TECHNISCHE UNIVERSITÄT MÜNCHEN

TUM School of Management –
Lehrstuhl für Finanzmanagement und Kapitalmärkte

Three Essays on Finance and Labor

Daniel Bias

Vollständiger Abdruck der von der TUM School of Management der Technischen Universität München zur Erlangung des akademischen Grades eines

Doktors der Wirtschaftswissenschaften
(Dr. rer. pol.)

genehmigten Dissertation.

Vorsitzender: Univ.-Prof. Dr. Sebastian Schwenen

Prüfer der Dissertation:
1. Univ.-Prof. Dr. Christoph Kaserer
2. Univ.-Prof. Dr. Reiner Braun

Die Dissertation wurde am 18.02.2020 bei der Technischen Universität München eingereicht und durch die TUM School of Management am 15.04.2020 angenommen.

Thesis advisor: Christoph Kaserer Daniel Bias

Three Essays on Finance and Labor

Abstract

This dissertation examines three research questions on finance and labor. First, I[1] study the role of customer concentration in the upstream propagation of idiosyncratic firm-level shocks utilizing major labor strikes. I find that strike-hit customers impose substantial output loss on their suppliers. The negative effect on suppliers increases with suppliers' dependence on disrupted customers. These results show that customer concentration increases the vulnerability of production networks to idiosyncratic firm-level shocks. Second, I investigate how outside directorships of CEOs influence their managerial decision-making. For this purpose, I analyze how CEOs react after they observe, as directors of another firm, a labor strike that is plausibly exogenous to their firm. I find that CEOs increase cash holdings shortly afterwards. In the long run, CEOs agree to higher wages during contract negotiations and manage to reduce strike risk. These results suggest that outside directorships can facilitate both behavioral biases and observational learning. Finally, I examine the relationship between firm size and wage inequality using a large linked employee-establishment-firm dataset covering German workers. I find that wage inequality between and within firms increases monotonously with firm size. Decomposing wage inequality into workforce composition and non-composition effects reveals that workforce composition is responsible for most of the size-inequality relationship within firms and between firms.

[1]In this dissertation, I use the term "I" in the introduction and conclusion. It does not necessarily refer to me directly since the second and third essay are based on joint work with my co-authors.

Drei Aufsätze zu Finance und Labor

Abstract

Diese Dissertation untersucht drei Forschungsfragen zu Finance und Labor. Zuerst erforsche ich die Rolle von Kundenkonzentration für die Schock-Übertragung von Kunden- zu Zuliefererfirmen. Hierbei nutze ich Streiks als idiosynkratische Schocks auf die Produktion aus. Ich zeige auf, dass Streiks bei Kunden zu einer stark negativen Umsatzentwicklung der Zulieferer führen. Die negativen Auswirkungen sind umso stärker, je höher das Abhängigkeitsverhältnis der Zulieferer zu den Kunden ist. Diese Ergebnisse stellen heraus, dass Kundenkonzentration die Anfälligkeit von Produktionsnetzwerken für idiosynkratische Firmenereignisse erhöht. Als zweites untersuche ich den Einfluss externer Aufsichtsratsmandate auf das Entscheidungsverhalten von CEOs. Zu diesem Zweck analysiere ich die Reaktion von CEOs auf das Beobachten eines Streiks als Aufsichtsrat einer anderen Firma. Ich komme zu dem Ergebnis, dass CEOs kurzfristig die liquiden Mittel ihrer Firma erhöhen. Langfristig stimmen sie höheren Löhnen in Gewerkschaftsverhandlungen zu und reduzieren das Risiko von Streiks. Diese Ergebnisse zeigen, dass externe Aufsichtsratsmandate sowohl zu Überreaktionen als auch zu Modell-Lernen bei CEOs führen können. Abschließend untersuche ich den Zusammenhang von Firmengröße und Lohnungleichheit auf Basis eines großen Linked-Employee-Establishment-Firm Datensatzes. Ich komme zu dem Ergebnis, dass Lohnungleichheit innerhalb und zwischen Firmen monoton mit der Firmengröße zunimmt. Eine Zerlegung der Lohnungleichheit zeigt, dass sich der Zusammenhang von Firmengröße und Lohnungleichheit (innerhalb und zwischen Firmen) überwiegend durch Unterschiede in der Zusammensetzung der Arbeitskräfte erklären lässt.

Overview

Contents

List of Tables

List of Figures

I dedicate this work to my mother, my father, and Anastasia. Thank you all for your endless love and support.

Acknowledgments

First of all, I thank Christoph Kaserer for supervising my dissertation, engaging with me in insightful discussions, and providing valuable feedback. I am also very thankful to him for giving me many opportunities to grow. He supported me very actively in presenting my research at conferences, in visiting other universities abroad, and in developing my scientific network.

I am very grateful to Chen Lin for hosting me several times at the University of Hong Kong, and Vidhan Goyal for hosting me at the Hong Kong University of Science and Technology. During my research visits at these universities, I received very helpful comments on my work.

I thank, in particular, my co-authors Chen Lin, Benjamin Lochner, and Thomas Schmid on projects that are presented in this dissertation. I also thank my co-authors on other projects: Claudia Guagliano, Martin Haferkorn, Michael Haimann, and Lisa Knauer. Further, I am especially grateful to Reiner Braun, Dalia Marin, Sebastian Mueller, Daniel Urban, and Sebastian Schwenen for their comments and advice. I also thank all my Ph.D. fellows at TUM and the participants of conferences and seminars that I attended.

Finally, I am grateful for the (financial) support of the Research Excellence Program by the TUM School of Management, the TUM Graduate School, and the German Academic Exchange Service.

A few years ago, the typical financial economist would have recoiled at the idea that finance may have something to do with labor relations or trade unions, and the typical labor economist would have dismissed as bizarre that shareholder protection or corporate governance might be relevant for labor economics.

Pagano and Volpin (2008)

0 Introduction

Firms require both financing and labor to produce goods or services. Nevertheless, historically financial economics and labor economics developed as separate fields with only a few in-between collaborations. While, typically, financial economists focused on the financial aspects of firms and less on their workforce (Matsa, 2018), labor economists put their focus on the "people" and less on the firms (Card, 2011). Recently, financial and labor economists started actively working at the intersection of their disciplines to provide novel insights on aspects that fall between the two fields (Pagano and Volpin, 2008).

This dissertation consists of three essays that examine research questions on finance and labor. In the first essay, I study the role of customer concentration in the upstream propagation of idiosyncratic firm-level shocks utilizing major labor strikes. In the second essay, I investigate how outside directorships of CEOs influence their managerial decision-making. For this purpose, I analyze how CEOs react after they observe, as directors of another firm, a labor strike that is plausibly exogenous to their firm. In the third essay, I examine the relationship between firm size and wage inequality using a large linked employee-establishment-firm dataset.

0.1 Research questions and designs

In each essay, I utilize a specific empirical strategy and dataset to examine the research question. In the following subsections, I outline the three research questions and research designs.

0.1.1 Customer concentration and upstream propagation of shocks: evidence from labor strikes

Recent literature explores whether firm-level shocks can be the origin of aggregate fluctuations. Gabaix (2011) points out that idiosyncratic shocks to very large firms do not average out and can generate aggregate volatility. Herskovic et al. (2018) bring this idea back to the firm level and develop a customer-supplier network model that links firm size distribution and firm-level volatility. In their model, the lower diversification of the customer base explains the higher volatility of smaller firms. I empirically test how customer concentration accelerates the propagation of firm-level idiosyncratic shocks from (large) customers to suppliers.

As disruptions of customers' production, I utilize a sample of 223 major labor strikes involving at least 1,000 workers between 1984 and 2013. These labor strikes are highly idiosyncratic events that occur at 110 large publicly listed U.S. firms that source intermediates from a large number of suppliers. The obligation of listed firms to report the names and associated sales of major customers allows me to identify suppliers and quantify their dependence on the strike-hit firms. The empirical challenge in this setting is that links between suppliers and customers with strike risk are not formed randomly. Hence, firms supplying to customers with strike risk might be different from the population of suppliers. To address this concern, I construct a sample of suppliers whose customers are exposed to strike risk. I find that about two-thirds of suppliers report a customer with strike risk.

I start by showing that major labor strikes at customers impose a substantial output loss on their suppliers. When a labor strike hits a customer, suppliers' sales growth experiences an average drop by 1.9 to 2.8 percentage points in the quarter of the strike and the following quarter. Hence, the disruptions of customers quickly transmit to their suppliers. Furthermore, suppliers' output loss increases with the severity of the labor strike, particularly when the strike lasts long.

To test the theoretical prediction that the output loss of suppliers should increase with their direct dependence on the disrupted customer, I quantify suppliers' dependence on the disrupted firm as the fraction of sales from the customer in the year before the strike. In line with the theoretical predictions, I find that suppliers that are less dependent on the strike-hit customers (10th percentile) experience a decrease by 0.7 percentage points in sales growth, while suppliers that are highly dependent (90th percentile) experience a drop by 4.2 percentage points. Next, I test the theoretical prediction that additional indirect links between suppliers and disrupted customers amplify suppliers' output loss. These additional indirect links exist if suppliers sell products to other companies whose business also depends to a large extent on the strike-hit firms. I find that suppliers' additional indirect links to strike-hit firms increase the negative effect on sales growth by a factor of three to four.

Finally, I explore the role of customers' firm size in the upstream propagation of shocks. Suppliers' direct dependence, suppliers' indirect dependence, and the number of suppliers all monotonically increase with customers' firm size. Comparing large customer firms (decile 8 and 9) to the largest customer firms (decile 10) shows a substantial increase, for the largest customer firms, in suppliers' dependence and the number of suppliers.

0.1.2 Outside directorships and CEO decision-making: evidence from labor strikes

The appointment of outside CEOs as directors is a frequent phenomenon. Fahlenbrach, Low and Stulz (2010) identify nearly 2,000 CEO outside directors in a sample of about 5,000 firms, which corresponds to five percent of all directors. In large firms, CEOs are even more common on boards of directors. Despite their relevance, little is known about how outside positions affect managerial decision-making.

I analyze how CEOs react when they observe, as directors, a labor strike at another firm. 790 strikes occurred at 303 different S&P 1500 firms between 1984 and 2011, enabling me to identify 215 events of CEOs who observed a strike as directors. A strike observation can make strike risk more salient for the CEO and lead to more precautionary behavior in the short term. The psychology literature has shown that the attention placed on an event or risk factor plays a key role in its perceived probability (e.g., Bordalo, Gennaioli and Shleifer, 2012). On the other hand, observing a labor strike as a director may enable the CEO to gain insider knowledge and learn about labor negotiations (e.g., Bikhchandani, Hirshleifer and Welch, 1998). In the long run, this may affect their own bargaining with labor.

There are three reasons why labor strikes are useful in analyzing how outside directorships influence CEOs' decision-making. First, strikes are very costly events for firms that require a lot of attention from management and the board. Second, strikes are idiosyncratic events that are plausibly exogenous for the firms of the CEOs. Third, in contrast to many other types of shocks, such as the sudden death of a CEO, strikes are not exogenous for strike-hit firms. Rather, they are related to prior management decisions. Because CEOs can influence labor relations and signaling at their own firms, observational learning is a plausible mechanism for strikes to affect CEOs' decision-making.

An empirical challenge in this study is that outside directorships of CEOs are endogenously formed (Fahlenbrach, Kim and Low, 2018). CEOs who serve as directors at strike-hit firms may be different from the general population of CEOs. Although my focus on time-series variation within firms already reduces this potential concern to some extent, I also use two other strategies to address it. First, I analyze not only the strike observation itself but also the severity of the strike. The severity of a strike depends on many factors and is difficult to predict. Second, my control group consists of firms with a CEO who also sits on the board of a strike-hit firm but did not during the strike. These firms are more comparable to my event firms than the average firm.

I start by analyzing whether CEOs who observe a strike (temporarily) overreact to the more salient labor risk. I focus on the precautionary behavior of CEOs and measure how they adjust quarterly cash holdings. I find that CEOs who observe a labor strike increase cash holdings by about 0.7 percentage points, which represents a 10% increase in relative terms. This effect is more pronounced for more severe strikes. When I analyze the time patterns of the cash changes, I find that the increase starts in the quarter of the experience, reaches a maximum four quarters thereafter, and then reverts. This pattern is consistent with an overreaction to more salient labor risk because the salience of the strike observation diminishes over time.

Next, I focus on possible long-term effects due to observational learning and analyze two aspects of labor relations: wage setting in labor contract negotiations and strikes as a consequence of failed labor negotiations. I find that CEOs tend to agree to higher wages in the years after their strike observation compared to other firms with contract settlements. At the same time, they manage to decrease the strike risk for their firm, especially after observing a severe strike. These findings indicate that strike-observing CEOs learn from insider information about labor negotiations.

0.1.3 Firm size, workforce composition, and wage inequality

Wage inequality increased sharply during the last decades (e.g., Katz and Autor, 1999). Only recently has the literature begun focusing on the role of firms (e.g., Mueller, Ouimet and Simintzi, 2017a; Bloom et al., 2018; Song et al., 2019). I analyze the role of firm size in wage inequality using a linked employee-establishment-firm dataset from Germany that covers approximately 51,000 individual firms, 115,000 establishments, and 11.6 million workers. The uniqueness of this dataset comes from the fact that it links administrative data about individual workers with firm-level information, such as accounting figures.

Wage inequality can arise either because wages within larger firms are more dispersed ("within-firm inequality") or because larger firms pay their workers higher average wages ("between-firm inequality"). Comparing the largest firms (decile 10) to the smallest firms (decile 1), I find that the largest firms have a 30% higher within-firm wage variance and pay 70% higher wages. This positive effect of firm size is not only present when comparing the top and bottom of the size distribution but also increases monotonously with firm size.

To better understand why wage inequality is more pronounced in larger firms, I decompose between- and within-firm wage inequality into workforce composition and non-composition effects. For the decomposition, I use parameter estimates of the model introduced by Abowd, Kramarz and Margolis (1999) following the implementation of Card, Heining and Kline (2013). Generally speaking, this approach exploits movers between establishments to disentangle the overall wage into different components (e.g., person and establishment fixed effects). I find that composition effects play an important role in higher wage inequality in larger firms; the fact that larger firms employ, on average, a more heterogeneous and higher-quality workforce explains approximately 65% of the within-firm inequality and 60% of the between-firm inequality.

Next, I test which theoretical explanations can explain the higher dispersion of wages in large firms after controlling for the workforce composition. I find that the higher dispersion of wages in larger firms after workforce composition effects disappears when I control for monitoring complexity proxies (i.e., managers per employee and geographical dispersion of establishments within firms). Therefore, more dispersed wages in larger firms seem to be strongly related to their higher monitoring complexity.

Lastly, I examine which factors can explain higher average wages in larger firms after controlling for the workforce composition (the so-called "large-firm wage premium", LFWP). I do not find any empirical support for the hypotheses that differences in monitoring complexity, profitability, or unionization levels are the main drivers. However, when I compare small and large firms located in the same region, which rely on the same labor market, the LFWP is reduced by one-third or one-half, depending on the specification. Therefore, it seems that factors related to local labor markets play some role in the existence of the LFWP.

0.2 Contributions

The three essays of this dissertation contribute to multiple strands of the literature. I briefly summarize the main contribution of each essay.

In the *first essay*, I provide empirical evidence that customer concentration accelerates the propagation of idiosyncratic shocks from customers to suppliers. When a customer is disrupted, suppliers' output loss increases with their direct dependence on the disrupted firm. Suppliers' output loss is largely amplified by additional indirect links between suppliers and the disrupted customer. It is the very large firms that have a large number of highly dependent suppliers. Overall, my results provide evidence that production networks are vulnerable to idiosyncratic shocks to very large firms.

I contribute to the very recent literature on the propagation of idiosyncratic shocks in firm-level input-output networks (e.g., Barrot and Sauvagnat, 2016) by studying in-depth the mechanism underlying the upstream propagation of shocks: customer concentration. Furthermore, I contribute to the literature on the role of very large ("granular") firms in firm-level and aggregate volatility. Particularly, my findings provide empirical evidence for the theoretical predictions of the model by Herskovic et al. (2018) that explains the higher volatility of smaller firms by a lower diversification of their customer base.

In the *second essay*, I analyze how CEOs react when they observe, as directors, a labor strike at another firm. Exploiting this setting, I show that outside directorships influence CEOs' decision-making. In the short run, their attention is directed towards risk factors that they observe at the director firm. This leads to a temporary overestimation of that risk factor and precautionary behavior. Since this precautionary behavior seems to be an irrational overreaction that is most likely not in the best interest of the firm, this mechanism could be interpreted as a "dark side" of CEO outside directorships. However, there is also a "bright side", since my findings suggest that CEOs learn from the insider knowledge they receive at the director firm. In the long run, this observational learning affects their own behaviors as CEOs.

Most importantly, I add to the literature on the appointment of CEOs as directors (e.g., Fahlenbrach, Low and Stulz, 2010). As a potential benefit of CEOs' outside positions, my study provides direct empirical evidence for the existence of learning effects. As potential reasons against CEO outside directorships, my results indicate that not only do time and effort spent by CEOs matter, but their decision-making can also be adversely affected due to behavioral biases.

In the *third essay*, I document that the largest firms (decile 10) have a 30% higher within-firm variance and pay 70% higher wages than the smallest firms (decile 1). The fact that large firms employ, on average, a more heterogeneous and higher-quality workforce explains about 65% of the differences in within-firm inequality. The corresponding figure for between-firm wage inequality is 60%. I find that the higher within-firm inequality in larger firms after workforce composition effects disappears when I control for monitoring costs. The higher average wages in larger firms after composition effects (large-firm wage premium", LFWP) cannot be explained by differences in monitoring complexity, profitability, or unionization. However, I find that the LFWP decreases by approximately one-third if I compare small and large firms operating in the same local labor market.

My findings complement the literature on higher dispersion of wages in larger firms (e.g., Mueller, Ouimet and Simintzi, 2017b) by showing that differences in the workforce composition can explain about 65% of differences in within-firm inequality between large and small firms, that within-firm inequality increases monotonously with firm size, and that differences in monitoring complexity can explain differences in within-firm inequality between small and large firms after controlling for workforce composition. I add to the literature on higher average wages in larger firms (e.g, Bloom et al., 2018) by testing theoretical explanations that require detailed firm-level data. Surprisingly, I do not find any evidence that rent-sharing, monitoring complexity, or unionization have a substantial impact on the LFWP.

0.3 Outline

The remainder of this dissertation proceeds as follows. In Chapter 1, I study the role of customer concentration in the upstream propagation of shocks. In Chapter 2, I investigate how outside directorships influence CEO decision-making. In Chapter 3, I analyze the relationship between firm size and wage inequality. Finally, in Chapter 4, I briefly summarize the main results and highlight their contributions and implications.

1

Customer Concentration and the Upstream Propagation of Idiosyncratic Shocks: Evidence from Labor Strikes

Authors:	Daniel Bias
First author:	Daniel Bias
Keywords:	Customer concentration, shock propagation, customer-supplier, network effects, labor strikes
JEL codes:	L14, L22, L23, L25, E32, G32, J52
Presentations:	-
Current status:	Working paper

Abstract:

This paper studies the role of customer concentration in the upstream propagation of idiosyncratic firm-level shocks. I utilize major labor strikes as idiosyncratic disruptions of large firms with multiple suppliers. I find that strike-hit customers impose a substantial output loss on their suppliers. The negative effect increases with suppliers' *direct* dependence on disrupted customers. Moreover, suppliers' output loss is amplified by additional *indirect* links that exist if suppliers sell products to other companies whose business also depends on the large disrupted customer. Overall, these results show that customer concentration increases the vulnerability of production networks to idiosyncratic firm-level shocks.

1.1 Introduction

Recent literature explores whether firm-level shocks can be the origin of aggregate fluctuations. Gabaix (2011) points out that idiosyncratic shocks to very large ("granular") firms do not average out and can generate aggregate volatility.[1] Herskovic et al. (2018) bring this idea back to the firm level and develop a customer-supplier network model that links firm size distribution and firm-level volatility. In their model, the lower diversification of the customer base explains the higher volatility of smaller firms. In this paper, I empirically test how customer concentration accelerates the propagation of firm-level idiosyncratic shocks from (large) customers to suppliers.

As disruptions of customers' production, I utilize a sample of 223 major labor strikes involving at least 1,000 workers between 1984 and 2013. These labor strikes are highly idiosyncratic events that occur at 110 large publicly listed U.S. firms that source intermediates from a large number of suppliers.[2] The obligation of listed firms to report the names and associated sales of major customers allows me to identify suppliers and quantify their dependence on the strike-hit firms.

An illustrative example for the upstream propagation of strikes and the role of customer concentration is the 40-day walkout by 48,000 workers against General Motors (GM) starting on September 19, 2019. The walkout forced an immediate shutdown of GM's American factories. After a few days, GM's factories in Mexico and Canada also had to shut down. At about the same time, GM's suppliers began to halt their productions and to consider temporary layoffs. Theoretically, the impact of the strike should vary from supplier to supplier,

[1] See also Carvalho et al. (2012); di Giovanni and Levchenko (2012); Carvalho and Gabaix (2013); di Giovanni, Levchenko and Mejean (2014, 2018); Grassi (2017); Baqaee and Farhi (2019); Carvalho and Grassi (2019).

[2] Gabaix (2011) considers major labor strikes as one example for an idiosyncratic productivity shock that can explain aggregate fluctuations.

depending on their reliance on GM. In the worst case, a supplier sells its products only to GM. Then, its entire sales directly depend on GM ("direct customer concentration"). For example, American Axle and Manufacturing took one of the strongest hits as 41% of the company's sales were concentrated on General Motors in 2018. Suppliers may suffer from the strike beyond their direct exposure to GM if they sell products to other companies whose business also depends to a large extent on the central customer GM ("indirect customer concentration"). For example, Methode Electronics manufactures subsystems for the automotive industry and sells them to both GM and its tiered suppliers.[3]

Consistent with the anecdotal evidence, I start by showing that major labor strikes at customers impose a substantial output loss on their suppliers. When a labor strike hits a customer, suppliers' sales growth experiences an average drop by 1.9 to 2.8 percentage points in the quarter of the strike and the following quarter. Hence, the disruptions of customers quickly transmit to their suppliers. Neither differences in firm size, age, and profitability nor local economic development nor industry-wide shocks are driving this result. In a robustness test, I show that suppliers' output loss increases with the severity of the labor strike, particularly when the strike lasts long. Another robustness test shows that firm-specific seasonal trends do not drive the results.

The empirical challenge in this setting is that links between suppliers and customers with strike risk are not formed randomly. Hence, firms supplying to customers with strike risk might be different from the population of suppliers. To address this concern, I construct a sample of suppliers whose customers are exposed to strike risk. I exploit the fact that strikes typically occur at the end of a labor contract between the firm and its union. Linking the information on expiring labor contracts to the supplier-customer network allows me to identify

[3]Another example is Universal Logistics, which provides customized transportation and logistic support. Its 2018 annual report states: "If the UAW and our automotive customers and their suppliers are unable to negotiate new contracts and our customers' plants experience slowdowns or closures as a result, our revenue and profitability could be negatively impacted."

suppliers whose customers are exposed to the risk of a labor strike. I find that about two-thirds of suppliers report a customer with an expiring labor contract.

I quantify suppliers' dependence on the disrupted firm as the fraction of sales from the customer in the year before the strike. This allows me to test the theoretical prediction that the output loss of suppliers should increase with their direct dependence on the disrupted customer. Using the lagged fraction of sales as a continuous treatment measure further mitigates endogeneity concerns. Suppliers' sales with customers is hard to adjust in the short-run (i.e., due to the requirement of relation-specific investments). Furthermore, it is rather unlikely that the occurrence of the strike is correlated with each suppliers' dependence on the strike-hit customer. In line with the theoretical predictions, I find that suppliers that are less dependent on the strike-hit customers (10th percentile) experience a decrease by 0.7 percentage points in sales growth, while suppliers that are highly dependent (90th percentile) experience a drop by 4.2 percentage points.

Next, I test whether additional indirect links between suppliers and disrupted customers amplify suppliers' output loss. These additional indirect links exist if suppliers sell products to other companies whose business also depends to a large extent on the strike-hit firms, as also illustrated in the example of GM. I refer to such a customer-supplier relationship as a supplier's central customer. To quantify the effect of additional indirect links between suppliers and strike-hit firms, I estimate a difference-in-differences setting. I find that suppliers' additional indirect links to strike-hit firms increase the negative effect on suppliers' sales growth by three to four times.

Finally, I explore the role of customers' firm size in the upstream propagation of shocks. Suppliers' direct dependence, suppliers' indirect dependence, and the number of suppliers all monotonically increase with customers' firm size. Comparing large customer firms (decile 8 and 9) to the largest customer firms (decile 10) shows a substantial increase of all three measures for the largest firms. On average, the largest customers account for 18% of suppliers' sales, have a likelihood of 28% to be a central customer for at least one supplier, and source intermediates from 16 publicly listed suppliers that classify them as a major customer.

To summarize, my empirical setting allows me to identify that customer concentration is of importance for the upstream propagation of idiosyncratic firm-level shocks. When a customer is disrupted, suppliers' output loss increases with their direct dependence on the disrupted firm. Suppliers' output loss is largely amplified by additional indirect links between suppliers and the disrupted customer. It is the very large firms that have a large number of highly dependent suppliers. Overall, my results provide evidence that production networks are vulnerable to idiosyncratic shocks to very large firms.

This paper contributes to several strands of the literature. Empirical evidence on the propagation of idiosyncratic shocks in firm-level input-output networks is still limited. Using information on publicly listed U.S. firms, Barrot and Sauvagnat (2016) show that input disruptions caused by natural disasters impose substantial output losses on direct customers when the disrupted suppliers produce specific inputs. Following the 2011 earthquake in Japan, Boehm, Flaaen and Pandalai-Nayar (2019) document the cross-border transmission of input disruptions for U.S. affiliates of Japanese multinationals. Exploiting the Japanese earthquake and an input-output network covering private and public firms, Carvalho et al. (2016) show that the disruptions of disaster-hit firms affect direct and indirect customers as well as suppliers. Building on their findings, I utilize information on public U.S. firms and labor strikes at large firms with multiple suppliers to study in-depth the mechanism underlying the upstream propagation of shocks: (i) direct customer concentration accelerates the upstream propagation of shocks, (ii) indirect customer concentration amplifies the negative effect on suppliers, (iii) production networks of firms with concentrated customer bases are vulnerable to idiosyncratic firm-level shocks.[4]

These findings also contribute to the literature on the role of granular firms in firm-level and aggregate volatility. Linking firm size distribution and firm-level volatility, Herskovic et al. (2018) develop a supplier-customer model that explains the higher volatility of smaller firms by a lower diversification of their customer base. My findings provide empirical evidence for the theoretical predictions of their model. In line with Gabaix (2011)'s granular view, I show that it is the very large customer firms that have a large number of suppliers with a substantial direct and indirect customer concentration. These findings are also related to papers studying the macroeconomic impacts of large and well-connected firms' failure in production networks (Baqaee, 2018; König, 2018; Taschereau-Dumouchel, 2019).[5]

My findings add to the literature in financial economics that studies how firms are affected by their customers. Recent studies find an effect of firms' (direct) customer concentration on their financing costs (Dhaliwal et al., 2016; Houston, Lin and Zhu, 2016; Campello and Gao, 2017). My results suggest that suppliers' additional indirect links to customers should also affect their financing costs. Moreover, my results are related to studies on how customers affect suppliers' corporate policies (Titman, 1984; Titman and Wessels, 1988; Shahrur, 2005; Chu, 2012). These papers find that suppliers hold less leverage if a customer must make relation-specific investments, if there is more competition among suppliers, and if suppliers' products are easy to substitute. My results support the alternative explanation that suppliers hold less leverage to be financially

[4] My results are consistent with Inoue and Todo (2019), who simulate the production losses caused by a Japanese mega earthquake in a supplier–customer network. They also find a substantial amplifying effect of indirect links referred to as "complex cycles in networks".

[5] In this paper, I assume the firm-level input-output network as given. Atalay et al. (2011), Chaney (2014), Carvalho and Voigtländer (2015), Boehm and Oberfield (2018), Lim (2018), Oberfield (2018), Tintelnot et al. (2018), and Acemoglu and Azar (2020) explicitly model the formation of networks.

flexible when shocks disrupt their customers.

Lastly, my work adds to the literature on the effect of labor strikes. Krueger and Mas (2004), Mas (2008), and Gruber and Kleiner (2012) show that strikes affect the quality of products and services. Becker and Olson (1986) and Dinardo and Hallock (2002) document adverse stock market reactions for strike-hit firms, and De Fusco and Fuess (1991) positive stock market reactions for competitors of strike-hit airlines. Evidence on the effect of strikes on supply chains is scarce. McHugh (1991) finds short-run productivity declines in supplier and customer industries. Persons (1995) documents the negative effect of automotive strikes on stock prices of steel suppliers. I show that strikes at customers impose substantial output losses on their suppliers.

1.2 Theory and empirical strategy

1.2.1 Simplified network model

To illustrate the role of customer concentration in the upstream propagation of shocks, I use a simplified example of the network model derived by Herskovic et al. (2018).[6] In their reduced form network model, the growth rate of a supplier s depends on its own idiosyncratic shock component and on the weighted average of growth rates of the firms c it sells its products to:

$$g_s = \mu_s + \gamma \sum_{c=1}^{N} w_{s,c} g_c + \varepsilon_s, \tag{1.1}$$

where γ denotes the rate of decay as a shock propagates through the network and $w_{s,c}$ the network weight. The network weight measures the supplier's dependence on customer c. It determines how strongly the growth rate of supplier

[6] Herskovic et al. (2018) derive a network model in which shocks are transmitted from customers to suppliers. They use this model to show the relationship between firm size distribution, firm volatility, and customer concentration.

s is influenced by customer c.

Assume a simplified example of three firms. A supplier firm s sells its goods to two customer firms $c1$ and $c2$. There is no economic link between firms $c1$ and $c2$. The growth rate of the three firms can be written as follows,[7]

$$\begin{aligned} g_{c1} &= \mu_{c1} \\ g_{c2} &= \mu_{c2} \\ g_s &= \mu_s + w_{s,c1}\mu_{c1} + w_{s,c2}\mu_{c2} \end{aligned} \tag{1.2}$$

Assume an idiosyncratic shock disrupts firm $c1$ (e.g., a major strike occurs). The production disruption of customer $c1$ affects supplier s by $w_{s,c1}\mu_{c1}$. Hence, the upstream propagation of a shock to supplier s depends on the firm's customer concentration. The more the supplier's customer base is concentrated on the disrupted customer, the stronger should be the negative effect on the supplier.

Now, assume that there is also an economic link between the two customers $c1$ and $c2$. $c2$ is selling parts of its output as intermediate to $c1$. Due to the additional indirect link of supplier s to $c1$ through $c2$, I refer to $c1$ as the supplier's central customer. In this case, the growth rate of the firms can be written as follows,

$$\begin{aligned} g_{c1} &= \mu_{c1} \\ g_{c2} &= \mu_{c2} + w_{c1,c2}\mu_{c2} \\ g_s &= \mu_s + w_{s,c1}\mu_{c1} + w_{s,c2}(\mu_{c2} + w_{c1,c2}\mu_{c2}). \end{aligned} \tag{1.3}$$

Then, the total effect of the production disruption at customer $c1$ on supplier s is $w_{s,c1}\mu_{c1} + w_{s,c2}w_{c1,c2}\mu_{c2}$. It consists of the direct effect $w_{s,c1}\mu_{c1}$ and the indirect effect $w_{s,c2}w_{c1,c2}\mu_{c2}$ through the business relation of the customers $c1$ and $c2$. Thus, I expect that additional indirect links between a supplier and the disrupted

[7]For the simplified illustration, I assume the rate of decay γ to be equal to one.

central customer will amplify the supplier's output loss. I refer to this as the effect from "indirect customer concentration".

1.2.2 Empirical strategy

As idiosyncratic shocks to customers' production, I use major labor strikes involving at least 1,000 workers. Such a major labor strike disrupts the production of the affected establishments and can halt the entire production of the strike-hit firm, for instance, if a major part of the workforce or a small group of essential workers walk out. The main focus of my analysis is not the impact of the production disruption on the customer firm itself but the propagation to its suppliers.

To identify firms supplying a major part of their output to the strike-hit firm (i.e., having a large network weight with the strike-hit firm), I rely on the obligation of publicly listed firms in the United States to disclose major customers accounting for ten percent or more of their sales. Some firms even report major customers accounting for a lower fraction of their sales (c.f., Atalay et al., 2011). I consider all firms reported as a major customer. Since I am interested in the impact of a production disruption at a long-term customer on the supplying firms, I require that a supplier reports the strike-hit firm as a major customer in the year before the strike and the year of the strike.

I analyze the effect of labor strikes at customers on the quarterly sales growth of its suppliers. I estimate the following regression model at the firm-quarter level on the supplier sample,[8]

$$\Delta ln(Sales)_{i,t,t-4} = \beta StrikeHitsCustomer_{i,t,t-1} + \alpha_i + \pi_t + \eta_{i,t-4} + \varepsilon_{i,t}, \tag{1.4}$$

[8] My empirical setting follows Barrot and Sauvagnat (2016) who use a similar approach to estimate the downstream propagation of suppliers affected by a natural disaster to their customers. Technically, it is a difference-in-differences framework at the firm level that assumes sales growth to be flat in the absence of treatment.

where $ln(Sales)_{i,t,t-4}$ is the sales growth of firm i between the current quarter and the same quarter one year before. The main variable of interest is $StrikeHitsCustomer_{i,t,t-1}$. It is a dummy that equals one if a strike hits a customer of the firm in the current quarter or the previous quarter. α_i denotes firm fixed effects, π_t year-quarter fixed effects, and $\eta_{i,t-4}$ a fixed effect that controls for a firm's number of customers in the previous year. The fixed effects are constructed as terciles of a supplier's number of customers in the year before. $\varepsilon_{i,t}$ is the error term. In some specifications, I add firm-level controls lagged by one year, industry-year fixed effects, and state-year fixed effects. I construct the lagged controls as terciles of firms' total assets, return on assets, and age interacted with year-quarter fixed effects. Standard errors in all regression models are clustered at the firm-level.

The empirical challenge in this setting is that links between suppliers and customers with strike risk are not formed randomly. Hence, firms supplying to customers with strike risk might be different from the population of suppliers. To address this concern, I limit my sample to suppliers whose customers are exposed to strike risk. For this, I exploit that strikes typically occur at the end of a labor contract between a firm and its union.[9] If a labor union wishes to modify or terminate the labor contracts, it must notify the Federal Mediation and Conciliation Services. Linking this information to the supplier-customer network allows me to identify suppliers whose customers are exposed to the risk of a labor strike. I find that about two-thirds of suppliers report (at least once) a customer with an expiring labor contract.

The network model, as presented in Section 1.2.1, predicts that the negative effect on suppliers increases with their dependence on the disrupted cus-

[9]The National Labor Relations Board distinguishes two classes of lawful objects for strikes. These are "economic strikers" (ES) and "unfair labor practice strikers" (ULPS). ES strike for some economic concessions of the employer (e.g., higher wages or shorter hours). ULPS protest against unfair labor practices committed by the employer (e.g., interfering with an employee's right to organize). ULPS are more strongly protected by law. While ES occur during contract negations, ULPS do not have to occur during contract negotiations.

tomer. To test this theoretical prediction, I quantify suppliers' dependence on the disrupted firm as the fraction of sales from the customer in the year before the strike. Using the lagged fraction of sales as continuous treatment measure ($StrikeHitsCustomer_{i,t,t-1,\%sales}$) further mitigates endogeneity concerns since suppliers' sales from customers evolve slowly over time (e.g., due to the requirement of relation-specific investments).

Next, I test the theoretical prediction that additional indirect links between suppliers and the disrupted customers amplify suppliers' output loss. These additional indirect links exist if suppliers sell products to other companies whose business also depends to a large extent on the strike-hit firms. I refer to such a customer-supplier relationship as a supplier's central customer.[10] For the classification of central customers, I assume that a firm is permanently a customer of a supplier in all years after the supplier reports the firm the first time as one of its customers. Based on this network definition, I construct a dummy variable indicating that the strike-hit firm is a supplier's central customer ($StrikeHitsCentralCustomer_{i,t,t-1}$). To estimate the effect of additional indirect links between suppliers and strike-hit firms, I estimate a difference-in-differences model:

$$\begin{aligned} \Delta ln(Sales)_{i,t,t-4} = {} & \beta_1 StrikeHitsCustomer_{i,t,t-1} + \\ & \beta_2 StrikeHitsCentralCustomer_{i,t,t-1} + \\ & \alpha_i + \pi_t + \eta_{i,t-4} + \varepsilon_{i,t}. \end{aligned} \tag{1.5}$$

This specification compares the output loss of suppliers that have both a direct link and an indirect link to the strike-hit firm and the output loss of suppliers with only a direct link to the strike-hit firm.

[10]While I assume the network structure as given, Carvalho and Voigtländer (2015) formally model the evolution of the business network. They show theoretically and empirically that firms are more likely to develop new input-relations among their suppliers' network neighborhood. This behavior can foster the formation of indirect links between suppliers and customers whose role I study in the upstream propagation of shocks.

1.3 Data

1.3.1 Sample and firm financials

Financial data is retrieved from the Compustat North America Fundamentals Quarterly database. The sample consists of non-financial firms over the period from 1983 to 2014. All continuous variables are winsorized at the 1st and 99th percentiles. The calculation of sales growth is adjusted for inflation using the Consumer Price Index by the Bureau of Labor Statistics (BLS).

1.3.2 Supplier-customer links

The Statement of Financial Accounting Standard (SFAS) No. 14 requires publicly listed U.S. firms to report their reliance on major customers. A major customer accounts for ten percent or more of a firm's sales. Firms shall disclose the names of these customers and the total amount of sales from each such customer.

I obtain information on the names of firms' major customers and the associated sales from the Compustat Segement database for the period between 1983 and 2013. To link the names of the major customers to Compustat firms, I use the linking table provided on the website of Jean Noël Barrot. Barrot and Sauvagnat (2016) apply a phonetic string matching to assign the names of major customers to Compustat identifiers.

1.3.3 Labor strikes

Since there is no comprehensive database on labor strikes, I consolidate information on strikes from the BLS, the Federal Mediation and Conciliation Service (FMCS), the National Mediation Board (NMB), and labor contract data by Cramton and Tracy (1992).[11] I include strikes in my sample if the information on the

[11]The labor contract data by Cramton and Tracy is available via the Inter-university Consortium for Political and Social Research (ICPSR 1020).

name of the strike-hit employer, the number of striking workers, the affected state, the beginning date, and the end date is available. The BLS defines a labor strike as major if at least 1,000 workers are involved. I follow their definition and consider only strikes above this threshold. Next, I manually link the major labor strikes to Compustat firms using the name of the strike-hit employer. Finally, I aggregate all labor strikes that start at a firm in the same calendar week into a single strike event.[12]

In total, I identify 223 major labor strikes at 110 individual Compustat firms in 207 firm-quarters.[13] Figure 1.1 presents the distribution of labor strikes over time and industries. Subfigure (a) presents the number of strikes per year. The largest number of strikes (26) occurred in 1986. The number of strikes per year has decreased over time. The average number of strikes per year was 15 in the 1980s, 8 in the 1990s, and 5 in the 2000s. In recent (out-of-sample) years, the number of major strikes is increasing again.[14] Subfigure (b) presents the number of strikes per industry division. The most labor strikes occurred in the manufacturing division with 151 (68%), followed by transportation, communication and utilities with 39 (18%), and retail trade with 21 strikes (10%). Subfigure (c) presents the number of strikes per broad manufacturing industry. Inside manufacturing, firms producing transportation equipment account for the majority of strikes with 84 walkouts, followed by firms producing industrial machinery/electronic equipment with 18 strikes, and firms producing metal/metal products with 13 strikes.

Table 1.1 presents descriptive statistics on the severity of the labor strikes. On average, a strike lasts for 44 days. The median value is 22 days. The average

[12] I conduct this aggregation of strikes by firm and calendar week since some strikes are reported on the firm level (aggregated over all establishments) and others on the establishment level. Otherwise, strikes reported separately for multiple establishments may distort the summary statistics of strikes.

[13] The dataset also includes lockouts. However, strikes represent the vast majority of the work stoppages.

[14] The BLS reports that 2018 was the year with the largest number of major labor strikes since 2007 and the largest number of workers involved since 1986.

Figure 1.1: Major labor strikes
This figure illustrates the distribution of 223 major labor strikes involving at least 1,000 workers over time and industries. Subfigure (a) presents the number of strikes per year. Subfigure (b) presents the number of strikes per industry division (using two-digit SIC codes). Subfigure (c) presents the number of strikes per broad manufacturing industry. The strikes occur between 1984 and 2013.

(a) Major labor strikes per year

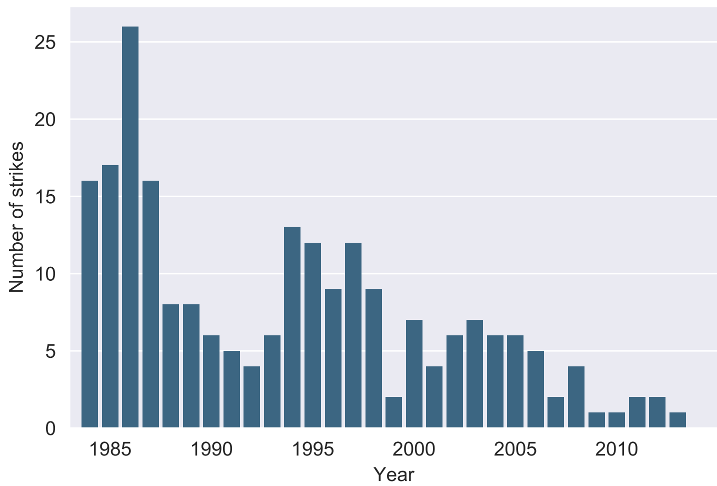

(b) Major labor strikes per industry division

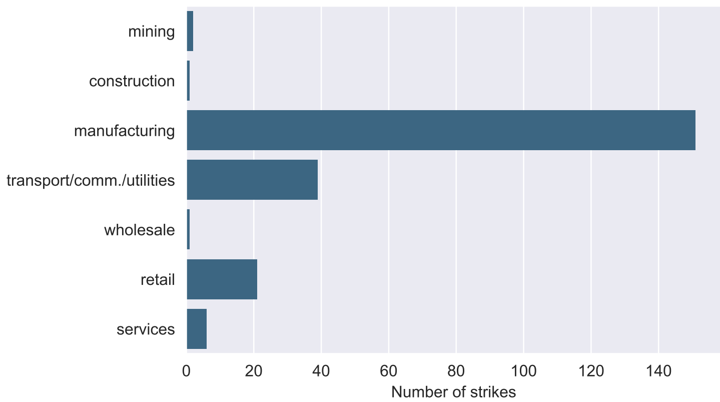

(c) Major labor strikes per broad manufacturing industry

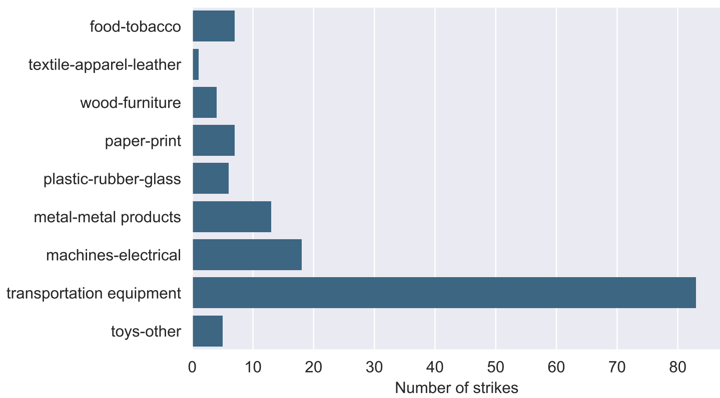

number of striking employees is 7,034. The average (median) fraction of striking employees relative to the firm's total workforce is 14% (4.9%). The ratio of idled employee-days to the firm's total number of employee-days per fiscal quarter has an average value of 36% and a median value of 24%.

Table 1.1: Descriptive statistics: major labor strikes

This table presents summary statistics of major labor strikes involving at least 1,000 workers. The labor strikes occur between 1984 and 2013. The 223 labor strikes occur at 110 firms and 207 firm-quarters. The 110 strike-hit firms consist of 92 firms from the customer sample and 40 firms from the supplier sample. 22 firms are both in the customer and the supplier sample as they are a major customer to a supplier firm and simulatenously depend on a major customer themselves. Reported are the number of observations (Obs), mean value (Mean), standard deviation (SD), 25% percentile (p25), median (p50), and 75% percentile (p75). A detailed description of all variables can be found in Appendix A.1.

	Obs	Mean	SD	p25	p50	p75
strike duration	223	43.572	70.516	8.000	22.000	49.000
striking employees	223	7034	11500	1500	3000	6700
striking emp/total emp	221	0.140	0.196	0.019	0.049	0.169
idled employee-days	223	271191	578808	24780	57500	168000
idled emp-days/total emp-days	223	0.360	0.334	0.089	0.244	0.544

1.3.4 Collective bargaining agreements

Information on expiring collective bargaining agreements with labor unions is obtained from the website of Thomas Holmes and the FMCS. Thomas Holmes filed a Freedom of Information Request with the FMCS for expiring labor contracts in the period between 1985 and 2003.[15] The FMCS reports information on expiring labor contracts between 2005 to 2013.

To link expiring labor contracts to Compustat firms, I apply a record linkage approach using various contract and employer characteristics of which the employer name, the industry and the state are the most important ones. Hence, I include only labor contracts that provide this information. Further, I consider only labor contracts covering at least one hundred workers. The starting point for the record linkage approach is then a dataset of 114,229 labor contracts.

[15] Holmes (2006) uses this data to identify unionized health care facilities.

The record linkage approach allows me to link 14,814 labor contracts to 1,544 Compustat firms.

1.3.5 Descriptive statistics

Tables 1.2 provides descriptive statistics on the samples. Panel A presents the supplier sample consisting of 132,587 firm-quarters and 4,118 individual firms. A supplier is included if it reports (at least once) a customer with an expiring labor contract.[16] These firms are then included in each year from three years before to three years after they report another firm as a customer in the Compustat Segment database. On average, a firm reports 1.20 customers in a given year, which account on average for 18.6% of its sales. The largest customer accounts on average for 14.9%. 5.3% of firms in a given year have a central customer that is a customer of the firm (direct link) and also a customer of the firm's other customers (indirect link). The main variable of interest is the sales growth relative to the same quarter in the previous year. The mean and median values are 7.5% and 5.1%. The probability that a strike hits a firm's customer is 1.3%, while the probability that a major labor strike occurs at a supplier is only 0.04%.

Panel B presents the customer sample consisting of 31,114 firm-years and 578 individual firms. A firm is included if it has at least one expiring labor contract.[17] These firms are included in the sample in each year from three years before to three years after they are reported as customer by a supplier. On average, a firm is reported by 3.59 suppliers as a customer. The average and median sales growth for the customers are 3.0% and 2.4%. The probability that one of the customer firms is hit by a major strike is 0.6%.

[16] The majority of suppliers (64%) report (at least once) a customer with an expiring labor contract. I exclude 2,312 suppliers (36%) as they are not exposed to the risk of a labor strike at a customer.

[17] About 22% of customer firms have at least one expiring labor contract. I exclude 2,101 firms (78%) that have no expiring labor contracts.

Table 1.2: Descriptive statistics: supplier and customer sample

This table presents summary statistics of the sample. Panel A presents the suppler sample, which consists of 132,587 firm-quarters and 4,118 individual firms. Panel B presents the customer sample, which consists of 31,114 firm-quarters and 578 indof variablesividual firms. Reported are the number of observations (Obs), mean value (Mean), standard deviation (SD), 25% percentile (p25), median (p50), and 75% percentile (p75). A detailed description of all variables can be found in Appendix A.1.

	Obs	Mean	SD	p25	p50	p75
Panel A: Supplier sample						
sales growth$_{t,t-4}$	132,587	0.075	0.511	-0.107	0.051	0.245
strike hits customer	132,587	0.013	0.113	0.000	0.000	0.000
strike hits firm	132,587	0.000	0.019	0.000	0.000	0.000
number of customers	132,587	1.195	1.270	0.000	1.000	2.000
%sales with customers	132,587	0.186	0.226	0.000	0.120	0.283
%sales with largest customer	132,587	0.149	0.181	0.000	0.110	0.210
central customer	132,587	0.053	0.224	0.000	0.000	0.000
real size	132,587	5.029	2.125	3.527	4.881	6.457
age	132,587	11.300	9.501	4.000	9.000	17.000
roa	132,587	0.035	0.277	0.002	0.103	0.170
Panel B: Customer sample						
sales growth$_{t,t-4}$	31,114	0.030	0.230	-0.047	0.024	0.098
strike hits firm	31,114	0.006	0.076	0.000	0.000	0.000
number of suppliers	31,114	3.586	8.794	0.000	1.000	3.000
real size	31,114	8.746	1.484	7.792	8.803	9.796
age	31,114	21.625	10.543	14.000	21.000	29.000
roa	31,114	0.140	0.068	0.099	0.136	0.179

1.4 Results

1.4.1 Effect on disrupted customers

I first analyze the effect of major labor strikes on the production of strike-hit firms. A labor strike involving at least 1,000 workers should harm the disrupted firms' output. Since firms' production numbers are not observable, I rely on quarterly sales as the output measure.

In the customer sample, I regress firms' sales growth relative to the same quarter in the previous year on dummies indicating whether a labor strike hits the firm in the current and each of the previous three quarters. Table 1.3 presents the results. In Column 1, the coefficient estimates on the dummies indicating that a strike occurs in the current quarter and the previous quarter are -4.1 and -3.1 percentage points, and statistically significant at the 5%-level. The coefficient estimates on the dummies indicating that a strike occurs two quarters before and three quarters before are -0.68 and -0.88 percentage points and are not statistically significant. In Column 2, I add lagged controls for firms' size, age, and profitability to the model. They are constructed as terciles of the lagged control variables and interacted with year-quarter dummies. The size of the coefficient estimates slightly decreases to -3.3 and -2.6 percentage points. In Columns 3 and 4, I introduce industry-year and state-year fixed effects. In the full model, the coefficient estimates are -4.0 and -3.0 percentage points. These results provide evidence that strike-hit firms perform worse than other firms in the same industry and the same state.

The purpose of the following section is to analyze the role of customer concentration in the upstream propagation of shocks. However, if customers that source from a larger number of suppliers or from more highly-dependent suppliers are, on average, hit by more severe labor strikes, I would mechanically

Table 1.3: Strikes and customers' sales growth

This table presents regression estimates of firms' sales growth relative to the same quarter in the previous year on dummies indicating whether a firm is hit by a labor strike in the current and each of the previous three quarters. All regressions include firm and year-quarter fixed effects. In Column 2, I control for firm-level characteristics (constructed as terciles of lagged size, age, and return on assets, respectively) interacted with year-quarter dummies. In Column 3, I include industry dummies using two-digit SIC codes interacted with year dummies. In Column 4, I include state dummies interacted with year dummies. Regressions contain all firm-quarters of the customer sample between 1983 and 2013. T-statistics presented in parentheses are based on robust standard errors clustered at the firm level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix A.1.

	(1)	(2)	(3)	(4)
strike hits firm$_t$	-0.041**	-0.033*	-0.031*	-0.040***
	(-2.46)	(-1.95)	(-1.91)	(-2.65)
strike hits firm$_{t-1}$	-0.031**	-0.026*	-0.020	-0.030**
	(-2.44)	(-1.91)	(-1.52)	(-2.28)
strike hits firm$_{t-2}$	-0.0068	-0.0075	-0.010	-0.018
	(-0.59)	(-0.61)	(-0.78)	(-1.20)
strike hits firm$_{t-3}$	-0.0088	-0.0057	-0.017	-0.021
	(-0.65)	(-0.41)	(-1.21)	(-1.36)
Firm FE	Yes	Yes	Yes	Yes
Year-quarter FE	Yes	Yes	Yes	Yes
Firm-level controls	No	Yes	Yes	Yes
Industry-year FE	No	No	Yes	Yes
State-year FE	No	No	No	Yes
Obs	31,114	31,108	31,107	31,107
Number - firms	578	578	578	578
R2	0.140	0.176	0.320	0.401

overestimate the role of customer concentration. To test this, I regress firms' sales growth on a dummy indicating whether a strike hits the firm in the current or the previous quarter interacted with the firm's number of suppliers, or respective with the average suppliers' dependence. Appendix A.2 presents the results. I do not find any indication that firms with more intensive supplier relationships are, on average, hit by more severe strikes.

1.4.2 Upstream propagation: effect on suppliers' sales

In this section, I analyze the effect of strike-hit customers on their suppliers' sales and the role of customer concentration in the upstream propagation of shocks.

1.4.2.1 Time dynamics of the upstream propagation

I start by exploring the time dynamics of the upstream propagation of production disruptions at customers. Therefore, I separately estimate for customers and suppliers a regression model with time dummies indicating the four quarters before and the four quarters following the major labor strike. Figure 1.2 illustrates the results. For the strike-hit customers, the coefficient estimates before the major labor strike are economically close to zero and statistically not significant. In line with the results in Section 1.4.1, the coefficient estimates for the dummies in the current and the previous quarter are negative and statistically significant. In the two quarters after that, the coefficient estimates are negative and not statistically significant. For the suppliers, the coefficient estimates indicating the four quarters before a strike are positive and not statistically significant. The coefficient estimates for the dummies indicating that a strike occurs at a customer in the current quarter, one quarter before and two quarters before, are negative and statistically significant. Hence, production disruptions at customers have an immediate effect on the depending suppliers, and this negative effect also prevails in the following two quarters.

Figure 1.2: Labor strikes and sales growth
This figure illustrates regression estimates of quarterly sales growth in the one year before and the one year following a labor strike for both the strike-hit major customers and their suppliers. Sales growth is the growth in sales relative to the same quarter in the previous year. The solid line connects coefficient estimates of sales growth on dummies indicating whether a firm is hit by a labor strike in the subsequent four quarters, the current quarter, and the previous three quarters. It is estimated on the customer sample. The dashed line connects coefficient estimates of sales growth on dummies indicating whether a firm's major customer is hit by a labor strike in the subsequent four quarters, the current quarter, and the previous three quarters. It is estimated on the supplier sample. Both regressions include firm and year-quarter fixed effects. The sample period spans 1983 to 2013. The quarter of the strike (t0) is highlighted in red. T-statistics presented in parentheses are based on robust standard errors clustered at the firm level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix A.1.

1.4.2.2 Baseline results

In the supplier sample, I estimate the regression stated in equation 1.4. The dependent variable is the firms' sales growth relative to the same quarter in the previous year. The variable of interest is the dummy *strike hits customer*$_{t,t-1}$ indicating that a labor strike occurs at a customer in the current quarter or the previous quarter. Table 1.4 presents the results. In Column 1, the coefficient estimate is -0.028 and statistically significant. In Column 2, I include lagged controls for firms' size, age, and profitability to the model. The size of the coefficient estimate slightly decreases to -0.019. In Columns 3, I add industry-year fixed effects to control for industry-wide shocks. The coefficient estimate is -0.024. In Column 4, I further introduce state-year fixed effects to control for local economic development. The coefficient estimate is -0.026. Overall, these results provide evidence that production disruptions at customers impose a substantial negative effect on suppliers' sales growth in the following two quarters. Affected suppliers' sales growth decreases by about one-third relative to the sample mean.

Table 1.4: Upstream propagation: baseline results

This table presents regression estimates of firms' sales growth relative to the same quarter in the previous year on a dummy indicating whether one of their major customers is hit by a labor strike in the current or in the previous quarter. All regressions include firm and year-quarter fixed effects as well as the firms' lagged number of major customers (constructed as terciles). In Column 2, I control for firm-level characteristics (constructed as terciles of lagged size, age, and return on assets, respectively) interacted with year-quarter dummies. In Column 3, I include industry dummies using two-digit SIC codes interacted with year dummies. In Column 4, I include state dummies interacted with year dummies. Regressions contain all firm-quarters of the supplier sample between 1983 and 2013. T-statistics presented in parentheses are based on robust standard errors clustered at the firm level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix A.1.

	(1)	(2)	(3)	(4)
strike hits customer$_{t,t-1}$	-0.028***	-0.019**	-0.024**	-0.026***
	(-3.18)	(-2.10)	(-2.56)	(-2.72)
Firm FE	Yes	Yes	Yes	Yes
Year-quarter FE	Yes	Yes	Yes	Yes
Number of customers	Yes	Yes	Yes	Yes
Firm-level controls	No	Yes	Yes	Yes
Industry-year FE	No	No	Yes	Yes
State-year FE	No	No	No	Yes
Obs	132,562	132,560	132,553	132,549
Number - firms	4093	4093	4092	4092
R2	0.154	0.177	0.214	0.232

In a robustness test, I show that suppliers' output loss increases with the severity of the labor strike, particularly when the strike lasts long. Appendix A.3

presents the results. Appendix A.4 shows a further robustness test to rule out that the results are driven by firm-specific seasonal trends uncaptured by the year-quarter fixed effects. First, I introduce firm dummies interacted with four dummies indicating the corresponding fiscal quarters. Second, I consider only firms whose fiscal quarters match the calendar quarters. The results remain very similar.

1.4.2.3 Suppliers' direct dependence on customers

I quantify suppliers' direct dependence as the fraction of sales from the customer in the year before the strike. This allows me to test the theoretical prediction that the output loss of suppliers should increase with their direct dependence on the disrupted customer. Instead of the dummy indicating that a strike hits the customer in the current quarter or the previous quarter, I use this fraction as a continuous treatment measure for suppliers' direct dependence on the strike-hit customer.

Table 1.5 presents the results. In line with the theoretical predictions, the coefficient estimates on the fraction of total sales from the strike-hit customers are negative and statistically significant, ranging from -0.12 to -0.091. These estimates indicate that suppliers that are less dependent on the strike-hit customers (10th percentile) experience a decrease by about 0.7 percentage points in sales growth, while suppliers that are highly dependent (90th percentile) experience a drop by about 4.2 percentage points.[18] Overall, these results show that the output loss is much stronger for suppliers that are more dependent on the strike-hit customers.

[18]The mean value of this fraction is 21%, the 10th percentile 7%, and the 90th percentile 40% for the suppliers whose customer is hit by a strike.

Table 1.5: Upstream propagation: suppliers' dependence on strike-hit customers

This table presents regression estimates of firms' sales growth relative to the same quarter in the previous year on a continuous variable measuring the lagged fraction of suppliers' total sales from the major customer that is hit by a labor strike in the current or the previous quarter. All regressions include firm and year-quarter fixed effects as well as the firms' lagged number of major customers (constructed as terciles). In Column 2, I control for firm-level characteristics (constructed as terciles of lagged size, age, and return on assets, respectively) interacted with year-quarter dummies. In Column 3, I include industry dummies using two-digit SIC codes interacted with year dummies. In Column 4, I include state dummies interacted with year dummies. Regressions contain all firm-quarters of the supplier sample between 1983 and 2013. T-statistics presented in parentheses are based on robust standard errors clustered at the firm level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix A.1.

	(1)	(2)	(3)	(4)
strike hits customer$_{t,t-1,\%sales}$	-0.12***	-0.091**	-0.10**	-0.11**
	(-2.63)	(-2.02)	(-2.26)	(-2.32)
Firm FE	Yes	Yes	Yes	Yes
Year-quarter FE	Yes	Yes	Yes	Yes
Number of customers	Yes	Yes	Yes	Yes
Firm-level controls	No	Yes	Yes	Yes
Industry-year FE	No	No	Yes	Yes
State-year FE	No	No	No	Yes
Obs	132,514	132,512	132,505	132,501
Number - firms	4093	4093	4092	4092
R2	0.154	0.177	0.214	0.232

1.4.2.4 Suppliers' direct and indirect dependence on customers

Next, I test the theoretical prediction that additional indirect links between suppliers and the disrupted customers amplify the suppliers' output loss. These additional indirect links exist if suppliers sell products to other companies whose business also depends on the strike-hit firms. For example, if the supplier sells products to both the strike-hit firm and one of the strike-hit firm's other suppliers. I refer to such a customer-supplier relationship as a supplier's central customer. I introduce a dummy variable indicating whether a strike hits the firm's central customer (*strike hits central customer*$_{t,t-1}$). Using a difference-in-differences setting, I compare the effect on suppliers whose central customer is hit by a strike and suppliers whose customer is hit by a strike.

Table 1.6 presents the results. The estimates on the baseline treatment dummy indicating whether a strike hits one of the firm's customer in the current or the previous quarter range from -2.3 percentage points to -1.4 percentage points. In Column 2, the coefficient estimate is marginally statistically insignificant (t-statistic of -1.47). For the dummy indicating whether a strike hits a central customer, the coefficient estimates are substantially larger, ranging from -7.0 to -5.4 percentage points. Hence, the output loss of suppliers with direct and additional indirect links increases by a factor of three to four compared to suppliers with only a direct link. Altogether, these results provide evidence that additional indirect links to the strike-hit firm substantially amplify suppliers' output loss.

1.4.3 Customers' firm size and the upstream propagation of shocks

In this section, I explore the role of customers' firm size in the upstream propagation of shocks. The previous results provide evidence that the propagation of shocks from customers to suppliers is accelerated by suppliers' direct de-

Table 1.6: Upstream propagation: suppliers with direct and indirect links to strike-hit customers

This table presents regression estimates of firms' sales growth relative to the same quarter in the previous year on a dummy indicating whether one of their major customers is hit by a labor strike in the current or in the previous quarter and a dummy indicating whether the strike-hit firm is a firm's central customer. A central customer is a firm's direct customer (direct link) and also a customer of the firms' other customers (indirect link). To identify such indirect links between a supplier and its customer, I assume that a firm is a permanent customer of a supplier in all year after the supplier reports the firm the first time as one of its customers. All regressions include firm and year-quarter fixed effects as well as the firms' lagged number of major customers (constructed as terciles). In Column 2, I control for firm-level characteristics (constructed as terciles of lagged size, age, and return on assets, respectively) interacted with year-quarter dummies. In Column 3, I include industry dummies using two-digit SIC codes interacted with year dummies. In Column 4, I include state dummies interacted with year dummies. Regressions contain all firm-quarters of the supplier sample between 1983 and 2013. T-statistics presented in parentheses are based on robust standard errors clustered at the firm level. ***, ** and * indicate significance on the 1%-, 5%- and 10%-levels, respectively. A detailed description of all variables can be found in Appendix A.1.

	(1)	(2)	(3)	(4)
strike hits customer$_{t,t-1}$	-0.023**	-0.014	-0.018*	-0.020**
	(-2.54)	(-1.47)	(-1.89)	(-2.01)
strike hits central customer$_{t,t-1}$	-0.054**	-0.057**	-0.064***	-0.070***
	(-2.20)	(-2.30)	(-2.60)	(-2.87)
Firm FE	Yes	Yes	Yes	Yes
Year-quarter FE	Yes	Yes	Yes	Yes
Number of customers	Yes	Yes	Yes	Yes
Firm-level controls	No	Yes	Yes	Yes
Industry-year FE	No	No	Yes	Yes
State-year FE	No	No	No	Yes
Obs	132,562	132,560	132,553	132,549
Number - firms	4093	4093	4092	4092
R2	0.154	0.177	0.214	0.232

pendence on the customers and suppliers' additional indirect dependence on the customers. Further, the aggregate effect of the propagation increases with customers' number of suppliers. To better understand the role of firm size, I sort firms in the customer sample into firm-size deciles and calculate the mean values of the three parameters for the respective deciles. Figure 1.3 illustrates the results.

Subfigure (a) presents the relationship between customers' size and suppliers' direct dependence. The mean value of suppliers' dependence is monotonically increasing with customers' firm size. Small- to medium-sized firms (deciles 1 to 7) account for 7 to 13% of suppliers' sales, larger firms (decile 8 and 9) for about 15%, and the largest firms (decile 10) for 18%. Subfigure (b) presents the relationship between customers' size and the likelihood to be a supplier's central customer. A central customer is a firm's direct customer and also a customer of the firm's other customers. The mean value of the likelihood to be a supplier's central customer increases monotonically with firm size and shoots up for the very largest firms. Small- to medium-sized firms have, on average, a likelihood of 5% and below, larger firms of 8 to 9% and the largest firms of 28%. Subfigure (c) presents the relationship between customers' size and the number of suppliers. It follows a very similar pattern. Small to medium-sized firms have, on average, 1 to 3 publicly listed suppliers, larger firms 4 to 5, and the largest firms 16.[19] These results indicate the vulnerability of production networks to idiosyncratic shocks to very large firms.

1.5 Conclusion

This paper studies the role of customer concentration in the upstream propagation of idiosyncratic firm-level shocks. Using customer-supplier links reported

[19] In untabulated results, I formally test the differences between very large customers (decile 10) and large customers (decile 8 and 9). They are statistically significant.

Figure 1.3: Firm size and upstream propagation of shocks
This figure illustrates the role of customers' firm size for the upstream propagation of shocks. Firms in the customer sample are sorted into deciles of real firm size. Subfigure (a) presents the mean direct dependence of suppliers per customer size decile. The direct dependence of a customer's suppliers is measured by the firm-level mean of suppliers' fraction of sales from the firm. Subfigure (b) presents the mean likelihood to be (at least one) supplier's central customer per customer-size decile. A central customer is a firm's direct customer (direct link) and also a customer of the firms' other customers (indirect link). To identify such indirect links between a supplier and its customer, I assume that a firm is a permanent customer of a supplier in all years after the supplier reports the firm as one of its customers for the first time. Subfigure (c) presents mean customers' number of suppliers per customer-size decile. A detailed description of all variables can be found in Appendix A.1.

(a) Customer size and suppliers' direct dependence

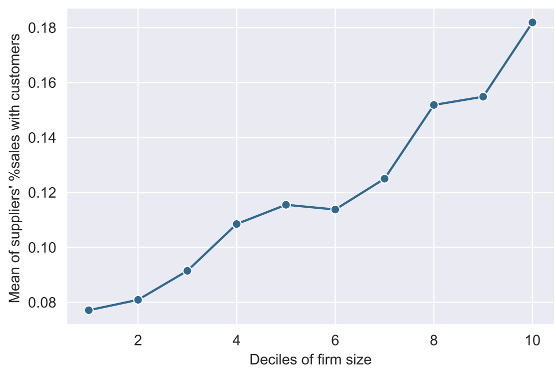

(b) Customer size and suppliers' indirect dependence

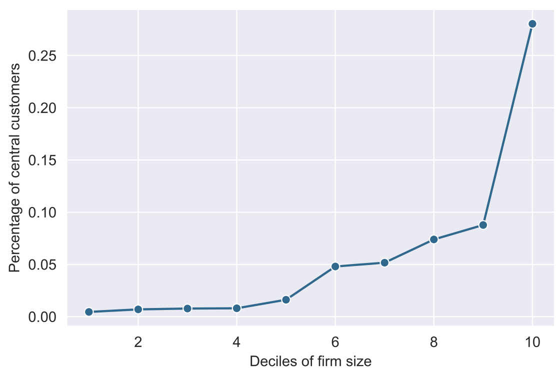

(c) Customer size and number of suppliers

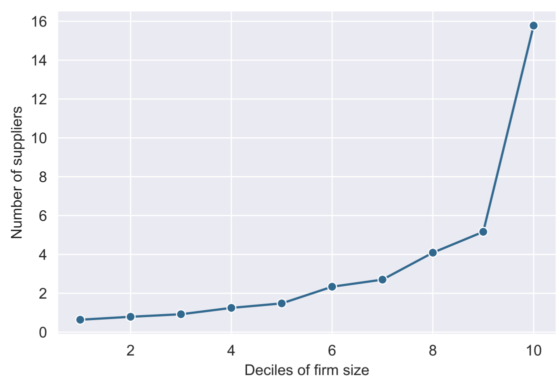

by publicly listed U.S. firms, I show that strike-hit customers impose a substantial output loss on their suppliers. Suppliers experience an average drop in sales growth by 1.9 to 2.8 percentage in the quarter of the strike and the following quarter, which equals a decline by about one-third. In line with theoretical predictions, the negative effect on suppliers' sales growth increases with their direct dependence on the strike-hit customers. Next, I show that suppliers' output loss increases by a factor of three to four if the suppliers sell products directly to the disrupted firm and to other customers whose business also depends on the central disrupted customer. Finally, I show that suppliers direct dependence on customers, the indirect dependence on central customers, and the number of suppliers all monotonically increase with the customers' firm size. Comparing large customer firms to the largest customer firms shows a substantial increase in all three measures for the largest firms. Overall, my results provide evidence that production networks are vulnerable to idiosyncratic shocks to very large firms.

Fools say that they learn by experience. I prefer to profit by others experience.

Otto von Bismarck, first German Chancellor

2

Do Outside Directorships Influence CEO Decision-Making? Evidence from Labor Strikes

Authors:	Daniel Bias, Thomas Schmid
First author:	Daniel Bias
Keywords:	Outside directorships, labor strikes, salience theory, observational learning, cash holdings
JEL codes:	G32, G34 J52, L14
Presentations:	American Finance Association Meeting 2017 (Chicago) Zuhai Finance International Conference 2018 (Zuhai)
Current status:	Working paper (https://dx.doi.org/10.2139/ssrn.2888378)

Abstract:

CEO outside directorships are an important phenomenon; however, little is known about their influence on managerial decision-making. We investigate how CEOs react after they observe, as directors of another firm, a labor strike that is plausibly exogenous to their firm. They increase cash holdings shortly afterwards, most likely due to an overreaction to the more salient strike risk. In the long run, CEOs adjust their labor negotiations. Consistent with observational learning, they offer higher wages during contract negotiations and manage to reduce strike risk. These results suggest that outside directorships can facilitate both behavioral biases and observational learning.

2.1 Introduction

The appointment of outside CEOs as directors is a frequent phenomenon. Fahlenbrach, Low and Stulz (2010) identify nearly 2,000 CEO outside directors in a sample of about 5,000 firms, which corresponds to five percent of all directors. In larger firms, CEOs are even more common on boards of directors. According to the SpencerStuart Board Index, 52% of CEOs in S&P 500 firms served on at least one additional public board in 2007, decreasing to 37% in 2017. Despite their relevance, little is known about how such outside positions affect managerial decision-making.

In this paper, we analyze how CEOs react when they observe, as directors, a labor strike at another firm. This observation can make strike risk more salient for the CEO and lead to more precautionary behavior in the short term. The psychology literature has shown that the attention placed on an event or risk factor plays a key role in its perceived probability (see, among others, Bordalo, Gennaioli and Shleifer, 2012, for an overview on salience theory). On the other hand, observing a labor strike as a director may enable the CEO to gain insider knowledge and learn about labor negotiations (see, among others, Bikhchandani, Hirshleifer and Welch, 1998, for an overview on observational learning). In the long run, this may affect their own bargaining with labor.

Consider the example of John E. Bryson, who served as CEO for Edison International in 2000 and simultaneously on the board of directors of Boeing Co. During this year, Mr. Bryson observed how the CEO of Boeing, Philip M. Condit, had to deal with a costly 41-day strike by about 22,000 workers. Cole (2000), in the Wall Street Journal, reported a controversial discussion in the boardroom of Boeing about this labor dispute. Observing this strike as a director could have affected Mr. Bryson's decision-making in two ways. First, it could have directed his attention to strike risk and made this risk factor more salient to him. Second, he might have learned from observing why the labor negotiations failed and how Mr. Condit handled the strike.

There are three reasons why labor strikes are useful in analyzing how outside directorships influence CEOs' decision-making. First, strikes are very costly events for firms that require a lot of attention from senior management and the board.[1] In our sample, we find that the median operating cash flow drops by about one-quarter when a firm is experiencing a strike. More generally, Becker and Olson (1986) documented the importance of strikes for the overall financial market. Krueger and Mas (2004), Mas (2008), and Gruber and Kleiner (2012) showed that strikes affect product and service quality.

Second, strikes are idiosyncratic events that are plausibly exogenous for the firms of the CEOs.[2] We show that strikes have no predictive power for the strike risk of firms in other industries. However, some correlations between strikes do occur within industries, as these firms often share the same union. We consequently exclude all events for which the CEO and director firm are from the same industry (and, in a robustness test, from the same state). For the remaining events, it is unlikely that the strike at the director firm has any direct effect on the strike risk of the CEO firm.

Third, in contrast to many other types of shocks, such as the sudden death of a CEO or weather-related losses, strikes are not exogenous for strike-hit firms. Rather, they are related to prior management decisions and, in many cases, long-term mismanagement of labor relations.[3] It has also been argued

[1]An example of a costly strike is the 45-day work stoppage by about 36,000 employees in the landline division of Verizon Communications. For the second quarter of 2016, Verizon reported $400m in incremental costs due to strike activities. Another example is the 40-day strike of the United Auto Workers union against General Motors in 2019, which cost the company between $3.8bn and $4bn.

[2]Gabaix (2011) considers strikes to be one possible idiosyncratic firm-level shock that explains aggregate business cycles.

[3]Imberman (1983) finds that "most strikes have little to do with money or benefits. Workers vote for a strike only when they are frustrated because their needs, wants, and ideas go unheard, unheeded, or unanswered. In such cases, workers see managers as adversaries. By understanding how employee dissatisfaction manifests itself in the early stages, managers have a chance to deal with the problems before they lead to a strike".

that sending credible signals about a firm's profitability is a key mechanism for avoiding the escalation of labor conflicts (e.g., Card, 1990a). Because CEOs can influence labor relations and signaling at their own firms, observational learning is a plausible mechanism for strikes to affect CEOs' decision-making.

Overall, 790 strikes occurred at 303 different S&P 1500 firms between 1984 and 2011, enabling us to identify 215 events of CEOs who observed a strike as directors. We start by analyzing whether CEOs who observe a strike (temporarily) overreact to the more salient labor risk. We focus on the precautionary behavior of CEOs and measure how they adjust quarterly cash holdings. In different contexts, the previous literature has shown that firms hold liquidity to protect themselves from adverse shocks (see Almeida et al., 2014, for a survey on this literature).[4]

An empirical challenge in our study is that outside directorships of CEOs are endogenously formed, and appointments of CEOs to boards of others firms are not random (Fahlenbrach, Kim and Low, 2018). CEOs who serve as directors at strike-hit firms may be different from the general population of CEOs (e.g., with regard to their level of risk aversion). Although our focus on time-series variation within firms already reduces this potential concern to some extent, we also use two other strategies to address it. First, we analyze not only the strike observation itself but also the severity of the strike. The severity of a strike depends on many factors and is difficult to predict, often even at the start of a labor dispute. Hence, strike severity is plausibly exogenous to factors related to a director-CEO. Second, our control group consists of firms with a CEO who also sits on the board of a strike-hit firm but did not during the strike. These firms, in which the CEO either left the board before the strike or joined afterwards, are more comparable to our event firms than the average firm.

[4]Bargaining provides an alternative mechanism for how strike observations could affect cash holdings. In this perspective, CEOs decrease cash holdings to improve their bargaining position with labor (e.g., Klasa, Maxwell and Ortiz-Molina, 2009; Matsa, 2010; Benmelech, Bergman and Enriquez, 2012; Myers and Saretto, 2016).

We find that CEOs who observe a labor strike increase cash holdings by about 0.7 percentage points, which represents a 10% increase in relative terms. This effect is more pronounced for strikes with a long duration, strikes in which a high percentage of the workforce participates, and strikes that lead to a high loss for the strike-hit firm. When we analyze the time pattern of the cash changes, we find that the increase starts in the quarter of the experience, reaches a maximum four quarters thereafter, and then reverts. This pattern is consistent with an overreaction to more salient labor risk because the salience of the strike observation diminishes over time.

Next, we focus on possible long-term effects due to observational learning and analyze two aspects of labor relations: wage setting in labor contract negotiations and strikes as a consequence of failed labor negotiations. The data on wage setting comes from the settlement summaries database of Bloomberg BNA, which provides details on the settlement terms of collective bargaining agreements (e.g., wage changes). We find that CEOs tend to agree to higher wages in the years after their strike observation compared to other firms with contract settlements. At the same time, they manage to decrease the strike risk for their firm, especially after observing a severe strike. Agreeing to higher wages is one likely channel through which they reduce strike risk, but they may also adjust other aspects of labor relation management that we cannot measure. These findings indicate that strike-observing CEOs learn from insider information about labor negotiations.

To summarize, we show that outside directorships influence CEOs' decision-making. In the short run, their attention is directed towards risk factors that they observe at the director firm. This leads to a temporary overestimation of that risk factor and precautionary behavior. Since this precautionary behavior seems to be an irrational overreaction that is most likely not in the best interest of the firm, this mechanism could be interpreted as a "dark side" of CEO outside directorships. However, there is also a “bright side”, since our findings suggest that CEOs learn from the insider knowledge they receive at the director firm. In the long run, this observational learning affects their own behaviors as CEOs.

We contribute to different strands of the literature. Most importantly, we add to the literature on the appointment of CEOs as directors. Fahlenbrach, Low and Stulz (2010), among others, argue that there is high demand to appoint CEOs as outside directors. On one hand, learning effects are potential benefits of outside directorships for the CEO firm (e.g., Booth and Deli, 1996; Carpenter and Westphal, 2001; Perry and Peyer, 2005). Our study provides direct empirical evidence for the existence of such learning effects due to outside directorships. On the other hand, time and efforts spent for outside directorships are potential reasons against CEO outside directorships and board busyness in general (e.g., Fich and Shivdasani, 2006; Field, Lowry and Mkrtchyan, 2013). Our results indicate that not only do time and effort spent by CEOs matter, but their decision-making can also be adversely affected due to behavioral biases that direct too much attention to risk factors they observed as directors.

Second, we contribute to the literature on salience and finance. Bordalo, Gennaioli and Shleifer (2012) and Bordalo, Gennaioli and Shleifer (2013), for instance, develop a theoretical model in which more salient extreme payoffs have a disproportionately high weight for the pricing of assets. More related to our study, Dessaint and Matray (2017) discuss evidence that managers overreact to salient risk by analyzing cash adjustments when hurricanes hit firms with close proximity. In contrast to their study, our focus lies on the board channel and a conceptually different risk factor: labor. Unlike natural disasters, strike risk can be influenced by CEOs because they are not exogenous. Rather, strikes are the result of failed negotiations. Thus, adjustments in bargaining and labor relation management can affect strike risk. This mechanism, which does not exist for exogenous risk factors, allows us to show that observing others’ behaviors not only directs CEOs' short-term attention to the involved risk factors but also enables them to learn in the long run.

Lastly, we also add to the literature on how managers' (past) experiences affect their own decision-making. Bernile, Bhagwat and Rau (2017), for instance, find a non-monotonic relationship between managers' early-life exposure to natural disasters and later risk-taking in their corporate decision-making. Dittmar and Duchin (2015) show that previous professional experiences of managers matter for corporate policies. Studies that demonstrate a relationship between previous experiences of economic conditions and managerial decision-making include Malmendier and Nagel (2011), Malmendier, Tate and Yan (2011), and Schoar and Zuo (2017). In contrast to those papers, we do not focus on experiences that affect managers directly, but their observations of others' behavior.

2.2 Theory and empirical strategy

2.2.1 Salience theory

Observing how another CEO deals with a major risk factor may direct the attention of the director-CEO towards this risk. The psychology literature has shown that the perceived likelihood of events increases if more attention is paid to them due to the so-called availability heuristic (Tversky and Kahneman, 1973, 1974).[5] Availability refers to how easily people can think of similar events or occurrences, and one factor that affects this availability is salience. As Tversky and Kahneman (1974) explain, "[...] the impact of seeing a house burning on the subjective probability of such accidents is probably greater than the impact of reading about a fire in the local paper" (p. 1127). In our context, this means that observing a labor strike as a director increases the subjective labor risk to

[5]See DellaVigna (2009) for a survey of the literature on how psychological biases can affect the decision-making of individuals and Malmendier (2018) for an overview on behavioral corporate finance.

a greater degree than reading or hearing about strikes at other firms.

Furthermore, Tversky and Kahneman state that "recent occurrences are likely to be relatively more available than earlier occurrences. It is a common experience that the subjective probability of traffic accidents rises temporarily when one sees a car overturned by the side of the road" (p. 1127). In our context, this means that CEOs' overestimation of labor risk should be a temporary bias. Over time, the availability of the strike observation decreases, and the salience of the labor risk becomes smaller. Thus, in the long run, the behavioral bias should disappear because the CEO's subjective overestimation of the labor risk reverts back.[6]

This over-estimation of the labor risk factor, due to the strike observation as director, can affect the decision-making of the CEO. In particular, more salient labor risk can lead to precautionary behavior in the CEO. For unionized firms, CEOs may act more cautiously and build up liquidity buffers because they overestimate the risk that their own firm will be hit by a strike.[7] Strikes are costly events that can have disastrous consequences for firms. As mentioned in the introduction, in our sample, the median operating cash flow drops by about one-quarter when a firm is experiencing a strike. Liquidity buffers can help to mitigate the adverse effects of strikes (e.g., by hiring replacement workers).

For non-unionized firms, CEOs can overestimate the strength of labor unions and the risk that their firms could become unionized in the near future. To attain legal recognition under the National Labor Relations Act, a union must

[6]See Dessaint and Matray (2017) for a similar argumentation in the context of CEOs who observe how a firm that is located in close proximity is hit by a hurricane. Although this does not affect the probability of this risk factor for the CEO firm, its perceived probability seems to increase temporarily, leading to increased cash holdings.

[7]The National Labor Relations Board distinguishes two classes of lawful objects for strikes. These are "unfair labor practice strikes" (ULPSs) and "economic strikes" (ESs). ULPSs protest against unfair labor practices committed by the employer (e.g., a refusal to bargain with the union). ESs pursue the objective of obtaining some economic concession (e.g., higher wages). While ESs occur during contract negotiations, ULPSs do not have to occur during contract negotiations. For this study, we do not consider other ways how labor can exert pressure on management, for example, via legal benefits of unionization (e.g., Campello et al., 2018).

win over 50% of the workers' votes in a secret ballot. After this, the employer is obligated to negotiate "in good faith" with the union.[8]

In both unionized and non-unionized firms, CEOs can overestimate the strike risk of unionized supplier or customer firms. Labor conflicts adversely affect not only strike-hit firms but also, sometimes, their customers or suppliers. McHugh (1991), for instance, finds that industries linked as suppliers or customers to strike-hit industries experience a stronger decline in productivity than the strike-hit industry itself. An example of a costly strike at a customer firm is Rockwell Collins, which is a supplier to Boeing Co. When Boeing Co. was hit by a 2-month strike of 27,000 employees in 2008, the quarterly sales of Rockwell Collins dropped by about 14%. American Axle is an example of a supply-side strike. This firm, which is a key supplier of General Motors, was hit by a strike of more than 50% of its workers in 2008. General Motors estimated that this strike cost them $2.6bn in the first and second quarters of that year.

2.2.2 Observational learning

Observational learning is the "influence resulting from rational processing of information gained by observing others" (Bikhchandani, Hirshleifer and Welch, 1998, p. 153).[9] Director-CEOs gain information when they observe how another CEO deals with a labor dispute. In contrast to firm outsiders, director-CEOs have access to insider information about the bargaining process and the adverse effects of the strike on the firm, among other things.

There are two possible mechanisms through which director-CEOs can learn from insider information. The first is related to the bargaining that occurred before the strike. In line with the view of Hicks (1963) that "the majority of actual strikes are doubtless the result of faulty negotiation" (p. 146), director-

[8]For a more detailed description of the legal recognition process of a union, please refer to DiNardo and Lee (2004).

[9]A related concept to observational learning is peer learning. See, among others, Shue (2013) and Ouimet and Tate (2020).

CEOs can infer that the bargaining of the CEO at the director firm was not optimal if it resulted in a strike, learn from the unsuccessful negotiations, and adjust their own bargaining strategy with labor in order to avoid strikes at their own firm.

For a highly simplified illustration, assume that a CEO, Barbara, sits on the board of another firm. The CEO of this other firm, Aaron, is currently negotiating with a union about a new labor contract. There is an optimal level of bargaining, i.e., an optimal offer to the union, that minimized the labor cost of the firm while at the same time avoids a strike (denoted by "X").[10] This optimal offer X depends on firm-level factors (denoted by "F") such as its profitability or the costliness of a strike, and union-level factors (denoted by "U"), such as the union's willingness and ability to call for a strike. Aaron has complete information on F, but incomplete information on U. Thus, he must infer U from the union's imperfect signals (S). Because X is not known with certainty, Aaron chooses a bargaining level x based on F and S.

The outcome of the negotiations is binary and either no strike (new contract) or strike. We assume that no strike is always preferable for the firm. The outcome depends on Aaron's bargaining level x. If he bargains too hard and his offer is too low ($x > X$), the union will call for a strike. Barbara observes this outcome as a director and receives the information that Aaron chose an x that is too high. She then processes this information and learns that the optimal level X is smaller than Aaron's x. When she herself negotiates with a union in the future, her estimate of the optimal level is smaller and she rationally chooses a lower x (for the same levels of F and S).[11] Her rational decision to bargain less

[10]Several economic papers model strike incidents as outcomes of the bargaining process between unions and firms (e.g., Ashenfelter and Johnson, 1969; Hayes, 1984; Hart, 1989; Card, 1990b; Cramton and Tracy, 1992).

[11]This illustration could also be extended to explain why more severe strikes have a stronger effect on decision-making. CEO-director Barbara could infer that Aaron negotiated much too hard if the strike is very severe ($x >> X$). Thus, she would reduce her estimate of the optimal bargaining level more substantially than if a mild strike had been observed.

hard and increase the offer to the union leads to higher wage increases and a lower strike probability at her own firm.[12]

Second, director-CEOs can learn from insider information about the true cost of strikes. Some information on the financial cost of strikes for firms is publicly available. For example, General Motors announced that the recent 40-day strike by United Auto Workers cost them between $3.8bn and $4bn. However, firms have incentives to under-report the costliness of strikes to discourage unions from striking in the future. Furthermore, there may be long-term consequences that are difficult to observe for outsiders. These can include reputational damages, if suppliers or customers were adversely affected by the strike or negative consequences on the workforce due, for example, to conflicts between workers who participated in the strike with those who did not.

If the strike-observing CEOs update their beliefs about the true cost of strikes, they will rationally adjust their bargaining with unions. Using the previous illustration, Barbara would increase her beliefs about the costliness of a strike for the firm (which is part of F). This would affect her bargaining behavior: she would bargain less and improve her offer to the union, leading to higher wage increases and a lower probability of a strike occurring.

Both learning mechanisms suggest that strike-observing CEOs increase their offers to unions to reduce the probability of striking when negotiating with labor at their own firm. Since labor negotiations are infrequent events that mainly occur at the end of a labor contract (the average contract length is three years in our sample[13]), those effects typically materialize only in the long term. To summarize, salience theory suggests that CEOs increase cash holdings

[12]The opposite learning effect occurs if Aaron does not negotiate hard enough ($x < X$). In this case, Barbara observes that Aaron bargained too little and offered to much to the union, and she would reduce her offer to the union, increasing the strike probability. However, successful labor negotiations may not be very visible for directors because they require less attention from the board.

[13]Chava, Danis and Hsu (2020) find an average labor contract duration of three years in the full sample of settlement summaries from Bloomberg BNA between 1988 and 2016.

temporarily after observing a labor strike, while observational learning can affect CEOs' bargaining with labor in the long run and lead to higher wage increases and lower strike risk.

2.2.3 Empirical strategy

In general, strikes are idiosyncratic events, and the shock to a strike-hit firm is plausibly exogenous to the firm of the director-CEO. However, this may not necessarily be true if the CEO firm and the strike-hit firm are from the same industry, in a supplier-customer relationship, or located in the same state. To rule out potential correlations of strike activities within industries, we exclude all CEO strike observations for which the strike-hit firm and the CEO firm operate in the same Fama-French 12 industry. Using our strike dataset, we show that strikes have predictive power for the strike risk of firms in the same industry, but not for those in other industries (cf. Section 2.4.3 and Table 2.4). Furthermore, we generally exclude events if the CEO and the strike-hit firm are in a customer-supplier relationship to avoid direct spillover effects of labor strikes on the CEO firm. We also show that our results are unaffected by the exclusion of events in which both firms are from the same state.

The main empirical challenge for our study is that outside directorships of CEOs are endogenous. There are reasons why some CEOs serve as directors on the boards of other firms, whereas others do not (Fahlenbrach, Kim and Low, 2018). This leads to the concern that director-CEOs may be different from the general population of CEOs, for example, with regard to their level of risk aversion. Furthermore, those CEOs who serve on the board of a strike-hit firm may have different characteristics than non-director CEOs or director-CEOs at firms that are not hit by a strike.

In our empirical analysis, we analyze how a CEO's decision-making changes over time after a strike observation. Thus, we do not directly compare the behaviors of director-CEOs and non-director-CEOs, which reduces concerns about their (potential) different characteristics to some extent. Nevertheless, different characteristics could lead to different time dynamics, and we follow two strategies to address this potential concern.

First, the severity of strikes is often difficult to predict, even at the start of a labor dispute.[14] Although CEOs who serve as directors at strike-hit firms may be different from the general population of CEOs, the severity of a strike is plausibly exogenous to director-CEO characteristics. We use the strike duration as our main measure for severity because even a small group of employees can cause serious problems for a company.[15] As alternative measures, we use the ratio of striking employees to total employees, the idled employee-days during the strike, and the changes in return on assets and operating cash flow following the labor strike.

Second, we use a control group that consists of firms whose CEOs also sit on the board of directors of another firm that is hit by a strike, but this strike occurred before or after the CEO served on the board of directors. The selection process for the outside directorship of both the strike-observing CEOs and the control CEOs should be comparable in this specification because the only difference is the time at which they joined or left the board. In our robustness tests, we alternatively estimate the effect only for treated firms.

To investigate CEOs short-term precautionary behavior after observing a strike, we use quarterly cash holdings as the dependent variable. Cash holdings are defined as cash and cash equivalents over total assets. We measure how these cash holdings change between the eight quarters before the strike observation

[14]There are several reasons why predicting the severity of a strike is difficult. First, collective bargaining in the U.S. usually takes place not at the firm level but the establishment level. Each local union can only call a strike with the consent of the majority of all represented employees. Second, although the upper bound of the number of striking employees is fixed at the beginning of a strike, its duration is still unforeseeable. Third, the severity of a union's labor strike may be amplified if other unions or even customers decide to honor picket lines.

[15]Consider a pilot strike that shuts down the whole company despite a relatively small fraction of striking employees (e.g., Comair in 2001 or Spirit Airlines in 2010).

(pre-period, q-8 to q-1) and the following eight quarters (post-period, q0 to q7). We focus on this relatively short window because we want to capture the CEOs' reaction directly after the strike observation while the labor risk is still salient to them. Our main regressions model, which includes these windows around the strike observations and the control firms, can be written as:

$$Cash_{i,q} = \alpha_i + \beta StrikeObs_{i,q} + X_{i,q-1} + \pi_q + \psi_{ind,y} + \eta_{ind,q'} + \varepsilon_{i,q}, \tag{2.1}$$

where i denotes a firm, q a fiscal year-quarter, q' a fiscal quarter, y a fiscal year, and *ind* an industry. α_i is a firm-fixed effect, $StrikeObs_{i,q}$ is a dummy that equals one in the quarter of the strike observation and the seven quarters thereafter, $X_{i,q-1}$ is a vector of lagged control variables, π_q are quarter-year fixed effects that control for general time patterns, $\psi_{ind,y}$ are industry times year fixed effects that control for any industry-specific yearly time dynamics, $\eta_{ind,q'}$ are industry-quarter fixed effects that control for any industry-specific seasonality, and $\varepsilon_{i,q}$ is the error term.[16] In a robustness test, we additionally add state times year fixed effects. The industry classification follows the Fama-French 12 industries classification. Robust standard errors are clustered at the firm level.

In the second step, we replace the strike observation dummy with the duration of the observed strikes to capture the heterogeneous treatment effects, depending on the severity of the observed strike. To show the robustness of our results, we apply alternative measures for the severity of a strike. These measures are the ratio of striking employees to total employees, the idled employee-days during the strike, and the changes in return on assets and operating cash flow following the labor strike.

To measure potential long-term effects from observational learning, we an-

[16]Please note that the illustration in Equation 2.1 is slightly simplified. For firms with more than one event, we estimate a separate α and ε for each event. Thus, the firm-event fixed effect for firm i and event e is $\alpha_{i,e}$ and the corresponding error term is $\varepsilon_{i,e,q}$.

alyze whether strike-observing CEOs adjust their labor relations management. We focus on the potential changes in the wage outcomes of collective bargaining agreements and long-term strike risk at the CEO firms. Since contract settlements and strikes are infrequent events, we use yearly instead of quarterly data, consider all available data before the strike observation, and use all other S&P 1500 firms with contract settlements (for wage regressions) or at least one strike (for strike regressions) as control groups. Furthermore, we are less interested in short-term effects, and more in long-term changes. Because of that, we examine the impact of a strike experience on wages and strike risk for time horizons of two, four, six and eight years after the strike observation. We set the strike observation dummy to one from the year of the strike observations until the end of the respective time period. We estimate the following two equations for wage changes and strike risk, respectively:

$$\Delta Wage_{i,y} = \alpha_i + \beta StrikeObs_{i,y} + X_{i,y-1} + \pi_y + \psi_{ind,y} + \varepsilon_{i,y}. \tag{2.2}$$

$$StrikeDummy_{i,y} = \alpha_i + \beta StrikeObs_{i,y} + X_{i,y-1} + \pi_y + \varepsilon_{i,y}. \tag{2.3}$$

$\Delta Wage_{i,y}$ is the wage change in the first year of the collective bargaining agreement, measured in percentage points. $StrikeDummy_{i,y}$, is a dummy variable that is set to one if a strike begins at the firm in the respective fiscal year and zero otherwise. Due to the binary nature of this variable, we use a conditional logit regression to estimate Equation 2.3.

2.3 Data

2.3.1 Sample and firm financials

Our sample consists of S&P 1500 firms between 1983 and 2012. Since the S&P 1500 did not exist before 1994, we use the index constituents in the year 1994 for the 1983-1993 period. We exclude firms from the financial services industry (SIC industry codes 6000-6999) and observations with missing information on cash and short-term investments, total assets, and common equity. Financial data is retrieved from the Compustat North America Fundamentals Quarterly database. We construct the Fama-French 12 dummies using a firm's historic 4-digit SIC industry code. All continuous variables are winsorized at the 1st and 99th percentiles. For the analysis of cash holdings, we impute missing values of lagged market leverage and return on assets by the next available value from the four previous quarters.

2.3.2 Strikes

Because there exists no comprehensive database on labor strikes, we combine information from the Bureau of Labor Statistics (BLS), the Federal Mediation and Conciliation Service (FMCS), the National Mediation Board (NMB), and labor contract data by Cramton and Tracy (1992). The BLS provides data on major labor strikes involving 1,000 or more workers. The FMCS collects information on labor strikes of any size. The NMB handles labor strikes in the railroad and airline industries. Cramton and Tracy provide a sample of 6,641 contract negotiations between 1969 and 1990 involving 1,000 or more workers.[17] Their dataset also includes information on strike incidents.

The minimum information that we require on a strike is the employer name,

[17]The labor contract data by Cramton and Tracy is available via the Inter-university Consortium for Political and Social Research (ICPSR 1020).

the number of striking workers, the state, the beginning date, and the end date. Further, we exclude smaller strikes with fewer than 100 striking employees and remove obvious data errors. Next, we manually match the strike events to the S&P 1500 firms using the employer name and validate the matching outcomes using press articles, 10-K filings, and online media. Finally, we group all strike events at one firm that start within the same calendar week into a single strike event.[18] This procedure leads to a sample of 790 labor strikes in the 1984-2011 period.[19]

2.3.3 CEOs observing a strike as director at another firm

To identify CEOs who observe a labor strike as a director at another firm, we require information on the CEOs and directors of S&P 1500 firms. Our primary data source is the BoardEx database, which provides comprehensive information on boards from 1999 onward. Since BoardEx researches the full biographical information of individuals, it also offers some information on CEO and director positions before 1999. As an alternative source before 1999, we manually screen the 10-ks and proxy statements of strike-hit firms for CEOs who serve as directors on the board during the strike period.

Additionally, we impose the following restrictions: (i) The CEO must be appointed as CEO before they observe the strike and must remain the CEO of the firm for at least one year after the strike observation. (ii) The strike-hit firm and the CEO firm must not stand in a customer-supplier relationship. We obtain information on customer-supplier relationships from Barrot and Sauvagnat

[18]We group the strike events by firm and calendar week since some strikes in the data are reported on firm level (aggregated over all establishments) and others on establishment level. Otherwise, strikes reported separately for multiple establishments may distort the summary statistics (e.g., strike duration). This grouping does not affect the number of CEO strike observations.

[19]The dataset also includes lockouts. However, strikes represent the vast majority of our events. Less than 5% of the CEO strike observations are related to lockouts. Similarly, Cramton and Tracy (1992) report that lockouts represent less than 4% of the labor disputes (strikes plus lockouts) in their dataset from 1970 to 1989.

(2016).[20] (iii) The strike-hit firm and the CEO firm must also not operate in the same Fama-French 12 industry. (iv) An event is only considered if there are at least 15 quarters between two strike observations. This restriction has two purposes. First, it helps to avoid repeated events, e.g., due to a long strike or a major labor dispute that leads to a series of consecutive strikes at a director firm. Second, it ensures that we have a clean pre-observation time period.

In total, we identify 201 CEOs of 185 S&P 1500 companies who observe 215 strikes at another firm between 1985 and 2010.[21] In Figure 2.1, each line connects the headquarters of a CEO firm (●) and a strike-hit director firm (■). Subfigure (a) illustrates the CEO strike observations for the whole sample period. Subfigure (b) illustrates the 14 CEO strike observations in the 1995 fiscal year and Subfigure (c) the 4 CEO strike observations in the 2009 fiscal year. From those graphs, it is obvious that the distance between the CEO firm's and director firm's headquarters varies considerably. In some cases, the two firms are located in the same state or adjacent states; however, very long distances are not uncommon.

2.3.4 Control groups

CEOs who observe a strike as a director of another firm may be different from the average CEOs, for example, with regard to their abilities or attitudes. To reduce concerns that differences in CEOs characteristics bias our results, we choose a control group of firms whose CEOs also serve on the board of strike-hit firms for our short-term cash regressions. The only difference to the strike-observing CEOs is that they either left the board before the strike or joined afterwards. Hence, the CEOs of the control control group likely undergo a similar selection

[20] They apply a matching algorithm to assign a firm's key customer names from Compustat Segment Files to Compustat identifiers. Please refer to their paper for further details on the matching approach.

[21] This period is shorter than 1983 to 2012 for two reasons. First, we need a pre-period of eight quarters to estimate the change in cash holdings after a CEO strike observation. Second, the post-period after the strike observation is 8 quarters.

Figure 2.1: Geographic location of the CEO firm and the strike-hit director firm
This figure presents the geographic locations of the CEO firms' headquarters (●) and the strike-hit director firms' headquarters (■). Subfigure (a) illustrates the 215 strike observations by CEOs between 1985 and 2010. We exclude 5 CEO strike observations from this illustration because either the CEO firm's or the director firm's headquarters is not located in the United States. Subfigure (b) illustrates the 14 strike observations by CEOs in 1995. Subfigure (c) illustrates the 4 strike observations by CEOs in 2009. A detailed description of all variables can be found in Appendix B.1.

(a) Strike observations by CEOs between 1985 and 2010

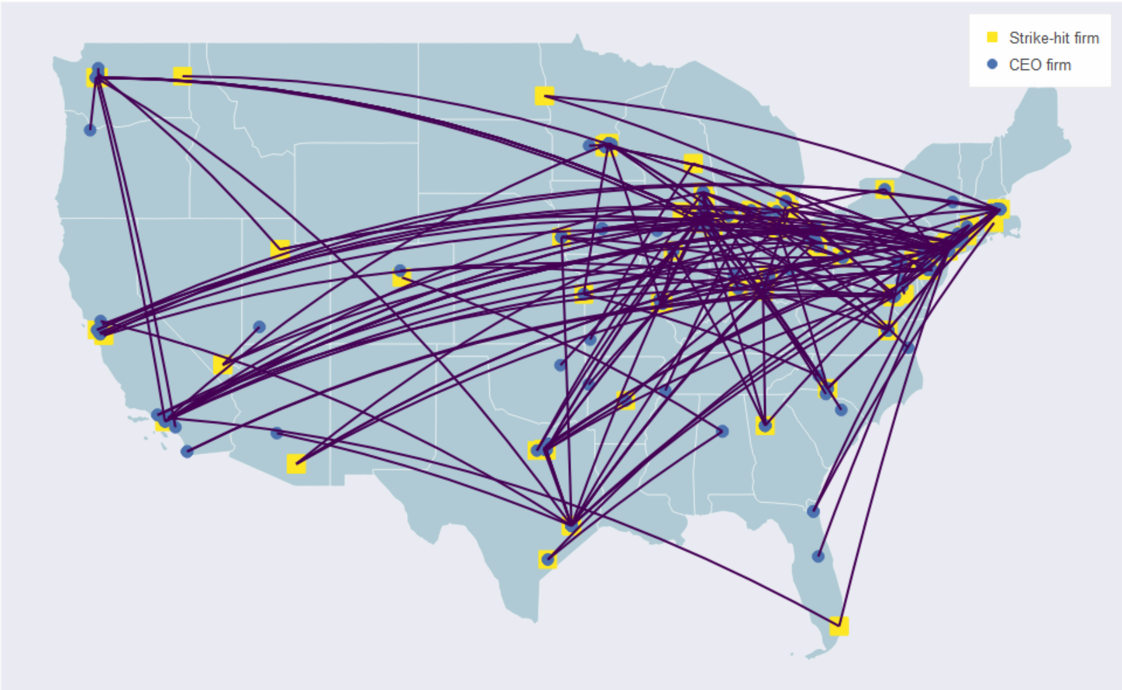

(b) Strike observations in 1995

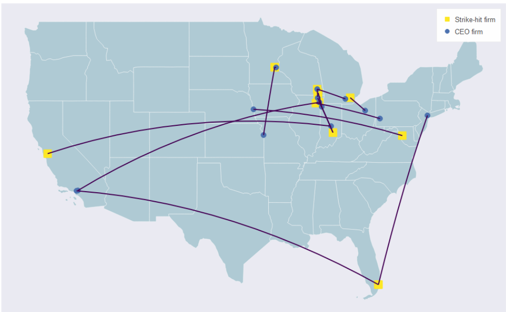

(c) Strike observations in 2009

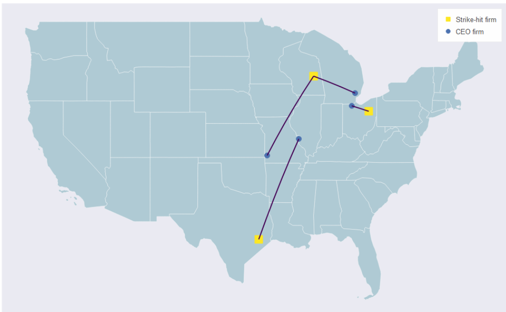

process as the strike-observing CEOs. This procedure leads to 96 control firms that account for 8,149 firm-year-quarter observations.

For the long-term wage and strike-risk regressions, we use all other S&P 1500 firms with contract settlements (for wage regressions) or at least one strike (for strike regressions) as control groups for two reasons. First, these firms are more natural controls since they also have contract settlements or strikes. Second, the number of firms in our short-term control group with contract settlements or strikes is too small to allow a meaningful empirical estimation.

2.3.5 Settlement of collective bargaining agreements

For the analysis of the effect of a CEO's strike observation on labor negotiations, we obtain information on collective bargaining agreements from the Settlement Summaries database of Bloomberg BNA. The dataset includes information on the employer, the employer's geographic location, the employer's SIC and NAICS codes, the union, the settlement date, the contract expiration date, and brief summaries of the agreed-upon wage changes and other provisions of the contract.

This dataset is available from 1988 on. Our initial sample contains about 23,000 collective bargaining agreements in the period from 1988 to 2012. To identify collective bargaining agreements of S&P 1500 firms, we apply a record linkage approach using the employer/firm name, the industry codes, and the geographic location. Afterwards, we manually review whether the linked collective bargaining agreement belongs to the firm. In total, we identify 2,844 collective bargaining agreements that belong to 453 firms.[22]

[22] The papers by Yi (2019) and Qiu and Zhang (2019) also link collective bargaining agreements to Compustat firms. However, they do not use the Settlement summaries data but the expiring labor contracts from Bloomberg BNA. This data product covers the universe of expiring labor contracts but does not provide information on the settlement terms. Yi (2019) identifies 377 firms that have at least one expiring labor contract with 500 or more employees during the 1995-2014 period. Qiu and Zhang (2019) identify 551 firms with at least one expiring labor contract with 100 or more employees during the 2005-2010 period.

To quantify the wage change under a labor contract, we follow the approach of Chava, Danis and Hsu (2020) who use the wage increase over the first year of each contract as a proxy for the total increase in wages. We develop an algorithm to extract the first-year wage increase from the text describing the wage changes. If a firm settles more than one labor contract per firm-year, we calculate the mean first-year wage increase over all contracts. Panel B of Table 2.1 provides descriptive statistics. We observe at least one labor contract settlement in 1,622 firm years. The mean wage increase in the first year is about 3.02%, and the mean contract length is about 1,280 days or 3.5 years.

2.3.6 Industry unionization

We obtain annual industry unionization rates for the 1984-2012 period from the Union Membership and Coverage Database by Barry Hirsch and David Macpherson.[23] This database provides the annual industry unionization rates for Census Industry Classification (CIC). Using a crosswalk table to NAICS, we map industry unionization rates to our sample firms.

2.3.7 Descriptive statistics

Table 2.1 presents descriptive statistics for the quarterly sample that we use to study cash adjustments after a CEO's strike observation and the yearly sample that we use to analyze the long-term effect on labor relations management. Panel A presents summary statistics on the quarterly sample, which consists of 11,505 firm-quarters and 281 individual firms. Based on the 215 observed strikes, we create a strike observation dummy that equals one in the quarter of the strike observation and the seven quarters after that. In total, the strike observation dummy is equal to one for 14.7% of our sample (1,692 observations). Cash

[23]Publicly available at www.unionstats.com; see Hirsch and Macpherson (2003) for a detailed description of the database.

holdings are defined as cash and cash equivalents over total assets. The mean (median) cash holdings are 7.3% (3.8%). To study the effect on wages and long-term strike risk, we move to a yearly sample of all S&P 1500 firms. Panel B provides descriptive statistics for the yearly sample that covers all S&P 1500 firms.

Table 2.1: Descriptive statistics: quarterly and yearly sample

This table presents summary statistics of the quarterly sample used for the analysis of the short-term effect on cash holdings and the yearly sample used for the analysis of the long-term effect on labor relations. Panel A presents summary statistics on the quarterly sample which consists of 11,505 firm-quarters and 281 individual firms (185 treated and 96 control firms). Panel B presents summary statistics on the yearly sample consisting of all S&P 1500 firms. Reported are the number of observations (Obs), mean value (Mean), standard deviation (SD), 25% percentile (p25), median (p50), 75% percentile (p75). A detailed description of all variables can be found in Appendix B.1.

	Obs	Mean	SD	p25	p50	p75
Panel A: Quarterly sample (short-term analysis)						
strike observation	11,505	0.147	0.354	0.000	0.000	0.000
cash	11,505	0.073	0.089	0.014	0.038	0.100
real size	11,505	8.318	1.545	7.211	8.286	9.439
market leverage	11,282	0.246	0.187	0.099	0.209	0.367
roa	11,497	0.029	0.022	0.016	0.026	0.039
Panel B: Yearly sample (long-term analysis)						
cash	35,319	0.127	0.156	0.019	0.062	0.178
real size	35,319	7.437	1.622	6.297	7.344	8.534
market leverage	34,675	0.223	0.199	0.049	0.178	0.352
roa	34,876	0.101	0.091	0.059	0.097	0.145
no. employees	34,809	2.052	1.260	1.065	1.872	2.868
ind. unionization	34,049	0.152	0.143	0.040	0.102	0.229
strike-hit firm	35,319	0.185	0.388	0.000	0.000	0.000
strike dummy	35,319	0.019	0.135	0.000	0.000	0.000
right-to-work law	35,319	0.293	0.455	0.000	0.000	1.000
settlement dummy	35,319	0.046	0.209	0.000	0.000	0.000
#settlements	1,622	1.753	1.514	1.000	1.000	2.000
Δwage	1,471	3.021	2.702	1.900	3.000	3.714
workers under settlement	1,495	4203	15841	390	1000	2500
contract duration	1,607	1280	428	1095	1096	1461

2.3.8 Strike characteristics

Descriptive statistics of the 790 strikes between 1984 and 2011 are presented in Panel A of Table 2.2. The strike duration has a mean of 51 days and a median of 24 days. The mean (median) number of striking employees relative to the total workforce is 6.2% (1.8%). Panel B provides summary statistics on the 215 strikes that are observed by CEOs as directors at other firms. Our baseline measure for the severity of an observed labor strike is its duration. The mean strike duration of the observed strikes is slightly longer at 65 days, but the median value is very similar to all strikes (29 days). Furthermore, the mean observed strike involves a lower percentage of striking employees (5.1%); but again, this difference disappears for median values (1.7%).

Table 2.2: Descriptive statistics: strikes

This table presents summary statistics of all strikes between 1984 and 2011 in Panel A and strikes that are observed by a CEO as director at another firm in Panel B. Reported are the number of observations (Obs), mean value (Mean), standard deviation (SD), 25% percentile (p25), median (p50), 75% percentile (p75). A detailed description of all variables can be found in Appendix B.1.

	Obs	Mean	SD	p25	p50	p75
Panel A: All strikes						
strike duration	790	51.157	78.926	10.000	24.000	55.000
striking employees	790	2187.002	7024.992	215.000	409.500	1225.000
striking emp/total emp	790	0.062	0.119	0.007	0.018	0.051
idled employee-days	790	91935	316495	3910	11700	39000
$\Delta roa_{t0,t-4}$	623	-0.004	0.015	-0.009	-0.001	0.003
$\Delta ocf_{t0,t-4}$	485	-0.005	0.021	-0.016	-0.004	0.006
Panel B: Strikes observed by CEOs						
strike duration	215	65.512	110.145	14.000	29.000	55.000
striking employees	215	2809.239	9524.629	250.000	472.000	1200.000
striking emp/total emp	215	0.051	0.092	0.006	0.017	0.048
idled employee-days	215	196068	773086	6000	14400	49800
$\Delta roa_{t0,t-4}$	188	-0.004	0.014	-0.009	-0.001	0.003
$\Delta ocf_{t0,t-4}$	159	-0.004	0.020	-0.015	-0.004	0.005

Figure 2.2 illustrates the strike statistics. Subfigure (a) presents the development of the strike frequency and strike statistics over time. The mean number of strikes per year is 28. The general trend over our sample period is a falling number of labor strikes. Whereas we observe 546 strikes in the first half of our sample (1984-1997), 244 strikes occur in the second half (1998-2011). Despite the overall trend, strike risk varies considerably from year to year. Subfigures (b), (c), and (d) explore the intensity of strikes over time. They present the number of striking employees relative to the total workforce, the strike duration

in days, and the total sum of idled employee days, respectively. Graphically, there is no incidence of a decrease in strike intensity over our sample period. In contrast, it seems to be the case that strikes have been gaining intensity since the mid-1990s.

Subfigures (e), (f), and (g) illustrate the idiosyncratic character of the strike events. Subfigure (e) presents the distribution of the strikes across industries and time. Although strikes are concentrated in the manufacturing industry, firms from all industries are hit by strikes. In relative terms, manufacturing firms account for 46% of the strike events. If we use a more granular industry definition in Subfigure (f) to explore the distribution of strikes within the manufacturing industry, it turns out that most sub-industries are affected by strikes. Thus, strikes are not concentrated in any particular industry; they are common in all manufacturing firms. Subfigure (g) shows the geographical distribution of the strikes. It presents the number of strikes per state and 10 firms over the 1984-2011 period.[24] In the median state, 10 firms are hit by 7 strikes over our sample period of 28 years. Except for South Dakota, we observe at least one strike in each state. Overall, there seems to be no evidence for a geographic clustering of strikes in particular states.

[24]In detail, the score is calculated as the total number of strikes per state, divided by the median number of sample firms headquartered in that state, multiplied by 10. The geographic location of a strike is observed on the establishment level. Strikes that affect establishments in multiple states are not considered.

Figure 2.2: Strikes between 1984 and 2011
This figure illustrates the basic statistics of strikes. Subfigures (a) to (d) show the development of the strike frequency and several strike characteristics over time. Subfigure (e) shows the strike frequency in Fama-French 12 industries and over time. Subfigure (f) shows the strike frequency of manufacturing firms in Fama-French 48 industries and over time. Subfigure (g) presents the number of strikes in a state per 10 firms between 1984 and 2011. The score is calculated as the total number of strikes that occurred in a state divided by the median number of sample firms headquartered in that state, multiplied by 10. The geographic location of strike events is observed on the establishment level. Strikes that involve establishments from multiple states are excluded. A detailed description of all variables can be found in Appendix B.1.

(a) Strike frequency

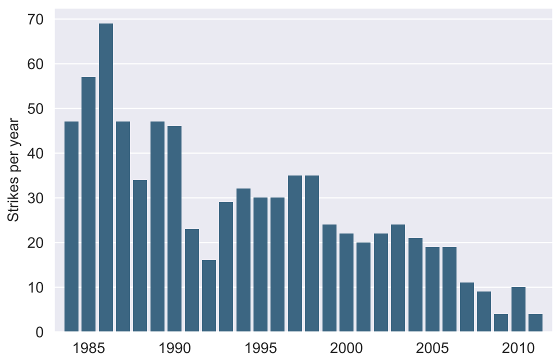

(b) Striking employees/total employees

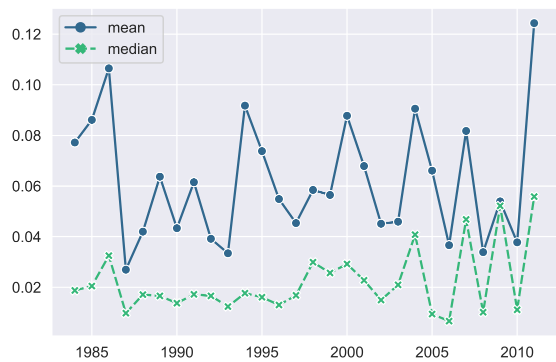

(c) Strike duration

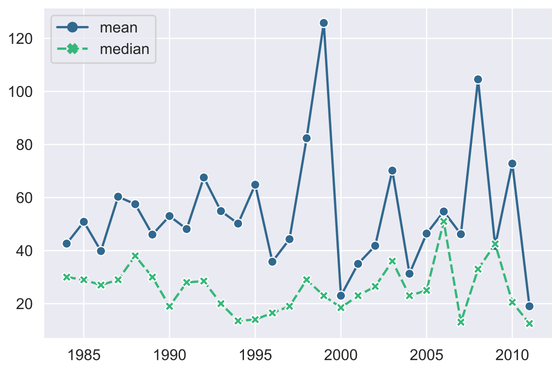

(d) Idled employee-days

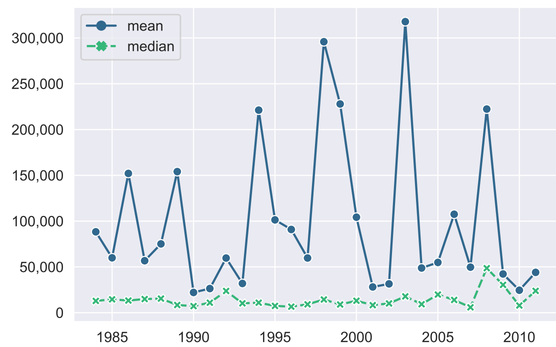

continued on next page

(e) Strike frequency in Fama-French 12 industries

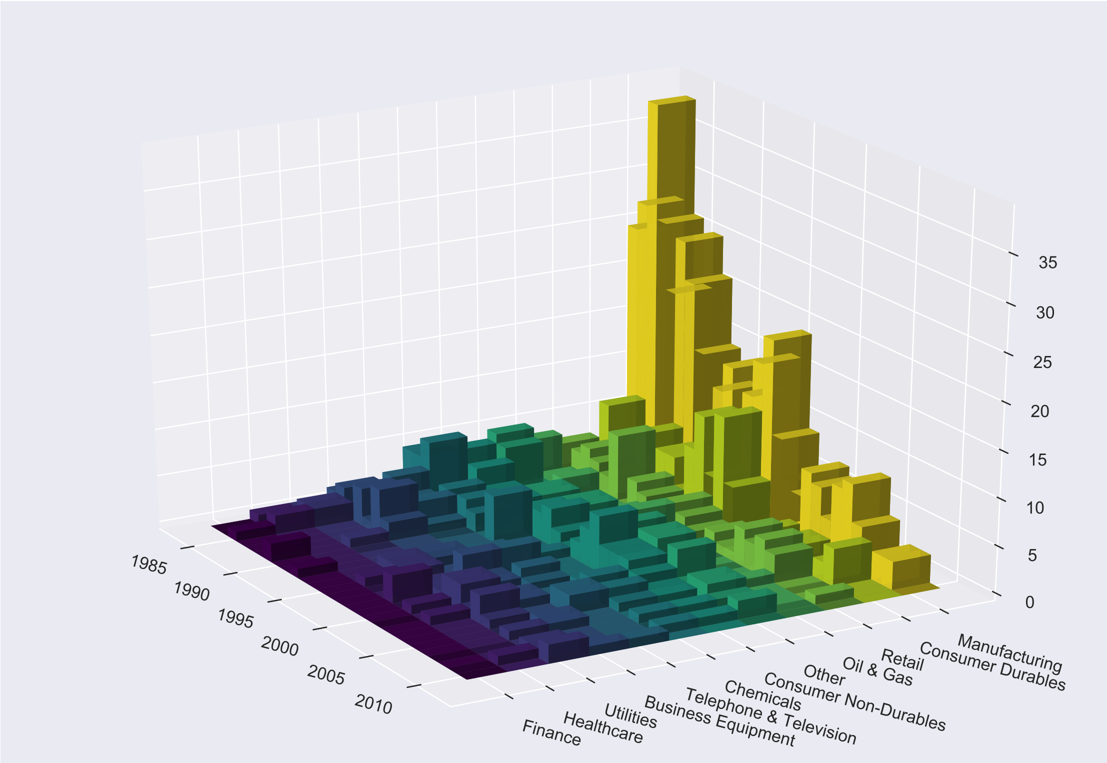

(f) Strike frequency of manufacturing firms in Fama-French 48 industries

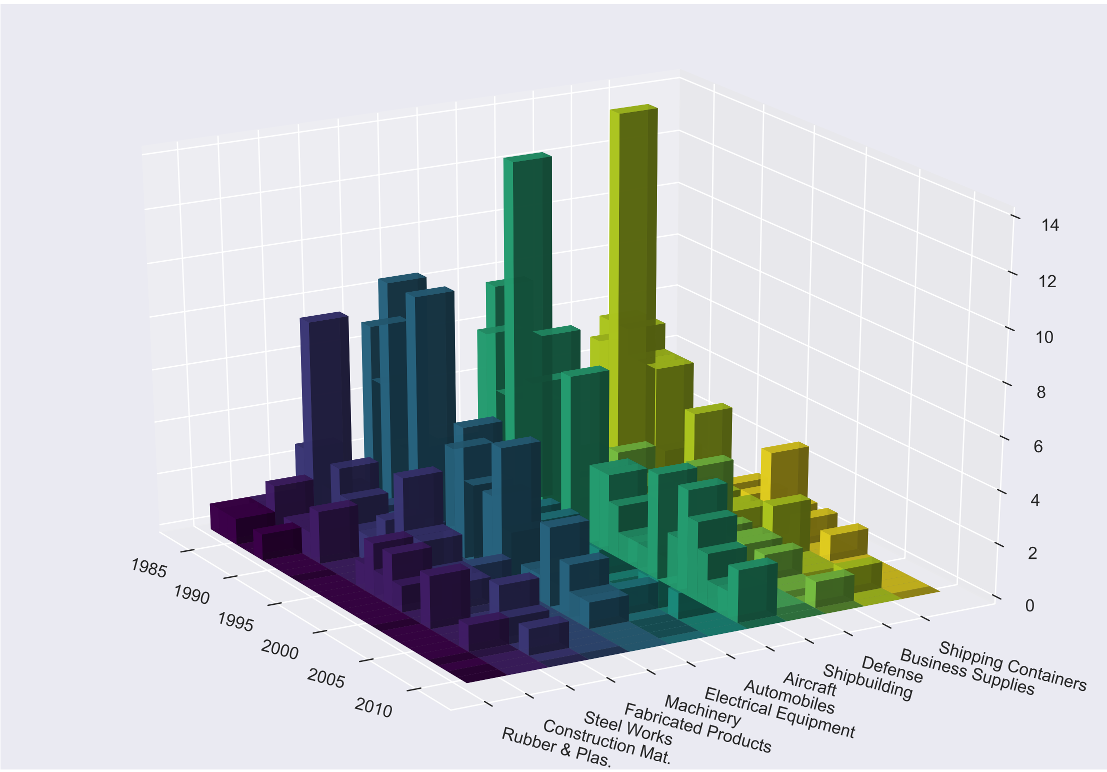

continued on next page

(g) Number of strikes per state and 10 firms

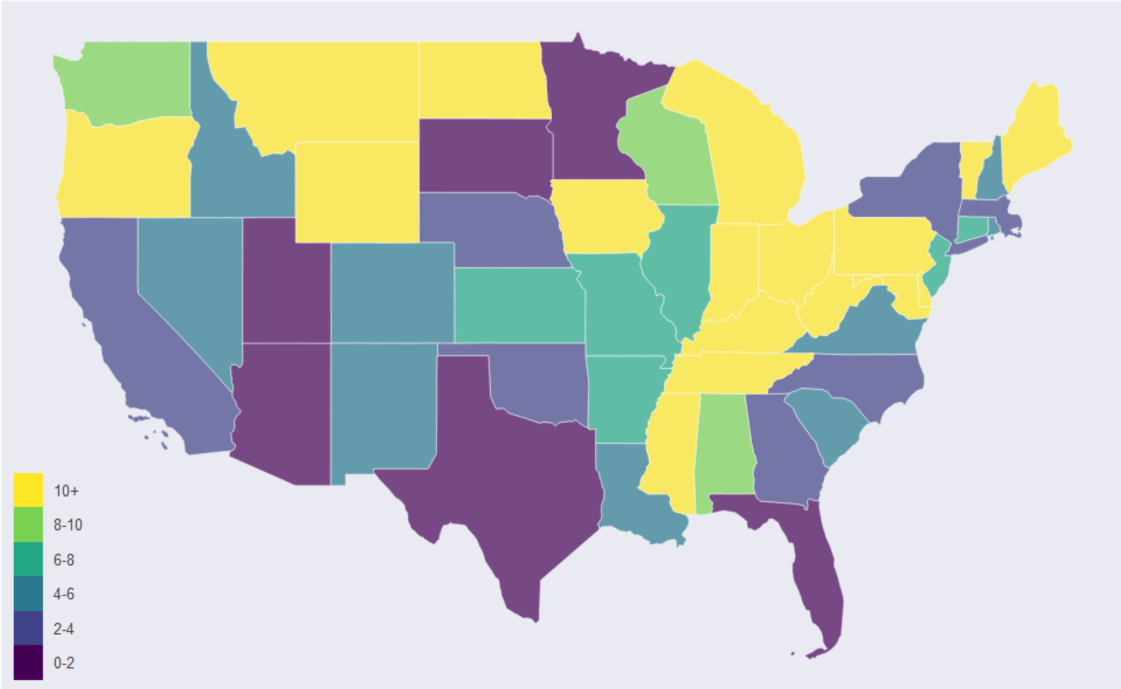

2.4 Salience of labor risk: short-term effects on cash holdings

Salience theory predicts a temporary increase in precautionary behavior after the CEO's attention is directed towards labor risk (see Section 2.2.1 for more details). We investigate short-term adjustments made to quarterly cash holdings after a CEO's strike observation.

2.4.1 Baseline results

In Table 2.3, Column 1, we estimate the baseline model of Equation 1 without firm-level control variables. This model investigates how firms whose CEO observes a labor strike change their cash holdings relative to the control group. The firm-fixed effects ensure that we compare cash changes within firms over time. Furthermore, the year-quarter fixed effects control for general time patterns, the industry-year fixed effects ensure that we essentially compare treated firms to control firms from the same industry and in the same year, and the industry-quarters fixed effects adjust for any industry-specific seasonality. The coefficient estimate is 0.0066 (t-stat of 2.04). Controlling for lagged real size, lagged market leverage, and lagged return on assets in Column 2 slightly increases the coefficient estimate to 0.0070. These estimates correspond to a relative cash increase of about 10%.

Table 2.3: Strike observation and cash holdings

This table presents estimates from regressions of cash holdings on the strike observation dummy and the severity of the observed strike. The strike observation dummy is set to one in the quarter of the CEO strike observation and the seven quarters thereafter. The logarithm of the observed strike's duration is used to measure its severity. For treated firms, we include the eight quarters before the CEO strike observation (q-8 to q-1) as the pre-period and the eight quarters afterwards as post-period (q0 to q7). All regressions include firm, fiscal-quarter-year, industry-year, and industry-quarter fixed effects. In Columns 2 and 4, we also include the lagged real size, the lagged market leverage and the lagged return on assets. T-statistics based on Huber/White robust standard errors clustered by firms are presented in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix B.1.

	(1)	(2)	(3)	(4)
strike observation	0.0066**	0.0070**		
	(2.04)	(2.25)		
strike severity$_{duration}$			0.0022**	0.0024***
			(2.51)	(2.88)
lag of real size		-0.023***		-0.023***
		(-3.73)		(-3.73)
lag of market leverage		-0.070***		-0.070***
		(-2.62)		(-2.62)
lag of return on assets		0.31		0.31
		(1.58)		(1.58)
Firm FE	Yes	Yes	Yes	Yes
Year-quarter FE	Yes	Yes	Yes	Yes
Industry-year FE	Yes	Yes	Yes	Yes
Industry-quarter FE	Yes	Yes	Yes	Yes
Obs	11,504	11,162	11,504	11,162
Number - firms	281	280	281	280
Number - events	215	213	215	213
R2	0.633	0.660	0.633	0.660

In Column 3, we replace the strike observation dummy with the logarithm of the strike's duration. This specification allows for heterogeneous treatment effects, depending on the severity of the observed strike. The coefficient estimate for the strike duration is 0.0022 (t-value of 2.51). This estimate implies that CEOs who observe a mild strike (10th percentile) increase cash holdings by 0.0039, whereas the increase is 0.011 for CEOs observing a severe strike (90th percentile). In Column 4, we also include firm-level control variables, which leads to a marginally higher coefficient estimate of 0.0024 (t-value of 2.88).

2.4.2 Time dynamics

Next, we analyze the time dynamics of the cash changes around the CEO strike observations. Salience theory predicts that the overestimation of labor risk occurs shortly after the strike observation and decreases over time. Figure 2.3 graphically illustrates the development of cash holdings from four quarters before to seven quarters after the strike observation. Subfigure (a) presents the development of mean cash holdings of firms whose CEOs observe a strike. Subfigure (b) presents the coefficient estimates from a panel regression of cash holdings on time dummies of the strike observation from four quarters before to seven quarters after the strike observation using our main sample. The model specification follows Column 1 of Table 2.3.

Cash holdings show no clear time pattern before the strike observation. Consistent with precautionary behavior after the strike observation, CEOs start to increase cash holdings in the quarter of their strike observation (q0). Cash holdings reach a peak after four quarters (q4), which indicates that labor risk is more salient for CEOs within the first year after their strike observation, and they continue to build up liquidity during this time period. In the following quarters (q5 to q7), cash holdings decline and approach their original level. These time dynamics of cash changes around CEO strike observations are in line with a temporary overestimation of labor risk due to salience. This conclusion is also supported by the coefficient estimates in Appendix B.2.

Figure 2.3: Strike observation and cash holdings: time dynamics
This figure illustrates how cash holdings change around a strike observation, from four quarters before (q-4) to seven quarters after the strike observation by the CEO (q+7). Subfigure (a) presents the development of the mean cash holdings of the firms whose CEO observes a strike. In (q-4), the mean cash holdings is set to one. Subfigure (b) presents the coefficient estimates from a panel regression of cash holdings on time dummies of the strike observation. The base period is from eight to five quarters before the strike observation. The model specification follows Column 1 of Table B.2. Error bars show the 90% confidence bounds. The period from the quarter of the strike observation (q0) to seven quarters thereafter (q+7) is highlighted in red. A detailed description of all variables can be found in Appendix B.1.

(a) Mean cash holdings

(b) Coefficient plot

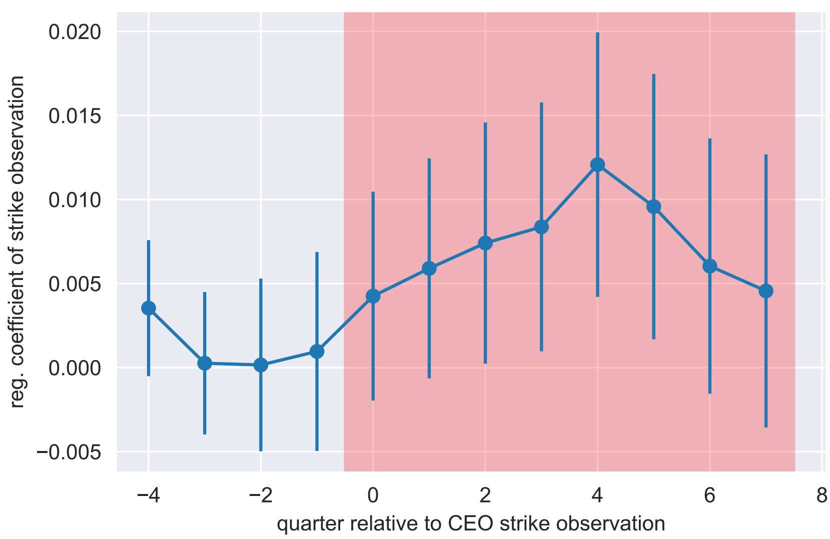

2.4.3 Robustness

One potential concern is that actual strikes affect the strike risk at other firms. In this case, our results may not be related to the salience of labor risk and its effect on the decision-making of strike-observing CEOs, but to a change in strike risk at the CEO firm. We show how strikes affect strike risk in Table 2.4. We present the results from the logit models in Columns 1 and 2. In Columns 3 and 4, we add firm-fixed effects and estimate conditional logit models. Across all models, only strikes in the same industry and firms' number of employees have a positive impact on strike risk. These two factors may be explained by firms sharing the same union within an industry and by unions targeting larger firms as a signal of strength. Most importantly for our purpose, we find that strikes in other industries do not predict strike risk (t-values of below 0.4 in all models). Thus, changes in strike risk cannot explain our previous results, which are based on CEOs who observe a strike in a firm from a different industry.[25] We also find that previous strikes at the firm, the industry-level unionization, and right-to-work laws affect strike risk in the logit models, but their impact is insignificant once we add firm-fixed effects.

For the next robustness test, we use alternative measures of strike severity. Table 2.5 presents the results. The duration of a labor strike is our baseline measure for the severity of a labor strike. However, there are other characteristics that can make a strike harmful for a company. The alternative measures we use are the ratio of striking employees to total employees (Column 1), the logarithm of the number of idled-employee days (Column 2), the absolute change of return on assets from the first quarter of the strike to the same quarter one year before (Column 3; if the change in the return on assets is not negative, it is set to zero), and the absolute change of return on assets from the first quarter of the strike

[25] As explained in Section 2.2.3, we always exclude events for which the CEO and the director firm are from the same industry or in a supplier-customer relationship.

Table 2.4: Strike risk

This table presents estimates from regressions of a strike dummy on industry, state, and firm characteristics. The dependent variable is a dummy that equals one if a labor strike begins at the firm in the respective fiscal year and zero otherwise. All models include year fixed effects. The estimated model is a logit model in Columns 1 and 2 and a conditional logit model in Columns 3 and 4. The sample is constructed on the firm-year level and ranges from 1985 to 2011. In Columns 1 and 2, the sample contains the firm-years of all S&P 1500 firms. In Columns 3 and 4, the sample consists of firms that are hit by at least one strike during the sample period. T-statistics based on Huber/White robust standard errors clustered by firms are presented in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix B.1.

	(1)	(2)	(3)	(4)
lag of strike dummy	1.63***	0.98***	-0.28*	-0.34**
	(9.27)	(5.75)	(-1.92)	(-2.41)
lag of ind. #strikes	0.59***	0.55***	0.39***	0.33***
	(5.95)	(5.74)	(3.03)	(2.67)
lag of other ind. #strikes	0.12	-0.085	0.10	0.100
	(0.39)	(-0.27)	(0.28)	(0.28)
lag of #strikes in state	0.0013	0.038	0.093	0.080
	(0.013)	(0.38)	(0.83)	(0.73)
lag of ind. unionization	3.07***	2.69***	-0.34	-0.18
	(9.11)	(7.90)	(-0.45)	(-0.24)
right-to-work law	-0.40**	-0.35**	0.66	0.46
	(-2.35)	(-2.07)	(1.45)	(1.09)
lag of no. employees		0.68***		0.71***
		(16.1)		(4.48)
lag of roa		-1.35*		0.74
		(-1.89)		(0.65)
Firm FE	No	No	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Obs	29,309	29,309	5,216	5,216
Number - firms	2407	2407	270	270
R2	0.15	0.21	0.049	0.060

to the same quarter one year before (Column 4; if the change in the operating cash flow is not negative, it is set to zero). For all four alternative measures, we obtain coefficient estimates that are positive and statistically significant. Hence, in line with our baseline measure, the cash increase is more pronounced for observed strikes that are more severe.

Table 2.5: Robustness: alternative measures for severity of the observed strikes

This table presents estimates from regressions of cash holdings on alternative measures for the severity of the strike that a CEO observes. As alternative measures, we use the ratio of striking employees to total employees (Column 1), the logarithm of idled employee-days (Column 2), the absolute value of the decline in return on assets in the first quarter of the strike compared to the same of quarter of the last year (Column 3; if the change of return on assets is positive it is set to zero), and the absolute value of the decline in operating cash flow in the first quarter of the strike compared to the same of quarter of the last year (Column 4; if the change of operating cash flow is positive it is set to zero). All models include firm, fiscal-quarter-year, industry-year, and industry-quarter fixed effects, as well as the lagged real size, the lagged market leverage, and the lagged return on assets. T-statistics based on Huber/White robust standard errors clustered by firms are presented in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix B.1.

	(1)	(2)	(3)	(4)
strike severity$_{\%idled}$	0.086**			
	(2.19)			
strike severity$_{idled-days}$		0.00082**		
		(2.57)		
strike severity$_{\Delta roa}$			0.90***	
			(3.54)	
strike severity$_{\Delta ocf}$				0.62***
				(3.27)
Controls	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Year-quarter FE	Yes	Yes	Yes	Yes
Industry-year FE	Yes	Yes	Yes	Yes
Industry-quarter FE	Yes	Yes	Yes	Yes
Obs	11,162	11,162	10,947	10,730
Number - firms	280	280	280	280
Number - events	213	213	186	159
R2	0.660	0.660	0.660	0.656

Next, we use only treated firms, exclude events if the CEO firm and the strike-hit firm are located in the same state, and add state-year fixed effects to the baseline models. Table 2.6 presents the results. In Panel A, we run our main analysis using only firms whose CEOs observe a labor strike during our sample period (treated firms). In this specification, quarters in which CEOs do not observe strikes are used as control observations. The coefficient estimates are similar albeit slightly lower than our baseline results. In Panel B, we remove 31 events for which the headquarters of the CEO firm and the strike-hit firm are in the same state. We find no evidence that strikes in the same state have a predictive power for strike risk; however strikes at firms in the same region could have spillover effects that are unrelated to the board channel. The coefficient estimates for both the strike observation dummy and the strike duration are virtually unchanged in this specification. In Panel C, we add state-year fixed effects to our baseline specification. All models indicate that a strike observation leads to higher cash holdings, and the coefficient estimates are even higher than in our baseline specification.

Table 2.6: Robustness: alternative sample, same-state events, and state-year fixed effects

This table presents estimates from regressions of cash holdings on the strike observation dummy and the severity of the observed strike. In Panel A, we only consider firms with a CEO strike observation. In Panel B, we exclude all events for which the CEO and director firm are in the same state. In Panel C, we add state-year fixed effects. T-statistics based on Huber/White robust standard errors clustered by firms are presented in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix B.1.

	(1)	(2)	(3)	(4)
Panel A: Only treated firms				
strike observation	0.0052*	0.0057**		
	(1.95)	(2.08)		
strike severity$_{duration}$			0.0019**	0.0020**
			(2.37)	(2.47)
Controls	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes
Year-quarter FE	Yes	Yes	Yes	Yes
Industry-year FE	Yes	Yes	Yes	Yes
Industry-quarter FE	Yes	Yes	Yes	Yes
Obs	3,345	3,278	3,345	3,278
Number - firms	185	184	185	184
Number - events	215	213	215	213
R2	0.862	0.869	0.862	0.869
Panel B: Exclusion of events in same state				
strike observation	0.0074**	0.0079**		
	(2.12)	(2.36)		
strike severity$_{duration}$			0.0023**	0.0024***
			(2.44)	(2.81)
Controls	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes
Year-quarter FE	Yes	Yes	Yes	Yes
Industry-year FE	Yes	Yes	Yes	Yes
Industry-quarter FE	Yes	Yes	Yes	Yes
Obs	11,260	10,925	11,260	10,925
Number - firms	281	279	281	279
Number - events	184	183	184	183
R2	0.629	0.657	0.629	0.657

Continued on next page

Table 2.6 continued

	(1)	(2)	(3)	(4)
Panel C: Additional state-year fixed effects				
strike observation	0.0099***	0.0088**		
	(2.75)	(2.45)		
strike $severity_{duration}$			0.0032***	0.0029***
			(3.26)	(2.92)
Controls	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes
Year-quarter FE	Yes	Yes	Yes	Yes
Industry-year FE	Yes	Yes	Yes	Yes
Industry-quarter FE	Yes	Yes	Yes	Yes
State-year FE	Yes	Yes	Yes	Yes
Obs	11,059	10,726	11,059	10,726
Number - firms	272	271	272	271
Number - events	209	207	209	207
R2	0.749	0.767	0.749	0.767

2.5 Observational learning: long-term effects on labor relations

Observational learning predicts that CEOs learn from a strike observation and adjust their own labor relation management (see Section 2.2.2 for more details). We investigate long-term changes to employee wages in labor contract negotiations and strike risk.

2.5.1 Labor contract negotiations and employee wages

We first focus on wage changes from collective bargaining agreements. Due to the limited number of contract settlements per quarter, we use yearly instead of quarterly data for this analysis. Furthermore, we expect that effects from observational learning would occur more in the long term. Because of that, and also due to the fact that the average labor contract duration is about 3.5 years in our sample, we examine the effect on wage changes over horizons of two, four, six, and eight years after the CEO strike observation. We set the strike observation dummy to one from the year of the strike observations until the end of the respective time period. We compare the change in wages of firms whose CEO observed a strike to control-group firms in the same industry and year.

The dependent variable for this analysis is the wage change in the first year of the collective bargaining agreement. If a firm settles more than one labor contract per year, we calculate the mean wage change. The variable of interest is the duration of the observed labor strike as a measure for its severity. All models include firm, year, and industry-year fixed effects, as well as the lagged number of employees, the lagged market leverage, and the lagged return on assets. The sample consists of firms for which we observe at least one labor contract settlement.

Table 2.7 presents the results for strike severity, that is, the continuous treatment effect. Column 1 shows the results for the two-year horizon. The coefficient estimate on the strike duration is 0.28 (t-stat of 2.16). This coefficient estimate implies that CEOs who observe strikes with a median duration increase wages by about one percentage point more than CEOs without strike observation. Since the average wage increase is approximately three percent, this additional wage increase is economically meaningful. Over the time horizon of four years, the coefficient estimate is very similar at 0.28. For the even longer periods of six and eight years, we find a positive and statistically significant effect with slightly lower coefficient estimates of 0.23 and 0.19, respectively. The qualitatively similar results for the strike observation dummy are shown in Appendix B.3. Overall, this test provides evidence that CEOs who observed a labor strike agree to higher wages during contract negotiations at their own firms.

Table 2.7: Strike observation and the long-term effect on wage changes

This table presents estimates from regressions of the wage change in the first year of new labor contracts (measured in percentage points) on the severity of the observed strike. We consider a time period of two years (Column 1), four years (Column 2), six years (Column 3), and eight years (Column 4) after the strike observation. All models include firm, year and industry-year fixed effects. The sample is constructed on the firm-year level, and includes firms for which we observe the settlement of at least one labor contract. T-statistics based on Huber/White robust standard errors clustered by firms are presented in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix B.1.

	(1)	(2)	(3)	(4)
	2 years	4 years	6 years	8 years
strike severity$_{duration,longterm}$	0.28**	0.28**	0.23**	0.19*
	(2.16)	(2.27)	(2.01)	(1.84)
lag of no. employees	0.75	0.55	0.40	0.30
	(1.41)	(1.11)	(0.91)	(0.74)
lag of market leverage	-0.96	-1.07	-0.73	-0.45
	(-0.67)	(-0.80)	(-0.58)	(-0.40)
lag of roa	-3.59	-2.82	-1.77	-1.20
	(-0.86)	(-0.72)	(-0.46)	(-0.35)
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Industry-year FE	Yes	Yes	Yes	Yes
Obs	919	968	1,013	1,072
Number - firms	216	229	234	241
Number - events	28	36	39	40
R2	0.612	0.607	0.597	0.589

2.5.2 Strike risk

Next, we analyze whether observing strikes at another firm reduces the risk of failed labor negotiations at the CEO's firm. If CEOs collect valuable information on the bargaining process of firms and unions, we also expect them to implement long-term measures to reduce the strike risk at their firms. Like the wage analysis, we test the effect on strike risk over time horizons of two, four, six, and eight years after the CEO strike observation. To compare the strike risk of a firm before and after the CEO observes a strike, we estimate conditional logit regressions.

Table 2.8 presents the results. The dependent variable is a strike dummy indicating the start of a labor strike in the respective firm year. All models are conditional logit models with firm and year fixed effects. The sample consists of all firms that had at least one labor strike because the conditional logit model excludes all firms for which the outcome variable does not vary at all. Hence, the total number of yearly observations (CEO strike observations) is 4,134 (70) for the shortest time horizon of two years and up to 4,568 (77) for the longest time horizon of eight years. The variable of interest is strike duration as a measure for the severity of the observed strike. Column 1 presents the results for two years after the strike observation. The coefficient estimate is negative (-0.048) but not statistically significant. For the four years after the CEO strike observation, the coefficient estimate is -0.095 (t-stat of -1.37). Over the six- and the eight-year periods, the coefficient estimates were -0.13 and -0.14, respectively, and they were both statistically significant at the 5% level. These coefficients can be interpreted as follows: On average, in the six (eight) years after the CEO strike observation, the odds for being strike-hit is 0.88 (0.87) times the odds before the strike observation.[26]

[26] Appendix B.4 presents the analysis of strike risk using the CEO strike observation dummy. As for strike severity, the coefficient estimate becomes more negative for longer time horizons, despite not being statistically significant. For six and eight years, the coefficient estimates are

Table 2.8: Strike observation and the long-term effect on strike risk

This table presents estimates from regressions of a strike dummy on the severity of the observed strike. We consider a time period of two years (Column 1), four years (Column 2), six years (Column 3), and eight years (Column 4) after the strike observation. The dependent variable is a strike dummy that is set to one if a labor strike begins at a firm in the respective fiscal year and zero otherwise. The sample is constructed on the firm-year level and consists of firms that are hit by at least one strike during the sample period. T-statistics based on Huber/White robust standard errors clustered by firms are presented in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix B.1.

	(1)	(2)	(3)	(4)
	2 years	4 years	6 years	8 years
strike severity$_{duration,longterm}$	-0.048	-0.095	-0.13**	-0.14**
	(-0.64)	(-1.37)	(-2.05)	(-2.24)
lag of strike dummy	-0.27*	-0.27*	-0.29**	-0.31**
	(-1.81)	(-1.88)	(-1.99)	(-2.11)
lag of ind. #strikes	0.30**	0.30**	0.30**	0.32**
	(2.19)	(2.26)	(2.33)	(2.47)
lag of other ind. #strikes	0.066	0.087	0.083	0.13
	(0.17)	(0.23)	(0.22)	(0.35)
lag of #strikes in state	0.074	0.075	0.084	0.073
	(0.59)	(0.61)	(0.71)	(0.62)
lag of ind. unionization	0.18	0.13	-0.16	0.034
	(0.19)	(0.15)	(-0.18)	(0.039)
right-to-work law	0.61	0.56	0.63	0.60
	(1.33)	(1.25)	(1.47)	(1.41)
lag of no. employees	0.67***	0.66***	0.70***	0.70***
	(3.68)	(3.72)	(4.00)	(4.08)
lag of roa	0.15	-0.11	-0.096	0.048
	(0.12)	(-0.086)	(-0.081)	(0.041)
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Obs	4,134	4,277	4,433	4,568
Number - firms	247	248	252	254
Number - events	70	72	75	77
R2	0.052	0.052	0.056	0.058

This test shows that a firm's long-term strike risk is reduced after the CEO observed a (severe) strike at another firm. It supports the view that CEOs may utilize their insider knowledge from the failed bargaining with labor at the director firm to adjust their own labor negotiations and implement measures that lead to a reduction in the long-term strike risk.

2.6 Conclusion

More than one-third of CEOs in S&P 500 firms serve as directors at other firms. One potential benefit of such outside positions is that CEOs might learn from observations they make at the director firm. On the other hand, the required time and potential distractions from their core tasks are potential reasons against CEO outside directorships. This paper contributes to this discussion by providing direct empirical evidence on the effects of outside directorships on observational learning and distractions due to behavioral biases. For this purpose, we analyze how CEOs' precautionary behavior and labor relation management changes after they observed a labor strike as a director of another firm. We focus on labor strikes for three reasons. First, they are exogenous to the CEO firm. We show that strikes do not affect strike risk at other firms once we exclude events from the same industry. Second, labor disputes are very costly for firms and require substantial attention from the board. Third, in contrast to many other types of shocks, such as the sudden death of a CEO or weather-related losses, strikes are not exogenous for the strike-hit firm, enabling CEOs to influence the strike risk at their firms.

We identify 215 events of CEOs who observe a strike as a director at another firm. We first document that CEOs increase cash holdings shortly after the strike. The cash increase starts in the quarter of the experience, reaches a maximum four quarters thereafter, and then reverts. These results are consistent

-0.35 (t-stat of -1.44) and -0.38 (t-stat of -1.63), respectively.

with predictions from salience theory: CEOs temporarily overestimate strike risk when their attention is directed to this risk factor. This precautionary behavior is most likely a behavioral bias that is not in the best interest of the firm. These results imply that behavioral biases due to overreaction to salient risks can be a "dark side" of CEO outside directorships (besides time and efforts). However, there is also a "bright side" of outside directorships, as we find that CEOs adjust their bargaining with labor during labor contract negotiations, which is consistent with observational learning. We show that CEOs tend to agree to higher wages in the years after their strike observation and manage to decrease the strike risk for their firm in the long run. These findings indicate that strike-observing CEOs can learn from their insider information about negotiations with labor.

3

Firm Size, Workforce Composition, and Wage Inequality

Authors: Daniel Bias, Chen Lin, Benjamin Lochner, Thomas Schmid

First author: Daniel Bias

Keywords: Wage inequality, workforce composition, rent-sharing, efficiency wages

JEL codes: L25, J31, G30

Presentations: 23rd Theories & Methods in Macroeconomics Conference (Nuremberg)
EEA-ESEM 2019 (Manchester)
3rd Dale T. Mortensen Conference 2019 (Aarhus)
Seminar in Macroeconomics University of Konstanz 2019 (Konstanz)
Perspectives on (Un-)Employment Ph.D. Workshop 2020 (Nuremberg)

Current status: Working paper

Abstract:

Using a matched employee-establishment-firm dataset covering German workers, we find that wage inequality increases monotonously with firm size. Firms in the top decile pay 70% higher wages than the firms in the bottom decile, and their within-firm wage variance is 30% larger. Decomposing wage inequality into workforce composition and non-composition effects reveals that composition is responsible for 65% of the size-inequality relationship within firms and for 60% between firms. Higher wage variance in larger firms is largely explainable by more heterogeneous job characteristics and higher employee monitoring complexity. Higher wages in larger firms are not related to differences in profitability, monitoring complexity, or unionization levels, but different job characteristics and local labor markets play some role. Analyzing establishment size within firms reveals that larger establishments show more variation in workforce quality but do not pay an economically significant wage premium.

3.1 Introduction

A recent study by the International Labour Organization found that the bottom half of workers were paid only 6.4% of global wages. This wage inequality, which increased sharply during the last decades (e.g., Katz and Autor, 1999), is today a key political issue. President Barack Obama (2013) even described inequality as the "the defining issue of our time". Only recently has the literature begun focusing on the role of firms, which ultimately set wages, for wage inequality in the economy (e.g., Mueller, Ouimet and Simintzi, 2017a; Bloom et al., 2018; Song et al., 2019).

In this paper, we analyze the role of firm size for wage inequality. Wage inequality can arise either because wages within larger firms are more dispersed ("within-firm inequality") or because larger firms pay their workers higher average wages ("between-firm inequality"). For our analysis, we use a linked employee-establishment-firm dataset from Germany that covers approximately 51,000 individual firms, 115,000 establishments, and 11.6 million workers.[1] The uniqueness of this dataset comes from the fact that it links administrative data about individual workers with firm-level information, such as accounting figures.

Comparing the largest firm decile in terms of total assets to the bottom decile, we find that the largest firms have a 30% higher within-firm wage variance and pay 70% higher wages.[2] However, this positive effect of firm size is not only present when comparing the top and bottom of the size distribution but also increases monotonously with firm size. We use total assets as our main

[1]Worker-level information is based on administrative data that originates from the German social security system. The worker-level data and the employee-establishment-firm link is provided by the Institute for Employment Research (Institut für Arbeitsmarkt- und Berufsforschung, IAB). Firm-level data is obtained from historical versions of Bureau van Dijk's Amadeus database.

[2]Firms in the largest decile have on average approximately EUR 150 million in total assets and 500 employees, while those in the bottom decile have approximately EUR 0.6 million in total assets and 50 employees.

size proxy because it is most common in the financial economics literature and captures all resources available to the firm. However, for comparison, we also follow the majority of the inequality literature and use employee-based size proxies. Despite similar results for within-firm inequality, employee-based size measures tend to underestimate the effect of size on between-firm inequality.

Next, we graphically explore the development of average firm size and wage inequality over time. We find that firm size increased substantially between 1995 and 2007 (i.e., the year before the financial crisis). This increase is not only observable for the largest firms but also for smaller firms. For instance, the average size of both the largest 50 firms in Germany and firms ranked between 501 and 1000 in terms of size increased by about 70%. In the same time period, wage inequality approximately doubled, with a stronger increase in between-firm inequality compared to within-firm inequality.[3] Therefore, although this evidence is suggestive, it is consistent with the view that increases in average firm size contributed to the increasing wage inequality in the last decades.

To better understand why wage inequality is more pronounced in larger firms, we decompose between- and within-firm wage inequality into workforce composition and non-composition effects. Composition effects capture differences in workers' wages between small and large firms that exist due to differences in workforce quality. Non-composition effects are related to differences in wages for workers of the same quality. For the decomposition, we use parameter estimates of the model introduced by Abowd, Kramarz and Margolis (1999) (henceforth AKM) following the implementation of Card, Heining and Kline (2013) (henceforth CHK).[4] Generally speaking, this approach exploits movers between establishments to disentangle the overall wage into different components

[3]These patterns changed after the financial crises, with a decrease in firm size across the whole size distribution and a decreasing trend for wage inequality.

[4]The estimation is carried out on the universe of full-time jobs held by workers aged 20-60 years from 2010 to 2017. The approach follows CHK with some modifications (e.g., it also include female workers; see Sections 3.2.1 and 3.3.3 for more details).

(e.g., person and establishment fixed effects). We then use these AKM components to decompose the mean firm wage and the within-firm wage variance into workforce composition and non-composition effects. We find that composition effects play an important role in higher wage inequality in larger firms; the fact that larger firms employ, on average, a more heterogeneous and higher-quality workforce explains approximately 65% of the within-firm inequality and 60% of the between-firm inequality.

Next, we try to explain why there is a positive relationship between firm size and the different components of wage inequality. For composition effects, we hypothesize that larger firms have a higher-quality workforce and greater heterogeneity because their job characteristics differ from those of smaller firms. It is often argued that larger firms have a greater capital intensity and that capital-skill complementaries exist (Hamermesh, 1980; Oi, 1983; Schmidt and Zimmermann, 1991; Lallemand, Plasman and Rycx, 2005). Thus, larger firms can make better use of employees' skills and are more likely to employ high-skilled workers. Furthermore, the outsourcing of jobs that are not related to firms' core business and that often require workers with lower skills (such as food, cleaning, security, and logistics) became increasingly common over the past decades (Goldschmidt and Schmieder, 2017). If larger firms are more likely to engage in outsourcing, this reduces their reliance on low-skilled workers. Both effects can shift the average level of workforce quality upwards. Assuming that outsourcing is not completely possible, that larger firms also need to rely on low- and medium-skilled workers to support operations, and that tasks in larger firms are more specialized, we also expect to see greater heterogeneity of jobs in larger firms.

We find empirical support for this theoretical prediction: When we control for the fraction of (highly) complex jobs and the concentration of occupations within a firm, the impact of firm size on workforce quality dispersion becomes much weaker. This indicates that the higher dispersion of workforce quality in larger firms is strongly related to differences in job characteristics between large and small firms. Furthermore, we find that measures for job complexity explain approximately one-quarter of the between-firm composition effect. Thus, differences in job characteristics explain at least partially the higher mean workforce quality in larger firms.

For the more dispersed wages after composition effects, we hypothesize that larger firms face more severe monitoring problems, which increases employees' opportunities for shirking. In order to reduce shirking behavior, larger firms can pay higher wages as an incentive mechanism ("efficiency wage hypothesis").[5] Furthermore, more severe monitoring problems in larger firms may not only lead to higher average wages, but the wages could also be more dispersed if larger firms use more tournaments (i.e., wage differentials between hierarchy levels) and/or bonus payments (i.e., performance-based incentive payments) to incentivize workers.

We find that the higher dispersion of wages in larger firms after workforce composition effects disappears when we control for monitoring complexity proxies (i.e., managers per employee and geographical dispersion of establishments within firms). Therefore, more dispersed wages in larger firms seem to be strongly related to their higher monitoring complexity. However, we do not find empirical support for the idea that higher average wages in larger firms are related to their monitoring complexity.

Several explanations for higher average wages in larger firms after composition effects (the so-called "large-firm wage premium", LFWP) have been discussed in existing literature.[6] First, as described in the previous paragraph, the efficiency

[5]Theoretical motivation for this idea derives from Eaton and White (1983) and Shapiro and Stiglitz (1984), among others. Krueger and Summers (1988) are among the first to provide empirical evidence for the existence of noncompetitive wages.

[6]Possible explanations that are not covered in our discussion include differences between large versus small firms regarding their ability to screen workers' quality (Garen, 1985), working conditions (Brown and Medoff, 1989), governance structures (Pagano and Volpin, 2005;

wage explanation argues that larger firms face higher monitoring complexity and pay higher wages to incentivize workers. Second, larger firms may generate higher rents by exploiting their market power, and they may share some of their higher rents with employees ("rent-sharing hypothesis").[7] Third, larger firms may have higher unionization rates (enabling employees to extract a higher fraction of rents) and a higher threat of unionization (Dickens and Lang, 1985). Fourth, larger firms may cluster in regions with high economic activity and thus a high demand for labor.

We do not find any empirical support for the view that differences in monitoring complexity, profitability, or unionization levels are the main drivers behind the existence of the LFWP. Neither controlling for monitoring complexity nor profitability nor industry-level unionization rates has a material impact on the relationship between firm size and the average wage. However, when we compare small and large firms located in the same region that rely on the same labor market (i.e., those from the same industry), the LFWP is reduced by one-third or one-half, depending on the specification. Therefore, it seems that factors related to local labor markets play some role in the existence of the LFWP.

To summarize, our tests indicate that more heterogeneous job characteristics in larger firms and higher employee monitoring complexity explain their higher within-firm wage inequality, to a high degree. Between-firm inequality, that is, higher average wages in larger firms, are more challenging to explain. Differences in job characteristics can partly explain larger firms' higher-quality workforce, but a significant part of the higher average workforce quality in larger firms remains even after controlling for this factor. Local labor market factors

Cronqvist et al., 2009), ownership structures (Ellul, Pagano and Schivardi, 2018), internal labor markets (Tate and Yang, 2015), or other forms of employee participation, such as employee stock ownership plans (Kim and Ouimet, 2014).

[7]In this context, Christofides and Oswald (1992) find that real wage is positively related to past industry performance, which is consistent with the rent-sharing hypothesis. Related to this finding, Abowd and Lemieux (1993) use product market competition to show that rent-sharing considerations affect workers' wages. Card et al. (2018) provide a comprehensive discussion of the literature linking firm productivity to wages.

seem to explain some of the higher average wages in larger firms after composition effects, but a substantial part remains unexplained. While we acknowledge that our proxies (especially for monitoring) are far from being perfect, we find that efficiency wages, rent-sharing, and unionization are unlikely to be the main drivers behind this LFWP.

To better understand the role of firm size for wage inequality, we also investigate how the size of establishments within firms affects wages and their dispersion. When we compare establishments of the same firm, we find that larger establishments pay more dispersed wages. This result is mainly driven by more heterogeneous workforce quality in larger establishments. Similar to our firm-level results, it thus seems that establishment size is positively related to the heterogeneity of jobs. Although there is some evidence that larger establishments employ a higher-quality workforce, the economic significance of this result is low compared to differences between large and small firms. After controlling for composition effects, larger establishments within the same firm pay only marginally higher wages than smaller establishments; compared to differences between firms, the size effect within firms is reduced by 80%. Thus, it seems that larger firms pay higher average wages, relatively independently of the size of a particular establishment within that firm. By contrast, the size of an establishment within a (large) firm plays a very significant role in the dispersion of wages and especially for the dispersion of workforce quality.

Our results contribute to several strands of the literature. Most importantly, we add to the literature that investigates how firms affect wage inequality. Despite the recent literature highlighting the role of firms for the increase in wage inequality[8], it is still not well understood how and especially why firm size affects wage inequality.

[8]See, among others, Alvarez et al. (2018) for Brazil, Card, Heining and Kline (2013) for Germany, Card et al. (2018) for Portugal, Faggio, Salvanes and Van Reenen (2010) for the U.K., and Song et al. (2019) for the U.S.

Regarding the higher dispersion of wages in larger firms, Davis et al. (1991) were among the first to document a positive correlation between plant size and wage dispersion. Mueller, Ouimet and Simintzi (2017a) and Mueller, Ouimet and Simintzi (2017b) find that the overall within-firm wage inequality is more pronounced in larger and more profitable firms, but they do not decompose wage inequality into workforce composition and non-composition effects. We complement their analysis by showing that composition effects account for about 65% of the within-firm pay inequality. Song et al. (2019) investigate the contribution of firms to the increase in wage inequality. Although they do not focus on the role of firm size, they show that the increase of within-firm wage inequality happened mainly in mega firms with more than 10,000 employees. Our paper complements their analysis by showing that wage inequality increases monotonously with firm size (not only in mega firms) and by providing (some) explanations for the positive relationship between firm size and wage inequality.

For between-firm wage inequality, there is a large body of literature that documents higher average wages in larger firms (e.g., Brown and Medoff, 1989; Oi and Idson, 1999; Bloom et al., 2018). However, this literature often lacks firm-level data, such as profitability or establishment structures, which makes the testing of several potential explanations for higher wages in larger firms (e.g., rent-sharing) challenging. Our contribution to this literature is to test whether such firm-level factors help explain the LFWP. Somewhat surprisingly, our results do not show any evidence that rent-sharing, monitoring complexity, or unionization have a substantial impact on the LFWP.

Our results also contribute to the inequality literature by documenting the following stylized facts. First, average firm size increased between 1995 and the financial crisis not only for the largest firms but across the whole size distribution. This is in line with the recently documented rise of superstar firms (Autor et al., 2020) and the finding that large public U.S. firms tripled their average size in the last two decades (Grullon, Larkin and Michaely, 2019). Second, our firm-level data allows us to use firm capital as a size measure.[9] Capital-based and employee-based firm size measures lead to relatively similar results for within-firm wage inequality, but employee-based proxies seem to underestimate the effect of size on between-firm wage inequality. Third, between-firm wage differentials are mainly driven by firm size, not establishment size. However, within-firm wage inequality is mainly driven by establishment size, not firm size. These findings show that it is important to distinguish between establishment and firm size when analyzing the relationship between size and wage inequality.

3.2 Data

3.2.1 Employer-employee data

We obtain administrative linked employer-employee data provided by the Research Data Centre (FDZ) of the Institute for Employment Research (Institut für Arbeitsmarkt- und Berufsforschung, IAB) that belongs to the German Federal Employment Agency (Bundesagentur für Arbeit). The Integrated Employment Biographies (IEB) data originates from records of the German social security system. The data includes total earnings and days worked at each job in a year, as well as information on education, occupation, industry, and part-time or full-time status.[10] Further, the data includes parameter estimates of the AKM model.

The data preparation by the FDZ follows the steps conducted by CHK. The starting point is the universe of full-time jobs held by workers aged 20-60 from

[9]The fall of the labor share in the largest firms (Autor et al., 2020) highlights the importance of differentiating between firm size based on capital and labor.

[10]For further details on the dataset, please refer to the technical report by Antoni, Ganzer and vom Berge (2016).

2010 to 2017.[11] While CHK examine male workers in West Germany, the data cover both male and female workers in East and West Germany. Marginal employment and apprenticeship are excluded. As in CHK, the FDZ focuses on the main job held by each worker in a given year, that is the job with the highest total wage sum (including bonus payments). For all these jobs, the average daily wage is calculated by dividing the total wage sum by the total duration of the main job. One important limitation is that wages are only reported up to a time- and region-specific threshold—the contribution assessment ceiling ("Beitragsbemessungsgrenze"). The FDZ follows the procedure suggested by Dustmann, Ludsteck and Schönberg (2009) and CHK and imputes the upper tail of the wage distribution by running a series of Tobit regressions, allowing for a maximum degree of heterogeneity by fitting the model separately for gender, time, education levels, and eight five-year age groups. Missing and inconsistent information in the education variable is imputed using the methodology proposed in Fitzenberger, Osikominu and Völter (2006).

The dataset consists of over 165 million worker-establishment years, almost 32 million unique workers, and more than 2.7 million establishments.

3.2.2 Firm-establishment data

The linked employer-employee data provides only information on employees and establishments. For example, individual i is employed at establishment j. To add information on the firm structure, we use the novel ORBIS-ADIAB dataset. It provides a linking table between establishment identifiers by the Institute for Employment Research and the firm identifiers by the Bureau van Dijk (BvD). For example, establishment identifiers j_1 and j_2 belong to the firm k. Firm-level financial information from 2010 to 2016 is then derived from the Amadeus database by BvD.

[11] The administrative data originates from the social security system; therefore, the IEB data does not include employment spells of civil servants and self-employed workers.

The process of linking establishments to firms is documented in full detail by Antoni et al. (2018). The most important variables for the record linkage are the establishment and the company name, the legal form, the industry code, and the postal code. The record linkage is carried out separately for the years 2014 and 2016. For the years 2010 to 2013 and the year 2015, we then make the assumption that the latest link of an establishment to a firm is still valid. In the resulting linking table, 3.8% of establishment-years are mapped to multiple firms. This is because the record linkage approach is not able to uniquely identify one firm or the establishment undergoes an actual ownership change. We exclude the establishment-years that are not uniquely assigned to one firm.

3.2.3 Sample construction

We restrict the sample to firms with more than 20 employees in the IAB data to ensure that statistics about wage variations within firms are meaningful. We also exclude from our sample firms for which we do not obtain information on total assets and sales.[12] The final sample covers a total of 51,515 firms with 115,632 establishments and approximately 11.6 million individual workers in the 2010-2016 period.

3.2.4 Descriptive statistics

Table 3.1 provides descriptive statistics on the firm level. Unscaled financial variables are consumer price index-adjusted, and all continuous variables are winsorized at the 1st and 99th percentiles. We report firm mean values in the 2010-2016 period. The average firm employs 64 workers in Germany and 75 employees worldwide. It has mean total assets of EUR 5.4 million, sales of EUR

[12]Further, we exclude establishments for which AKM components cannot be estimated. The AKM approach identifies establishment fixed effects by workers moving between establishments and is estimated relative to an omitted reference establishment. The estimation is performed on the largest set of establishments that are connected through worker transitions.

273,000 per employee, an operating profit of EUR 26,000 per employee, and a return on assets of 13%.

Table 3.1: Descriptive statistics on the firm level

This table presents descriptive statistics on the firm-level. Reported are the number of observations (Obs), mean value (Mean), standard deviation (SD), median (p50), mean weighted by firms' total German employees (Wgted Mean), and standard deviation weighted by firms' total German employees (Wgted SD). A detailed description of all variables can be found in Appendix C.1.

	Obs	Mean	SD	P50	Wgted Mean	Wgted SD
Panel A: Firm characteristics						
log(total assets)	51,515	15.506	1.611	15.246	17.305	1.965
log(employees$_{firm}$)	51,515	4.163	0.888	3.906	5.698	1.281
log(employees$_{world}$)	48,937	4.312	1.104	4.111	5.718	1.336
sales/employees	48,937	0.273	0.533	0.136	0.320	0.534
profit/employees	29,795	0.026	0.079	0.011	0.026	0.074
profit/total assets	30,570	0.129	0.153	0.114	0.111	0.136
ind. unionization$_{empl}$	51,495	0.387	0.138	0.400	0.410	0.137
ind. unionization$_{estab}$	51,495	0.148	0.088	0.120	0.148	0.090
no. establishments	51,515	1.565	2.008	1.000	4.520	5.481
median(distance)	51,467	0.686	1.551	0.000	1.999	2.317
sd(distance)	51,467	0.430	1.060	0.000	0.988	1.241
manager/employees	51,496	0.039	0.063	0.018	0.045	0.061
semi-skilled/employees	51,496	0.123	0.183	0.045	0.119	0.187
skilled/employees	51,496	0.454	0.314	0.496	0.424	0.300
complex/employees	51,496	0.109	0.138	0.062	0.120	0.139
highly complex/employees	51,496	0.080	0.127	0.033	0.097	0.127
no. occupations	51,515	12.474	8.468	10.143	25.940	14.877
hhi(occupations)	51,515	0.309	0.196	0.250	0.255	0.192
Panel B: Within-firm variance of log wage and AKM components						
var(log wage)	51,515	0.103	0.054	0.093	0.113	0.051
var(person FE)	51,515	0.091	0.045	0.084	0.095	0.040
var(estab. FE)	51,515	0.001	0.002	0.000	0.002	0.003
var(Xb)	51,515	0.013	0.009	0.010	0.015	0.008
var(residual)	51,515	0.014	0.010	0.011	0.018	0.011
Panel C: Firm mean of log wage and AKM components						
log wage	51,515	4.455	0.332	4.470	4.607	0.374
person FE	51,515	4.190	0.213	4.171	4.275	0.237
firm FE	51,515	0.309	0.164	0.323	0.380	0.177
Xb	51,515	-0.044	0.035	-0.036	-0.050	0.032
residual	51,515	0.001	0.006	0.001	0.002	0.005

3.3 Method

3.3.1 Between- and within-wage inequality

$y_t^{i,j,k}$ denotes the log of the real daily wage of worker i employed by establishment j belonging to firm k in year t.[13]. Using information on the firm structure, we construct

$$\bar{y}^k = \frac{1}{N_k} \sum_t \sum_i y_t^{i,j,k}$$
$$var_k(y_t^{i,j,k}) = \frac{1}{N_k} \sum_t \sum_i (y^{i,k} - \bar{y}^k)^2, \tag{3.1}$$

where $\bar{y}^k$ is the mean wage of firm k, $var_k(y_t^{i,j,k})$ the variance of wages within firm k, and N_k the total number of employee-year observations of firm k in our sample. In parts of our analysis, we run an analysis at the establishment level. Next, we construct

$$\bar{y}^j = \frac{1}{N_j} \sum_t \sum_i y_t^{i,j}$$
$$var_k(y_t^{i,j}) = \frac{1}{N_j} \sum_t \sum_i (y^{i,j} - \bar{y}^j)^2, \tag{3.2}$$

where $\bar{y}^j$ is the mean wage of establishment j, $var_k(y_t^{i,j})$ is the variance of wages within establishment j, and N_j is the total number of employee-year observations of establishment j in our sample.

3.3.2 Wage inequality and size

We are interested in the greater variation of wages within larger firms compared to smaller firms ("within-firm inequality") and the wage differences be-

[13]For notational convenience, we suppress the dependence of subscript j on worker i and t, such that j = J(i,t), as well as the dependence of subscript k on worker i, establishment j and t, such that k = K(i,j,t).

tween smaller and larger firms ("between-firm inequality"). To establish these relationships, we regress, in our baseline models, the firm mean wage and the within-firm variance in wages on firm size:

$$\begin{aligned} var_k(y_t^{i,j,k}) &= \beta size_k + \varepsilon_k \\ \bar{y}^k &= \beta size_k + \varepsilon_k, \end{aligned} \tag{3.3}$$

where all models are weighted by firms' total number of (German) employees. We cluster standard errors on the firm level. Our baseline measure of firm size is total assets. In Appendices C.7 and C.8, we show the robustness of our main results to the use of alternative employee-based measures, total number of German employees and worldwide employees.

To analyze the role of establishment size within firms, we estimate regression models using establishment-level data, for example,

$$var_k(y_t^{i,j}) = \beta size_j + \pi_k + \varepsilon_j, \tag{3.4}$$

where π_k denotes the firm fixed effect to compare only establishments within one firm. We measure establishment size by the establishments' number of employees. All models on the establishment level are weighted by the establishments' number of employees. Standard errors are still clustered on the firm level.

3.3.3 AKM-type regression model

The implementation follows the CHK implementation of the model introduced by AKM. The following regression model is estimated in the 2010-2017 period,

$$y_t^{i,j} = \alpha^i + \psi^j + \beta X_t^i + r_t^{i,j}, \tag{3.5}$$

where $y_t^{i,j}$ denotes the log average daily wage of individual i in year t, α^i denotes the person fixed effect, ψ^j an establishment fixed effect, X_t^i denotes an index of time-varying observable characteristics, and $r_t^{i,j}$ denotes a residual. X_t^i includes an unrestricted set of year dummies and quadratic and cubic terms in age fully interacted with educational attainment.[14]

The estimation is done on the largest connected set of establishments that are linked by worker transitions within the 2010-2017 period. Importantly, the estimation is carried out on the universe of the full-time linked employer-employee dataset for Germany and not only on the subset for which we also observe the firm structure. The data preparation follows the steps conducted by CHK. The most relevant modifications are that the data also cover workers in East Germany and female workers (for further details, see Section 3.2.1).

3.3.4 Decomposition of firm mean wage and within-firm variance of wages

We use the parameter estimates from the AKM-type regression to decompose the firm mean wage ($\bar{y}^k$) and within-firm variance of wages ($var_k(y_t^{i,j,k})$). Ignoring time-varying working characteristics βX_t^i for now, the firm-level mean can be decomposed into

$$\bar{y}^k = \bar{\alpha}^k + \bar{\psi}^k, \tag{3.6}$$

where means are calculated over all person-year observations of firm k in our sample. $\bar{\alpha}^k$ is the composition part, measuring the average worker quality identified from time-invariant worker characteristics (the worker fixed effects). $\bar{\psi}^k$ is the firm-level mean of the establishment fixed effects, to which we refer as "firm fixed effect". It reflects firm-specific pay policies, such as rent-sharing (e.g., Card et al., 2018) or efficiency wages (e.g., Krueger and Summers, 1988).

Song et al. (2019) extend the standard AKM/CHK variance decomposition to

[14] As in CHK, the age variable is normalized to 40 years. See Card et al. (2018) for a discussion.

differentiate between a within-firm component and a between-firm component. We adopt the within-firm component to our setting, in which we consider the firm-establishment structure. Again ignoring time-varying working characteristics βX_t^i for now, we can decompose the within-firm variance into

$$\begin{aligned} var_k(y_t^{i,j,k}) &= var_k(\alpha^i - \bar{\alpha}^k) + var_k(\psi^j - \bar{\psi}^k) \\ &\quad + 2cov_k(\alpha^i - \bar{\alpha}^k, \psi^j - \bar{\psi}^k) + var_k(r_t^{i,j}), \end{aligned} \tag{3.7}$$

where the variances are calculated over all person-year observations of firm k in our sample. The first component is attributable to within-firm differences in worker quality. The second part measures differences in establishment-specific pay policies, such as an establishment-size dependent wage premium (e.g., Brown and Medoff, 1989). The third part reflects sorting of workers within the firm, for example, highly paid workers might be sorting to high-paying establishments of a firm. The last term is the idiosyncratic wage component that is orthogonal to all other components (Card, Heining and Kline, 2013). We interpret the last term as within-firm variation of the idiosyncratic worker wage premium.

3.4 Results

3.4.1 Firm size and wage inequality

As a first step, we investigate how firm size[15] affects wage inequality within and between firms. We show this in a graphical analysis and a regression analysis.

Figure 3.1 illustrates the relationship between wage inequality and firm size. For this figure, we sort all sample firms into deciles according to their total assets. Each decile consists of approximately 5,100 firms. Subfigure (a) focuses

[15] As explained before, we use the total assets of a firm as our main size proxy. In Section 3.4.2, we will compare this size measure to employee-based proxies.

on within-firm wage inequality and presents the mean variance of wages within firms in each decile. Additionally, we show the mean within-firm variances of the AKM components.[16]

We find a positive correlation between firm size and the within-firm variance. For the bottom decile of firms, the wage variance is on average 0.089, while it is 0.119 for the top decile. In relative terms, this equals an increase of 34%.[17] The findings that the variance of the person fixed effects accounts for approximately 85% of the wage variance and that both move nearly parallel suggest that the greater variation in wages within larger firms is mainly related to their higher heterogeneity in terms of workforce quality. The next most important component of overall wage variance is the variance of the residual. The variance of the residual includes all aspects that are not captured by the AKM model and can be interpreted as the variance of idiosyncratic worker premiums. However, the graphical analysis fails to provide evidence that the variance of the residual increases monotonously with firm size. Rather, the relationship between firm size and the variance of the residual seems to follow a U-shape. Therefore, factors that are not captured by the other AKM components do not seem to drive the positive relationship between firm size and within-firm wage inequality.

Subfigure (b) focuses on between-firm wage inequality and presents the average wage paid by a firm for each size decile. Additionally, we again show the decomposition of the overall wage into the AKM components. In line with the prior literature, we find that the average wage is increasing substantially over the size groups. While firms in the bottom size decile pay wages that are 35% below the sample mean, firms in the top decile pay wages that are 35% above the sample mean. Thus, the difference in terms of average wages between firms

[16] However, for simplicity, we only show the within-firm variances of the AKM components and ignore their covariances. Therefore, this figure does not show a full decomposition of the within-firm variance (see Appendix C for a full decomposition).

[17] Using the difference between the 90th percentile and the 10th percentile as alternative measures for wage dispersion within firms, we find that the relative increase between the bottom and top decile is 26% (results unreported).

Figure 3.1: Firm size and wage inequality
This figure illustrates the relationship between firm size and wage inequality. We sort firms into deciles according to total assets. The firm size deciles are plotted on the x-axis. Subfigure (a) presents within-firm variance of log wage, person fixed effect, establishment fixed effect, Xb, and residual for each decile. Subfigure (b) presents firm mean wage and its decomposition into AKM components for each decile. A detailed description of all variables can be found in Appendix C.1.

(a) Wage inequality within firms

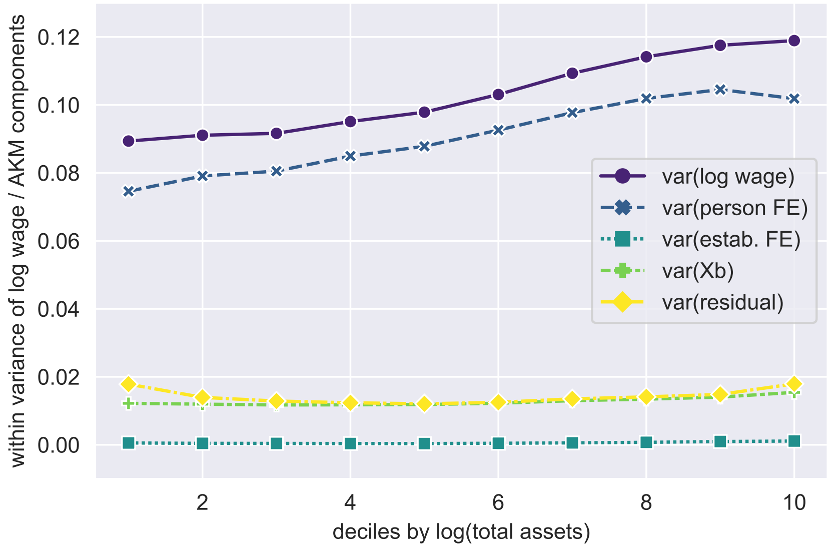

(b) Wage inequality between firms

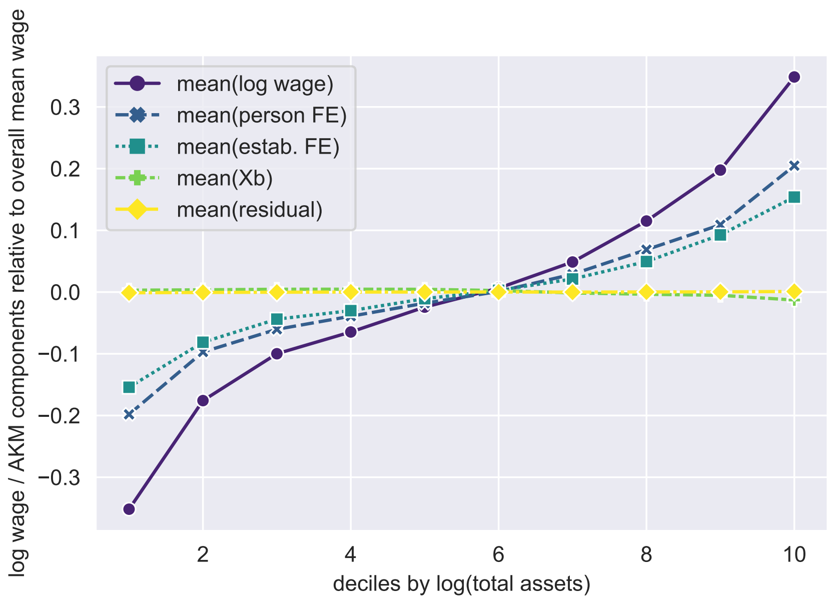

in the bottom and top decile is approximately 70%. The decomposition shows that about 40% of this differences is due to workforce composition effects and the remaining 30% is due to firm-specific wage premiums (i.e., the premium paid to a hypothetical worker with average observable and unobservable characteristics).

Table 3.2 presents the results of the corresponding regression analysis. In Panel A, we use the within-firm variance of wages and AKM components as dependent variables. The independent variable in all regressions is the log of a firm's total assets (see Equation 3.3). For the within-firm variance of wages, the regression coefficient is 0.0054. This implies that a firm would increase its within-firm inequality by 4.8% in relative terms if its size doubled.[18] Analyzing the AKM components reveals that the person fixed effects, i.e., the workforce composition, account for 65% of the positive relationship between size and wage variance within firms. Approximately one-fifth of the wage variance is explained by the variance of the residual, that is, the variance of the idiosyncratic worker premiums that are not related to workforce composition or other components in the AKM.[19]

In Panel B, we focus on wage inequality between firms and use the firm mean wage and its AKM components as dependent variables. The independent variable in all regressions is the log of a firm's total assets (see Equation 3.3). For the firm mean wage, the regression coefficient is 0.12. This implies that a firm would increase its average (log) wage by 12% in relative terms if its size doubled. The coefficient for the person fixed effect is 0.069. Thus, 60% of the positive relationship between firm size and average wages can be explained by workforce composition effects, that is, higher workforce quality in larger

[18] Note that employee-weighted mean of the within-firm variance is 0.113.

[19] Please note that the within-firm variance of log wages can be decomposed into within-firm variances of the AKM components and their covariances. The regression coefficients of the within-firm variances of the AKM components do not add up to the regression coefficient of the within-firm variance of wages in our table because the covariances of the AKM components are not reported.

Table 3.2: Wage inequality and firm size

This table presents regressions of wage inequality within and between firms on firm size, which is measured by total assets. In panel A, the dependent variables are the within-firm variance of log wage, person fixed effect, establishment fixed effect, Xb, and residual from the AKM-type regression. In Panel B, the dependent variables are the firm mean of log wage, person fixed effect, establishment fixed effect, Xb, and residual from the AKM-type regression. For the firm mean of residual, the coefficient on total assets is omitted or is virtually zero. This is why we report it as not relevant (n.r.). Regressions are weighted by firms' total number of German employees. T-statistics based on Huber/White robust standard errors clustered by firms are presented in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix C.1.

	(1)	(2)	(3)	(4)	(5)
Panel A: Within-firm inequality					
	var(log wage)	var(person FE)	var(estab. FE)	var(Xb)	var(residual)
log(total assets)	0.0054***	0.0035***	0.00033***	0.00056***	0.00091***
	(8.94)	(8.76)	(6.36)	(8.04)	(5.65)
Obs	51,515	51,515	51,515	51,515	51,515
R2	0.04	0.03	0.05	0.02	0.03
Panel B: Between-firm inequality					
	log wage	person FE	firm FE	Xb	residual
log(total assets)	0.12***	0.069***	0.053***	-0.0018***	n.r.
	(21.02)	(17.44)	(21.05)	(-6.70)	n.r.
Obs	51,515	51,515	51,515	51,515	51,515
R2	0.40	0.32	0.34	0.01	n.r.

firms. The coefficient for the firm fixed effect, which is the average across all establishment fixed effects of a firm, is 0.053. Thus, firm-specific wage premiums account for approximately 40% of the higher wages in larger firms.[20]

3.4.2 Measurement of firm size

Our main measure of firm size is total assets. The majority of the previous literature on wage inequality used employees to measure (firm) size. This is at least partly related to the fact that accounting data, such as total assets, is normally not available in datasets that cover individual workers. Our dataset includes firm-level accounting data from the Amadeus database that allows us to use total assets as our main proxy for firm size. We use this measure because it is the most commonly used size proxy[21] and it captures all resources available to the firm.

Nevertheless, we also use the number of worldwide employees of a firm according to the Amadeus database or the number of employees in Germany according to the data provided by the IAB as size proxies. First, we plot the relationship between wage firm size and wage inequality in Figure 3.2 and use total assets or the employee-based measures to sort firms into size deciles. Subfigure (a) shows that the employee-based size measures lead to very similar results as total assets for within-firm wage inequality. For between-firm wage inequality in Subfigure (b), we find that the inequality-size relationship is substantially weaker when using employee-based size measures. While there is a 70% difference between firms in the top and bottom size decile according to their total assets, this difference is only approximately 30% when we sort firms

[20]For the residual, the coefficient on total assets is omitted or is virtually zero. Therefore, we report it as not relevant (n.r.). Please note that we only present the first two true digits of coefficient estimates. Nevertheless, the regression coefficients of the AKM components without rounding add up to the regression coefficient of the firm mean wage.

[21]Dang, Li and Yang (2018) analyze the 100 most cited papers from top economics, finance, and accounting journals that use firm size measures in the area of empirical corporate finance. Among those 100 papers, 87 use a single size measure, which is total assets in 49 papers, market capitalization in 20 papers, sales in 16 papers, and number of employees in 2 papers.

by their number of employees.

We also repeat our regressions from Table 3.2 using the employee-based size measures. The results, which are reported in Appendices C.7 and C.8, are not substantially different from our main regression results. However, we again find that the economic magnitude in the between-firm regressions is smaller when using employee-based size measures. Overall, these results indicate that the choice of the size proxy does matter in the context of wage inequality and that employee-based proxies can underestimate the impact of firm size on between-firm wage inequality.

3.4.3 Development of firm size and wage inequality over time

To shed some light on the question of whether increases in firm size over the last decades can help to explain the increase in wage inequality, we analyze the development of firm size and wage inequality over time. This test uses an extended sample period from 1995 to 2015.[22] Firm size is measured by firms' total assets. For firm size, we do not only focus on the largest firms but also investigate how the size of smaller and more typical firms changed over time. Therefore, we distinguish between firms ranked 1-50, 51-100, 101-250, 251-500, and 501-1000 in terms of firm size in each year.

Subfigure (a) of Figure 3.3 shows how the average size of firms in these bins developed over time. We find that firm size increased substantially between 1995 and the financial crisis. This increase is not only observable for the largest firms but also for smaller firms. For instance, the average size of both the largest 50 firms in Germany and firms ranked between 501 and 1000 in terms of size increased by about 70%. Between 2007 and 2010, firm size declined for all bins, and it remained relatively constant after 2010. The development of within-firm, between-firm, and total wage inequality is shown in Subfigure (b). Similar to the

[22] In our other tests, we focus on the years 2010 to 2016 (see Sections 3.2.2 and 3.2.1 for a detailed explanation).

Figure 3.2: Firm size and wage inequality: capital- and employee-based firm size measures
This figure compares the relationship between firm size and wage inequality using capital- and employee-based firm size measures. We sort firms into deciles according to firms' total assets, total German employees, and total worldwide employees. The firm size deciles are plotted on the x-axis. Subfigure (a) presents the within-firm variance log wage of deciles. Subfigure (b) presents the firm mean wage of the deciles. A detailed description of all variables can be found in Appendix C.1.

(a) Wage inequality within firms

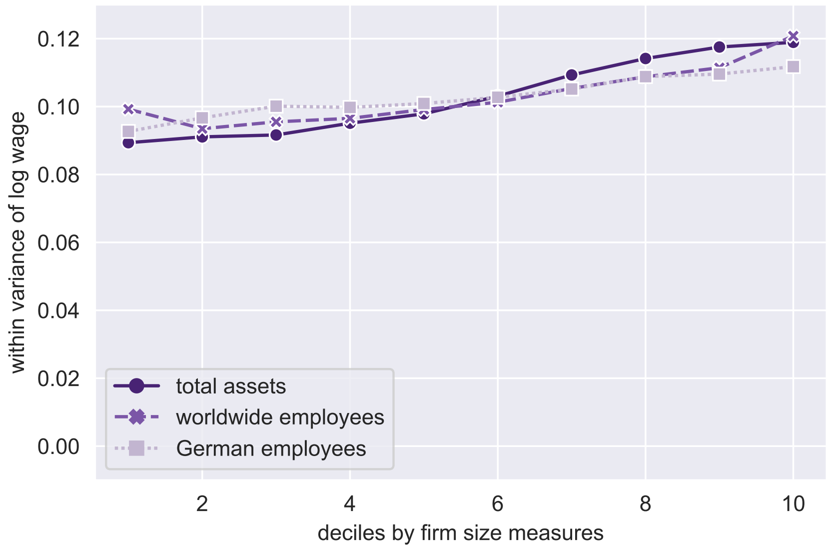

(b) Wage inequality between firms

development of firm size, we find an upward trend of wage inequality between 1995 and 2008. The overall wage inequality approximately doubled in this time period, which was mostly driven by increasing between-firm inequality. Between 2009 and 2010, wage inequality is relatively constant, and it shows a declining trend between 2011 and 2015. One possible explanation for the sharp decline between 2014 and 2015 is the introduction of the mandatory minimum wage in Germany on January 1, 2015.

When comparing the development of average firm size and wage inequality, it is remarkable that both increased in a relatively similar manner until the financial crisis. Both stopped their upward trend since the financial crisis. While the development of firm size stagnated, the wage inequality even decreased. Although this evidence is suggestive, it is consistent with the view that increases in average firm size contributed to the increasing wage inequality in the last decades.

3.4.4 What explains higher wage inequality in larger firms?

In this section, we test possible reasons for the positive relationship between firm size and wage inequality. In Section 3.4.1, we have established that higher within-firm wage inequality in larger firms is driven by the larger heterogeneity of their workforce quality and—to a smaller extent—higher variance of the idiosyncratic worker premiums. The between-firm inequality stems from higher average workforce quality in larger firms and firm-specific wage premiums. We explain each of these four components in separate subsections below.

3.4.4.1 Within-firm inequality: heterogeneity of worker quality

Our main hypothesis is that larger firms have a higher heterogeneity of workforce quality because their job characteristics differ from those in smaller firms. First, because larger firms can make better use of workers' skills due to their

Figure 3.3: Development of firm size and wage inequality over the 1995–2015 period
This figure illustrates the development of firm size and wage inequality in the 1995–2015 period. We choose this period because accounting information on German firms is very limited before 1995, and we do not (yet) have accounting information on the universe of German firms for the years 2016 and 2017. Subfigure (a) presents the relative change of firm size, which is measured by total assets, compared to 1995. In each year, we distinguish between firms ranked 1–50, 51–100, 101–250, 251–500, and 501–1000 in terms of firm size. Here, we rely on the universe of German firms from the Amadeus database. Subfigure (b) presents the within-firm, the between-firm, and the overall variation of log wages over time. For this subfigure, we assume that the linking table of establishments to firms, provided by the ORBIS-ADIAB dataset (see Section 3.2.2), is valid in the 1995–2015 period.

(a) Relative change of total assets compared to 1995

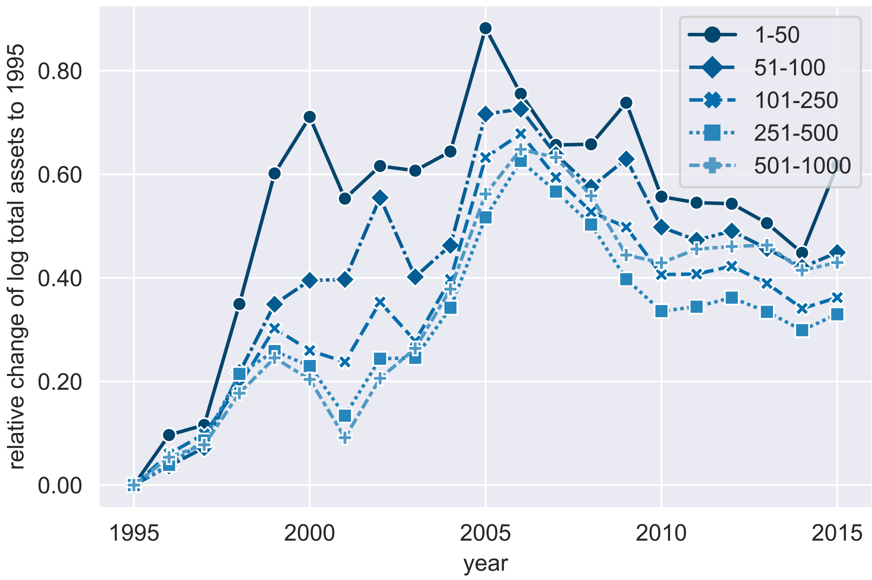

(b) Within-firm, between-firm, and overall wage inequality

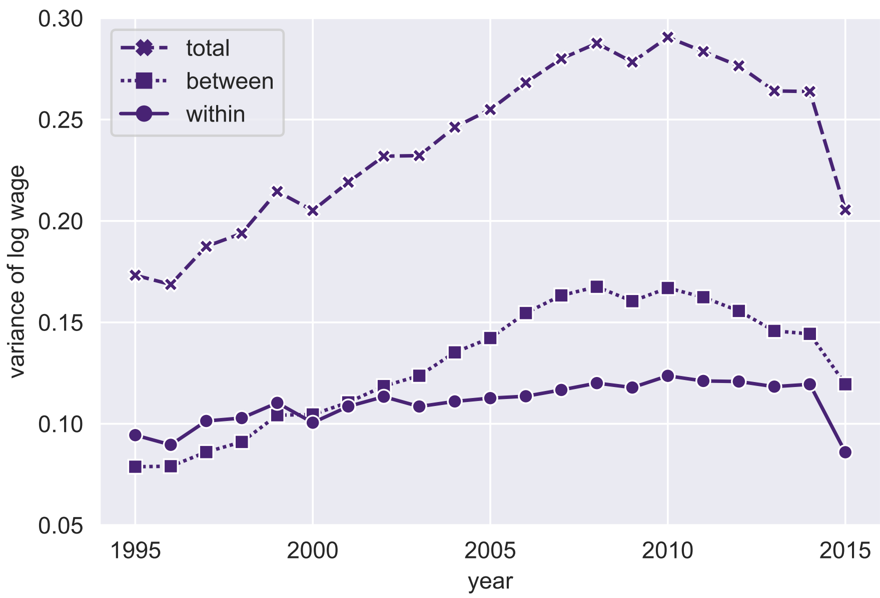

higher capital intensity (Hamermesh, 1980; Oi, 1983; Schmidt and Zimmermann, 1991; Lallemand, Plasman and Rycx, 2005), we expect that occupations that require (very) high skills are clustered in these firms. At the same time, larger firms still need to rely on occupations that require medium- or low-skilled workers to support their operations. Second, despite recent outsourcing trends (Goldschmidt and Schmieder, 2017), larger firms are more likely to perform tasks that are not their core business in-house due to economies of scale (e.g., tax reporting). Third, occupations in larger firms are more focused on a specific task because firms can divide the tasks into smaller parts, creating higher levels of specialization (Bryson, Erhel and Salibekyan, 2019). All these factors likely increase the heterogeneity of occupations in larger firms and—as a consequence—the heterogeneity of workforce quality.

Empirically, we use three different proxies to measure the dispersion of occupations within firms. The first is the number of employees at a firm who perform complex tasks standardized by a firm's total employees.[23] The second measure is closely related to the first and captures the fraction of highly complex occupations. The last proxy uses the number of different occupations within a firm to measure the concentration of occupations. For this purpose, we calculate the Herfindahl index at the firm level based on the share of occupations within the firm.[24] We plot how these three proxies and the number of occupations are related to firm size in Figure C.4. The fraction of (highly) complex occupations and the number of occupations increase with firm size. For example, firms in the top size decile have on average about 25 different occupations, while the corresponding number for firms in the bottom decile is only 8. Not surprisingly, the concentrations of occupations as measured by the Herfindahl index

[23]The Classification of Occupations 2010 (KldB 2010) makes it possible to differentiate between occupational activities according to four degrees of complexity, which are unskilled or semi-skilled, specialist, complex specialist, and highly complex activities. The level of requirement is given in the last digit of the 5-digit occupation code.

[24]For this calculation, an occupation is measured by the first 3 digits of the Classification of Occupations 1988 (KldB 1988).

falls with firm size. Therefore, job characteristics seem to differ substantially between small and large firms, potentially explaining the higher variation in workforce quality in the latter.

As a next step, we include these three proxies as control variables in our baseline regressions. Table 3.3 presents the results. For comparison, Column 1 shows the baseline model, regressing the within-firm variance of the person fixed effects on total assets. The coefficient estimate is 0.0035. In Column 2, we add the share of employees carrying out unskilled to semi-skilled tasks to the model. The coefficient estimate on firm size remains unchanged. In Column 3, we include the share of employees performing highly complex tasks. The coefficient estimate on firm size drops to 0.0029. In Column 4, we add the Herfindahl index of occupations to the model. The coefficient on firm size is 0.0022 in this model, which is also considerably smaller than in the baseline model. The effects of the additional variables are as expected; a higher share of complex occupations increases the heterogeneity of the workforce quality, while the opposite is true for a higher concentration of occupations. The last column presents a specification in which we include all three proxies simultaneously. The coefficient on firm size is now 0.0018, which is approximately 50% smaller than in our baseline model. Therefore, differences in job characteristics seem to be a key driver for the greater worker heterogeneity in larger firms.

3.4.4.2 Within-firm inequality: variance of idiosyncratic worker premiums

The residual of the AKM model can be interpreted as idiosyncratic worker premium. We hypothesize that larger firms exhibit a higher variance of the idiosyncratic worker premium due to more severe monitoring problems. Because monitoring is more complex in larger firms, employees have more opportunities for shirking (e.g., Eaton and White, 1983; Shapiro and Stiglitz, 1984). In order to reduce shirking behavior, larger firms can use more tournaments

Table 3.3: Within-firm inequality: explaining greater worker heterogeneity in larger firms

The dependent variable is the within-firm variance of the person fixed effect. Column 1 presents the baseline model regressing the variance of the person fixed effect on total assets. In columns 2 to 5, we add control variables for the dispersion of occupations within firms to the baseline model. These are the number of employees performing unskilled to semi-skilled tasks standardized by total German employees, the number of employees performing highly complex tasks standardized by total German employees, and the Herfindahl index as a concentration measure of occupations. Regressions are weighted by firms' total number of German employees. T-statistics based on Huber/White robust standard errors clustered by firms are presented in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix C.1.

	(1)	(2)	(3)	(4)	(5)
log(total assets)	0.0035***	0.0035***	0.0029***	0.0022***	0.0018***
	(8.76)	(9.07)	(6.11)	(5.64)	(4.30)
semi-skilled/employees		0.00088			0.0076**
		(0.25)			(2.29)
highly complex/employees			0.052***		0.049***
			(8.57)		(8.16)
hhi(occupations)				-0.036***	-0.034***
				(-11.64)	(-9.87)
Obs	51,515	51,496	51,496	51,515	51,496
R2	0.03	0.03	0.06	0.06	0.08

(i.e., wage differentials between hierarchy levels) and/or bonus payments (i.e., performance-based incentive payments) to incentivize workers. Thus, wages can be more dispersed in larger firms, even after accounting for their higher heterogeneity of workforce quality.

To measure monitoring complexity, we use three different proxies. Our first measure is a firm's number of managers per employees, which captures the relationship between (supervising) managers and normal workers. We assume that a high fraction of managers, relative to workers, indicates higher monitoring complexity. The other two measures we use exploit the geography of a firm's establishments. Monitoring is likely more complex if the median distance between establishments and the headquarters is higher and if their location is more dispersed (e.g., Giroud, 2013). Thus, we use the median distance and the standard deviation of the distances as second and third measures. We plot how these three proxies are related to firm size in Figure C.5. All three show a positive relationship with firm size. For example, the fraction of managers is about two percent in the bottom firm size decile, while it is almost ten percent in the top decile. Although we acknowledge that these proxies are far from being perfect, they indicate that monitoring complexity increases with firm size.

Next, we include these three proxies as control variables in our baseline regressions (see Table 3.4). The dependent variable is the within-firm variance of idiosyncratic worker premiums, but we exclude managers from this calculation.[25] For comparison, Column 1 shows the baseline model, regressing the variance of idiosyncratic worker premiums on total assets. The coefficient estimate is 0.00079. In Column 2, we include the ratio of managers to employers as control variable. In this specification, the coefficient on total assets drops by approximately 40% to 0.00049. In Column 3, we control for the median distance between the firm headquarters and the establishments. Here, the coefficient estimate drops even further to 0.00022. Column 4 add the standard deviation of this distance to the model. The coefficient of total assets is slightly lower than in the baseline model (0.00052). In the last column, we control for all three proxies of monitoring complexity at the same time. In this specification, the coefficient estimate on size is negative (-0.000042) and statistically insignificant. The coefficient estimates for our monitoring complexity proxies are as expected as all three have a positive impact on the variance of idiosyncratic worker premiums. Overall, this test provides evidence that the greater variation of idiosyncratic worker premiums in larger firms is related to their higher monitoring complexity.

3.4.4.3 Between-firm inequality: worker quality

We hypothesize that larger firms have a higher (average) workforce quality because their job characteristics differ from those in smaller firms. The argument

[25] Managers are excluded to mitigate the concern that the higher share of managers itself and not the higher monitoring complexity in larger firms is driving the greater variation of the AKM residuals. The results are very similar if we do not exclude managers.

Table 3.4: Within-firm inequality: explaining greater variation of idiosyncratic wage premiums in larger firms

The dependent variable is the within-firm variance of the residual. Managers are excluded from this analysis to mitigate the concern that the higher share of managers itself and not the higher monitoring complexity in larger firms is driving the greater variation of the AKM residuals. Column 1 presents the baseline model regressing the variance of the residual on total assets. In columns 2 to 5, we add control variables for monitoring complexity to the baseline model. These are the number of managers employed divided by German employees, the median distance between a firm's headquarters and establishments, and the standard deviation of this distance. Regressions are weighted by firms' total number of German employees. T-statistics based on Huber/White robust standard errors clustered by firms are presented in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix C.1.

	(1)	(2)	(3)	(4)	(5)
log(total assets)	0.00079***	0.00049***	0.00022**	0.00052***	-0.000042
	(5.16)	(2.98)	(2.20)	(3.62)	(-0.40)
manager/employees		0.045***			0.044***
		(10.73)			(11.70)
median(distance)			0.0011***		0.0011***
			(8.93)		(5.09)
sd(distance)				0.0013***	7.3e-06
				(8.23)	(0.02)
Obs	51,515	51,496	51,467	51,467	51,457
R2	0.02	0.08	0.07	0.04	0.13

here is similar as for their higher heterogeneity of workforce quality. First, capital-skill complementarities lead to a higher share of occupations which are (highly) complex and require workers with high skill levels. Second, larger firms may be more likely to engage in outsourcing, which reduces occupations which require lower skills. Examples include food, cleaning, security, and logistics tasks (see Goldschmidt and Schmieder, 2017, for more details on outscoring trends in Germany). Both factors can lead to a higher share of "complex" occupations in larger firms, which shifts the average level of workforce quality upwards.

We use three proxies to measure the complexity of occupations. These are the fraction of employees conducting unskilled to semi-skilled, complex, and highly complex tasks. As expected, larger firms have a lower fraction of employees performing unskilled to semi-skilled tasks and a higher fraction of employees performing (highly) complex tasks (see Figure C.4).

In Table 3.5, we include these variables as controls in our baseline regressions. For comparison, Column 1 shows the baseline model, regressing the mean worker quality, measured by the average person fixed effect in a firm, on total assets. The coefficient in the baseline model is 0.069. In Column 2, we add the share of employees carrying out an unskilled or semi-skilled task to the model. The coefficient on total assets decreases to 0.060. Next, we include the share of employees carrying out complex tasks. The coefficient on total assets decreases to 0.062. In Column 3, we add the share of employees performing highly complex tasks. This reduces the coefficient on total assets to 0.059. The effects of our measures for job complexity are as expected, with a higher average worker quality in firms that have a higher share of (highly) complex occupations. Column 5 presents a joint model that includes all three control variables simultaneously.[26] The coefficient of total assets is 0.051 in this model, which represents approximately a one-quarter reduction compared to the baseline model. Hence, differences in the complexity of occupations can partially explain the higher workforce quality in larger firms.

3.4.4.4 Between-firm inequality: large-firm wage premium

Previously, we documented that firm-specific wage premiums account for slightly less than half of the higher wages in larger firms. This means that larger firms pay higher wages even after accounting for workforce composition effects and the other factors in the AKM. Several explanations for this LFWP haven been discussed in the existing literature. We focus on efficiency wages, rent-sharing, unionization, and local labor market effects.

First, the efficiency wage explanation argues that larger firms face higher monitoring complexity and pay higher wages to incentivize workers. In Panel A of Table 3.6, we control for the same proxies for monitoring complexity as in Sec-

[26] The level of employees' requirements is based on the last digit of the 5-digit Classification of Occupations 2010. It distinguishes four levels of requirement: unskilled to semi-skilled, skilled, complex, and highly complex. The omitted category in Column 5 is employees performing skilled tasks.

Table 3.5: Between-firm inequality: explaining higher worker quality in larger firms

The dependent variable is the firm mean of the person fixed effect. Column 1 presents the baseline model regressing the mean person fixed effect on total assets. In columns 2 to 5, we add control variables for the complexity of occupations to the baseline model. These are the number of employees performing unskilled to semi-skilled tasks standardized by total German employees, the number of employees performing complex tasks standardized by total German employees, and the number of employees performing highly complex tasks standardized by total German employees. Regressions are weighted by firms' total number of German employees. T-statistics based on Huber/White robust standard errors clustered by firms are presented in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix C.1.

	(1)	(2)	(3)	(4)	(5)
log(total assets)	0.069***	0.060***	0.062***	0.059***	0.051***
	(17.44)	(14.79)	(14.64)	(13.23)	(11.49)
semi-skilled/employees		-0.41***			-0.34***
		(-16.59)			(-12.88)
complex/employees			0.51***		0.26***
			(10.52)		(5.89)
highly complex/employees				0.74***	0.56***
				(15.38)	(14.90)
Obs	51,515	51,496	51,496	51,496	51,496
R2	0.32	0.42	0.41	0.47	0.57

tion 3.4.4.2. However, neither the fraction of managers per employee nor the median distance between establishments nor its standard deviation lead to a substantial reduction of the coefficient for total assets. This result indicates that monitoring complexity is not the main driver behind the LFWP.

Table 3.6: Between-firm inequality: explaining the large-firm wage premium

The dependent variable is the firm mean of the establishment fixed effect. Column 1 presents the baseline model regressing the mean establishment fixed effect on total assets. In columns 2 to 5, we add control variables and fixed effects to the baseline model. In panel A, we add control variables for monitoring complexity. These are the managers per employees, the median distance between firm headquarters and establishments, and the standard deviation of the distance. In panel B, we include measures for firms' rent. These are the profit per employees, the sales per employees, and the profit standardized by total assets. In panel C, we control for industry unionization. In detail, we add the industry unionization rate based on employees, its interaction term with total assets, the industry unionization rate based on establishments, and its interaction term with total assets to the model. For the interaction terms, the variables are centered at the mean. In Panel D, we add industry, county, and industry-county fixed effects to control for local effects. Panels A to D are based on the firm-level dataset, and regressions are weighted by firms' total number of employees. In Panel E, we test the role of local effects on establishment level. This allows us to control for the industry of the establishment, the county of the establishment, and the combination of establishment industry and county. Establishment-level regressions are weighted by establishments' number of employees. T-statistics based on Huber/White robust standard errors clustered by firms are presented in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix C.1.

	(1)	(2)	(3)	(4)	(5)
Panel A: Efficiency wages					
log(total assets)	0.053***	0.050***	0.059***	0.053***	0.055***
	(21.05)	(18.75)	(38.76)	(20.03)	(33.41)
manager/employees		0.43***			0.44***
		(6.93)			(8.47)
median(distance)			-0.011***		-0.018***
			(-4.39)		(-4.23)
sd(distance)				-0.00016	0.018***
				(-0.08)	(3.40)
Obs	51,515	51,496	51,467	51,467	51,457
R2	0.34	0.37	0.36	0.35	0.40
Panel B: Rent-sharing					
log(total assets)	0.053***	0.051***	0.051***	0.052***	0.050***
	(21.05)	(14.24)	(18.28)	(16.22)	(13.57)
profit/employees		0.16***			0.034
		(3.45)			(0.77)
sales/employees			0.029***		0.030***
			(4.35)		(4.10)
profit/total assets				0.038*	0.023
				(1.87)	(1.06)
Obs	51,515	29,795	48,937	30,570	29,795
R2	0.34	0.26	0.35	0.26	0.27

Continued on next page

Table 3.6 continued

	(1)	(2)	(3)	(4)	(5)
Panel C: Industry unionization					
log(total assets)	0.053***	0.053***	0.050***	0.052***	0.051***
	(21.05)	(21.25)	(23.07)	(20.23)	(20.73)
ind. unionization$_{empl}$		-0.044	-0.16***		
		(-1.27)	(-8.03)		
log(total assets) x unionization$_{empl}$			0.064***		
			(3.98)		
ind. unionization$_{estab}$				-0.18***	-0.27***
				(-3.76)	(-8.52)
log(total assets) x unionization$_{estab}$					0.044**
					(2.24)
Obs	51,515	51,495	51,495	51,495	51,495
R2	0.34	0.35	0.36	0.35	0.35
Panel D: Local effects at the firm level					
log(total assets)	0.053***	0.040***	0.050***	0.034***	0.033***
	(21.05)	(21.19)	(40.58)	(37.64)	(38.10)
Industry FE	No	Yes	No	Yes	Yes
County FE	No	No	Yes	Yes	Yes
Industry x county FE	No	No	No	No	Yes
Obs	51,515	51,477	51,515	51,477	38,129
R2	0.34	0.58	0.49	0.69	0.82
Panel E: Local effects at the establishment level					
log(total assets)	0.054***	0.037***	0.049***	0.033***	0.027***
	(22.06)	(14.58)	(21.79)	(15.84)	(8.57)
Estab. industry FE	No	Yes	No	Yes	Yes
Estab. county FE	No	No	Yes	Yes	Yes
Estab. ind. x county FE	No	No	No	No	Yes
Obs	115,632	115,591	115,046	115,016	85,251
R2	0.33	0.59	0.45	0.69	0.83

Second, larger firms may generate higher rents by exploiting their market power, and they may share some of their higher rents with employees. To measure firms' rents, we use three accounting-based measure—profits per employee, sales per employee, and profits to total assets. When we explore the relationship between firm size and these three measures in Figure C.6, we find that profits and sales per employee increase with firm size. For profits to total assets, however, we find lower values for larger firms. When we control for the measures in Panel B of Table 3.5, none of these proxies has a substantial impact on the coefficient for total assets. Therefore, we cannot find any evidence that rent-sharing is the driving force behind the LFWP.

Third, larger firms may have higher unionization rates, which enables employees to extract a higher fraction of rents, and a higher threat of unionization (Dickens and Lang, 1985). Unfortunately, we cannot observe unionization at the firm level. Therefore, we must measure unionization as the fraction of unionized employees and the fraction of unionized establishments in an industry. To allow for a different effect of industry unionization on larger firms, we include not only industry unionization in our model but also the interaction term with firm size. However, the coefficient estimate on firm size in Panel C of Table 3.5 ranges in all models from 0.050 to 0.053, which is very close to the baseline estimate of 0.053. Therefore, unionization is also unlikely the main reason for the existence of a LFWP.

Fourth, larger firms may cluster in regions with high economic activity and thus a high demand for labor. This can lead to excess demand for labor and increase local wages. To test whether such effects have an impact on the LFWP, we adjust our baseline model and only compare large and small firms in the same region and industry. If the clustering of large firms in particular regions plays an important role in the LFWP, we would expect that the size effect is much weaker after including industry and region fixed effects in the model. The results in Panel D of Table 3.5 show that the inclusion of industry fixed effects, especially in combination with county fixed effects, reduces the impact of firm size on the LFWP by about one-third (from 0.053 to 0.033).

A potential concern with these results is that firm-level regressions could have limited validity in this context because the county where the firm's headquarters is located is not necessarily the county in which the majority of the operations takes place. Furthermore, different establishments of the firm may source their workforce from different labor markets, depending on their role in the firm (e.g., R&D center vs. sales office). Therefore, we repeat the analysis at the establishment level in Panel E. The results are even stronger for this specification; the coefficient for firm size drops by approximately 50% when we include county fixed effects, industry fixed effects, and their interaction. Overall, these results indicate that local effects can help to partially explain the existence of the LFWP.

3.4.5 Establishment size within firms and wage inequality

To better understand the role of firm size for wage inequality, we now investigate how the variation in the size of establishments within firms affect wages. Theoretically, two mechanisms can lead to the previously documented higher and more dispersed wages in larger firms. First, the size of the overall firm may matter for wage policies of firms. In this perspective, large firms would pay higher and/or more dispersed wages, independent of the size of a particular establishment within that firm. Second, establishment size may matter for wages, and every establishment of a firm may have its own wage policy, depending on its size. In this perspective, only large establishments within firms would pay higher and/or more dispersed wages. If this holds true, our previous results for larger firms would be driven by their, on average, larger establishments. Disentangling firm size versus establishment size as potential drivers behind larger firms' higher and more dispersed wages will shed more light on the reasons for differences in wage policies.

Empirically, we conduct regressions at the establishment level for this test and include firm fixed effects. This firm fixed effect ensures that the results are driven by variations in the size of establishments within firms. Because information on total assets is not available at the establishment level, we use the number of employees in a particular establishment as a size proxy in this test. Furthermore, we have to exclude all firms that only operate a single establishment due to the lack of within-firm establishment size variation. Appendix C.2 presents summary statistics at the establishment level.

3.4.5.1 Establishment size within firms and the variation in wages

We first analyze how size differences of establishments within firms affect the variation in wages. For this purpose, we regress the within-establishment variance in the wage and the AKM components on establishment size. Table 3.7 presents the results. In Panel A, the dependent variable is the within-establishment variance in wages. For comparison, we start by estimating a model without firm fixed effects in Column 1. When we add the firm fixed effect in Column 2, the coefficient estimate for establishment size increases by more than one-third (from 0.0089 to 0.014). To control for industry or region-specific pay policies within a firm, Columns 3 and 4 add interactions of firm fixed effects with the establishment-industry and county fixed effects. The coefficients are slightly lower in these specifications (0.012 and 0.012). These results indicate that variations in establishment size within firms have an even greater impact on wage inequality than cross-sectional firm-size variations.

In Panel B, the dependent variable is the within-establishment variance in the person fixed effect, which is the most important AKM component for the within-firm variance of wages (cf., Section 3.4.1). We find that adding a firm fixed effect to the model increases the effect of establishment size on the variation in the person fixed effect (0.011 instead of 0.0067). Adding controls for industry and region reduces the coefficient again, but it still stays above the baseline model. This result is broadly consistent with our prior explanation for the higher variation in workforce quality in larger firms because larger establishments within firms likely have a greater variety of occupations.

The variation in idiosyncratic wage premiums is investigated in Panel C. We find that the firm fixed effect reduces the effect of establishment size (0.0011 instead of 0.0017). Interestingly, the coefficient for establishment size decreases even further when we control for industry (to 0.00094) and region (0.00077). Again, these findings are consistent with the monitoring complexity explanation

Table 3.7: Wage inequality within establishments and establishment size

The dependent variables are the within-establishment variance of log wage in Panel A, the within-establishment variance of the person fixed effect in Panel B, and the within-establishment variance of the residual in Panel C. In column 1, the dependent variables are regressed on establishments' number of employees. In column 2, we add firm fixed effects to the model to compare only establishments within firms. In columns 3, we include firm fixed effects interacted with establishment industry fixed effects to compare establishments within firms and in the same industry. In columns 4, we include firm fixed effects interacted with establishment county fixed effects to compare establishments within firms and in the same county. The sample is constructed at the establishment level and only consists of establishments that belong to a firm with at least two establishments. Regressions are weighted by establishments' number of employees. T-statistics based on Huber/White robust standard errors clustered by firms are presented in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix C.1.

	(1)	(2)	(3)	(4)
Panel A: Variance(log wage)				
log(employees$_{estab}$)	0.0089***	0.014***	0.012***	0.012***
	(9.58)	(24.90)	(20.86)	(8.88)
Firm FE	No	Yes	Yes	Yes
Firm FE x Estab. industry FE	No	No	Yes	No
Firm FE x Estab. county FE	No	No	No	Yes
Obs	72,246	72,092	61,767	26,209
R2	0.05	0.64	0.72	0.79
Panel B: Variance(person FE)				
log(employees$_{estab}$)	0.0067***	0.011***	0.0097***	0.0087***
	(9.58)	(27.37)	(21.41)	(16.86)
Firm FE	No	Yes	Yes	Yes
Firm FE x Estab. industry FE	No	No	Yes	No
Firm FE x Estab. county FE	No	No	No	Yes
Obs	72,242	72,085	61,762	26,207
R2	0.05	0.68	0.72	0.82
Panel C: Variance(residual)				
log(employees$_{estab}$)	0.0017***	0.0011***	0.00094***	0.00077***
	(9.90)	(8.14)	(7.15)	(2.60)
Firm FE	No	Yes	Yes	Yes
Firm FE x Estab. industry FE	No	No	Yes	No
Firm FE x Estab. county FE	No	No	No	Yes
Obs	72,242	72,085	61,762	26,207
R2	0.08	0.69	0.75	0.83

for higher variance in idiosyncratic wage premiums in larger firms. Although larger establishments are also more difficult to monitor than smaller ones, the main difficulty arises from firm size rather than establishment size.

3.4.5.2 Establishment size within firms and the level of wages

The next step is to investigate how size differences of establishments within firms affect average wages. For this purpose, we regress the average wage of an establishment and the AKM components on establishment size. Table 3.8 presents the results. In Panel A, the dependent variable is the average wage. For comparison, we start by estimating a model without firm fixed effects in Column 1. When we add the firm fixed effect in Column 2, the coefficient estimate for establishment size drops substantially (from 0.054 to 0.026). To control for industry- or region-specific pay policies within a firm, Columns 3 and 4 add interactions of firm fixed effects with establishment-industry and county fixed effects. While the industry-fixed effect has little impact, the coefficient for establishment size drops to 0.0064 (less than 20% of the baseline coefficient) and becomes statistically insignificant when we only compare small and large establishments in the same region. This result indicates that firm size rather than establishment size is the main driver behind higher wages in larger firms and that regional factors matter greatly for the average establishment wage.

In Panels B and C, we repeat these analyses for the establishment mean of the person fixed effect (i.e., workforce quality) and the establishment fixed effect (i.e., establishment wage premium). We again find that the inclusion of the firm fixed effect in Column 2 reduces the coefficient for establishment size (from 0.033 to 0.023), but the decrease is less pronounced than for the overall wage in Panel A. When we only compare establishments in the same region in Column 4, the coefficient goes down to 0.0030 and becomes statistically insignificant. This result suggest that firms use workers of similar quality in the same region,

Table 3.8: Wage inequality between establishments and establishment size

The dependent variables are the establishment mean of log wage in Panel A, the person fixed effect in Panel B, and the establishment fixed effect in Panel C. In column 1, the dependent variables are regressed on establishments' number of employees. In column 2, we add firm fixed effects to the model to compare only establishments within firms. In columns 3, we include firm fixed effects interacted with establishment-industry fixed effects to compare establishments within firms and in the same industry. In columns 4, we include firm fixed effects interacted with establishment-county fixed effects to compare establishments within firms and in the same county. The sample is constructed at the establishment level and only consists of establishments that belong to a firm with at least two establishments. Regressions are weighted by establishments' number of employees. T-statistics based on Huber/White robust standard errors clustered by firms are presented in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix C.1.

	(1)	(2)	(3)	(4)
Panel A: Log wage				
log(employees$_{estab}$)	0.054***	0.026***	0.024***	0.0064
	(7.64)	(13.15)	(10.88)	(0.77)
Firm FE	No	Yes	Yes	Yes
Firm FE x Estab. industry FE	No	No	Yes	No
Firm FE x Estab. county FE	No	No	No	Yes
Obs	74,111	74,088	63,572	27,214
R2	0.40	0.93	0.95	0.96
Panel B: Person FE				
log(employees$_{estab}$)	0.033***	0.023***	0.020***	0.0030
	(7.17)	(11.82)	(8.76)	(0.54)
Firm FE	No	Yes	Yes	Yes
Firm FE x Estab. industry FE	No	No	Yes	No
Firm FE x Estab. county FE	No	No	No	Yes
Obs	74,111	74,088	63,572	27,214
R2	0.30	0.86	0.90	0.93
Panel C: Establishment FE				
log(employees$_{estab}$)	0.023***	0.0055***	0.0055***	0.0033
	(6.97)	(4.89)	(3.92)	(1.01)
Firm FE	No	Yes	Yes	Yes
Firm FE x Estab. industry FE	No	No	Yes	No
Firm FE x Estab. county FE	No	No	No	Yes
Obs	74,111	74,088	63,572	27,214
R2	0.34	0.91	0.92	0.95

independent of the size of an establishment. Panel C shows the results for the establishment wage premium. The firm fixed effect in Column 2 leads to a substantial reduction in the establishment size coefficient (from 0.023 to 0.0055). Similar to before, the coefficient is even smaller and statistically insignificant after the inclusion of county-fixed effects. This finding shows that the LFWP does not depend on the size of a particular establishment within the firm, at least if regional factors are taken into account. This finding fits well with the previously discussed (partial) explanation of the LFWP, which is that larger firms are more likely to be located in regions with high economic activity and a high demand for labour. If a large firm has its headquarters in such a location, it is plausible to assume that the majority of their establishments are located in this region as well, no matter of their size. Therefore, firm size rather than establishment size matters most for the LFWP.

3.4.6 Firm structure and wage inequality

Our results emphasize the importance of information on firm structures. Without this information, researchers may falsely classify wage heterogeneity across establishments within firms as between-firm inequality. To analyze this effect, we conduct a decomposition of the overall wage inequality into a between and a within component using firm-level and establishment-level information. The procedure and formulas are presented in greater detail in Appendix C. Figure 3.4 presents the results. We find that the average downward bias of the within component to the overall wage variation is approximately 2.5 percentage points during our sample period from 2010 to 2016.

Figure 3.4: Decomposition of overall wage inequality: firm versus establishment structure
This figure compares the decomposition of overall wage inequality into a between and within component using information on firm and establishment structure in the 2010–2017 period. Subfigure (a) presents the decomposition. Subfigure (b) presents the share of the within component to the overall variation based on firm-level information and on establishment-level information. A detailed description of all variables can be found in Appendix C.1.

(a) Decomposition into between/within **(b)** Within/total variation

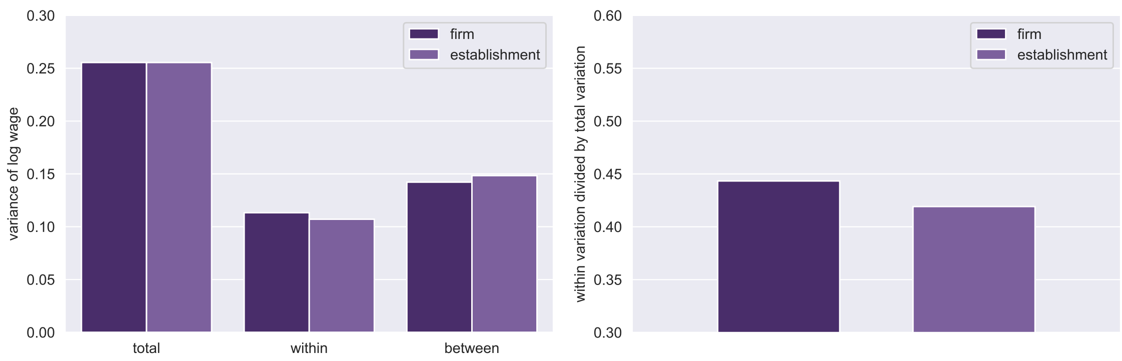

3.5 Conclusion

In this paper, we analyze the role of firm size for within-firm and between-firm wage inequality. Using a linked employee-establishment-firm dataset from Germany that covers approximately 51,000 individual firms, 115,000 establishments, and 11.6 million workers, we first find that firm size has a strong positive correlation with both types of wage inequality.

Second, we decompose wage inequality into workforce composition and non-composition effects. Composition effects are responsible for 65% of the relationship between firm size and within-firm wage inequality. The corresponding figure for between-firm wage inequality is 60%. Thus, the finding that larger firms pay more unequal and on average higher wages is strongly linked to differences in workforce composition between large and small firms.

Third, we go one step further and try to explain why firm size has a positive impact on the different components of wage inequality. For within-firm wage inequality, we find that higher wage variance in larger firms is largely explainable by more heterogeneous job characteristics and higher employee monitoring complexity. Higher average wages in larger firms are more challenging to explain. Different job characteristics play some role in the higher average workforce quality in larger firms, but they cannot fully explain the large-firm effect. Differences in profitability, monitoring complexity, or unionization levels cannot explain higher wages in larger firms after composition effects (the so-called large-firm wage premium). If anything, local labor markets play some role, but again they cannot fully explain the large-firm effect.

Fourth, we exploit variations in establishment size within firms to shed more light on the relationship between (firm) size and wage inequality. Larger establishments of a firm pay more unequal wages, which is mainly driven by their higher variation in workforce quality. However, larger establishments do not pay an economically significant wage premium. This indicates that all establishments of large firms pay a wage premium, relatively independent of their particular size.

Overall, these results show that firms in general and their size in particular matter for wage inequality. However, our findings also suggest that there is little direct (or causal) connection between firm size and wage inequality. Rather, differences in wage dispersion and average wages between small and large firms are (at least partly) explainable by factors that are correlated with firm size ("omitted variables"). These are job characteristics, monitoring complexity, and potentially clustering of larger firms in specific regions. However, even in the absence of a direct connection between firm size and wage inequality, these results imply that wage inequality will increase if firms in an economy become larger. This is not restricted to the top end of the size distribution (so-called mega firms) but also holds true for size increases of smaller and more representative firms.

4

Conclusion

The three essays of this dissertation examine research questions on finance and labor. In the first essay, I study the role of customer concentration in the propagation of idiosyncratic firm-level shocks from customers to suppliers. As idiosyncratic disruptions, I utilize major labor strikes. In the second essay, I investigate how outside directorships of CEOs influence their managerial decision-making. For this purpose, I analyze how CEOs react after they observe, as directors of another firm, a labor strike. In the third essay, I examine the relationship between firm size and wage inequality using a large matched employee-establishment-firm dataset. In this chapter, I briefly summarize the main results of the three essays and highlight their contributions and implications.

In the *first essay*, I study how customer concentration accelerates the upstream propagation of idiosyncratic firm-level shocks. As idiosyncratic disruptions, I utilize 223 major labor strikes at 110 large publicly listed U.S. firms that source intermediates from a large number of suppliers. I show that strike-hit customers impose a substantial output loss on their suppliers. On average, suppliers' drop in sales growth is 1.9 to 2.8 percentage points in the quarter of strike and the following quarter, which equals a decline by about one-third. Hence, the disruptions of customers quickly transmit to their suppliers.

In line with theoretical predictions, the negative effect on suppliers' sales growth increases with their direct dependence on the strike-hit customers. Furthermore, I show that suppliers' output loss increases by a factor of three to four if the suppliers sell products to other customers whose business also depends on the central disrupted customer. Finally, I document that suppliers' direct dependence on customers, suppliers' indirect dependence on central customers, and the number of suppliers all monotonically increase with customers' firm size.

Altogether, the findings of the first essay suggest that customer concentration is a major risk of a firm. To assess firms' customer concentration, it is not only necessary to consider the direct customer concentration but also to take into account indirect links between firms' customers. For the aggregate level, my firm-level findings suggest that customer concentration increases the vulnerability of production networks. In particular, production networks are vulnerable to idiosyncratic shocks that hit very large firms with multiple highly-dependent suppliers.

The first essay contributes to multiple strands of the literature. First, I add to the literature on shock propagation in firm-level input-output networks (e.g., Barrot and Sauvagnat, 2016; Carvalho et al., 2016) by studying in-depth how direct and indirect customer concentration accelerate the upstream propagation of shocks. Second, my findings contribute to the literature on the role of very large ("granular") firms in firm-level and aggregate volatility. I provide empirical evidence for the theoretical predictions of the model by Herskovic et al. (2018)

that explains higher volatility of smaller firms by a lower diversification of their customer base. Furthermore, I point out that it is the very large firms that have a large number of suppliers with a high direct and an even higher indirect customer concentration. These findings are also related to papers studying the macroeconomic impacts of firms' failure in production networks (e.g., Baqaee, 2018).

Third, the first essay adds to the literature in financial economics that studies how firms are affected by their customers. Recent studies find an effect of firms' (direct) customer concentration on their financing costs (e.g., Campello and Gao, 2017). My results suggest that suppliers' additional indirect links to customers should also affect their financing costs. Moreover, my results are related to studies on how customers affect suppliers' corporate policies (e.g., Chu, 2012). These papers find that suppliers hold less leverage if a customer must make relation-specific investments, if there is more competition among suppliers, and if suppliers' products are easy to substitute. My results support the alternative explanation that suppliers hold less leverage to be financially flexible when shocks disrupt their customers. Lastly, the first essay adds to the literature on the effect of labor strikes (e.g., McHugh, 1991; Persons, 1995) by showing, on the firm-level, that strikes at customers impose substantial output losses on their suppliers.

In the *second essay*, I investigate how outside directorships of CEOs influence their managerial decision-making. For this purpose, I analyze how CEOs' precautionary behavior and labor relation management changes after they observed a labor strike as a director of another firm. I focus on labor strikes for three reasons. First, they are exogenous to the CEO firm. I show that strikes do not affect strike risk at other firms once I exclude events from the same industry. Second, labor disputes are very costly for firms and require substantial attention from the board. Third, in contrast to many other types of shocks, such as the sudden death of a CEO or weather-related losses, strikes are not exogenous for the strike-hit firm, enabling CEOs to influence the strike risk at their firms.

I identify 215 events of CEOs who observe a strike as a director at another firm. I first document that CEOs increase cash holdings shortly after the strike. This effect is more pronounced for strikes with a long duration, strikes in which a high percentage of the workforce participates, and strikes that lead to a high loss for the strike-hit firm. When I analyze the time pattern of the cash changes, the cash increase starts in the quarter of the experience, reaches a maximum four quarters thereafter, and then reverts.

In the long run, I find that CEOs tend to agree to higher wages in the years after their strike observation compared to other firms with contract settlements. At the same time, they manage to decrease the strike risk for their firm, especially after observing a severe strike. Agreeing to higher wages is one likely channel through which they reduce strike risk, but they may also adjust other aspects of labor relation management that I cannot measure. These findings indicate that strike-observing CEOs learn from insider information about labor negotiations.

The findings of the second essay contribute to the ongoing discussion on CEOs' outside directorships. The short-term cash increase of CEOs who observe a labor strike is consistent with the predictions from salience theory (e.g., Dessaint and Matray, 2017): CEOs temporarily overestimate strike risk when their attention is directed to this risk factor. This result implies that behavioral biases due to overreaction to salient risks can be a "dark side" of CEO outside directorship, besides the time and effort spent by CEOs (e.g., Fich and Shivdasani, 2006; Field, Lowry and Mkrtchyan, 2013). However, there is also a "bright side" of outside directorships. My empirical findings indicate that strike-observing CEOs can learn from observations they make at the director firm. Furthermore, the second essay adds to the literature on how managers' (past) experiences affect their own decision-making (e.g., Malmendier, Tate and

Yan, 2011) by showing that managers' observation of others' behavior affects their decision-making.

In the *third essay*, I analyze the role of firm size in within-firm and between-firm wage inequality. I use a novel linked employee-establishment-firm dataset with detailed information on individual workers and firms. It covers approximately 51,000 individual firms, 115,000 establishments, and 11.6 million workers from Germany.

I first find that firm size shows a strong positive correlation to both types of wage inequality. Second, I find that differences in workforce composition between small and large firms are responsible for 65% of the relationship between firm size and within-firm wage inequality. The corresponding figure for between-firm wage inequality is 60%. Therefore, the finding that larger firms pay more unequal and on average higher wages is strongly linked to differences in workforce composition between large and small firms.

Third, I try to explain why firm size has a positive impact on the different components of wage inequality after controlling for workforce composition. For within-firm wage inequality, I find that higher wage variance in larger firms after composition effects is largely explainable by higher employee monitoring complexity. Higher average wages in larger firms after composition effects (the so-called large-firm wage premium) are more challenging to explain. Neither differences in profitability nor monitoring complexity nor unionization levels can explain the large-firm wage premium. If anything, local labor markets play some role, but again they can also not fully explain the large-firm effect.

Overall, these results show that firms in general and their size in particular matter for wage inequality. However, the findings also suggest that there is little direct (or causal) link from firm size to wage inequality. Rather, differences in wage dispersion and average wages between small and large firms are (at least partly) explainable by factors that are correlated with firm size ("omitted variables"). These are workforce composition, job characteristics, monitoring complexity, and potentially clustering of larger firms in specific regions. Yet, even in the absence of a direct link from firm size to wage inequality, these results imply that wage inequality will increase if firms in an economy become larger. This is not restricted to the top end of the size distribution (so-called mega firms), but also holds true for size increases of smaller and more representative firms.

The results of the third essay contribute to several strands of the literature. Most importantly, I add to the literature which investigates how firms affect wage inequality (e.g., Card, Heining and Kline, 2013; Song et al., 2019). The findings complement the literature on higher dispersion of wages in larger firms (e.g., Mueller, Ouimet and Simintzi, 2017b) by showing that differences in the workforce composition can explain about 65% of differences in within-firm inequality between large and small firms, that within-firm inequality increases monotonously with firms size, and that differences in monitoring complexity can explain the relationship between firm size and within-firm inequality after composition effects. I add to the literature on higher average wages in larger firms (e.g, Bloom et al., 2018) by testing theoretical explanations that require detailed firm-level data. I do not find any evidence that rent-sharing, monitoring complexity, or unionization have a substantial impact on the large-firm wage premium.

To conclude, the three essays of this dissertation examine research questions on finance and labor. The findings motivate several avenues for future research. The first essay suggests that idiosyncratic shocks to very large firms can cause great harm to production networks. Given the trends towards superstar firms and rising market concentration, it seems relevant to identify systemically important firms outside the financial sector. The second essay provides empirical evidence that CEOs can learn from their observations as outside directors at other firms. Since more than one-third of CEOs in S&P 500 firms serve as directors at other firms, it is relevant—from a shareholder perspective—to further explore the benefits of CEOs' outside positions and whether these benefits can outweigh the distractions. Finally, the third essay sheds light on the role of firm size in within-firm and between-firm pay inequality. As the availability of linked employee-firm datasets is increasing, it is interesting to examine whether inequality can also affect firm-level outcomes.

References

Abowd, J. A., Lemieux, T., 1993. The Effects of Product Market Competition on Collective Bargaining Agreements: The Case of Foreign Competition in Canada. The Quarterly Journal of Economics 108 (4), 983–1014.

Abowd, J. M., Kramarz, F., Margolis, D. N., 1999. High Wage Workers and High Wage Firms. Econometrica 67 (2), 251–333.

Acemoglu, D., Azar, P. D., 2020. Endogenous Production Networks. Econometrica 88 (1), 33–82.

Almeida, H., Campello, M., Cunha, I., Weisbach, M. S., 2014. Corporate Liquidity Management: A Conceptual Framework and Survey. Annual Review of Financial Economics 6 (1), 135–162.

Alvarez, J., Benguria, F., Engbom, N., Moser, C., 2018. Firms and the Decline in Earnings Inequality in Brazil. American Economic Journal: Macroeconomics 10 (1), 149–189.

Antoni, M., Ganzer, A., vom Berge, P., 2016. Sample of integrated labour market biographies (SIAB) 1975-2014. IAB FDZ Datenreport 04/2016 EN.

Antoni, M., Koller, K., Laible, M.-C., Zimmermann, F., 2018. Orbis-ADIAB: From record linkage key to research dataset: Combining commercial company data with administrative employer-employee data. IAB FDZ Methodenreport 04/2018 EN.

Ashenfelter, O., Johnson, G. E., 1969. Bargaining Theory, Trade Unions, and Industrial Strike Activity. American Economic Review 59 (1), 35.

Atalay, E., Hortacsu, A., Roberts, J., Syverson, C., 2011. Network structure of production. Proceedings of the National Academy of Sciences 108 (13), 5199–5202.

Autor, D., Dorn, D., Katz, L. F., Patterson, C., Van Reenen, J., 2020. The Fall of the Labor Share and the Rise of Superstar Firms*. The Quarterly Journal of Economics.

Baqaee, D. R., 2018. Cascading Failures in Production Networks. Econometrica 86 (5), 1819–1838.

Baqaee, D. R., Farhi, E., 2019. The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem. Econometrica 87 (4), 1155–1203.

Barrot, J.-N., Sauvagnat, J., 2016. Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks *. The Quarterly Journal of Economics 131 (3), 1543–1592.

Becker, B. E., Olson, C. A., 1986. The Impact of Strikes on Shareholder Equity. Industrial and Labor Relations Review 39 (3), 425.

Benmelech, E., Bergman, N. K., Enriquez, R. J., 2012. Negotiating with Labor under Financial Distress. The Review of Corporate Finance Studies 1 (1), 28–67.

Bernile, G., Bhagwat, V., Rau, P. R., 2017. What Doesn't Kill You Will Only Make You More Risk- Loving: Early-Life Disasters and CEO Behavior. The Journal of Finance 72 (1), 167–206.

Bikhchandani, S., Hirshleifer, D., Welch, I., 1998. Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades. Journal of Economic Perspectives 12 (3), 151–170.

Bloom, N., Guvenen, F., Smith, B. S., Song, J., von Wachter, T., 2018. The Disappearing Large-Firm Wage Premium. AEA Papers and Proceedings 108, 317–22.

Blossfeld, H.-P., 1987. Labor-Market Entry and the Sexual Segregation of Careers in the Federal Republic of Germany. American Journal of Sociology 93 (1), 89–118.

Boehm, C. E., Flaaen, A., Pandalai-Nayar, N., 2019. Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 Tōhoku Earthquake. The Review of Economics and Statistics 101 (1), 60–75.

Boehm, J., Oberfield, E., 2018. Misallocation in the Market for Inputs: Enforcement and the Organization of Production. NBER Working Paper.

Booth, J. R., Deli, D. N., 1996. Factors affecting the number of outside directorships held by CEOs. Journal of Financial Economics 40 (1), 81–104.

Bordalo, P., Gennaioli, N., Shleifer, A., 2012. Salience Theory of Choice Under Risk. The Quarterly Journal of Economics 127 (3), 1243–1285.

Bordalo, P., Gennaioli, N., Shleifer, A., 2013. Salience and Asset Prices. American Economic Review 103 (3), 623–628.

Brown, C., Medoff, J., 1989. The Employer Size-Wage Effect. Journal of Political Economy 97 (5), 1027–1059.

Bryson, A., Erhel, C., Salibekyan, Z., 2019. Perceptions of non-pecuniary job quality using linked employer-employee data. European Journal of Industrial Relations - forthcoming.

Campello, M., Gao, J., 2017. Customer concentration and loan contract terms. Journal of Financial Economics 123 (1), 108–136.

Campello, M., Gao, J., Qiu, J., Zhang, Y., 2018. Bankruptcy and the Cost of Organized Labor: Evidence from Union Elections. The Review of Financial Studies 31 (3), 980–1013.

Card, D., 1990a. Strikes and Wages: A Test of an Asymmetric Information Model. The Quarterly Journal of Economics 105 (3), 625–659.

Card, D., 1990b. Strikes and Wages: A Test of an Asymmetric Information Model. The Quarterly Journal of Economics 105 (3), 625.

Card, D., 2011. The Return of the Firm to Labor Economics. Presentation at the 16th Annual Meeting of the Society of Labor Economists (SOLE 2011).

Card, D., Cardoso, A. R., Heining, J., Kline, P., 2018. Firms and Labor Market Inequality: Evidence and Some Theory. Journal of Labor Economics 36 (S1), S13–S70.

Card, D., Heining, J., Kline, P., 2013. Workplace Heterogeneity and the Rise of West German Wage Inequality*. The Quarterly Journal of Economics 128 (3), 967–1015.

Carpenter, M. A., Westphal, J. D., 2001. The strategic context of external network ties: Examining the impact of director appointments on board involvement in strategic decision making. Academy of Management Journal 44 (4), 639–660.

Carvalho, V., Gabaix, X., 2013. The Great Diversification and its Undoing. American Economic Review 103 (5), 1697–1727.

Carvalho, V., Voigtländer, N., 2015. Input Diffusion and the Evolution of Production Networks. Unpublished Working Paper.

Carvalho, V. M., Acemoglu, D., Ozdaglar, A., Tahbaz-Salehi, A., 2012. The Network Origins of Aggregate Fluctuations. Econometrica 80 (5), 1977–2016.

Carvalho, V. M., Grassi, B., 2019. Large Firm Dynamics and the Business Cycle. American Economic Review 109 (4), 1375–1425.

Carvalho, V. M., Nirei, M., Saito, Y. U., Tahbaz-Salehi, A., 2016. Supply Chain Disruptions: Evidence from the Great East Japan Earthquake. Unpublished Working Paper.

Chaney, T., 2014. The Network Structure of International Trade. American Economic Review 104 (11), 3600–3634.

Chava, S., Danis, A., Hsu, A., Feb. 2020. The Economic Impact of Right-to-Work Laws: Evidence from Collective Bargaining Agreements and Corporate Policies. Journal of Financial Economics.

Christofides, L. N., Oswald, A. J., 1992. Real Wage Determination and Rent-Sharing in Collective Bargaining Agreements. The Quarterly Journal of Economics 107 (3), 985–1002.

Chu, Y., 2012. Optimal capital structure, bargaining, and the supplier market structure. Journal of Financial Economics 106 (2), 411–426.

Cole, J., 2000. Boeing Board Encounters Dissent Over Firm's Handling of Strike. The Wall Street Journal (2000-04-14), https://www.wsj.com/articles/SB955665440468278747.

Cramton, P., Tracy, J., 1992. Strikes and Holdouts in Wage Bargaining: Theory and Data. American Economic Review 82 (1), 100–121.

Cronqvist, H., Heyman, F., Nilsson, M., Svaleryd, H., Vlachos, J., 2009. Do Entrenched Managers Pay Their Workers More? The Journal of Finance 64 (1), 309–339.

Dang, C., Li, Z. F., Yang, C., 2018. Measuring firm size in empirical corporate finance. Journal of Banking & Finance 86, 159–176.

Davis, S. J., Haltiwanger, J., Katz, L. F., Topel, R., 1991. Wage Dispersion between and within U.S. Manufacturing Plants, 1963-86. Brookings Papers on Economic Activity. Microeconomics 1991, 115.

De Fusco, R. A., Fuess, S. M., 1991. The Effects of Airline Strikes on Struck and Nonstruck Carriers. ILR Review 44 (2), 324–333.

DellaVigna, S., 2009. Psychology and Economics: Evidence from the Field. Journal of Economic Literature 47 (2), 315–372.

Dessaint, O., Matray, A., 2017. Do managers overreact to salient risks? Evidence from hurricane strikes. Journal of Financial Economics 126 (1), 97–121.

Dhaliwal, D., Judd, J. S., Serfling, M., Shaikh, S., 2016. Customer concentration risk and the cost of equity capital. Journal of Accounting and Economics 61 (1), 23–48.

di Giovanni, J., Levchenko, A. A., 2012. Country Size, International Trade, and Aggregate Fluctuations in Granular Economies. Journal of Political Economy 120 (6), 1083–1132.

di Giovanni, J., Levchenko, A. A., Mejean, I., 2014. Firms, Destinations, and Aggregate Fluctuations. Econometrica 82 (4), 1303–1340.

di Giovanni, J., Levchenko, A. A., Mejean, I., 2018. The Micro Origins of International Business-Cycle Comovement. American Economic Review 108 (1), 82–108.

Dickens, W. T., Lang, K., 1985. A Test of Dual Labor Market Theory. American Economic Review 75 (4), 792–805.

Dinardo, J., Hallock, K. F., 2002. When Unions "Mattered": The Impact of Strikes on Financial Markets, 1925–1937. ILR Review 55 (2), 219–233.

DiNardo, J., Lee, D. S., 2004. Economic Impacts of New Unionization on Private Sector Employers: 1984-2001. The Quarterly Journal of Economics 119 (4), 1383–1441.

Dittmar, A., Duchin, R., 2015. Looking in the Rearview Mirror: The Effect of Managers' Professional Experience on Corporate Financial Policy. The Review of Financial Studies 29 (3), 565–602.

Dustmann, C., Ludsteck, J., Schönberg, U., 2009. Revisiting the German Wage Structure *. The Quarterly Journal of Economics 124 (2), 843–881.

Eaton, C., White, W. D., 1983. The Economy of High Wages: An Agency Problem. Economica 50 (198), 175.

Ellul, A., Pagano, M., Schivardi, F., 2018. Employment and wage insurance within firms: Worldwide evidence. The Review of Financial Studies 31 (4), 1298–1340.

Faggio, G., Salvanes, K. G., Van Reenen, J., 2010. The evolution of inequality in productivity and wages: panel data evidence. Industrial and Corporate Change 19 (6), 1919–1951.

Fahlenbrach, R., Kim, H., Low, A., 2018. The Importance of Network Recommendations in the Director Labor Market. Unpublished Working Paper.

Fahlenbrach, R., Low, A., Stulz, R. M., 2010. Why do firms appoint CEOs as outside directors? Journal of Financial Economics 97 (1), 12–32.

Fich, E. M., Shivdasani, A., 2006. Are Busy Boards Effective Monitors? The Journal of Finance 61 (2), 689–724.

Field, L., Lowry, M., Mkrtchyan, A., 2013. Are busy boards detrimental? Journal of Financial Economics 109 (1), 63–82.

Fitzenberger, B., Osikominu, A., Völter, R., 2006. Imputation Rules to Improve the Education Variable in the IAB Employment Subsample. Schmollers Jahrbuch : Journal of Applied Social Science Studies / Zeitschrift für Wirtschafts- und Sozialwissenschaften 126 (3), 405–436.

Gabaix, X., 2011. The Granular Origins of Aggregate Fluctuations. Econometrica 79 (3), 733–772.

Garen, J. E., 1985. Worker Heterogeneity, Job Screening, and Firm Size. Journal of Political Economy 93 (4), 715–739.

Giroud, X., 2013. Proximity and Investment: Evidence from Plant-Level Data *. The Quarterly Journal of Economics 128 (2), 861–915.

Goldschmidt, D., Schmieder, J. F., 2017. The Rise of Domestic Outsourcing and the Evolution of the German Wage Structure*. The Quarterly Journal of Economics 132 (3), 1165–1217.

Grassi, B., 2017. IO in I-O: Size, Industrial Organization, and the Input-Output Network Make a Firm Structurally Important. Unpublished Working Paper.

Gruber, J., Kleiner, S. A., 2012. Do strikes kill? Evidence from New York State. American Economic Journal: Economic Policy 4 (1), 127–157.

Grullon, G., Larkin, Y., Michaely, R., 2019. Are US Industries Becoming More Concentrated?*. Review of Finance 23 (4), 697–743.

Hamermesh, D. S., 1980. Substitution and Labor Market Policy. Challenge 22 (6), 44–47.

Hart, O., 1989. Bargaining and Strikes. The Quarterly Journal of Economics 104 (1), 25.

Hayes, B., 1984. Unions and Strikes with Asymmetric Information. Journal of Labor Economics 2 (1), 57–83.

Herskovic, B., Kelly, B., Lustig, H., Van Nieuwerburgh, S., 2018. Firm Volatility in Granular Networks. Unpublished Working Paper.

Hicks, J. R., 1963. The Theory of Wages. Palgrave Macmillan UK, London.

Hirsch, B. T., Macpherson, D. A., 2003. Union membership and coverage database from the Current Population Survey (CPS): Note. Industrial and Labor Relations Review 56 (2), 349–354.

Holmes, T., 2006. Geographic Spillover of Unionism. Unpublished Working Paper.

Houston, J. F., Lin, C., Zhu, Z., 2016. The Financial Implications of Supply Chain Changes. Management Science 62 (9), 2520–2542.

Imberman, W., 1983. Who strikes-and Why? Harvard Business Review 61 (6), 18–28.

Inoue, H., Todo, Y., 2019. Firm-level propagation of shocks through supply-chain networks. Nature Sustainability 2 (9), 841–847.

Katz, L. F., Autor, D. H., 1999. Chapter 26 Changes in the wage structure and earnings inequality. Handbook of Labor Economics 3 PART (1), 1463–1555.

Kim, E. H., Ouimet, P. P., 2014. Broad-Based Employee Stock Ownership: Motives and Outcomes. The Journal of Finance 69 (3), 1273–1319.

Klasa, S., Maxwell, W. F., Ortiz-Molina, H., 2009. The strategic use of corporate cash holdings in collective bargaining with labor unions. Journal of Financial Economics 92 (3), 421–442.

König, M. D., 2018. Aggregate Fluctuations in Adaptive Production Networks. Unpublished Working Paper.

Krueger, A., Mas, A., 2004. Strikes, Scabs, and Tread Separations: Labor Strife and the Production of Defective Bridgestone/Firestone Tires. Journal of Political Economy 112 (2), 253–289.

Krueger, A. B., Summers, L. H., 1988. Efficiency Wages and the Inter-Industry Wage Structure. Econometrica 56 (2), 259–293.

Lallemand, T., Plasman, R., Rycx, F., 2005. Why do large firms pay higher wages? Evidence from matched worker-firm data. International Journal of Manpower 26 (7-8), 705–723.

Lim, K., 2018. Endogenous Production Networks and the Business Cycle. Unpublished Working Paper.

Malmendier, U., 2018. Behavioral Corporate Finance. In: Bernheim, B. D., DellaVigna, S., Laibson, D. (Eds.), Handbook of Behavioral Economics: Applications and Foundations 1. Vol. 1. Norh Holland, pp. 277–379.

Malmendier, U., Nagel, S., 2011. Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?*. The Quarterly Journal of Economics 126 (1), 373–416.

Malmendier, U., Tate, G., Yan, J., 2011. Overconfidence and Early-Life Experiences: The Effect of Managerial Traits on Corporate Financial Policies. The Journal of Finance 66 (5), 1687–1733.

Mas, A., 2008. Labour Unrest and the Quality of Production: Evidence from the Construction Equipment Resale Market. The Review of Economic Studies 75 (1), 229–258.

Matsa, D. A., 2010. Capital Structure as a Strategic Variable: Evidence from Collective Bargaining. The Journal of Finance 65 (3), 1197–1232.

Matsa, D. A., 2018. Capital Structure and a Firm's Workforce. Annual Review of Financial Economics 10 (1), 387–412.

McHugh, R., 1991. Productivity Effects of Strikes in Struck and Nonstruck Industries. Industrial and Labor Relations Review 44 (4), 722.

Mueller, H. M., Ouimet, P. P., Simintzi, E., 2017a. Wage Inequality and Firm Growth. American Economic Review 107 (5), 379–383.

Mueller, H. M., Ouimet, P. P., Simintzi, E., 2017b. Within-Firm Pay Inequality. The Review of Financial Studies 30 (10), 3605–3635.

Myers, B. W., Saretto, A., 2016. Does Capital Structure Affect the Behavior of Nonfinancial Stakeholders? An Empirical Investigation into Leverage and Union Strikes. Management Science 62 (11), 3235–3253.

Obama, B., 2013. Remarks on the Economy. The Washington Post (2013, December 4), http://wapo.st/IEiroQ.

Oberfield, E., 2018. A Theory of Input-Output Architecture. Econometrica 86 (2), 559–589.

Oi, W. Y., 1983. Heterogeneous Firms and the Organization of Production. Economic Inquiry 21 (2), 147–171.

Oi, W. Y., Idson, T. L., 1999. Chapter 33 Firm size and wages. In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics, 1st Edition. Vol. 3. North-Holland, pp. 2165–2214.

Ouimet, P., Tate, G., 2020. Learning from Coworkers: Peer Effects on Individual Investment Decisions. The Journal of Finance 75 (1), 133–172.

Pagano, M., Volpin, P., 2005. Managers, Workers, and Corporate Control. The Journal of Finance 60 (2), 841–868.

Pagano, M., Volpin, P. F., 2008. Labor and Finance. Unpublished Working Paper.

Perry, T., Peyer, U., 2005. Board Seat Accumulation by Executives: A Shareholder's Perspective. The Journal of Finance 60 (4), 2083–2123.

Persons, O. S., 1995. The Effects of Automobile Strikes on the Stock Value of Steel Suppliers. Industrial and Labor Relations Review 49 (1), 78.

Qiu, J., Zhang, Y., 2019. Wage Rigidity and Debt Financing: Evidence From Labor Contract Renewal During the Financial Crisis. Unpublished Working Paper.

Schmidt, C. M., Zimmermann, K. F., 1991. Work Characteristics, Firm Size and Wages. The Review of Economics and Statistics 73 (4), 705–710.

Schoar, A., Zuo, L., 2017. Shaped by Booms and Busts: How the Economy Impacts CEO Careers and Management Styles. The Review of Financial Studies 30 (5), 1425–1456.

Shahrur, H., 2005. Industry structure and horizontal takeovers: Analysis of wealth effects on rivals, suppliers, and corporate customers. Journal of Financial Economics 76 (1), 61–98.

Shapiro, C., Stiglitz, J. E., 1984. Equilibrium Unemployment as a Worker Discipline Device. American Economic Review 74 (3), 433–444.

Shue, K., 2013. Executive Networks and Firm Policies: Evidence from the Random Assignment of MBA Peers. The Review of Financial Studies 26 (6), 1401–1442.

Song, J., Price, D. J., Guvenen, F., Bloom, N., von Wachter, T., 2019. Firming Up Inequality*. The Quarterly Journal of Economics 134 (1), 1–50.

Taschereau-Dumouchel, M., 2019. Cascades and Fluctuations in an Economy with an Endogenous Production Network. Unpublished Working Paper.

Tate, G., Yang, L., 2015. The Bright Side of Corporate Diversification: Evidence from Internal Labor Markets. The Review of Financial Studies 28 (8), 2203–2249.

Tintelnot, F., Kikkawa, A. K., Mogstad, M., Dhyne, E., 2018. Trade and Domestic Production Networks. NBER Working Paper.

Titman, S., 1984. The effect of capital structure on a firm's liquidation decision. Journal of Financial Economics 13 (1), 137–151.

Titman, S., Wessels, R., 1988. The Determinants of Capital Structure Choice. The Journal of Finance 43 (1), 1–19.

Tversky, A., Kahneman, D., 1973. Availability: A heuristic for judging frequency and probability. Cognitive Psychology 5 (2), 207–232.

Tversky, A., Kahneman, D., 1974. Judgment under Uncertainty: Heuristics and Biases. Science 185 (4157), 1124–1131.

Yi, I., 2019. Slashing Liquidity Through Asset Purchases: Evidence from Collective Bargaining. Unpublished Working Paper.

A
Chapter 1

APPENDIX A

Table A.1: Definition of variables

Variable	Description
Main variables	
strike hits firm	Dummy indicating whether a firm is hit by a labor strike. Source: Own calculation.
strike hits customer	Dummy indicating whether one of a firm's major customers is hit by a labor strike. Source: Own calculation.
strike hits central customer	Dummy indicating whether one of a firm's major customers, which is also a customer of the firm's other customers, is hit by a labor strike. Source: Own calculation.
sales growth	Sales growth relative to same quarter in the previous year ($\Delta ln(sale_{t,t-4})$). Source: Compustat (CS).
Strike characteristics	
duration	Number of days that the labor strike lasts. Source: BLS, CT, NMB and CS.
striking employees	Number of employees participating in labor strike. Source: Bureau of Labor Statistics (BLS), Cramton and Tracy (1992) (CT), Federal Mediation and Conciliation Service (FMCS) and National Mediation Board (NMB).
striking emp/total emp	Number of employees participating in labor strike divided by the strike-hit firm's total number of employees (emp). Source: BLS, CT, FMCS, NMB and CS.
idled employee-days	Number of employees idling during labor strike multiplied by duration of the labor strike. Source: BLS, CT, FMCS and NMB.
idled emp-days/total emp-days	Idled employee-days relative to the employee-days in a quarter. If this ratio lies above one (a firm's entire workforce is on strike for more than one quarter), it is set to one. Source: BLS, CT, FMCS and NMB.
Firm-level control variables	
number of customers	Number of major customers that a firm reports in the Compustat Segement Database. The variable is lagged by one year. Source: CS.
%sales with customers	Fraction of a firm's sales that the major customers account for. The variable is lagged by one year. Source: CS.
%sales with largest customer	Fraction of a firm's sales that the largest customer accounts for. The variable is lagged by one year. Source: CS.
central customer	Dummy variable indicating whether a firm has a central customer that is the firm's major customer (direct link) and also a customer of the firm's other customers (indirect link). The variable is lagged by one year. Source: CS
number of suppliers	Number of suppliers that that report the firm as a major customer in the Compustat Segement Database. The variable is lagged by one year. Source: CS.
real size	Natural logarithm of total assets ($ln(at)$) in m of 2012 USD. The variable is lagged by one year. Source: CS.

continued on next page

Appendix A.1 continued

Variable	Description
age	Number of years since a firm's initial public offering (IPO). If the IPO date is missing, age is defined as the number of years since the firm has been in the Compustat database. The variable is lagged by one year. Source: CS.
roa	Operating income after depreciation divided by total assets ($\frac{oiadp}{at}$). The variable is lagged by one year. Source: CS.

CS stands for Compustat, BLS for Bureau of Labor Statistics, CT for labor contract data from Cramton and Tracy (1992), FMCS for Federal Mediation and Conciliation Service, and NMB for National Mediation Board.

Table A.2: Strikes and customers' sales growth: supplier characteristics

This table presents regression estimates of firms' sales growth relative to the same quarter in the previous year on a dummy indicating whether a firm is hit by a labor strike in the current or in the previous quarter interacted with the firm's number of suppliers (constructed as terciles of the lagged value, Columns 1 and 2), and interacted with the firm's suppliers' dependence (constructed as terciles of the lagged mean value of suppliers' fraction of sales with the firm, Columns 3 and 4). All regressions include firm and year-quarter fixed effects. In Columns 2 and 4, I include controls for firm-level characteristics (constructed as terciles of lagged size, age, and return on assets, respectively) interacted with year-quarter dummies, industry dummies using two-digit SIC codes interacted with year dummies, and state dummies interacted with year dummies. Regressions contain all firm-quarters of the customer sample between 1983 and 2013. T-statistics presented in parentheses are based on robust standard errors clustered at the firm level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix A.1.

	(1)	(2)	(3)	(4)
supplier characteristics	number		%sales	
strike hits firm$_{t,t-1}$	-0.057***	-0.043***	-0.061**	-0.043*
	(-3.14)	(-2.65)	(-2.23)	(-1.92)
strike hits firm$_{t,t-1}$ x tercile 2	0.055	0.040	0.033	0.035
	(1.09)	(0.97)	(1.00)	(1.24)
strike hits firm$_{t,t-1}$ x tercile 3	0.022	0.0093	0.032	0.0029
	(0.88)	(0.39)	(0.99)	(0.10)
tercile 2	-0.018**	-0.012	-0.016***	-0.00037
	(-2.48)	(-1.47)	(-2.61)	(-0.049)
tercile 3	-0.032***	-0.020**	-0.015**	0.0043
	(-3.88)	(-2.04)	(-2.11)	(0.55)
Firm FE	Yes	Yes	Yes	Yes
Year-quarter FE	Yes	Yes	Yes	Yes
Firm-level controls	No	Yes	No	Yes
Industry-year FE	No	Yes	No	Yes
State-year FE	No	Yes	No	Yes
Obs	31,114	31,107	31,114	31,107
Number - firms	578	578	578	578
R2	0.142	0.402	0.141	0.401

Table A.3: Upstream propagation: strike severity

This table presents regression estimates of firms' sales growth relative to the same quarter in the previous year on continuous variables measuring the severity of the labor strike that hits one of their major customers in the current or in the previous quarter. The strike severity is measured by the the logarithm of the strike duration in days (Column 1), the ratio of idling employees to the customer's total number of employees (Column 2), the logarithm of the idled employee-days (Column 3), and the idled employee-days relative to the employee-days in a quarter (Column 4; if this ratio lies above one, it is set to one). All regressions include firm and year-quarter fixed effects, controls for firms' lagged number of major customers (constructed as terciles) as well as firm-level characteristics (constructed as terciles of lagged size, age, and return on assets, respectively) interacted with year-quarter dummies, industry dummies using two-digit SIC codes interacted with year dummies, and state dummies interacted with year dummies. Regressions contain all firm-quarters of the supplier sample between 1983 and 2013. T-statistics presented in parentheses are based on robust standard errors clustered at the firm level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix A.1.

	(1)	(2)	(3)	(4)
strike hits customer$_{t,t-1,duration}$	-0.010***			
	(-3.07)			
strike hits customer$_{t,t-1,\%idled}$		-0.028		
		(-0.66)		
strike hits customer$_{t,t-1,idled-days}$			-0.0022***	
			(-2.74)	
strike hits customer$_{t,t-1,\%idled-days}$				-0.073***
				(-3.10)
Firm FE	Yes	Yes	Yes	Yes
Year-quarter FE	Yes	Yes	Yes	Yes
Number of customers	Yes	Yes	Yes	Yes
Firm-level controls	Yes	Yes	Yes	Yes
Industry-year FE	Yes	Yes	Yes	Yes
State-year FE	Yes	Yes	Yes	Yes
Obs	132,549	132,549	132,549	132,549
Number - firms	4092	4092	4092	4092
R2	0.232	0.232	0.232	0.232

Table A.4: Upstream propagation: seasonal trends

This table presents regression estimates of firms' sales growth relative to the same quarter in the previous year on a dummy indicating whether one of their major customers is hit by a labor strike in the current or the previous quarter. In Panel A, I include firm dummies interacted with fiscal-quarter dummies to control for firm-specific seasonal trends. In Panel B, I limit the sample to firms whose fiscal quarters match the calendar quarters. All regressions include firm and year-quarter fixed effects as well as the firms' lagged number of major customers (constructed as terciles). In Column 2, I control for firm-level characteristics (constructed as terciles of lagged size, age, and return on assets, respectively) interacted with year-quarter dummies. In Column 3, I include industry dummies using two-digit SIC codes interacted with year dummies. In Column 4, I include state dummies interacted with year dummies. Regressions contain all firm-quarters of the supplier sample between 1983 and 2013. T-statistics presented in parentheses are based on robust standard errors clustered at the firm level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix A.1.

	(1)	(2)	(3)	(4)
Panel A: Firm-fiscal-quarter fixed effects				
strike hits customer$_{t,t-1}$	-0.028***	-0.018*	-0.024**	-0.026***
	(-3.02)	(-1.96)	(-2.44)	(-2.59)
Firm-fiscal-quarter FE	Yes	Yes	Yes	Yes
Year-quarter FE	Yes	Yes	Yes	Yes
Number of customers	Yes	Yes	Yes	Yes
Firm-level controls	No	Yes	Yes	Yes
Industry-year FE	No	No	Yes	Yes
State-year FE	No	No	No	Yes
Obs	131,790	131,788	131,780	131,775
Number - firms	3959	3959	3958	3957
R2	0.168	0.192	0.229	0.247
Panel B: Fiscal quarter matches calendar quarter				
strike hits customer$_{t,t-1}$	-0.040***	-0.026**	-0.033***	-0.029**
	(-3.61)	(-2.20)	(-2.78)	(-2.37)
Firm FE	Yes	Yes	Yes	Yes
Year-quarter FE	Yes	Yes	Yes	Yes
Number of customers	Yes	Yes	Yes	Yes
Firm-level controls	No	Yes	Yes	Yes
Industry-year FE	No	No	Yes	Yes
State-year FE	No	No	No	Yes
Obs	81,779	81,775	81,764	81,760
Number - firms	2676	2676	2675	2675
R2	0.169	0.197	0.244	0.269

B
Chapter 2

APPENDIX B

Table B.1: Definition of variables

Variable	Description
Main variables	
strike observation	Dummy which equals one if a CEO observes a labor strike as a director at another firm. Source: Own calculation.
strike severity$_{duration}$	Logarithm of the strike duration as a measure for the severity of the observed strike. Source: Own calculation.
strike severity$_{\%idled}$	Ratio of striking employees to total employees of the strike-hit firm as a measure for the severity of the observed strike. Source: Own calculation.
strike severity$_{idled-days}$	Logarithm of the idled employee-days as a measure for the severity of the observed strike. Source: Own calculation.
strike severity$_{\Delta roa}$	Absolute value of the decline in return on assets in the first quarter of the strike compared to the same of quarter of the last year (if the change of return on assets is positive it is set to zero) as a measure for the severity of the observed strike. Source: Own Calculation, Compustat (CS).
strike severity$_{\Delta ocf}$	Absolute value of the decline in operating cash flow (standardized by total assets) in the first quarter of the strike compared to the same of quarter of the last year (if the change of operating cash flow is positive it is set to zero) as a measure for the severity of the observed strike. Source: Own Calculation, CS.
cash	Cash and cash equivalents (cheq) scaled by total assets (atq). Source: Compustat (CS).
Δwage	Change of wages, measured in percentage points, in the first year of the new labor contract. If we observe more than one labor contract settlement in the firm-year, we report the mean value of all settlements. Source: Bloomberg BNA.
strike dummy	Dummy which equals one if a labor strike begins at the firm in the respective fiscal year (y). Source: Own calculation.
Strike characteristics	
strike duration	Number of days that the labor strike lasts. Source: BLS, CT, NMB and CS.
striking employees	Number of employees participating in labor strike. Source: Bureau of Labor Statistics (BLS), Cramton and Tracy (1992) (CT), Federal Mediation and Conciliation Service (FMCS) and National Mediation Board (NMB).
striking emp/total emp	Number of employees participating in labor strike divided by the strike-hit firm's total number of employees (emp). Source: BLS, CT, FMCS, NMB and CS.
idled employee-days	Number of employees idling during labor strike times duration of the labor strike. Source: BLS, CT, FMCS and NMB.
$\Delta roa_{t0,t-4}$	Difference in roa between the first quarter of strike and the same quarter one year before. Source: CS.

continued on next page

Appendix B.1 continued

Variable	Description
$\Delta ocf_{t0,t-4}$	Difference in operating cash flow (standardized by total assets) between the first quarter of the strike and the same quarter one year before. Source: CS.
Financial variables	
real size	Natural logarithm of total assets (atq) in m of 2012 USD. Source: CS.
market leverage	Total debt divided by total debt plus market capitalization (prccq * cshoq). Total debt includes current and long-term liabilities (dlcq + dlttq). Source: CS.
roa	Operating income after depreciation divided by total assets ($\frac{oiadpq}{atq}$). Source: CS.
Labor-related variables	
no. employees	Natural logarithm of number of workers. Source: CS.
ind. unionization	Average fraction of industry workforce covered by collective bargaining in % in the calendar year. Source: Unionstats of Barry Hirsch and David Macpherson (www.unionstats.com).
ind. #strikes	Number of strikes that occured in the firm's Fama-French 12 industry in the calendar year. Source: Own calculation.
other ind. #strikes	Number of strikes that occured outside the firm's Fama-French 12 industry in the calendar year. Source: Own calculation.
#strikes in state	Number of strikes that occured in the state of the firm's headquarters in the calendar year. Source: Own calculation.
right-to-work law	Dummy which equals one if the headquarters are located in a state with right-to-work law in place. Source: National Conference of State Legislatures.
Settlement of labor contracts	
settlement dummy	Dummy which equals one if we observe the settlement of at least one labor contract during the respective firm-year. Source: Bloomberg BNA.
#settlements	Number of settled labor contracts. Source: Bloomberg BNA.
workers under settlement	Number of workers under the labor contract. If we observe more than one labor contract settlement in the firm-year, we report the mean value of all settlements. Source: Bloomberg BNA.
contract duration	Duration of the new labor contract in days. If we observe more than one labor contract settlement in the firm-year, we report the mean value of all settlements. Source: Bloomberg BNA.

CS stands for Compustat, BLS for Bureau of Labor Statistics, CT for labor contract data from Cramton and Tracy (1992), FMCS for Federal Mediation and Conciliation Service, NMB for National Mediation Board, BX for Boardex and BNA for Bureau of National Affairs.

Table B.2: Robustness: time dynamics

This table presents estimates from regressions of cash holdings on time dummies from four quarters before to seven quarters after a CEO's strike observation (Columns 1 and 2), and alternatively the duration of a strike as measure for its severity (Columns 3 and 4). We use the same fixed effects and control variables as in Table 2.3. For treated firms, we include the eight quarters before the CEO strike observation (q-8 to q-1) as the pre-period and the eight quarters afterwards as post-period (q0 to q7). The base period is from eight to five quarters before the strike observation. Figure 2.3 illustrates the coefficient estimates of Column 1. T-statistics based on Huber/White robust standard errors clustered by firms are presented in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix B.1.

	(1)	(2)	(3)	(4)
	dummy	dummy	duration	duration
strike observation/severity$_{q-4}$	0.0035	0.0026	0.0011*	0.00097
	(1.44)	(1.06)	(1.65)	(1.45)
strike observation/severity$_{q-3}$	0.00026	-0.0014	-0.00011	-0.00041
	(0.10)	(-0.52)	(-0.16)	(-0.60)
strike observation/severity$_{q-2}$	0.00016	-0.00060	-0.000023	-0.00016
	(0.051)	(-0.19)	(-0.028)	(-0.20)
strike observation/severity$_{q-1}$	0.00096	0.000098	0.00042	0.00033
	(0.27)	(0.027)	(0.43)	(0.34)
strike observation/severity$_{q0}$	0.0043	0.0037	0.0011	0.0011
	(1.13)	(1.02)	(1.15)	(1.17)
strike observation/severity$_{q+1}$	0.0059	0.0052	0.0015	0.0015
	(1.49)	(1.35)	(1.48)	(1.52)
strike observation/severity$_{q+2}$	0.0074*	0.0069	0.0021*	0.0021*
	(1.70)	(1.62)	(1.78)	(1.83)
strike observation/severity$_{q+3}$	0.0084*	0.0074*	0.0031**	0.0030**
	(1.87)	(1.69)	(2.43)	(2.45)
strike observation/severity$_{q+4}$	0.012**	0.012***	0.0040***	0.0043***
	(2.53)	(2.75)	(3.08)	(3.53)
strike observation/severity$_{q+5}$	0.0096**	0.0098**	0.0030**	0.0032***
	(2.00)	(2.12)	(2.41)	(2.76)
strike observation/severity$_{q+6}$	0.0060	0.0065	0.0023*	0.0024**
	(1.31)	(1.45)	(1.79)	(2.00)
strike observation/severity$_{q+7}$	0.0046	0.0055	0.0020	0.0024*
	(0.93)	(1.16)	(1.50)	(1.86)
Controls	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes
Year-quarter FE	Yes	Yes	Yes	Yes
Industry-year FE	Yes	Yes	Yes	Yes
Industry-quarter FE	Yes	Yes	Yes	Yes
Obs	11,504	11,162	11,504	11,162
Number - firms	281	280	281	280
Number - events	215	213	215	213
R2	0.633	0.660	0.634	0.661

Table B.3: Robustenss: strike observation dummy and the long-term effect on wage changes

This table presents estimates from regressions of the wage change in the first year of new labor contracts (measured in percentage points) on dummies indicating a CEO strike observation in the last two years (Column 1), four years (Column 2), six years (Column 3), and eight years (Column 4). All models include firm, year and industry-year fixed effects. The sample is constructed on the firm-year level and includes firms for which we observe the settlement of at least one labor contract. T-statistics based on Huber/White robust standard errors clustered by firms are presented in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix B.1.

	(1)	(2)	(3)	(4)
	2 years	4 years	6 years	8 years
strike observation$_{longterm}$	1.10**	0.97**	0.76**	0.69**
	(2.52)	(2.35)	(1.98)	(1.99)
lag of no. employees	0.77	0.58	0.41	0.32
	(1.47)	(1.17)	(0.95)	(0.79)
lag of market leverage	-0.93	-1.06	-0.71	-0.45
	(-0.65)	(-0.79)	(-0.57)	(-0.40)
lag of roa	-3.37	-2.56	-1.53	-1.03
	(-0.78)	(-0.63)	(-0.39)	(-0.29)
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Industry-year FE	Yes	Yes	Yes	Yes
Obs	919	968	1,013	1,072
Number - firms	216	229	234	241
Number - events	28	36	39	40
R2	0.612	0.606	0.596	0.588

Table B.4: Robustness: strike observation dummy and the long-term effect on strike risk

This table presents estimates from regressions of a strike dummy on dummies indicating a CEO strike observation in the last two years (Column 1), four years (Column 2), six years (Column 3), and eight years (Column 4). The dependent variable is a strike dummy that is set to one if a labor strike begins at a firm in the respective fiscal year and zero otherwise. The sample is constructed on the firm-year level and consists of firms that are hit by at least one strike during the sample period. T-statistics based on Huber/White robust standard errors clustered by firms are presented in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix B.1.

	(1)	(2)	(3)	(4)
	2 years	4 years	6 years	8 years
strike observation$_{longterm}$	-0.015	-0.23	-0.35	-0.38
	(-0.051)	(-0.89)	(-1.44)	(-1.63)
lag of strike dummy	-0.27*	-0.27*	-0.28**	-0.31**
	(-1.81)	(-1.87)	(-1.96)	(-2.09)
lag of ind. #strikes	0.30**	0.30**	0.30**	0.32**
	(2.19)	(2.26)	(2.33)	(2.46)
lag of other ind. #strikes	0.067	0.088	0.083	0.13
	(0.18)	(0.24)	(0.22)	(0.35)
lag of #strikes in state	0.073	0.075	0.085	0.073
	(0.58)	(0.62)	(0.72)	(0.62)
lag of ind. unionization	0.16	0.13	-0.16	0.031
	(0.17)	(0.15)	(-0.18)	(0.037)
right-to-work law	0.61	0.55	0.63	0.60
	(1.34)	(1.25)	(1.47)	(1.41)
lag of no. employees	0.67***	0.66***	0.70***	0.70***
	(3.67)	(3.71)	(3.99)	(4.08)
lag of roa	0.13	-0.12	-0.098	0.048
	(0.11)	(-0.096)	(-0.083)	(0.041)
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Obs	4,134	4,277	4,433	4,568
Number - firms	247	248	252	254
Number - events	70	72	75	77
R2	0.052	0.051	0.056	0.057

C
Chapter 3

APPENDIX C

Table C.1: Definition of variables

Variable	Description
Wage and AKM componentens	
log wage	Imputed real log daily wage. Source: IEB (Integrated Employment Biographies).
person FE	Person fixed effect from the AKM-type regression. The implementation and interpretation of the AKM-type regression is explained in detail in Section 3.3.3.
establishment FE	Establishment fixed effect from the AKM-type regression. The implementation and interpretation of the AKM-type regression is explained in detail in Section 3.3.3.
firm FE	Employee-weighted average of a firm's establishment fixed effects from the AKM-type regression.
Xb	Combination of life cycle and aggregate factors from the AKM-type regression. The implementation and interpretation of the AKM-type regression is explained in detail in Section 3.3.3.
residual	Residual from the AKM-type regression. The implementation and interpretation of the AKM-type regression is explained in detail in Section 3.3.3.
Firm/establishment size	
log(total assets)	Natural logarithm of total assets (toas) in m of 2013 EUR. Source: Amadeus.
log($employees_{firm}$)	Natural logarithm of the firm's total employees in Germany. Source: BHP (Betriebshistorik Panel), IAB.
log($employees_{world}$)	Natural logarithm of the firm's worldwide employees. Source: Amadeus.
log($employees_{estab}$)	Natural logarithm of the establishment's employees. Source: BHP, IEB.
Other firm characteristics	
sales/employees	Total sales standardized by total number of worldwide employees in m of 2013 EUR per employee ($\frac{sales}{employees_{world}}$). Source: Amadeus.
profit/employees	Earnings before interest, taxes, depreciation, and amortization standardized by total number of worldwide employees in m of 2013 EUR per employee ($\frac{ebta}{employees_{world}}$). Source: Amadeus.
profit/total assets	Earnings before interest, taxes, depreciation, and amortization standardized by total assets ($\frac{ebta}{toas}$). Source: Amadeus.
ind. $unionization_{empl}$	Fraction of unionized employees in the industry sector (WZ2008) in 2014. Source: Federal Statistical Office.
ind. $unionization_{estab}$	Fraction of unionized establishments in the industry sector (WZ2008) in 2014. Source: Federal Statistical Office.
no. establishments	Firm's number of establishments. Source: Oribs-ADIAB.
median(distance)	Median of logarithm of distance (in kilometers) from county of establishment to county of firm headquarters (ln(distance + 1)). If establishment and headquarters are located in same county, the distance is zero. Source: Own calculation.

continued on next page

Appendix C.1 continued

Variable	Description
sd(distance)	Standard deviation of logarithm of distance (in kilometers) from county of establishment to county of firm headquarters (ln(distance + 1)). If all establishment are in same county as headquarters or the firm has only one establishment, it set to zero. Source: Own calculation.
manager/employees	Number of employed managers (according to (Blossfeld, 1987)) divided by total employees. Source: BHP, IEB.
semi-skilled/employees	Number of (full-time and part-time) employees who perform unskilled or semi-skilled tasks divided by total (full-time and part-time) employees. This variable is based on the last digit of the 5-digit Classification of Occupations 2010 and is available from 2011 on. Source: BHP, IEB.
skilled/employees	Number of (full-time and part-time) employees who perform skilled tasks divided by total (full-time and part-time) employees. This variable is based on the last digit of the 5-digit Classification of Occupations 2010 and is available from 2011 on. Source: BHP, IEB.
complex/employees	Number of (full-time and part-time) employees who perform complex tasks divided by (full-time and part-time) total employees. This variable is based on the last digit of the 5-digit Classification of Occupations 2010 and is available from 2011 on. Source: BHP, IEB.
highly complex/employees	Number of (full-time and part-time) employees who perform highly complex tasks divided by (full-time and part-time) total employees. This variable is based on the last digit of the 5-digit Classification of Occupations 2010 and is available from 2011 on. Source: BHP, IEB.
no. occupations	Number of occupations according to the first three digits of the Classification of Occupations 1988 (KldB 1988). Source: IEB.
hhi(occupations)	Herfindahl index as a concentration measure of employees' occupations according to the first three digits of the Classification of Occupations 1988 (KldB 1988). Source: IEB.
Establishment and firm characteristics used to construct fixed effects	
county_{firm}	County ("Landkreis") of firm's headquarters. Source: Amadeus.
county_{estab}	County ("Landkreis") of establishment. Source: BHP.
industry_{firm}	3-digits of the classifcation of economic activities WZ 2008. Source: Amadeus.
industry_{estab}	3-digits of the classifcation of economic activities WZ 2008. Source: BHP.

BHP stands for Betriebshistorik Panel provided by the Institute of Employment Research, IEB for Integrated Employment Biographies provided by the Institute of Employment Research, and Amadeus for the Amadeus database by Bureau van Dijk.

Table C.2: Descriptive statistics on the establishment level

This table presents descriptive statistics on the establishment-level. Reported are the number of observations (Obs), mean value (Mean), standard deviation (SD), median (p50), mean weighted by establishments' number of employees (Wgted Mean), and standard deviation weighted by establishments' number of employees (Wgted SD). A detailed description of all variables can be found in Appendix C.1.

	Obs	Mean	SD	P50	Wgted Mean	Wgted SD
Panel A: Firm and establishment characteristics						
log(total assets)	115,632	16.675	1.992	16.677	17.327	1.960
log(employees$_{estab}$)	115,632	2.916	1.584	3.247	5.079	1.221
log(employees$_{firm}$)	115,632	5.154	1.383	4.989	5.716	1.277
log(employees$_{world}$)	111,801	5.323	1.479	5.412	5.738	1.336
distance	114,995	2.594	2.589	2.652	1.340	2.239
semi-skilled/employees	115,600	0.120	0.184	0.035	0.119	0.188
skilled/employees	115,600	0.465	0.314	0.517	0.424	0.300
(complex/employees)$_{estab}$	110,649	0.135	0.167	0.076	0.156	0.151
(h. complex/employees)$_{estab}$	110,649	0.105	0.142	0.054	0.130	0.147
Panel B: Within-establishment variance of log wage and AKM components						
var$_{estab}$(log wage)	113,763	0.089	0.064	0.077	0.106	0.056
var$_{estab}$(person FE)	113,763	0.075	0.053	0.067	0.089	0.043
var$_{estab}$(Xb)	113,763	0.012	0.011	0.008	0.014	0.009
var$_{estab}$(residual)	113,763	0.015	0.014	0.011	0.018	0.013
Panel C: Establishment mean of log wage and AKM components						
log wage$_{estab}$	115,632	4.455	0.365	4.461	4.607	0.387
person FE$_{estab}$	115,632	4.188	0.251	4.164	4.275	0.252
estab. FE$_{estab}$	115,632	0.312	0.186	0.328	0.380	0.184
Xb$_{estab}$	115,632	-0.046	0.047	-0.037	-0.050	0.037
residual$_{estab}$	115,632	0.001	0.009	0.000	0.002	0.007

Figure C.3: Differences between smaller and larger firms: firm size
We sort firms into deciles according to total assets. The firm size deciles are plotted on the x-axis. The variable on the y-axis is stated in the title of each subfigure. A detailed description of all variables can be found in Appendix C.1.

(a) Log of total assets

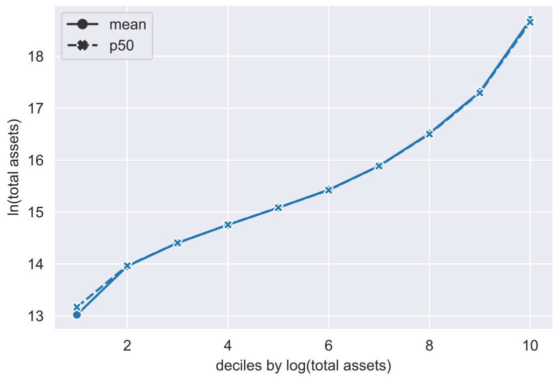

(b) Total assets in m of 2013 EUR

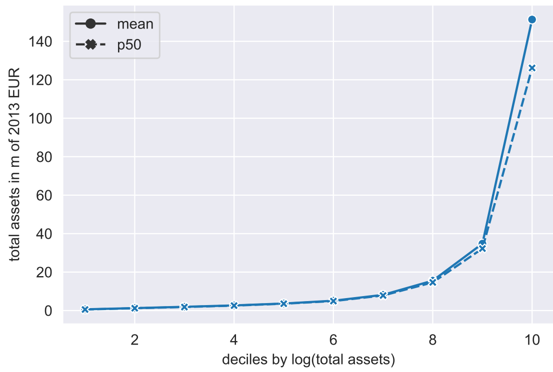

(c) Log of German employees

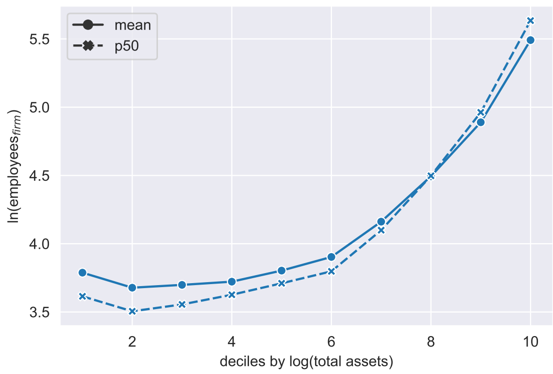

(d) German employees

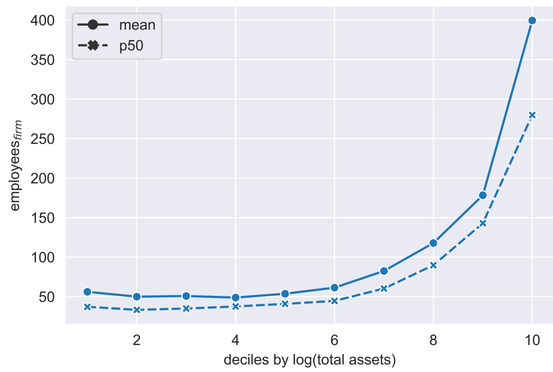

(e) Log of worldwide employees

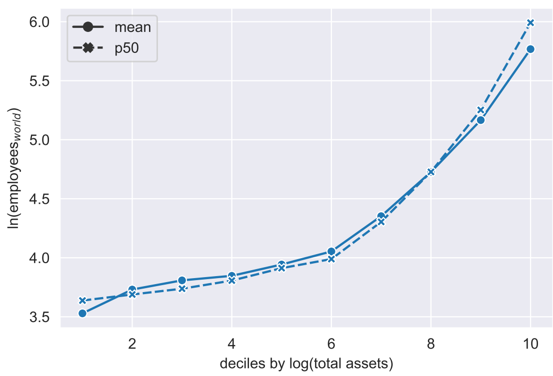

(f) Worldwide employees

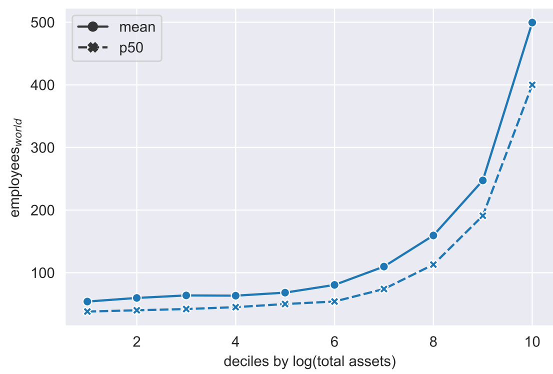

Figure C.4: Differences between smaller and larger firms: job characteristics
We sort firms into deciles according to total assets. The firm size deciles are plotted on the x-axis. The variable on the y-axis is stated in the title of each subfigure. A detailed description of all variables can be found in Appendix C.1.

(a) Semi-skilled/employees

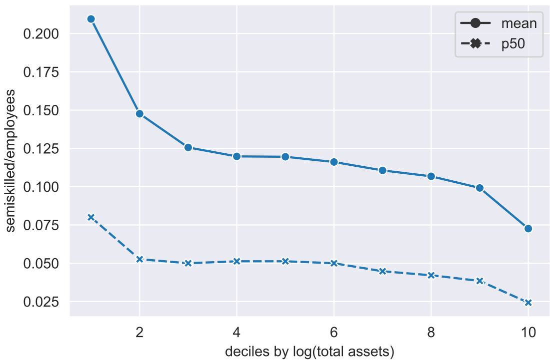

(b) Skilled/employees

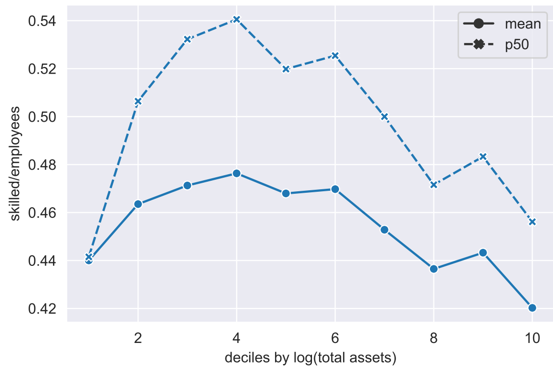

(c) Complex/employees

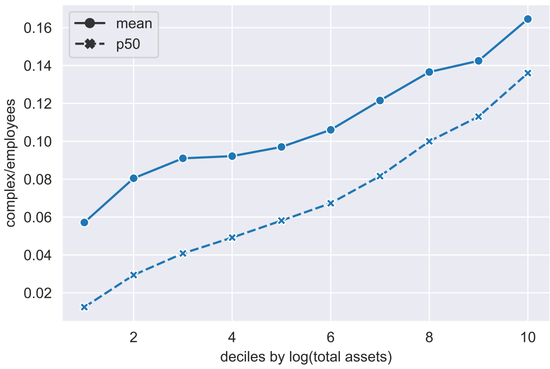

(d) Highly complex/employees

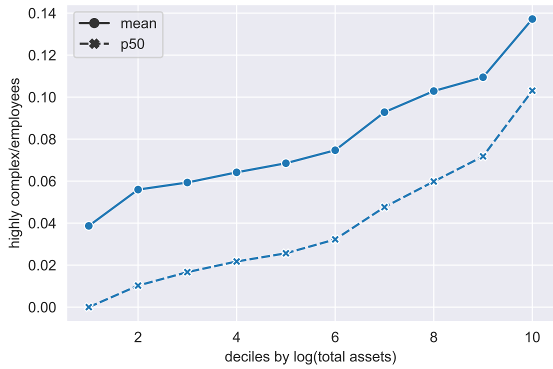

(e) Number of occupations

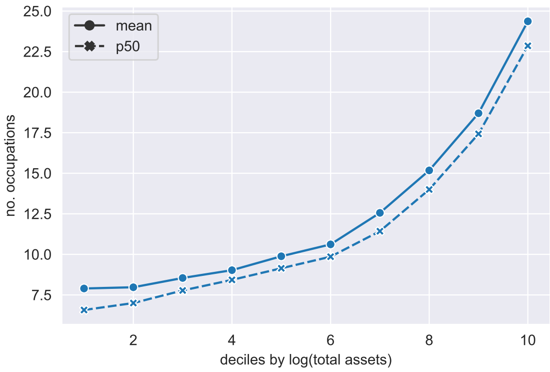

(f) Herfindahl index of occupations

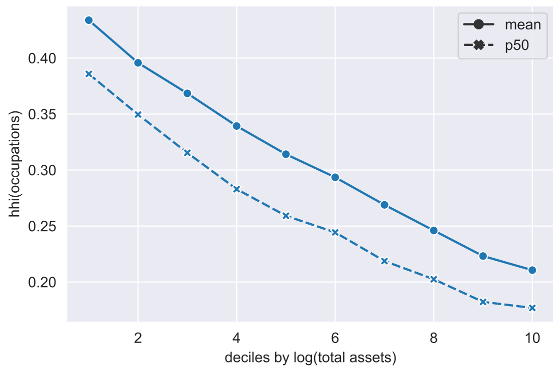

Figure C.5: Differences between smaller and larger firms: monitoring complexity
We sort firms into deciles according to total assets. The firm size deciles are plotted on the x-axis. The variable on the y-axis is stated in the title of each subfigure. A detailed description of all variables can be found in Appendix C.1.

(a) Manger/employees

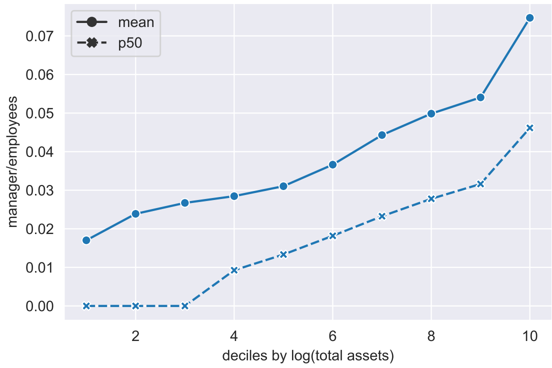

(b) Median(distance)

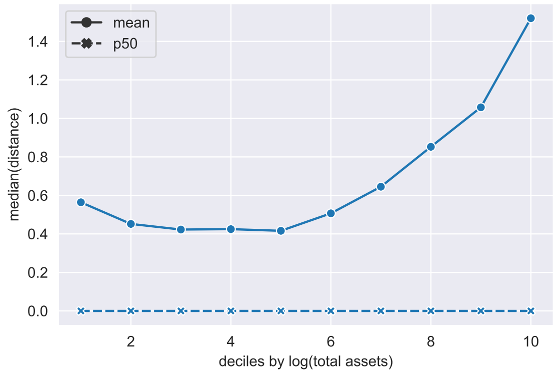

(c) Sd(distance)

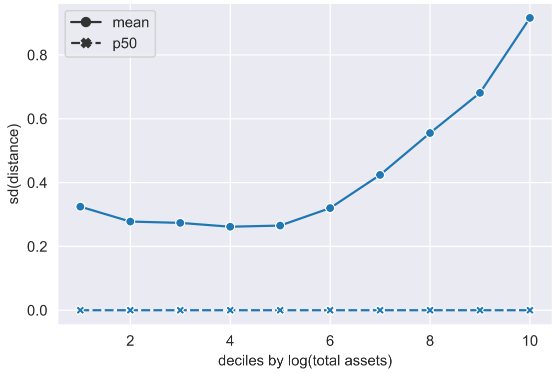

Figure C.6: Differences between smaller and larger firms: rent sharing
We sort firms into deciles according to total assets. The firm size deciles are plotted on the x-axis. The variable on the y-axis is stated in the title of each subfigure. A detailed description of all variables can be found in Appendix C.1.

(a) Profit/employees in m of EUR 2013

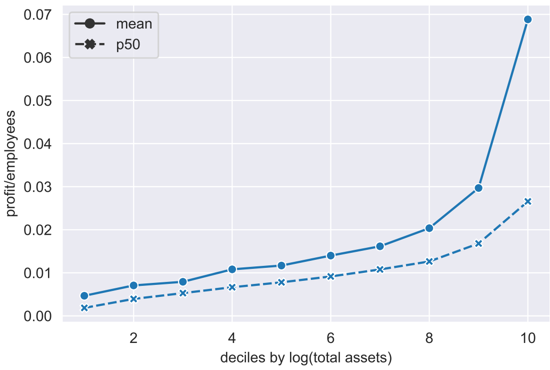

(b) Sales/employees in m of EUR 2013

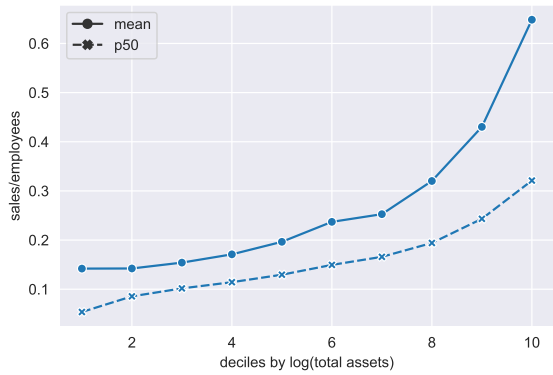

(c) Profit/total assets

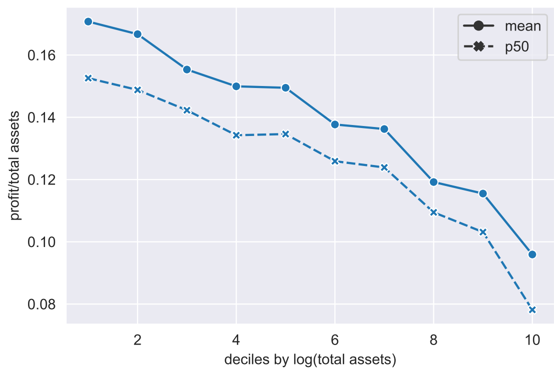

Table C.7: Wage inequality and firm size: measured by worldwide employees

This table presents regressions of wage inequality within and between firms on firm size, which is measured by total number of worldwide employees. In panel A, the dependent variables are the within-firm variance of log wage, the person fixed effect, the establishment fixed effect, the Xb, and the residual from the AKM-type regression. In Panel B, the dependent variables are the firm mean of log wage, person fixed effect, establishment fixed effect, Xb, and residual from the AKM-type regression. For the firm mean of residual, the coefficient on total number of worldwide employees is omitted or is virtually zero. This is why we report it as not relevant (n.r.). Regressions are weighted by firms' total number of German employees. T-statistics based on Huber/White robust standard errors clustered by firms are presented in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix C.1.

	(1)	(2)	(3)	(4)	(5)
Panel A: Within-firm inequality					
	var(log wage)	var(person FE)	var(estab. FE)	var(Xb)	var(residual)
log(empl$_{world}$)	0.0064***	0.0038***	0.00058***	0.00080***	0.0015***
	(7.51)	(6.88)	(8.25)	(8.04)	(6.70)
Obs	48,937	48,937	48,937	48,937	48,937
R2	0.03	0.02	0.07	0.02	0.03
Panel B: Between-firm inequality					
	log wage	person FE	firm FE	Xb	residual
log(empl$_{world}$)	0.11***	0.061***	0.048***	-0.0024***	n.r.
	(12.61)	(10.41)	(14.01)	(-6.54)	n.r.
Obs	48,937	48,937	48,937	48,937	48,937
R2	0.15	0.12	0.13	0.01	n.r.

Table C.8: Wage inequality and firm size: measured by total German employees

This table presents regressions of wage inequality within and between firms on firm size, which is measured by total number of German employees. In panel A, the dependent variables are the within-firm variance of log wage, person fixed effect, establishment fixed effect, Xb, and residual from the AKM-type regression. In Panel B, the dependent variables are the firm mean of log wage, person fixed effect, establishment fixed effect, Xb, and residual from the AKM-type regression. For the firm mean of residual, the coefficient on total number of German employees is omitted or is virtually zero. This is why we report it as not relevant (n.r.). Regressions are weighted by firms' total number of German employees. T-statistics based on Huber/White robust standard errors clustered by firms are presented in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. A detailed description of all variables can be found in Appendix C.1.

	(1)	(2)	(3)	(4)	(5)
Panel A: Within-firm inequality					
	var(log wage)	var(person FE)	var(estab. FE)	var(Xb)	var(residual)
log($empl_{firm}$)	0.0064***	0.0023***	0.00071***	0.0011***	0.0023***
	(6.12)	(3.31)	(8.42)	(10.15)	(9.70)
Obs	51,515	51,515	51,515	51,515	51,515
R2	0.03	0.01	0.09	0.03	0.07
Panel B: Between-firm inequality					
	log wage	person FE	firm FE	Xb	residual
log($empl_{firm}$)	0.11***	0.062***	0.049***	-0.0031***	n.r.
	(10.53)	(8.62)	(11.82)	(-7.29)	n.r.
Obs	51,515	51,515	51,515	51,515	51,515
R2	0.14	0.11	0.13	0.02	n.r.

Table C.9: Full decomposition of within-firm and within-establishment variances

This table presents the full decomposition of the within variance on firm- and establishment-level into the AKM components. Panel A presents the decomposition of the within-firm variance. Panel B presents the decomposition of the wihin-establishment variance. Presented are the employee-weighted mean values of within variances and covariances. A detailed description of all variables can be found in Appendix C.1.

	Obs	Var Component	Share
Panel A: Within-firm variance			
var(log wage)	51,515	0.113	1.000
var(person FE)	51,515	0.095	0.844
var(estab. FE)	51,515	0.002	0.014
var(Xb)	51,515	0.015	0.130
var(residual)	51,515	0.018	0.155
2cov(person FE, Xb)	51,515	-0.018	-0.158
2cov(person FE, residual)	51,515	0.002	0.019
2cov(Xb, residual)	51,515	-0.000	-0.002
2cov(person FE, estab. FE)	51,515	-0.000	-0.002
2cov(estab. FE, Xb)	51,515	0.000	0.000
2cov(estab. FE, residual)	51,515	0.000	0.000
Panel B: Within-establishment variance			
var_{estab}(log wage)	113,763	0.106	1.000
var_{estab}(person FE)	113,763	0.089	0.841
var_{estab}(Xb)	113,763	0.014	0.136
var_{estab}(residual)	113,763	0.018	0.166
$2cov_{estab}$(person FE, Xb)	113,763	-0.017	-0.158
$2cov_{estab}$(person FE, residual)	113,763	0.002	0.019
$2cov_{estab}$(Xb, residual)	113,763	-0.000	-0.002

Decomposition of overall wage inequality into a between and a within component

We decompose the overall variance of log daily wage into between-firm and within-firm variation. As we observe information on firms and establishments, we can also compare the firm-level decomposition to the decomposition into a between-establishment and a within-establishment component ignoring firm-level information.

The overall variance of wage can be decomposed into a between- and a within-firm component

$$var(y_t^{i,j,k}) = var(\bar{y}_t^k) + \sum_k w_k \times var_k(y_t^{i,j,k} | i \in j, j \in k), \tag{C.1}$$

where $var(\bar{y}_t^k)$ is the between-firm variance of firm mean wage, and the second term is the employment-weighted mean of within-firm variance in employee wage. w_k denotes the employment share of firm k in the sample. Alternatively, we can decompose the overall variance into a between- and within-establishment component ignoring the firm-level information

$$var(y_t^{i,j,k}) = var(\bar{y}_t^j) + \sum_j w_j \times var_j(y_t^{i,j,k} | i \in j). \tag{C.2}$$

This allows us to explore the difference in the share of within variation to overall variation using firm and establishment information,

$$\frac{\sum_j w_j \times var_j(y_t^{i,j,k} | i \in j)}{var(y_t^{i,j,k})} - \frac{\sum_k w_k \times var_k(y_t^{i,j,k} | i \in j, j \in k)}{var(y_t^{i,j,k})}. \tag{C.3}$$

The comparison is interesting because the identification of firm structures was not possible in German administrative data until the release of the ORBIS-ADIAB dataset (see Section 3.2.2) and gains importance given the increasing complexity of firm structures.