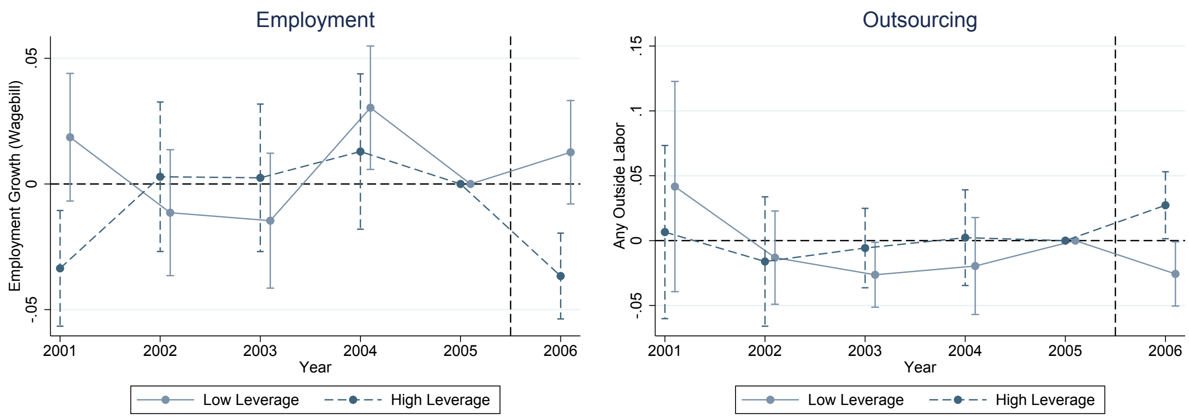

Figure 3.7. Employment and Outsourcing for High- and Low-Leverage Firms
The dependent variables are the change in the worker wage bill and an indicator for any outsourcing. The markers in the figures show coefficients from regression analogous to the main specification in which high-leverage X treatment X pre-boycott term was split into several terms for each year separately. Coefficients for low-leverage firms are for firms exposed to the boycott (treated) with leverage below median. Coefficients for high- versus low-leverage are differential effects for high-leverage firms within treated firms. The brackets denote 90% confidence intervals.

corresponding graphical evidence for the outsourcing adjustment. While the difference between high-leverage firms in the treatment and control group was flat and insignificant before the boycott, after the boycott exposed high-leverage firms are significantly more likely to use outsourcing.[19]

The data on outsourcing that we use in columns 1 and 2 of Table 3.8 come directly from the firm's profit and loss statement, so the results are based only on total level of costs associated with outsourcing. To provide more details on what services firms outsource, we use supplementary data on outsourcing activities from a firm survey on purchased intermediate goods and services collected by Statistics Denmark.[20] The data cover 1,221 manufacturing firms from our main

[19]The accounting information allows us also to disentangle outsourcing into hiring of agency workers and other freelancers. Analysis conducted for the indicators of these two separate types of outsourcing yields results with magnitudes similar in both, but the results become less precise when analyzed separately.

[20]For another recent application using this outsourcing survey to analyze firms' choices between intermediate inputs and in-house production, see Chan (2017).

Table 3.8. Outsourcing Response

All regressions include firm fixed effects, industry-year fixed effects, and binary variables for each year interacted with indicators for high leverage. The main independent variable is triple interaction of exposure to the boycott (treatment), having high leverage, and post-boycott period (year 2006). Dependent variables are indicator for any outsourcing (column 1), amount of outsourcing expenses as the share of total sales (column 2), as well as indicators for outsourcing in the area of information and communication technologies, marketing and engineering (columns 3-5). Standard errors are clustered at the industry level. The bottom row presents mean of dependent variables in the pre-boycott period. Joint p-val row presents a p-value from the F-test for significance of the high-leverage firms' response (the sum of baseline coefficient for low-leverage firms and differential effect for high-leverage firms).

	(1)	(2)	(3)	(4)	(5)
	Any	Outsourcing		Any Outsourcing	
	Outsourcing	(% Sales)	ICT	Marketing	Engineering
Treatment	-0.0213	-0.0017	-0.0067	-0.0232	-0.0170
X 2006	(0.018)	(0.001)	(0.009)	(0.018)	(0.027)
Treatment X	0.0509**	0.0047*	0.0429**	0.0541***	0.0120
High X 2006	(0.020)	(0.003)	(0.020)	(0.018)	(0.046)
Obs	76,826	76,826	5,309	5,309	5,309
R-squared	0.032	0.035	0.040	0.032	0.023
Firms	15,208	15,208	1,230	1,230	1,230
Joint p-val	0.0197	0.287	0.035	0.003	0.882
Sample Avg 01-05	0.747	0.0144	0.965	0.930	0.614

*** p<0.01, ** p<0.05, * p<0.1

sample. We construct several variables indicating non-zero spending on outsourcing of tasks in a given category. We then use our main specification (3.1) to see which categories are responsible for the overall increase observed in the profit and loss statement.

We present selected results in columns 3 to 5 of Table 3.7. We find that high-leverage firms significantly increase outsourcing in the areas of information and communication technology, as well as marketing, but not engineering.[21] Table 3.13 in the Appendix reports results for other

[21] The baseline rates of outsourcing in these categories are very high, but we nonetheless find that the remaining firms that did not outsource before start doing it after the boycott, bringing the share of high leverage firms which use e.g. information and communication technologies outsourcing close to 100%.

groups of activities. In general, these newly outsourced activities are more likely to correspond to tasks performed by high-skill workers, e.g., lawyers, accountants, or IT specialists, which is consistent with the decrease in high-skilled workers' share in total employment reported in column 6 of Table 3.6. Outsourcing high-skill services is further consistent with substituting non-core activities at these manufacturing firms: the results do not show any significant response in engineering services, suggesting that these firms remain focused on their core competencies.

3.5. Main Mechanism and Alternative Explanations

In this section, we shed more light on the underlying mechanism for the main results. We first address the concern that leverage may not only capture the role of financial constraints at the time of the boycott, but could also serve as a proxy for other unobserved differences across firms, in particular differences in firm size, industry, and product variety. In the second step, we provide additional evidence supporting our main results and financial constraints interpretation. Analyzing cash holdings, we show that being a cash-rich firm has qualitatively the same effect as having low-leverage, but cash holdings are relatively small and can only explain a small part of our findings. Finally, we show that a higher share of long-term debt maturing soon after the boycott (similarly to Almeida et al. 2011) yields effects that are quantitatively similar to our main findings for high-leverage firms.

Table 3.9 analyzes the role of other firm characteristics that may influence firms' responses during the crisis. To ensure that the documented effect of leverage is not driven by these characteristics, we extend the main specification in equation (3.1) by adding additional control variables. Notice that Table 3.9 presents only the triple-interaction coefficients. This is because

Table 3.9. Alternative Explanations for Differential Adjustment of High- and Low-Leverage Firms

Each panel in the table reports results for a separate regression. All regressions are analogous to the main specification but enriched with an interaction of exposure indicator, year fixed effects, and the main additional regressor of interest. In panels A-B the additional regressor is a continuous measure of firm size, defined as total employment or total sales. In panel C, the additional regressor is the number of products exported by the firm. In panels D-E we include measures of managerial quality based on managers' average education and pay. The main independent variable is the triple interaction of exposure to the boycott (treatment), having high leverage, and post-boycott period (year 2006). Notice that coefficients for Treatment X 2006 are not reported, because after additional regressors are added, there is no single coefficient that identifies the effect for low-leverage firms. In all regressions, standard errors are clustered at the industry level.

	(1)	(2	(3)	(4)	(5)	(6)
Model Extension	Δln(Sales)	ΔNew Products (HS6)	ΔNew Export Products (HS6)	ln(1+ Investment)	Δln(FTE Employment)	Any Outsourcing
Panel A: Employment X Year FE X Exposed						
Treated X 2006	-0.0523*	-0.3330**	-0.2695**	-0.1964	-0.0552***	0.0572***
X High Leverage	(0.026)	(0.162)	(0.120)	(0.157)	(0.018)	(0.019)
Panel B: Sales X Year FE X Exposed						
Treated X 2006	-0.0619**	-0.3063*	-0.2526**	-0.2041	-0.0530***	0.0529***
X High Leverage	(0.030)	(0.150)	(0.109)	(0.154)	(0.017)	(0.019)
Panel C: Number of Export Products X Year FE X Exposed						
Treated X 2006	-0.0504*	-0.3569**	-0.2739**	-0.2371	-0.0587***	0.0525**
X High Leverage	(0.027)	(0.165)	(0.131)	(0.159)	(0.018)	(0.019)
Panel D: Manager Education X Year FE X Exposed						
Treated X 2006	-0.0181	-0.4558**	-0.3594**	-0.2443	-0.0306	0.0421
X High Leverage	(0.024)	(0.188)	(0.143)	(0.259)	(0.022)	(0.030)
Panel E: Manager Pay X Year FE X Exposed						
Treated X 2006	-0.0098	-0.4733**	-0.3791***	-0.2708	-0.0294	0.0316
X High Leverage	(0.027)	(0.175)	(0.134)	(0.249)	(0.023)	(0.025)

*** p<0.01, ** p<0.05, * p<0.1

after including interactions of exposure, year, and various firm characteristics, the effect for low-leverage firms is not identified by a single coefficient.

Panels A and B show that the differences between high- and low-leverage firms are not explained by differences in firm size. Specifically, we include measures of firm size in 2005, defined by total employment and sales, as well as interaction terms for size, the exposure indicator, and each year of the sample period. The results from this model are quantitatively similar to the main results. This is an important robustness check because Table 3.1 shows some differences in firm size between these two groups among exposed firms.

Panel C provides an alternative check for whether highly leveraged firms suffer more because they have fewer opportunities to redirect sales, measured by the number of export products before the boycott. This regression also rejects product variety as a main factor in the differences between high- and low-leverage firms. Most importantly, existing differences in product mix do not explain the results on product innovation and the extensive margin of trade in section 3.4.2.

Panels D and E alleviate the concern that leverage is correlated with manager quality. To this end, we leverage the occupational information in the matched employer-employee data to identify managers and their characteristics (Friedrich, 2016). We use two proxies for manager's quality: average educational level and average pay of a manager in the firm. We interact these proxies with year fixed effects and exposure indicators. The results remain similar to the main specification, although the estimates of employment reduction and increase in outsourcing become imprecise.

In the second step, we provide additional direct evidence of the liquidity mechanism, supporting our financial constraints interpretation of the main results. First, another way of looking at a firm's liquidity, ignored by our measure of leverage, is the size of cash holdings. If the effects we document are driven by lack of liquidity, analyzing the effect of cash holdings should yield opposite patterns. At the same time, we want to understand to what extent having high leverage is correlated with low cash holdings.

Table 3.10 first analyzes the role of cash ratio, defined as cash holdings relative to total liabilities. Panel A uses an indicator for high versus low cash ratio, where the cutoff is the median cash ratio among all exposed and non-exposed exporters in 2005. The results suggest that firms with higher cash holdings as a share of total liabilities increase their sales more and are insignificantly less likely to use outsourcing. This is consistent with the shielding role of cash holdings in the face of adverse shocks. Yet, these results are not precisely estimated.

Panel B replaces the binary indicator for high versus low cash ratio with the continuous measure of cash relative to liabilities in 2005. This additional variation provides more power to find statistically significant patterns of more labor hoarding and less outsourcing among exposed firms with higher cash holdings. Yet, comparing the magnitude of these effects with the main results for high versus low leverage indicates that cash constraints explain only a small part of the adjustment pattern. The average cash ratio among exposed firms is 5% of total liabilities, with a standard deviation of 10%. This implies that an increase in the cash ratio by one standard deviation yields, for example, a muted employment reduction by 0.1%, compared to the average effect of -3.5 to -6% in Table 3.7.

We do not find evidence that a higher cash ratio improves a firm's ability to innovate. That lack of effect and the small cash buffer effect on employment and outsourcing outcomes may be partially explained by the fact that cash holdings of most firms are small compared to the size of the economic shock. This suggests that what really matters for increased innovation is being able to increase debt to finance these activities.

We provide additional suggestive evidence for the mechanism of financial constraints in Panel C of Table 3.10. Unfortunately, we do not observe loan-level data on maturity. Instead, we can only use changes in net debt stocks in the years before the boycott to approximate for maturity after the boycott starts. In particular, we define firms as likely to face a high stock of maturing long-term debt in 2006 if less than half of their stock of long-term debt at the end of 2005 is accounted for by an increase in long-term debt during 2005. As a result, these firms

Table 3.10. Results: Cash Holdings and Debt Maturity

Each panel in the table reports results for a separate regression. All regressions include firm fixed effects, industry-year fixed effects, and the main regressor of interest interacted with year indicators. In panels A-B the main additional regressor is a measure of cash ratio, either as a binary indicator for above-median ratio (panel A) or as a continuous measure (panel B). In panel C, the main regressor of interest is an indicator for high share of long-term debt maturing soon, which takes value 1 if less than half of all long-term debt of the firm can be accounted for by an increase in long-term debt in 2005. The main independent variable is the triple interaction of exposure to the boycott (treatment), having high leverage, and post-boycott period (year 2006). In all regressions, standard errors are clustered at the industry level.

Model Extension	(1) Δln(Sales)	(2 ΔNew Products (HS6)	(3) ΔNew Export Products (HS6)	(4) ln(1+ Investment)	(5) Δln(FTE Employment)	(6) Any Outsourcing
Panel A: Cash Ratio						
Treatment X 2006	-0.0073	0.4007*	0.2804*	0.1765	-0.0135	0.0125
	(0.016)	(0.197)	(0.159)	(0.113)	(0.017)	(0.012)
Treatment X 2006	0.0511**	-0.0925	0.0283	-0.0108	0.0254	-0.0259
X High Cash	(0.021)	(0.195)	(0.149)	(0.134)	(0.021)	(0.024)
Panel B: Continuous Cash/Liabilities						
Treatment X 2006	0.0190	0.3949*	0.3128*	0.1285*	-0.0156	0.0157
	(0.014)	(0.205)	(0.175)	(0.073)	(0.010)	(0.011)
Treatment X 2006	-0.0027	-0.0285	-0.0184	0.0348	0.0118**	-0.0126*
X Cash Ratio	(0.005)	(0.032)	(0.036)	(0.037)	(0.005)	(0.007)
Panel C: High Share of Long-Term Debt Maturing Soon						
Treatment X 2006	0.0238	0.4104*	0.3750**	0.1645*	0.0105	-0.0221
	(0.016)	(0.206)	(0.162)	(0.084)	(0.011)	(0.017)
Treatment X 2006	-0.0147	-0.0781	-0.1365	0.0079	-0.0220*	0.0426*
X Debt Maturing Soon	(0.019)	(0.213)	(0.193)	(0.083)	(0.012)	(0.021)

*** p<0.01, ** p<0.05, * p<0.1

with a lower share of recent increase in long-term debt face a higher liquidity constraint during the boycott, when a larger share of their debt will mature.

Panel C of Table 3.10 replaces the indicator for high leverage with an indicator for high stock of long-term debt that will mature soon during the boycott. We find a significant reduction in employment for firms with a higher share of maturing long-term debt. This response goes along with a significant increase in outsourcing of labor services. In contrast, those firms whose debt is not maturing soon, i.e. those who are less constrained, significantly increase investment, product innovation and new market entry. The noisy nature of our measure and the fact that long-term debt is only one part of total obligations of firms explains the smaller effects and reduced precision that we find in Panel C compared to the main analysis.

In sum, these exercises first highlight that leverage does not simply capture differences across exposed firms based on firm size, industry, product variety, or manager quality. Instead, consistent with the additional results on cash holdings and debt maturity, the effect of leverage is likely to reflect financial constraints.

3.6. Robustness

We base our results on a simple and intuitive research design: we categorize firms as low- or high-leverage and compare those who were exposed to the boycott, controlling for outcomes of firms with the same leverage status that were not exposed to the boycott. In this section, we present key robustness checks for our analysis related to the degree of exposure and the definition of leverage.

Table 3.11 reports the results of robustness checks for our main outcomes of interest: sales, export markets, product innovation, investment, employment, and outsourcing. Each panel of the table reports results for a separate specification.[22] In panels A-C, we use alternative measures of leverage. Panel A replaces the binary indicator for high- versus low-leverage before the boycott by the continuous measure of total debt over total assets before the boycott. This specification yields highly significant results that have similar magnitude as our main results. In panel B, we limit our attention to debt related to financial leverage, specifically using short- and long-term bank debt rather than total liabilities. The results are again similar to the main

[22] In Appendix Table 3.16 we further show that the results are quantitatively similar to the main findings in specifications for firm outcomes in levels, allowing for group-specific trends.

Table 3.11. Robustness Results

Each panel in the table reports results for a separate regression. All regressions present results analogous to the main specification, but with original leverage or exposure measure substituted with an alternative. Indicator for high total leverage is substituted with a continuous measure of leverage (panel A), indicator for high financial leverage (panel B), or indicator for high leverage defined using industry-level median (panel C). Indicator for any significant exposure to Muslim countries is substituted with a continous measure in exposure (as % of sales in panel D and % of exports in panel E). Panel F presents the specification in which continuous measure of exposure (% of sales) is included as a control variable. In panel G, the indicator of exposure uses only Arab countries for defining boycott-affected export destinations. The main independent variable is the triple interaction of exposure to the boycott (treatment), having high leverage, and post-boycott period (year 2006). In all regressions, standard errors are clustered at the industry level.

Model Extension		(1) Δln(Sales)	(2 ΔNew Products (HS6)	(3) ΔNew Export Products (HS6)	(4) ln(1+ Investment)	(5) Δln(FTE Employment)	(6) Any Outsourcing
Panel A:	Treated X 2006	0.0312	1.1028***	0.9159***	0.3665*	0.0485***	-0.0524*
Continuous		(0.037)	(0.375)	(0.299)	(0.185)	(0.014)	(0.028)
Leverage	Treated X 2006	-0.0242	-1.2234**	-1.0349***	-0.2911	-0.0770***	0.0824**
	X Leverage	(0.052)	(0.458)	(0.373)	(0.248)	(0.023)	(0.030)
Panel B:	Treated X 2006	0.0228	0.4817**	0.3824**	0.2407***	0.0113	-0.0049
Financial		(0.014)	(0.209)	(0.163)	(0.086)	(0.009)	(0.019)
Leverage	Treated X 2006	-0.0182	-0.2519*	-0.1921	-0.1510	-0.0312**	0.0159
	X High Leverage	(0.019)	(0.143)	(0.143)	(0.179)	(0.015)	(0.029)
Panel C:	Treated X 2006	0.0303**	0.5547**	0.4380**	0.2770***	0.0246***	-0.0129
Leverage		(0.012)	(0.228)	(0.187)	(0.094)	(0.006)	(0.018)
by Industry	Treated X 2006	-0.0344	-0.4711**	-0.3627**	-0.2117	-0.0591***	0.0318*
	X High Leverage	(0.023)	(0.177)	(0.139)	(0.149)	(0.014)	(0.018)
Panel D:	Exposure X 2006	0.1281***	0.3869	0.3338	0.1383	0.0114	-0.0375
Continuous		(0.044)	(0.729)	(0.572)	(0.169)	(0.048)	(0.032)
Exposure	Exposure X 2006	-0.1884*	-0.3914	-0.5156*	-0.2623	-0.1916***	0.0500
(% Exports)	X High Leverage	(0.098)	(1.082)	(0.261)	(0.333)	(0.047)	(0.045)
Panel E:	Exposure X 2006	0.4600***	0.3316	1.1690	0.5750	-0.1475	-0.1668
Continuous		(0.150)	(1.858)	(1.157)	(0.556)	(0.099)	(0.123)
Exposure	Exposure X 2006	-0.7678*	-2.4239	-2.3800***	-1.7602	-0.7495	0.2208
(% Sales)	X High Leverage	(0.398)	(2.149)	(0.811)	(1.513)	(0.714)	(0.294)
Panel F:	Treated X 2006	0.0347**	0.6435***	0.4861**	0.3264***	0.0466***	-0.0181
Control for		(0.015)	(0.230)	(0.194)	(0.104)	(0.014)	(0.019)
Exposure	Treated X 2006	-0.0516*	-0.4873**	-0.3916**	-0.2415	-0.0614***	0.0503**
	X High Leverage	(0.026)	(0.192)	(0.143)	(0.159)	(0.019)	(0.020)
Panel G:	Treated X 2006	0.0567***	0.5019*	0.4274*	0.3143**	0.0195***	-0.0088
Arab		(0.017)	(0.285)	(0.219)	(0.140)	(0.006)	(0.021)
Countries	Treated X 2006	-0.0739*	-0.3581	-0.3273*	-0.1347	-0.0375	0.0391
as Treatment	X High Leverage	(0.038)	(0.244)	(0.181)	(0.195)	(0.022)	(0.028)

specification. Panel C defines high- versus low-leverage within industry, and shows that the main results are not explained by differences in leverage across industries.

Panels D-F focus on the size of the shock. In panel D, we replace the binary treatment indicator for boycott exposure with a continuous measure of the share of exports to Muslim countries in 2005. The results show that more exposed firms see higher decreases in sales and employment, and introduce fewer new export products. The point estimates for the number of new products, investment, and outsourcing are also similar to those from the main specification, but these results are not precisely estimated. Panel E replicates the results of Panel D with an alternative measure of exposure, i.e., the share of export to Muslim countries in total sales (as opposed to total exports). The typical values for this measure are very different (the average share in total exports is 21% versus share in sales of 3.9%, obtained after dropping outlier firms with very large exposures), which changes the magnitudes of the coefficients. Yet, the effects of one standard deviation change remain similar, which confirms that the choice of exposure measure is not driving our results. In Panel F, we include the continuous exposure measure as a control variable, and the results are unchanged. Finally, in Panel G, we redefine the treatment indicator by using only previous exports to Arab countries. This specification ignores informal boycotts in many other countries, thereby slightly understating the main effects and reducing precision.

3.7. Conclusion

This paper shows how financial leverage influences firms' responses to an unexpected demand shock. Our results highlight the importance of capital structure in determining the ability to adapt to a changing environment and, as a consequence, the importance of flexibility considerations in firms' capital structure decisions. Responding to the boycott, low-leverage firms are able to increase their investment and add new product categories in their non-Muslim destination markets to counteract the negative demand shock, thereby avoiding a decrease in sales and employment. For high-leverage firms, the innovation response is muted and total sales and employment decline. Instead, highly leveraged firms attempt to increase their operational flexibility by substituting their employees with outsourcing.

We conclude with two limitations of our study: First, an important missing element which determines the relationship between capital structure and flexibility is the ability for debt renegotiation. When a firm is hit by a negative shock and needs to flexibly adjust, are creditors willing to renegotiate their debt contracts? Data limitations prevent us from analyzing this question in this study, but more work is needed to extend existing evidence (Gilson et al., 1990). Second, we note that all our results about firms' adaptation to this shock capture short-term responses within one year after the boycott started. Even though the lack of innovation response suggests adverse effects in the longer term, we do not take a strong stand on the long-run costs or benefits of these differential adjustments. It is possible that financial distress forces firms to engage in necessary change (Jensen, 1986), which can benefit them in the long run. Boycott-driven structural change is confirmed by Friedrich (2016), who shows that a small share of exposed firms decide to de-layer and systematically change their occupational hierarchy, and internal wage structure. We view these questions about the relationship between financial distress, organizational change and persistent consequences on competitive advantage as a promising avenue for future research.

Appendix

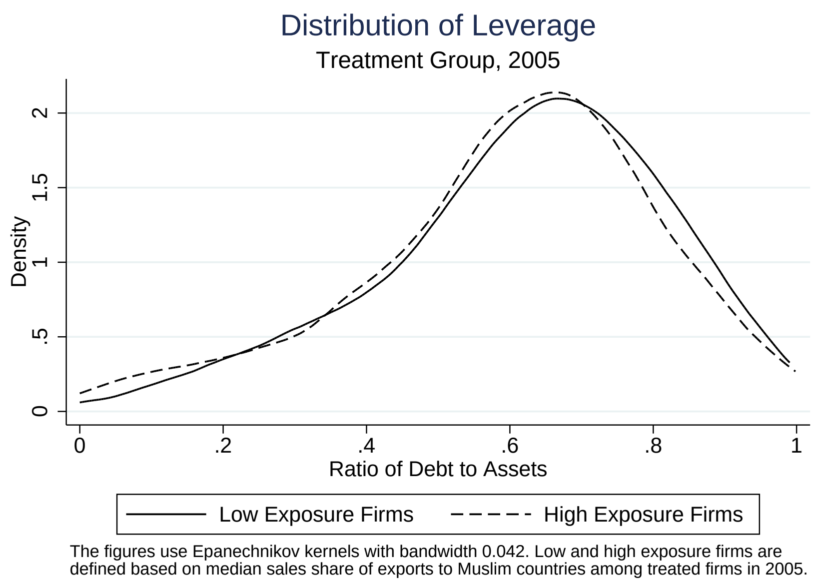

Figure 3.8. Distribution of Leverage for Firms with Low and High Exposure to the Boycott

The figure shows the distribution of leverage (defined as share of total liabilities in total assets in 2005) for firms exposed to the boycott. High exposure refers to firm with above-median share of exports sold to Muslim countries in 2005.

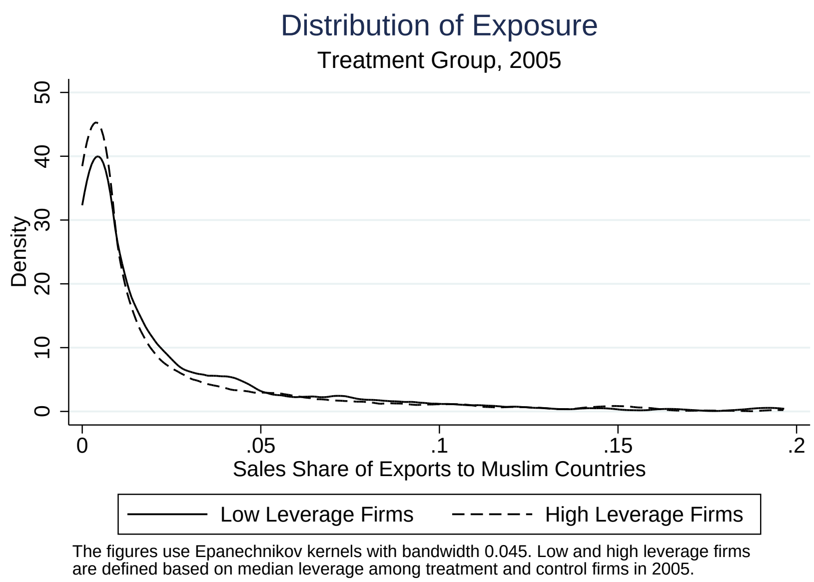

Figure 3.9. Distribution of Exposure to Muslim Countries (Measured as Share of Exports to Muslim Countries in Sales) for Firms with Low Leverage

The figure shows the distribution of an alternative measure of exposure to Muslim countries: share of exports to Muslim countries in total sales. High leverage corresponds to above-median value of share of total liabilities and total assets. Exposure and leverage are calculated based on values from 2005.

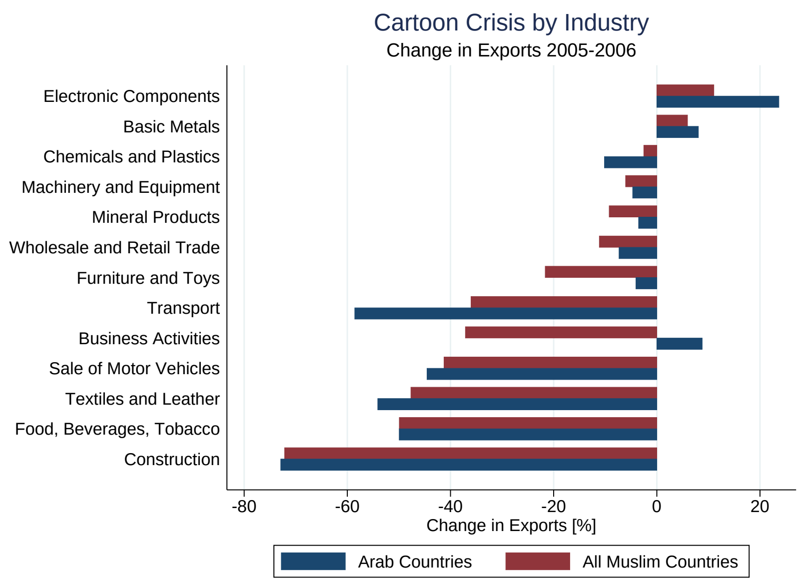

Figure 3.10. Boycott across Industries
The figure shows the relative change in the value of total export to countries affected by the boycott by industry. Red bars define affected countries as the set of countries with more than 50% Muslims. Blue bars define affected countries as Arab countries.

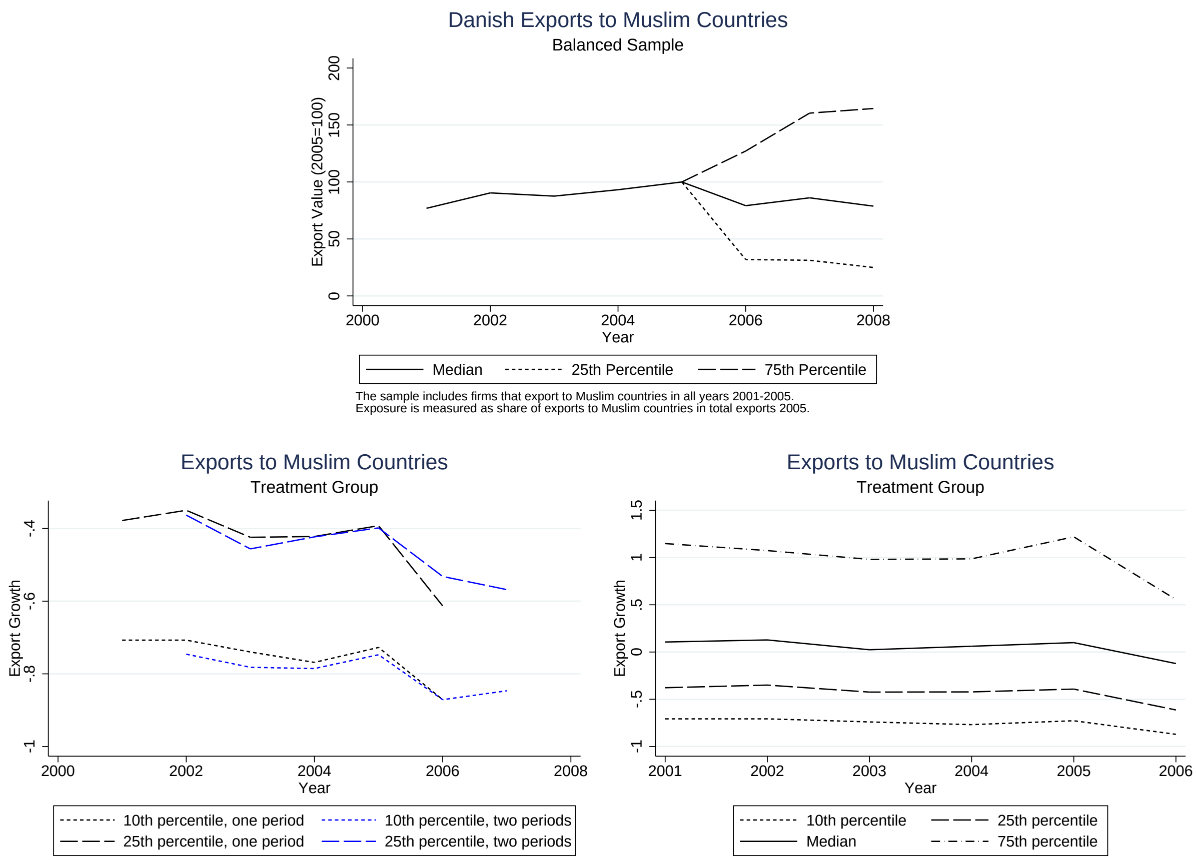

Figure 3.11. Exports to Muslim Countries for Danish Firms

The figure shows patterns of export evolution for Danish firms affected by the boycott. The top panel shows the evolution of 25th, 50th and 75th percentile of firms (computed based on the value of exports to Muslim countries in 2007; 2005 value normalized to 100) for firms with stable presence in Muslim countries, defined as positive export in each year during 2001-2005 period. The middle panel shows the change in exports relative to one or two years ago for 10th and 25th percentile of firms (sorted based on the magnitude of export drop in each year). The bottom panel shows the year-to-year growth of firms' export to Muslim countries for various percentiles of firm (defined based on the magnitude of export change in each year).

Table 3.12. Redirecting Sales: New Export Markets and Products (Level Specification)

All regressions include firm fixed effects, industry-year fixed effects, and binary variables for each year interacted with indicators for high leverage. The main independent variable is triple interaction of exposure to the boycott (treatment), having high leverage and post-boycott period (year 2006). Dependent variables are changes of number of export destinations, and number of exported products defined as non-zero flows in 6-, 4- and 2-digits product category in the HS system. The bottom row presents the mean of the level of dependent variables in the pre-boycott period. In all regressions, standard errors are clustered at the industry level. Joint p-val row presents a p-value from the F-test for significance of the high-leverage firms' response (the sum of baseline coefficient for low-leverage firms and differential effect for high-leverage firms).

	(1)	(2)	(3)	(4)
	ΔExport Destinations	ΔNum Exp Products (6-digit)	ΔNum Exp Exp Products (4-digit)	ΔNum Exp Exp Products (2-digit)
Treated X 2006	-1.5318***	1.3201*	1.2783***	0.4108***
	(0.285)	(0.666)	(0.382)	(0.107)
Treated X High X 2006	0.1699	-0.6954	-0.7997	-0.2478*
	(0.361)	(0.764)	(0.482)	(0.133)
Observations	53,910	53,910	53,910	53,910
R-squared	0.032	0.017	0.026	0.032
Firms	13,307	13,307	13,307	13,307
Joint p-val	3.66e-05	0.156	0.152	0.0597
Sample Avg 01-05	7.810	11	7.878	4.149

*** p<0.01, ** p<0.05, * p<0.1

Table 3.13. Outsourcing Response - Detailed Analysis

All regressions include firm fixed effects, industry-year fixed effects, and binary variables for each year interacted with indicators for high leverage. The main independent variable is triple interaction of exposure to the boycott (treatment), having high leverage, and post-boycott period (year 2006). The data sample is limited to 1,221 firms for which the results of outsourcing survey are available. We define 12 categories of outsourced services by grouping several related activity codes. The bottom row presents averages of dependent variables in the pre-boycott period. In all regressions, standard errors are clustered at the industry level. Joint p-val row presents a p-value from the F-test for significance of the high-leverage firms' response (the sum of baseline coefficient for low-leverage firms and differential effect for high-leverage firms).

	(1) Transport	(2) ICT	(3) Accounting & Legal	(4) Engineering	(5) Marketing	(6) HR & Training
Treated X 2006	0.0072*	-0.0067	-0.0019	-0.0170	-0.0232	-0.0233**
	(0.004)	(0.009)	(0.011)	(0.027)	(0.018)	(0.011)
Treated X High X 2006	0.0247*	0.0429**	0.0522**	0.0120	0.0541***	0.0514
	(0.014)	(0.020)	(0.020)	(0.046)	(0.018)	(0.037)
R-squared	0.028	0.040	0.024	0.023	0.032	0.028
Joint p-val	0.026	0.035	0.015	0.882	0.003	0.401
Sample Avg 01-05	0.985	0.965	0.965	0.614	0.930	0.870

	(7) Security	(8) Cleaning	(9) Food	(10) Consulting	(11) Construction & Repairs	(12) Sales Commission
Treated X 2006	0.0261	0.0117	0.0302	-0.0112	-0.0018	0.0001
	(0.019)	(0.017)	(0.022)	(0.034)	(0.007)	(0.021)
Treated X High X 2006	-0.0022	-0.0023	-0.0313	0.0754	0.0381	0.0208
	(0.041)	(0.026)	(0.040)	(0.046)	(0.023)	(0.038)
R-squared	0.022	0.030	0.030	0.029	0.026	0.021
Joint p-val	0.513	0.663	0.973	0.039	0.074	0.501
Sample Avg 01-05	0.671	0.900	0.614	0.658	0.972	0.383
Firms	1,230	1,230	1,230	1,230	1,230	1,230
Observations	5,309	5,309	5,309	5,309	5,309	5,309

*** p<0.01, ** p<0.05, * p<0.1

Table 3.14. Relationships in the Supply Chain

All regressions include firm fixed effects, industry-year fixed effects, and binary variables for each year interacted with indicators for high leverage. The main independent variable is triple interaction of exposure to the boycott (treatment), having high leverage, and post-boycott period (year 2006). Dependent variables are log-changes of short term debt (to suppliers and other creditors); log-changes of receivables (current, long-term, and subgroups of current receivables: work in progress, finished goods, and other); and log-change of inventories. The bottom row presents mean of levels of dependent variables in the pre-boycott period (in thousands of DKK). In all regressions, standard errors are clustered at the industry level. Joint p-val row presents a p-value from the F-test for significance of the high-leverage firms' response (the sum of baseline coefficient for low-leverage firms and differential effect for high-leverage firms).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Δln(Short-Term Debt)				Δln(Receivables)			Δln(Inventory)
	To Suppliers	To Other	Current	Long-Term	Ongoing Work	Goods	Other	Total
Treatment X	0.0816**	-0.0005	0.0748**	0.0609	0.2925***	0.0851**	0.0394	0.0493
2006	(0.033)	(0.029)	(0.030)	(0.100)	(0.068)	(0.034)	(0.060)	(0.041)
Treatment X	-0.0930***	-0.0672	-0.0901**	-0.1865*	-0.0793	-0.1570***	0.0166	-0.1062
High X 2006	(0.032)	(0.041)	(0.035)	(0.109)	(0.185)	(0.044)	(0.047)	(0.064)
Observations	76,110	76,263	76,711	47,357	18,467	75,060	74,966	43,480
R-squared	0.020	0.015	0.008	0.082	0.109	0.009	0.021	0.040
Firms	15,189	15,201	15,203	13,327	7,844	15,136	15,162	13,601
Joint p-val	0.716	0.0548	0.623	0.159	0.211	0.0102	0.347	0.510
Sample	7266	20318	20121	3795	2702	11284	8567	7451
Avg 01-05								

*** p<0.01, ** p<0.05, * p<0.1

Table 3.15. Response of Liabilities: Placebo Test

All regressions include firm fixed effects, industry-year fixed effects, and binary variables for each year interacted with indicators for high leverage. In contrast to the main analysis, the sample period is 2001-2004 and high leverage is defined based on the debt to assets ratio at the end of 2002. The treatment indicator is unchanged from the main results, i.e. the treatment group includes all firms with non-negligible exports to Muslim countries in 2005. The main independent variable is triple interaction of exposure to the boycott (treatment), having high leverage, and a dummy for 2004. Dependent variables are the log-changes in total debt (column 1), short-term debt to suppliers (column 2), other short-term debt (column 3), long-term debt (column 4), and 1 + long-term debt (column 5). The bottom row presents the mean level of the dependent variables over 2001-2003 (in millions of DKK). In all regressions, standard errors are clustered at the industry level (53 industries). Joint p-val row presents a p-value from the F-test for significance of the high-leverage firms' response (the sum of baseline coefficient for low-leverage firms and differential effect for high-leverage firms).

	(1)	(2)	(3)	(4)	(5)
	Δln(Debt)	Δln(Short-Term Debt)		Δln(Long -Term Debt)	Δln(1+Long -Term Debt)
		To Suppliers	To Other		
Treatment X 2004	-0.0428*	-0.0232	-0.0215	-0.0632	-0.1808
	(0.025)	(0.036)	(0.030)	(0.045)	(0.125)
Treatment X High 2002 X 2004	0.0502	0.0769*	0.0440	0.0628	-0.0378
	(0.048)	(0.044)	(0.058)	(0.061)	(0.206)
Obs	52,276	51,808	51,889	31,131	52,287
R-squared	0.017	0.024	0.015	0.010	0.004
Firms	14,613	14,592	14,598	10,516	14,615
Joint p-val	0.831	0.153	0.609	0.992	0.075
Sample Avg 01-03	37.6	7.1	19.3	18.1	11.2

*** p<0.01, ** p<0.05, * p<0.1

Table 3.16. Robustness Results: Level Specifications

Each panel in the table reports results for a separate regression. All regressions present results analogous to the main specification, but using outcomes in levels and controlling for group-specific linear trends. The main independent variable is the triple interaction of exposure to the boycott (treatment), having high leverage, and post-boycott period (year 2006). In all regressions, standard errors are clustered at the industry level.

	(1) ln(Exports)	(2) ln(Sales)	(3) ln(Debt)	(4) New Products (HS6)	(5) New Export Products (HS6)	(6) ln(1+ Investment)	(7) Δln(FTE Employment)	(8) Any Outsourcing
Treated X 2006	-0.1701**	0.0085	0.0139	0.4412**	0.3503*	0.2936***	-0.011	-0.0086
	(0.049)	(0.014)	(0.020)	(0.206)	(0.190)	(0.119)	(0.014)	(0.021)
Treated X 2006	-0.0619	-0.0373	-0.0324	-0.2754*	-0.2595*	-0.2383	-0.0458**	0.0505**
X Leverage	(0.052)	(0.045)	(0.033)	(0.157)	(0.149)	(0.158)	(0.023)	(0.020)
Observations	61,625	76,826	76,805	10,258	10,258	76,826	76,826	76,826
R-Squared	0.009	0.048	0.031	0.021	0.021	0.047	0.036	0.032
Joint p-val	0.000	0.444	0.461	0.177	0.339	0.716	0.002	0.063

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 3.17. Coefficients from Figures from the Main Text

The table presents the coefficients depicted in the figures in the main text. All regressions include firm fixed effects, industry-year fixed effects, and binary variables for each year interacted with indicators for high leverage. The main independent variables are triple interaction of exposure to the boycott (treatment), having high or low leverage, and the post-boycott period (year 2006). Dependent variables are changes in the number of new products and newly exported product, value of investment as % of lagged assets, log-change of total wage bill, an indicator for any outsourcing, and an indicator for any leasing (operational or financial). In all regressions, standard errors are clustered at the industry level.

	(1)	(2)	(3)	(4)
	ΔNum New Prod (6-digit)	ΔNum New Exp Prod (6-digit)	Δlog(Wages Total)	Any Outsourcing
Treatment X 2001	0.0289	0.0411	0.0186	0.0417
X Low Leverage	(0.268)	(0.181)	(0.015)	(0.049)
Treatment X 2002	0.2960	0.2873	-0.0114	-0.0131
X Low Leverage	(0.361)	(0.284)	(0.015)	(0.022)
Treatment X 2003	-0.0635	-0.0768	-0.0146	-0.0263*
X Low Leverage	(0.260)	(0.168)	(0.016)	(0.015)
Treatment X 2004	0.0570	0.1649	0.0303*	-0.0196
X Low Leverage	(0.102)	(0.105)	(0.015)	(0.023)
Treatment X 2006	0.6067**	0.5194**	0.0126	-0.0257*
X Low Leverage	(0.263)	(0.214)	(0.013)	(0.015)
Treatment X 2001	0.1440	0.0672	-0.0336**	0.0066
X High Leverage	(0.200)	(0.125)	(0.014)	(0.041)
Treatment X 2002	0.4875***	0.2917***	0.0029	-0.0161
X High Leverage	(0.147)	(0.100)	(0.018)	(0.030)
Treatment X 2003	0.0577	0.0306	0.0024	-0.0057
X High Leverage	(0.141)	(0.111)	(0.018)	(0.019)
Treatment X 2004	-0.2093	-0.0204	0.0129	0.0023
X High Leverage	(0.124)	(0.082)	(0.019)	(0.022)
Treatment X 2006	0.1459	0.1173	-0.0366***	0.0273*
X High Leverage	(0.142)	(0.131)	(0.010)	(0.016)
Observations	10,258	10,258	74,826	76,826
R-squared	0.028	0.030	0.015	0.032
Firms	2369	2369	14963	15208

*** p<0.01, ** p<0.05, * p<0.1

References

Abeliansky, A. and K. Prettner (2017). Automation and demographic change. Working Paper.

Acemoglu, D. and D. Autor (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of Labor Economics*, Volume 4, pp. 1043–1171. Elsevier.

Acemoglu, D., D. Dorn, G. H. Hanson, B. Price, et al. (2014). Return of the solow paradox? IT, productivity, and employment in US manufacturing. *American Economic Review 104*(5), 394–99.

Acemoglu, D. and P. Restrepo (2018a). Demographics and automation. Working Paper.

Acemoglu, D. and P. Restrepo (2018b). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review 108*(6), 1488–1542.

Acemoglu, D. and P. Restrepo (2019). Robots and jobs: Evidence from US labor markets. Forthcoming in Journal of Political Economy.

Acharya, V. and Z. Xu (2017). Financial dependence and innovation: The case of public versus private firms. *Journal of Financial Economics 124*(2), 223–243.

Acharya, V. V., R. P. Baghai, and K. V. Subramanian (2013). Wrongful discharge laws and innovation. *The Review of Financial Studies 27*(1), 301–346.

Agarwal, S., G. Amromin, I. Ben-David, S. Chomsisengphet, T. Piskorski, and A. Seru (2017). Policy intervention in debt renegotiation: Evidence from the home affordable modification program. *Journal of Political Economy 125*(3), 654–712.

Agrawal, A. K., J. Gans, and A. Goldfarb (Eds.) (2019). *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Alfaro, I., N. Bloom, and X. Lin (2018). The finance uncertainty multiplier. Technical report, National Bureau of Economic Research.

Almeida, H., M. Campello, B. Laranjeira, and S. Weisbenner (2011). Corporate debt maturity and the real effects of the 2007 credit crisis. *Critical Finance Review 1*(1), 3–58.

Amiti, M. and D. E. Weinstein (2011). Exports and financial shocks. *The Quarterly Journal of Economics 126*(4), 1841–1877.

Amore, M. D., C. Schneider, and A. Žaldokas (2013). Credit supply and corporate innovation. *Journal of Financial Economics 109*(3), 835–855.

AMRON, C. (2009). AMRON-SARFiN report on real estate and mortgage market in poland.

AMRON, C. (2015). AMRON-SARFiN report on real estate and mortgage market in poland.

Andrade, G. and S. N. Kaplan (1998). How costly is financial (not economic) distress? evidence from highly leveraged transactions that became distressed. *The Journal of Finance 53*(5), 1443–1493.

Asquith, P., R. Gertner, and D. Scharfstein (1994). Anatomy of financial distress: An examination of junk-bond issuers. *The Quarterly Journal of Economics 109*(3), 625–658.

Autor, D. (2019). Work of the past, work of the future. Working Paper.

Autor, D. H., F. Levy, and R. J. Murnane (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics 118*(4), 1279–1333.

Autor, David, H., D. Dorn, and G. H. Hanson (2013). The China syndrome: Local labor market effects of import competition in the United States. *American Economic Review 103*(6), 2121–68.

Babina, T. and S. T. Howell (2019). Innovation investment and labor mobility: Employee entrepreneurship from corporate R&D. Working Paper.

Baghai, R., R. Silva, V. Thell, and V. Vig (2020). Talent in distressed firms: Investigating the labor costs of financial distress. *Available at SSRN 2854858*.

Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics 131*(4), 1593–1636.

Bednarzik, R., A. Kern, and J. J. Hisnanick (2017). Displacement and debt: The role of debt in returning to work in the period following the great recession. *Working Paper*.

Bena, J. and K. Li (2014). Corporate innovations and mergers and acquisitions. *The Journal of Finance 69*(5), 1923–1960.

Bena, J., H. Ortiz-Molina, and E. Simintzi (2018). Shielding firm value: Employment protection and process innovation. Working Paper.

Bernanke, B. S. and M. Gertler (1995). Inside the black box: the credit channel of monetary policy transmission. *Journal of Economic perspectives 9*(4), 27–48.

Bernstein, A. (2018). Household debt overhang and labor supply. *Working Paper*.

Bessen, J. E. (2016). How computer automation affects occupations: Technology, jobs, and skills. Working Paper.

Bessen, J. E., M. Goos, A. Salomons, and W. Van den Berge (2019). Automatic reaction-what happens to workers at firms that automate? Working Paper.

Bloom, N., M. Draca, and J. Van Reenen (2016). Trade induced technical change? the impact of Chinese imports on innovation, IT and productivity. *The Review of Economic Studies 83*(1), 87–117.

Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016). Consumption inequality and family labor supply. *American Economic Review 106*(2), 387–435.

Börsch-Supan, A. (2003). Labor market effects of population aging. *Labour 17*, 5–44.

Bresnahan, T. F., E. Brynjolfsson, and L. M. Hitt (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *The Quarterly Journal of Economics 117*(1), 339–376.

Brown, J. and D. A. Matsa (2016a). Boarding a sinking ship? an investigation of job applications to distressed firms. *The Journal of Finance 71*(2), 507–550.

Brown, J. and D. A. Matsa (2016b). Locked in by leverage: Job search during the housing crisis. *Working Paper*.

Brynjolfsson, E. and L. M. Hitt (2000). Beyond computation: Information technology, organizational transformation and business performance. *Journal of Economic Perspectives 14*(4), 23–48.

Brynjolfsson, E. and A. McAfee (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & Company.

Brynjolfsson, E., D. Rock, and C. Syverson (2019). Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. In *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Buchak, G., G. Matvos, T. Piskorski, and A. Seru (2018). Fintech, regulatory arbitrage, and the rise of shadow banks. *Journal of Financial Economics 130*(3), 453–483.

Busse, M. (2002). Firm financial condition and airline price wars. *RAND Journal of Economics 33*(2), 298–318.

Caggese, A., V. Cuñat, and D. Metzger (2019). Firing the wrong workers: Financing constraints and labor misallocation. *Journal of Financial Economics 133*(3), 589–607.

Calza, A., T. Monacelli, and L. Stracca (2013). Housing finance and monetary policy. *Journal of the European Economic Association 11*(suppl_1), 101–122.

Cesarini, D., E. Lindqvist, M. J. Notowidigdo, and R. Östling (2017). The effect of wealth on individual and household labor supply: evidence from Swedish lotteries. *American Economic Review 107*(12), 3917–46.

Chaboud, A. P., B. Chiquoine, E. Hjalmarsson, and C. Vega (2014). Rise of the machines: Algorithmic trading in the foreign exchange market. *The Journal of Finance 69*(5), 2045–2084.

Chan, M. (2017). How substitutable are labor and intermediates? *Unpublished working paper. University of Minnesota*.

Chava, S., A. Oettl, A. Subramanian, and K. V. Subramanian (2013). Banking deregulation and innovation. *Journal of Financial economics 109*(3), 759–774.

Chava, S. and M. R. Roberts (2008). How does financing impact investment? the role of debt covenants. *The journal of finance 63*(5), 2085–2121.

Chen, Z., J. Harford, and A. Kamara (2019). Operating leverage, profitability, and capital structure. *Journal of financial and quantitative analysis 54*(1), 369–392.

Cheng, H., R. Jia, D. Li, and H. Li (2019). The rise of robots in China. *Journal of Economic Perspectives 33*(2), 71–88.

Chetty, R. and A. Szeidl (2007). Consumption commitments and risk preferences. *The Quarterly Journal of Economics 122*(2), 831–877.

Chevalier, J. A. (1995a). Capital structure and product-market competition: Empirical evidence from the supermarket industry. *The American Economic Review*, 415–435.

Chevalier, J. A. (1995b). Do lbo supermarkets charge more? an empirical analysis of the effects of lbos on supermarket pricing. *The Journal of Finance 50*(4), 1095–1112.

Chodorow-Reich, G. (2014). The employment effects of credit market disruptions: Firm-level evidence from the 2008–9 financial crisis. *The Quarterly Journal of Economics 129*(1), 1–59.

Clemens, M. A., E. G. Lewis, and H. M. Postel (2018). Immigration restrictions as active labor market policy: Evidence from the mexican bracero exclusion. *American Economic Review 108*(6), 1468–87.

Cloyne, J., C. Ferreira, and P. Surico (2020). Monetary policy when households have debt: new evidence on the transmission mechanism. *The Review of Economic Studies 87*(1), 102–129.

Cohn, J. B. and M. I. Wardlaw (2016). Financing constraints and workplace safety. *The Journal of Finance 71*(5), 2017–2058.

Cornaggia, J., Y. Mao, X. Tian, and B. Wolfe (2015). Does banking competition affect innovation? *Journal of financial economics 115*(1), 189–209.

D'Acunto, F., R. Liu, C. Pflueger, and M. Weber (2018). Flexible prices and leverage. *Journal of Financial Economics 129*(1), 46–68.

D'Acunto, F., N. Prabhala, and A. G. Rossi (2019). The promises and pitfalls of robo-advising. *The Review of Financial Studies 32*(5), 1983–2020.

Dauth, W., S. Findeisen, J. Suedekum, and N. Woessner (2018). Adjusting to robots: worker-level evidence. Working Paper.

Davies, S. (1979). *The diffusion of process innovations*. CUP Archive.

Del Boca, D. and A. Lusardi (2003). Credit market constraints and labor market decisions. *Labour Economics 10*(6), 681–703.

Desai, M. A., C. F. Foley, and K. J. Forbes (2008). Financial constraints and growth: Multinational and local firm responses to currency depreciations. *The Review of Financial Studies 21*(6), 2857–2888.

Deshpande, M. (2016). The effect of disability payments on household earnings and income: Evidence from the ssi children's program. *Review of Economics and Statistics 98*(4), 638–654.

Di Maggio, M., A. Kermani, B. J. Keys, T. Piskorski, R. Ramcharan, A. Seru, and V. Yao (2017). Interest rate pass-through: Mortgage rates, household consumption, and voluntary deleveraging. *American Economic Review 107*(11), 3550–88.

Doms, M., T. Dunne, and K. R. Troske (1997). Workers, wages, and technology. *The Quarterly Journal of Economics 112*(1), 253–290.

Donaldson, J. R., G. Piacentino, and A. Thakor (2019). Household debt overhang and unemployment. *The Journal of Finance 74*(3), 1473–1502.

Dotan, A. and S. A. Ravid (1985). On the interaction of real and financial decisions of the firm under uncertainty. *The Journal of Finance 40*(2), 501–517.

Ellguth, P. and I. Kohaut, Susanne; Moller (2014). The IAB Establishment Panel - methodological essentials and data quality. *Journal for Labour Market Research 47*(1-2), 27–41.

Fahn, M., V. Merlo, and G. Wamser (2017). Relational contracts, the cost of enforcing formal contracts, and capital structure choice-theory and evidence. *CESifo Working Paper*.

Fazzari, S. M., R. G. Hubbard, and B. C. Petersen (1988). Financing constraints and corporate investment. *Brookings Papers on Economic Activity 1988*(1), 141–206.

Flodén, M., M. Kilström, J. Sigurdsson, and R. Vestman (2017). Household debt and monetary policy: Revealing the cash-flow channel.

Foley, C. F. and K. Manova (2015, August). International Trade, Multinational Activity, and Corporate Finance. *Annual Review of Economics 7*(1), 119–146.

Fort, T. C., J. R. Pierce, and P. K. Schott (2018). New perspectives on the decline of US manufacturing employment. *Journal of Economic Perspectives 32*(2), 47–72.

Fortin, N. M. (1995). Allocation inflexibilities, female labor supply, and housing assets accumulation: Are women working to pay the mortgage? *Journal of Labor Economics 13*(3), 524–557.

Fos, V., A. Liberman, and C. Yannelis (2017). Debt and human capital: Evidence from student loans. *Available at SSRN 2901631*.

Fresard, L. (2010). Financial strength and product market behavior: The real effects of corporate cash holdings. *The Journal of finance 65*(3), 1097–1122.

Friedrich, B. U. (2016). Trade shocks, firm hierarchies and wage inequality. *Unpublished working paper. Yale University*.

Garriga, C., F. E. Kydland, and R. Šustek (2017). Mortgages and monetary policy. *The Review of Financial Studies 30*(10), 3337–3375.

Gilson, S. C., K. John, and L. H. Lang (1990). Troubled debt restructurings: An empirical study of private reorganization of firms in default. *Journal of financial economics 27*(2), 315–353.

Giroud, X. and H. M. Mueller (2016). Firm leverage, consumer demand, and employment losses during the great recession. *The Quarterly Journal of Economics*.

Graetz, G. and G. Michaels (2018). Robots at work. *Review of Economics and Statistics 100*(5), 753–768.

Graham, J. R. and C. R. Harvey (2001). The theory and practice of corporate finance: evidence from the field. *Journal of Financial Economics 60*(2-3), 187–243.

Granja, J. and S. Moreira (2020). Product innovation and credit market disruptions. *Available at SSRN 3477726*.

Gross, D. B. and N. S. Souleles (2002). Do liquidity constraints and interest rates matter for consumer behavior? evidence from credit card data. *The Quarterly journal of economics 117*(1), 149–185.

Gu, L., D. Hackbarth, and T. Li (2018). Inflexibility and leverage. *Available at SSRN 3296926*.

Han Kim, E., H. Kim, Y. Li, Y. Lu, and X. Shi (2019). How do equity offerings affect technology adoption, employees and firm performance? Working Paper.

Heckman, J. (1974). Life cycle consumption and labor supply: An explanation of the relationship between income and consumption over the life cycle. *The American Economic Review 64*(1), 188–194.

Hedlund, A., F. Karahan, K. Mitman, and S. Ozkan (2017). Monetary policy, heterogeneity, and the housing channel.

Hiller, S., P. J. Schröder, and A. Sø rensen (2014). Foreign demand shifts and the reallocation of exports - evidence from the danish cartoon crisis. *Aarhus University*.

Hombert, J. and A. Matray (2017). The real effects of lending relationships on innovative firms and inventor mobility. *The Review of Financial Studies 30*(7), 2413–2445.

Huber, K. (2018). Disentangling the effects of a banking crisis: evidence from German firms and counties. *American Economic Review 108*(3), 868–98.

Iacoviello, M. (2005). House prices, borrowing constraints, and monetary policy in the business cycle. *American economic review 95*(3), 739–764.

Imbens, G. W., D. B. Rubin, and B. I. Sacerdote (2001). Estimating the effect of unearned income on labor earnings, savings, and consumption: Evidence from a survey of lottery players. *American economic review 91*(4), 778–794.

Jäger, K. (2016). Eu klems growth and productivity accounts 2017 release, statistical module1. *Description of methodology and country notes for Spain*.

Jappelli, T. and L. Pistaferri (2010). The consumption response to income changes. *Annu. Rev. Econ. 2*(1), 479–506.

Jappelli, T. and A. Scognamiglio (2018). Interest rate changes, mortgages, and consumption: evidence from italy. *Economic Policy 33*(94), 183–224.

Jensen, H. (2008). The mohammed cartoons controversy and the boycott of danish products in the middle east. *European Business Review 20*(3), 275–289.

Jensen, M. C. (1986). Agency costs of free cash flow, corporate finance, and takeovers. *American Economic Review 76*(2), 323–329.

Kanatas, G. and J. Qi (2016). Outsourcing and financing decisions in industry equilibrium. *Review of Finance 20*(6), 2247–2271.

Kartashova, K. and X. Zhou (2020). How do mortgage rate resets affect consumer spending and debt repayment? evidence from canadian consumers.

Kashyap, A. K. and J. C. Stein (2000). What do a million observations on banks say about the transmission of monetary policy? *American Economic Review 90*(3), 407–428.

Kerr, W. R. and R. Nanda (2015). Financing innovation. *Annual Review of Financial Economics 7*, 445–462.

Kini, O., J. Shenoy, and V. Subramaniam (2017). Impact of financial leverage on the incidence and severity of product failures: Evidence from product recalls. *The Review of Financial Studies 30*(5), 1790–1829.

Koch, M., I. Manuylov, and M. Smolka (2019). Robots and firms. Working Paper.

Kopczuk, W. (2012). The Polish business "flat" tax and its effect on reported incomes: a Pareto improving tax reform. *Working Paper*.

Krieger, J., D. Li, and D. Papanikolaou (2018). Developing novel drugs. *Available at SSRN 3095246*.

Kuzmina, O. (2018). Employment flexibility and capital structure: Evidence from a natural experiment. *Columbia Business School Research Paper* (13-69).

La Cava, G., H. Hughson, G. Kaplan, et al. (2016). The household cash flow channel of monetary policy. *Reserve Bank of Australia, Research Discussion Papers 12*.

Lewis, E. (2011). Immigration, skill mix, and capital skill complementarity. *The Quarterly Journal of Economics 126*(2), 1029–1069.

MacKay, P. (2003). Real flexibility and financial structure: An empirical analysis. *The Review of Financial Studies 16*(4), 1131–1165.

Mandelker, G. N. and S. G. Rhee (1984). The impact of the degrees of operating and financial leverage on systematic risk of common stock. *Journal of Financial and Quantitative Analysis 19*(1), 45–57.

Manova, K. (2013). Credit constraints, heterogeneous firms, and international trade. *Review of Economic Studies 80*(2), 711–744.

Mao, Y. and J. J. Wang (2018). Access to finance and technological innovation: Evidence from antebellum America. *Kelley School of Business Research Paper* (18-21).

McElheran, K. (2018). Economic measurement of AI. Working Paper.

Medina, L. and F. Schneider (2017). Shadow economies around the world: New results for 158 countries over 1991-2015. *CESifo Working Paper*.

Mian, A., K. Rao, and A. Sufi (2013). Household balance sheets, consumption, and the economic slump. *The Quarterly Journal of Economics 128*(4), 1687–1726.

Mian, A. and A. Sufi (2014). What explains the 2007–2009 drop in employment? *Econometrica 82*(6), 2197–2223.

Michaels, G. and X. Zhi (2010, July). Freedom fries. *American Economic Journal: Applied Economics 2*(3), 256–81.

Modigliani, F. and M. H. Miller (1963). Corporate income taxes and the cost of capital: a correction. *The American Economic Review 53*(3), 433–443.

Moll, B., L. Rachel, and P. Restrepo (2019). Uneven growth: Automation impact on income and wealth inequality. Working Paper.

Moon, S. K. and G. M. Phillips (2020). Outsourcing supply contracts, human capital and firm capital structure. *Management Science Forthcoming*.

Novy-Marx, R. (2011). Operating leverage. *Review of Finance 15*(1), 103–134.

Okun, A. M. (1963). *Potential GNP: its measurement and significance*. Yale University, Cowles Foundation for Research in Economics.

Paravisini, D., V. Rappoport, P. Schnabl, and D. Wolfenzon (2015). Dissecting the effect of credit supply on trade: Evidence from matched credit-export data. *The Review of Economic Studies 82*(1), 333–359.

Peterson, M. A. (1994). Cash flow variability and firm's pension choice: A role for operating leverage. *Journal of Financial Economics 36*(3), 361–383.

Raj, M. and R. Seamans (2019). AI, labor, productivity, and the need for firm-level data. In *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Rajan, R. and L. Zingales (1998). Financial dependence and growth. *American Economic Review 88*(3), 559–86.

Reinartz, S. J. and T. Schmid (2016). Production flexibility, product markets, and capital structure decisions. *The Review of Financial Studies 29*(6), 1501–1548.

Rothstein, J. and C. E. Rouse (2011). Constrained after college: Student loans and early-career occupational choices. *Journal of Public Economics 95*(1), 149–163.

Schmucker, A., J. Eberle, A. Ganzer, J. Stegmaier, and M. Umkehrer (2018). Establishment history panel 1975-2016. FDZ Data Report.

Schwab, K. (2017). *The Fourth Industrial Revolution*. Crown Business.

Serfling, M. (2016). Firing costs and capital structure decisions. *The Journal of Finance 71*(5), 2239–2286.

Sharpe, S. A. (1994). Financial market imperfections, firm leverage, and the cyclicality of employment. *The American Economic Review 84*(4), 1060–1074.

Simintzi, E., V. Vig, and P. Volpin (2014). Labor protection and leverage. *The Review of Financial Studies 28*(2), 561–591.

Smeets, V. and F. Warzynski (2013). Estimating productivity with multi-product firms, pricing heterogeneity and the role of international trade. *Journal of International Economics 90*(2), 237–244.

Sorkin, I. (2017). Ranking firms using revealed preference. *The Quarterly Journal of Economics 133*(3), 1331–1393.

Van Horne, J. C. (1977). *Financial management and policy, 4th Ed.* Prentice Hall.

WEF (2018). The future of jobs. World Economic Forum Report.

Xu, S.-J. (2018). Skilled labor supply and corporate investment: Evidence from the H-1B visa program. Working Paper.

Zeira, J. (1998). Workers, machines, and economic growth. *The Quarterly Journal of Economics 113*(4), 1091–1117.

Zhang, M. B. (2019). Labor-technology substitution: Implications for asset pricing. *The Journal of Finance 74*(4), 1793–1839. Forthcoming in The Journal of Finance.