Essays on Innovation and Competition

An Analysis of Codetermination, Price Disclosure, and Tax Pass-Through

Inauguraldissertation
zur Erlangung des akademischen Grades
Doctor oeconomiae publicae (Dr. oec. publ.)
an der Volkswirtschaftlichen Fakultät
an der Ludwig-Maximilians-Universität München

2023

vorgelegt von
Alina Sagimuldina

Referentin:	Prof. Dr. Dr. h.c. Monika Schnitzer
Korreferent:	Prof. Dr. Martin Watzinger
Promotionsabschlussberatung:	12. Juli 2023

Tag der mündlichen Prüfung: 3. Juli 2023

Namen der Berichterstatter: Monika Schnitzer, Martin Watzinger, Florian Englmaier

Acknowledgements

First and foremost, I would like to thank my supervisor Monika Schnitzer. I am deeply grateful for her unwavering support, guidance, and continuous encouragement over the past years. Her expertise, thoughtful comments, and sharp advice have been key to the completion of this dissertation. Furthermore, I would like to thank my second supervisor Martin Watzinger for always taking the time to provide me with helpful advice. His enthusiasm for innovation research served as a continuous source of inspiration. I would also like to thank Florian Englmaier, who kindly agreed to join my dissertation committee.

I am very grateful to Alessandra Allocca and Markus Nagler. Their mentoring, willingness to discuss and give feedback on early-stage ideas, and sincere support were instrumental in many moments. I am also grateful to Anna Gumpert for all that I have learned from her rigor and clarity of thought.

My colleagues at the Munich Graduate School of Economics, at the Microeconomic Determinants of Labor Productivity Program, and especially at the Seminar for Comparative Economics were always helpful and made graduate school much nicer. I have fond memories of the time spent together in courses, at the office, or over lunch with Maximilian Schader, Kathrin Wernsdorf, Cristina Rujan, Agustín Faure, Marie-Louise Arlt, and Xuan Teng. I also wish to thank Ines Steinbach, Julia Zimmermann, and the student assistants at the Seminar for Comparative Economics for their outstanding help.

My co-authors Robin Mamrak, Felix Montag, Monika Schnitzer, and Christoph Winter deserve special credit for this dissertation. I learned a lot from working with them and would have not been able to realize any of the projects without the insights I gained from our joint work and the discussions.

Besides, I am happy to have been a part of the "GRK 1928 – Microeconomic Determinants of Labor Productivity" graduate program and want to thank the program coordinator Carsten Eckel for all of his efforts. I gratefully acknowledge the Research Training Group GRK 1928 and the Deutsche Forschungsgemeinschaft for their financial support. This dissertation also benefited from numerous comments and discussions at conferences and seminars.

My deepest gratitude goes to my family, especially my parents, for their love and support throughout this journey. Without it, none of this would have been possible.

Alina Sagimuldina
July 2023

Contents

List of Tables

List of Figures

Preface

Innovation and competition policies are prominent in today's political discourse. Many developed economies experience a slowdown in productivity growth, geoeconomic fragmentation, and a rise in concentration (Covarrubias, Gutiérrez, and Philippon, 2019), an increase in markups (De Loecker and Eeckhout, 2021; De Loecker, Eeckhout, and Unger, 2020), and a rise in firm profits (e.g., De Loecker, Eeckhout, and Unger, 2020). Governments in the EU and the U.S. have intensified efforts to stimulate research and development activity with public funding programs such as Horizon Europe and the 2022 CHIPS and Science Act, and politicians, as well as competition authorities, increasingly advocate for more stringent antitrust and competition regulations (The Economist, 2022).

Innovation and market power have far-reaching implications. Technological progress drives long-run productivity growth (Romer, 1986). Market power has been shown to increase inequality (Baker and Salop, 2015), reduce labor share (Autor, Dorn, Katz, Patterson, and Van Reenen, 2020), lead to missing innovation (Watzinger, Fackler, Nagler, and Schnitzer, 2020), and slow down productivity growth (Olmstead-Rumsey, 2022), with a broad consensus of the benefits to promote competition. As policymakers are likely to make decisions in a setting with increasingly limited resources, it is particularly important to understand what advances innovation, how to design policies that promote competition, and how competition affects the effectiveness of other policies.

This dissertation sheds light on aspects that relate to the determinants of innovation as well as the determinants and effects of competition. The chapters of this thesis are self-contained and can be read independently. The first chapter strives to inform about what affects innovation and automation innovation in a firm, with a particular focus on the role of firm governance. The second chapter is concerned with what determines the effectiveness of the policy that aims to foster competition, particularly the price transparency policy. The third chapter turns to studying how imperfect consumer information about prices determines the intensity of competition and how this affects the pass-through of commodity taxes.

The first chapter studies how shared governance of firms, or worker representation on corporate boards, affects innovation and particularly automation innovation in a firm. Theoretical predictions on this effect are ambiguous: shared governance may improve information flow between employees and capital owners and increase worker productivity (Freeman and Medoff, 1985; Freeman and Lazear, 1995), but may also lead to a hold-up problem, delayed decision-making, and disinvestment (Grout, 1984; Jensen and Meckling, 1979). In addition, worker representatives may raise concerns about the potential labor-displacing effects of new technologies. Empirically, the effect is difficult to estimate because the decision to give employees decision-making rights or the legal requirement to do so is often endogenous.

To circumvent these issues, I use a sharp legal change to shared governance in Germany. On August 10, 1994, the German parliament enacted the reform of the Stock Corporation Act. As a result of the last-minute political compromise, the reform abruptly abolished the so-called codetermination, a mandate to allocate a one-third share of supervisory board seats to worker-elected representatives, for smaller stock firms that incorporate on the reform day or afterward. Importantly, stock firms incorporated before the reform remained subject to the mandate and could not evade it by reincorporating after the reform. Other changes applied to all stock firms with up to 500 employees, independent of their incorporation date. The reform also did not affect alternative ways of firm-level worker representation or firms with other legal forms, such as limited liability companies (LLCs).

The unexpected and rigid cohort-based change to codetermination under the reform gives rise to quasi-experimental variation in shared governance across stock firms that I use in the empirical design. I collect information on the incorporation date and legal form of firms that incorporate in Germany within a five-year window around the reform date and merge it with patent data. To estimate the effect of codetermination on innovation, I compare patent evolution at stock firms incorporated shortly before and shortly after the August 10, 1994 cutoff to the patent evolution at untreated LLCs incorporated shortly before and after the reform in a difference-in-differences design. Using the difference among peer cohorts of never treated LLCs additionally accounts for incorporation time, cohort, or age effects.

The empirical analysis shows that the shared governance of firms leads to a strong decline in automation innovation. The effect on non-automation and general innovation is also negative but smaller in size and imprecise. Whereas the effect on automation innovation is driven by a decline in more valuable patents, this is different for non-automation. For automation, the effect is the largest for patents that build on or are closely related to science. I find that automation patents at codetermined stock firms rely on a smaller and older base of the prior art, and their inventor teams become smaller. For non-automation, the corresponding effects are much smaller and the confidence intervals do not exclude zero.

The chapter provides novel empirical evidence on how codetermination affects firm technological progress, an important but thus far understudied potential effect of shared governance. This complements the existing evidence in the literature on the effect of shared governance on other firm and worker-level outcomes (e.g., Jäger, Schoefer, and Heining, 2021, Blandhol, Mogstad, Nilsson, and Vestad, 2020, and Harju, Jäger, and Schoefer, 2021). It also shows that worker representation on corporate boards is likely to be an additional determinant of automation innovation, along with worker wages (Dechezleprêtre, Hémous, Olsen, and Zanella, 2022), labor supply (Danzer, Feuerbaum, and Gaessler, 2020), and demographic change (Acemoglu and Restrepo, 2022). Overall, this chapter suggests that policies that change the shared governance of firms could affect the direction of firm innovation and, in particular, divert it away from automation.

The second chapter, which is based on joint work with Felix Montag and Christoph Winter, changes its focus from firm governance and its effect on innovation to an analysis of a pro-competitive policy. We focus on mandatory price disclosure, which is becoming a popular tool to make markets more competitive. Before introducing mandatory price disclosure, it is crucial to understand what its effect is going to be in the particular setting. We ask what determines the price effect of mandatory price disclosure in a setting where consumers are imperfectly informed about prices. In particular, we study how the effect of providing price information to consumers depends on how well informed they are beforehand.

Our theoretical analysis contributes to the literature on mandatory price disclosure by deriving novel predictions about how it affects prices in the context of the Varian (1980) model. On the supply side, there are sellers that sell a homogeneous good and set prices. On the demand side, there are fully informed *shoppers* that know all prices, as well as uninformed *non-shoppers* that visit a seller at random. We model mandatory price disclosure as leading to an increase in the share of *shoppers*. We assume that price information coming from mandatory price disclosure always reaches a fixed number of consumers, irrespective of whether these are *shoppers* or *non-shoppers*. Theoretically, we show that the more uninformed consumers there are prior to the introduction of the policy, the larger is the reduction in prices it induces.

We test the predictions in the context of the introduction of mandatory price disclosure in the German retail fuel market. Two features of the setting make it particularly suitable for this analysis: First, we observe high-frequency, station-level price changes for Germany and France before and after the introduction of this price transparency policy. Second, mandatory price disclosure was introduced simultaneously for diesel and gasoline. On average, consumers buying gasoline are less informed about prices than consumers buying diesel. Consumers can also not substitute between fuel types. Since the same fuel stations sell both types of fuel, there are no supply-side differences between fuel types. We use a difference-in-differences design to estimate the price effect of mandatory price disclosure for each fuel type, where fuel stations in Germany are part of the treatment group and fuel stations in France are in the control group.

We find that mandatory price disclosure decreases prices for all fuels but that this decrease is larger for gasoline, which has a less informed consumer base, than for diesel. The difference in treatment effects is particularly strong in the five months after the introduction of mandatory price disclosure. Thereafter, the treatment effect stabilizes at a lower level for both fuel types. Finally, we show that follow-on information campaigns, such as local radio reports about fuel prices, can further intensify the treatment effect.

The chapter contributes to the empirical literature on price transparency policies by studying a novel mechanism of how mandatory price disclosure affects prices. In this context, our analysis highlights the importance of the share of consumers informed about prices before the policy introduction. Our findings relate to the empirical literature that studies other mechanisms behind the effect of mandatory price disclosure, such as its role in stabilizing collusion (Albæk, Møllgaard, and Overgaard, 1997; Luco, 2019) or inducing credibility to price-based advertising (Ater and Rigbi, 2023). Overall, this chapter suggests that mandatory price disclosure is most effective in markets where few consumers are well-informed before its introduction and that complementary information campaigns can be used to strengthen the effect of the policy.

The third chapter, which is based on joint work with Felix Montag and Monika Schnitzer, turns to studying what determines competition and how this affects commodity tax pass-through. Understanding how and when firms pass through taxes to consumers is fundamental for the design of optimal tax policy. Pass-through determines the corrective effect of Pigouvian taxes, the effectiveness of unconventional fiscal policy to stimulate the economy and the distributional consequences of any commodity tax. At the same time, even though consumers are rarely fully informed about prices, we know fairly little about tax pass-through when consumers have imperfect price information. In the third chapter, we ask how market power caused by imperfect consumer information about prices affects the pass-through of commodity prices.

Theoretically, we adapt the Stahl (1989) model to the analysis of tax pass-through. Similarly to the Varian (1980) model, the framework features fully informed *shoppers* and uninformed *non-shoppers*. *Shoppers* know all prices and *non-shoppers* can search for prices sequentially. With this framework, we introduce a novel notion of price sensitivity of demand to the analysis of tax pass-through in oligopolistic markets. The larger the share of *shoppers*, the higher is the average price sensitivity of consumers and the more it pays for sellers to compete with their choice of prices. How well consumers are informed about prices therefore determines the equilibrium intensity of competition in the market.

We derive two theoretical predictions about how competition affects pass-through: First, the more price sensitive consumers are on average, the higher is the pass-through rate. This is different to how another common notion of price sensitivity, the price elasticity of demand, affects pass-through. A classic result under perfect competition is that the higher the price elasticity of demand, the lower the pass-through rate (e.g., Weyl and Fabinger, 2013). Second, there is a hump-shaped relationship between the number of sellers and pass-through.

To test our predictions empirically, we exploit a temporary decline in the value-added tax and a carbon price introduction in the German retail fuel market. A key feature of this setting is that we can separately study fuel products that differ in how well their consumers are informed about prices. We use a unique dataset containing the universe of price changes at fuel stations in Germany and France and estimate pass-through of the tax changes to diesel and gasoline prices in a difference-in-differences design. As predicted by the theory, we find that pass-through is higher for diesel, which is used by on average more price sensitive consumers, than for gasoline. We also find a hump-shaped relationship between pass-through and the number of fuel stations in a local market.

By showing how price sensitivity affects pass-through when consumers are imperfectly informed, we shed light on a novel explanation of what determines tax pass-through and extend the existing empirical literature (e.g., Miravete, Seim, and Thurk, 2018, Nakamura and Zerom, 2010, and Hollenbeck and Uetake, 2021). As price sensitivity affects pass-through differently from the price elasticity of demand in markets with perfect competition, our study suggests that imperfect consumer information should be accounted for in the design of optimal Pigouvian taxes and when predicting the distributional consequences of a tax.

In summary, this dissertation provides new insights into firm innovation and the driving forces behind the determinants and effects of competition. Evidence from these analyses may hopefully contribute to the design of economic policies accounting for the increasing importance of innovation and competition and help make efficient use of scarce resources.

Chapter 1

Does Labor on the Board Affect Firm Innovation? Evidence from Codetermination in Germany

1.1 Introduction

Some form of shared governance in the private sector exists in many EU countries.[1] In a system of industrial relations with shared governance, employees contribute to a firm's decision-making via organizations such as works councils or through direct participation in corporate boards. The alternative is the system without shared governance, which is prevalent in the U.S. and other liberal market economies. However, proposals for worker representation have also been included in political discussions in some of these countries, such as the U.K. and the U.S.[2]

One important consideration behind the evaluation of shared governance and the decision to introduce it is its potential effect on a firm's technological progress. Theoretical predictions on the relationship between the shared governance of firms and technological progress are ambiguous. Shared governance may improve information exchange between capital owners and employees and increase worker productivity, hence, helping firms progress (Freeman and Med-

[1] See, for instance, Conchon (2011) for an overview of the board-level employee representation regulations in the EU countries.

[2] See, e.g., The Economist (2020) article "Deutschland AG rethinks workers' role in management" for further information on recent trends in shared governance.

off, 1985; Freeman and Lazear, 1995). On the other hand, it may lead to a hold-up problem, delayed decision-making, and disinvestment (Grout, 1984; Jensen and Meckling, 1979).[3] Employee representatives may also raise concerns about the potential labor-displacing effects of new technologies.[4]

In this chapter, I empirically study how the shared governance of firms affects innovation and particularly automation innovation in a firm. I focus on a reform in Germany that in 1994 abruptly abolished the mandate for worker representation on firm supervisory boards (also known as codetermination) for newly incorporated stock firms but locked in already incorporated stock firms with the mandate.[5] The reform gives rise to quasi-experimental variation in shared governance across stock firms, as the cohort-based change to codetermination was an unexpected result of a last-minute political compromise. I exploit firm-level data on legal forms and incorporation dates to determine whether a firm is a stock corporation or a limited liability company (LLC) and is affected by the reform and merge it with patent data.[6] Empirically, I use a difference-in-differences framework to estimate the effect of codetermination on innovation and compare the innovation outcomes at stock firms incorporated shortly before and after the reform to the outcomes at their peer cohorts of never treated LLCs. The identifying assumption is that without the 1994 reform, the difference in the number of employee adjusted patents between stock firms that incorporate shortly before and shortly after the reform would be identical to the corresponding difference between LLCs.

I find that shared governance has a strong negative effect on automation innovation. This evidence is consistent with the theory that increased worker representation may prevent firms from investing in labor-saving technologies. Depending on the fixed effects used in the estimation, codetermination reduces automation patenting by 76 to 79 percent. The effect is driven by more valuable automation patents, as reflected by their distance to science, number and age of patent backward citations, and inventor team size. The effect of shared governance on innovation in

[3]Alternatively, under certain conditions, high wages could encourage firms to produce more labor-saving innovation (Acemoglu, 2010).

[4]The Pew Research Center (2017) survey reports that about two-thirds of U.S. adults are concerned about the negative effects of job automation.

[5]*"Gesetz für kleine Aktiengesellschaften und zur Deregulierung des Aktienrechts"* (1994). The other form of worker representation at a firm level is through works councils. In what follows, codetermination refers only to worker representation on the board and not to works councils.

[6]Stock firms refer to "*Aktiengesellschaften*" and "*Kommanditgesellschaften auf Aktien*". LLCs refer to "*Gesellschaften mit beschränkter Haftung (GmbH)*".

general and non-automation innovation is also negative, but smaller in magnitude and imprecise. I find no sizable or statistically significant decline in non-automation patents that build on science. The effects on non-automation patent value, as reflected by backward citations and inventor team size, are also not distinguishable from zero. Overall, this suggests that employee representation on corporate boards shifts firm patenting away from automation but does not otherwise significantly affect innovation.

Like many other countries, Germany has a two-tier corporate governance system. The executive board is the managing body responsible for the day-to-day business of a firm. The supervisory board elects and monitors the executive board. It consists of shareholders and often workers. The supervisory board is encouraged to be actively involved in all fundamental firm decisions, such as strategic and financial planning, investments, outsourcing, and R&D expenditures, by the German Corporate Governance Code. In supervisory boards, worker representatives enjoy the same rights as the shareholders. Anecdotally, shareholders and worker-elected representatives aim to reach an agreement, and decisions in the supervisory board are made unanimously.[7]

The August 10, 1994 reform of the Stock Corporation Act abolished the codetermination mandate for stock firms with up to 500 employees incorporated on or after the reform date. Importantly, stock firms incorporated before the reform remained subject to the mandate and could not evade it by reincorporating after the reform. The cohort-based codetermination aspect of the reform was an unexpected result of a last-minute political compromise. The other changes applied to all stock corporations independent of their incorporation date.[8] The reform did not introduce changes to other forms of worker representation or to firms with other legal structures: independent of their incorporation date, LLCs are not mandated to appoint worker-elected representatives to their board if they have up to 500 employees. As the reform gives rise to variation in codetermination based on the incorporation date of a stock firm, it allows circumventing the need for granular information on worker representation at the firm level.

In this chapter, I exploit the grandfathering rule of the 1994 reform and compare patent evolution at stock firms incorporated shortly before (treated) and shortly after (untreated) the August 10, 1994 cutoff to the patent evolution at untreated LLCs incorporated shortly before and after

[7] See, for example, Gold (2011) for further details on the contribution of worker-elected representatives to board decisions.

[8] See, e.g., Raiser, Veil, and Jacobs (2015).

the reform in a difference-in-differences design. In the baseline set-up, I study the evolution of patents per 100 employees filed between 1998 and 2014 by stock firms and LLCs that incorporate within three years of the reform, and control for year by industry fixed effects. Using the difference among peer cohorts of never treated LLCs additionally accounts for incorporation time, cohort, or age effects.

To estimate the effect of codetermination on innovation, I combine a number of data sources: (i) data on characteristics of German firms, firm ownership, and board composition based on Bureau van Dijk's Orbis, (ii) a firm-level patent data set from Orbis Intellectual Property, (iii) a comprehensive data set on patents PATSTAT, (iv) the automation patent classification developed by Dechezleprêtre et al. (2022), and (v) information on the supervisory board composition of publicly listed firms that distinguishes worker-elected directors based on the Hoppenstedt Aktienführer.

I find that firms locked in with codetermination produce 76 to 79 percent fewer automation patents and 54 to 59 percent fewer patents per 100 employees in general.[9] The effect on automation innovation is economically and statistically significant. The effect of codetermination on innovation is negative but imprecise. These results persist to the estimation under a majority of alternative bandwidths around the reform date, using samples of firms that incorporate within one year and up to five years around the reform. As a robustness check, I estimate the effect on (automation) innovation using the placebo reform date on August 10, 1996 or 1997. Reassuringly, I find no sizable or statistically significant effect using the placebo reform dates.

To further evaluate the validity of the empirical design, I check whether stock firms strategically delay their incorporation until after the 1994 reform to evade codetermination. I do not find evidence that firms manipulate their incorporation date in a McCrary (2008) test of the density continuity around the cutoff date. In addition, I check whether the industry composition or composition of firms based on their legal form changes after the reform, and find little evidence for compositional shifts around the reform. This is in line with the sharp and unexpected grandfathering in codetermination under the 1994 law and its binding nature for the locked-in stock corporations.[10]

[9]The values depend on the fixed effects used in the estimation.

[10]Several shareholders attempted to challenge an arbitrary grandfathering cutoff date in court, suggesting the reform is binding in the affected stock firms. See, e.g., BVerfG, the decision of the second Chamber of the First Senate from January 9, 2014 (1 BvR 2344/11).

In the next step, I study the mechanisms. I decompose the average effects by patent distance to science and estimate the effects on the number and age of patent backward citations as well as the size of inventor teams. Krieger, Schnitzer, and Watzinger (2022) show that patents closer to science are on average more valuable, have larger inventor teams and build on larger and younger prior art. I find that patents that closely relate to science drive the negative effect on automation innovation. On average, automation patents in codetermined stock firms include 29.4 percent fewer and 62 percent older backward citations, and their inventor teams are on average 37.3 percent smaller. This is different for non-automation, where the effects are closer to zero and not statistically significant.

This chapter contributes to the literature on the effects of worker representation in corporate boardrooms on firm and worker-level outcomes. While other studies show how changes to shared governance affect a variety of these outcomes, evidence on the innovation effect of such changes remains scarce. There is an expanding literature that uses micro-level administrative data and quasi-experimental variation in worker representation on the board.[11] Most recent studies report zero or small positive effects of codetermination. The most closely related paper is Jäger, Schoefer, et al. (2021), which analyzes the effect of codetermination using the 1994 reform to the Stock Corporation Act in Germany. They find that codetermination has no effect on wages, rent sharing, or labor share, and show some evidence for an increase in firm capital formation. Blandhol et al. (2020) use administrative data from Norway and find no effect of worker representation on corporate boards on worker compensation and other firm outcomes. Harju et al. (2021) study the effect of worker representation on board or supervisory councils using Finnish administrative data and disentangle the effect of worker voice from worker decision-making authority. Kim, Maug, and Schneider (2018) find that workers in quasi-parity firms are protected from layoffs at the expense of lower wages. I contribute to this literature by providing new evidence on the effect of shared governance on firm innovation and its direction.

This work is also related to the literature that studies the relationship between organized labor and innovation. There is a vast literature on the economic impacts of unions. DiNardo and

[11] There is also a vast literature that exploits variation in codetermination based on firm employment (comparing firms with one-third and one-half worker share) and firm industry or reports correlations in the German setting of shared governance (FitzRoy and Kraft, 1993; Gorton and Schmid, 2004; Svejnar, 1981). In the survey of this literature, Conchon (2011) reports that 10 studies report positive effects of codetermination, 9 studies find negative effects, and 11 studies find zero effects.

Lee (2004) report the small effects of unionization on business survival, employment, output, and productivity. Most studies find a negative relationship between unionization and R&D investment or innovation (Acs and Audretsch, 1988; Bradley, Kim, and Tian, 2016; Connolly, Hirsch, and Hirschey, 1986; Hirsch and Link, 1987). This study is one of the first to provide quasi-experimental evidence on the effect of worker representation at a firm level, where organized labor may exert a more direct influence on innovation, and to distinguish between automation and non-automation innovation.[12]

Finally, this chapter contributes to the literature on the determinants of automation innovation. This study shows that the corporate governance of firms, particularly worker representation on corporate boards, is likely to be an additional determinant of firm automation innovation. Prior empirical research suggests that low- and high-skilled worker wages (Dechezleprêtre et al., 2022), labor supply (Danzer et al., 2020), and demographic change (Acemoglu and Restrepo, 2022) play a role in the decision to produce automation innovation. I extend this literature by showing that the shared governance of firms is associated with a decline in the number and value of automation patents and changes the way in which automation innovation is organized.

The remainder of this chapter is structured as follows: Section 1.2 describes the institutional setting. Section 1.3 gives an overview of the data and provides some descriptive evidence. Section 1.4 discusses the empirical design. Section 1.5 includes the empirical results and Section 1.6 presents the mechanisms. Section 1.7 concludes.

1.2 Institutional Background

1.2.1 Firm governance

Germany has a two-tier corporate governance system. The executive board is a managing body. It is responsible for day-to-day operations in a firm. The supervisory board elects, monitors and advises the executive board. It also determines the size of the executive board and the required

[12]To the best of my knowledge, Kraft, Stank, and Dewenter (2011) is the only other study that analyzes the effect of board-level worker representation on firm innovation. They exploit a reform in Germany that in 1976 extended a mandate of one-half worker share on the board to firms with more than 2,000 employees. In a comparison of patent evolution at 148 larger and smaller stock firms, they find that stock firms with one-half worker share file slightly more patents as compared to stock firms with one-third worker share or without codetermination.

qualifications of its members. The supervisory board consists of shareholders and often worker-elected representatives.

Whereas the executive board is responsible for the development of the firm strategy and its implementation, it is advised to coordinate these decisions with the supervisory board. The German Corporate Governance Code encourages an open dialogue between the executive and the supervisory board. The executive board is mandated to inform the supervisory board comprehensively and in a timely manner on fundamental issues relevant to the company. The supervisory board may at any time require the executive board to provide additional information.

In turn, the supervisory board is advised to be actively involved in all fundamental firm decisions, such as strategic and financial planning, investments, outsourcing and R&D expenditures. Certain transactions or investments may be subject to approval by the supervisory board. According to the German Stock Corporation Act, the supervisory board is required to meet regularly and at least four times every calendar year.[13]

Results of a survey of the supervisory board members of German firms, conducted by I.M.U. (2021) in 2019, confirm that the role of the supervisory board is likely to extend beyond the supervision of the executive board.[14] A majority of the survey respondents report that the supervisory board provides some advice to the management. The respondents rank "medium- to long-term strategy", "restructuring and reorganization", and "impact of technological change on business model" along with "annual financial statements" among the top topics discussed by the board. Anecdotally, this suggests that the supervisory board tends to engage in questions beyond its immediate responsibility and not only monitor but also advise the firm management.

Oftentimes, the supervisory board includes worker-elected representatives. This so-called codetermination is one of the forms through which employees can contribute to decision-making in a firm.[15] Jäger, Noy, and Schoefer (2022b) give a detailed overview of how shared governance operates in Germany. The employees' representatives are elected in a general, secret, equal,

[13]The supervisory board of unlisted companies is required to meet at least two times a year.

[14]I.M.U. is the Institute for Codetermination and Corporate Governance that operates within the Hans Böckler Foundation in Germany. In total, 506 worker-elected representatives from supervisory boards of German firms participated in the survey.

[15]Another form of employee representation is through a works council, which is a shop-floor representation institution. The rights of the works councils are stipulated by *BetrVG* (the 1972 Works Constitution Act). Employees in companies with five or more workers have a right to demand the formation of a works council. Works councils can negotiate directly with the employer and have rights on firm decisions related to, e.g., occupational safety, work hours, and organizational changes.

and typically direct election by all employees entitled to vote.[16] In supervisory boards of larger companies, two to three employee seats are reserved for members of a trade union present in the company. The remaining seats are distributed among blue-collar workers, salaried employees and executive employees in proportion to their groups' presence in the company.

All worker representatives enjoy equal rights with shareholder representatives on the board and nearly all must be employees of the firm. Gold, Kluge, and Conchon (2010) report experiences of the employee representatives on supervisory boards in German firms in a series of interviews. Anecdotally, shareholders and worker-elected representatives aim to reach an agreement and decisions in the supervisory board are made unanimously. One of the representatives comments: "I don't feel in a minority or any kind of inferiority. Both sides [worker-elected representatives and shareholders] try to achieve unanimity". In addition, most of the interviewed employee representatives insist that they prioritize the interests of the employees: "I have to focus on getting as much as I can for the employees... We [worker representatives] obviously have reasons, mainly to do with preserving jobs, why we assent to a particular decision or do not oppose it." Appendix A.1 provides further anecdotal evidence on the role of worker representatives in the supervisory board.

The codetermination mandate varies from zero to full parity and depends on the firm legal form, incorporation date, and the number of employees. It was first introduced in 1951 in iron, coal and steel-producing industries, where it mandates full parity for firms in the sector with more than 1,000 employees.[17] The supervisory board has an equal number of employee and employer representatives. In case of a tie during a vote on the board, an external member who is neutral holds a decision (§§ 1, 4 *MontanMitbestG* 1951).

In 1952, codetermination was extended to all firms that have at least 500 employees and required that one third of supervisory board seats are occupied by worker representatives.[18] State-owned enterprises and companies with fewer than 500 employees were exempted from codetermina-

[16] Employee representatives in companies with more than 8,000 employees are elected by delegates unless the employees entitled to vote decide on the direct election (§9 *MitbestG*).

[17] Jäger, Noy, and Schoefer (2022a) and Silvia (2013) provide further details on the historical development of codetermination in Germany. In particular, they report that the aftermath of the Second World War provided favorable conditions to the labor groups in Germany. As heavy industries supplied the machinery during the War and its leaders supported the Nazi regime, the Allied forces aimed to democratize these industries and provide more decision-making power to the employees instead of the industrialists.

[18] Today, this applies to firms that have more than 500 employees (§1 *DrittelbG* 2004).

tion, with the exception of stock corporations that were not family-owned (§§ 76, 77 *BetrVG* 1952).[19]

The last major extension of codetermination was introduced in 1976. It mandated all firms that have more than 2,000 employees and operate outside of the coal, iron and steel producing sector to 'quasi-parity'. In these companies, employees have a right to 50 percent representation on their supervisory boards (§§ 1, 7 *MitbestG* 1976).[20]

In sum, in all firms with between 501 and 2,000 employees, worker-elected representatives constitute one third of the supervisory board, independent of the firm legal form. The share increases to one-half for all firms with more than 2,000 employees, and there is full parity for firms with more than 1,000 employees in the iron, coal and steel sector. In smaller firms with up to 500 employees, only stock firms are mandated to preserve one-third share of supervisory board seats to employee representatives, unless these are family-owned.

1.2.2 1994 Reform of the Stock Corporation Act

In early 1994, parties in the German parliament started discussing some changes for the smaller stock firms to make the legal form of the stock corporation more attractive to small and medium-sized enterprises. Incorporating as a stock firm provided enterprises with the possibility to raise capital on the stock exchange, but at the same time required costly compliance with strict and extensive formal, accounting, and publicity regulations. At the time, small and medium-sized enterprises instead disproportionately opted for the legal form "GmbH" (in what follows, limited liability company or LLC).

The political discussions on potential changes to smaller stock firms started in February 1994. Codetermination was one of the points on which the political parties had divergent stances. The governing conservative-liberal coalition (CDU/CSU and FDP) proposed to abolish the codetermination mandate for all smaller stock firms so that this regulation is in line with the absence of the mandate for smaller LLCs and does not deter small and medium-sized enterprises from

[19] A stock firm is family-owned if it has a single shareholder who is a natural person or if all shareholders are related or related by marriage (§76 *BetrVG* 1952).

[20] In contrast to the firms in the coal and steel sector, if a vote in the supervisory board results in a tie, then a supervisory board chairman, who is a shareholder representative, holds a decisive vote (§29 *MitbestG* 1976). The 1952 and 1976 Acts do not apply to companies that directly and predominantly serve political, charitable, educational, scientific, artistic, or religious purposes.

choosing stock corporation as their legal form. The center-left opposition of Social Democrats (SPD), which held a majority in the Federal Council (*Bundesrat*), insisted on preserving the codetermination mandate for all small stock firms.

In late May 1994, the two sides agreed to preserve the codetermination mandate but only for the already incorporated small stock firms. The motivation for grandfathering was that the stock firms incorporated before the reform already learned how to operate with employee representatives on their boards. As the grandfathering aspect of the codetermination mandate was a rather late and unexpected result of a political compromise, it seems unlikely that firms expected the cohort-based change to codetermination. Section 1.4 provides further evidence on a potential strategic delay of incorporation by stock firms around the reform date.[21]

The final changes were implemented with the August 10, 1994 reform of the Stock Corporation Act. The reform changed the codetermination mandate for newly incorporated smaller stock firms. Stock firms with up to 500 employees are no longer mandated to have worker-elected representatives in the supervisory board if they incorporate on or after August 10, 1994. At the same time, it permanently locked in the already incorporated stock firms with one-third worker share on the board mandate.[22] In the empirical design, I exploit the grandfathering aspect of the reform and compare outcomes at stock firms incorporated slightly before the reform to those incorporated slightly after to estimate the effect of codetermination on innovation.

Importantly, only the change to codetermination was cohort-based while other changes applied to all smaller stock firms independent of their incorporation date. The reform did not introduce changes to firms with other legal forms or to other forms of worker representation. For *all* stock firms with up to 500 employees, it strengthened statutory autonomy concerning the appropriation of profits, allowed one-person founding, and introduced simplifications regarding the general meeting.

Despite additional changes to small stock corporations introduced with the reform, commentators suggested that the abolition of codetermination was the main change and the other policies

[21] *Drucksache* 12/6721 and 12/7848 of the *Deutscher Bundestag* (1994) report the initially proposed bill and the follow-up compromise grandfathering suggestion. The *Deutscher Bundestag* plenary proceedings 12/208, 12/233, and 672 report the minutes of plenary meetings.

[22] *"Gesetz für kleine Aktiengesellschaften und zur Deregulierung des Aktienrechts"* (1994)

were secondary.[23] Moreover, as the other changes applied to all smaller stock firms, these are netted out by the first difference when comparing different cohorts of firms in the empirical design.

Finally, stock firms that incorporate before the reform and are locked in with the codetermination mandate cannot easily evade it. A stock firm that temporarily changes its legal form and re-incorporates as a stock firm after the reform date is not released from the worker representation mandate on its supervisory board.[24] The grandfathering feature of codetermination for small stock firms was challenged in the Federal Constitutional Court on grounds of unequal treatment, which further illustrates the binding nature of the grandfathering. The Court ruled that the grandfathering remains in force.[25]

1.3 Data

In this Section, I discuss the data sources. First, I provide an overview of the data on firms, patents, and patent characteristics. Second, I present the summary statistics and descriptive evidence on worker share on the board and patent evolution.

1.3.1 Characteristics of firms and patents

My primary data set contains firm-level information on the number of patents filed every year between 1998 and 2014 by stock firms and LLCs incorporated in Germany in a five-year window around the August 10, 1994 reform.[26]

In order to obtain information on firms, I use the Bureau van Dijk's Orbis database. Bureau van Dijk's Orbis provides data on firms' financial characteristics and ownership structure and is based on official company reports and business registers. Importantly for this study, Orbis includes information on smaller and unlisted enterprises and can be merged with other data sets.

[23] See, e.g., "Nicht nur weiße Salbe", a commentary reported in *Frankfurter Allgemeine Zeitung* on May 27, 1994, which argues that change to codetermination is the only significant change to smaller stock firms introduced with the 1994 reform.

[24] See, for example, Raiser et al. (2015) for further details.

[25] *BVerfG*, the decision of the second Chamber of the First Senate from January 9, 2014 (1 BvR 2344/11).

[26] Patenting activity is limited by 2014 as this is the last available year with comprehensive coverage in the PATSTAT edition in use.

I collect information on firms' characteristics such as incorporation date, legal form, industry (4-digit NACE code) and employment for German firms that incorporate between 1989 and 1999.

I classify firms into stock corporations and LLCs using their legal form as reported in Orbis. Stock firms refer to "*Aktiengesellschaften*" and "*Kommanditgesellschaften auf Aktien*" and LLCs refer to "*Gesellschaften mit beschränkter Haftung (GmbH)*".[27] I exclude firms that serve political, charitable, educational, scientific, artistic or religious purposes as well as state-owned firms, media enterprises and fully family-owned firms, as these are mostly never subject to codetermination independent of their legal form and incorporation date.[28] I also drop firms located in East Germany. Appendix A.2 reports further details on sample construction.

Using the information on the firm legal form and date of incorporation, I classify the firms that incorporate within a five-year window around the 1994 reform into four groups: stock firms incorporated before August 10, 1994 (*treated*), stock firms incorporated on or after August 10, 1994 (*control*), and LLCs incorporated before the reform or on the reform day or afterward (additional control).

To estimate the effect on innovation, I merge the firm data with the patent data from PATSTAT. The matching process proceeds as follows. I first collect all patents filed by firms in the sample using Orbis Intellectual Property (Orbis IP) database.[29] Since Orbis IP tends to be less comprehensive than PATSTAT, I merge patent publications from the former to the latter and retrieve the most frequent "person id" among patents filed by each firm in the sample. "Person id" is the identification number for the standardized name of patent applicants and inventors in PATSTAT. I then collect all patent applications that correspond to the most frequent "person id" for each sample firm from PATSTAT, and this is the final set of firm-level patents in the data set.

As I am also interested in classifying patents into automation and non-automation, I use automation classification developed by Dechezleprêtre et al. (2022). The classification aims to detect patents that are potentially labor-substituting and allow for the replacement of employ-

[27] I classify a firm as a stock firm if its national legal form as reported in Orbis is "Public limited company – AG", "Limited partnership by shares – KGaA" or "Limited liability company & partnership by shares – GmbH & Co. KGaA". I classify a firm as an LLC if its national legal form is "Limited liability company – GmbH" or "Limited liability company & partnership – GmbH & Co. KG".

[28] These are excluded from the sample based on their primary industry code, firm name or Orbis ownership data. Fully family-owned firms are exempted from codetermination if they have fewer than 500 employees.

[29] Orbis and Orbis IP use a unified firm identifier.

ees in certain tasks. The method proceeds in two steps. First, the authors use the full text of patent applications from the European Patent Office and search for automation-related keywords in each patent application. Second, they compute the share of patents with at least one automation-related keyword in each technology category (as defined by CPC/IPC codes) and classify technology categories as automation based on this share. As such, they first classify patent technology categories and then patents. The two-step procedure uses the combined wording of patents within a technology class for the classification instead of relying on keywords in a single patent, which could give a weak signal for whether a patent corresponds to automation. In Appendix A.4, I classify patents directly using automation keywords and report the results using this alternative classification.

Finally, to study the mechanisms behind the effect of codetermination on innovation, I consider additional characteristics of patents such as backward and forward citations, inventor team size, and patent distance to science. I retrieve patent citation and inventor information from PATSTAT. The distance to science measure is based on patent-to-article and patent-to-patent citation linkages and follows Ahmadpoor and Jones (2017). Krieger et al. (2022) and Poege, Harhoff, Gaessler, and Baruffaldi (2019) show that patents closer to science have on average a higher private value.

Table 1.1 presents summary statistics of the data. The baseline sample contains 496 stock firms and 17,243 LLCs that incorporate within three years of the 1994 reform. Stock firms are on average larger in size and more innovative, as reflected by the median number of employees and the average number of patents. Stock firms that incorporate after the reform have on average a larger stock of patents and automation patents between 1998 and 2014 as compared to stock firms that incorporate before the reform. When adjusted to employment and averaged by year, this difference becomes less sizable for patents and more sizable for automation patents and persists if one considers the corresponding differences for LLCs.

Table 1.1: Summary Statistics

	Stock firms pre-reform	Stock firms post-reform	LLCs pre-reform	LLCs post-reform
Mean # patents, 1998-2014	1.274	2.472	.675	.849
Mean # automation patents, 1998-2014	.197	.396	.119	.147
Mean # patents per 100 empl	.063	.096	.046	.049
Mean # automation patents per 100 empl	.006	.022	.008	.010
Median # employees	50	46	27	28
N	208	288	8119	9124

Notes: "pre-reform" and "post-reform" refer to firms in Germany incorporated before and on or after the August 10, 1994 reform, respectively. The sample includes stock corporations and LLCs incorporated within three years of the reform. "Mean # patents, 1998-2014" and "mean # automation patents, 1998-2014" refer to the total number of firm patents or automation patents filed between 1998 and 2014, averaged by firm legal form and pre- or post-reform incorporation. "Mean # patents per 100 empl" and "mean # automation patents per 100 empl" refer to the mean number of firm patents or automation patents filed between 1998 and 2014 per year and per 100 employees, averaged by firm legal form and pre- or post-reform incorporation.

1.3.2 Worker share on the board

The 1994 reform gives rise to the variation in employee representation on company boards that I exploit in this study. One of the related concerns is that stock firms permanently locked in with codetermination do not (fully) comply with the one-third worker share mandate. Alternatively, stock firms that incorporate after the reform could appoint worker-elected representatives to their supervisory board even without a legal obligation to do so.

In the following, I follow Jäger, Schoefer, et al. (2021) and present evidence that the reform shifted the supervisory board composition for the affected stock firms using Hoppenstedt Aktienführer data. The Hoppenstedt Aktienführer provides yearly information on all listed German stock companies between 1956 and 2018.[30] For each year and every listed stock firm, it reports the names and roles of supervisory board members, which allows differentiation between worker representatives and shareholders on the board. I consider stock firms that incorporate within five years of the 1994 reform and use the data on supervisory board composition from the 1990s, as there is a structural break in reporting that starts in the 2000s. I further exclude firms located in East Germany and all firms that incorporate in the year 1994 since the Hoppenstedt Aktienführer reports only the year of firm incorporation and not the full date. Finally, I compute the yearly worker share on the board of every firm and keep all firm-years in which at least one

[30] I access the digitized version of the Aktienführer Data Archive which is maintained by the University of Mannheim, available at `https://digi.bib.uni-mannheim.de/aktienfuehrer/data/index.php`.

third of supervisory board members have non-missing information on their role (shareholder vs. employee representative).

Figure 1.1: Empirical worker share on the board

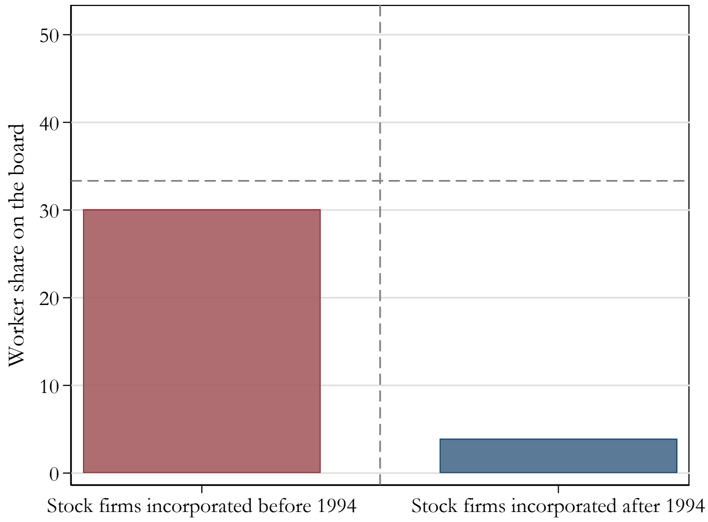

Notes: The Figure shows the average percentages of worker representatives on supervisory boards across listed stock firms that have up to 500 employees and incorporate within 5 years before (left) or after (right) the August 10, 1994 reform. The data on supervisory board composition is from the Hoppenstedt Aktienführer and is based on the 1990s editions, as there is a structural break in reporting in the 2000s. The Figure is based on all firm-years where at least one third of supervisory board members have non-missing information on their role (worker representative vs. shareholder). The sample excludes stock firms that incorporate in the year 1994, as the Hoppenstedt Aktienführer does not report the full incorporation date, and stock firms located in East Germany.

Figure 1.1 shows the average share of worker representatives on supervisory boards of smaller stock firms, i.e. stock firms with up to 500 employees, separately for the firms incorporated before and after the 1994 reform. The horizontal dashed line marks a one-third worker share. The share at stock firms locked in with codetermination is nearly one-third. The share at stock firms that incorporate after the reform and are no longer subject to the mandate is close to zero. The firms appear to comply with the codetermination mandate. In the absence thereof, the firms do not seem to allocate a sizable share of seats to employees.[31] This in line with the binding nature of the reform, as described in Section 1.2.[32]

[31] Appendix A.2 reports the analogous figure for stock firms that have more than 500 employees. There is no sizable difference in the worker share on the board among larger stock firms that incorporate before and after the reform. This is in line with the fact that the 1994 reform affected only smaller stock firms.

[32] As an additional intervention check, in Appendix A.2 I study how the reform changes the supervisory board composition.

1.3.3 Descriptive evidence

Before moving to the econometric analysis, I present some descriptive evidence on the evolution of patent and automation patent number in stock firms and LLCs.

Figure 1.2 shows the evolution of the number of patents per 100 employees between 1998 and 2014, averaged by the firm legal form and incorporation shortly before (blue) or after the reform date (red). On average, firms that incorporate after the 1994 reform appear to file a higher number of patent applications per employee. Figure 1.2 also suggests that stock firms that incorporate on the reform day or after (*control*) are more innovative than those that incorporate before the reform (*treated*) (solid), after accounting for the corresponding difference in LLCs (dash).[33]

As I am also interested in the effect of codetermination on automation, Figure 1.3 shows the mean automation patent stock by a quarter of firm incorporation relative to the reform date, separately for stock firms and LLCs. The vertical dashed line marks the reform date. Every marker in the figure corresponds to the average number of automation patents per 100 employees that firms file in total between 1998 and 2014, for a given legal form and quarter of incorporation. In addition to the mean automation patent stock, the figure includes fitted lines separately for stock firms and LLCs and the incorporation period before or after the reform. There is little change in automation patenting between LLCs that incorporate before and after the 1994 reform (light blue line). This is different for stock firms. Figure 1.3 suggests that stock firms that incorporate before the reform, locked in with codetermination, file on average fewer automation patents per employee (red). Figure 1.3 also shows that this difference persists in the majority of the bandwidths around the reform date.

This is only descriptive evidence. In the empirical specification, I further test whether the results remain similar when I use econometric analysis.

[33]Figure 1.2 shows that the trends in patent evolution between 2008 and 2011 are not identical. As a robustness check, in Appendix A.4 I re-estimate the baseline specification where I restrict the time period to the window between 1998 and 2007. The results remain robust to using this shorter time window.

Figure 1.2: Evolution of employee-adjusted number of patents

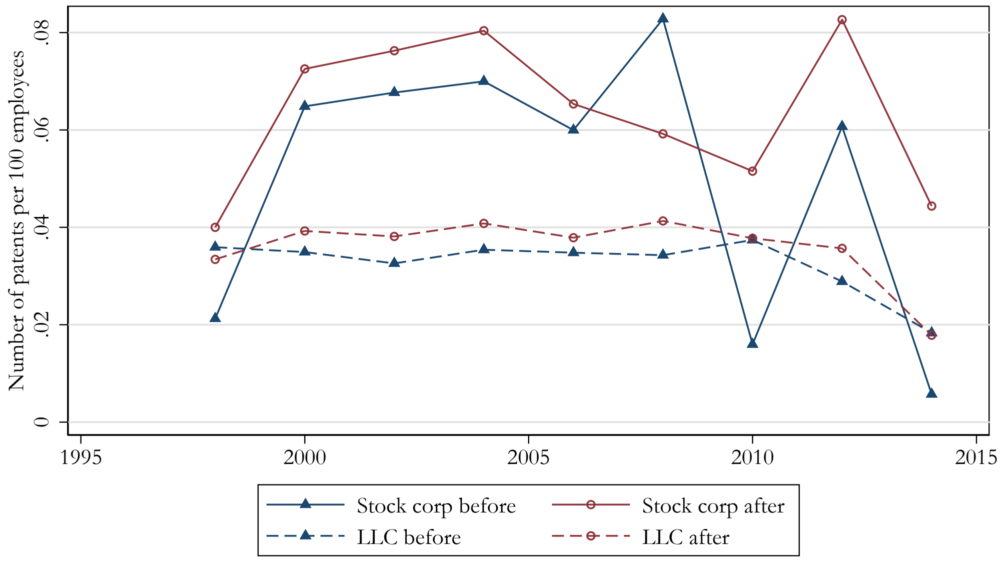

Notes: The Figure shows the evolution of patent number per 100 employees, averaged by firm legal status and incorporation date before or after the reform for the period between 1998 and 2014. The solid lines show the patent evolution for stock corporations, and the dashed lines show the analogous evolution for LLCs. The sample contains stock firms and LLCs that incorporate within three years around the 1994 reform.

Figure 1.3: Mean automation patent stock by quarter of firm incorporation

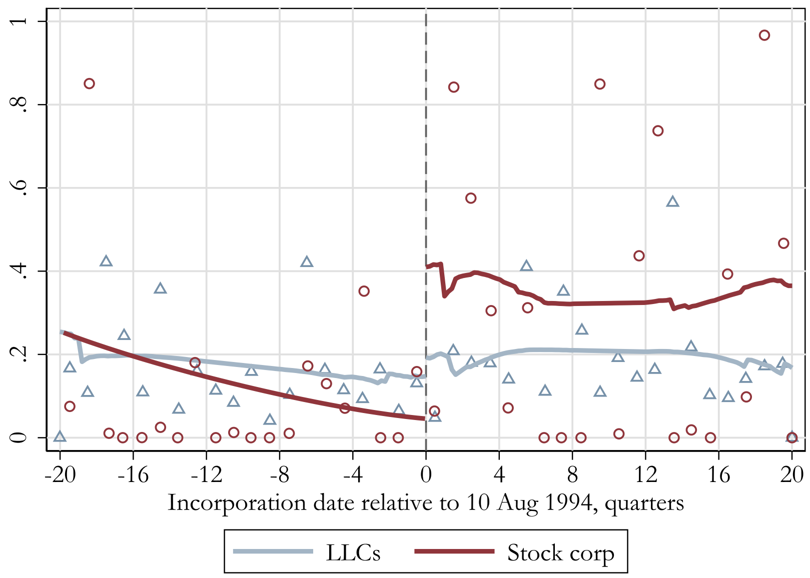

Notes: The Figure shows the mean automation patent stock by a quarter of firm incorporation relative to the reform date, separately for stock firms and LLCs. Every marker shows the average automation patent stock per 100 employees that the firms incorporated in a given quarter and with a given legal form file on average between 1998 and 2014 (in total). The vertical dashed line marks the reform date. The Figure includes fitted lines, separately for stock firms and LLCs and for the incorporation period before and after the reform date. The sample contains stock firms and LLCs that incorporate within a three-year window around the 1994 reform.

1.4 Empirical Strategy

In an ideal world, a regression that relates firm patenting to the presence or a share of employee representatives on the firm board would estimate the effect of shared governance on firm innovation. However, the estimate from such regression is likely to be biased as firms endogenously choose whether or which share of worker representatives to appoint to their boards. More productive firms may have more resources to elect board directors among employees and outside of the usual candidates' pool. Higher performing firms may also have on average stronger works council presence, which may exert pressure to appoint employee representatives to the firm board.

To address this concern, I exploit the 1994 reform and the legal change to worker representation on corporate boards of newly incorporated smaller stock firms.[34]

1.4.1 The effect of codetermination

I start by estimating the following simple difference specification:

$$E[Y_{it}|X_{it}] = exp[\beta_0 + \beta_1 Pre\text{-}reform_{it} + X'_{it}\gamma + \epsilon_{it}], \tag{1.1}$$

where Y_{it} is a number of patents or automation patents a stock firm i files in a year t per 100 employees, $Pre\text{-}reform_{it}$ denotes an indicator variable that is one for firms that incorporate before the 1994 reform and are locked in with the codetermination mandate, and X_{it} are year or year by two-digit NACE industry fixed effects. In this simple difference specification, I only consider stock firms that incorporate within three years before and three years after the August 10, 1994 reform, and compare the measure of innovation output at these firms between 1998 and 2014.[35]

The dependent variable, the number of firm patents or automation patents per 100 employees, is skewed and nonnegative.[36] I estimate Equation 1.1 using a Poisson pseudo-likelihood regres-

[34] The analysis is placed in the same setting and is partially based on Jäger, Schoefer, et al. (2021).

[35] In all years between 1998 and 2014, I assume a firm has zero patents if I do not observe firm patent applications in PATSTAT and a firm is reported in the 2020 Orbis database snapshot.

[36] 15 percent of stock firms that incorporate within three years of the 1994 reform in the sample file at least one patent application between 1998 and 2014.

sion with multiple fixed effects following Correia, Guimarães, and Zylkin (2020).[37] I cluster standard errors at a firm level.

The coefficient β_1 in Equation 1.1 captures a difference in the 1998 to 2014 number of firm patents or automation patents per 100 employees between the stock firms that incorporate shortly before and shortly after the 1994 reform, and so are either permanently with or without the codetermination mandate. All changes that the smaller stock firms undergo due to the 1994 reform apply to all these firms independently of the incorporation date, and only the codetermination mandate varies between the two groups. In addition, year or industry by year fixed effects capture shocks or industry-year specific shocks, such as the development of a breakthrough technology upon which follow-on innovation builds intensively, that identically affect stock firms incorporated before and after the reform.

The stock firms that incorporate within three years before and after the reform could still differ in innovation outcomes due to the factors related to the differential time of incorporation or firm age. To account for incorporation time, cohort, and age effects, I use the cohorts of limited liability companies incorporated within a three-year window around the 1994 reform as an additional control.

I estimate the following difference-in-differences specification:

$$E[Y_{it}|X_{it}] = exp[\beta_0 + \beta_1 Pre\text{-}reform_{it} \times Stock_{it} + \beta_2 Pre\text{-}reform_{it} + \beta_3 Stock_{it} + X_{it}'\gamma + \epsilon_{it}], \quad (1.2)$$

where $Stock_{it}$ is an indicator variable that is one for stock firms and zero for LLCs, and the other variables are as in Equation 1.1. β_2 captures the effect of incorporating before the reform and absorbs factors related to incorporation time as business cycle effects. β_3 captures the baseline effect of incorporating as a stock firm. This reflects the differences in patenting among stock firms and LLCs. The coefficient β_1 corresponds to the effect of legally mandated worker representation (one-third worker share) on supervisory boards of stock firms.

I estimate Equation 1.2 using a Poisson pseudo-likelihood regression with multiple fixed effects. Similar to the simple difference specification, in the baseline I consider the firms that incorporate within a three-year window around the reform, and compare innovation outcomes between 1998

[37] Appendix A.4 reports the results using inverse hyperbolic sine transformation and OLS regression.

and 2014. In the estimation, I control for year or year by two-digit NACE industry fixed effects and cluster standard errors at a firm level.

The identifying assumption is that without the 1994 reform, the difference in the number of employee-adjusted (automation) patents between stock firms that incorporate shortly before and shortly after the 1994 reform would be identical to the corresponding difference between limited liability companies. Thus, the assumption is not that stock firms and LLCs are similar along dimensions related to firm innovation. Table 1.1 shows that stock firms and LLCs do differ. Stock firms are on average larger in size and more innovative, as measured by the median number of employees and the average number of patents. Rather, the trends in patenting should be parallel.[38]

1.4.2 Validity of the empirical design

One of the potential concerns with the empirical design could be that stock firms strategically delay incorporation until after the August 10, 1994 reform date to evade codetermination. To test for this, in Figure 1.4 I plot the density for newly incorporated stock firms by month of firm incorporation relative to the 1994 reform date, before and after the reform cutoff date. The vertical dashed line marks the reform date. There is no evidence of bunching in the density of incorporation as a stock firm after the 1994 reform.

Figure 1.4 further reports a McCrary (2008) test of continuity of the density at the 1994 reform cutoff date against the alternative of a jump in the density function. A statistically significant discontinuity estimate would suggest that there is a discontinuity in the density around the reform cutoff date and that firms manipulate their incorporation date. The McCrary discontinuity estimate is 0.243 with a standard error of 0.416. There is little evidence that stock firms strategically delay their incorporation date. This is in line with the 1994 reform proceedings, which indicate that the cohort-based nature of a change to codetermination was a rather unexpected

[38]An alternative empirical design is to compare patent evolution at treated and control stock firms before and after they incorporate. This is not feasible as only 4 percent of stock firms file at least one patent application before incorporation. It is also not clear whether stock firms do not patent or are not yet established in firm-years with zero patent observations before their incorporation, as Orbis reports only the incorporation date.

result of a last-minute political compromise. In addition, the change to codetermination came into force the day after the reform was made public.[39]

Figure 1.4: Strategic delay of incorporation (McCrary test)

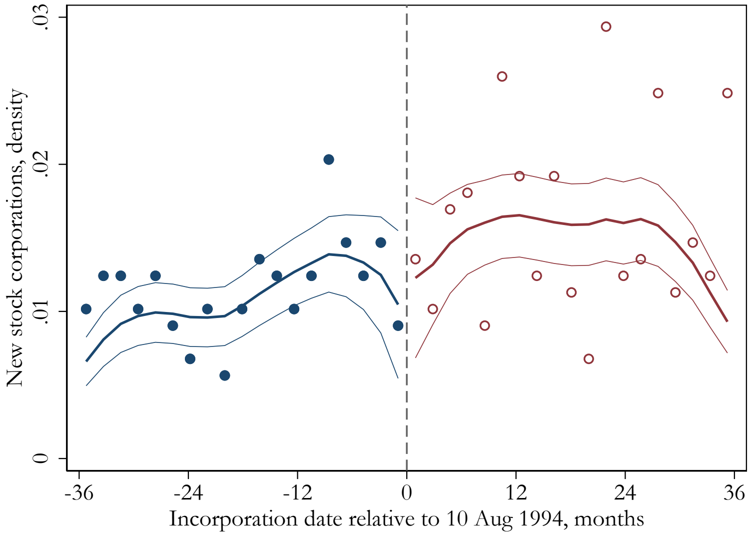

Notes: The Figure shows the density for newly incorporated stock firms by month of incorporation and relative to the 1994 reform date, before and after the cutoff date. The vertical dashed line marks the August 10, 1994 reform cutoff date. The Figure additionally reports a McCrary test (2008) of density continuity at the reform cutoff against the alternative of a jump in the density function. The sample of stock firms is based on firms that incorporate as a stock corporation within a three-year window of the 1994 reform and which are reported in the 2020 Orbis database snapshot. The stock firms located in East Germany are excluded.

Another potential concern with the empirical design could be that more firms choose stock corporation instead of LLC as their legal form after the reform. Figure 1.5 shows the percent of stock firms in a sample of both newly incorporated stock firms and LLCs, by time of incorporation relative to the August 10, 1994 reform. The vertical line marks again the reform cutoff date. The probability of incorporating as a stock corporation does not appear to change significantly around the reform cutoff date. In addition, I test for this in a more formal way and run a regression of a dummy of incorporation as a stock firm on the indicator for incorporation before the reform, a time trend and the interaction of the two. I find a small trend towards incorporating as a stock firm over time, but no level shift or trend change between the pre- and

[39]The reform was promulgated on August 9, 1994. All changes stipulated by the reform came into force on August 10, 1994.

post-reform period.[40] There seems to be little evidence that firms disproportionately select into stock corporation status instead of LLC legal form after the reform date.[41]

Figure 1.5: Selection into stock corporation status

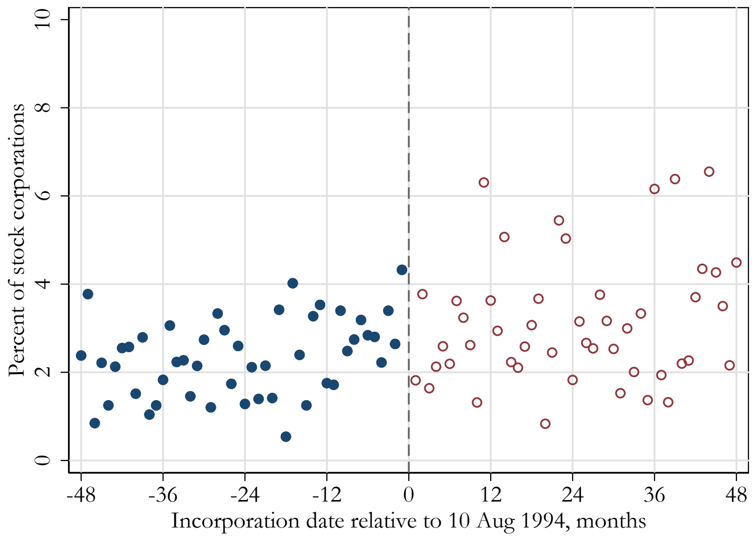

Notes: The Figure shows the percent of firms that incorporate as a stock corporation in a sample of both newly incorporated stock firms and LLCs, by month of incorporation and relative to the 1994 reform date. The vertical line marks the reform cutoff date. The sample of stock firms is based on firms that incorporate as a stock corporation or an LLC within a four-year window of August 10, 1994 and which are reported in the 2020 Orbis database snapshot. The stock firms and LLCs located in East Germany are excluded.

As robustness, I estimate Equation 1.2 using alternative bandwidths around August 10, 1994, running from a one-year and up to a five-year window around the reform date. While I cannot test directly the identifying assumption behind Equation 1.2, I further estimate the baseline specification using a placebo reform date on August 10, 1996 or 1997 and restricting the sample to the always untreated stock firms along with their peer cohorts of LLCs. This informs about potentially differential trends between the stock firms and LLCs or the lifecycle effects, for example, whether older stock firms are always less innovative than younger stock firms as compared to the analogous difference between the slightly older and younger cohorts of LLCs. I also estimate the codetermination effect on automation innovation using alternative automation classification, which is based on automation keywords, and restricting the automation patent sample to process patents. In addition, I repeat the baseline difference-in-differences analysis

[40]The output table is reported in Appendix A.3.

[41]This also suggests that the reform seemingly did not achieve its goal of making the stock legal form more attractive for firms as compared to the alternative form of an LLC.

using the inverse hyperbolic sine transformation of the outcome variable and the OLS regression. The results are reported in Appendix A.4 and are in line with the main findings.

Finally, to address potential compositional changes among firms after the reform, I control for year by two-digit industry fixed effects in the baseline specification. Appendix A.3 further reports the change to the industry composition of firms after the reform using a simple difference and a difference-in-differences specification. Depending on the specification, the estimates have a p-value of .491 and .381 in an F-test and are jointly not significant. The reform does not appear to lead to a significant change in the industry composition of firms.

1.5 Results

Descriptive evidence suggests that stock firms locked in with codetermination file a smaller number of employee-adjusted automation patents, whereas there is a less sizable negative effect on the number of patents overall.

In this Section, I present the results for the average treatment effect of worker representation in firm boards on firm innovation, distinguishing between patents, automation and non-automation patents. I start with an overview of the results using a simple difference specification estimated in a sample of stock firms before moving to the difference-in-differences design. I then show how codetermination effects vary when estimated using different bandwidths around the 1994 reform date.

1.5.1 Effect of codetermination

Table 1.2 reports the estimation results using the simple difference specification described in Equation 1.1. The estimation is based on the sample of stock firms that incorporate within three years pre- and post-reform. Columns (1) and (2) include the effect of codetermination on the number of firm patents per 100 employees with a control for year and year by two-digit NACE industry fixed effects, respectively. Columns (3) and (4) include the results where the sample of firm patents is restricted to automation, similarly with a control for year or year by industry fixed effects. Columns (5) and (6) report the estimates for the effect of codetermination on non-automation innovation, correspondingly with year or year by industry fixed effects.[42]

The results in Columns (1) and (2) of Table 1.2 show that the shared governance of firms is associated with a decline in the number of patents a firm files per 100 employees. The number of patents per 100 employees decreases by 52 percent when estimated with year fixed effects, and the effect declines in magnitude to 44 percent when estimated with year by industry fixed effects.[43] The estimates are sizable, but are also imprecise and not statistically significant at a 10 percent significance level.

Decomposing the overall codetermination effect into the effect on automation and non-automation in Columns (3) to (6) suggests that the negative effect is primarily driven by a negative change in automation innovation. Columns (3) and (4) show that the number of firm automation patents per 100 employees declines by 74 and 70 percent, depending on the fixed effects in estimation, in the stock firms permanently locked in with the shared governance. These estimates are both economically and statistically significant.

Table 1.2 further shows that there are no large differences in the estimates between the estimation with year and year by industry fixed effects. This is in line with the absence of a significant change to the industry composition after the 1994 reform date.

Next, I estimate the effect of shared governance on firm innovation using the difference-in-differences design described in Equation 1.2. Table 1.3 shows the estimates for the effect on firm innovation and the decomposition into automation and non-automation, where the sample additionally includes the cohorts of LLCs incorporated within a three-year window around the 1994 reform date. Columns (1) and (2) report estimates of the codetermination effect on firm number of patents per 100 employees. Columns (3) and (4) restrict the sample to automation patents. Columns (5) and (6) report the effects, where the sample is restricted to non-automation patents. In Columns (1), (3) and (5), the estimation includes a control for year fixed effects. Correspondingly, Columns (2), (4) and (6) include the results from the estimation with year by two-digit industry fixed effects.

[42]Patent applications that are not classified as automation following Dechezleprêtre et al. (2022) classification are assumed to be related to non-automation.

[43]Table 1.2 and Table 1.3 report the results using a Poisson pseudo-likelihood regression. ($\exp[\hat{\beta}] - 1$) converts the estimates to elasticities, where $\hat{\beta}$ is the coefficient reported in the table.

Table 1.2: Effect of shared governance on firm innovation (simple difference)

	Patents		Automation patents		Non-automation	
	(1)	(2)	(3)	(4)	(5)	(6)
Pre-reform	-0.741	-0.572	-1.344***	-1.189**	-0.655	-0.486
	(0.483)	(0.479)	(0.489)	(0.497)	(0.520)	(0.512)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Year × industry FE	No	Yes	No	Yes	No	Yes
Observations	8,432	4,306	8,432	2,378	8,432	3,750
Log-likelihood	-3970.395	-2537.038	-710.957	-407.271	-3509.205	-2211.015

Notes: Columns (1) and (2) include estimates of the effect of codetermination on the yearly number of firm patent applications per 100 employees between 1998 and 2014. Columns (3) and (4) include estimates of the effect on the number of firm automation patents per 100 employees. Columns (5) and (6) report estimates of the effect on the number of firm non-automation patents per 100 employees. The estimation is based on Equation 1.1 and uses the sample of stock firms incorporated within a three-year window around August 10, 1994. All estimates are from a Poisson pseudo-likelihood regression. Standard errors are clustered at a firm level and are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Netting out incorporation time, cohort and age effects with an additional control group of LLCs does not lead to a sizable change in the results. Columns (1), (2), (5) and (6) in Table 1.3 show that employee representation on firm boards is associated with a decline in firm innovation in general and non-automation innovation. These effects are sizable but are again imprecise and not statistically significant at a 10 percent significance level. The negative effect on automation innovation is economically and statistically significant. Columns (3) and (4) show that shared governance decreases employee-adjusted number of firm automation patents by 76 to 79 percent.

Finally, similarly to Table 1.2, there are no sizable changes between the results in Columns (1), (3) and (5), which include year fixed effects, and Columns (2), (4) and (6), which additionally control for year by industry fixed effects.

The results using a simple difference specification and difference-in-differences design are quite similar. In the remainder of this chapter, the estimates are always based on the difference-in-differences design as described in Equation 1.2.

Overall, I find that codetermination leads to a reduction in the number of firm employee-adjusted automation patents. The effect on general innovation and non-automation is also negative, but smaller in size and imprecise.

Table 1.3: Effect of shared governance on firm innovation (DID)

	Patents		Automation patents		Non-automation	
	(1)	(2)	(3)	(4)	(5)	(6)
Pre-reform × Stock	-0.782	-0.882	-1.433**	-1.571**	-0.701	-0.795
	(0.534)	(0.536)	(0.646)	(0.653)	(0.559)	(0.560)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Year × industry FE	No	Yes	No	Yes	No	Yes
Observations	301,563	248,505	301,563	181,949	301,563	243,123
Log-likelihood	-112743	-96500	-30841	-25230	-89304	-76964

Notes: Columns (1) and (2) include estimates of the effect of codetermination on the yearly number of firm patent applications per 100 employees between 1998 and 2014. Columns (3) and (4) include estimates of the effect on the number of firm automation patents per 100 employees. Columns (5) and (6) report estimates of the effect on the number of firm non-automation patents per 100 employees. The estimation is based on Equation 1.2 and uses the sample of both stock firms and LLCs incorporated within a three-year window around August 10, 1994. All estimates are from a Poisson pseudo-likelihood regression. Standard errors are clustered at a firm level and are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

1.5.2 Different bandwidths

Table 1.2 and Table 1.3 show the results using a sample of firms that incorporate within a three-year window around August 10, 1994. In the following, I report estimates for the effect of codetermination on firm innovation using alternative bandwidths around the reform date.

Figure 1.6 shows the effect of codetermination on the number of firm patents per 100 employees, using the sample of firms that incorporate within 12 to 60 months around August 10, 1994. I estimate the effects with a difference-in-differences design as described in Equation 1.2 and using Poisson pseudo-likelihood regression. All underlying specifications include year by two-digit NACE industry fixed effects, and standard errors are clustered at a firm level.

Figure 1.6 reports the estimates under different bandwidths along with the 90 percent confidence intervals. The baseline estimate, reported in Table 1.3, is highlighted in black color.[44] The estimates remain relatively stable and similar to the baseline for a majority of alternative bandwidths around the reform date. The confidence intervals mostly do not exclude zero.

[44]The figure reports the coefficient estimates of pseudo-maximum likelihood Poisson specification and not the associated elasticities.

Figure 1.6: Effect of shared governance on firm innovation

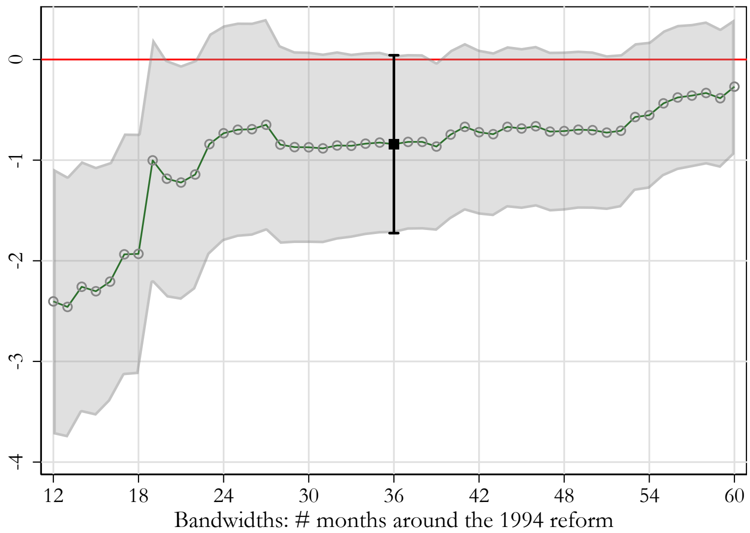

Notes: The Figure shows the effect of codetermination on the number of patents per 100 employees, using different bandwidths around the 1994 reform date. The Figure includes 90% confidence intervals along with the estimates. The effects are estimated using a difference-in-differences design, where a difference in the number of patents per 100 employees filed between 1998 and 2014 at stock firms affected and not affected by codetermination is compared to the corresponding difference for their peer cohorts of not affected LLCs. All specifications include year by two-digit NACE industry fixed effects. Standard errors are clustered at a firm level. Black highlights the baseline estimate reported in Table 1.3.

Figure 1.7 and Figure 1.8 show the analogous results for the effects of codetermination on the number of firm automation and non-automation employee-adjusted patents. Under a majority of alternative bandwidths, the estimates remain stable and similar to the baseline as reported in Table 1.3 and highlighted in black in the figures. Figure 1.7 shows that employee representation on supervisory boards is mostly associated with a negative and statistically significant effect on firm automation innovation. Figure 1.8 shows that the effect of shared governance on non-automation innovation is also negative but noisy under a majority of bandwidths. In Figure 1.8, the confidence intervals mostly do not exclude zero.

Appendix A.4 reports analogous results when the estimation is based on a simple difference specification. Overall, the baseline effects reported in Table 1.3 remain robust to a majority of alternative bandwidths around the reform when using both difference-in-differences and a simple difference specification.

In Appendix A.4, I report additional evidence on the robustness of the results: (1) I estimate the baseline Equation 1.2 using a placebo reform date on August 10, 1996 or 1997; (2) report the

Figure 1.7: Effect of shared governance on automation innovation

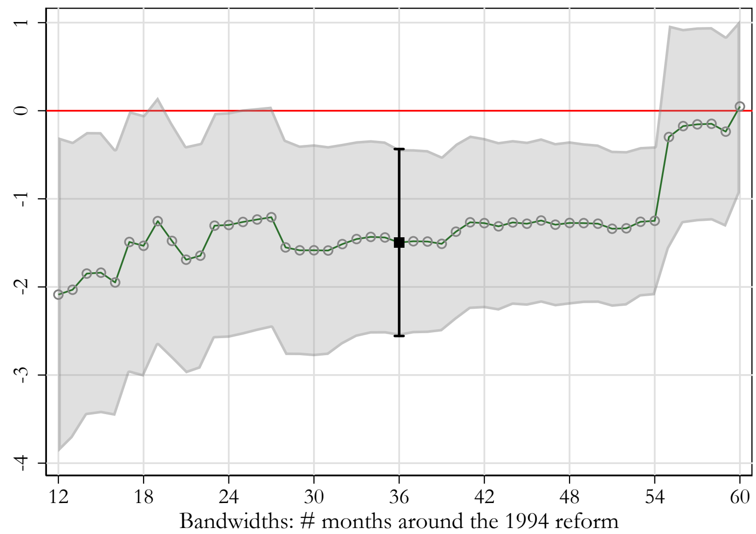

Notes: The Figure shows the effect of codetermination on the number of automation patents per 100 employees, using different bandwidths around the 1994 reform date. The Figure includes 90% confidence intervals along with the estimates. The effects are estimated using a difference-in-differences design and following Equation 1.2. All specifications include year by two-digit NACE industry fixed effects. Standard errors are clustered at a firm level. Black highlights the baseline estimate reported in Table 1.3.

effects on firm automation innovation using an alternative, keyword-based, automation classification; (3) estimate the codetermination effect based on the sample of automation patents that is constrained to process automation; and (4) use the inverse-hyperbolic sine transformation of the outcome variables and estimate the effects with an OLS regression.[45] The results hold in all of these alternative specifications, and there are no sizable effects of the placebo reforms on firm (automation/non-automation) innovation.

[45]One of the concerns with the estimation of the codetermination effect on firm automation innovation could be that firms license away their automation patents and do not use the newly patented automation technologies in-house. In one of the robustness checks, I restrict the sample of automation patents to process automation innovation, which likely serves as a better proxy for automation innovation developed for internal use (Danzer et al., 2020; Klepper, 1996).

Figure 1.8: Effect of shared governance on non-automation innovation

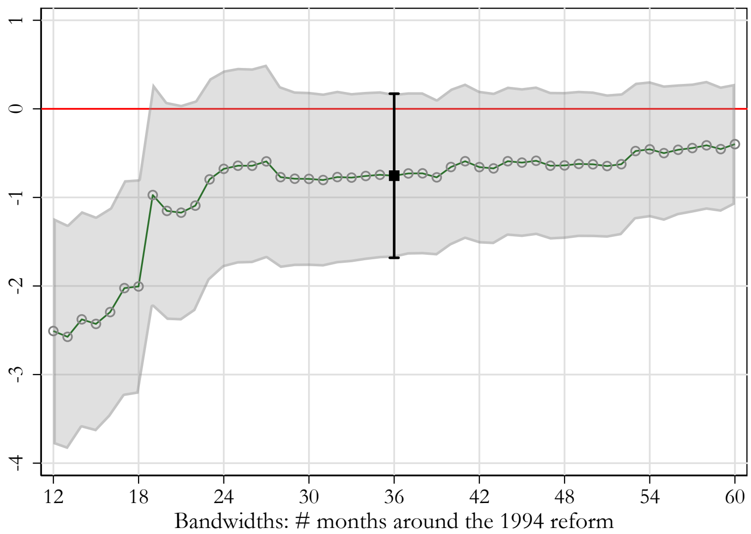

Notes: The Figure shows the effect of codetermination on the number of non-automation patents per 100 employees, using different bandwidths around the 1994 reform date. The Figure includes 90% confidence intervals along with the estimates. The effects are estimated using a difference-in-differences design and following Equation 1.2. All specifications include year by two-digit NACE industry fixed effects. Standard errors are clustered at a firm level. Black highlights the baseline estimate reported in Table 1.3.

1.6 Mechanisms

Section 1.5 showed that the shared governance of firms is associated with a negative, but not statistically significant, effect on innovation and non-automation innovation, and with an economically and statistically significant negative effect on automation innovation. In this Section, I provide further evidence on how these effects vary along patent relatedness to science, number and age of patent backward citations, and the size of inventor teams behind these patents. This helps us understand the mechanisms behind the effect of shared governance on general innovation and automation innovation.

1.6.1 Distance to science

The negative effects of codetermination on innovation and automation innovation capture a decline in the number of patents but not necessarily in their quality or value.

To understand how codetermination affects patent value, I begin by studying the effect heterogeneity with respect to patent distance to science. As patents closer to science are on average more valuable (Krieger et al., 2022; Poege et al., 2019), this provides a first insight into whether the shared governance of firms also leads to a decline in the value of (automation) patents.

To classify how closely patents relate to science, I use the patent-to-article citation data by Marx and Fuegi (2020) and build on the patent distance to science measure developed by Ahmadpoor and Jones (2017). Patents that directly cite a scientific article have a distance of one. Patents that do not cite a scientific article directly but cite a patent that cites a scientific article have a distance of two, and so on. I apply this distance to science measure to every patent filed by stock firms and LLCs in my sample between 1998 and 2014. I then compute the total number of patents filed by every firm in each year between 1998 and 2014 by their distance to science and estimate the codetermination effect following Equation 1.2 on the number of employee-adjusted (automation/non-automation) patents with a distance to science of one, two, three, and four or above.

Figure 1.9: Effect of shared governance on firm innovation by distance to science

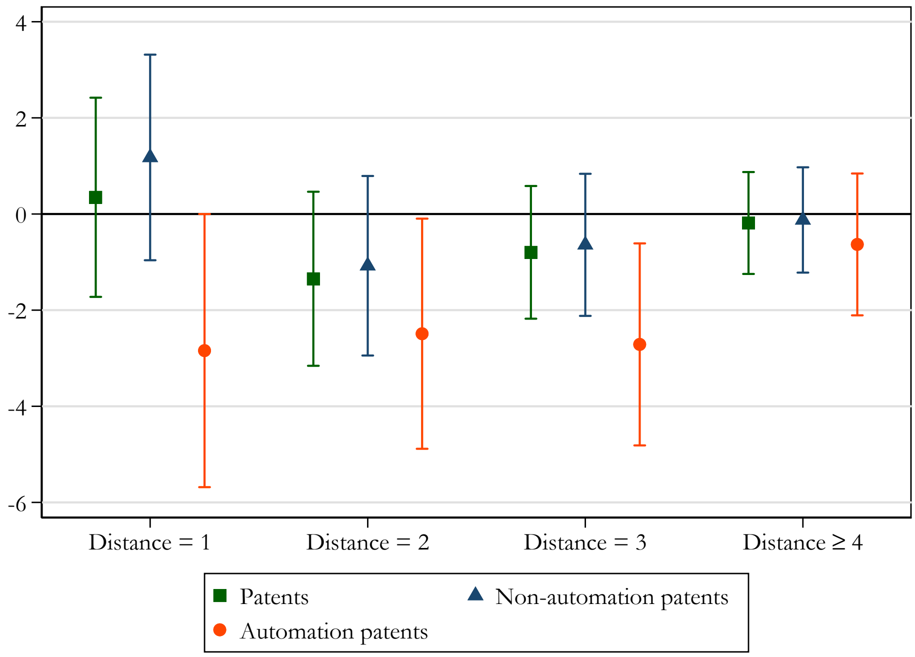

Notes: The Figure shows the codetermination effect on the number of patents, automation and non-automation patents per 100 employees and filed between 1998 and 2014, separately for patents with a distance to science of one, two, three, and four and above. Effects are estimated using a difference-in-differences design, as described by Equation 1.2, and Poisson pseudo-likelihood regression. All underlying specifications include year by two-digit industry fixed effects. Standard errors are clustered at a firm level. The Figure reports 95 percent confidence intervals along with the estimates.

Figure 1.9 reports the effect of shared governance on firm employee-adjusted number of patents, automation and non-automation patents, separately for patents that build directly on science and are more distant from science.[46] I estimate the effects using a difference-in-differences design and Poisson pseudo-likelihood regression. All underlying specifications include year by two-digit industry fixed effects, and standard errors are clustered at a firm level. Figure 1.9 reports 95 percent confidence intervals along with the estimates.[47]

Figure 1.9 shows that the negative effect of codetermination on automation innovation is driven by the decline in automation patents that build directly on science or are related to science. Shared governance of firms leads to the largest decline in automation patents that are science-based and directly cite a scientific article (distance of one). The effects on automation patents with a distance to science of two and three are similarly economically and statistically significant. This is different for the effects on the number of all employee-adjusted firm patents and non-automation patents. Figure 1.9 shows that the effect on non-automation patents that build directly on science is positive but imprecise.

Overall, I find that shared governance of firms is associated with less science-based automation patenting and more science-based non-automation patenting. This suggests that codetermination leads to a decline in the number of particularly valuable automation patents, as patents that directly build on science are on average more valuable (Arora, Belenzon, and Suh, 2022; Krieger et al., 2022; Poege et al., 2019).

1.6.2 Backward citations

Next, I study how codetermination affects the number and age of backward citations of patents filed by firms locked in with the mandate.

Figure 1.10 shows the effect of shared governance on the average number of backward citations per patent for firm automation patents filed between 1998 and 2014. Figure 1.11 shows the anal-

[46] The estimates for a distance of one are based on the full sample of stock firms and LLCs and their patents filed between 1998 and 2014. In the estimation of the effect on the number of patents with a distance of two, I restrict the sample to firm-year observations with no distance of one patents. Similarly, the estimation for patents with a distance of three (four) excludes firm-year observations with a non-zero distance of one and two (one, two, and three) patents. This implies that the samples used in the estimation of the results displayed in Figure 1.9 change depending on the distance to science.

[47] Figure 1.9 reports the coefficient estimates of pseudo-maximum likelihood Poisson regression and not the associated elasticities.

ogous results for non-automation. Both figures report estimates with the 90 percent confidence intervals, using different bandwidths around the August 10, 1994 reform. Black highlights the baseline estimate in a sample of stock firms and LLCs that incorporate within three years pre- and post-reform. The estimation follows the difference-in-differences specification as described in Equation 1.2 and includes a control for year by two-digit industry fixed effects. I cluster standard errors at a firm level. Finally, I restrict the sample to firm-year observations with a non-zero number of automation or non-automation patents. The estimates thus reflect results on the intensive margin.

Figure 1.10 shows that codetermined stock firms tend to file automation patents with a smaller number of backward citations. The estimates are noisily estimated but are negative under a majority of bandwidths. The baseline estimate indicates that the shared governance of firms leads to a 29.4 percent decline in the average number of backward citations per patent. This is mostly not the case for non-automation patents. Figure 1.11 shows that under a majority of bandwidths, the estimates for non-automation patents are close to zero. In the baseline, shared governance of firms is associated with a 2.7 percent decline in the average number of backward citations per non-automation patent.[48]

Table 1.4 shows the effect of codetermination on the average age of backward citations in patents filed by codetermined stock firms between 1998 and 2014.[49] The estimation follows the baseline difference-in-differences specification from Equation 1.2. Columns (1) and (2) include the effect of shared governance on the average age of backward citations in all patents, with year or year by two-digit industry fixed effects. Columns (3) and (4) restrict the sample to automation patents. Columns (5) and (6) show the results for non-automation patents. I exclude firm-year observations with zero (automation/non-automation) patents and estimate the effects on the intensive margin.

The results in Columns (3) and (4) of Table 1.4 show that codetermination leads to an increase in the average age of backward citations for automation patents. The average age of backward citations in automation patents increases by 62 percent when the estimation includes year by industry fixed effects.[50] The effect is sizable and statistically significant. Columns (2) and (6)

[48]Appendix A.5 reports the table with the corresponding estimates for a baseline three-year window around the reform.

[49]Age of backward citations is defined as the difference between filing dates of citing and cited patents.

[50]The average age of backward citations in automation patents filed by untreated stock firms is 12 years.

Figure 1.10: Effect of shared governance on average number of backward citations per patent, automation

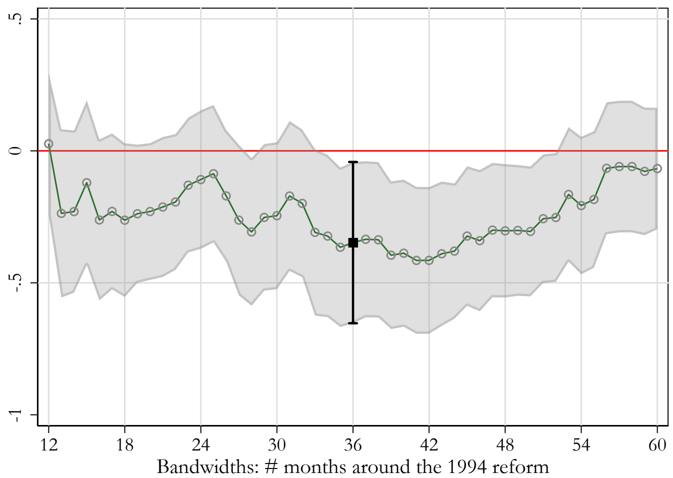

Figure 1.11: Effect of shared governance on average number of backward citations per patent, non-automation

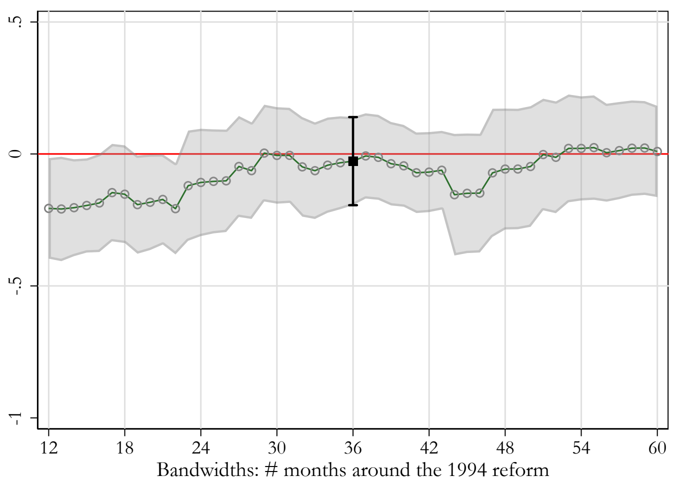

Notes: The Figure shows the effect of codetermination on the average number of backward citations per patent for automation patents in the top panel and non-automation patents in the bottom panel, using different bandwidths around the August 10, 1994 reform. The Figure reports 90% confidence intervals along with the estimates. The effects are estimated using a difference-in-differences design as described in Equation 1.2. All specifications include year by two-digit NACE industry fixed effects and are estimated using pseudo-maximum likelihood Poisson regression. Standard errors are clustered at a firm level.

Table 1.4: Effect of shared governance on average age of backward citations

	Patents		Automation patents		Non-automation	
	(1)	(2)	(3)	(4)	(5)	(6)
Pre-reform × Stock	0.043	0.181	0.201	0.483**	0.060	0.182
	(0.113)	(0.114)	(0.258)	(0.196)	(0.111)	(0.119)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Year × industry FE	No	Yes	No	Yes	No	Yes
Observations	4,953	4,788	1,347	1,200	4,395	4,235
Log-likelihood	-19598.349	-17656.508	-5271.987	-4235.365	-17643.692	-15696.781

Notes: Columns (1) and (2) include estimates of the effect of codetermination on the average age of backward citations in firm patent applications filed between 1998 and 2014. Columns (3) and (4) include estimates of the effect on the average age of backward citations in automation patents. Columns (5) and (6) report estimates of the effect on the average age of backward citations in non-automation patents. Firm-year observations with zero (automation/non-automation) patents are excluded. The estimation is based on Equation 1.2 and uses the sample of both stock firms and LLCs incorporated within a three-year window around August 10, 1994. All estimates are from a Poisson pseudo-likelihood regression. Standard errors are clustered at a firm level and are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

show that the effect for all patents and non-automation patents is much smaller. Non-automation patents filed by the codetermined stock firms include backward citations that are on average 6.2 to 19.9 percent older, depending on fixed effects in the estimation. These effects are not statistically significant at a 10 percent significance level.

These findings relate again to patent value. Patents closer to science and more valuable patents tend to have a higher number of backward citations and rely on a larger base of prior art. In addition, their cited prior art tends to be younger, as measured by the age of backward citations (Harhoff, Scherer, and Vopel, 2003; Krieger et al., 2022). This suggests that codetermination leads to a decline in the value of automation patents filed by the locked-in stock firms, as reflected by the effects on the number and age of patent backward citations. This is not so for non-automation patents, where the effects are not statistically significant and closer to zero.[51]

1.6.3 Size of the inventor teams

Finally, I study how codetermination affects the way innovation in a firm is organized by estimating the effect on the inventor team size.

[51] Appendix Table A.9 reports the effect of codetermination on the number of forward citations per patent, which further informs about the effect on patent quality. The results are in line with the findings related to patent value.

Table 1.5 reports the effect of shared governance on the average inventor team size for patents filed by the codetermined stock firms between 1998 and 2014.[52] Columns (1) and (2) include the effects on the average number of inventors per patent for all patents. Columns (3) to (6) decompose the effect into average size of inventor teams for automation and non-automation patents. The estimation follows the difference-in-differences specification from Equation 1.2. I again exclude all firm-years with zero (automation/non-automation) patents and report the results on the intensive margin.

The results in Columns (3) and (4) show that codetermination leads to a 20.1 to 37.3 percent decline in the average inventor team size for automation patents.[53] The effects are sizable and statistically significant at a 1 percent significance level when estimated with year by industry fixed effects. The effect of codetermination on the average number of inventors for all patents and non-automation patents is smaller. Columns (2) and (6) show that firms locked in with codetermination tend to have on average 12 percent fewer inventors per patent or non-automation patent. The estimates are not statistically significant and the 95 percent confidence intervals do not exclude zero.

Table 1.5 shows that stock firms locked in with codetermination reduce the size of the inventor teams working on automation technologies. At the same time, there is no significant effect on the size of inventor teams who work on non-automation technologies.

Overall, to understand the mechanisms behind the effects of codetermination, I studied the heterogeneity with patent distance to science and estimated the effects on the number and age of patent backward citations as well as the size of the inventor teams. Several findings emerge. First, the negative effect of codetermination on automation innovation is driven by patents that are closer to science. Second, automation patents in codetermined firms tend to build on smaller and older prior art. Third, they are produced by on average smaller inventor teams. The main takeaway is that these results point to a decline in the value of automation patents and a change in the way automation innovation is organized in codetermined firms.[54] This is not so for non-automation, where the effects are smaller in magnitude and imprecise. The additional evidence

[52]The number of inventors per patent includes all inventors who contribute to a patent and not only those directly employed by filing firms.

[53]On average, automation patents filed between 1998 and 2014 by control stock firms have 2.2 inventors per patent.

[54]Krieger et al. (2022) show that patents closer related to science have on average larger inventor teams.

Table 1.5: Effect of shared governance on average size of the inventor teams

	Patents		Automation patents		Non-automation	
	(1)	(2)	(3)	(4)	(5)	(6)
Pre-reform × Stock	-0.131	-0.131	-0.224	-0.467***	-0.149	-0.130
	(0.118)	(0.124)	(0.219)	(0.180)	(0.128)	(0.134)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Year × industry FE	No	Yes	No	Yes	No	Yes
Observations	5,307	5,142	1,454	1,308	4,704	4,549
Log-likelihood	-8126.466	-7684.008	-2315.727	-1962.954	-7245.954	-6820.216

Notes: Columns (1) and (2) include estimates of the effect of codetermination on the average number of inventors in firm patent applications filed between 1998 and 2014. Columns (3) and (4) include estimates of the effect on the average number of inventors in automation patents. Columns (5) and (6) report estimates of the effect on the average number of inventors in non-automation patents. Firm-year observations with zero (automation/non-automation) patents are excluded. The estimation is based on Equation 1.2 and uses the sample of both stock firms and LLCs incorporated within a three-year window around August 10, 1994. All estimates are from a Poisson pseudo-likelihood regression. Standard errors are clustered at a firm level and are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

thus suggests that worker representatives in firm boards seem to divert the firm focus away from automation but do not lead to sizable changes in non-automation innovation.

1.7 Conclusion

In this chapter, I study the effect of the shared governance of firms on the direction of firm innovation. While there exists literature that shows how changes to corporate governance affect a variety of firm outcomes, evidence on the innovation effect of such changes remains scarce.

I focus on the policy that mandates worker representation on corporate boards. I exploit the 1994 reform that lifted the mandate for newly incorporated small stock firms but permanently locked in already incorporated stock firms with shared governance in Germany. Empirically, I show that the shared governance of firms has a strong negative effect on automation innovation and a smaller and imprecisely estimated negative effect on innovation in general and non-automation. For automation, the effect is driven by patents that build on science or are closely related to science. I provide evidence that automation patents at codetermined stock firms build on a smaller and older base of prior art, and their inventor teams become smaller. The corresponding effects for non-automation patents are much smaller and the confidence intervals do not exclude zero.

Thus, worker representatives on corporate boards seem to direct firm innovation away from automation technologies. The analysis suggests that policies that affect firm governance could have an effect on the direction of firm innovation.

Chapter 2

Whom to Inform About Prices? Evidence From the German Fuel Market

2.1 Introduction

Mandatory price disclosure (MPD) is becoming a popular policy tool to make markets more competitive.[1] Studies estimating the local treatment effect of mandatory price disclosure on prices find mixed results.[2] So far, there is limited evidence about why mandatory price disclosure sometimes lowers prices and sometimes does not. However, before introducing MPD, it is crucial to understand what its effect is going to be in the particular setting.

In this chapter, we ask what determines the price effect of mandatory price disclosure. More specifically, we focus on two key elements: How well consumers are informed prior to MPD, as well as the persistence of the price effects of MPD. Using a theoretical model with imperfect price information among consumers, we study how the share of uninformed consumers before mandatory price disclosure affects the price effect of MPD. We test the predictions in the context of the introduction of MPD in the German retail fuel market. There are two features of the setting that make it particularly suitable for this analysis: First, we observe high-frequency,

This chapter is based on joint work with Felix Montag and Christoph Winter (Montag, Sagimuldina, and Winter, 2023).

[1]MPD was introduced in numerous sectors, such as supermarkets, retail fuel, cement, or healthcare, and in many countries, such as Israel, Austria, Germany, Chile, Denmark, or the United States.

[2]See, for example, Luco (2019), who finds that mandatory price disclosure increased retail margins in the Chilean fuel market and Ater and Rigbi (2023), who find that mandatory price disclosure decreased prices at Israeli supermarkets.

station-level price changes for Germany and France before and after the introduction of MPD. Second, MPD was introduced simultaneously for diesel and gasoline. On average, consumers buying gasoline are less informed about prices than consumers buying diesel. Consumers can also not substitute between fuel types. Since the same fuel stations sell both types of fuel, there are no supply-side differences between fuel types. We use a difference-in-differences design to estimate the price effect of MPD for each fuel type. Fuel stations in Germany are part of the treatment group, whereas fuel stations in France are in the control group. Finally, we study whether follow-on local radio reports about fuel prices can intensify the treatment effect.

Several findings emerge: Theoretically, we show that the more uninformed consumers there are prior to the introduction of MPD, the larger is the reduction in prices that it induces. Empirically, we find that MPD decreases prices for all fuels but that this decrease is larger for gasoline, which has a less informed consumer base, than for diesel. In the German retail fuel market, MPD decreases gasoline prices by around 2.7 percent and diesel prices by around 1.8 percent. The difference in treatment effects is particularly strong in the five months after the introduction of MPD. Thereafter, the treatment effect stabilizes at between 1 and 2 percent for diesel and gasoline. Since the level of gasoline prices is higher than the level of diesel prices, the long-term effect of MPD in terms of Eurocents is higher for gasoline than for diesel. Finally, follow-on information treatments through local radio reports about prices can intensify the treatment effect. Overall, this suggests that MPD is most effective in markets where few consumers are well-informed before its introduction and that complementary information campaigns can increase the effect of MPD.

The theoretical analysis builds on Varian (1980). On the supply side, there are sellers that sell a homogeneous good and set prices. On the demand side, there are fully informed *shoppers* that know all prices, as well as uninformed *non-shoppers* that visit a seller at random. All consumers inelastically demand a single unit of the homogeneous good. In equilibrium, sellers set prices by randomizing according to a mixed strategy. Informed *shoppers* know all prices in the market, always buy from the lowest-price seller and pay the minimum price. Uninformed *non-shoppers* visit a single seller, observe its price and decide whether to purchase at that price or not purchase at all.

We model MPD as leading to an increase in the share of *shoppers*. Sellers always know all prices and are thus not directly affected by MPD.[3] We assume that price information coming from MPD always reaches a fixed number of consumers, irrespective of whether these are *shoppers* or *non-shoppers*. The ex ante share of *shoppers* thus affects how MPD changes prices in two ways: First, it affects the *marginal* effect of MPD on prices. Second, it affects how many *non-shoppers* become *shoppers* through MPD.

In the empirical application, we study the introduction of the Market Transparency Unit for Fuels (MTU) in Germany. Since September 2013, all fuel stations in Germany have to report all price changes in real-time to a central database. This aggregates the information and allows information service providers to defuse this information to consumers (e.g., via smartphone applications). This policy was recommended by the German Federal Cartel Office (2011) after diagnosing that a lack of price information on the consumer side was responsible for a lack of competition between fuel stations.

The station-level price reports to the MTU form the basis of our data set. To estimate the price effects of MPD we also need price data for fuel stations in Germany before the introduction of mandatory price disclosure. Here, we leverage that there already existed some smartphone applications prior to MPD that allowed users to self-report fuel prices, which were then collected and diffused to users in a similar fashion to the price information from the MTU.[4] We have access to the pre-MPD price notifications by users for one of these apps. This includes 20.5 million price notifications between the 1 September 2012 and the 31 August 2013. For the control group, we exploit the fact that there exists a similar database containing fuel prices at all fuel stations in France since 2007.

We use a synthetic difference-in-differences (SDID) design to estimate the price effects of mandatory price disclosure (see Arkhangelsky, Athey, Hirshberg, Imbens, and Wager, 2021). Similar to regular difference-in-differences, the treatment effect is estimated by isolating the change in prices after the introduction of MPD in the treatment group that is not present in the

[3]There is a rich theoretical literature on how improving price transparency on the producer side can stabilize collusion (see, for example, Green and Porter, 1984 or Kühn and Vives, 1995). It is likely the reason why MPD led to higher prices in the Danish concrete industry (Albæk et al., 1997) and the Chilean gasoline market (Luco, 2019). Our application is different in that producers already invested heavily in observing their competitors' prices before MPD (German Federal Cartel Office, 2011).

[4]The usage of these apps before MPD was considerably lower than after its introduction. This is why the introduction of MPD led to an important change in the the information set of consumers.

control group. Similar to synthetic control methods, the unit and time period weights in the control group are optimized as to best match pre-trends in the treatment group. Arkhangelsky et al. (2021) report that SDID performs weakly better than synthetic control and difference-in-differences methods.

By comparing the effect of MPD on gasoline and diesel prices, we can test the prediction about how the pre-MPD level of consumer information affects the price effect of MPD. A key feature of the setting is that the same fuel stations sell both types of fuel at the same pump. Besides the fuel type, the overall product (e.g., the shopping experience or the location) is identical. The key difference between gasoline and diesel is that these are bought by consumers that differ in their incentives to acquire information about prices and so in their ex ante information levels. In Germany, cars with diesel engines are driven by consumers that drive on average twice as many kilometers per year as gasoline buyers.[5] Buying a car with a diesel engine is a fixed cost investment to lower marginal costs.

Already prior to MPD the incentives to become informed about fuel prices and further reduce the price per liter was higher for diesel drivers. Using data on the user-reported price notifications before MPD, we show that the reporting intensity was higher for diesel than for gasoline. Using user-level search data after the MPD introduction, we show that the intensity of usage remained higher for diesel than for gasoline. Both of these pieces of evidence are consistent with our theoretical modeling of MPD.

To further strengthen the robustness of our main results, we rely on alternative identification strategies with which we can study the same theoretical mechanisms. First, we rely on an alternative information shock in which we study the local price effects of regular local radio stations that start reporting the lowest fuel prices in their reception area at some point after MPD. This also sheds light on the question of whether policymakers have any additional levers to ensure that the effect of MPD is persistent. Second, we use alternative identification strategies, where we isolate stations 20 to 100 kilometers from the Franco-German border or study differences in the treatment effect for local monopolists as compared to stations in competitive markets.

This chapter makes two main contributions. First, we derive empirically verifiable theoretical predictions on the role of ex ante consumer information for the effect of mandatory price dis-

[5] This is based on the figures from *Verkehr in Zahlen 2018* for the years 2013 and 2014.

closure policies. We build on the theoretical model of imperfect consumer information about prices by Varian (1980). We extend this framework by modeling how MPD affects consumers, accounting for how many consumers are informed *shoppers* ex ante. This yields an unequivocal prediction in which the magnitude of the price effect of MPD monotonically decreases in the ex ante share of *shoppers*. In contrast, there is no monotonic relationship between the ex ante share of *shoppers* and the price effect of a marginal increase in the share of *shoppers*. Thus, tailoring the modeling of the information shock to match how MPD works in practice allows to obtain an unambiguous theoretical prediction.

Second, we extend the existing empirical literature on price transparency policies by studying a novel mechanism of how MPD affects prices. In this context, our analysis highlights the importance of the share of consumers informed about prices before MPD. Importantly, we also show how the effect of MPD evolves over time and how complementary information campaigns can be used to strengthen the effect of MPD. Our findings relate to Albæk et al. (1997) and Luco (2019), who find that increasing price transparency in homogeneous goods markets led to an increase in prices. Since price transparency can also affect information on the supply side, this suggests that in those cases it seems to have stabilized collusion. In contrast, the German retail fuel market already had very high supply-side price transparency even before MPD. Ater and Rigbi (2023) find that MPD for Israeli supermarkets led to more intense competition, because low-price supermarket chains used MPD-enabled price comparisons to lend credibility to their price-based advertising campaigns. Brown (2019b) and Brown (2019a) study the demand- and supply-side responses of increasing price transparency in the U.S. health care market. Rossi and Chintagunta (2016) study how mandating fuel stations on Italian motorways to post the prices of rivals affects prices. There are important differences in the design of this policy as compared to the MTU.[6] However, their simulated price effect of the price disclosure policy leads to results that are of a similar magnitude to our findings. Martin (2020) studies how limiting the publicly distributed prices only to a subset of cheapest fuel stations affects equilibrium prices.

Finally, this chapter relates to an extensive empirical literature that analyzes pricing decisions for retail fuel.[7] There is an extensive empirical literature that studies the role of imperfect information in these markets (see, for example, Chandra and Tappata, 2011, Pennerstorfer, Schmidt-

[6]The policy only applies to the highly restrictive sample of motorway fuel stations. It also only allows drivers to discover rival prices once they reached a particular station, as opposed to seeing all prices online.

[7]Eckert (2013) provides an overview of the earlier literature on pricing in fuel markets.

Dengler, Schutz, Weiss, and Yontcheva, 2020, Byrne and de Roos, 2017 or Byrne and de Roos, 2022). In contrast, Houde (2012) emphasizes the role of spatial differentiation as opposed to imperfect information. Byrne and de Roos (2019) and Assad, Clark, Ershov, and Xu (2020) study how humans and algorithms learn to tacitly coordinate on softer competition and higher prices. Although understanding pricing decisions and the source of price dispersion in fuel markets is interesting in and of itself, Genakos and Pagliero (2022) and Montag, Sagimuldina, and Schnitzer (2021) show how this affects the pass-through of commodity taxes and thus has broader implications for the effectiveness of other policy tools.

The remainder of this chapter is structured as follows: Section 2.2 outlines the theoretical model. Section 2.3 describes the institutional setting and the data. Section 2.4 provides descriptive evidence on the price effects of MPD. Section 2.5 presents the empirical design and Section 2.6 includes the empirical results. Section 2.7 concludes.

2.2 Theoretical Model

We begin by theoretically shedding light on the effects of mandatory price disclosure policies in a context where consumers are imperfectly informed about prices. In our analysis MPD can be seen as synonymous with any exogenous information shock that makes prices at all sellers perfectly visible for some consumers. However it is different to changes in the visibility of prices at only some sellers or changes in price transparency endogenously chosen by sellers (e.g., through advertising).

Due to the structure of the market in the empirical application and the nature of the information shock, we place the analysis in the context of the Varian (1980) information model. Our focus lies on showing how the share of ex ante informed consumers affects the price effects of MPD.

2.2.1 Setup

The model features sellers and consumers. Sellers sell a homogeneous good and set prices. Consumers can be divided into two groups: *shoppers*, who know all prices and buy from the lowest-price seller, and *non-shoppers*, who draw a single seller at random, observe its price, and can only decide between buying and not buying at that price. Mandatory price disclosure leads to an exogenous increase in the share of *shoppers* in the population of consumers.

On the demand side, there is a unit mass of atomistic consumers that each inelastically demand a single unit of the good. The valuation of the good is the same across consumers and is denoted by υ. A fraction ϕ of consumers are *shoppers*. They know all prices and always buy from the lowest price seller. If there is a tie, shoppers are shared equally by the lowest price sellers.[8] A fraction $1 - \phi$ of consumers are *non-shoppers*.

On the supply side, there is a fixed and exogenous number of symmetric sellers. Each seller produces the homogeneous good at a marginal cost of production normalized to zero. We denote the number of firms by N, and sellers are indexed by i. Sellers form expectations about rival prices and choose a pricing strategy to maximize expected profits.

Finally, we need to model the impact of mandatory price disclosure. By enabling the creation of smartphone applications with which consumers can access all price information instantaneously at no cost beyond using the application, mandatory price disclosure converts some consumers from uninformed *non-shoppers* to fully informed *shoppers*. Furthermore, mandatory price disclosure is likely to lead to more than just a marginal increase in the share of informed consumers. How many consumers can be converted from being uninformed *non-shoppers* to being fully-informed *shoppers* depends on how many consumers are already fully informed prior to MPD. We therefore assume that MPD increases the share of fully informed *shoppers* by $\Delta_\phi(1 - \phi_0)$, where Δ_ϕ is the size of the information shock and ϕ_0 is the ex ante share of *shoppers*.

These two components are essential to model the effect of MPD. Δ_ϕ captures how large the information shock is (e.g., whether the existence of the measure is widely advertised). In contrast, $1 - \phi_0$ captures how many uninformed consumers there still are that could be informed by the policy. For example, if most consumers are already *shoppers* prior to the policy, even a heavily advertised MPD policy cannot lead to a large increase in the share of *shoppers*. Intuitively, the functional form of the information technology is such that MPD leads to information about prices being sent to a random subset of the population of consumers. Δ_ϕ determines how many consumers receive this message. $1 - \phi_0$ captures how many of these are turned into *shoppers* because they receive the message.

[8] In practice, there are no ties when there are no mass points in pricing strategies.

We search for the equilibrium pricing strategy by solving for the Nash Equilibrium of the game. Thereafter, we show how MPD affects equilibrium prices.

2.2.2 Equilibrium price distribution

There exists no equilibrium in pure strategies. Instead, there is a unique symmetric mixed strategy Nash equilibrium, which is characterized by the density function $F(p_i)$ and the closed and bounded support $[\underline{p}, p_r]$. p_r is the reservation price of non-shoppers and $\underline{p}$ is the minimum of the support from which a seller draws prices in the symmetric Nash equilibrium. In equilibrium, *shoppers* always buy from the lowest price seller and *non-shoppers* buy from the seller that they visit at random. Details on the derivation of these objects can be found in Appendix B.1.

Non-shoppers draw a single seller and observe its price. They purchase the good so long as the price is weakly below their valuation υ. Their reservation price p_r is thus equal to υ. Since no one purchases at a price above υ, no seller charges a price above υ in equilibrium and all *non-shoppers* buy the good at the randomly drawn seller.

The remaining equilibrium objects are derived using two equiprofit conditions that are based on the fact that in the symmetric mixed strategy Nash equilibrium, any price that a seller sets with positive probability should yield the same expected profit. A seller that sets the reservation price sells to its share of *non-shoppers*. A seller that sets the lowest price among all sellers sells to all *shoppers* and to its share of *non-shoppers*.[9] We solve for the minimum element of the support from which sellers draw prices in equilibrium, $\underline{p}$, by setting the expected profit under that price equal to the expected profit under the reservation price. We then derive the equilibrium density function by setting the expected profit under a price p_i equal to that under the reservation price.

The minimum element of the support from which sellers draw prices in equilibrium is

$$\underline{p} = \frac{\upsilon}{\frac{\phi N}{1-\phi} + 1}.$$

[9] There are no mass points in the equilibrium pricing strategies.

The cumulative density function from which sellers draw prices in equilibrium is

$$F(p_i) = 1 - (\frac{\upsilon - p_i}{p_i}\frac{1-\phi}{N\phi})^{\frac{1}{N-1}} .$$

In equilibrium, the expected profit of seller i is

$$E[\pi_i] = \upsilon\frac{1-\phi}{N} .$$

We can define two further objects, the expected price and the expected minimum price. Since *non-shoppers* always buy from the seller that they visit at random, the expected price reflects the average price paid by *non-shoppers*. In turn, since fully informed *shoppers* always buy from the lowest price seller, the expected minimum price corresponds to the average price paid by *shoppers*.

The expected price is

$$E[p] = \underline{p} + (\frac{1-\phi}{N\phi})^{\frac{1}{N-1}} \int_{\underline{p}}^{\upsilon} (\frac{\upsilon - p}{p})^{\frac{1}{N-1}} dp .$$

The expected minimum price is

$$E[p_{min}] = \frac{1-\phi}{\phi}(\upsilon - E[p]) .$$

2.2.3 Effect of mandatory price disclosure

Let us now turn to how mandatory price disclosure affects the equilibrium price distribution. We begin by highlighting how the share of fully informed *shoppers* affects the equilibrium price distribution. Since the reservation price of non-shoppers corresponds to their valuation of the good υ, this remains unaffected. We thus focus on how the minimum element of the support from which sellers draw prices, $\underline{p}$, and the equilibrium density function, $F(p_i)$, are affected when the share of shoppers ϕ increases.

Lemma 2.1. *With $0 < \phi < 1$, for any $\hat{\phi} > \phi$ the minimum element of the support of the equilibrium pricing strategy $\hat{\underline{p}} < \underline{p}$ and the Nash equilibrium pricing strategy with $\hat{\phi}$ first-order stochastically dominates (FOSD) the pricing strategy with ϕ, i.e. $\hat{F}(p) \geq F(p)\ \forall\ p$.*

This means that when $0 < \phi < 1$ and the share of *shoppers* ϕ increases, the minimum element of the support from which sellers draw prices decreases. Thus, the support of prices from which firms draw in equilibrium shifts to lower prices. At the same time, for each price on this support, the likelihood that a drawn price is lower than said price increases if ϕ increases.

When ϕ converges to zero, the Nash equilibrium converges to a degenerate distribution at the monopoly price. In this case, the monopoly price corresponds to the reservation price of non-shoppers, which is equal to the valuation of the good υ. When ϕ converges to one, so nearly all consumers in the market are fully informed about prices of all sellers, the Nash equilibrium converges to a degenerate distribution at the marginal cost (i.e., zero), which is the full-information Bertrand equilibrium.

Since an increase in the share of fully informed consumers in the market leads to a shift of the equilibrium density function towards lower prices, and to the downward shift of the minimum price a seller may choose in equilibrium, $E[p]$ and $E[p_{min}]$ also decrease. This means that when consumers become on average more informed, the average price paid by shoppers and the average price paid by non-shoppers decline and the expected price paid decreases for all consumers.

After establishing that more fully informed *shoppers* always lead to lower prices, we want to understand how the size of the effect of MPD varies with market conditions (i.e., the ex ante share of *shoppers*). That is, we want to understand how the effect of Δ_ϕ on equilibrium prices varies with ϕ_0.

Proposition 2.1. *With* $0 < \Delta_\phi < 1$ *and* $\phi = \phi_0 + \Delta_\phi(1 - \phi_0)$, *for any* $\hat{\phi}_0 > \phi_0$ *the change in the minimum element of the support of the equilibrium pricing strategy due to* Δ_ϕ *is* $\Delta\hat{\underline{p}} > \Delta\underline{p}$, *and the Nash equilibrium pricing strategy is such that* $\frac{\partial^2 F(p)}{\partial \Delta_\phi \partial \phi_0} < 0$.

The proof can be found in Appendix B.1. This means that the shift in the equilibrium price distribution towards lower prices due to the information shock Δ_ϕ is smaller in magnitude for markets with a higher ex ante share of *shoppers*. The effect of the information shock on the minimum element of the support of the equilibrium pricing strategy is also smaller when there are more *shoppers* before MPD. Figure 2.1 illustrates how the effect of MPD varies with the ex ante share of *shoppers* graphically.

Figure 2.1: Effect of the information shock on the equilibrium pricing strategy

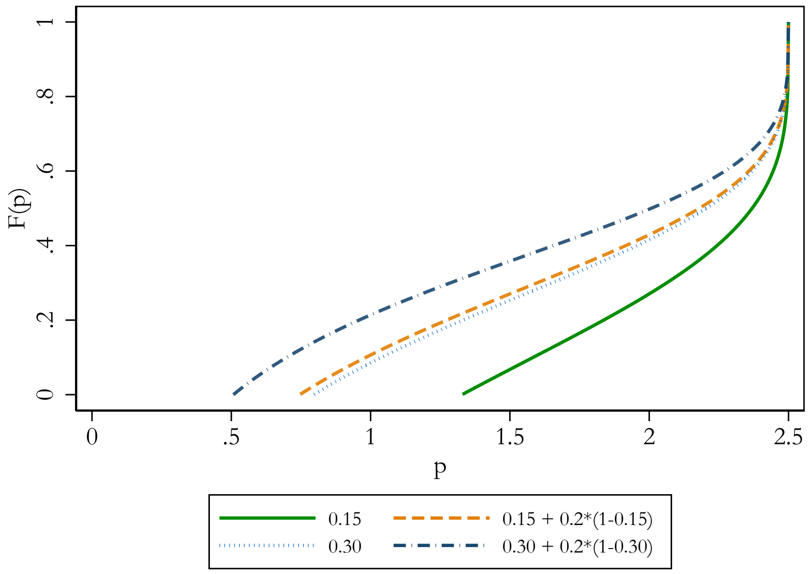

Note: The Figure shows simulation results of how the distribution from which sellers draw prices in the symmetric Nash equilibrium changes if the information shock Δ_ϕ hits the market. Parameter values: $\upsilon = 2.5$, $N = 5$, $\phi_{01} = 0.15$, $\phi_{02} = 0.30$ and $\Delta_\phi = 0.20$. The solid line and the short-dashed line capture the equilibrium price distribution when the ex ante share of *shoppers* is at 15% and 30%, respectively. The long-dashed line and the dot-dashed line show the corresponding density functions after the information shock of 0.2 hits the market. The information shock shifts the equilibrium price distribution towards lower prices, and the downward shift is larger in magnitude when the ex ante share of informed consumers is lower.

MPD shifts the entire distribution of prices more strongly towards lower prices if there are few *shoppers* ex ante. Therefore, the same holds true for the expected price, paid by *non-shoppers* in expectation, and the expected minimum price, paid by *shoppers* in expectation.

2.3 Institutional Setting

In the empirical application we study how mandatory price disclosure affects equilibrium prices in the German retail fuel market.

2.3.1 The German retail fuel market

Retail fuels are products with a very high degree of homogeneity within their product category. Although filling stations also sell other products, we focus our attention on the sale of fuel.

The two main fuel categories are diesel and gasoline. Consumers cannot substitute between the two in the short-term, as vehicles can only either run on one or the other type. In our analysis, we focus on gasoline with an octane rating of 95 and an ethanol share of 5 percent (also referred to as *E5*), as well as on diesel, which were correspondingly used in 56 and 29 percent of passenger vehicles with combustion engines in Germany in 2013.[10]

On the demand side, diesel and gasoline motorists differ in how much they drive. Diesel motorists tend to drive longer distances. According to the figures from *Verkehr in Zahlen 2018*, in 2013 to 2014 drivers of diesel passenger vehicles drove on average $20,500$ kilometers, whereas drivers of gasoline passenger vehicles on average drove only $11,000$ kilometers per year.

A potential explanation for why diesel motorists are more frequent drivers could be that buying a diesel vehicle is considered as a fixed cost investment to incur lower marginal costs afterwards. Diesel vehicles tend to be more expensive than gasoline vehicles, however, the per liter price for diesel fuel is consistently lower than that for gasoline. Motorists who expect to drive longer distances can therefore self-select into paying more upfront for a diesel vehicle in order to save

[10] This is based on 2013 statistics from *Verkehr in Zahlen 2018* and *Bundesverband der deutschen Bioethanolwirtschaft 2013*.

on fuel costs later on. Diesel motorists are thus likely to have a higher incentive to search for lower fuel price and be on average more informed about prices than gasoline motorists.

One could still argue that since diesel vehicles are oftentimes used for business purpose, diesel motorists may actually be less prone to search for lower prices. However, this is not a valid concern in our case. As of January 2013, out of 12.6 million diesel passenger vehicles in circulation in Germany, 4.6 million vehicles, including those with gasoline and diesel engine, were in use for commercial purpose. This means that at least 63 percent of diesel vehicles are owned and operated by private individuals (Kraftfahrt-Bundesamt, 2013). Among the remaining 37 percent of diesel vehicles used for business purpose, some drivers receive a lump-sum or a per mile fuel allowance or are self-employed, which creates additional incentives to save on fuel costs. Thus, many diesel vehicles being used for commercial purpose does not invalidate our observation that diesel motorists are on average more price sensitive than gasoline drivers.[11]

On the supply side, the retail fuel market in Germany is vertically organized. In the upstream market, crude oil is refined into retail products. These are sold and distributed to the downstream market, where filling stations sell the retail products to motorists. According to the German Federal Cartel Office (2011), concentration is high in both the upstream and downstream markets. Furthermore, some firms are vertically integrated, whereas others are not.

2.3.2 Mandatory price disclosure

Before the introduction of MPD, consumers were much less informed about prices than firms and hence found it difficult to assess the competitiveness of a particular fuel price. In the absence of an information clearinghouse, consumers faced significant search costs. To find the prices of all potential sellers, a motorist would need to drive to all stations.[12]

A market investigation ending in 2011 led the German Federal Cartel Office (GFCO) to find that prices in regional fuel markets had been higher than under functioning competition. After the market investigation, the GFCO and the German Monopolies Commission concluded that

[11] In Section 2.4, we provide further descriptive evidence which suggests that diesel drivers are on average more informed about fuel price than gasoline drivers both before and after MPD.

[12] There were already some apps that allowed users to self-report fuel prices, which were then collected and diffused to users in a similar fashion to the price information from MPD, but the usage of these apps before MPD was considerably lower than after its introduction.

a lack of price transparency on the consumer side caused the lack of competition and therefore called for an increase in price transparency in the downstream market. In 2012, parliament passed a law which set out the creation of the market transparency unit for petrol under the management of the GFCO and on 12 September 2013 the operation of the MTU began. The MTU is an information clearinghouse that collects prices in real-time and allows app creators to diffuse the information to users. It hence provides consumers access to live price data and increases price transparency.

2.3.3 Data

Our core data set contains station-level prices and retail margins for the universe of fuel stations in Germany and France for the years 2013 and 2014. We supplement this with consumer search data from a major fuel price app provider in Germany after mandatory price disclosure.

Prices, retail margins and fuel station characteristics

Our primary data set contains station-level prices and retail margins for *E5* gasoline and diesel on weekdays at 5 pm between 12 April 2013 and 31 August 2014 in Germany.[13] Throughout most of our analyses we use the station-level gross retail price, which includes taxes and duties, as an outcome variable. In order to estimate heterogeneities in the treatment effect, we add station characteristics such as information on the brand, address and geographic coordinates to our data set.

To illustrate how the MTU affects fuel stations, we carry out some analyses using retail margins as an outcome variable. We compute retail margins by subtracting the share of the price of crude oil that goes into the production of diesel or gasoline from the net retail price using the daily crude oil price at the port of Rotterdam.[14] Although these retail margins still contain different cost types, such as the cost of refining or transportation costs, the main source of input cost variation, the price of crude oil, is eliminated.

[13]We choose prices at 5 pm since this is the time around which most fuel is bought in Germany. More details on daily price cycles and purchase patterns are included in Appendix B.2.

[14]For a detailed description of the calculation of prices and margins, see Appendix B.2.

A novel and unique feature of our data is that we have rich station-level price information *before* the introduction of MPD. At that time, some smartphone apps existed that allowed their users to self-report station-level fuel prices. Although the usage of these apps was only a fraction of the usage of price comparison apps after MPD and the publicity that came with it, the pre-MTU apps contain rich price information. We use price data for the pre-MPD period supplied by one of the leading apps collecting self-reported prices. This data set comprises 17 million price reports for more than 13,500 stations between 1 January and 12 September 2013. Although the MTU went into operation on 12 September 2013, we only have access to its data from the 1 October 2013 onwards. Since our self-reported pre-MPD data only goes until the 12 September 2013, the period in between is not subject of our analysis.

For most days in the pre-MPD period, we have prices for more than 80% of fuel stations.[15] In case the reporting of prices is not random, selection could harm the validity of our estimation results. The most natural selection mechanism is that fuel stations themselves report prices onto the apps when they are low to attract *shoppers*. At the same time, they could refrain from posting prices when they are high in order not to discourage consumers from driving to their fuel station and discover the price. In this case, prices in our sample before MPD should be downward-biased. However, since we find that prices decreased after the introduction of MPD, this selection mechanism would work against us, and our estimates can be seen as a lower bound.

Another concern could be that the composition of fuel stations changed in our sample before and after the introduction of MPD. Table 2.1 presents summary statistics of our data. As can be seen in Panel A, the composition of fuel stations does not change significantly between the pre- and post-MPD periods concerning the share of integrated stations, the share of oligopoly stations or the number of competitors in local fuel markets. A detailed split of fuel stations by brand before and after the MPD introduction can be found in Table B.1 of Appendix B.2.1. Overall, the composition of brands is very similar.

[15]The daily number of fuel stations with price reports and the number of daily price changes are reported in Figures B.2 and B.3 in Appendix B.2. We exclude days after the MTU introduction from our analysis, where the number of price changes compared to the previous day drop by more than 40%. Since we observe the universe of price changes after the introduction of the MTU, and the average number of daily price changes is usually stable, we conclude that these days are affected by technical difficulties. In total, this affects ten days during the 15 months of data used from the MTU.

Table 2.1: Summary statistics

A. Station characteristics			
	D pre-MTU	D post-MTU	France
Number of Stations	13,782	14,606	9,224
Share of integrated stations	59%	57%	
Share of oligopoly stations	47%	46%	
Median # comp. (5 km catchments)	4	3	2
Share of local monopolists	15%	15%	19%
B. Prices and Margins			
	D pre-MTU at 5 p.m.	D post-MTU at 5 p.m.	France at 5 p.m.
Mean price, gasoline	1.60	1.50	1.54
Mean retail margin, gasoline	0.08	0.05	0.10
Mean daily spread, gasoline	0.09	0.07	0.14
Mean price, diesel	1.41	1.33	1.34
Mean retail margin, diesel	0.11	0.09	0.10
Mean daily spread, diesel	0.09	0.08	0.13

Notes: "D pre-MTU" and "D post-MTU" refer to fuel stations in Germany before and after the introduction of the MTU, respectively. The pre-MTU phase goes from 1 January 2013 until 12 September 2013. The post-MTU phase goes from 1 October 2013 until 31 December 2014. For France, all figures are for the full period 1 January 2013 until 31 December 2014. The average daily spread is measured as the average of the difference between the retail margin at the 95^{th} percentile and the 5^{th} on each day.

The largest share of the retail price for fuel in Germany consists of taxes and input costs. To analyze the share of the fuel price that can be influenced by fuel stations, we further analyze the effect on retail margins. First, we subtract taxes and levies to compute net fuel prices. Thereafter, we subtract the daily crude oil price at the port of Rotterdam to obtain retail margins.

Since January 2007, all fuel stations in France selling more than $500\,m^3$ of fuel per year have to report all price changes to a government agency similar to the MTU in Germany. Regular checks are carried out and fines imposed on fuel stations that do not comply with this rule. The French government makes all price information since 2007 publicly available on a government website.[16] We thus observe the universe of price changes of these fuel stations in France for our observation period. The data is regarded to be of very high quality and has previously been used by other researchers.[17]

[16]https://www.prix-carburants.gouv.fr/rubrique/opendata/, last accessed March 2021.

[17]Gautier and Saout (2015), for example, use this data to study the speed at which market prices of refined oil are transmitted to retail petrol prices.

The data set contains a list of notifications with the price, the type of fuel, the address and geographic coordinates of the fuel stations and the opening times. In contrast to the data of the MTU in Germany, it does not contain any information on the brand of the station or any other company-related information.

To compute retail margins, we also need a measure for input prices in France. Similarly to Germany, we use daily market prices for crude oil at the port of Rotterdam as a proxy for ex-refinery prices in France.

Local radio reports

After the introduction of mandatory price disclosure, some local radio stations started broadcasting local fuel prices over the air. Since some of the radio stations only started broadcasting prices at a time after the introduction of MPD, we exploit these introductions to study the effect of a follow-on information shock on prices. To facilitate the data collection, we restrict this analysis to the German state of Bavaria.

There are 381 radio stations in Germany broadcasting via short-wave out of which 83 are active in Bavaria. Among these, we identified 60 radio stations that could potentially broadcast fuel prices, which we contacted. Among these stations, we identified four local radio stations that broadcasted local fuel prices (e.g., the three lowest price fuel stations in their reception area) more than once a day at some point after the introduction of MPD in 2013 and 2014 and know the exact period of time of these broadcasts. We merge this information with data on the geographic availability of radio stations which we received from *fmlist.org*.

Search data, Google trends, and app usage

We complement our data set with information that paints a fuller picture of who is informed about prices, salience of the information, and its usage over time.

First, we use a data set that includes search queries in 2015 from a major smartphone app displaying fuel prices to users in Germany. For each search query there is a unique searcher device ID, as well as a time stamp and the fuel type that was searched for. We can therefore analyze how the extensive and intensive margins of search differ between the fuel types.

Second, we analyze information from Google trends on keywords surrounding the MTU. This tells us when public attention for the measure is particularly high and so when salience of the price information is high.

Third, we have data on the monthly usage of three major price comparison applications in Germany starting in May 2014.

2.4 Descriptive Evidence

Before moving to the econometric analysis, let us present some descriptive evidence to analyze the interplay between the level of ex ante price information, the usage of the price information, and the price effect of mandatory price disclosure.

2.4.1 Consumer information

According to the industry description in Section 2.3 and the theoretical assumptions on the effect of MPD, we would expect drivers fueling their cars with diesel to be more informed before and after the introduction of MPD.

Differences in price notifications by fuel type in the period before MPD provides suggestive evidence for differences in the information levels between fuel types. Intuitively, since fuel prices for price comparison apps before MPD were self-reported by users, motorists that report more prices are also likely to use this price information more. To proxy for how informed diesel and gasoline motorists were before MPD, we adjust the daily number of diesel and gasoline price reports to the number of diesel and gasoline vehicles in circulation in Germany.[18] Figure 2.2 shows the daily number of price notifications per 1,000 vehicles in circulation for each day in Germany between September 2012 and August 2013. The number of diesel price notifications per diesel car in circulation is about 64 percent higher than that of gasoline notifications. This strongly suggests that before MPD, diesel motorists were on average more informed about prices than gasoline drivers.

[18] From the count of price notifications, we drop all instances when *E5* gasoline, *E10* gasoline and diesel prices are reported during the same minute and for the same station, since this likely reflects self-reporting of prices by

Figure 2.2: Price notification patterns, pre-MPD (Germany)

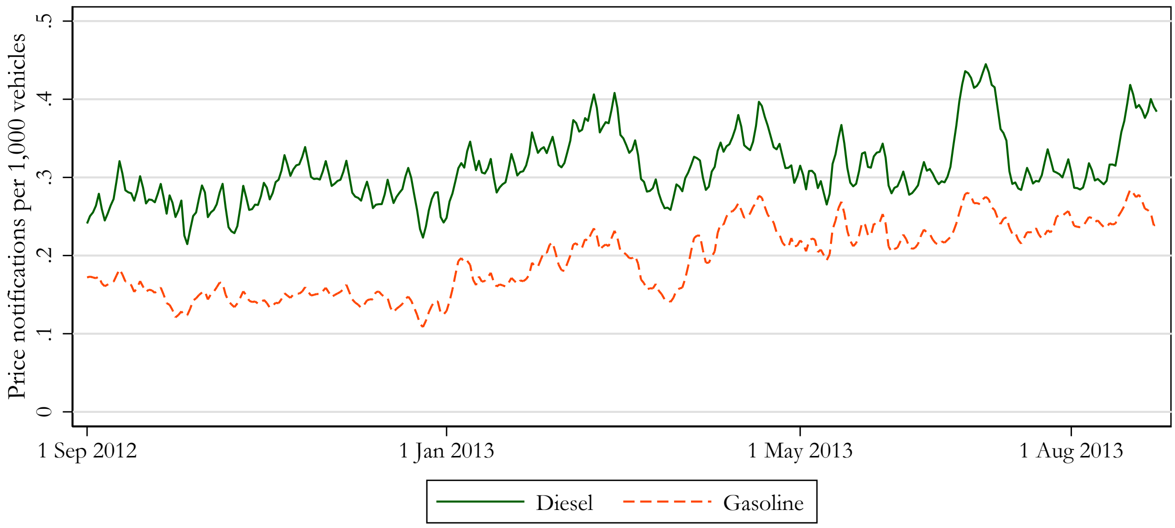

Notes: The Figure shows the daily number of self-reported price notifications by fuel type to a major German smartphone app per 1,000 diesel or gasoline vehicles in circulation. The data is available from September 2012 to August 2013. The solid line corresponds to the notification intensity for diesel. The dashed line corresponds to the notification intensity for gasoline.

After the introduction of MPD, self-reporting of prices became obsolete. Information on differences in app usage between users searching for prices for different fuel types can nevertheless provide evidence on relative differences in the information levels. To this end, we use data on search queries from a major fuel price app provider in Germany in 2015. Figure 2.3 shows the number of daily unique users searching for gasoline and diesel prices per 1,000 vehicles of the particular fuel type in circulation. The data is available for January to May 2015 and October to December 2015. The number of unique searchers (as opposed to the number of searches) captures the extensive margin of information usage and is thus similar to capturing differences in information through the share of *shoppers* in the theoretical model. Similarly to the pre-MPD pattern, the number of searchers is consistently higher for diesel than for gasoline prices.

Next, we investigate the intensive margin of price search, namely whether there are differences in the number of price searches per diesel or gasoline user. Figure 2.4 shows the average number of daily searches per unique user for diesel and gasoline. As becomes clear from the figure, there are no systematic differences in the number of searches between fuel types.

stations and not by motorists. 16 percent of all price notifications are individual reports for either gasoline or diesel price.

Figure 2.3: Unique daily price searchers by fuel type, post-MPD (Germany)

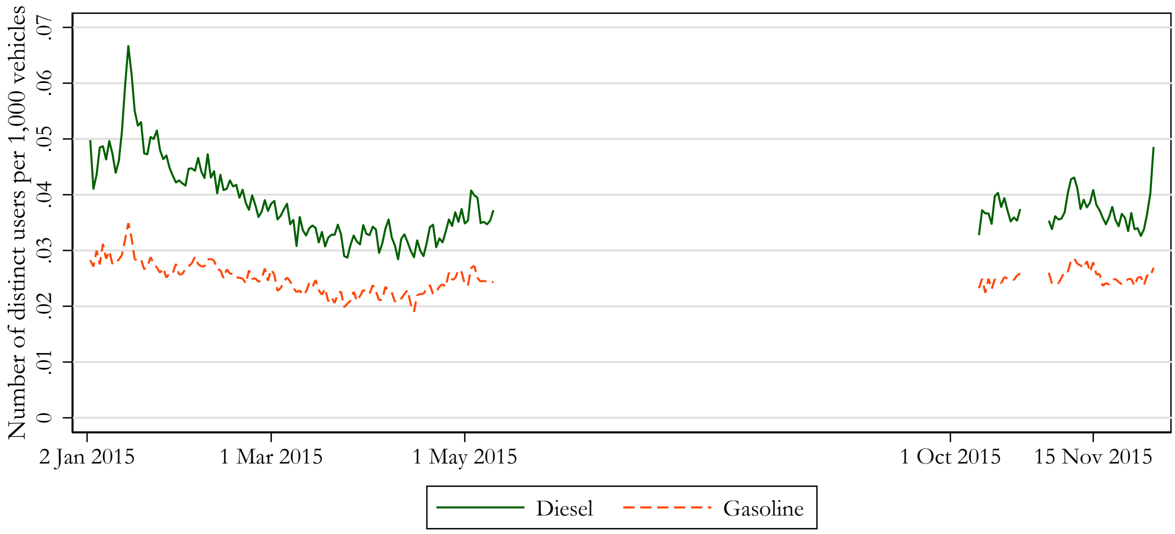

Notes: The Figure shows the daily number of distinct users who search for diesel or gasoline price in Germany in 2015, per 1,000 diesel or gasoline vehicles in circulation.

Figure 2.4: Average daily search number per user by fuel type, post-MPD (Germany)

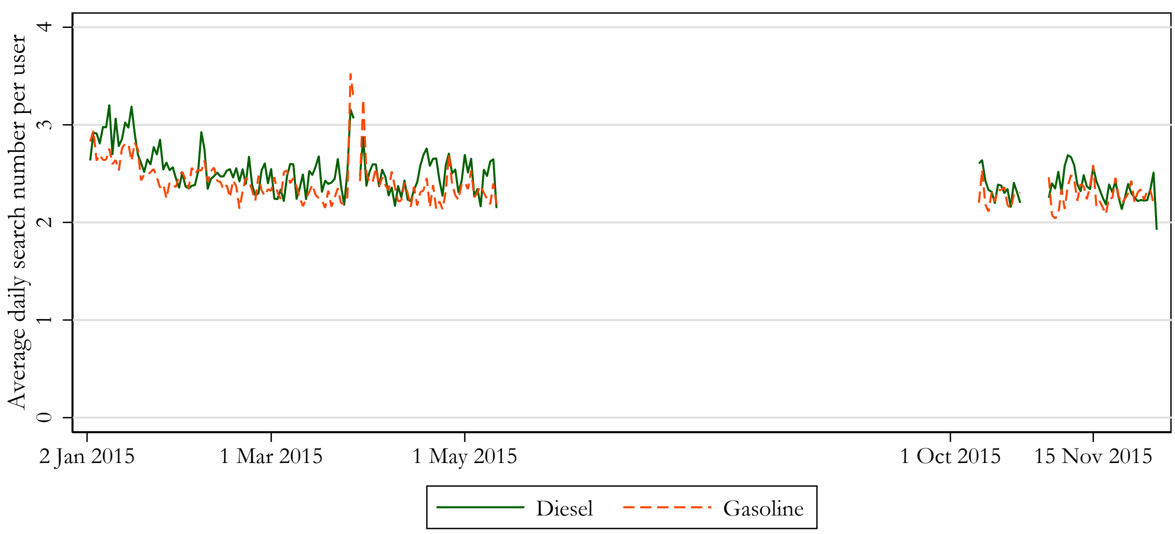

Notes: The Figure shows the daily number of price searches by fuel type at a major German smartphone app per 1,000 diesel or gasoline vehicles in circulation. The data is available for January to May and October to December 2015. The solid line corresponds to the search intensity for diesel. The dashed line corresponds to the search intensity for gasoline.

Before and after the introduction of MPD there is strong evidence suggesting that diesel drivers are systematically more informed about prices than gasoline drivers. This is driven by the extensive margin (i.e., a higher share of informed diesel drivers) as opposed to the intensive margin (i.e., informed diesel drivers knowing more than informed gasoline drivers). Thus, more diesel than gasoline drivers decide to become informed but conditional on becoming informed, the search behavior appears to be similar.

To understand the usage of the price data made available to consumers by MPD over time, we analyze two pieces of evidence. The first is shown in Figure 2.5, which plots the search indicator for different keywords surrounding the MTU, fuel prices and price comparison apps on Google in Germany between January 2013 and December 2014. These are indexed such that 100 corresponds to the week-keyword combination that has the most search queries. Searches for all keywords peak in mid-September, when operations of the MTU began. Whereas searches for the MTU itself declined again quickly, searches for "Tankstellen App" (fuel station app), "Benzinpreis App" (fuel price app), or "Benzinpreisvergleich" (fuel price comparison) remain high until mid-January 2014.

The second piece of evidence is included in Figure 2.6, which shows the evolution of monthly page impressions for three mobile price comparison applications for which data is available starting in April 2014. Although these three mobile applications are only a fraction of the German mobile fuel price comparison market, they together have more than 70 million page impressions in December 2014. This shows that mobile price comparison applications were widely used. Usage per app also appears to be relatively stable between April 2014 and October 2014 for *Clever Tanken* and *T-Mobile Tanken*.

Figure 2.5: Evolution of Google searches for MPD-related search terms in Germany

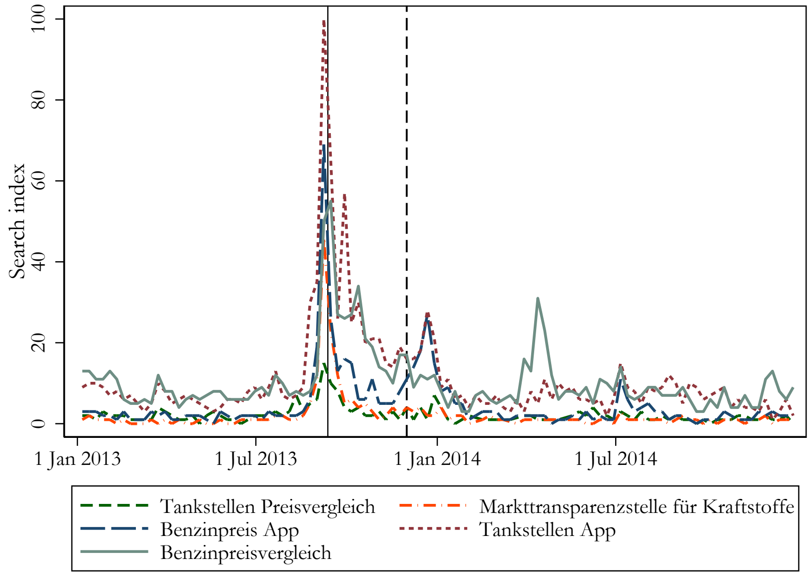

Notes: The figure shows the evolution of Google searches in Germany between 1 January 2013 and 31 December 2014 for MPD-related keywords. Searches are indexed such that 100 corresponds to the moment in time and keyword with the highest number of searches during the observation period. The search terms are "Tankstellen Preisvergleich" (fuel station price comparison), "Marttransparenzstelle für Kraftstoff" (market transparency unit for fuel), 'Benzinpreis App" (fuel price app), 'Tankstellen App" (fuel station app), and "Benzinpreisvergleich" (fuel price comparison). The vertical solid line marks the beginning of the MTU test phase. The vertical dashed line marks the beginning of the MTU full-scale phase.

Figure 2.6: Monthly page impressions

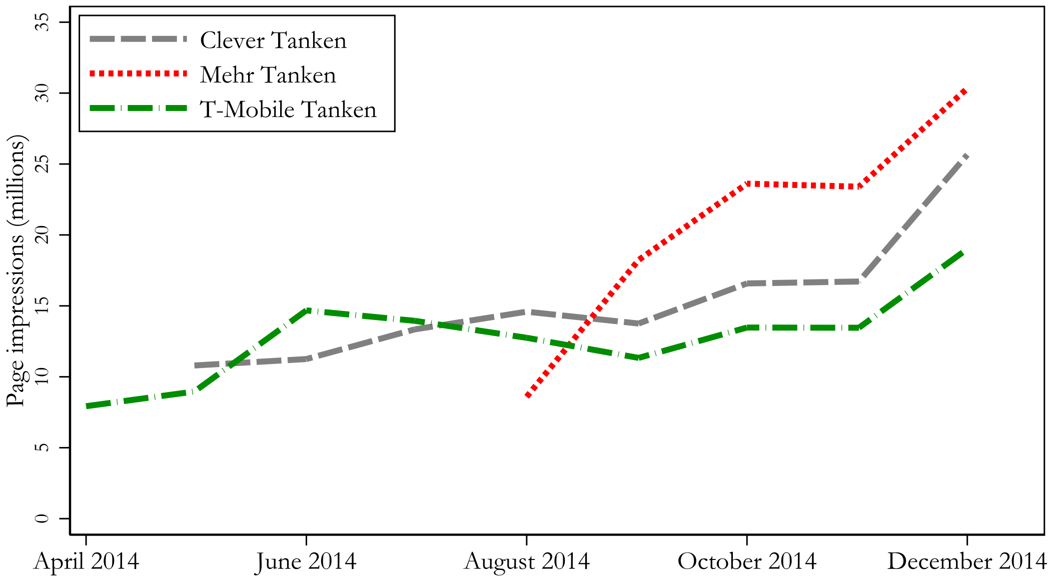

Notes: The Figure shows the evolution of monthly page impressions for three popular mobile price comparison applications. Each line begins when data for the particular app becomes available and ends at the end of our sample period, in December 2014.

2.4.2 Price effect of mandatory price disclosure

To study the effect of mandatory price disclosure on diesel and gasoline prices we begin by comparing how the difference between prices in Germany and France evolves over time for diesel and gasoline, respectively. Figure 2.7 shows the evolution of gross prices in Germany relative to France between April 2013 and September 2014 for diesel and gasoline. The solid line plots the difference in daily diesel prices between Germany and France, demeaned by the average difference prior to MPD. The dashed line plots the same for gasoline.

Figure 2.7: Evolution of the difference in gross prices between Germany and France

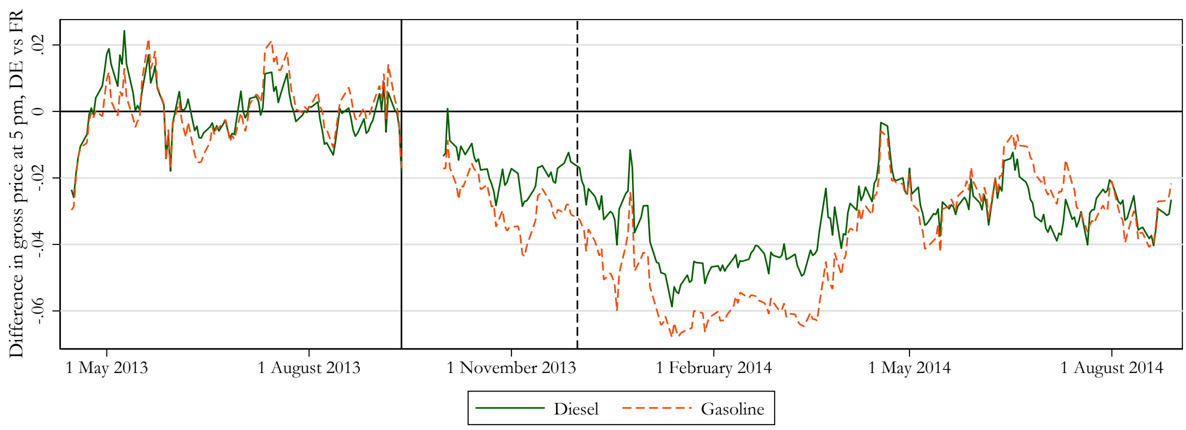

Notes: The solid line shows the evolution of the difference in daily diesel prices between Germany and France, demeaned by the corresponding average difference prior to MPD. The dashed line shows the evolution of the analogous difference in gasoline prices. The vertical solid line marks the beginning of the MTU test phase. The vertical dashed line marks the beginning of the MTU full-scale phase.

Before MPD, the difference in gross prices between Germany and France oscillates around zero for both types of fuel. After MPD, it appears as though prices fall more strongly for gasoline than for diesel. The effect of MPD appears to be strongest in January 2014, stagnate thereafter and then become weaker but still existant after May 2014.

Relating this to the descriptive evidence on consumer information, it appears as though the price effect of MPD is stronger for gasoline, where we expect a lower share of ex ante informed consumers. This is in line with the theoretical prediction in Proposition 2.1. The strength of the treatment effect of mandatory price disclosure also appears to coincide with the public attention devoted to fuel price comparison apps shown in Figure 2.5. This suggests that public attention to this information and active usage are key to fully exploit the potential of MPD.

2.5 Empirical Strategy

After providing descriptive evidence on the effect of MPD, we test whether the descriptive results withstand more rigorous econometric analysis. In our main specification we use station-level fuel prices in Germany and France and a synthetic difference-in-difference strategy to estimate the price effects of MPD for diesel and gasoline. We test the robustness of the results and how these relate to the theoretical model by estimating the price effect of follow-on radio reports that enhance the diffusion of price information.

2.5.1 The effect of mandatory price disclosure

To estimate the average effect of mandatory price disclosure on fuel prices, we use a synthetic difference-in-differences (SDID) framework in which we compare log fuel prices at stations in Germany to those in France, before and after MPD.

The synthetic difference-in-differences is a method recently proposed by Arkhangelsky et al. (2021). It combines the advantages of difference-in-differences with those of synthetic control methods. Similarly to difference-in-differences, SDID estimates the treatment effect by comparing the difference in outcomes of a treatment and a control group before and after the treatment, and relies on the parallel trends assumption. Similarly to the synthetic control method, SDID re-weighs units in the control group to make pre-trends in outcomes as similar as possible to those of the treatment group. Arkhangelsky et al. (2021) report that SDID performs weakly better than synthetic control and difference-in-differences methods.

The estimation proceeds in two steps. In the first step, we compute weights for the control units and for the pre-treatment time periods. SDID unit weights are designed to minimize the difference in pre-trends of outcomes between exposed and unexposed units prior to the treatment. SDID time weights are set to balance time periods before and after the treatment for the control units and emphasize pre-treatment time periods most predictive of the post-treatment ones. In the second step, we estimate the treatment effect with the use of the unit and time weights from the first step.[19] Standard errors are computed via the jackknife method.[20]

Specifically, we solve the following minimization problem:

$$(\hat{\beta}^{sdid}, \hat{\mu}, \hat{\alpha}, \hat{\gamma}) = \underset{\beta,\mu,\alpha,\gamma}{\arg\min} \left\{ \sum_{i=1}^{N} \sum_{t=1}^{T} (Y_{it} - \mu - \alpha_i - \gamma_t - MPD_{it}\beta)^2 \, \hat{w}_i^{sdid} \hat{\tau}_t^{sdid} \right\} \tag{2.1}$$

where $\hat{\beta}$ corresponds to the estimated effect of the MTU introduction, and $\hat{w}_i$ and $\hat{\tau}_t$ are SDID unit and time weights. Y_{it} is the logarithm of the fuel price at station i and week t. α_i and γ_t are fuel station and week fixed effects. The variable MPD_{it} is a dummy that equals one for treated units after the treatment. These are fuel stations in Germany after the introduction of the MTU.[21]

Estimation of the treatment effect with SDID requires a balanced panel. We compute weekly average fuel prices and restrict our sample to fuel stations in Germany and France that have no missing weekly price observations.[22] This is the case for 47% of stations in Germany and 94% of stations in France. Since we estimate the effect of MPD using this restricted sample, in Appendix B.3 we report the results estimated using regular difference-in-differences when we use the full, unbalanced panel and daily price observations. The results hold.

To study the effect of MPD over time, we estimate the parameters of the following regression model:

$$ln(p_{it}) = \sum_{j=-5}^{11} \beta_j MPD_{it} + \alpha_i + \gamma_t + \epsilon_{it}, \tag{2.2}$$

where $ln(p_{it})$ is the logarithm of the weekly average fuel price at station i. β captures the effect of the mandatory price disclosure starting five months before its introduction and up to eleven

[19] In Appendix B.2, we show the geographic distribution of control stations that receive a disproportionately higher unit weight in estimation via SDID. These stations are scattered throughout France and do not appear to cluster in a particular region. Therefore potential clustering of control stations due to re-weighting by SDID does not affect our results.

[20] The jackknife method produces a conservative estimate of the variance in large panels with a high number of treated units. We use the jackknife method instead of bootstrapping as the latter is too computationally intensive in this case.

[21] We solve the minimization problem using the *synthdid* package in R developed by Arkhangelsky et al. (2021).

[22] We employ weekly average fuel prices since a high share of stations in Germany have at least one day without a reported fuel price during the time period used in the estimation of the treatment effect.

months after. The regression is weighted by the SDID unit and time weights, and we control for fuel station and week fixed effects.

2.5.2 France as a control group

We identify the effect of MPD using the evolution of fuel prices at fuel stations in France as a comparison. Two assumptions need to be met to identify the effect of MPD in our framework: The first is that there cannot be any other transitory shocks affecting fuel stations in France and Germany differently before and after the introduction of MPD other than MPD itself. The second is that there are no spillovers from the treatment onto the control group. Subsequently, we provide evidence that suggests that both assumptions hold.

The station fixed effects capture time-invariant differences between fuel stations in France and Germany. The week fixed effects capture transitory shocks that affect French and German fuel stations equally. Due to its similarities in size, wealth and geographic location, as well as our narrow observation period, there should not be any additional transitory demand and supply shocks that affect France and Germany differently. We nevertheless discuss the most obvious candidates.

Important transitory demand shocks in the retail fuel market are school and public holidays, as well as local economic shocks. School and public holidays in France and Germany are highly correlated. In addition, since holidaymakers in Europe often cross several countries on the way to their holiday destination and France and Germany are popular holiday destinations and important transit countries, they are usually hit similarly and at the same time by these demand shocks.

Transitory supply shocks affect fuel stations much in the same way. Due to their geographic proximity, fuel stations in France and Germany procure most of their fuel from similar sources. Furthermore, the European Single Market and the Schengen Agreement mean customs, border controls or other regulatory hurdles do not restrict arbitrage possibilities between the two countries. To nevertheless ensure the elimination of any transitory shocks to input prices and to restrict our analysis to the share of the fuel price that can be affected by fuel stations, we additionally use retail margins as outcome variables. These retail margins are net of taxes, levies and the wholesale price of Brent oil in Rotterdam on a given day.

Also, fuel stations in France constitute a good control group because there were no important regulatory changes in the French fuel market over our observation period. The impact of the introduction of mandatory price disclosure in 2007 should have stabilized by 2013 and thus not affect different French fuel stations differently over our observation period. In contrast to other countries, France, like Germany, did not restrict its fuel stations in their price-setting behavior other than by imposing mandatory price disclosure.[23]

One might be worried that there may still be idiosyncratic developments, which add random noise to the data and thus lead to an underestimation of the absolute value of the effects. We therefore, re-run our analysis for a sub-sample of the data around the Franco-German border, for which the economic conditions should be similar due to geographic proximity. First, we restrict our analysis to fuel stations that are 100 kilometers left and right to the border. Fuel stations in the treatment and control groups are thus in the same economic area and only exposed to common transitory shocks. Second, to eliminate any potential spillover effects, we drop all fuel stations that are less than 20 kilometers left and right of the border. We are left with a Donut-SDID, where stations on both sides of the border are geographically close, but stations that are potentially subject to spillover effects are dropped.

Finally, a potential concern could be that the drop in the price of crude oil in the second half of 2014 could bias our results. For the analysis of fuel prices and retail margins where we control for station and week fixed effects, this would require the pass-through of input prices to change differently for the treatment and the control group over time. This is unlikely to be a concern because most of our analysis only uses data until 31 August 2014, whereas the largest share of the decrease in the price of crude oil occurred between October and December 2014. We also directly account for potentially differential pass-through of oil cost shocks by including an interaction of the country indicator with the crude oil price in our estimation.

Furthermore, our data set allows us to robustly estimate the treatment effect using different treatment groups and different identification strategies. Two analyses are of particular interest, as the approaches are very different to the strategies used to obtain the main results: In the first,

[23]In 2011, Austria, for example, introduced a rule banning fuel stations from raising prices more than once a day.

we treat local monopolists in Germany as the control group and all other German stations as the treatment group.[24] In the second, we use country-level weekly fuel prices for all countries in the European Union and treat Germany as the treatment group and all other countries as the control. The results are reported in Appendix B.3 and are in line with our main findings.

2.5.3 Radio reports

As discussed in Section 2.3, some local radio stations started broadcasting local fuel prices over the air after the introduction of MPD. This allows us to test the robustness of our main result. If MPD increases the share of fully informed *shoppers*, thereby decreasing prices, then local radio reports should further increase the share of *shoppers*, thereby leading to a further local decrease in prices.

To limit the burden on data collection, we restrict the analysis of radio reports to the German state of Bavaria.[25] As described in Section 2.3, we identify four stations that have segments that recur at least daily and in which they broadcast the prices at the cheapest fuel stations in the reception area. We discard two of the radio stations because they already broadcasted the lowest fuel prices amongst those called in by their listeners before MPD started. We exclude all fuel stations in their reception areas from the analysis, as they are treated throughout the observation period. The two remaining radio stations are *Radio Arabella*, which started its broadcast on 25 April 2014 and *Extra-Radio*, which started its broadcasts on 2 February 2014.

Figure 2.8 shows the reception areas of *Radio Arabella* and *Extra-Radio*. For each fuel station we know whether, on a particular day, it is within the reception area of a radio station broadcasting prices or not.

Using a difference-in-differences design, we estimate the following fixed effects regression model:

$$ln(p_{it}) = \beta_0 + \beta_1 Radio_{it} + \alpha_i + \gamma_t + \epsilon_{it} \tag{2.3}$$

[24]The empirical literature analysing price dispersion in retail fuel markets considers different geographic market definitions. For example, Chandra and Tappata (2011) consider a 1 mile as well as a 2 miles radius, while Barron, Taylor, and Umbeck (2004), Hosken, McMillan, and Taylor (2008) and Lewis (2008) consider a radius of 1.5 miles. We use different catchment sizes in further results in Appendix B.3.5.

[25]Fuel stations in the treatment and control groups are therefore also all in Bavaria.

Figure 2.8: Radio reception areas and fuel stations in Bavaria

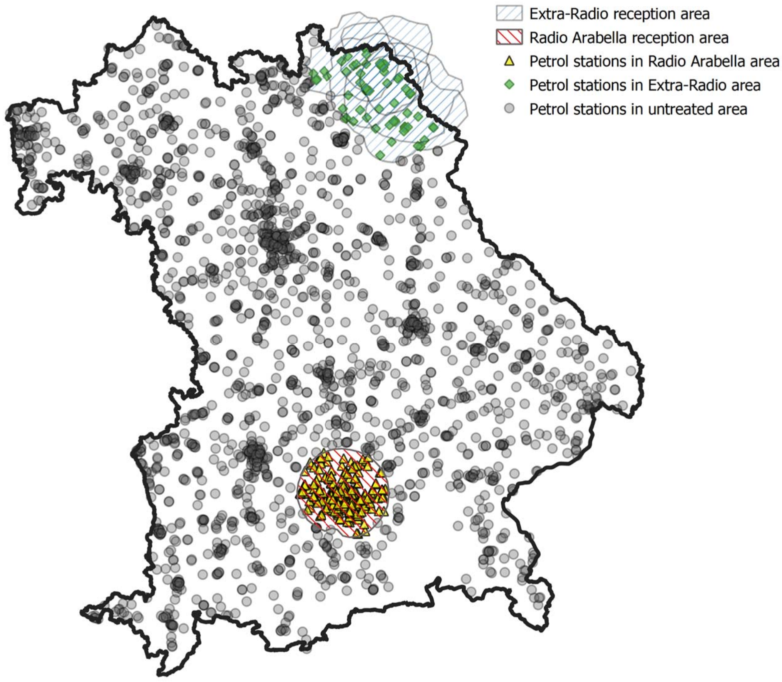

where $ln(p_{it})$ corresponds to the logarithm of the gross price for diesel or gasoline at station i at time t and $Radio_{it}$ is a dummy equal to one if fuel station i lies in the reception area of a radio station broadcasting local fuel prices at date t. α_i are fuel station fixed effects, and γ_t are date fixed effects.

We can thus exclude that fuel stations in the control group are affected by reports of radio stations we have not surveyed. We restrict our analysis to the period October 2013 until September 2014, which is the twelve months after the beginning of the test phase of the MTU.

To estimate the effect of radio reports on fuel prices we need to ensure that there are no spillovers of radio reports onto fuel stations in the control group and that the decision of radio stations to report was not because they anticipated evolutions in their local market that would also affect fuel prices.

There are two possibilities which could lead to spillover effects between the treatment and control groups: First, motorists outside of the reception area of the radio station could listen to the radio station via the internet. Second, commuters driving through the reception area of the radio station could update their information set by listening to the broadcasts and change their behavior accordingly after leaving the reception area. Both of these threats to identification are unlikely to be strong. Radio stations were still predominantly listened to via short-wave in 2013 and 2014. In particular, in more rural areas, mobile internet reception was still weak, making it difficult to listen to radio via the internet when on the road. Furthermore, although commuters learn something about the distribution of prices by listening to the radio, which may still be valuable outside the reception area, the value of this information is likely decreasing with distance to the reception area. In any event, both concerns lead to the control group being partially treated and would thus lead us to underestimate the treatment effect.

Another potential threat to identification could be that radio stations anticipated a trend that would create local demand for reports about fuel prices and that also affected fuel prices. This seems unlikely. After multiple interviews with program directors we learned that the decision of broadcasting fuel prices is not based on a market analysis but rather based on the fit of such a segment to the existing program.

We now turn to the radio stations that define our treatment group. We consider radio reports about fuel prices by *Extra-Radio*, which broadcasts in and around Hof, a city in North-Eastern Bavaria, close to the Czech border, and *Radio Arabella*, which is a radio station broadcasting in and around Munich. Whereas *Extra-Radio* broadcasted the lowest fuel prices in its reception area daily between 2 February 2014 and 5 March 2017, *Radio Arabella* started reporting the lowest prices several times a day on 25 April 2014 and reports are still ongoing at the time of writing.

The presence of a country border is important. In particular, the reception area of *Extra-Radio* is very close to the border with the Czech Republic, the focal city Hof being less than 10 kilometers away from the border. Since Germany and the Czech Republic are both members of the Schengen Area, there are no border controls and shopping in the neighboring country is frequent. Due to lower taxes and levies, fuel prices are consistently 20 Eurocent lower in the Czech Republic. It therefore seems plausible that independent of price reports by radio stations or smartphone apps, price-sensitive consumers always buy fuel in the Czech Republic, whereas only inelastic consumers buy from fuel stations treated by *Extra-Radio*. We would therefore expect that reports by *Extra-Radio* have little to no effect on fuel prices. To test this hypothesis, we estimate the regression model for both radio stations separately. In each of these regressions we exclude fuel stations within the reception area of the other radio station from the control group.

2.6 Results

2.6.1 Effect of mandatory price disclosure by fuel type

Table 2.2 includes the main estimation results. Columns (1) and (2) include the effect of MPD on the logarithm of fuel prices for gasoline and diesel, respectively, using the full sample of French and German fuel stations. Columns (3) and (4) include results where the sample is restricted to fuel stations 20 to 100 kilometers away from the Franco-German border.[26]

The main takeaway from these results is that MPD is successful at decreasing prices and that its effectiveness is higher for gasoline than for diesel. In line with the theoretical predictions and the descriptive evidence the effect of MPD is larger when the share of ex ante informed consumers is lower. Since the same fuel stations offer diesel and gasoline, supply side characteristics cannot explain these differences in the effect of the MTU across the two fuel types.

Table 2.2: Effect of MPD on the logarithm of gross prices

	Gasoline	Diesel	Gasoline	Diesel
	(1)	(2)	(3)	(4)
MPD	-0.027***	-0.018***	-0.029***	-0.021***
	(0.0005)	(0.0004)	(0.001)	(0.001)
95% Confidence interval	[-0.028, -0.026]	[-0.019, -0.018]	[-0.032, -0.027]	[-0.023, -0.019]
Week FE	*Yes*	*Yes*	*Yes*	*Yes*
Station FE	*Yes*	*Yes*	*Yes*	*Yes*
Observations	632,884	751,219	49,539	55,517

Notes: Columns (1) and (2) include estimates of the effect of MPD on log weekly prices for gasoline and diesel, respectively, using all fuel stations in Germany and France. Columns (3) to (4) include the same estimates for a restricted sample of fuel stations 20 to 100 kilometers away from the Franco-German border. The observation periods goes from 15 April 2013 to 31 March 2014. Standard errors are computed using the jackknife method and are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[26] The results are robust to changes to the distance thresholds. We provide estimates for alternative thresholds in Appendix B.3.2.

Figure 2.9 shows the time-varying effects of mandatory price disclosure on the logarithm of weekly average gross prices for gasoline and diesel. After the start of MPD prices decline for both fuel types, however more strongly for gasoline than for diesel. The largest effect of MPD is in January 2014. This also coincides with the end of widespread public attention for the MTU and price comparison apps, as seen in Figure 2.5. Following this period of high attention, the effect of MPD becomes smaller in magnitude again but remains stable. This is in line with evidence that there is a stable and continuous use of price comparison apps after April 2014. The MPD induced price effect stabilizes at approximately the same percentage point for diesel and gasoline. As the price level of gasoline is higher than for diesel, the long-term price effect in Eurocents is stronger for gasoline than for diesel.

Figure 2.9: Time-varying effect of MPD on the logarithm of gross prices

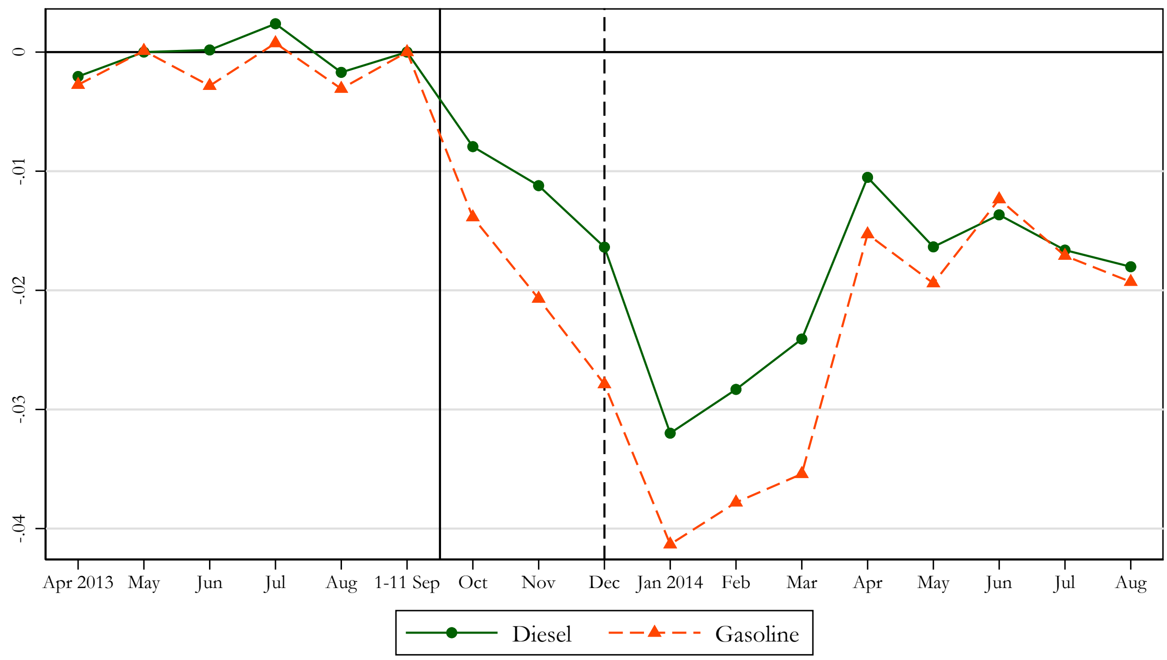

Notes: The Figure shows time-varying treatment effects of MPD on log weekly prices for gasoline and diesel between April 2013 and August 2014. The vertical solid line marks the beginning of the MTU test phase. The vertical dashed line marks the beginning of the MTU full-scale phase.

In Appendix B.3, we demonstrate the robustness of our results. First, we use the full, unbalanced sample of fuel stations and a regular difference-in-differences estimator. Second, we estimate the Donut-SDID using alternative distance thresholds. Third, we control for an interaction of the crude oil price and a country dummy, to allow for differential pass-through of crude oil shocks in each country. Fourth, we estimate the effect of MPD on retail margins. Fifth, we focus only on stations in Germany and use local monopolists, whose competitive environment did not change as a result of MPD, as a control group. Sixth, we use country-level weekly average prices for all 27 countries in the European Union from the Weekly Oil Bulletin, using Germany as the treatment group and all other countries as a control group to estimate the effect of MPD for diesel and gasoline. Our results hold in all of these alternative specifications.

2.6.2 Radio reports

In Table 2.3 we report the results from regressing the logarithm of prices on the existence of local radio reports about fuel prices. Columns (1) and (2) include the results of the effect of reports by *Extra-Radio* and *Radio Arabella* on gasoline prices. Columns (3) and (4) include the results for diesel.

We find that whereas reports by *Radio Arabella* lead to lower fuel prices, this is not the case for reports by *Extra-Radio*. This is consistent with our expectation, since the reception area of *Extra-Radio* lies on the border to the Czech Republic, where fuel is significantly cheaper, and so radio reports do not add any relevant information for price sensitive consumers. Overall, we find that where follow-on radio reports add further information for consumers, they lead to a further decrease in prices.

Table 2.3: Effect of radio reports on the logarithm of gross prices

	Gasoline		Diesel	
	(1)	(2)	(3)	(4)
Treatment group:	Extra-Radio	Arabella	Extra-Radio	Arabella
Radio reports	0.003	-0.002***	0.002	-0.005***
	(0.003)	(0.0004)	(0.002)	(0.0004)
Date FE	*Yes*	*Yes*	*Yes*	*Yes*
Station FE	*Yes*	*Yes*	*Yes*	*Yes*
Observations	350,655	452,481	355,928	457,559
Adjusted R^2	0.694	0.705	0.625	0.643

Notes: There are 70 fuel stations in the reception area of *Extra-Radio* and 585 fuel stations in the reception area of *Radio Arabella*. Columns (1) and (3) compare log prices for gasoline and diesel, respectively, at fuel stations in the reception areas of *Extra-Radio* to other fuel stations in Bavaria before and after the beginning of radio reports. Columns (2) and (4) do the same for radio reports by *Radio Arabella*. Standard errors, clustered at the fuel station level, are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

2.7 Conclusion

In this chapter, we study the determinants of the price effect of mandatory price disclosure. Theoretically, we derive novel predictions about how MPD affects prices in the context of the Varian (1980) model. We show that the magnitude of the price effect of MPD monotonically decreases in the share of consumers that are well informed about prices ex ante.

Empirically, we study the price effect of mandatory price disclosure in the German retail fuel market. Overall, we find that MPD led to lower prices. There are two important mechanisms that we uncover in our empirical analysis: First, we confirm the theoretical prediction that the effect of MPD is stronger for markets where there are fewer ex ante well informed consumers (i.e., gasoline). Second, we find that the magnitude of the price effect of MPD declines over time, before staying constant at between 1 and 2 percent for diesel and gasoline. Since the gasoline price level is higher than of diesel, this means that there is a higher long-run effect of MPD on gasoline prices in terms of Eurocents. At the same time, follow-on information campaigns, such as local radio reports about fuel prices, appear to be able to strengthen the effect of MPD.

There are two implications for policy that we draw from this analysis: First, assessing the level of consumer information prior to mandatory price disclosure is essential. If few consumers are well informed, mandatory price disclosure can lead to important price reductions. Should most consumers already be well informed, the pro-competitive potential of MPD is limited. Second, making price information available may not be sufficient to reap the pro-competitive benefits. We find that when public attention to the policy declines, so do the price effects of MPD. However since local radio reports are able to deliver a pro-competitive follow-on information shock, policymakers could achieve the same by regularly pushing for large-scale information adoption through public information campaigns.

Chapter 3

Does Tax Policy Work When Consumers Have Imperfect Price Information? Theory and Evidence

3.1 Introduction

Understanding how and when firms pass through taxes to consumers is fundamental for the design of optimal tax policy. Pass-through determines the corrective effect of Pigouvian taxes, the effectiveness of unconventional fiscal policy to stimulate the economy and the distributional consequences of any commodity tax. Chetty, Looney, and Kroft (2009) show that under perfect competition, the pass-through of taxes decreases the less salient they are as part of the price paid by consumers. Weyl and Fabinger (2013) extend the theoretical analysis of pass-through to oligopolistic markets with perfect information.[1] They find that pass-through decreases in the aggregate price elasticity of demand. Although consumers are rarely omniscient about prices, little is known about how pass-through behaves when consumers have imperfect price information.

This chapter is based on joint work with Felix Montag and Monika Schnitzer (Montag, Sagimuldina, and Schnitzer, 2021).

[1] Miravete et al. (2018) apply this analysis to the estimation of the Laffer curve under oligopolistic competition empirically.

In this chapter, we propose a new theoretical framework to analyze commodity tax pass-through in oligopolistic markets where consumers have imperfect information about prices. We derive theoretical predictions about the pass-through rates as a function of the information consumers have about market prices and as a function of the number of sellers. We find that the more consumers are well informed about prices, the higher is the pass-through rate. We also show that there is a hump-shaped relationship between the number of sellers and pass-through. To test our predictions empirically, we study heterogeneities in the pass-through of a tax decrease and a subsequent tax increase in the German retail fuel market. We show that, as predicted by the theory, pass-through increases in how well consumers are informed about prices. We also find evidence for a hump-shaped relationship between pass-through and the number of fuel stations in a local market.

For our theoretical analysis, we adapt the consumer search model by Stahl (1989) to the analysis of tax pass-through. This model distinguishes between fully informed shoppers (who know all prices) and uninformed non-shoppers (who can search for prices sequentially). This framework allows us to introduce a novel notion of price sensitivity of demand to the analysis of tax pass-through: The larger the number of informed consumers, the more it pays for sellers to compete for them with their choice of prices. Price sensitivity of demand, as experienced by sellers, therefore depends on how many consumers have access to an information clearinghouse and are thus perfectly informed.

In equilibrium, firms set prices by randomizing according to a mixed strategy. Informed shoppers know all prices in the market, always buy from the lowest-price seller and therefore pay the minimum price. Uninformed non-shoppers draw the first price for free and then pay a search cost to draw more prices. In equilibrium, prices are chosen such that they do not search and thus pay the first price they draw. From an ex ante point of view, informed shoppers pay the expected minimum price, while uninformed non-shoppers pay the expected price.

The model has two key predictions about how competition affects pass-through. First, the larger the share of price sensitive consumers, the higher is the pass-through rate to all prices. Second, the larger the number of firms in the market, the larger is the pass-through rate to the expected minimum price, paid by informed shoppers. In contrast, the pass-through rate to the expected price, paid by uninformed non-shoppers, first increases and then decreases in the number of sellers. The latter effect can be explained by the fact that above a certain threshold, as more firms are active in the market, it becomes less and less likely for a particular firm to attract shoppers and so firms are more likely to charge a higher price and only serve uninformed non-shoppers. Thus, in a context with imperfect information about prices, a larger number of sellers does not monotonically lead to a more competitive outcome. How pass-through to the average price paid by consumers in the market varies with the number of firms depends on the share of informed and uninformed consumers in the market. These predictions are true for the pass-through of ad-valorem taxes, per unit taxes, as well as symmetric marginal cost shocks.

Next, we test our theoretical predictions by studying two important tax changes in the German retail fuel market. As part of the fiscal response to the COVID-19 pandemic, the German government announced a six-month temporary value-added tax (VAT) reduction on 3 June 2020, taking effect on 1 July on most products, including fuel. On 1 January 2021, the VAT rate returned back to its original level. At the same time, the government introduced a carbon tax on fuel.[2] We estimate pass-through of the tax decrease as well as the two tax increases to diesel and gasoline prices using a unique dataset containing the universe of price changes at fuel stations in Germany and France before and after the policy change.

To estimate pass-through, we use the synthetic difference-in-differences (SDID) recently introduced by Arkhangelsky et al. (2021). This method combines the advantages of difference-in-differences (DID) and synthetic control (SC). To analyze how price sensitivity affects pass-through, we compare daily prices of the three main fuel types sold at fuel stations in Germany and France.

There is strong evidence suggesting that diesel drivers are on average more price sensitive than drivers fueling gasoline. Frequent drivers tend to use diesel cars. On average, diesel car users drive twice as many kilometres per year than gasoline drivers. By buying a car with a more expensive diesel engine, they make a fixed cost investment to decrease their marginal cost of driving. This suggests that diesel drivers have a greater incentive to become informed about fuel prices.[3] Using data on search querries from a smartphone app displaying fuel prices to users,

[2]For simplicity, we will frequently refer to the policy change on 1 July 2020 as the tax decrease and the change on 1 January 2021 as the tax increase.

[3]Johnson (2002) made a similar argument for why diesel drivers are more price sensitive.

we confirm empirically that the search intensity among diesel drivers is higher. Within gasoline, the evidence strongly suggests that customers of *E5* are less price sensitive than *E10* customers.

We find that the pass-through rate of the tax decrease (tax increase) is 79 (92) percent for diesel, whereas it is 52 (75) percent for *E10* and 34 (69) percent for *E5*. As predicted by the theoretical model, the higher the price sensitivity of consumers, the higher the pass-through rate. Since the same stations sell all three types of fuel, unobserved station characteristics cannot explain these differences.

Finally, we use the geolocation and brand information of fuel stations to compute the number of rival fuel stations within a local market. We then estimate how the pass-through rate varies with the number of rival stations. Consistent with our theoretical predictions, we find that the pass-through rate first increases and then decreases in the number of rival fuel stations within a local market. Empirically, this relationship seems to disappear when pass-through is very high.

Our chapter makes two main contributions. First, we introduce a novel notion of price sensitivity to the theoretical analysis of pass-through in oligopolistic markets. How well consumers are informed about prices affects the equilibrium intensity of competition in the market. We find that the more price sensitive consumers are on average, the higher is the pass-through rate. This is different to how another common notion of price sensitivity, the price elasticity of demand, affects pass-through. A classic result under perfect competition is that the higher the price elasticity of demand, the lower the pass-through rate. Weyl and Fabinger (2013) show that this result extends to models with imperfect competition.[4] Our notion of price sensitivity is different, in that there is no aggregate quantity response of consumers. Instead, we capture how likely it is that consumers seek out buying their fixed quantity from the cheapest seller.[5]

In contrast to the context studied by Chetty et al. (2009), in our context taxes are salient for all buyers. Thus, the pass through in our model is not a function of salience as in Chetty et al. (2009)'s context of perfect price competition, but a function of price sensitivity of consumers,

[4]More precisely, this holds true for the market-level price elasticity of demand. In oligopolistic markets, a higher price elasticity of demand decreases pass-through via an aggregate quantity response and increases pass-through by intensifying competition. Weyl and Fabinger (2013) show that which of these effects is larger depends on the relative elasticities of demand and supply and the curvature of demand. Previous work (see, e.g., Stern, 1987 or Hamilton, 1999) studied tax pass-through in a Cournot model. All of these studies focus on settings with perfect information. Instead, we focus on settings where consumers have imperfect information about prices.

[5]This can be thought of as the price sensitivity of the residual demand that a particular seller faces, whilst market demand remains unchanged.

which in turn affects the intensitiy of price competition. Chetty et al. (2009) shows that consumers underreact to commodity taxes if they are not salient. Increasing tax salience in Chetty et al. (2009)'s context and increasing consumer information about the sum of price and taxes when there is imperfect competition as in our model therefore have opposite effects on pass-through.

Second, we provide novel empirical evidence on the determinants of commodity tax pass-through and relate them to our theoretical predictions. A unique feature of our empirical setting is that close to all fuel stations sell all three types of fuel. This allows us to disentangle the two different aspects of imperfect competition: the fact that consumers are imperfectly informed about prices and the fact that the market is oligopolistic with a small number of competitors. We can therefore test how the pass-through rate differs for consumer groups that differ in their price sensitivity whilst holding the network of stations constant. We can also test how pass-through varies when we hold the price sensitivity constant and vary the number of competitors. In contrast to the previous literature, our setting allows us to disentangle these two mechanisms empirically within the same study. Finally, studying a tax decrease and a subsequent tax increase six months later strengthens the robustness of our results.

To the best of our knowledge, there are no previous empirical studies that combine the analysis of these two mechanisms. Furthermore, our explanation as to why pass-through increases when consumers are better informed is new to the literature. Reassuringly, our theoretical framework can encompass and reconcile previous empirical observations. Duso and Szücs (2017) find that cost pass-through is higher for competitive electricity tariffs, which consumers need to actively seek out, than for default tariffs. Kosonen (2015) finds that after a VAT decrease, Finnish hairdressers cut prices more for advertised services. Genakos and Pagliero (2022) find that tax pass-through by fuel stations on isolated Greek islands increases in the number of stations. Miller, Osborne, and Sheu (2017) find that cost pass-through in the cement industry decreases in the number of competitors. In our model, we predict a hump-shaped relationship between the pass-through rate and the number of competitors, which means that both empirical results can be consistent with our model. Kopczuk, Marion, Muehlegger, and Slemrod (2016) find no strong correlation between industry concentration and pass-through of diesel taxes. They therefore conclude that market power is unlikely to play an important role in explaining pass-through. Our results suggest that with imperfect price information, concentration may not be a good proxy for competition.

More generally, we extend a growing empirical literature on pass-through of tax or cost changes. There are numerous studies that, as an intermediate or final step, estimate average pass-through rates.[6] However, few investigate their determinants. Notable exceptions are Miravete et al. (2018), Hollenbeck and Uetake (2021) and Nakamura and Zerom (2010), who study the interplay between pass-through and market power. Miravete et al. (2018) show empirically that market power reduces pass-through and therefore changes the Laffer curve. Not accounting for non-competitive pricing thus leads to an ineffective tax policy. Hollenbeck and Uetake (2021) find that imperfect competition and log-convex demand is responsible for over-shifting in the legal marijuana industry. Nakamura and Zerom (2010) find that exchange rate pass-through is reduced by local costs and markup adjustments. Our study differs in that we analyze how informational frictions on the consumer side determine pass-through. This also gives policymakers a possible angle on how to increase the pass-through rate, for example by mandating price transparency.[7]

Within our setting, we can also study the speed of, and asymmetries in, pass-through. Like Benzarti, Carloni, Harju, and Kosonen (2020), we find higher pass-through for the tax increase than for the tax decrease. Using monthly sales data for home appliances, Büttner and Madzharova (2021) show that VAT pass-through is full and relatively fast. Similarly, Fuest, Neumeier, and Stöhlker (2020) find full pass-through of the 2020 German temporary VAT reduction at supermarkets of the Rewe Group.[8] Our results indicate that although pass-through of both tax changes is fast, it remains incomplete even two months after the tax change.

[6] Some studies focus on particular industries, such as energy markets (see, e.g., Fabra and Reguant, 2014, Kopczuk et al., 2016, Li and Stock, 2019 or Ganapati, Shapiro, and Walker, 2020) or sin products (see, e.g., Dubois, Griffith, and O'Connell, 2020, Harding, Leibtag, and Lovenheim, 2012 or Conlon and Rao, 2020). Others estimate the average pass-through rate across a large number of industries (see, e.g., Benedek, De Mooij, Keen, and Wingender, 2019). The findings of these studies are mixed, as they include evidence for under-shifting (e.g. Benzarti and Carloni, 2019, Carbonnier, 2007), full pass-through (e.g. Benedek et al., 2019) and over-shifting (e.g. Besley and Rosen, 1999).

[7] Luco (2019), Ater and Rigbi (2023) and Montag and Winter (2020) study the effect of different mandatory price disclosure policies and find mixed results.

[8] Jacob, Müller, and Wulff (2021) find higher pass-through of the corporate tax by fuel stations in municipalities with fewer stations. This differs from unit and ad-valorem taxes as the corporate tax is levied on profits, with a partial deductibility of costs.

Our results not only inform policymakers aiming to set optimal Pigouvian taxes, but also the use of unconventional fiscal policy to stimulate the economy. This describes the use of temporary tax cuts or pre-announced tax increases to stimulate inflation by targeting household expectations directly.[9] For temporary tax cuts to stimulate inflation expectations and consumption, consumers need to expect that prices will rise after the tax increases again. This is most likely the case if the temporary tax cut and the pre-announced tax increase are passed-through to consumers. Since we find that pass-through increases in the price sensitivity of consumers, our results indicate that targeting such measures at markets where the price sensitivity of consumers is high can increase the cost effectiveness of unconventional fiscal policy.

Finally, we extend the empirical literature on pricing in retail fuel markets. Whereas Houde (2012) models fuel stations as differentiated by station locations but abstracts from imperfect information, recent studies found that models of imperfect information and consumer search are well-suited to explain empirical findings in retail fuel markets.[10] We extend this literature by combining a theoretical model with incomplete information and granular data on fuel prices to study the pass-through of taxes in retail fuel markets.[11]

The remainder of this chapter is structured as follows: Section 3.2 outlines the theoretical model, Section 3.3 describes the industry, Section 3.4 gives an overview of the data and presents descriptive evidence, Section 3.5 discusses the empirical strategy, Section 3.6 presents the estimation results and Section 3.7 concludes.

3.2 Theoretical Model

Our aim is to analyze theoretically how pass-through varies with the price sensitivity of consumers and the number of sellers. We therefore set up a model where firms sell a homogeneous

[9] See, for example, D'Acunto, Hoang, and Weber (2018), or D'Acunto, Hoang, and Weber (2022).

[10] These include Chandra and Tappata (2011), Byrne and de Roos (2017), Byrne and de Roos (2022) or Pennerstorfer et al. (2020).

[11] There is a large empirical literature on cost pass-through in retail fuel markets using error correction models and testing the rockets-and-feathers hypothesis, which focuses on asymmetric pass-through of increases and decreases (e.g. Bachmeier and Griffin, 2003, Deltas, 2008 or Verlinda, 2008) and the speed of pass-through (e.g. Johnson, 2002). Most of these studies do not provide a theoretical explanation for their findings. A notable exception is Borenstein, Cameron, and Gilbert (1997), who show that asymmetric pass-through could either be explained by tacit collusion or by imperfect information. For a review of the literature, see Eckert (2013). Furthermore, Deltas and Polemis (2020) shows that many of the conclusions from studies using error correction models to estimate pass-through rates may strongly depend on research design and data features.

good to consumers who are either fully informed about prices or can search for lower prices. The model is based on the rich literature on consumer search in industrial organization, and in particular on the model by Stahl (1989). We extend this model by introducing marginal costs and an ad-valorem tax in order to be able to analyze tax pass-through.

3.2.1 Setup

There is a mass M of consumers. Each consumer has the same valuation v for the homogeneous good and inelastically demands one unit of the product. A fraction ϕ of consumers are fully informed shoppers and $1 - \phi$ are non-shoppers, who can search sequentially. Shoppers know prices of all sellers and therefore always buy from the lowest price seller. If there is a tie, shoppers are shared equally among the lowest price sellers. Non-shoppers only know the distribution of prices and draw a first price for free. They can then choose to randomly draw prices of additional sellers at an incremental search cost s, in the hope of finding a lower price. Non-shoppers buy the good as soon as the price that they draw is weakly below their reservation price p_r, at which non-shoppers are indifferent between accepting the price and drawing a new price at search cost s, because the expected price savings of drawing another price are equal to the search cost s.

On the supply side, there is an infinite number of symmetric firms that can potentially enter the market. Each firm can enter the market for a fixed and sunk cost F and produce at a constant marginal cost of c. The number of entrants is denoted by N and firms are indexed by i. Finally, sales are subject to an ad-valorem tax τ.

The game proceeds in two stages. In the first stage, firms decide whether to enter the market. In the second stage, sellers first choose prices and consumers then make search and purchase decisions. To find the subgame perfect Nash equilibrium of the game, we solve it via backward induction.

Before proceeding any further, we should define some more notation. When discussing prices, we always refer to the price paid by consumers. We assume that sellers bear the initial incidence of a tax and then (partially) "pass through" the cost of the tax to consumers. It is a well known result from the theoretical literature that equilibrium prices should be equivalent, irrespective of whether the initial tax incidence is with buyers or sellers. The pass-through rate of marginal costs is $\rho_c = \frac{\partial p}{\partial c}$. Note, that the pass-through rate of a per unit tax is equivalent to the pass-through rate of marginal costs. The pass-through rate of the ad-valorem tax is

$$\rho_\tau = \frac{\partial p}{\partial \tau} \cdot \frac{1+\tau}{p} .$$

In the following, we focus on what determines the pass-through rate of the ad-valorem tax. As we show in Appendix C.1.3, the determinants of the pass-through rate of marginal costs or per unit taxes are qualitatively equivalent.

Finally, it is worth discussing the notion of price sensitivity in this model. Whereas many canonical models analyzing pass-through rates think of the sensitivity of consumers to prices in terms of the price elasticity of demand, our notion of price sensitivity is different. As described above, all consumers always inelastically demand a single unit of the good so long as the price is below their valuation. There is thus no response in the aggregate quantity if prices change.

Instead, we capture a different way of how consumers are sensitive to prices, namely through the share of shoppers ϕ and the incremental search cost of non-shoppers s. If there are more shoppers, then a larger share of consumers is going to buy from the lowest price seller for sure. This decreases the expected profit of setting a price that is not the lowest price in the market. If the search cost of non-shoppers is lower, then non-shoppers are more willing to search for lower prices. This decreases the reservation price of non-shoppers and also leads to lower prices.

3.2.2 Stage 2: Equilibrium price distribution

In the following, we characterize the equilibrium while the analysis of the model is relegated to Appendix C.1. There exists no pure strategy equilibrium in prices. There is a unique symmetric mixed strategy equilibrium where all sellers draw a price from the interval $[\underline{p}, p_r]$ according to the distribution $F(p_i)$, where p_r is the reservation price of non-shoppers and $\underline{p}$ is the minimum price a seller will charge. Shoppers always buy from the lowest price seller, whereas non-shoppers draw a single price and buy at this price. In equilibrium, non-shoppers do not search sequentially, because any price they draw is below their reservation price.

The symmetric equilibrium pricing strategy is characterized by the equilibrium objects p_r, $\underline{p}$ and $F(p_i)$. The reservation price of non-shoppers is

$$p_r = \begin{cases} E[p] + s & \text{if } E[p] + s < v \\ v & \text{otherwise} \end{cases} .$$

If searching sequentially is sufficiently cheap, the reservation price of non-shoppers is the sum of the expected price at the next draw and the search cost s. With relatively high search costs, the reservation price of non-shoppers is simply the valuation of the good v.

The minimum element of the support from which sellers draw prices in equilibrium is

$$\underline{p} = \frac{p_r}{\frac{\phi N}{1-\phi} + 1} + c\frac{1+\tau}{1 + \frac{1-\phi}{\phi N}} .$$

The cumulative density function of the equilibrium pricing strategy is

$$F(p_i) = 1 - (\frac{p_r - p_i}{p_i - c(1+\tau)} \frac{1-\phi}{N\phi})^{\frac{1}{N-1}} .$$

The expected second stage profits (i.e. excluding the fixed and sunk cost of entry) of a seller are

$$E[\pi_i] = (\frac{p_r}{1+\tau} - c)\frac{1-\phi}{N}M .$$

Two further objects are of interest for our analysis, namely the expected price and the expected minimum price. Since non-shoppers do not search in equilibrium, they always buy at the first price they draw and thus the expected price is also the average price paid by non-shoppers. In contrast, shoppers always buy from the lowest price seller and thus the expected minimum price is also the average price paid by shoppers.[12]

The expected price is

$$E[p] = \underline{p} + (\frac{1-\phi}{N\phi})^{\frac{1}{N-1}} \int_{\underline{p}}^{p_r} (\frac{p_r - p}{p - c(1+\tau)})^{\frac{1}{N-1}} dp .$$

[12]The average refers to the average price paid by shoppers and non-shoppers if this game is often repeated across time or space. At a given time and location there is, of course, only one minimum price and N prices.

The expected minimum price is

$$E[p_{min}] = \frac{1-\phi}{\phi}[p_r - E[p] + (p_r - c(1+\tau))c(1+\tau)\int_{\underline{p}}^{p_r} \frac{1}{(p-c(1+\tau))^2}F(p)dp] \,.$$

3.2.3 Stage 1: Equilibrium entry

Entry occurs so long as the expected second stage profits of the entrant are greater or equal to the fixed and sunk cost of entry F. No further entry occurs if the next potential entrant cannot expect to recoup her entry costs.

The equilibrium number of entrants N^* will thus be such that

$$(\frac{p_r}{1+\tau} - c)\frac{1-\phi}{F}M - 1 < N^* \leq (\frac{p_r}{1+\tau} - c)\frac{1-\phi}{F}M \,. \tag{3.1}$$

Note that increasing the market size M (or decreasing the fixed cost F) directly translates into a higher number of active sellers and does not enter the equilibrium in any other way. At the same time, different numbers of active sellers lead to different intensities of competition. Thus, whenever we analyze how prices or pass-through vary with the number of active sellers we should think of this as variation in the local market size or the fixed cost of entry.

For the remainder of the analysis we will assume that there is no entry and treat the number of sellers as exogenous. This is because our empirical study is concerned with a short-term tax adjustment during which entry seems unlikely. In other applications it will make sense to endogenize the number of active sellers also for the analysis of pass-through. Unless otherwise stated, we focus on the case where $N^* \geq 2$, since for the informedness of consumers to matter there need to be at least two sellers active in the market.

3.2.4 Pass-through of an ad-valorem tax

We now turn to analyzing how ad-valorem taxes are passed through to consumers. We begin by studying how an increase in the ad-valorem tax τ affects the equilibrium pricing strategy. To simplify the analysis, we assume that the search cost s is sufficiently high, such that $p_r = \upsilon$.

We relax this assumption in Appendix C.1.5 and simulate how pass-through rates evolve with sequential search.[13] We show that qualitatively our results hold when search costs are low.

Since the reservation price now corresponds to the valuation of the good, only the minimum element of the support and the density of the pricing strategy are affected by a change in ad-valorem taxes.

Proposition 3.1. *With $0 < \phi < 1$, for any $\hat{\tau} > \tau$ the minimum element of the support of the equilibrium pricing strategy $\underline{\hat{p}} > \underline{p}$ and the Nash equilibrium pricing strategy with τ first-order stochastically dominates (FOSD) the pricing strategy with $\hat{\tau}$, i.e. $\hat{F}(p) \leq F(p) \quad \forall p$.*

This means that if the share of shoppers is strictly positive, an increase in the ad-valorem tax τ leads to a shift in the support of prices from which sellers draw in equilibrium towards higher prices. It also means that, for each price on this support, the likelihood that a drawn price is lower than said price decreases if the ad-valorem tax rate increases to $\hat{\tau}$.

Figure 3.1: Ad-valorem tax pass-through to the equilibrium pricing strategy

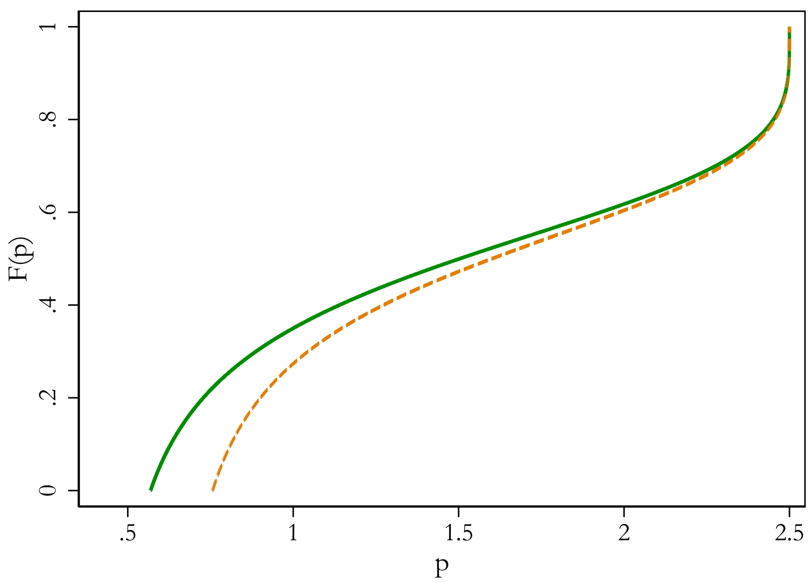

Note: The Figure shows simulation results of how the distribution from which sellers draw prices in the symmetric Nash equilibrium changes if the ad-valorem tax increases from τ to $\hat{\tau}$. The solid line corresponds to the distribution under τ. The dashed line corresponds to the distribution under $\hat{\tau}$. Parameter values: $v = 2.5$, $s = 0.75$, $c = 0.4$, $\tau = 0.1$ and $\hat{\tau} = 0.6$.

[13] An alternative simplification would be setting $N = 2$, which we consider to be less desirable for the purpose of this analysis.

As the share of shoppers converges to zero, the Nash equilibrium converges towards a degenerate distribution at the monopoly price, the classical result by Diamond (1971). The monopoly price corresponds to the valuation of the good, υ.

Since the minimum element of the support of prices and the density function monotonously move towards higher prices, other moments of interest, such as the expected price $E[p]$, which is the average price paid by non-shoppers, and the expected minimum price $E[p_{min}]$, which is the average price paid by shoppers, also increase. Thus, if ad-valorem taxes increase then the expected price paid increases for all consumers.

3.2.5 The effect of price sensitivity on the pass-through rate

We now turn to analyzing how the pass-through rate of an ad-valorem tax τ varies with the price sensitivity of consumers.

Proposition 3.2. *If the share of shoppers $\phi = 0$, pass-through of the ad-valorem tax $\rho_\tau = 0$. If $\phi = 1$, there is full pass-through, i.e. $\rho_\tau = 1$. As $\phi \rightarrow 1$, the pass-through rate $\rho_\tau \rightarrow 1$.*

Let us begin by analyzing two extreme cases. As we saw already, if there are no shoppers at all the Nash equilibrium is a degenerate distribution at the monopoly price, which is independent of the ad-valorem tax. Thus, if there are no shoppers, pass-through is zero. On the other hand, as the share of shoppers converges to one, the Nash equilibrium converges to the classical result by Bertrand (1883), where the Nash equilibrium is a degenerate distribution at $c(1 + \tau)$. Thus, if all consumers are shoppers, there is full pass-through of the ad-valorem tax.

Finally, for all values of ϕ between zero and one, we can show that the pass-through rate of the ad-valorem tax to the lower bound of the equilibrium price strategy is strictly increasing in ϕ. We can also show that the rate at which an increase in the ad-valorem tax from τ to $\hat{\tau}$ reduces the probability that a drawn price is below a particular price p, i.e. from $F(p)$ to $\hat{F}(p)$, strictly increases in the share of shoppers. Thus, the pass-through rate of the ad-valorem tax increases in the share of shoppers and converges to full pass-through as the share of shoppers converges to one.

3.2.6 The effect of the number of sellers on the pass-through rate

So far, we saw that a higher share of informed consumers increases the intensity of competition and leads to higher pass-through. However, the model also contains a second dimension of competition, the number of active sellers. This is considered more often in empirical applications, since it is more salient and easier to observe than the informedness of consumers. We therefore ask how pass-through varies with the number of active sellers.

Proposition 3.3. *With $0 < \phi < 1$, as $N \to \infty$ the pass-through of τ to the minimum element of the equilibrium price support converges to full pass-through, i.e. $\rho_{\tau,\underline{p}} \to 1$.*

As the number of sellers increases, competition for shoppers becomes more intense and so the minimum price that sellers consider charging in the symmetric Nash equilibrium converges towards $c(1 + \tau)$. As this occurs, the pass-through rate of the ad-valorem tax to $\underline{p}$ increases.

Showing how an increase in N affects the pass-through rate of ad-valorem taxes to $F(p)$, $E[p]$ and $E[p_{min}]$ analytically turns out to be more difficult. Instead, we resort to simulating how the pass-through rate varies with N.

In a setting without taxes or marginal costs but for a wider class of demand functions, Stahl (1989) shows that for a sufficiently high N', for $N > N'$ the equilibrium price distribution converges to a degenerate price distribution at the monopoly price as $N \to \infty$. At the same time, we know that as N increases from one to two, prices in the symmetric Nash equilibrium move from a degenerate distribution at the monopoly price to a competitive price distribution that includes prices below the monopoly price. Thus, the expected price first decreases and then increases again as $N \to \infty$. We also showed that as prices converge to the monopoly price, the pass-through rate converges to zero. Therefore, we expect the pass-through rate of ad-valorem taxes to $E[p]$ to first increase and then decrease as $N \to \infty$.

When we analyzed how pass-through varies with the share of shoppers, $E[p]$ (paid by non-shoppers) and $E[p_{min}]$ (paid by shoppers), as well as pass-through rates to these prices, always moved in the same direction. As $N \to \infty$, this is different. When s is sufficiently high such that $p_r = \upsilon$, $E[p_{min}]$ monotonously decreases in N and the pass-through rate of the ad-valorem tax to $E[p_{min}]$ monotonously increases.[14] This is because although each individual seller is more

[14] As we show in Appendix C.1.5, for some values of ϕ there is an intermediate range of values in which ρ_c to $E[p_{min}]$ decreases in N, after which it increases again. This is because $\underline{p}$ is a function of p_r.

likely to charge higher prices, with an increase in N and a decrease in $\underline{p}$, it is overall more likely that some seller will set a lower price to attract shoppers.

Figure 3.2: Pass-through of τ to $E[p]$

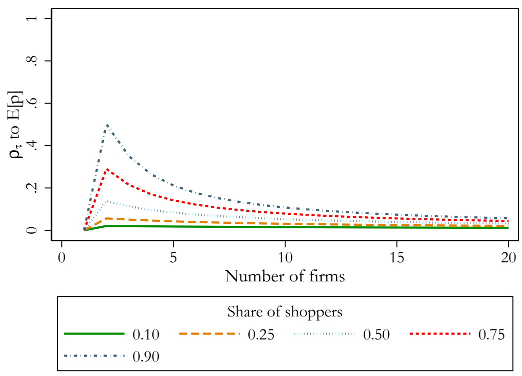

Parameter values: υ = 2.5, c = 0.4, τ = 0.2 and $\hat{\tau}$ = 0.22.

Figure 3.3: Pass-through of τ to $E[p_{min}]$

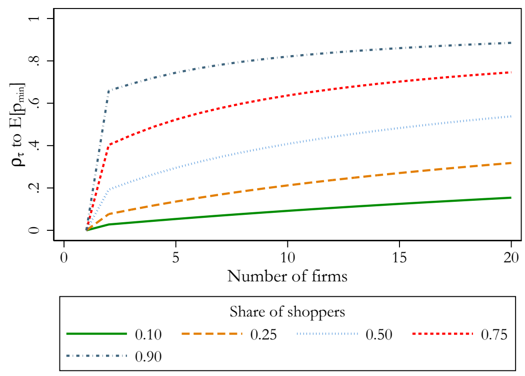

Parameter values: υ = 2.5, c = 0.4, τ = 0.2 and $\hat{\tau}$ = 0.22.

The simulation results in Figures 3.2 and 3.3 are in line with our expectations. As N increases, pass-through of the ad-valorem tax to the expected price first increases and then decreases. Pass-through to the expected minimum price always increases.

Finally, since prices paid by shoppers and non-shoppers evolve differently, we may be interested in how ad-valorem taxes are passed through to the expected average price paid by consumers in the markets. Fortunately, since both consumers types consume the same quantities and we know the share of each type of consumer, this can easily be considered.

Figure 3.4: ρ_τ to $E[\phi p_{min}+(1-\phi)p]$, $p_r = \upsilon$

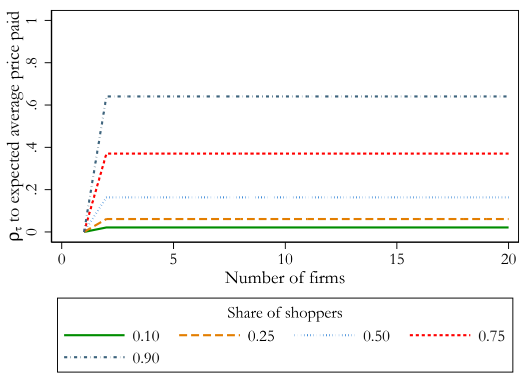

Parameter values: υ = 2.5, c = 0.4, τ = 0.2 and $\hat{\tau}$ = 0.22.

Figure 3.5: ρ_τ to $E[\phi p_{min} + (1 - \phi)p]$, p_r endogenous

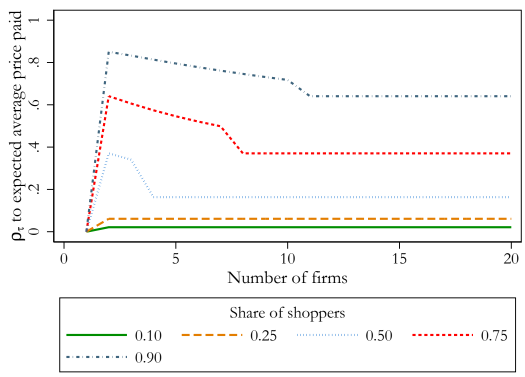

Parameter values: υ = 2.5, s = 0.75, c = 0.4, τ = 0.2 and $\hat{\tau}$ = 0.22.

The simulation in Figure 3.4 shows that when search costs are so high that $p_r = \upsilon$, pass-through of ad-valorem taxes first increases in N and then stays constant, because the decrease in pass-through to $E[p]$ and the increase in pass-through to $E[p_{min}]$ cancel each other out. Figure 3.5 shows that if search costs s are sufficiently low such that p_r is endogenous, pass-through to the expected average price paid first increases in N, then decreases in N and, as $p_r \to \upsilon$ when $N > 2$ and $N \to \infty$, ad-valorem tax pass-through remains constant when N is sufficiently large.

Thus far, when analyzing pass-through, we studied short-run responses in prices and thus held the number of sellers constant. Although our empirical application focuses on a temporary decrease in the VAT, and so is unlikely to induce entry, it is nevertheless worth discussing long-run responses. As we saw in the analysis of the entry stage in Section 3.2.3, an increase in the ad-valorem tax reduces the equilibrium number of sellers in the market. If the pre-change N is such that pass-through increases in N (i.e. very low levels of N), long-run pass-through is lower than short-run pass-through. If the pre-change N is such that pass-through decreases in N (i.e. sufficiently high N), long-run pass-through is higher than short-run pass-through.

3.3 The Retail Fuel Market

We now turn to the description of the retail fuel market in Germany. In 2019, total revenues from retail fuel sales were worth 92 billion Euro or approximately 3 percent of German GDP. In addition to its standalone value to the economy, this market has large externalities on the rest of the economy. Fuel prices are a key determinant of travel costs, commuting costs and, more broadly, the cost of personal transportation.

3.3.1 Diesel vs. gasoline

The first important distinction to make within fuels for passenger vehicles is between diesel and gasoline.[15] In Germany, diesel has a volume share of 44 percent of fuel for passenger vehicles

[15] Since fuel stations do not report prices for truck diesel to the Market Transparency Unit, we only focus on fuel prices for passenger vehicles.

with combustion engines and gasoline accounts for the remaining 56 percent.[16] Substituting between these two types of fuel is very costly, both on the demand and supply side.[17] In the short-term, these can be considered as separate markets.

Drivers of diesel and gasoline cars differ in how much they drive. Whereas only 32 percent of registered passenger vehicles in Germany have a diesel engine, compared to 66 percent that run on gasoline, frequent drivers often buy diesel cars.[18] On average, gasoline passenger vehicles drive $10,800$ kilometers, whereas diesel passenger vehicles drive $19,500$ kilometers per year.[19]

The reason why frequent drivers buy diesel cars whereas less frequent drivers buy cars with a gasoline engine is that buying a diesel car is more expensive, but the cost of fuel at the pump is lower. Buying a diesel car is therefore a fixed cost investment to lower the marginal cost of driving. Drivers that select into buying a diesel engine thus do so based on their cost sensitivity and their incentive to save on fuel costs due to the distances they drive every year.

We verify this claim using data on search queries in 2015 from a major smartphone app displaying fuel prices to users in Germany. Figure 3.6 shows the daily number of price searches by fuel type on a major German smartphone app per 1,000 diesel or gasoline vehicles in circulation. The ratio of price searches to the number of vehicles in circulation is around 54 percent higher for diesel than for gasoline. This shows that the search intensity among drivers of diesel-run vehicles is significantly higher than among drivers of gasoline-run vehicles. It therefore strongly suggests that diesel drivers are more price sensitive.

A frequently made observation is that commercial vehicles usually run on diesel and this may affect the average price sensitivity of drivers by fuel type. Although we showed that drivers of diesel vehicles search more, it is worth briefly discussing why commercial vehicles are not a concern. First, as of 1 January 2021 there were around 15 million passenger vehicles with a diesel engine, but, including those with a gasoline engine, only 5.1 million commercial pas-

[16] This is based on 2018 figures from *Verkehr in Zahlen 2019/2020*, published by the Federal Ministry of Transportation. To the best of our knowledge, these are the most recent administrative figures concerning the passenger vehicle market only.

[17] On the demand side, this would usually require buying a new vehicle. On the supply side, readjusting the ratio of diesel and gasoline made from a barrel of crude oil is possible, but only to a limited extent and at a high cost.

[18] This is based on April 2020 figures on registered passenger vehicles in Germany, published by the German Federal Motor Transport Authority.

[19] This is based on 2018 figures from *Verkehr in Zahlen 2019/2020*, published by the Federal Ministry of Transportation.

Figure 3.6: Consumer search patterns (Germany)

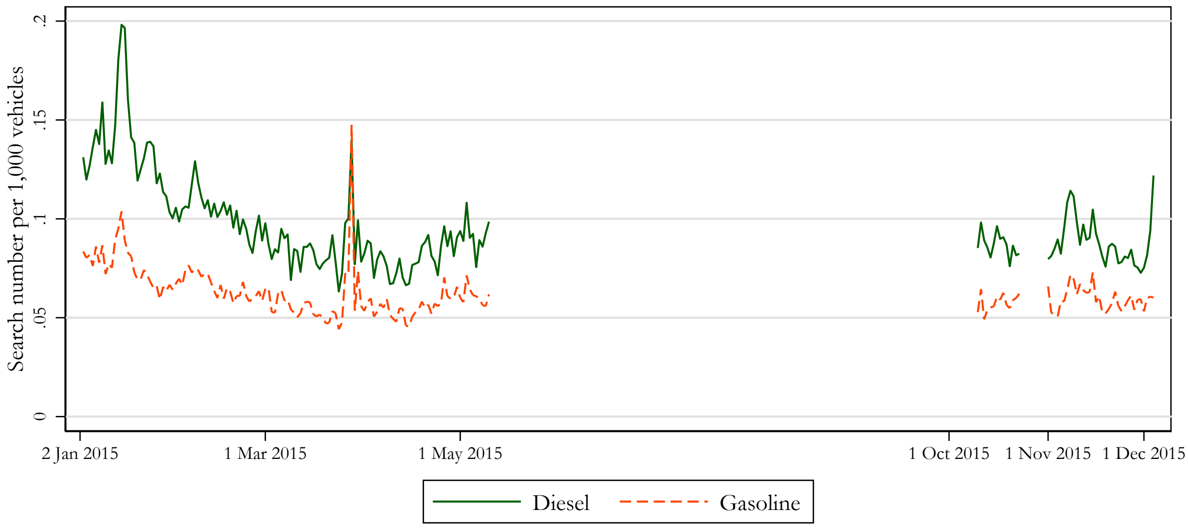

Notes: The Figure shows the daily number of price searches by fuel type on a major German smartphone app per 1,000 diesel or gasoline vehicles in circulation. The data is available for January to May and October to December 2015. The solid line corresponds to the search intensity for diesel. The dashed line corresponds to the search intensity for gasoline.

senger vehicles (Kraftfahrt-Bundesamt, 2021). At the very least, 66 percent of passenger cars with a diesel engine are therefore owned by private individuals. In addition, commercial vehicle drivers may also have an incentive to reduce fuel costs, such as those receiving a lump-sum (or distance-based) fuel allowance or those that are self-employed. The fact that many commercial vehicles run on diesel therefore does not call into question our finding that drivers of vehicles that run on diesel are, on average, more price sensitive than drivers of vehicles running on gasoline.

3.3.2 *E5* vs. *E10*

Within gasoline, there is differentiation according to the octane rating and the share of ethanol. Standard gasoline (commonly referred to as *Super*) has an octane rating of 95. It has a volume share of 95.4 percent of the gasoline market.[20] *Super Plus* accounts for the remaining volume

[20]This is based on 2019 figures from the monthly oil statistics, published by the Federal Office for Economic Affairs and Export Control.

and is gasoline with an octane rating of 98, required by some high-performance vehicles. We do not consider *Super Plus* for the remainder of our analysis.[21]

Within *Super*, we can further distinguish according to the ethanol share. Standard gasoline has a 5 percent share of ethanol and is thus commonly referred to as *E5*. In 2011, a new type of gasoline was introduced in Germany with a 10 percent ethanol share, referred to as *E10*. The aim of increasing the share of ethanol is to reduce greenhouse gas emissions and decrease the amount of fossil fuel used in transportation. Although *E5* and *E10* are not taxed differently, *E10* is usually around 4 to 5 Eurocent cheaper than *E5*. This is partly driven by the relative prices of crude oil and ethanol on the world market and partly by a minimum quota of biofuels that need to be sold by fuel stations every year.

After the introduction of *E10* in 2011, there was controversy about whether biofuels damage the engine. Although biofuels can pose a significant threat to the engine of a vehicle that is not certified to be compatible with *E10*, around 90 percent of gasoline-run vehicles, including all vehicles produced after 2012, are compatible with *E10*.[22] According to the German Automobile Association, *E10* is around 1.5 percent less efficient than *E5*.[23] This cannot fully account for the observed difference in *E5* and *E10* prices. All fuel stations in Germany are required to sell both types of fuel. Nevertheless, in 2019 *E5* still had a volume share of 85.6 percent within *Super* and *E10* only of 14.4 percent. Overall, many motorists who could buy less expensive *E10* choose not to do so and buy *E5* instead. Reasons for this could include preferences or a lack of information, which point towards a lower price sensitivity of *E5* customers compared to *E10* customers.

Recent findings by the German Automobile Association confirm this view. According to a survey conducted in Fall 2020, the most cited reason for fueling *E10* was its lower price (72 percent among respondents fueling *E10*), followed by concerns for the environment (37 percent). Amongst respondents stating that they do not fuel *E10*, the most cited reason not to

[21] *Super Plus* is a niche product in a different product market. Outside high-performance sports vehicles, most vehicles do not receive any additional benefit from fueling *Super Plus*. At the same time, it is always significantly more expensive than *Super* and the price difference can be up to 15 Eurocent at the same fuel station and time. This is also why fuel stations do not have to report prices of *Super Plus* to the Market Transparency Unit in Germany.

[22] A full list of compatible vehicles can be found at `https://www.dat.de/e10/`.

[23] See `https://www.adac.de/verkehr/tanken-kraftstoff-antrieb/benzin-und-diesel/e10-tanken/`.

do so were technical concerns (51 percent among respondents not fueling *E10*), followed by uncertainty about the cost and benefits (23 percent).[24]

Overall, the evidence therefore strongly suggests that among drivers of gasoline cars, the more price sensitive drivers become informed and buy *E10*, whereas the less price sensitive drivers buy *E5*.

3.3.3 Taxes and input costs

The largest share of the fuel price consists of taxes. A lump-sum energy tax of 65.45 Eurocents per liter is levied on gasoline (47.04 Eurocents per liter for diesel).[25] In addition, there is a 19 percent value-added tax which is levied on the net price of diesel and gasoline, including the energy tax. This value-added tax was temporarily reduced to 16 percent between July and December 2020. For simplicity, we will refer to this event as the “tax decrease”.

On 1 January 2021, at the same time as the value-added tax was raised back to 19 percent, the German Federal Government also introduced a carbon price of 25 Euro per emitted tonne of CO_2 on oil, gas and fuel. For *E5* and *E10*, this translates into a per unit tax of 6 Eurocents per liter (7.14 Eurocents including VAT). For diesel, the per unit tax is 6.69 Eurocents per liter (7.96 Eurocents including VAT).[26] Likewise, we will refer to this event as the “tax increase”. Since the increase in the VAT and the introduction of the carbon emissions price happened simultaneously and affected the same stations, we cannot separately identify the pass-through rate of the two. Instead, we jointly estimate their pass-through rate. This does not raise concerns regarding the theoretical predictions, as we showed that the predictions on the determinants of pass-through are qualitatively the same for ad-valorem taxes and per unit taxes.

Crude oil accounts for another important share of the fuel price and is the most important source of price fluctuations. A barrel (42 gallons) of crude oil can be refined into around 19 gallons of gasoline, 12 gallons of diesel, as well as 13 gallons of other products, such as jet fuel, petroleum

[24] The full survey results can be found at `https://www.adac.de/news/umfrage-e10-tanken/`.
[25] An additional fuel storage fee of 0.27 Eurocents per liter is levied on gasoline and 0.30 Eurocents per liter on diesel.
[26] Further details can be found in the “Brennstoff-Emissionshandelsgesetz” (2020 Fuel Emissions Trading Act).

coke, bitumen or lubricants.[27] Gasoline and diesel are the most valuable components of refined crude oil.

3.4 Data and Descriptive Evidence

We now turn to our empirical analysis. We begin by describing our dataset and then present descriptive evidence on the differences in pass-through between fuel types.

3.4.1 Data

Our dataset contains all price changes for close to all fuel stations in Germany and France, as well as several characteristics of these stations.[28] In Germany, stations report price changes in real-time to the Market Transparency Unit at the German Federal Cartel Office. Tankerkönig, a price comparison website, provides access to this data, as well as to station characteristics, to researchers.[29] Similarly, price changes in France have to be reported by stations to a government agency, which makes this data available to researchers.[30] Furthermore, we add data on the daily price of crude oil, the principal input product for diesel and gasoline, at the port of Rotterdam. Finally, we use data on daily regional mobility patterns from the COVID-19 Community Mobility Report provided by Google.

Our analysis of the tax decrease starts on 1 May 2020 and goes until 31 August 2020. For the tax decrease, we analyze data between 1 November 2020 and 28 February 2021. In this section, we report descriptive statistics for the analysis of the tax increase in summer 2020. We report the same descriptive statistics for the tax decrease in winter 2020/21 in Appendix C.2.

Using the data on price changes, we construct daily weighted average prices. Table 3.1 shows the summary statistics for the analysis of the tax reduction. The price level is generally higher in France than in Germany. Gross prices in France increase by around 5 to 6 Eurocent between the

[27] These are approximate shares which can vary by context and type of crude oil. The total volume of products refineries produce (output) is greater than the volume of crude oil that refineries process (input) because most of the products they make have a lower density than the crude oil they process. See `https://www.eia.gov/energyexplained/oil-and-petroleum-products/refining-crude-oil-inputs-and-outputs.php`.

[28] In France, fuel stations selling less than $500m^3$ per year are exempt from reporting price changes.

[29] See `https://creativecommons.tankerkoenig.de/`.

[30] See `https://www.prix-carburants.gouv.fr/rubrique/opendata/`.

Table 3.1: Summary statistics

	Germany pre-VAT cut	Germany post-VAT cut	France pre-VAT cut	France post-VAT cut
A. Station characteristics				
Number of stations	14,627	14,612	8,960	8,975
Median comp. nr. (5km markets)	4	4	2	2
Share of local monopolists	13%	13%	20%	19%
B. Prices, *E5*				
Mean price	1.21	1.27	1.30	1.36
Mean price net of taxes and duties	.36	.44	.40	.44
Mean retail margin	.13	.16	.17	.16
C. Prices, *E10*				
Mean price	1.18	1.23	1.27	1.32
Mean price net of taxes and duties	.34	.40	.39	.43
Mean retail margin	.11	.13	.16	.15
D. Prices, diesel				
Mean price	1.05	1.07	1.20	1.25
Mean price net of taxes and duties	.41	.45	.39	.43
Mean retail margin	.18	.17	.16	.15
E. Mobility data				
Retail & recreation	-22.2%	-2.4%	-32.4%	6.6%
Workplaces	-21.9%	-20.7%	-27.8%	-26.2%

Notes: "pre-VAT cut" and "post-VAT cut" refer to fuel stations in Germany and France before and after the reduction of the VAT rate, respectively. The pre-VAT phase goes from 1 May until 31 June 2020. The post-VAT phase starts on 1 July 2020.

pre- and post-tax cut periods. In Germany, gross prices increase by about 2 Eurocent for diesel and 5 to 6 Eurocent for *E5* and *E10*. At the same time, the increase in the net price in Germany is between 4 and 8 Eurocent, depending on the fuel type, which is larger than in France, and suggests that the tax reduction was not completely passed on to consumers.

We also calculate retail margins by subtracting taxes, duties and the share of the price of crude oil that goes into the production of diesel and gasoline, respectively.[31] Although these retail margins still contain different cost types, such as the cost of refining or transportation costs, the main source of input cost variation, the price of crude oil, is eliminated. Table 3.1 shows that retail margins declined by about 1 Eurocent for France after the tax reduction. Although there is a modest decrease in retail margins for diesel in Germany after the tax reduction, there is an increase in the retail margin of around 17.6 percent for *E10* and 20.4 percent for *E5*.[32]

[31]For a detailed description of the calculation of prices and margins, see Appendix C.2.

[32]Percentage changes are different from what you would calculate from the retail margins in the table because of rounding of margins in the table.

To capture regional changes in demand over time, we use the daily percentage change in visits to retail and recreation, as well as to the workplace, from the COVID-19 Community Mobility Report. With the former, we intend to capture local changes in the propensity of using a car for leisurely activities, including going on vacation. With the latter, we aim to capture local changes in the propensity to use a car for professional activities. Both of these variables are measured as the percentage change of activities compared to the median value for the corresponding day of the week during the five-week period 3 January to 6 February 2020. The data is disaggregated for 96 sub-regions in France and 16 regions in Germany. We use the geolocation of each fuel station to match the measures of local mobility to each station.

Table 3.1 shows that mobility patterns in France and Germany are similar. Whereas visits to retail and recreational facilities were around 22 to 32 percent lower in May to June compared to the baseline beginning of the year, in July to August, the number of such visits returned close to their pre-pandemic levels. At the same time, in both countries visits to workplaces were around 20 to 28 percent lower in May to August compared to the baseline.

Our dataset also contains a number of station characteristics, such as the exact geolocation, and, for Germany, the brand of a station. We use this data to measure the number of firms active in a local market. We define each market as a catchment area around a focal fuel station. We exploit the geolocation of each station to calculate the driving distance between stations using the road network.[33] Finally, we count the number of rival stations that are within a 3, 5 or 10 km catchment area around a focal station. Based on our market definition, we can also compute the share of stations that are without any competitor in their local market, i.e. the share of local monopolists. Table 3.1 shows that the median number of competing fuel stations within a 5 km catchment area is 4 in Germany and 2 in France. 13 percent of stations in Germany are local monopolists within a 5 km catchment area, compared to 19 to 20 percent in France.

We report summary statistics using the weights in the SDID in Appendix C.2. Results on average fuel prices, retail margins and stations characteristics remain analogous when stations in France are weighted by the SDID weights.

[33]By using the road network, we avoid classifying fuel stations that are close by air distance but not by road as competing with each other.

3.4.2 Descriptive evidence on heterogeneous pass-through

Before econometrically estimating pass-through of the tax changes on prices and retail margins, we study the pass-through of the policy changes descriptively. We can thereby gain first insights into whether pass-through differs between markets with very price sensitive consumers (diesel) and markets with less price sensitive consumers (*E5*). Let us begin by first looking at the VAT reduction on 1 July 2020.

Figure 3.7: Tax decrease: Price change as share of total tax change

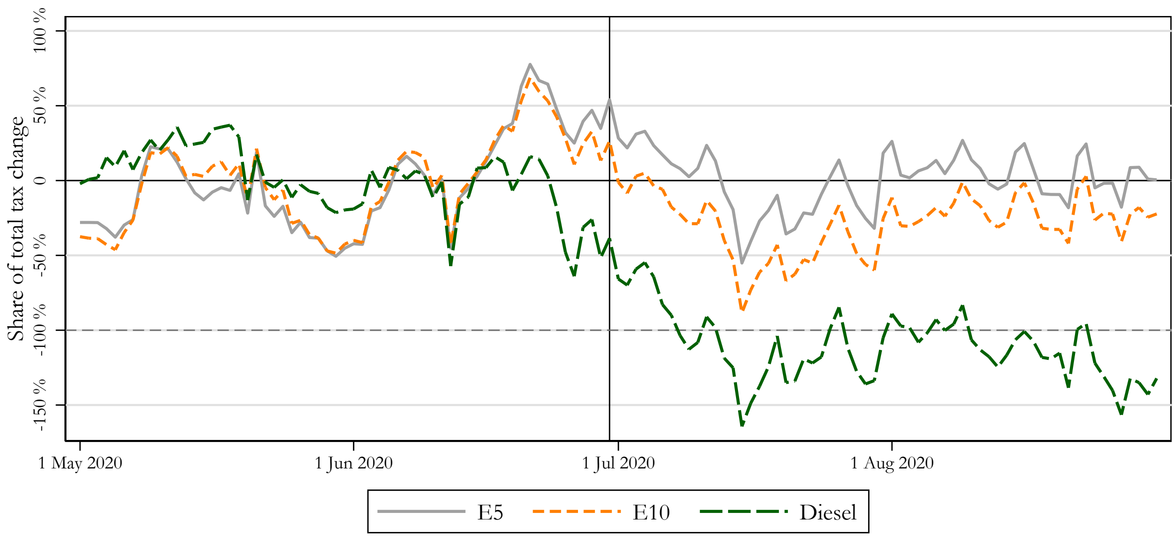

Notes: The solid line shows the nonparametric estimate of the daily average pass-through rate to prices for *E5*. The short-dashed and long-dashed lines show analogous estimates for *E10* and diesel, respectively. To estimate pass-through, we first subtract the average pre-period (1 May until 30 June 2020) price in Germany (France) from the daily average price in Germany (France). Next, we compute the difference between demeaned average prices in Germany and France. Finally, we divide this difference by 3 Eurocents for *E5* and *E10* and by 2.7 Eurocents for diesel, which would be the difference under full pass-through. The vertical solid line marks the starting date of the tax decrease. The horizontal dashed line indicates the full pass-through.

Figure 3.7 shows nonparametric estimates of the pass-through rate of the tax decrease to fuel prices. As we would expect, prior to the tax reduction, there is no pass-through of the tax decrease for any fuel type, as it has not yet occurred. The evolution of fuel prices evolves similarly for the three fuel types, which suggests that differences in pass-through rates after the tax decrease are not driven by pre-trends. The evolution of prices after the tax decrease suggests that pass-through was relatively fast, stabilized after around two weeks, and that it was highest for diesel and lowest for *E5*. The difference in pass-through between fuel types is in line with our theoretical prediction that pass-through increases if there are more price sensitive consumers in the market.

Although we can see that there are differences in the evolution of prices between France and Germany in the pre-period, these appear to be idiosyncratic. The findings described above can clearly be seen even before correcting for some of the idiosyncratic shocks. However, the absolute magnitudes of pass-through in this graph should be treated with caution and we provide more precise estimates of these in the following sections.

Figure 3.8: Tax decrease: Margin change as share of total tax change

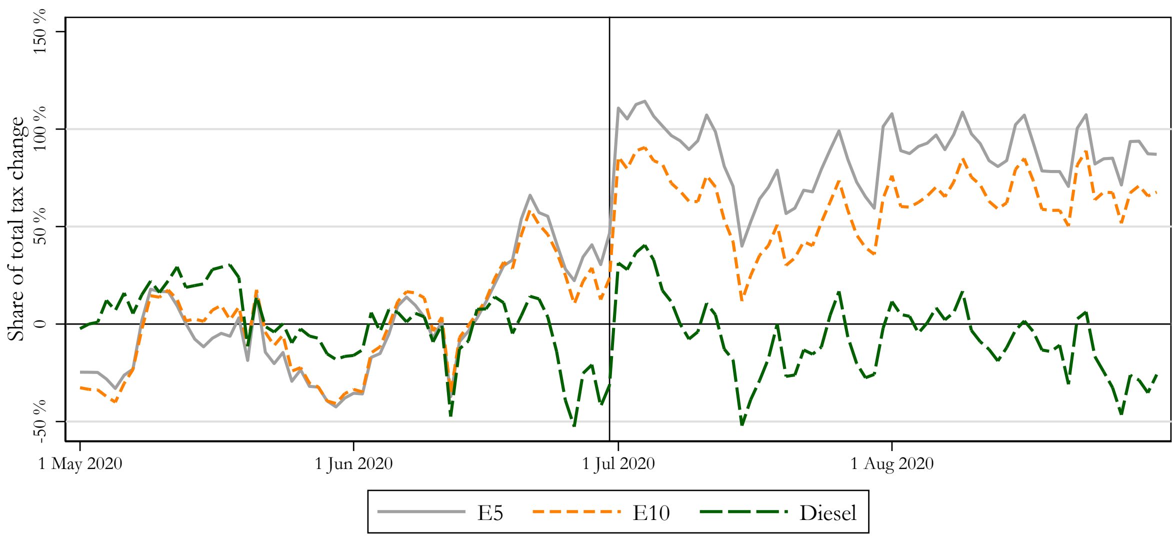

Notes: The solid line shows the nonparametric estimate of the daily average pass-through rate to retail margins for *E5*. The short-dashed and long-dashed lines show analogous estimates for *E10* and diesel, respectively. To estimate pass-through, we first subtract the average pre-period (1 May until 30 June 2020) retail margin in Germany (France) from the daily average retail margin in Germany (France). Next, we compute the difference between demeaned average retail margins in Germany and France. Finally, we divide this difference by 3 Eurocents for *E5* and *E10* and by 2.7 Eurocents for diesel, which would be the difference under full pass-through. The vertical solid line marks the starting date of the tax decrease.

Figure 3.8 plots the analogous graph for retail margins. Consistent with what we saw for prices, there is no pass-through of the tax decrease to retail margins prior to the tax decrease. In the post-period, retail margins appear to increase the most for *E5* and remain unchanged for diesel.

In Figure 3.9, we present nonparametric estimates of the pass-through rate by fuel type for the tax increase in winter 2020/21. As for the tax decrease, there is no anticipatory pass-through of the tax increase for most of the pre-increase period. In contrast to the tax decrease, there seem to be anticipatory effects in passing through the tax increases in the last two weeks of December. In our econometric analysis, we therefore drop the second half of December 2020, since this

Figure 3.9: Tax increase: Price change as share of total tax change

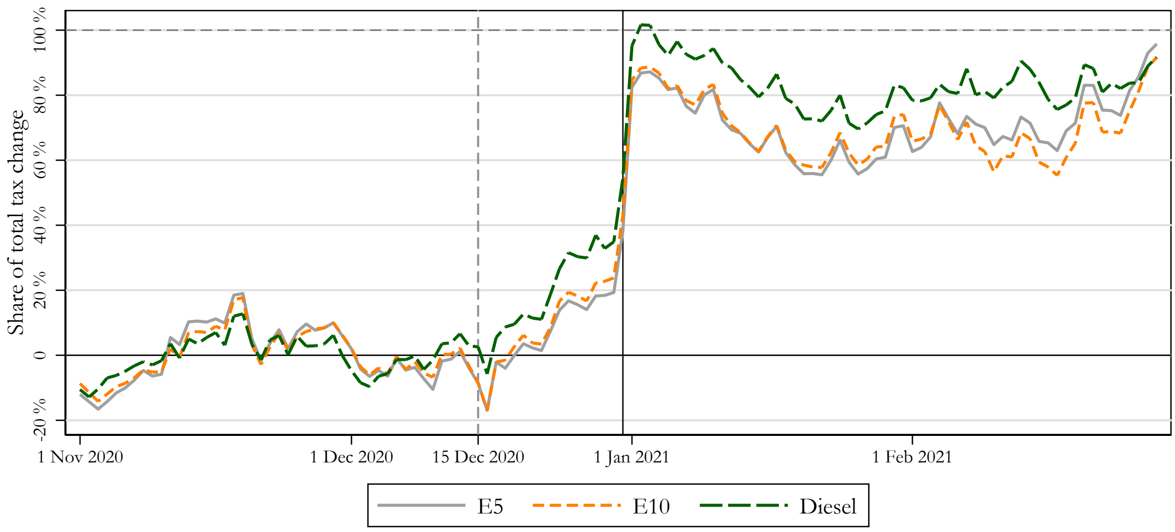

Notes: The solid line shows the nonparametric estimate of the daily average pass-through rate to prices for *E5*. The short-dashed and long-dashed lines show analogous estimates for *E10* and diesel, respectively. To estimate pass-through, we first subtract the average pre-period (1 November until 15 December 2020) price in Germany (France) from the daily average price in Germany (France). Next, we compute the difference between demeaned average prices in Germany and France. Finally, we divide this difference by 10 Eurocents for *E5* and *E10* and by 11 Eurocents for diesel, which would be the difference under full pass-through. The vertical solid line marks the starting date of the VAT increase and carbon emissions price in Germany. The horizontal dashed line indicates the full pass-through.

already appears to be partially treated. Finally, there is a sharp increase in the implied pass-through rate around 1 January 2021, after which this stays stable. Differences in pass-through between diesel and other types of fuel are very pronounced. As in summer 2020, pass-through appears to be highest for diesel. This is also consistent with our theoretical predictions. From the descriptive evidence, differences in pass-through between *E5* and *E10* seem less strong. We provide more precise estimates on this in the upcoming sections.

Figure 3.10: Tax increase: Margin change as share of total tax change

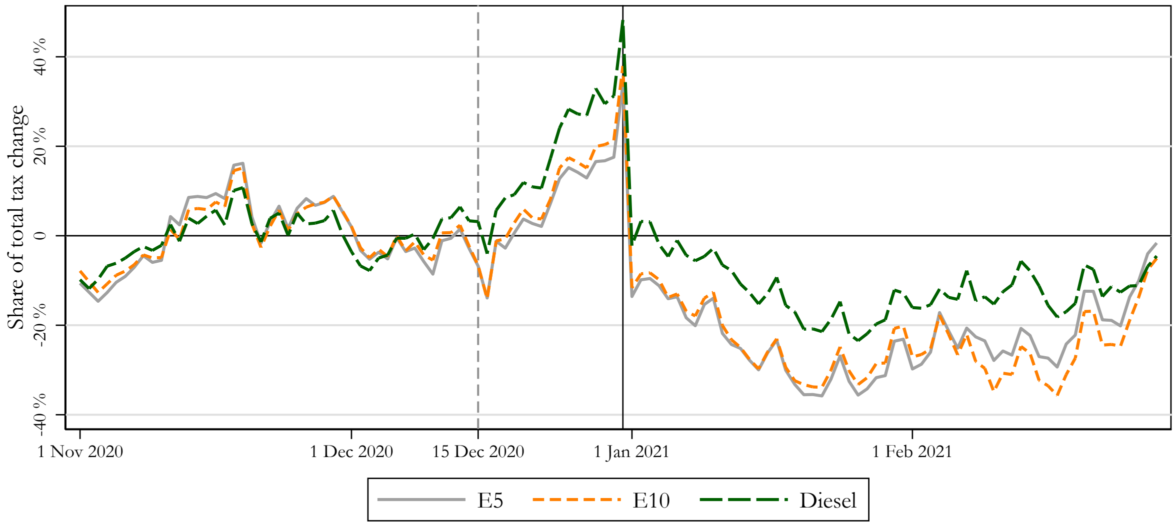

Notes: The solid line shows the nonparametric estimate of the daily average pass-through rate to retail margins for *E5*. The short-dashed and long-dashed lines show analogous estimates for *E10* and diesel, respectively. To estimate pass-through, we first subtract the average pre-period (1 November until 15 December 2020) retail margin in Germany (France) from the daily average retail margin in Germany (France). Next, we compute the difference between demeaned average retail margins in Germany and France. Finally, we divide this difference by 10 Eurocents for *E5* and *E10* and by 11 Eurocents for diesel, which would be the difference under full pass-through. The vertical solid line marks the starting date of the VAT increase and carbon emissions price in Germany.

Figure 3.10 shows how the tax increase is passed through to retail margins. Since stations begin increasing prices already in the second half of December 2020, even though the tax increase only occurred on 1 January 2021, there appears to be an increase in retail margins worth up to 30 percent of the subsequent tax change for diesel in the last week of December 2020 and around 20 percent for *E5* and *E10*. After the tax increase, the descriptive evidence suggests that the decrease in retail margins was lowest for diesel. This is consistent with what we see for prices.

The results in Figures 3.9 and 3.10 suggest that in the second half of December 2020, there are some anticipatory effects of the tax increase coming into effect on 1 January 2021 across all fuel types. A visual analysis of Figures 3.7 and 3.8 suggests that there could be anticipatory effects for *E5* and *E10* already in the second half of June 2020, but that these are less pronounced than in winter. Our preferred specification is therefore to account for anticipatory effects in winter but not in summer. In Appendix C.4, we show that our main empirical findings are robust to changing these assumptions. In Appendix C.1, we briefly discuss theoretically why anticipatory price increases could arise before a tax increase and a tax decrease.

3.5 Empirical Strategy

So far, we saw descriptively that pass-through of the two tax changes appears to be different across fuel types. At the same time, we saw that there were some idiosyncratic differences in the evolution of fuel prices between Germany and France. To cut through the noise and estimate pass-through rates, we use a synthetic difference-in-differences (SDID) strategy.

3.5.1 Synthetic difference-in-differences

The general idea of SDID is quite simple. As with difference-in-differences, we use fuel prices at French stations as the control group and so the treatment effect is the change in the difference between average fuel prices at fuel stations in Germany and France between pre- and post-treatment periods. In contrast to DID, weights of fuel stations in the control group, as well as weights of the pre-treatment periods are chosen as to match the pre-treatment trends in the treatment group.[34] In this sense it is similar to synthetic control methods. Arkhangelsky et al. (2021) report that SDID performs weakly better than DID and SC methods.

The estimation proceeds in two steps. In the first step, we compute the unit and time weights that minimize the difference in pre-treatment trends between the treated and control units and the difference in outcomes between pre- and post-treatment periods for the unexposed units. In

[34]On average, fuel prices are higher at stations in France than in Germany. Since SDID matches the pre-treatment trends in prices instead of the price level, as shown in Appendix C.4 control stations that receive a higher SDID unit weight are not clustered in a particular region in France.

the second step, we estimate a difference-in-differences model using the unit and time weights from the first step. We estimate standard errors using the jackknife method.[35]

To estimate the average pass-through rate of the tax changes on fuel prices, we compare stations in Germany and France, before and after the tax change. In particular, we solve the following minimization problem:

$$(\hat{\tau}^{sdid}, \hat{\mu}, \hat{\alpha}, \hat{\beta}) = \underset{\tau,\mu,\alpha,\beta}{\arg\min} \left\{ \sum_{i=1}^{N} \sum_{t=1}^{T} (Y_{it} - \mu - \alpha_i - \beta_t - Tax_{it}\tau)^2 \, \hat{w}_i^{sdid} \hat{\lambda}_t^{sdid} \right\} \tag{3.2}$$

where $\hat{\tau}^{sdid}$ is the estimated effect of the policy change, and $\hat{w}_i^{sdid}$ and $\hat{\lambda}_t^{sdid}$ are the SDID unit and time weights, respectively.[36] Y_{it} is the logarithm of the price of gasoline or diesel at fuel station i at date t, and Tax_{it} is a dummy variable that equals one for stations affected by the tax change at date t. For the analysis of the tax reduction, these are fuel stations in Germany from 1 July 2020 onwards. For the analysis of the subsequent tax increase, these are fuel stations in Germany from 1 January 2021 onwards. The variables α_i and β_t correspond to fuel station and date fixed effects, respectively.

To use the synthetic difference-in-differences method, we require a balanced sample. We therefore restrict our sample to fuel stations in France and Germany for which we have a price observation on every day in our sample. This is the case for 83 percent of fuel stations in Germany and 62 percent in France for the analysis of tax reduction, and for 83 percent of stations in Germany and 74 percent in France for the analysis of the tax increase. In Appendix C.4, we also estimate a DID model using the full, unbalanced sample.

Finally, we also want to assess the speed at which the tax changes are passed-through to consumers and verify that the parallel trends assumption holds. We therefore estimate time-varying effects of the tax changes using the following model:

$$ln(p_{it}) = \sum_{j=-k}^{8} \beta_j Tax_{it} + \mu_i + \gamma_t + \epsilon_{it} \tag{3.3}$$

[35] We use the jackknife method instead of bootstrapping, as the latter is computationally too intensive in our case. The jackknife method is a linear approximation of the bootstrap and gives a conservative estimate of the variance when the panel is large and the number of treated units is high.

[36] We estimate the model using the *synthdid* package by Arkhangelsky et al. (2021). A more detailed description of the algorithm can be found in Appendix C.3.

where $ln(p_{it})$ is the logarithm of the price of gasoline or diesel at fuel station i at date t. The regression is weighted by the SDID unit and time weights, and we control for fuel station and date fixed effects. The coefficient β_j captures the effect of the tax change in a period t on fuel prices in Germany in a week $t + j$, with $j \in [-k, 8]$.[37]

3.5.2 French fuel stations as a control group

To identify the effect of the tax change on fuel prices, two main assumptions must be satisfied. First, there should be no transitory shocks that would differentially affect fuel stations in Germany and France before and after the change in tax, other than the policy change itself. Second, there should be no spillover effects from the tax decrease or the tax increase in Germany onto the fuel market in France.

Station fixed effects control for any time-invariant differences between fuel stations in France and Germany, and date fixed effects capture the transitory shocks, such as fluctuations in the price of crude oil, that identically affect French and German stations. The two countries are similar in their geographic location, size, and wealth. Since in our analysis we also focus on relatively narrow windows around the reforms, this should alleviate concerns on transitory shocks differentially affecting French and German fuel stations.

To further strengthen our claim that the effects are not confounded by certain transitory shocks, we now discuss the most obvious candidates. On the demand side, public and school holidays in France and Germany are highly correlated. Travel restrictions put in place due to COVID-19 were lifted simultaneously in the two countries. Starting from 15 June 2020, residents of the Schengen Area and the United Kingdom could freely cross the territories of France and Germany again. Most holidaymakers within Europe typically travel across several countries in the EU, and as France and Germany are both popular travel destinations in close geographic proximity, demand shocks likely hit fuel stations in the two countries in a similar way.[38]

Transitory supply shocks should affect French and German fuel stations in a similar way. Due to their geographic proximity, the fuel stations in France and Germany procure most of their crude

[37] For the analysis of the tax reduction, $k = 7$. For the analysis of the tax increase, $k = 5$.

[38] In addition, we directly account for demand-related shocks by including regional information on the daily mobility to work and to retail and recreational places as control variables into our empirical specification. The results are reported in the Appendix.

oil from similar sources.[39] The two countries are also members of the European Single Market, which implies harmonized border checks, common customs policy, and identical regulatory procedures on the movement of goods within the EU.

No major reforms were implemented in France during our analysis period. In general, there are no fuel price-setting regulations in Germany and France, and both countries have mandatory disclosure of fuel prices, which reaffirms our choice of France as a suitable control group.

Furthermore, the SDID algorithm allows us to place higher weight on French fuel stations whose pre-trends are very similar to the pre-trends of stations in Germany and place lower weight on French stations whose pre-trends are very dissimilar. This should further alleviate any remaining concerns about French stations as a control group.

Finally, our analysis of the two episodes of a change in tax, the temporary VAT rate reduction in July 2020 and the subsequent increase in the VAT rate with simultaneous introduction of a carbon emissions price in January 2021, alleviates a concern that some confounding factor could drive the results. If we find similar heterogeneities in pass-through for the VAT increase in January 2021 as for the VAT decrease in July 2020, a transitory shock confounding our estimates in July 2020 would also have to be present in January 2021 and at that point work in the opposite direction. To illustrate this point: if we thought that we overestimate the pass-through rate for diesel in July 2020, because France is hit by a positive transitory demand in July 2020, which does not affect Germany, then also overestimating pass-through for diesel in January 2021 would now require France to be hit by a negative demand shock in January 2021, which does not affect Germany. Overall, this seems implausible. Finding consistent heterogeneities in pass-through rates between the July 2020 and January 2021 tax changes therefore suggests that we are robustly estimating actual differences in pass-through.

3.6 Results

In Section 3.2, we showed theoretically how the pass-through of a tax depends on the price sensitivity of consumers and the number of sellers. Descriptively, we showed that the hetero-

[39] We additionally account for potentially differential pass-through of oil cost shocks to fuel prices by allowing crude oil price affect fuel prices differently depending on the country. The results are reported in the Appendix.

geneities in the pass-through rate between fuel types are in line with our theoretical predictions. In this section, we provide further evidence on this and also study how pass-through depends on the number of sellers empirically.

3.6.1 Price sensitivity and tax pass-through

We first study how the pass-through of a tax varies with the price sensitivity of consumers. Theoretically, we showed that the higher the price sensitivity of consumers, the higher will be the pass-through rate of a tax. To test this prediction empirically, we estimate the effects of the tax changes in Germany on *E5*, *E10* and diesel prices, and compare the estimated pass-through rates across fuel types.

We begin our analysis of the tax changes by plotting their time-varying effects by fuel type.

Figure 3.11: Dynamic effect of the tax decrease on log fuel prices

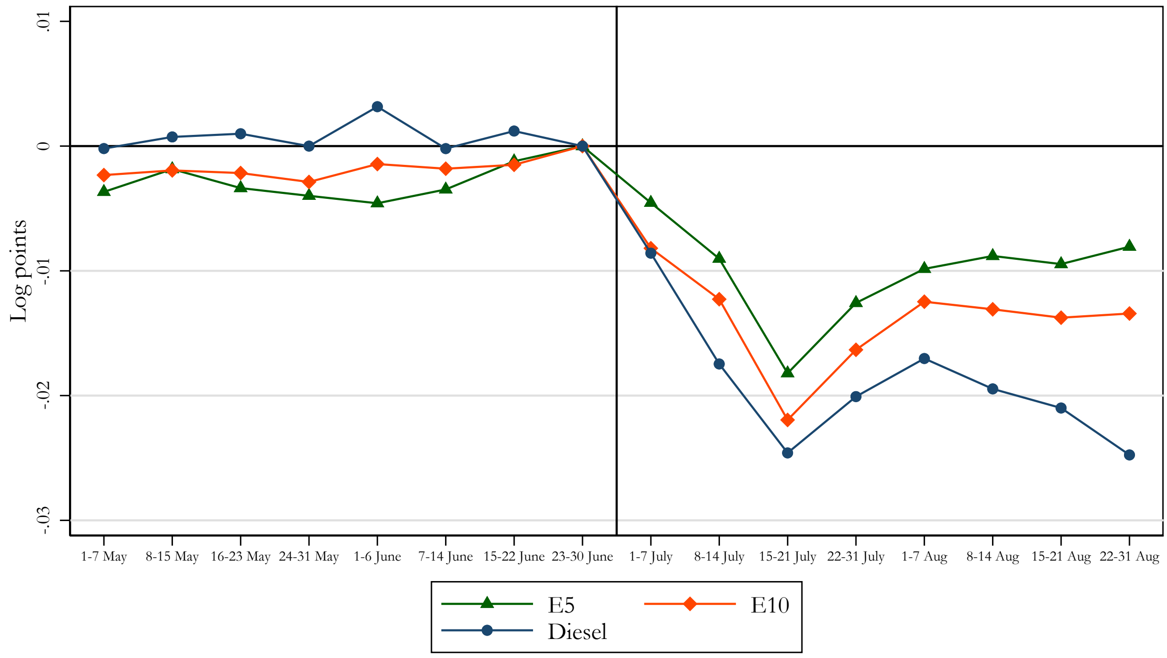

Notes: The graph shows the time-varying effect of the tax decrease on the log prices for *E5*, *E10* and diesel. The analysis period goes from 1 May until 31 August 2020. For the time-varying treatment effects, we estimate the model in Equation 3.3, weighted by the SDID unit and time weights. The vertical line marks the starting date of the tax decrease in Germany.

Figure 3.11 shows the time-varying effect of the tax decrease on the logarithm of prices for *E5*, *E10* and diesel.[40] The estimation is based on 1 May to 31 August 2020. The vertical line marks the beginning of the tax decrease in Germany. Prior to the tax reduction, the trends in log fuel prices are similar between France and Germany. After the tax reduction, log prices of all fuel types decline at fuel stations in Germany compared to fuel stations in France. The effect of the tax reduction is highest for diesel and lowest for *E5*, and is relatively fast. These results are consistent with the descriptive evidence and the theoretical predictions.

Figure 3.12: Dynamic effect of the tax increase on log fuel prices

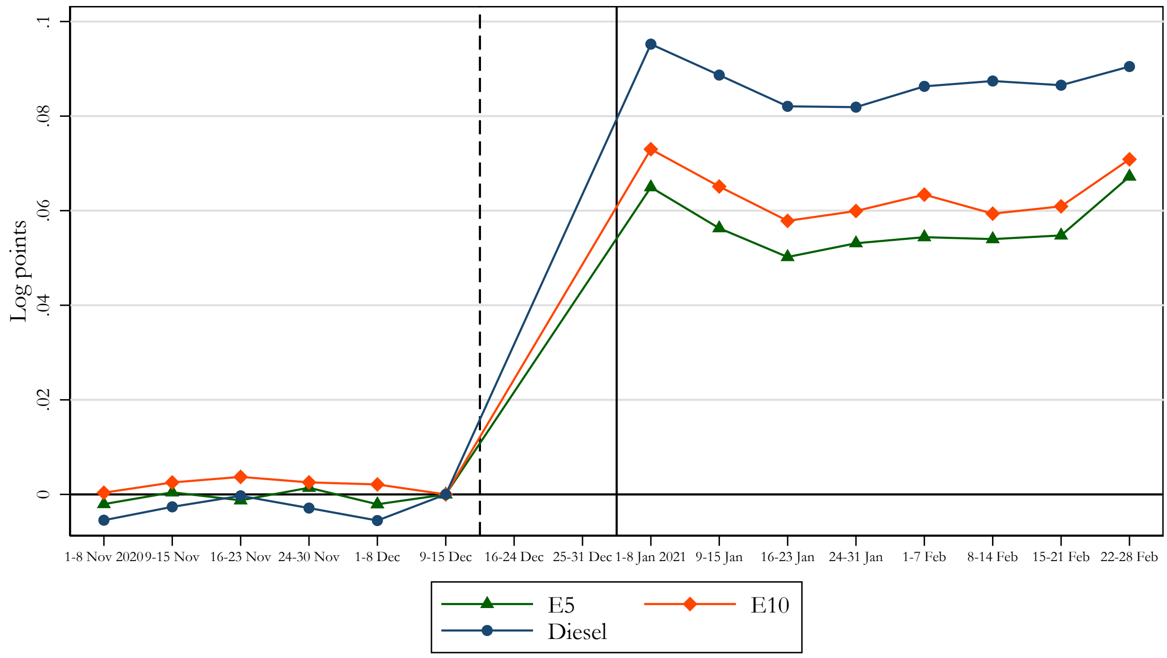

Notes: The graph shows the time-varying effect of the tax increase on log prices for *E5*, *E10* and diesel. The pre-treatment period goes from 1 November until 15 December 2020 and the post-treatment period from 1 January to 28 February 2021. For the time-varying treatment effects, we estimate the model in Equation 3.3, weighted by the SDID unit and time weights. The vertical solid line marks the beginning of the tax increase in Germany.

Figure 3.12 shows the time-varying effect of the tax increase. The analysis is based on the pre-treatment period of 1 November to 15 December 2020 and the post-treatment period of 1 January to 28 February 2021. As we saw in the descriptive analysis, there are anticipatory effects of the tax increase in the second half of December 2020. Since these days appear to be already partially treated, we drop them from the analysis.

[40]Figures with the time-varying effects on retail margins are reported in Appendix C.4.

Table 3.2: Effect of the tax change on log prices (percent)

	E5	E10	Diesel	E5	E10	Diesel
	(1)	(2)	(3)	(4)	(5)	(6)
Tax change	-.0085***	-.0130***	-.0199***	.0565***	.0627***	.0889***
	(.0013)	(.0013)	(.0015)	(.0015)	(.0019)	(.0020)
Pass-through rate	34%	52%	79%	69%	75%	92%
	[24%, 43%]	[42%, 62%]	[67%, 91%]	[66%, 73%]	[71%, 79%]	[88%, 96%]
Date fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Station fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,736,145	1,968,984	2,176,362	1,485,120	1,712,984	1,945,736

Notes: Columns (1) to (3) present average treatment effect estimates of the VAT reduction on E5, E10, and diesel log prices, respectively. Columns (1) to (3) use data from 1 May to 31 August 2020. Columns (4) to (6) present average treatment effect estimates of the VAT increase and CO_2 emissions tax on E5, E10, and diesel log prices, respectively. Columns (4) to (6) use data from 1 November to 15 December 2020 for pre-treatment period, and from 1 January to 28 February 2021 for post-treatment period. 95% confidence intervals on pass-through rates are reported in parentheses. Standard errors are computed using the jackknife method and are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Prior to the tax increase, the trends in the logarithm of fuel prices are similar between France and Germany. After the tax increase, log prices at fuel stations in Germany increase compared to those in France for all fuel types. Pass-through of the tax increase is almost immediate. Fuel prices increase by about 6 to 9 percent in the first week of January 2021 compared to the week ending on 15 December 2020. Similarly to our results for the tax reduction, the price increase is highest for diesel and lowest for *E5*, with *E10* in between.[41]

Next, we estimate the average treatment effect of the tax changes on the logarithm of prices for *E5*, *E10* and diesel. Table 3.2 shows the results of estimating the SDID model described in Equation 3.2. The outcome variable in all columns is the logarithm of price for each fuel type, including taxes and duties. Columns (1) to (3) correspond to the effect of the tax decrease. Columns (4) to (6) correspond to the effect of the subsequent tax increase. In all columns, we control for fuel station and date fixed effects.[42]

[41] Note, that relative pass-through rates cannot directly be inferred from Figure 3.12, as the percentage increase in prices for full pass-through is different between fuel types. We estimate pass-through rates in Table 3.2.

[42] In Appendix C.4, we show the geographic distribution of stations that receive a higher than average SDID unit weight in France for the case of the tax decrease and tax increase. Control stations with disproportionately higher SDID weights are scattered throughout France and do not appear to cluster in a particular region.

The results in Columns (1) to (3) of Table 3.2 show that the tax decrease led to a decline in prices of all fuel products. The average price for *E5* decreases by 0.85 percent after the tax reduction, whilst average prices for *E10* and diesel decrease by 1.3 and 1.99 percent, respectively.[43]

To estimate pass-through of the tax reduction, we start by considering the case of full pass-through. Under full pass-through, we expect prices for each fuel product to decrease by about 2.52 percent.[44] An estimated decline of 1.99 percent in diesel prices is therefore relatively close to full pass-through. Around 79 percent of the tax decrease is passed on to consumers who refuel with diesel. For *E10*, the pass-through rate is 52 percent. Finally, we estimate that 34 percent of the tax decrease is passed through to consumers of *E5*. For all fuel products, pass-through of the tax reduction is fast and relatively high, but incomplete.

The results in Columns (4) to (6) of Table 3.2 show the effect of the subsequent VAT rate increase and the introduction of a carbon price on log fuel prices. Since the increase in the VAT and the introduction of the carbon emissions price happened simultaneously and affected the same stations, we cannot separately identify the pass-through rate of the two. Instead, we jointly estimate their pass-through rate. This does not raise concerns regarding the theoretical predictions, as we showed that the predictions on the determinants of pass-through are qualitatively the same for ad-valorem taxes and per unit taxes.[45]

Columns (4) to (6) of Table 3.2 show that the tax increase led to an increase in prices of all fuel products. The average price of *E5* increases by about 5.65 percent, whereas *E10* and diesel prices increase by about 6.27 and 8.89 percent after the change in the VAT rate and carbon tax introduction, respectively.

Next, we estimate the pass-through rate of the tax increase. Under full pass-through, we would expect an increase in prices by 8.15 percent for *E5*, 8.37 percent for *E10* and 9.66 percent for diesel.[46] We find a joint pass-through rate of the tax increase of 69 percent for *E5*, 75 percent

[43]In Appendix C.4, we report the results when we additionally control for regional mobility for retail and recreational purposes and to workplaces, and allow the changes in the crude oil price to differentially affect fuel prices in France and Germany. Our results are robust to the inclusion of these additional controls.

[44]With a decrease in the VAT rate from 19 percent before the VAT decrease to 16 percent after the VAT decrease, this is $\frac{1.16-1.19}{1.19} * 100 \approx -2.52\%$.

[45]The only necessary adjustment is that we need to translate the per unit tax on carbon emissions into a percentage value, such that we can calculate how large the percentage increase in prices would be if the VAT rate and the carbon emissions tax were fully passed through.

[46]Under full pass-through, a change in the VAT rate from 16 to 19 percent would increase the fuel price by $\frac{1.19-1.16}{1.16} * 100 \approx 2.59$ percent. To estimate by what percentage the fuel price would increase if the carbon emissions

for *E10* and 92 percent for diesel. As for the tax decrease, pass-through is fast but incomplete and it is lowest for fuel types with fewer price sensitive consumers and higher for fuel types with more price sensitive consumers. In Appendix C.4, we report results for the tax decrease and tax increase using DID model. The ranking of pass-through rates across different fuel types and their magnitude remain robust to using this alternative specification.

As predicted by the theory, we find that the pass-through rate for diesel is highest and it is the lowest for *E5*. An advantage of our setting is that all fuel stations in Germany are required by law to sell all three types of fuel and so differences in the pass-through rates cannot be explained by supply-side factors, such as fuel station characteristics. Table 3.2 reports the 95 percent confidence interval on pass-through rates for the different fuel types. For both the tax decrease and subsequent tax increase, we can see that the difference between the pass-through rate for diesel and the two types of gasoline is statistically significant at the 5 percent level. Confidence intervals for the pass-through rate of *E5* and *E10* overlap, however, their ranges still strongly suggest that there is an important economic difference between the pass-through rates for *E5* and *E10*. Overall, our empirical results confirm the predictions in Proposition 3.2.

Across fuel types, the pass-through rate of the increase is above the pass-through rate of the decrease. Although this is not the focus of our study, these results are consistent with recent findings on asymmetric VAT pass-through by Benzarti, Carloni, et al. (2020).

Based on the descriptive price plots in Section 3.4, our preferred specification and the presented results so far correspond to accounting for anticipatory effect in winter 2020/21 but not in summer 2020. In Appendix C.4, we report results when we instead account for anticipatory effects in summer but not in winter. Even though pass-through estimates change when we use this alternative specification, the relationship between tax pass-through and price sensitivity is robust with respect to anticipatory effects. The pass-through remains highest for diesel and lowest for *E5*.

price was fully passed through, we divide the gross per liter price on carbon emissions for each fuel type by the average fuel price in Germany in the last week of 2020.

3.6.2 Number of sellers and tax pass-through

Finally, we study how the pass-through rate varies with the number of sellers in the market. In Section 3.2, we used simulations to show that theoretically there is a hump-shaped relationship between the number of sellers in the market and tax pass-through.

To verify this empirically, we study differences in the pass-through rate of the tax decrease across fuel stations with different numbers of competitors in their market. An important feature of our setting is that we can do this comparison within fuel type and so hold an important source of variation in price sensitivity fixed. We begin by estimating a pass-through rate for every station in Germany for each fuel type. For each station in Germany and fuel type, we estimate the model in Equation 3.2 adding an interaction term between the treatment period and the station's fixed effect.[47] The station-specific treatment effect is then the sum of the average treatment effect and this additional interaction. Finally, we group stations by the number of competitors in their market and calculate the average pass-through rate for each group.[48]

Figure 3.13 shows the relationship between the pass-through rate and the number of competitors of a focal station for *E5*. Each circle corresponds to the average pass-through rate for stations with a particular number of competitors within 5 km catchment area.[49] The size of a circle is proportional to the total number of stations with a given number of competitors. Figure 3.13 shows that the average pass-through is relatively low for local monopolists, and increases in the number of rivals, up to around six competitor stations. With more than six competitor stations, the average pass-through declines in the number of competitors.

Figure 3.14 shows the relationship between the pass-through rate and the number of competitors of a focal station for *E10*. Similar to *E5*, we observe a hump-shaped relationship between the pass-through rate and the number of competitors. The average pass-through rate is relatively low for local monopolists, peaks in the group of stations that have around six to eight competitors and then falls again in the number of competitors.

Figure 3.15 shows the relationship between the pass-through rate for diesel and the number of competitors of a focal station. In contrast to what we see for *E5* and *E10*, the relationship

[47] We use the same time and unit weights for each station-specific treatment effect and estimate this only once.

[48] In Appendix C.4, we show the analogous relationship between the pass-through rate of the tax increase and the number of competitors of a focal station.

[49] The pattern is similar for alternative radii.

Figure 3.13: Average pass-through by number of competitor stations, E5

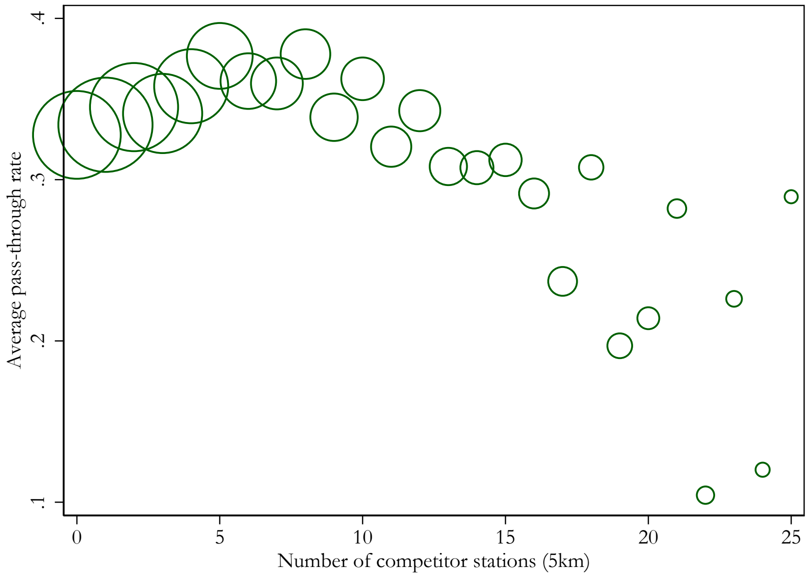

Notes: Each circle plots the average pass-through rate for a group of stations with a particular number of competitors within a 5 km catchment area. The number of competitor stations is trimmed at the top percentile.

Figure 3.14: Average pass-through by number of competitor stations, E10

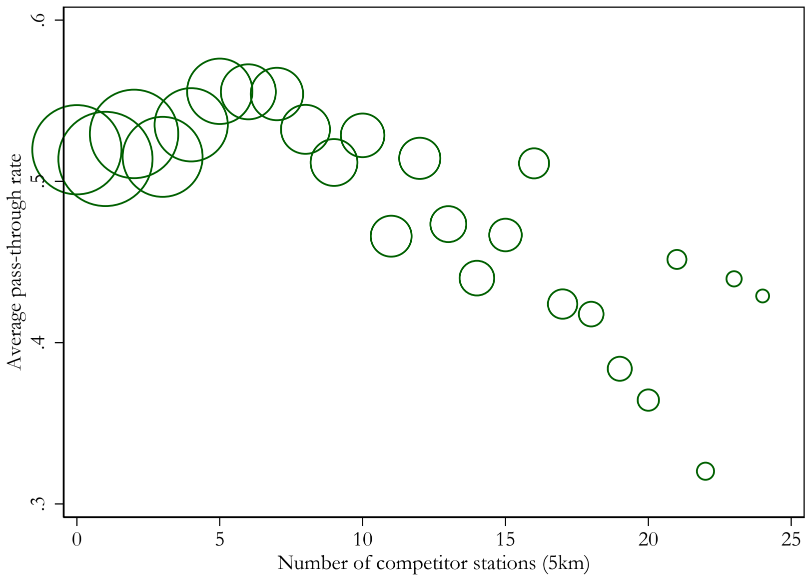

Notes: Each circle plots the average pass-through rate for a group of stations with a particular number of competitors within a 5 km catchment area. The number of competitor stations is trimmed at the top percentile.

between the pass-through rate and the number of competitors is mostly flat and, in parts, even increasing.

Since the theoretical model predicts a hump-shaped relationship between the number of sellers and the pass-through rate, one possibility could be that for diesel we only observe the upward-sloping part of the hump. Another possibility could be that the hump-shaped relationship becomes weaker for higher pass-through rates.

Figure 3.15: Average pass-through by number of competitor stations, diesel

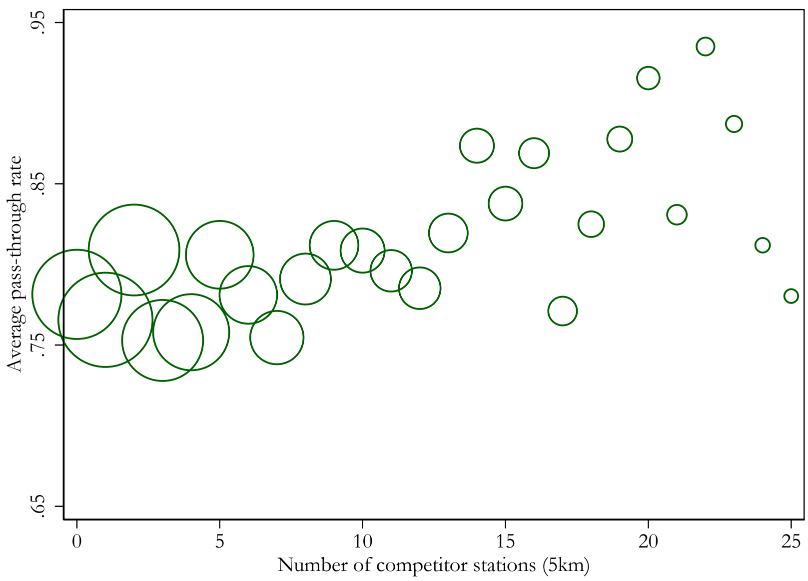

Notes: Each circle plots the average pass-through rate for a group of stations with a particular number of competitors within a 5 km catchment area. The number of competitor stations is trimmed at the top percentile.

We repeat this analysis for the tax increase in winter 2020/21 in Appendix C.4. For *E5*, we find a hump-shaped relationship as for the tax decrease. For *E10* and diesel, the relationship between the number of sellers and the pass-through rate is flat or even mildly increasing, as it was for diesel in summer 2020. This suggests that if pass-through is very high on average, the number of sellers has less of an impact on pass-through rates than if pass-through is at an intermediate level.

3.7 Conclusion

In this chapter, we investigated what determines pass-through of commodity taxes when consumers have incomplete information about prices. We began by setting up a theoretical search model in which there are some consumers that react strongly to lower prices and others that do not. By modelling the price sensitivity of consumers as the share of consumers that react strongly to lower prices, we introduced a novel notion of price sensitivity to the tax pass-through literature, which usually analyzes price sensitivity in the context of the price elasticity of demand. We show that this new way of modelling price sensitivity reverses the predictions on how price sensitivity affects pass-through. In our setting, the higher the price sensitivity of consumers, the higher the pass-through rate, because more price sensitive consumers let the market converge towards Bertrand competition.

In the second part of our analysis, we used data on fuel prices at all fuel stations in Germany and France to study how a temporary tax decrease and subsequent tax increase six months later, was passed through to consumers. In both cases, we find that pass-through is higher in markets with more price sensitive consumers.

These findings have important implications for economic policy. Whether the corrective goal of a Pigouvian tax or subsidy can be achieved hinges on whether the agents that should change their behavior also bear the incidence of the measure. Similarly, unconventional fiscal policy can only be effective in stimulating demand if consumers expect tax cuts to be passed through by firms. Finally, tax pass-through determines the distributional consequences of taxes and subsidies.

By showing how price sensitivity affects pass-through when consumers are imperfectly informed, we shed light on a novel explanation of what determines tax pass-through. Our findings are relevant beyond fuel markets and should be considered in any market where consumers do not know all prices. In these cases, policymakers should try to assess the extent to which information asymmetries exist, take these into consideration when predicting the effect of new taxes, and potentially accompany this with complementary measures targeting consumer behavior directly.

Appendices

Appendix A

Appendix to Chapter 1

A.1 Appendix to Section 1.2: Institutional Background

This Section provides some anecdotal evidence on the role of the supervisory board and the employee representatives on the board.

Gold et al. (2010) report experiences of the employee representatives on supervisory boards in German firms in a series of four interviews. Several common points emerge across the interviews. First, all employee representatives comment that they have the same rights and duties on the board as the shareholder representatives. Second, the interviewees report that their opinions are taken into the account by the shareholders and that the board generally aims to achieve a consensus: “We [worker representatives] have never felt much in a minority... I don’t think the shareholder side sees us as a minority”; “we [worker representatives and shareholders] operate on equal footing”; “I don’t feel in a minority or any kind of inferiority. Both sides try to achieve unanimity”. Third, most of the interviewed employee representatives insist that they prioritize the interests of the employees: “I have to focus on getting as much as I can for the employees... We [worker representatives] obviously have reasons, mainly to do with preserving jobs, why we assent to a particular decision or do not oppose it.” Fourth, the interviewees are persuaded that they have profoundly more company knowledge than the shareholders and are therefore able to improve decision-making on the board: “As employee representatives we have the big advantage, in contrast to the shareholders’ representatives, in that we know the structures and culture of the company intimately.”

Further, in a survey conducted by I.M.U. (2021) in 2019, 506 employee representatives in supervisory boards of German firms answered questions about the general conditions and the orientation of the supervisory work as well as the role of codetermination.[1] The survey results suggest that in the firms where the respondents work the supervisory board meets regularly, engages in topics beyond its immediate responsibility and mostly not only controls but also advises the firm management: 29 percent of the respondents report that the supervisory board holds medium- to long-term planning retreats in addition to its regular meetings; the respondents rank "medium- to long-term strategy", "restructuring and reorganization", and "impact of technological change on business model" along with "annual financial statements" among the top topics that are discussed by the board; 64 percent of the respondents report that the supervisory board provides some advice to the management. The employee representatives regard issues related to employment ('development of employment' and 'personnel policy') as both among the most important to them and where they have the most influence on the board.

Overall, the anecdotal evidence suggests that supervisory boards tend to discuss issues related to corporate strategy, firm organization, and technological change and advise the executive board on these topics. The role of the supervisory board is thus likely to go beyond the supervision of the management board. Employee representatives tend to actively participate in the discussions on the board, particularly when matters directly relate to employee interests, and can contribute to the decision-making at least to some extent.

A.2 Appendix to Section 1.3: Data

A.2.1 Sample restrictions

The sample construction proceeds as follows. In the first step, I collect Orbis data on firms that incorporate in Germany between 1989 and 1999 and are present in the 2020 database version. I then classify firms as stockholder firms and LLCs using Orbis national legal form information, where:

[1] I.M.U. is the Institute for Codetermination and Corporate Governance that operates within the Hans Böckler Foundation in Germany.

- Stock firm is a public limited company (AG), limited partnership by shares (KGaA), or limited liability company and partnership by shares (GmbH & Co. KGaA);
- LLC is a limited liability company (GmbH) or limited liability company and partnership (GmbH & Co. KG).

I keep firms classified as stock firms or LLCs and drop firms located in East Germany.

In the second step, I introduce further sample restrictions based on the types of firms or industries that have a heavy state involvement or are largely always exempt from codetermination (§1 *DrittelbG*).[2] The following firms or industries are excluded:

1. State-owned firms: to check whether a firm is state-owned, I use Orbis ownership data on current shareholders. A firm is classified as state-owned if more than 50 percent of its shares are owned by a shareholder with a type “Public authority, state, government” or “Public”, or by a shareholder with the name “KFW” (state-owned development bank in Germany).
2. Family-owned firms: stock firms with fewer than 500 employees and owned by a single family were already exempt from codetermination prior to the 1994 reform. I assume that a firm is family-owned if its individual shareholders with the same last name possess at least 99.99 percent of the shares. I use Orbis ownership data and shareholder last name information to classify firms as family-owned.
3. Subsidiaries of state-owned firms: these are firms with Domestic Ultimate Ownership link indicating more than 50 percent ownership by a government entity. In Orbis, this is the domestic ultimate ownership where the ‘DUOType’ variable is “Public authority, state, government”.
4. Branches and firms with fewer than 10 employees as they are exempt from the codetermination mandate.
5. Firms that were previously state-owned and privatized in the 1990s: these are the firms with links to Deutsche Bahn, Deutsche Telekom, or Deutsche Post DHL. I use Orbis company name information to see if a firm has a link to Deutsche Bahn, Deutsche Telekom, or Deutsche Post DHL.

[2]This mostly follows the sample restrictions that are used in Jäger, Schoefer, et al. (2021).

6. Not-for-profit firms: a firm is classified as not-for-profit if it has a “g” prefix in a company name, e.g. ‘gAG’ or ‘gGmbH’.
7. Industries with heavy state involvement, or with a large share of non-profit and media firms, which are mostly exempt from codetermination (§1 *DrittelbG*): NACE 35-39 (utilities), 490-492 (rail transport), 5813 (publishing), 60 (broadcasters), 72 (scientific), 84 (public administration), 85 (education) except 8553, 87-88 (charities), 94 (membership organizations), 97 (households as employers), 98 (private households), and 99 (extraterritorial bodies).

A.2.2 Worker share on the board

Figure A.1 shows the average share of workers on supervisory boards of publicly listed stock firms that incorporate within five years around the August 10, 1994 reform. The figure shows the share separately for firms that incorporate before and after the reform and for firms with up to 500 and more than 500 employees. The data on board composition is from the 1990s editions of the Hoppenstedt Aktienführer since there is a structural break in data reporting that starts in the 2000s.

Figure A.1 shows that there is no sizable difference in the share of workers on the board between larger listed stock firms that incorporate before and after the 1994 reform. Independent of the incorporation period, the listed stock firms with more than 500 employees have a slightly larger than one-third share of employee representatives on their boards. This is in line with the codetermination law in Germany, which mandates that firms with between 501 and 2,000 employees assign one-third share and firms with more than 2,000 employees allocate one-half share of the board seats to employee representatives.

Figure A.1: Empirical worker share on the board

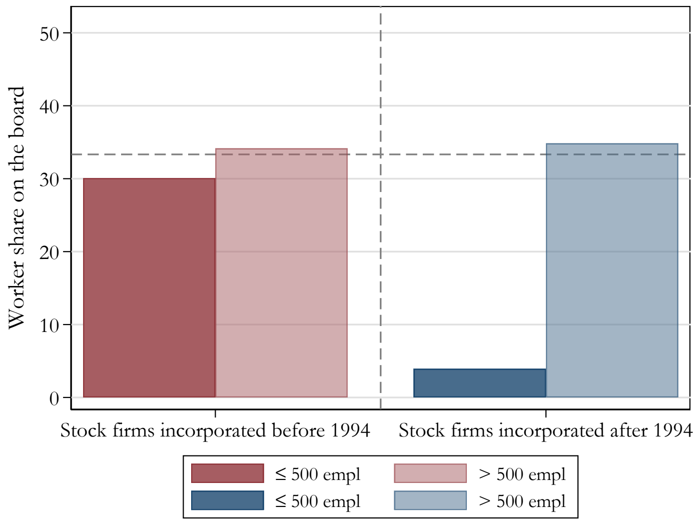

Notes: The Figure shows average share of worker representatives on supervisory boards of listed stock firms that incorporate within 5 years around the August 10, 1994 reform, separately for firms with up to 500 employees and with more than 500 employees. The data on supervisory board composition is from the Hoppenstedt Aktienführer and is based on the 1990s editions, as there is a structural break in reporting in the 2000s. The Figure is based on all firm-years where at least one third of supervisory board members have non-missing information on their role (worker representative vs. shareholder). The sample excludes stock firms that incorporate in the year 1994, as the Hoppenstedt Aktienführer does not report the full incorporation date, and stock firms located in East Germany.

A.2.3 Supervisory board composition

As an additional intervention check, I follow Jäger, Schoefer, et al. (2021) and study how the 1994 reform affects the supervisory board composition.[3]

To retrieve information on supervisory board members, I use the 2022 snapshot of the Bureau van Dijk's Orbis database. Bureau van Dijk's Orbis provides data on the names and roles of board members, directors, and senior managers. The data is a cross-section at the individual-company level. I identify supervisory board members through "Aufsichtsrat" or "Supervisory Board" in their role or department and keep data on individuals who are members of supervisory boards in firms in my sample.

In the next step, I classify supervisory board members along gender, nobility status, and Ph.D. or professorship status. The classification proceeds as follows. Gender is identified through the gender indicator in the data set. To determine nobility status, I use individuals' names and label board members as belonging to the nobility if their name contains "von", "v.", "Graf",

[3] The analysis follows the working paper version of Jäger, Schoefer, et al. (2021).

"Gräfin", "Baron", "Baronin", "Freiherr", "Frhr", "Freifrau", "Frfr", or "zu". Further, "Prof", "Professor", "Doktor" or "Dr." in the name indicates Ph.D. or professorship status. After identifying these demographic characteristics, I aggregate the information to the firm level.

Table A.1 shows how the 1994 reform affects the supervisory board composition along gender, nobility and Ph.D. / Professor status. The estimation follows the difference-in-differences specification and uses the sample of stock firms and LLCs that incorporate within three years around the reform. The estimation in all columns includes two-digit industry fixed effects.

The results in Columns (1) and (2) of Table A.1 suggest that codetermination has a positive but imprecise effect on the probability of having at least one woman on the supervisory board (12.2 pp increase) and the share of women on the board (5.5 pp increase). This could, to some extent, reflect the codetermination requirement to appoint at least one female employee representative to the supervisory board in establishments with more than 50 percent female employment (§76 *BetrVG* 1952).

Column (3) shows that codetermination reduces the probability of having at least one supervisory board member with a nobility title by 8.8 pp. Column (4) shows that codetermination also has a negative but imprecisely estimated effect on the share of board members belonging to the nobility. As around 0.1 percent of the population in Germany has a nobility title and the corresponding average percent in supervisory boards of control stock firms is 1.32, the decline further suggests that locked-in stock firms comply with the codetermination mandate. Finally, Columns (5) and (6) show that codetermination does not lead to significant changes in the likelihood or share of supervisory board members with a doctorate degree or professorship status.

Overall, the results in Table A.1 are in line with the binding nature of the 1994 reform and the descriptive evidence on the worker share in supervisory boards of firms with and without shared governance presented in Figure 1.1.

Table A.1: Effect of shared governance on supervisory board composition

	𝟙(Women > 0)	Share women	𝟙(Nobility > 0)	Share nobility	𝟙(Dr./Prof. > 0)	Share Dr./Prof.
	(1)	(2)	(3)	(4)	(5)	(6)
Pre-reform × Stock	0.122	0.055	-0.088***	-0.011	0.029	-0.040
	(0.081)	(0.035)	(0.031)	(0.007)	(0.076)	(0.045)
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	645	645	645	645	645	645
R^2	0.090	0.075	0.050	0.042	0.098	0.095

Notes: The Table reports the effects of codetermination on supervisory board composition, with outcome variables indicated in each column. The estimation follows the difference-in-differences specification and uses the sample of stock firms and LLCs incorporated within a three-year window around August 10, 1994. Estimation in all columns controls for two-digit NACE industry fixed effects. Standard errors are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.3 Appendix to Section 1.4: Empirical Strategy

A.3.1 Industry composition of firms

Figure A.2 shows estimates with 95 percent confidence intervals for the effect of the August 10, 1994 reform on the industry composition of firms. Each estimate corresponds to a separate regression of an industry category indicator on firm incorporation before the reform dummy in a sample of stock firms (simple difference) or on incorporation before the reform dummy, the stock firm indicator and the interaction of the two in a sample of both stock firms and LLCs (difference-in-differences). The estimates have a p-value of .491 (simple difference) or .381 (DID) in an F-test and are jointly not significant. The reform does not appear to lead to a significant change in the industry composition of firms.

Figure A.2: Industry composition of firms

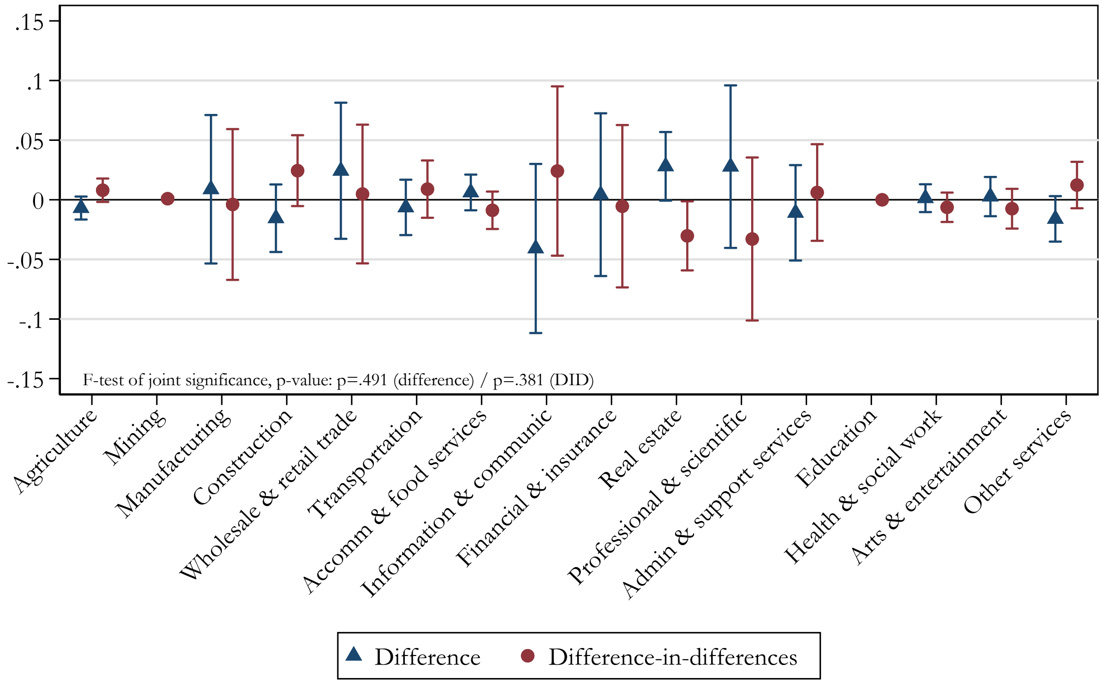

Notes: The Figure shows estimates for the change in the industry composition due to the reform using a simple difference and a difference-in-differences specification. The Figure includes 95% confidence intervals along with the estimates. Depending on the specification, the estimates have a p-value of .491 or .381 in an F-test and are not jointly significant. The estimation uses the sample of firms that incorporate within three years around August 10, 1994.

A.3.2 Selection into stock corporation

Table A.2 reports results for a test of whether firms select into a particular legal form due to the 1994 reform. The underlying data is presented in Figure 1.5. The estimates correspond to the regression of incorporation as a stock firm indicator on incorporation date (relative to August 10, 1994), an indicator for incorporation before the reform, and the interaction of the two in a sample of stock firms and LLCs. The estimation in Column (2) includes two-digit NACE industry fixed effects.

The results in Table A.2 suggest that firms do not disproportionately choose the stock corporation as their legal form after the 1994 reform date.

Table A.2: Selection into stock corporation status

	(1)	(2)
	$\mathbb{1}$(Stock firm > 0)	$\mathbb{1}$(Stock firm > 0)
Incorp date	0.0023*	0.0010
	(0.0014)	(0.0013)
Pre-reform	0.0020	0.0006
	(0.0044)	(0.0043)
Incorp date × pre-reform	0.0003	0.0006
	(0.0019)	(0.0018)
Industry FE	No	Yes
Observations	23,282	23,282
Adjusted R^2	0.001	0.057

Notes: The Table reports results for a test of whether firms select into a certain legal status due to the reform. The estimates correspond to the OLS regression of being incorporated as a stock firm indicator on incorporation date (relative to August 10, 1994), an indicator for whether a firm incorporates pre-reform, and the interaction of the two. Standard errors are clustered at a firm level and are reported in parentheses. Sample: 22,648 LLCs and 634 stock firms that incorporate within a four-year window around the 1994 reform.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.4 Appendix to Section 1.5: Results

A.4.1 Placebo reforms

Table A.3 shows the effects of the placebo reforms on the employee-adjusted number of patents, automation and non-automation patents. Columns (1) to (3) include the estimates for the effect of the August 10, 1996 placebo reform. Columns (4) to (6) include the estimates for the effect of the August 10, 1997 placebo reform. The estimation follows the baseline difference-in-differences specification as described in Equation 1.2 and uses the sample of stock firms and LLCs that incorporate within two years around the placebo reform date. Thus, it includes only never treated stock firms, i.e. the stock firms that are not subject to the codetermination mandate. All estimates are from a Poisson pseudo-likelihood regression. The estimation in all columns includes year by two-digit industry fixed effects.

Table A.3 shows that both placebo reforms do not have significant effects on the employee-adjusted number of patents, automation and non-automation patents. The results in Columns (1) to (3) are larger in magnitude but are imprecisely estimated. The estimates in Columns (4) to (6) are close to zero and not statistically significant. The results in Table A.3 suggest that the placebo reforms do not affect firm innovation.

Table A.3: Effect of placebo reforms on firm innovation

	Placebo reform: August 10, 1996			Placebo reform: August 10, 1997		
	Patents	Automat	Non-autom	Patents	Automat	Non-autom
	(1)	(2)	(3)	(4)	(5)	(6)
Pre-placebo × Stock	0.823	0.233	0.935	-0.020	0.083	-0.096
	(0.693)	(0.844)	(0.713)	(0.564)	(0.723)	(0.586)
Year × industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	148,453	110,013	146,737	159,106	118,636	156,684
Log-likelihood	-61804.44	-14369.18	-50527.02	-66035.95	-13971.26	-55691.41

Notes: The Table reports estimates of the effect of the placebo reform date on August 10, 1996 in Columns (1) to (3) and of the placebo reform date on August 10, 1997 in Columns (4) to (6). Columns (1) and (4) include estimates of the effect of the placebo reforms on the yearly number of firm patent applications per 100 employees filed between 2000 and 2014. Columns (2) and (5) include estimates of the effect on the number of firm automation patents per 100 employees. Columns (3) and (6) report estimates of the effect on the number of firm non-automation patents per 100 employees. The estimation is based on Equation 1.2 and uses the sample of both stock firms and LLCs that incorporate within a two-year window around the respective placebo reform date. All estimates are from a Poisson pseudo-likelihood regression. Standard errors are clustered at a firm level and are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.4.2 Estimation using alternative automation classification and process patents

Table A.4 shows the effect of shared governance on automation innovation using an alternative, keyword-based automation classification. A patent is classified as an automation patent if its abstract contains at least one of the following keywords: “automat”, “execut”, “inform”, “detect”, “input”, “system”, or “display”. The list of the keywords is based on Mann and Püttmann (2021) who manually classify 560 patents into automation and use machine learning with this training sample to eventually classify all USPTO patents granted between 1976 and 2014. The firm-level correlation between the number of automation patents using the keyword-based automation measure and the classification that categorizes patent technology categories first (the baseline) is at 83 percent in the sample.

Column (1) of Table A.4 includes the effect of codetermination on the number of automation patents per 100 employees filed between 1998 and 2014 and estimated using OLS regression. Column (2) of Table A.4 reports the analogous effect using a Poisson pseudo-likelihood regression. The results in Columns (3) and (4) include the codetermination effect on the number of automation patents and on the number of citation-weighted automation patents filed between 1998 and 2014, respectively, which are estimated using a Poisson pseudo-likelihood regression. The estimation in all columns is based on the baseline difference-in-differences specification and includes year by two-digit industry fixed effects. Table A.4 shows that shared governance has a sizable albeit imprecisely estimated negative effect on automation innovation when the keyword-based automation classification is in use.

Table A.5 shows the effect of codetermination on the number of employee-adjusted patents, automation and non-automation patents when the sample is constrained to process patents. Process patents may correspond more closely to innovations that are used within a firm instead of being licensed away (Klepper, 1996; Danzer et al., 2020). A patent is classified as a process patent if it contains “method”, “process”, or “procedure” in its claims text. The results in all columns are estimated using the baseline difference-in-differences specification as described in Equation 1.2, with either year or year by industry fixed effects. Columns (3) and (4) of Table A.5 show that shared governance of firms leads to a sizable negative effect on automation innovation. Columns (2) and (6) of Table A.5 show that the effect on innovation and non-automation innovation is negative but smaller in magnitude and imprecise. The results remain robust to the restriction of the sample to process patents.

Table A.4: Effect of shared governance on automation innovation using keyword-based automation classification

	(1)	(2)	(3)	(4)
Pre-reform × Stock	-0.012	-0.642	-0.801	-1.269
	(0.011)	(0.579)	(0.663)	(0.791)
Year × industry FE	Yes	Yes	Yes	Yes
Observations	301,546	202,679	202,679	202,679
Mean outcome	0.019			
Log-likelihood		-31055.359	-22186.297	-69773.006
Adjusted R^2	0.000			

Notes: The automation keywords are automat, execut, inform, detect, input, system, and display. The estimation is based on Equation 1.2 and uses the sample of stock firms and LLCs incorporated within a three-year window around August 10, 1994. Column (1) includes estimate of the effect of codetermination on the yearly number of firm automation patents per 100 employees filed between 1998 and 2014, using OLS regression, and Column (2) reports the analogous estimate, using a Poisson pseudo-likelihood regression. Column (3) includes estimate of the effect on the number of firm automation patents filed between 1998 and 2014, using a Poisson pseudo-likelihood regression. Column (4) reports estimate of the effect on the number of firm citation-weighted automation patents filed between 1998 and 2014, using a Poisson pseudo-likelihood regression. Effects in percent based on Poisson regression: -47% in Column (2); -55% in Column (3); -72% in Column (4). Standard errors are clustered at a firm level and are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.5: Effect of shared governance on process innovation

	Patents		Automation patents		Non-automation	
	(1)	(2)	(3)	(4)	(5)	(6)
Pre-reform × Stock	-0.660	-0.799	-1.283	-1.462*	-0.566	-0.701
	(0.671)	(0.674)	(0.786)	(0.809)	(0.692)	(0.694)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Year × industry FE	No	Yes	No	Yes	No	Yes
Observations	301,563	224,036	301,563	165,741	301,563	215,192
Log-likelihood	-61712.982	-51402.052	-23165.741	-18300.296	-45696.111	-38243.936

Notes: Columns (1) and (2) include estimates of the effect of codetermination on the yearly number of firm patent applications per 100 employees between 1998 and 2014. Columns (3) and (4) include estimates of the effect on the number of firm automation patents per 100 employees. Columns (5) and (6) report estimates of the effect on the number of firm non-automation patents per 100 employees. The estimation is based on Equation 1.2 and uses the sample of stock firms and LLCs that incorporate within a three-year window around August 10, 1994. All estimates are from a Poisson pseudo-likelihood regression. The sample of patents is constrained to process patents. A patent is classified as a process patent if it contains 'method', 'process' or 'procedure' in its claims text. Standard errors are clustered at a firm level and are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.4.3 Estimation using OLS regression and IHS-transformed outcome

The baseline results for the effect of codetermination on innovation, automation and non-automation innovation are estimated using a Poisson pseudo-likelihood regression. In the following, I re-estimate the difference-in-differences specification from Equation 1.2 using OLS regression and the inverse hyperbolic sine transformation of the outcome variable.

Table A.6 shows the effect of shared governance on the number of patents, automation and non-automation patents that firms file per 100 employees between 1998 and 2014, using OLS regression. Table A.7 shows the analogous effects using OLS regression and inverse hyperbolic sine transformation of the respective outcome variables. The estimation of the results in both tables follows the difference-in-differences specification and uses a sample of both stock firms and LLCs.

Column (4) in Table A.6 and Column (4) in Table A.7 show that codetermination leads to a statistically significant and sizable decrease in automation innovation. Columns (2) and (6) in Table A.6 and Table A.7 show that the effect on innovation and non-automation innovation is also negative but imprecise. The results remain robust to using OLS regression and IHS-transformed outcome variables.

Table A.6: Effect of shared governance on firm innovation, estimated using OLS

	Patents		Automation patents		Non-automation	
	(1)	(2)	(3)	(4)	(5)	(6)
Pre-reform × Stock	-0.072	-0.081	-0.018*	-0.021**	-0.056	-0.062
	(0.053)	(0.053)	(0.010)	(0.010)	(0.048)	(0.048)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Year × industry FE	No	Yes	No	Yes	No	Yes
Observations	301,563	301,546	301,563	301,546	301,563	301,546
Adjusted R^2	0.000	0.002	0.000	0.000	0.000	0.002
Mean outcome	0.069	0.069	0.015	0.015	0.054	0.054

Notes: Columns (1) and (2) include estimates of the effect of codetermination on the yearly number of firm patent applications per 100 employees between 1998 and 2014. Columns (3) and (4) include estimates of the effect on the number of firm automation patents per 100 employees. Columns (5) and (6) report estimates of the effect on the number of firm non-automation patents per 100 employees. The estimation is based on Equation 1.2 and uses the sample of stock firms and LLCs that incorporate within a three-year window around August 10, 1994. All estimates are from an OLS regression. Standard errors are clustered at a firm level and are reported in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.7: Effect of shared governance on firm innovation, IHS-transformed outcome

	Patents		Automation patents		Non-automation	
	(1)	(2)	(3)	(4)	(5)	(6)
Pre-reform × Stock	-0.012	-0.012	-0.006*	-0.007*	-0.006	-0.007
	(0.015)	(0.015)	(0.004)	(0.004)	(0.014)	(0.014)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Year × industry FE	No	Yes	No	Yes	No	Yes
Observations	301,563	301,546	301,563	301,546	301,563	301,546
Adjusted R^2	0.000	0.017	0.001	0.006	0.000	0.015

Notes: Columns (1) and (2) include estimates of the effect of codetermination on the yearly number of firm patent applications per 100 employees between 1998 and 2014. Columns (3) and (4) include estimates of the effect on the number of firm automation patents per 100 employees. Columns (5) and (6) report estimates of the effect on the number of firm non-automation patents per 100 employees. The estimation is based on Equation 1.2 and uses the sample of both stock firms and LLCs incorporated within a three-year window around August 10, 1994. All columns use the inverse hyperbolic sine transformation of the respective dependent variable. All estimates are from an OLS regression. Standard errors are clustered at a firm level and are reported in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.4.4 Estimation using 1998 to 2007 time period

The descriptive evidence presented in Figure 1.2 of Section 1.3 showed that trends in the evolution of the employee-adjusted number of patents are not always parallel throughout the 1998 to 2014 time period, in particular after 2007. In the following, I re-estimate the baseline codetermination effects on firm innovation, automation and non-automation innovation when the time period is restricted to 1998 to 2007.

Table A.8 shows the effect of shared governance on the number of patents, automation and non-automation patents filed by firms per 100 employees between 1998 and 2007. The estimation in all columns follows the baseline difference-in-differences specification as described in Equation 1.2 and uses a Poisson pseudo-likelihood regression.

Columns (3) and (4) in Table A.8 show that codetermination leads to a sizable negative effect on automation innovation. Columns (2) and (6) show that the effect on innovation and non-automation innovation is also negative but smaller in magnitude and imprecise. The results remain similar to using the 1998 to 2014 time period.

Table A.8: Effect of shared governance on firm innovation, 1998 to 2007

	Patents		Automation patents		Non-automation	
	(1)	(2)	(3)	(4)	(5)	(6)
Pre-reform × Stock	-0.524	-0.615	-1.047	-1.171*	-0.472	-0.566
	(0.621)	(0.625)	(0.690)	(0.689)	(0.677)	(0.681)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Year × industry FE	No	Yes	No	Yes	No	Yes
Observations	177,390	147,199	177,390	117,808	177,390	142,578
Log-likelihood	-69774.637	-59805.892	-25518.223	-21011.876	-51948.479	-44813.841

Notes: Columns (1) and (2) include estimates of the effect of codetermination on the yearly number of firm patent applications per 100 employees between 1998 and 2007. Columns (3) and (4) include estimates of the effect on the number of firm automation patents per 100 employees. Columns (5) and (6) report estimates of the effect on the number of firm non-automation patents per 100 employees. The estimation is based on Equation 1.2 and uses the sample of stock firms and LLCs that incorporate within a three-year window around August 10, 1994. All columns use patent data between 1998 and 2007. All estimates are from a Poisson pseudo-likelihood regression. Standard errors are clustered at a firm level and are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.4.5 Codetermination effect under different bandwidths: simple difference specification

Section 1.5 reports the codetermination effect on innovation, automation and non-automation innovation using different bandwidths around the 1994 reform when the estimation follows the difference-in-differences specification. In the following, I include the codetermination effects under different bandwidths, which are estimated using the simple difference specification from Equation 1.1.

Figure A.3 shows the estimates for the codetermination effect on the number of patents filed by firms per 100 employees between 1998 and 2014, using a simple difference specification. Figure A.4 and Figure A.5 include the estimates for the effect on the employee-adjusted number of automation and non-automation patents, respectively, similarly using a simple difference regression.

Figure A.4 shows that the negative effect of shared governance on automation innovation is sizable and statistically significant for a majority of the bandwidths around the 1994 reform. Figure A.3 and Figure A.5 show that the codetermination effect on innovation and non-automation innovation is negative but the confidence intervals do not exclude zero for a majority of the bandwidths. The results remain similar to the difference-in-differences estimates presented in Figure 1.6, Figure 1.7 and Figure 1.8.

Figure A.3: Effect of shared governance on innovation

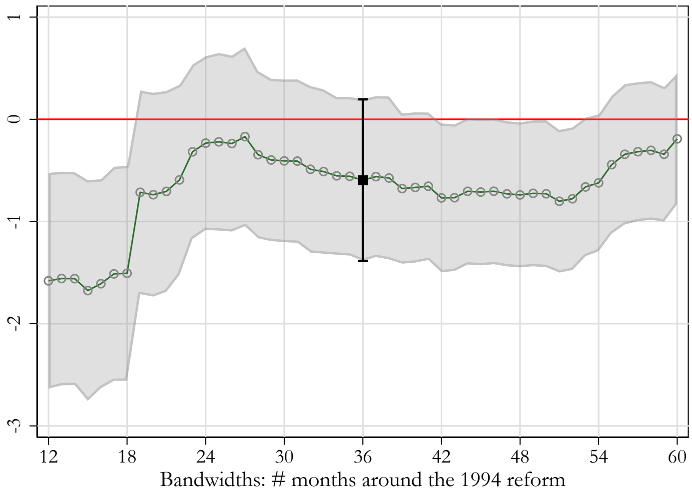

Notes: The Figure shows the effect of codetermination on the number of firm patents per 100 employees, using different bandwidths around the August 10, 1994 reform. The Figure includes 90% confidence intervals along with the estimates. The effects are estimated using a simple difference design in a sample of stock firms and Poisson pseudo-likelihood regression, as described by Equation 1.1. All specifications include year by two-digit NACE industry fixed effects. Standard errors are clustered at a firm level. Black highlights the baseline simple difference estimate reported in Table 1.2.

Figure A.4: Effect of shared governance on automation innovation

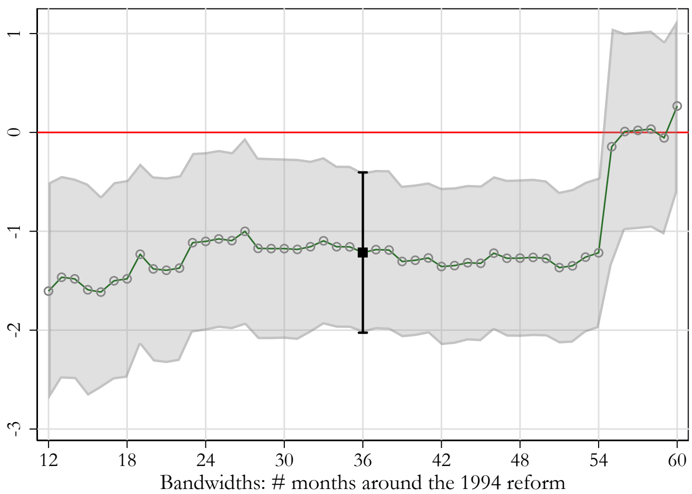

Notes: The Figure shows the effect of codetermination on the number of firm automation patents per 100 employees, using different bandwidths around the August 10, 1994 reform. The Figure includes 90% confidence intervals along with the estimates. The effects are estimated using a simple difference design in a sample of stock firms and Poisson pseudo-likelihood regression, as described by Equation 1.1. All specifications include year by two-digit NACE industry fixed effects. Standard errors are clustered at a firm level. Black highlights the baseline simple difference estimate reported in Table 1.2.

Figure A.5: Effect of shared governance on non-automation innovation

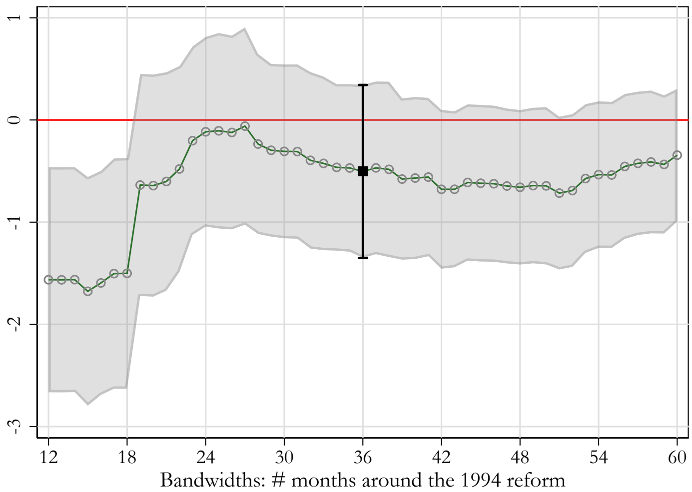

Notes: The Figure shows the effect of codetermination on the number of firm non-automation patents per 100 employees, using different bandwidths around the August 10, 1994 reform. The Figure includes 90% confidence intervals along with the estimates. The effects are estimated using a simple difference design in a sample of stock firms and Poisson pseudo-likelihood regression, as described by Equation 1.1. All specifications include year by two-digit NACE industry fixed effects. Standard errors are clustered at a firm level. Black highlights the baseline simple difference estimate reported in Table 1.2.

A.5 Appendix to Section 1.6: Mechanisms

A.5.1 Effect of codetermination on innovation: average citations

Section 1.6 suggests that codetermination reduces the value of automation patents, as reflected by the effect heterogeneity along patent distance to science and the effect on the number and age of patent backward citations. In the following, I study the effect on the number of forward citations per patent, which informs about how codetermination changes patent quality.

Table A.9 presents the estimates for the effect of shared governance on the number of forward citations per patent. Columns (1) to (2) include the estimates for all patents, and Columns (3) to (6) differentiate between automation and non-automation patenting. The estimation in all columns follows the difference-in-differences specification and uses a Poisson pseudo-likelihood regression.

Column (4) in Table A.9 shows that codetermination reduces the number of forward citations per patent by 67 percent for automation patents. Column (6) shows that the effect on the average number of forward citations for non-automation patents is close to zero and not statistically significant. The findings suggest that codetermination strongly reduces the quality of automation innovation but does not lead to a change in the quality of non-automation patents.

Table A.9: Effect of shared governance on firm innovation (average forward citations)

	Patents		Automation patents		Non-automation	
	(1)	(2)	(3)	(4)	(5)	(6)
Pre-reform × Stock	-0.018	-0.028	-1.059*	-1.114**	0.081	0.072
	(0.404)	(0.393)	(0.579)	(0.567)	(0.437)	(0.422)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Year × industry FE	No	Yes	No	Yes	No	Yes
Observations	301,563	248,505	301,563	181,949	301,563	243,123
Log-likelihood	-69714.334	-57300.536	-25293.555	-20769.614	-63259.344	-51935.412

Notes: Columns (1) and (2) include estimates of the effect of codetermination on the yearly number of forward citations per patent for firm patents filed between 1998 and 2014. Columns (3) and (4) include estimates of the effect on the yearly average number of forward citations for automation patents. Columns (5) and (6) report estimates of the effect on the yearly average number of forward citations for non-automation patents. The number of forward citations corresponds to the number of citations that patents receive within the first three years after their application date, including the year of application. The estimation is based on Equation 1.2 and uses the sample of both stock firms and LLCs incorporated within a three-year window around August 10, 1994. All estimates are from a Poisson pseudo-likelihood regression. Standard errors are clustered at a firm level and are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.5.2 Backward citations

Figure 1.10 and Figure 1.11 in Section 1.6 show the estimates for the codetermination effect on the average number of backward citations per patent for automation and non-automation patents, respectively, using different bandwidths around the 1994 reform. In the following, Table A.10 presents the corresponding estimates for the baseline three-year bandwidth around the reform. The estimation in all columns follows the difference-in-differences specification and uses a Poisson pseudo-likelihood regression.

Table A.10 shows that shared governance of firms leads to a decline in the average number of backward citations per patent for automation patents but not for non-automation patents. Column (4) shows that, on average, automation patents at codetermined stock firms include 29.4 percent fewer backward citations per patent. Column (6) shows that the effect of codetermination on the average number of backward citations per patent for non-automation patents is nearly zero and not statistically significant.

Table A.10: Effect of shared governance on average number of backward citations

	Patents		Automation patents		Non-automation	
	(1)	(2)	(3)	(4)	(5)	(6)
Pre-reform × Stock	-0.096	-0.096	-0.352**	-0.348*	-0.044	-0.027
	(0.102)	(0.095)	(0.167)	(0.186)	(0.105)	(0.102)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Year × industry FE	No	Yes	No	Yes	No	Yes
Observations	4,985	4,818	1,360	1,213	4,420	4,259
Log-likelihood	-15928.621	-14019.866	-4123.447	-3376.496	-14310.239	-12464.924

Notes: Columns (1) and (2) include estimates of the effect of codetermination on the average number of backward citations in firm patent applications filed between 1998 and 2014. Columns (3) and (4) include estimates of the effect on the average number of backward citations in automation patents. Columns (5) and (6) report estimates of the effect on the average number of backward citations in non-automation patents. Firm-year observations with zero (automation/non-automation) patents are excluded. The estimation is based on Equation 1.2 and uses the sample of both stock firms and LLCs incorporated within a three-year window around August 10, 1994. All estimates are from a Poisson pseudo-likelihood regression. Standard errors are clustered at a firm level and are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Appendix B

Appendix to Chapter 2

B.1 Appendix to Section 2.2: Theoretical Model

B.1.1 Equilibrium price distribution

Lemma 3.1. *Given some exogenous number of entrants N, there is no pure strategy Nash equilibrium.*

Proof. Suppose all sellers set some price p above marginal cost which is normalized to zero. Then each firm sells to its share of *non-shoppers* and *shoppers*. This cannot be an equilibrium since a seller could profitably deviate by marginally decreasing the price to $p - \epsilon$ and capture all the *shoppers*.

Suppose now that in equilibrium all sellers set a price at the marginal cost normalized to zero, i.e. $p_i = 0$ for any $i \in \{1, ..., N\}$. This cannot be an equilibrium since a seller could profitably deviate by increasing its price above the marginal cost, which will still allow to sell to its share of *non-shoppers* and make a positive profit.

Finally, suppose that one seller sets a lower price with all other sellers choosing the same higher price. This cannot be an equilibrium since the lowest price seller could profitably deviate by marginally increasing its price and still capture all the *shoppers*.

□

Lemma 3.2. *There are no mass points in the equilibrium pricing strategies.*

Proof. Suppose that in equilibrium some price p is charged with positive probability by the sellers. This means that there is a positive probability of a tie at this price. In this case, a seller has an incentive to deviate from p to $p - \epsilon$, which is set with the same probability, since undercutting other sellers allows the deviating seller to capture all *shoppers* and increase its profits. Thus, charging any price with positive probability cannot be an equilibrium.[1]

□

Lemma 3.3. *There is a symmetric mixed strategy Nash equilibrium, in which firms draw prices from $[\underline{p}, p_r]$ according to the density function $F(p_i)$, where the reservation price p_r is*

$$p_r = \upsilon\,.$$

The minimum price which firms may set in equilibrium is

$$\underline{p} = \frac{\upsilon}{\frac{\phi N}{1-\phi} + 1}\,.$$

The cumulative density function from which firms draw prices in equilibrium is

$$F(p_i) = 1 - \left(\frac{\upsilon - p_i}{p_i}\frac{1-\phi}{N\phi}\right)^{\frac{1}{N-1}}\,.$$

The expected profit of a firm i in equilibrium is

$$E[\pi_i] = \upsilon\frac{1-\phi}{N}\,.$$

The expected price is

$$E[p] = \underline{p} + \left(\frac{1-\phi}{N\phi}\right)^{\frac{1}{N-1}} \int_{\underline{p}}^{\upsilon} \left(\frac{\upsilon - p}{p}\right)^{\frac{1}{N-1}} dp\,.$$

[1] See Varian (1980) for a detailed proof.

The expected minimum price is

$$E[p_{min}] = \frac{1-\phi}{\phi}[p_r - E[p]]\,.$$

Proof. We begin with the reservation price. Since *non-shoppers* visit a seller at random and purchase a unit of the good if its price is below their reservation price, the reservation price corresponds to the valuation of the good υ by *non-shoppers*. No firm sets a price above the reservation price of *non-shoppers*.

Next, we derive the minimum price which firms may set in equilibrium, $\underline{p}$. For that, we utilize the equiprofit condition in the mixed strategy Nash equilibrium. The expected profit that a firm receives from setting the minimum price $\underline{p}$ should be the same as the expected profit from setting the reservation price p_r:

$$E[\pi(\underline{p})] = E[\pi(p_r)]\,. \tag{B.1}$$

Since there are no mass points in equilibrium pricing strategies, a firm that sets the minimum price $\underline{p}$ sells to all *shoppers* and its share of *non-shoppers*. A firm that sets the reservation price p_r only sells to its share of *non-shoppers*. The equiprofit condition can then be rewritten as

$$\underline{p}(\phi + \frac{1-\phi}{N}) = p_r \frac{1-\phi}{N}\,. \tag{B.2}$$

Simplifying this expression and replacing the reservation price with υ, we can solve for the minimum element of the support of prices $\underline{p}$:

$$\underline{p} = \frac{\upsilon}{\frac{\phi N}{1-\phi} + 1}\,. \tag{B.3}$$

To derive the equilibrium density function, we again use the equiprofit condition, namely that in the symmetric mixed strategy Nash equilibrium any price that a seller sets with positive probability should yield the same expected profit, i.e.

$$E[\pi(p_i)] = E[\pi(p_r)] \qquad \forall \quad p_i \in [\underline{p}, p_r]\,. \tag{B.4}$$

A firm that sets the price p_i has the lowest price among all sellers with the probability $(1 - F(p_i))^{n-1}$. In this case, a firm i sells to all *shoppers* and to its share of *non-shoppers*. With the probability $1 - (1 - F(p_i))^{n-1}$, a firm that sets the price p_i is not the lowest price seller in the market. In this case, it sells the product only to its share of *non-shoppers*. Finally, if a firm i chooses the reservation price $p_r = \upsilon$, it sells the product to its share of *non-shoppers*.

We can now rewrite the equiprofit condition as

$$p_i(\phi + \frac{1-\phi}{N})(1 - F(p_i))^{N-1} + p_i(\frac{1-\phi}{N})(1 - (1 - F(p_i))^{N-1}) = p_r\frac{1-\phi}{N}\,. \tag{B.5}$$

Simplifying this expression and solving for $F(p_i)$, we derive that the equilibrium density function from which sellers draw prices from the interval $[\underline{p}, p_r]$ is

$$F(p_i) = 1 - (\frac{\upsilon - p_i}{p_i}\frac{1-\phi}{N\phi})^{\frac{1}{N-1}}\,. \tag{B.6}$$

The reservation price p_r, the minimum price $\underline{p}$ and the equilibrium density function $F(p_i)$ uniquely define the symmetric mixed strategy Nash equilibrium of the game, assuming that there is a fixed and exogenous number of firms N in the market.

We can now compute the expected profit that each seller obtains in equilibrium, which by the equiprofit condition is identical to the expected profit from setting the reservation price $p_r = \upsilon$:

$$E[\pi_i] = E[\pi(p_r)] = \upsilon\frac{1-\phi}{N}\,. \tag{B.7}$$

Finally, we can derive the expected price, which is the average price paid by *non-shoppers*, and the expected minimum price, which is the average price paid by *shoppers*.

The expected price is

$$E[p] = \int_{\underline{p}}^{p_r} pf(p)dp = p_r - \int_{\underline{p}}^{p_r} F(p)dp\,. \tag{B.8}$$

Inserting the equilibrium density function $F(p)$ and the reservation price $p_r = \upsilon$, and simplifying yields

$$E[p] = \underline{p} + (\frac{1-\phi}{N\phi})^{\frac{1}{N-1}} \int_{\underline{p}}^{\upsilon} (\frac{\upsilon - p}{p})^{\frac{1}{N-1}} dp\,.$$

The expected minimum price is

$$E[p_{min}] = \int_{\underline{p}}^{p_r} p f_{min}(p) dp\,,$$

where the probability density function of the minimum price is

$$f_{min}(p) = N(1 - F(p))^{N-1} f(p)\,. \tag{B.9}$$

After inserting the equilibrium density function $F(p)$ into the above expression, we can simplify the probability density function of the minimum price to

$$f_{min}(p) = \frac{p_r - p}{p} \frac{1-\phi}{\phi} f(p)\,. \tag{B.10}$$

We can now substitute $f_{min}(p)$ into the expression for the expected minimum price:

$$E[p_{min}] = \int_{\underline{p}}^{p_r} p f_{min}(p) dp = \int_{\underline{p}}^{p_r} p \frac{p_r - p}{p} \frac{1-\phi}{\phi} f(p) dp\,,$$

which after simplification is equivalent to

$$E[p_{min}] = \frac{1-\phi}{\phi} \left[\int_{\underline{p}}^{p_r} p_r f(p) dp - E[p] \right]\,.$$

Finally, after further simplification, the expected minimum price becomes

$$E[p_{min}] = \frac{1-\phi}{\phi} \left[\upsilon - E[p] \right]\,.$$

□

B.1.2 Omitted proofs in Section 2.2

Proof of Lemma 2.1. Let us begin by analyzing how a change in the share of *shoppers* affects the minimum price which firms may set in equilibrium. Recall that in equilibrium

$$\underline{p} = \frac{\upsilon}{\frac{\phi N}{1-\phi} + 1}\,.$$

Then, for $0 < \phi < 1$, the derivative of the minimum price with respect to the share of *shoppers* ϕ is strictly negative:

$$\frac{\partial \underline{p}}{\partial \phi} = -\frac{\upsilon N}{(\phi N + 1 - \phi)^2} < 0\,.$$

Next, we study how the share of *shoppers* affects the equilibrium price distribution. We therefore derive the derivative of the cumulative density function with respect to ϕ:

$$\frac{\partial F(p)}{\partial \phi} = \frac{1}{N(N-1)\phi^2}\frac{\upsilon - p}{p}\left[\frac{\upsilon - p}{p}\frac{1-\phi}{N\phi}\right]^{\frac{1}{N-1}-1} > 0\,.$$

Thus, with $0 < \phi < 1$, for any $\hat{\phi} > \phi$, $\hat{F}(p) \geq F(p) \quad \forall p \in [\underline{p}, p_r]$.

□

Proof of Proposition 2.1. We first study how an information shock affects the minimum price that sellers may set in equilibrium. We assume that after the information shock, the share of fully informed consumers is $\phi = \phi_0 + \Delta_\phi(1-\phi_0)$, where ϕ_0 is the ex ante share of fully informed shoppers and $\Delta_\phi(1-\phi_0)$ captures an increase in the share of informed consumers due to the shock Δ_ϕ.

Then, taking the first order derivative of the minimum element of the support of the equilibrium pricing strategy with respect to the information shock Δ_ϕ, we obtain

$$\frac{\partial \underline{p}}{\partial \Delta_\phi} = -\frac{\upsilon N}{(\frac{\phi N}{1-\phi} + 1)^2}\frac{(1-\phi_0)(1-\phi) + \phi(1-\phi_0)}{(1-\phi)^2} < 0\,.$$

We can simplify this to obtain

$$\frac{\partial \underline{p}}{\partial \Delta_\phi} = -\frac{\upsilon N(1-\phi_0)}{(\phi(N-1)+1)^2} < 0\,.$$

The minimum price that sellers may set in equilibrium strictly declines in the information shock Δ_ϕ.

We now take the first order derivative of the above expression with respect to the ex ante share of fully informed consumers in the market ϕ_0:

$$\frac{\partial^2 \underline{p}}{\partial \Delta_\phi \partial \phi_0} = \upsilon N \frac{(\phi(N-1)+1)^2 + 2(1-\phi_0)(1-\Delta_\phi)(N-1)(\phi(N-1)+1)}{(\phi(N-1)+1)^4} > 0\,.$$

We can simplify this to obtain

$$\frac{\partial^2 \underline{p}}{\partial \Delta_\phi \partial \phi_0} = \upsilon N \frac{1+(2-\phi)(N-1)}{(\phi(N-1)+1)^3} > 0\,.$$

This means that the information shock Δ_ϕ leads to a stronger downward shift in the minimum price that sellers choose in equilibrium when ex ante consumers are on average less informed.

Next, we study how the magnitude of the effect of the information shock varies with the ex ante share of fully informed consumers for the equilibrium density function. We start by taking the first order derivative of the equilibrium density function with respect to Δ_ϕ:

$$\frac{\partial F(p)}{\partial \Delta_\phi} = \frac{\upsilon - p}{pN(N-1)} \left(\frac{\upsilon - p}{p} \frac{1-\phi}{N\phi} \right)^{\frac{1}{N-1}-1} \frac{1-\phi_0}{\phi^2} > 0\,.$$

This means that an information shock that increases the share of informed consumers in the market shifts the equilibrium density function from which firms draw prices towards lower prices.

To analyze how ex ante share of shoppers affects the magnitude of this downward shift in prices, we take the first order derivative of the above expression with respect to the initial level of the share of shoppers ϕ_0 and simplify to obtain

$$\frac{\partial^2 F(p)}{\partial \Delta_\phi \partial \phi_0} = -\frac{1}{N-1} \left(\frac{\upsilon - p}{pN} \right)^{\frac{1}{N-1}} \left(\frac{1-\phi}{\phi} \right)^{\frac{1}{N-1}-1} \left(\frac{1}{N-1} + (1+\phi_0)\big(1-\Delta_\phi\big) \right) < 0\,. \tag{B.11}$$

Thus for any $\hat{\phi}_0 > \phi_0$, $\Delta\underline{\hat{p}} > \Delta\underline{p}$ and $\frac{\partial F(p)}{\partial^2 \Delta_\phi \partial \phi_0} < 0 \quad \forall p$.

□

B.2 Appendix to Section 2.3: Institutional Setting

B.2.1 Retail margins and fuel station characteristics in Germany

Figure B.1 shows the distribution of fuel stations in Germany over our sample period. Fuel stations are spread across the country and clustered around urban areas.

Figure B.1: Distribution of fuel stations across Germany

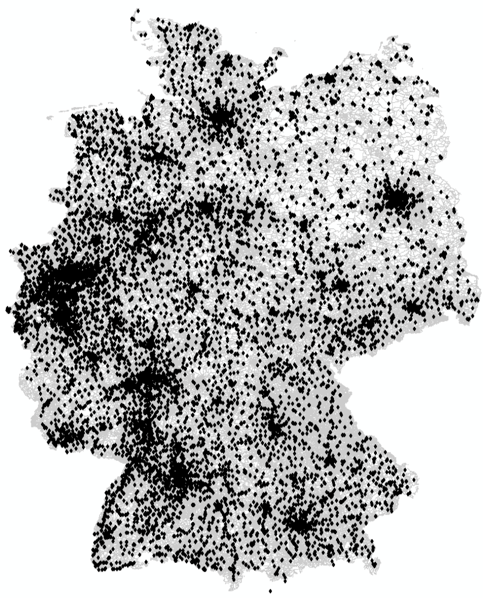

Note: The Figure shows the geographic distribution of fuel stations in Germany.

Table B.1 shows the share of the vertically integrated firms, as well as the share of non-integrated firms before and after the MTU introduction. Overall, the brand composition is very similar before and after the introduction of the MTU.

Although there are no restrictions on the number of times fuel stations can change prices in France or Germany, there are strong differences in the number of times they do. Whereas fuel stations in Germany change their prices on average four times a day over our observation

Table B.1: Share of stations in percent by brand

	Pre-MTU	Post-MTU
Aral	21.1	18
Shell	13.9	14.2
Esso	5.1	5.3
Total	7.3	4.6
Jet	5.2	4.6
Orlen	4.9	4.2
Agip	1.8	3.1
Hem	3.2	2.8
OMV	2.7	2.2
Non-integrated	34.9	41

Notes: The "Pre-MTU" column shows the share of fuel stations by brand in the sample for Germany before the introduction of the MTU. The "Post-MTU" column shows the share of fuel stations by brand in the sample for Germany after the introduction of the MTU. We consider all fuel stations that have at least one price entry in the sample before or after the MTU introduction, respectively.

period, French fuel stations change prices less than once a day.[2] Since we do not observe volume data, we cannot compute volume-weighted average fuel prices or retail margins over the day. We could thus either pick a particular time of day at which to measure prices and margins or calculate a simple average of prices and margins at different times of the day. Since fuel prices in France stay fairly constant during the day, either approach should lead to a similar result for France. The frequent price changes in Germany however, make it important to select the right time for which to calculate fuel prices and retail margins.

We choose to use prices at 5 pm in our analysis, and we construct retail margins based on these prices. A representative survey among motorists commissioned by the German Ministry for Economic Affairs and Energy (2018) in 2016 found that around 60 percent of respondents buy fuel between 4 pm and 7 pm, of which two-thirds buy fuel between 5 pm and 6 pm. At the same time, less than 5 percent of respondents buy fuel before 10 am.[3] The German Ministry for Economic Affairs and Energy (2018) furthermore documents daily price cycles with high prices in the morning, which fall over the day and rise again in the evening at around 8 pm.[4]

[2]This is consistent with findings by Haucap, Heimeshoff, Kehder, Odenkirchen, and Thorwarth (2017) for Germany and Gautier and Saout (2015) for France.

[3]The daily fuelling patterns are described in detail in Figure B.5 in Appendix B.2.1.

[4]This is consistent with pricing patterns in the data described in Figure B.6 in Appendix B.2.1.

This suggests that consumers are aware of these price cycles and fuel during the low price period in the late afternoon.[5] To gauge the effect of introducing mandatory price disclosure on consumers, it is therefore sensible to focus on fuel prices and retail margins at times where consumers buy fuel in large volumes.

In the estimation with SDID, we use weekly fuel prices. We compute the weekly fuel prices by averaging Monday to Friday prices at 5 pm. We exclude weekend prices from the analysis.

Figure B.2 shows the daily number of fuel stations for which the price panel contains a price entry at 5 pm. There is no structural break in the daily number of fuel stations for which there is an entry in the price panel before and after the MTU introduction. For most days in the pre-MTU period, we have prices for approximately $12,000$ fuel stations in our panel. This number stays approximately the same after the introduction of the MTU and only increases to around $13,500$ at the end of February 2014, when reporting issues of Total and Esso stop.[6] At any point in time over the observation period, our panel therefore includes prices for most of the approximately $14,700$ fuel stations in Germany.

Figure B.2: Number of fuel stations with positive price reports at 5pm

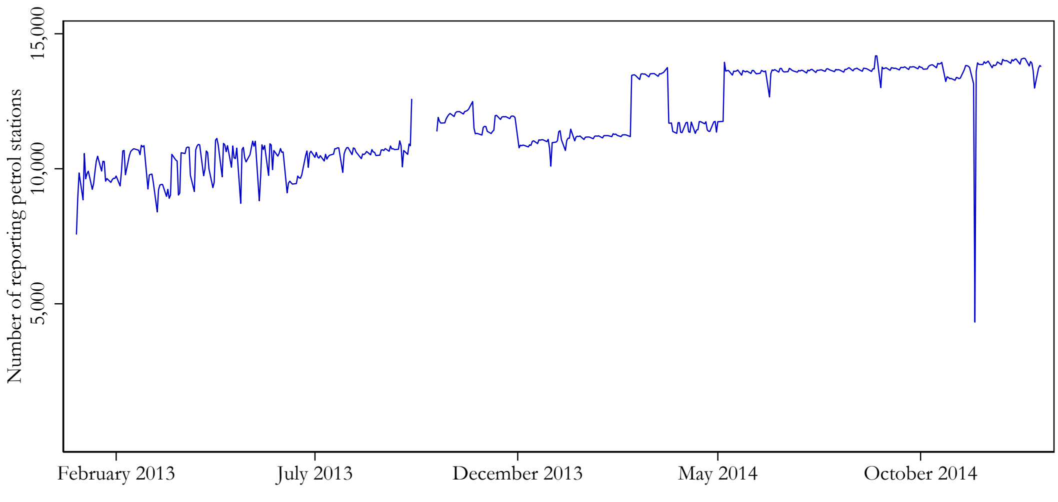

Notes: The Figure shows the average daily number of fuel stations with a positive price report at 5 pm in Germany in our sample.

[5]There are numerous newspaper articles on intertemporal price dispersion during our observation period, which suggest that consumers are aware of these patterns.

[6]Total and Esso report normally in October 2013. Esso reports only a very limited amount of prices between November 2013 and mid-February 2014. Total only reports a very limited amount of prices between December 2013 and mid-February 2014. Both experienced reporting issues in April 2014, after which they returned to full reporting.

Figure B.3 shows that there are fewer price changes per day in our data prior to the MTU introduction than after the MTU was introduced. This is because whereas after the introduction of the MTU we observe the universe of price changes in Germany, before the introduction of the MTU we only observe the subset of prices that was reported by users to the app.

Figure B.3: Number of daily price changes

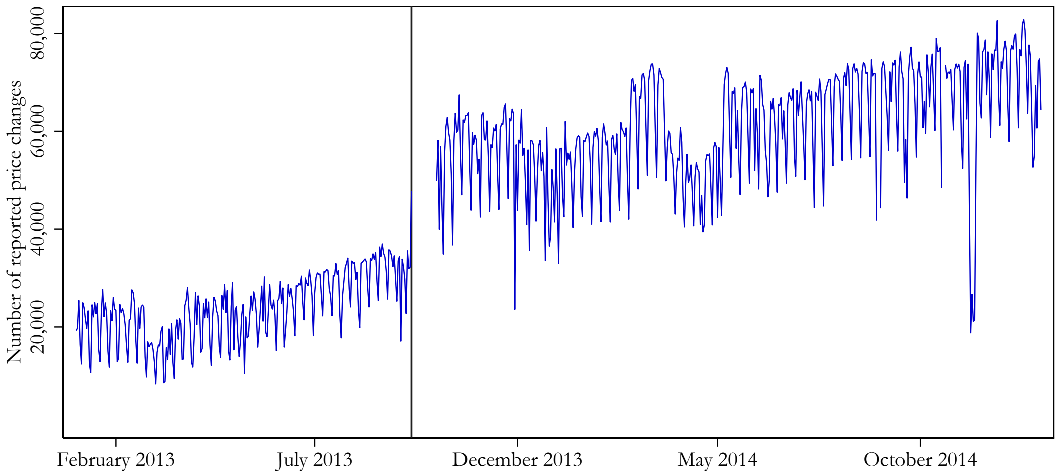

Notes: The Figure shows the average daily number of price changes in Germany in our data. In the pre-MTU period consecutive reports of the same price are not considered a price change.

Figure B.4 shows the number of notifications of price changes over the day, before and after the introduction of the MTU. Whereas before the introduction of the MTU there is a notification every time a user of the app reports a price, after the MTU there is a notification every time that there is a price change.

Figure B.5 shows the hourly fuelling patterns as reported in a representative survey among drivers commissioned by the German Federal Ministry of Economic Affairs. As discussed in Section 2.3, the majority of drivers buy fuel between 5 pm and 7 pm, whereas only very few drivers buy fuel in the morning.

The fuelling patterns are also consistent with price patterns reported in Figure B.6. Whereas gasoline and diesel prices are highest in the morning, they fall during the day until the early evening and start rising again at around 8 pm.

Figure B.4: Notification patterns over the day

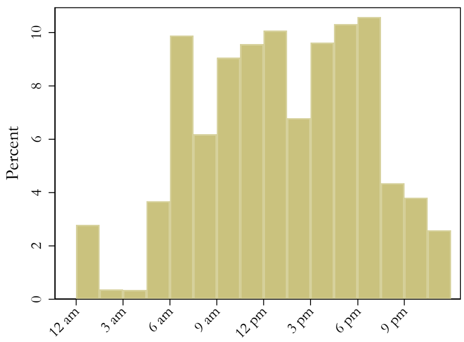

(a) Pre-MTU

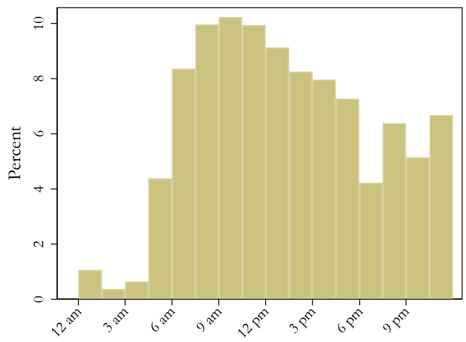

(b) Post-MTU

Notes: Panel (a) shows the share of price notifications in our data set for every hour of the day for the pre-MTU period. Panel (b) shows the share of price notifications in our data set for every hour of the day for the post-MTU period. Pre-MTU, each price report by users notifying a price change to the information service provider is a price notification. Post-MTU, each price change notified by fuel stations to the MTU is a price notification.

Figure B.5: Daily fuelling patterns

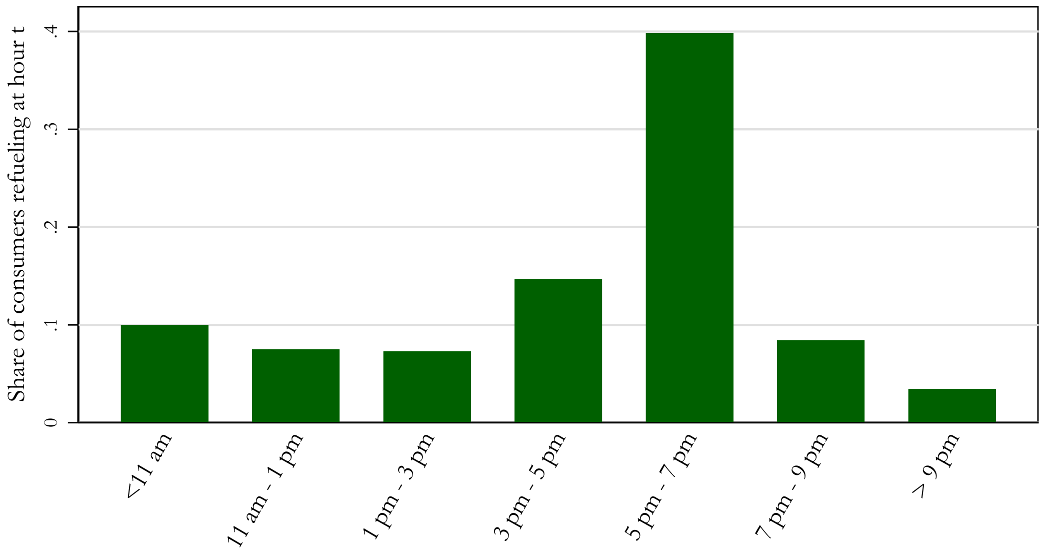

Notes: The Figure shows the average fuelling patterns by German motorists over the day. Data is based on a representative survey among drivers commissioned by the German Federal Ministry of Economic Affairs.

Figure B.6: Daily price patterns

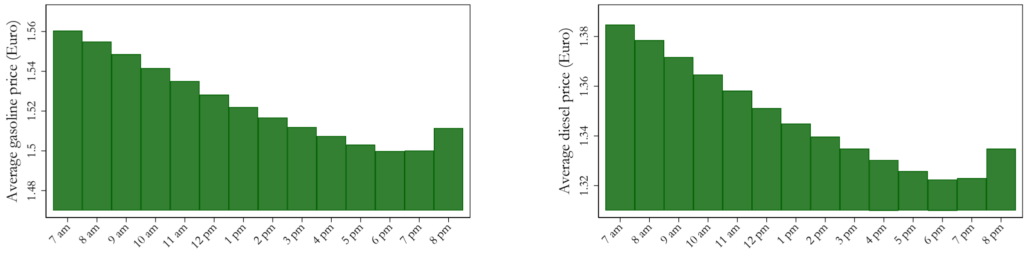

(a) Gasoline **(b)** Diesel

Notes: Panel (a) shows the average gasoline price for every hour between 7 am and 8 pm in Germany between 2013 and 2014. Panel (b) shows the average diesel price for every hour between 7 am and 8 pm in Germany between 2013 and 2014.

B.2.2 Distribution of fuel stations by SDID unit weights in France

Figure B.7 shows the geographic distribution of stations in France. Stations that receive a disproportionately high unit weight in the SDID estimation following Equation 2.1 either for gasoline or diesel are highlighted in the figure. The disproportionately weighted stations in the control group scatter throughout France. This means that potential geographic clustering via re-weighting by SDID unit weights does not affect our results.

Figure B.7: Geographic distribution of fuel stations by SDID unit weights, France

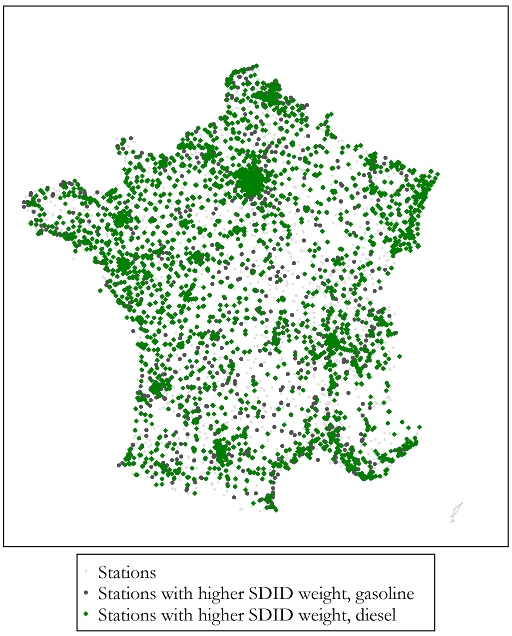

Notes: The Figure shows the geographic distribution of fuel stations in France. Stations that receive a disproportionally high unit weight in the SDID estimation are highlighted.

B.3 Appendix to Section 2.6: Results

In this Section we provide further empirical evidence on the average effect of the MTU on gasoline and diesel prices in Germany. It shows that our results in Section 2.6 are robust to using alternative specifications.

B.3.1 Difference-in-differences analysis

Since estimation by SDID requires a balanced panel, we additionally report the average treatment effect of the MTU introduction on log gross fuel prices using difference-in-difference analysis based on the full, unbalanced panel. Specifically, we estimate the following model:

$$Y_{it} = \beta_0 + \beta_1 MPD_{it} + \mu_i + \gamma_t + \epsilon_{it}, \tag{B.12}$$

where Y_{it} corresponds to the log gross fuel price at station i at date t and MPD_{it} is a dummy equal to one, if a fuel station i has to report its prices to the MTU at date t. This affects all fuel stations in Germany after the 1 October 2013. μ_i are fuel station fixed effects, and γ_t are date fixed effects.

Table B.2 reports the effects of the MTU introduction using Equation B.12. The outcome variable in all columns is logarithm of gross prices, and the estimation is based on data from 15 April 2013 to 31 March 2014. The results in Columns (1) and (2) of Table B.2 are based on the full, unbalanced panel. Columns (3) and (4) report estimates when we only use data on stations located within 20 to 100 km from the Franco-German border.

Table B.2 shows that the introduction of MPD led to a decline in prices of 3.0% to 3.1% for gasoline and 2.4% to 2.8% for diesel. The effects are economically and statistically significant, and, similarly to the results estimated via SDID, remain larger for gasoline.

Table B.2: Effect of MPD on the logarithm of gross prices

	Gasoline	Diesel	Gasoline	Diesel
	(1)	(2)	(3)	(4)
MPD	-0.030***	-0.024***	-0.031***	-0.028***
	(0.0002)	(0.0002)	(0.001)	(0.001)
Date FE	*Yes*	*Yes*	*Yes*	*Yes*
Station FE	*Yes*	*Yes*	*Yes*	*Yes*
Observations	4,110,958	4,706,894	357,816	387,949
Adjusted R^2	0.830	0.806	0.815	0.743

Notes: Columns (1) and (2) include estimates of the effect of MPD on log daily prices for gasoline and diesel, respectively, using all fuel stations in Germany and France. Columns (3) to (4) include the same estimates for a restricted sample of fuel stations 20 to 100 kilometers away from the Franco-German border. The observation periods goes from 15 April 2013 to 31 March 2014. Standard errors are clustered at the fuel station level and are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.3.2 Donut-SDID analysis

Figure B.8 illustrates the identification strategy for the Donut-SDID analysis graphically. To compare stations in economic regions that are as comparable as possible across countries, we restrict the panel to stations within 100 kilometers of the Franco-German border. Fuel stations that are less than 20 kilometers away from the Franco-German border are not considered, because these could be in direct competition to each other and so spillovers of the treatment effect could occur. This would threaten the stable unit treatment value assumption. Each point in Figure B.8 thus represents a fuel station, either in France or in Germany, which is 20 to 100 kilometers away from the border.

In Table B.3, we re-estimate the Donut-SDID regression for the analysis period 15 April 2013 until 31 March 2014 using different distances to the Franco-German border. We find that the results are robust to changing distance thresholds and the average effect of the MTU introduction is always larger for gasoline.

Figure B.8: Fuel stations 20 to 100 kilometers from the Franco-German border

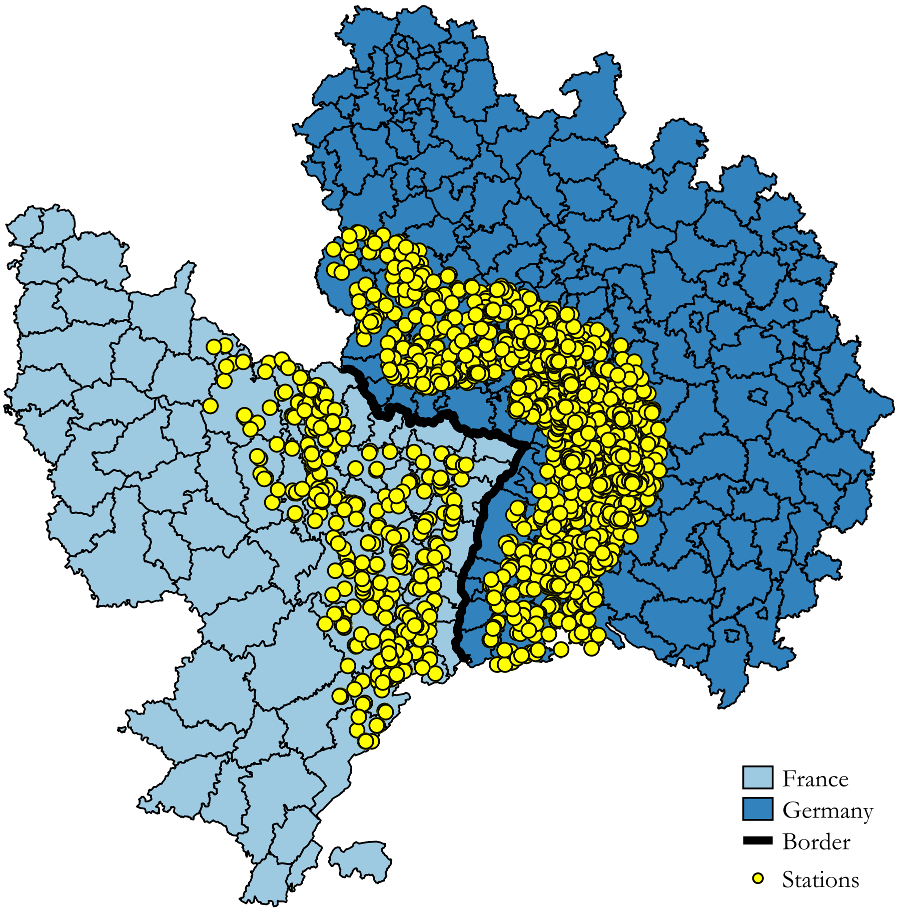

Notes: The thick, solid line represents the Franco-German border. Each point on the right of the border represents a fuel station in Germany, which is 20 to 100 kilometers away from the border. Each point on the left side of the border represents a fuel station in France, which is 20 to 100 kilometers away from the border. These are the fuel stations considered in our Donut-SDID analysis, when they have no missing weekly price observations.

Table B.3: Effect of MPD on the logarithm of gross prices using alternative donuts

	Gasoline	Diesel	Gasoline	Diesel	Gasoline	Diesel
	(1)	(2)	(3)	(4)	(5)	(6)
MPD	-0.031***	-0.021***	-0.029***	-0.020***	-0.028***	-0.020***
	(0.001)	(0.002)	(0.001)	(0.002)	(0.001)	(0.001)
95% CI	[-0.034, -0.028]	[-0.025, -0.018]	[-0.032, -0.026]	[-0.023, -0.017]	[-0.031, -0.026]	[-0.022, -0.018]
Week FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Station FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Observations	9,408	11,711	20,874	24,843	37,338	42,875

Columns (1) and (2) include estimates of the effect of MPD on log weekly prices for gasoline and diesel, respectively, using a restricted sample of fuel stations 20 to 40 kilometers away from the Franco-German border. Columns (3) and (4) include the same estimates for fuel stations 20 to 60 kilometers away from the border. Columns (5) and (6) include the same estimates for fuel stations 20 to 80 kilometers away from the border. The observation periods goes from 15 April 2013 to 31 March 2014. Standard errors are computed using the jackknife method and are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.3.3 Estimation with control for crude oil price

As discussed in Section 2.5, crude oil price experienced a sizable decline in the second half of 2014. The fluctuations in the price of crude oil could bias our estimates of the MTU effects if input costs were passed through differentially between stations in Germany and France. Even though we restrict our analysis to August 2014 in our main empirical specification, we additionally estimate the effect of the MTU introduction by directly allowing the differential pass-through of oil cost shocks between stations in Germany and France.

Table B.4 shows the effect of the MTU introduction on log gross weekly average gasoline and diesel price when we control for the indicator of stations in Germany interacted with the crude oil price at the port of Rotterdam. Columns (1) and (2) use the full balanced panel, and Columns (3) and (4) restrict the sample to stations located within 20 to 100 km from the Franco-German border. The effects are estimated via SDID, and all columns use data between 15 April 2013 and 31 March 2014. In addition to allowing for the differential pass-through of the input cost shocks between stations in Germany and France, we control for fuel station and time fixed effects.

Columns (1) and (2) in Table B.4 show that the introduction of the mandatory price disclosure led to the decrease in weekly average prices of 4.2% for gasoline and 1.8% for diesel. When the sample is restricted to the Donut-SDID, the corresponding estimates indicate a decline of 4.2% for gasoline and 2.4% for diesel. Overall, the magnitude of the MTU effect and its ranking with respect to the two fuel types remain robust to allowing for differential pass-through of the crude oil price between stations in Germany and France.

Table B.4: Effect of MPD on the logarithm of gross prices

	Gasoline	Diesel	Gasoline	Diesel
	(1)	(2)	(3)	(4)
MPD	-0.042***	-0.018***	-0.042***	-0.024***
	(0.011)	(0.003)	(0.008)	(0.001)
95% Confidence interval	[-0.064, -0.020]	[-0.023, -0.012]	[-0.057, -0.026]	[-0.027, -0.021]
Germany × crude oil price	*Yes*	*Yes*	*Yes*	*Yes*
Week FE	*Yes*	*Yes*	*Yes*	*Yes*
Station FE	*Yes*	*Yes*	*Yes*	*Yes*
Observations	632,884	751,219	49,539	55,517

Notes: Columns (1) and (2) include estimates of the effect of MPD on log weekly prices for gasoline and diesel, respectively, using all fuel stations in Germany and France. Columns (3) and (4) include the same estimates for a restricted sample of fuel stations 20 to 100 kilometers away from the Franco-German border. The observation periods goes from 15 April 2013 to 31 March 2014 and include a control for the interaction of an indicator for Germany with the crude oil price at the port of Rotterdam. Standard errors are computed using the jackknife method and are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.3.4 Effect of the MTU introduction on retail margins

Table B.5 shows the effects of the MTU introduction on retail margins, estimated using the SDID model in Equation 2.1. The outcome variable in all columns is weekly average retail margins, and the estimation is based on data from 15 April 2013 to 31 March 2014. All columns include fuel station and week fixed effects.

Results in Columns (1) and (2) show that the mandatory price disclosure led to the decrease in weekly average retail margins by 3.4 and 1.9 Eurocent for gasoline and diesel, respectively. In Columns (3) and (4), we restrict the analysis to stations within 20 to 100 km from the Franco-German border. Using this Donut-SDID, Columns (3) and (4) show that after the MTU introduction weekly average retail margins decline by 3.7 Eurocent for gasoline and 2.3 Eurocent for diesel. The effect of the MTU introduction is statistically and economically significant, and is larger for gasoline.

Table B.5: Effect of MPD on retail margins

	Gasoline	Diesel	Gasoline	Diesel
	(1)	(2)	(3)	(4)
MPD	-3.357***	-1.930***	-3.663***	-2.286***
	(0.071)	(0.043)	(0.121)	(0.111)
95% Confidence interval	[-3.495, -3.218]	[-2.014, -1.846]	[-3.900, -3.426]	[-2.502, -2.069]
Week FE	*Yes*	*Yes*	*Yes*	*Yes*
Station FE	*Yes*	*Yes*	*Yes*	*Yes*
Observations	632,884	751,219	49,539	55,517
Mean retail margin	8.36	10.77	8.51	11.20

Notes: Columns (1) and (2) include estimates of the effect of MPD on weekly average retail margins for gasoline and diesel, respectively, using all fuel stations in Germany and France. Columns (3) and (4) include the same estimates for a restricted sample of fuel stations 20 to 100 kilometers away from the Franco-German border. The observation periods goes from 15 April 2013 to 31 March 2014. Standard errors are computed using the jackknife method and are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.3.5 Local monopolists as a control group

Driving to another fuel station is costly and hence retail fuel markets are usually segmented geographically. We define local markets as driving distance catchment areas around a focal station. We assume that stations that do not face competition from another station in their catchment area act as local monopolists. Like in the analysis of Albæk et al. (1997) for the cement industry, these local monopolists are unaffected by increasing transparency and can therefore serve as a control group.

In Table B.6, we report the results of an estimation strategy in which we analyse the effect of the MTU on logarithm of gross prices of fuel stations in Germany for gasoline and diesel. We compare fuel stations in Germany, which have at least one competing fuel station in their catchment area to fuel stations that are local monopolists, and we estimate the effects via difference-in-differences approach. Only fuel stations that are of a different brand are considered as competitors. Whereas we consider local monopolists as untreated by the introduction of the MTU, because consumers have no alternative in the vicinity and can thus not act upon the new information, stations that have a competitor in their market are considered treated. In Columns (1) and (4), we define a local monopolist as not having any other station within a 1 kilometer radius. We find a treatment effect of 0.04 to 0.1 percent, however, according to this definition 64% of

Table B.6: Effect of MPD on the logarithm of gross prices (local monopolies)

	Gasoline			Diesel		
	(1)	(2)	(3)	(4)	(5)	(6)
MPD	-0.001***	-0.002***	-0.002***	-0.0004	-0.001***	-0.002***
	(0.0002)	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0004)
Date FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Station FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Observations	2,619,823	1,589,155	1,301,738	2,645,827	1,605,201	1,315,465
Share local monopolists	64.3%	42.3%	29.4%	64.3%	42.3%	29.4%
Adjusted R^2	0.813	0.815	0.815	0.662	0.669	0.669

Columns (1) and (4) include estimates of the effect of MPD on log prices for gasoline and diesel, respectively, using fuel stations that are local monopolists within 1 kilometer as the control group and all other stations as the treatment group. Columns (2) and (5) repeat the same analyses for a 3 kilometer radius. Columns (5) and (6) repeat the same analyses for a 5 kilometer radius. The observation periods goes from 15 April 2013 to 31 March 2014. Standard errors are clustered at the fuel station level and are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

fuel stations in Germany are local monopolists. We thus consider broader markets in Columns (2) and (3) for gasoline and in Columns (5) and (6) for diesel. In Columns (2) and (5), we define local monopolists as not having a competing station within a 3 kilometers radius. We drop all fuel stations with a competitor within a 3 kilometers radius, but without a competitor within a 1 kilometer radius from the control group, as these are local monopolists according to the market definition in Column (1) and (4). We find a treatment effect of 0.1 to 0.2 percent using 3 kilometers catchment areas. In Columns (3) and (6), we repeat this analysis for 5 kilometers catchment area and find a similar treatment effect to Columns (2) and (5). Overall, our results are consistent with Lemus and Luco (2021), who find that mandatory price disclosure reduced the time to reach a new equilibrium for oligopoly markets, but not for local monopolies.

Overall, the average effect of the MTU that we find using this specification is consistent with our estimates for the average effect of the MTU using France as a control group. The treatment effect of the MTU remains larger for the ex ante less informed consumer group. We are likely to underestimate the treatment effect using the local monopolist identification strategy, since consumers in monopoly markets are likely also partially treated by the MTU. It therefore makes sense that the magnitude of the effect that we find using local monopolists is smaller than when comparing gross fuel prices in Germany and France.

B.3.6 Difference-in-differences analysis: European countries as a control

To test the validity of France as a counterfactual, we also estimate the effect of the MTU introduction on fuel prices in Germany using 26 other European countries as a control group.[7] To do so, we use information on country-level weekly average net gasoline and diesel prices that are reported by the European Commission in the Weekly Oil Bulletin.

Table B.7 shows the effects of the MTU introduction on the logarithm of net gasoline and diesel prices, using a difference-in-differences strategy. As in our main analysis, the estimation is based on data between 15 April 2013 and 31 March 2014 and we control for week and country fixed effects in all columns. In Columns (3) and (4), we additionally control for the crude oil price at the port of Rotterdam interacted with country indicators, which allows for differential pass-through of oil cost shocks across countries.

Table B.7 shows that when we use other European countries as a control, the MTU introduction led to a decline of 3.0% to 3.3% for gasoline and 1.5% to 1.8% for diesel. The ranking of the effects with respect to the fuel types and their magnitude remain robust to using this alternative control group.

Table B.7: Effect of MPD on the logarithm of net prices

	Gasoline	Diesel	Gasoline	Diesel
	(1)	(2)	(3)	(4)
MPD	-0.033***	-0.018***	-0.030***	-0.015***
	(0.003)	(0.003)	(0.006)	(0.005)
Country × crude oil price	*No*	*No*	*Yes*	*Yes*
Date FE	*Yes*	*Yes*	*Yes*	*Yes*
Country FE	*Yes*	*Yes*	*Yes*	*Yes*
Observations	1,258	1,258	1,258	1,258
Adjusted R^2	0.868	0.836	0.879	0.860

Notes: Columns (1) and (2) include estimates of the effect of MPD on log net prices for gasoline and diesel, respectively, using Germany as a treatment group and all other EU countries as a control. Columns (3) to (4) include additional interactions between the crude oil price and an indicator variable for each country. Robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[7] Austria, Belgium, Bulgaria, Croatia, Cyprus, Czechia, Denmark, Estonia, Finland, France, Greece, Hungary, Ireland, Italy, Latvia, Lithuania, Luxembourg, Malta, Netherlands, Poland, Portugal, Romania, Slovakia, Slovenia, Spain, and Sweden form the control group.

Appendix C

Appendix to Chapter 3

C.1 Appendix to Section 3.2: Theoretical Model

C.1.1 Stage 2: Equilibrium price distribution

Lemma 3.4. *There is no pure strategy Nash equilibrium in prices in the second stage if $N \geq 2$ sellers entered the market in the first stage.*

Proof. Suppose that all N sellers chose to set the same price strictly above the constant marginal cost c. Then, all sellers receive a share $\frac{1}{N}$ of shoppers and non-shoppers. This cannot be a stable equilibrium because all sellers have an incentive to marginally undercut the common price and attract all shoppers. All sellers setting the price at the constant marginal cost c can also not be a stable equilibrium because sellers can profitably deviate by setting a higher price and only serving uninformed consumers.

Finally, suppose that sellers play pure strategies in which at least one seller chooses a lower price than the other sellers. This seller then serves all shoppers, as well as its share of uninformed consumers. This cannot be an equilibrium because the lowest price seller can always marginally increase its price without losing the shoppers to another seller.

□

Lemma 3.5. *There are no mass points in the equilibrium pricing strategies.*

Proof. Suppose that any price was played with positive probability. This would mean that there is a positive probability of a tie for shoppers at that price. This cannot be an equilibrium because a seller could profitably deviate from that strategy by charging a marginally lower price with the same probability and capture all shoppers in that case.[1]

□

Lemma 3.6. *There is a unique symmetric mixed strategy Nash equilibrium where all sellers draw a price from the distribution $F(p_i)$ on the interval $[\underline{p}, p_r]$, where*

$$\underline{p} = \frac{p_r}{\frac{\phi N}{1-\phi} + 1} + c\frac{1+\tau}{1 + \frac{1-\phi}{\phi N}}\,,$$

$$p_r = \begin{cases} E[p] + s & \text{if } E[p] + s < \upsilon \\ \upsilon & \text{otherwise} \end{cases}\,, \text{ and}$$

$$F(p_i) = 1 - \left(\frac{p_r - p_i}{p_i - c(1+\tau)}\frac{1-\phi}{N\phi}\right)^{\frac{1}{N-1}}\,.$$

The expected second stage profits (i.e. excluding the fixed and sunk cost of entry) of a seller are

$$E[\pi_i] = \left(\frac{p_r}{1+\tau} - c\right)\frac{1-\phi}{N}M\,.$$

The expected price is

$$E[p] = \underline{p} + \left(\frac{1-\phi}{N\phi}\right)^{\frac{1}{N-1}} \int_{\underline{p}}^{p_r} \left(\frac{p_r - p}{p - c(1+\tau)}\right)^{\frac{1}{N-1}} dp\,.$$

The expected minimum price is

$$E[p_{min}] = \frac{1-\phi}{N\phi}[p_r - E[p] + (p_r - c(1+\tau))c(1+\tau)\int_{\underline{p}}^{p_r} (p - c(1+\tau))^2 F(p)dp]\,.$$

Proof. We begin by deriving the reservation price of non-shoppers, p_r. Non-shoppers can search sequentially at an incremental search cost s. A necessary condition for search to oc-

[1] For a more detailed proof, see Varian (1980).

cur, irrespective of the price initially drawn, is that the sum of the expected price at the next draw and the sequential search cost does not exceed the valuation of the good. If this is fulfilled, non-shoppers with a particular first draw of p search as long as the expected gain of searching is greater than s. Thus, search occurs so long as

$$s < p - \int_{\underline{p}}^{p_{max}} pf(p)dp\,. \tag{C.1}$$

The reservation price of non-shoppers is such that they are exactly indifferent between continuing to search and buying at that price. No consumer buys at a price above the reservation price of non-shoppers. At the same time, sellers that do not sell to shoppers want to charge non-shoppers their reservation price. The maximum of the support of prices from which sellers draw in equilibrium is therefore $p_{max} = p_r$. Following Stahl (1989), a consistent reservation price $p_r \leq \upsilon$ must therefore satisfy

$$H(p_r; \phi, N, s) \equiv p_r - \int_{\underline{p}}^{p_r} pf(p)dp - s = 0\,. \tag{C.2}$$

Stahl (1989) shows that H has a unique root or none at all for a general class of demand functions which include linear demand. Thus, in this case there is no other symmetric mixed strategy Nash equilibrium of the pricing game.

As explained before, if the sum of the expected price at the next draw and the sequential search cost exceed the valuation υ, search never occurs. In this case, the reservation price is simply the valuation of the good. The equilibrium reservation price of non-shoppers is thus

$$p_r = \begin{cases} E[p] + s & \text{if } E[p] + s < \upsilon \\ \upsilon & \text{otherwise} \end{cases}\,. \tag{C.3}$$

Since it is never an equilibrium strategy for any seller to choose a price above the reservation price of non-shoppers, there is no sequential search in equilibrium.

Next, we turn to finding the lowest price sellers may draw in equilibrium, $\underline{p}$. Any price drawn with positive probability in equilibrium should yield the same expected profit. The expected profit of setting the price at $\underline{p}$ therefore has to equal the expected profit of setting the reservation price, thus

$$E[\pi(\underline{p})] = E[\pi(p_r)]\,. \tag{C.4}$$

Since we established that there are no mass points in the equilibrium pricing strategies, the probability of a tie is zero. A seller setting its price at $\underline{p}$ will therefore attract all shoppers and its share of non-shoppers that randomly visit its store. A seller setting its price at p_r will never attract any shoppers and only serve its share of non-shoppers. We can therefore re-write the expected profits as

$$(\frac{\underline{p}}{1+\tau} - c)(\phi + \frac{1-\phi}{N})M = (\frac{p_r}{1+\tau} - c)\frac{1-\phi}{N}M\,. \tag{C.5}$$

We can simplify this expression and re-arrange it to yield an expression for the lowest price sellers may draw in equilibrium

$$\underline{p} = \frac{p_r}{\frac{\phi N}{1-\phi} + 1} + c\frac{1+\tau}{1 + \frac{1-\phi}{\phi N}}\,. \tag{C.6}$$

The last ingredient necessary to characterize the distribution from which sellers draw prices in equilibrium is the density function of the distribution. To derive the density function, we can again exploit the equiprofit condition that

$$E[\pi(p_i)] = E[\pi(p_r)] \qquad \forall \quad p_i \in [\underline{p}, p_r]\,. \tag{C.7}$$

With probability $(1 - F(p_i))^{N-1}$ a seller choosing price p_i has the lowest price of all N sellers and will thus sell to all shoppers and its share of non-shoppers. With probability $1 - (1 - F(p_i))^{N-1}$ there is another seller charging a lower price and thus seller i only sells to its share of non-shoppers. Expected profits can be written as

$$\begin{aligned}(\frac{p_i}{1+\tau} - c)(\phi + \frac{1-\phi}{N})(1 - F(p_i))^{N-1}M + (\frac{p_i}{1+\tau} - c)(\frac{1-\phi}{N})(1 - (1 - F(p_i))^{N-1})M = \\ (\frac{p_r}{1+\tau} - c)\frac{1-\phi}{N}M\,.\end{aligned} \tag{C.8}$$

We can solve this equation for the equilibrium density function according to which each seller i draws its prices from the support $[\underline{p}, p_r]$

$$F(p_i) = 1 - \left(\frac{p_r - p_i}{p_i - c(1+\tau)}\frac{1-\phi}{N\phi}\right)^{\frac{1}{N-1}}. \tag{C.9}$$

For a given number of entrants N and a given set of exogenous parameters, Equations C.3, C.6 and C.9 uniquely identify the symmetric mixed strategy Nash equilibrium in prices.

We can derive the expected second stage profit of each seller i in this equilibrium. Since the expected profit of each seller in the symmetric equilibrium is the same for any price chosen with positive probability, the expected profit of seller i drawing a price from the equilibrium price distribution is

$$E[\pi_i] = E[\pi(p_r)] = \left(\frac{p_r}{1+\tau} - c\right)\frac{1-\phi}{N}M. \tag{C.10}$$

Finally, we can derive the expected prices paid by non-shoppers and shoppers, namely the expected price and the expected minimum price.

The expected price is

$$E[p] = \int_{\underline{p}}^{p_r} p f(p) dp = p_r - \int_{\underline{p}}^{p_r} F(p) dp, \tag{C.11}$$

after integrating by parts. We can then insert the equilibrium price distribution and simplify the expression, which yields

$$E[p] = \underline{p} + \left(\frac{1-\phi}{N\phi}\right)^{\frac{1}{N-1}} \int_{\underline{p}}^{p_r} \left(\frac{p_r - p}{p - c(1+\tau)}\right)^{\frac{1}{N-1}} dp.$$

To derive the expected minimum price we begin by setting up the probability density function of the minimum price. This can be written as

$$f_{min}(p) = N(1 - F(p))^{N-1} f(p). \tag{C.12}$$

After inserting $F(p)$ and simplifying the expression, this yields

$$f_{min}(p) = \frac{p_r - p}{p - c(1+\tau)}\frac{1-\phi}{\phi} f(p). \tag{C.13}$$

The expected minimum price is then

$$E[p_{min}] = \int_{\underline{p}}^{p_r} p f_{min}(p)dp = \int_{\underline{p}}^{p_r} p \frac{p_r - p}{p - c(1+\tau)} \frac{1-\phi}{N\phi} f(p)dp\,. \tag{C.14}$$

After adding and subtracting $c(1+\tau)$ in the numerator of the first fraction and further simplifications, we get that

$$E[p_{min}] = \frac{1-\phi}{\phi}[\int_{\underline{p}}^{p_r} p \frac{p_r - c(1+\tau)}{p - c(1+\tau)} f(p)dp - E[p]]\,.$$

Finally, we can use integration by parts and rearrange terms to get the following expression for the expected minimum price:

$$E[p_{min}] = \frac{1-\phi}{\phi}[p_r - E[p] + (p_r - c(1+\tau))c(1+\tau)\int_{\underline{p}}^{p_r} \frac{1}{(p - c(1+\tau))^2} F(p)dp]\,.$$

□

C.1.2 Stage 1: Equilibrium entry

Lemma 3.7. *Under free entry and with a sufficiently large number of symmetric potential entrants, such that the number of potential entrants always exceeds the number of firms that can be supported by the market, in equilibrium an integer number of N^* firms enter the market, such that*

$$(\frac{p_r}{1+\tau} - c)\frac{1-\phi}{F}M - 1 < N^* \leq (\frac{p_r}{1+\tau} - c)\frac{1-\phi}{F}M\,.$$

Proof. Suppose that there is a large number of symmetric firms which are sequentially asked whether they want to enter the market at the fixed and sunk cost F, knowing how many firms decided to enter before them. Firms are going to decide to enter the market so long as their expected second stage profits are at least as high as the fixed and sunk cost F. In equilibrium, the first N firms asked to enter will accept and firm $N+1$ and all firms following thereafter will reject if, and only if, the expected second stage profits of firms $1, ..., N$ are equal to F or higher and the expected second stage profits of firm $N+1$ are lower than F.

To derive the condition for the equilibrium number of firms entering the market, we use the expression for the expected second stage profit of firm i in Equation C.10. We calculate the expected second stage profits with N and $N+1$ entrants and re-arrange these to yield a condition on the equilibrium number of entrants. In equilibrium, an integer number of N firms enter the market, such that

$$(\frac{p_r}{1+\tau} - c)\frac{1-\phi}{F}M - 1 < N^* \leq (\frac{p_r}{1+\tau} - c)\frac{1-\phi}{F}M \,. \tag{C.15}$$

□

C.1.3 Pass-through of marginal costs

Next, we analyze how marginal costs or per unit taxes are passed through to consumers. Many of the results and intuitions regarding ad-valorem taxes directly translate to marginal costs (or per unit taxes).

Proposition 3.4. *With $0 < \phi < 1$, for any $\hat{c} > c$ the minimum element of the support of the equilibrium pricing strategy $\underline{\hat{p}} > \underline{p}$ and the Nash equilibrium pricing strategy with c first-order stochastically dominates (FOSD) the pricing strategy with $\hat{c}$, i.e. $\hat{F}(p) \leq F(p) \quad \forall p$.*

Analogous to the explanation for ad-valorem taxes, this means that if the share of shoppers is strictly positive, an increase in c leads to a shift in the support of the prices from which sellers draw in equilibrium towards higher prices. Furthermore, for each price on the equilibrium pricing support the likelihood that a drawn price is below said price decreases if marginal costs increase from c to $\hat{c}$.

As for the pass-through of ad-valorem taxes, the pass-through of marginal costs converges to zero as the share of shoppers converges to zero.

Since the minimum element of the support of prices and the density function monotonously move towards higher prices, other moments of interest, such as the expected price $E[p]$ and the expected minimum price $E[p_{min}]$ also increase.

We now turn to analyzing how the pass-through rate of marginal costs or per unit taxes vary with the price sensitivity of consumers and the number of active sellers.

Figure C.1: Marginal cost pass-through to the equilibrium pricing strategy

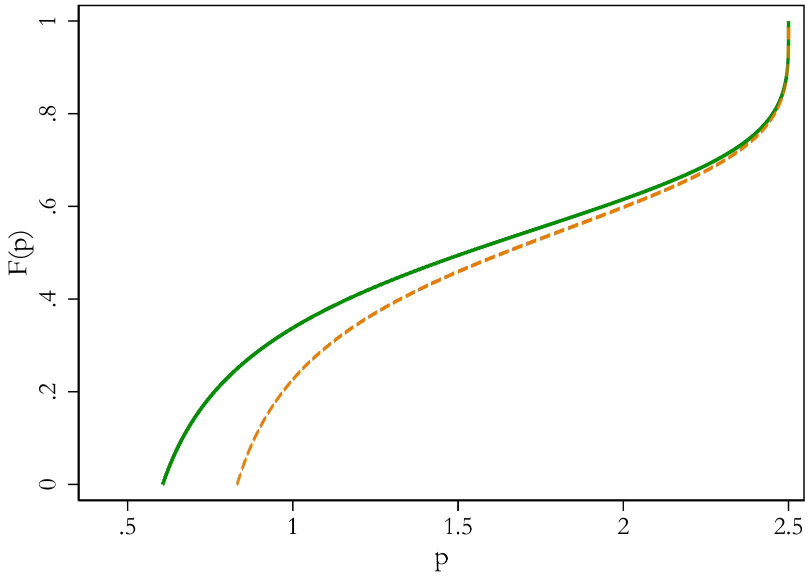

Note: The Figure shows simulation results of how the distribution from which sellers draw prices in the symmetric Nash equilibrium changes if marginal costs increase from c to $\hat{c}$. Parameter values: $\upsilon = 2.5$, $s = 0.75$, $\tau = 0.2$, $c = 0.4$, and $\hat{c} = 0.6$.

Proposition 3.5. *If the share of shoppers $\phi = 0$, marginal cost pass-through $\rho_c = 0$. If $\phi = 1$, there is full pass-through, i.e. $\rho_c = 1 + \tau$. As $\phi \to 1$, the pass-through rate $\rho_c \to 1 + \tau$.*

We can begin by looking at the cases when there are no shoppers and when there are only shoppers. If there are no shoppers, all sellers choose the monopoly price and pass-through of marginal costs is zero. If all consumers are shoppers, there is full pass-through of marginal costs or per unit taxes.[2]

For all values of ϕ between zero and one, we can show that the pass-through rate of marginal costs to the lower bound of the equilibrium price strategy is strictly increasing in the share of shoppers. We can also show that the rate at which an increase in marginal costs from c to $\hat{c}$ reduces the probability that a drawn price is below a particular price p, i.e. from $F(p)$ to $\hat{F}(p)$, strictly increases in the share of shoppers. Thus, the pass-through rate of marginal costs increases in the share of shoppers.

[2]Although an increase in the marginal cost from c to $\hat{c}$ leads to an increase of $(\hat{c} - c)(1 + \tau)$ to consumers, we would still classify this case as full pass-through (instead of over-shifting) since the producer price only increases by $\hat{c} - c$.

Let us now consider how the pass-through of marginal costs varies with the number of active sellers. As we will see, all of our results and intuitions with respect to ad-valorem tax pass-through extend to marginal costs.

Proposition 3.6. *With* $0 < \phi < 1$*, as* $N \to \infty$ *the pass-through of* c *to the minimum element of the equilibrium price support converges to full pass-through, i.e.* $\rho_{c,\underline{p}} \to 1 + \tau$.

As the number of sellers increases, competition for shoppers becomes fiercer and the pass-through rate of marginal costs to $\underline{p}$ increases.

Furthermore, we also expect pass-through of marginal costs to $E[p]$ to first increase and then decrease, whereas pass-through to $E[p_{min}]$ should always increase as $N \to \infty$. The same reasoning as laid out for ad-valorem taxes applies.

Figure C.2: Pass-through of c to $E[p]$

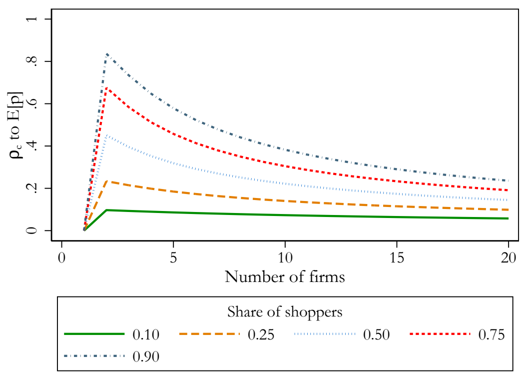

Parameter values: $\upsilon = 2.5$, $\tau = 0.2$, $c = 0.4$ and $\hat{c} = 0.44$.

Figure C.3: Pass-through of c to $E[p_{min}]$

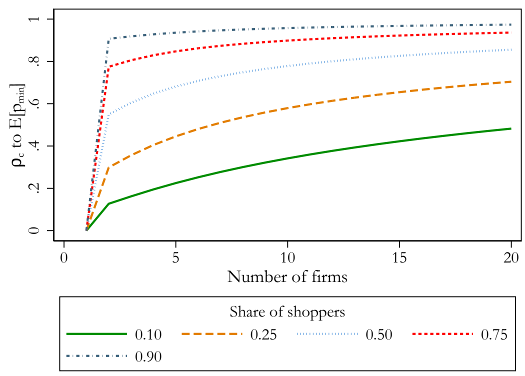

Parameter values: $\upsilon = 2.5$, $\tau = 0.2$, $c = 0.4$ and $\hat{c} = 0.44$.

The simulation results in Figures C.2 and C.3 are very similar to those for ad-valorem tax pass-through. As N increases, pass-through of c to the expected price first increases and then decreases. Pass-through to the expected minimum price always increases.

Finally, we consider how c is passed through to the expected average price paid by consumers in the markets.

Figure C.4: ρ_c to $E[\phi p_{min} + (1-\phi)p]$, $p_r = \upsilon$

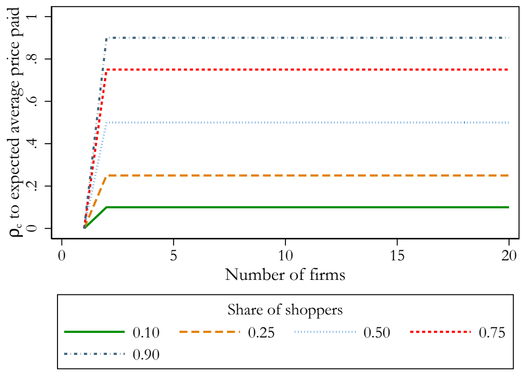

Parameter values: $\upsilon = 2.5$, $\tau = 0.2$, $c = 0.4$ and $\hat{c} = 0.44$.

Figure C.5: ρ_c to $E[\phi p_{min} + (1-\phi)p]$, p_r endogenous

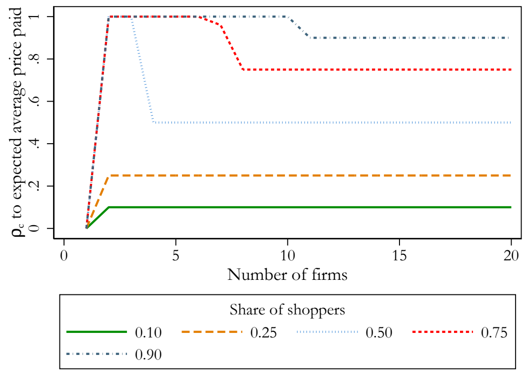

Parameter values: $\upsilon = 2.5$, $s = 0.75$, $\tau = 0.2$, $c = 0.4$ and $\hat{c} = 0.44$.

The simulation in Figure C.4 shows that when sequential search costs are so high that $p_r = \upsilon$, pass-through of marginal costs first increases in N and then stays constant, because the decrease in pass-through to $E[p]$ and the increase in pass-through to $E[p_{min}]$ cancel each other out. Figure C.5 shows that if sequential search costs s are sufficiently low such that p_r is endogenous, pass-through to the expected average price paid first increases in N, then decreases in N and, as $p_r \to \upsilon$ when $N > 2$ and $N \to \infty$, marginal cost pass-through remains constant when N is sufficiently large.

C.1.4 Proof of Propositions

Proof of Proposition 1. First, we assess the pass-through of τ to $\underline{p}$ if $0 < \phi < 1$.[3] Taking the first derivative with respect to τ, we find that

$$\frac{\partial \underline{p}}{\partial \tau} = c\left(1 + \frac{1-\phi}{\phi N}\right)^{-1} > 0\,.$$

Thus, with $0 < \phi < 1$, pass-through of τ to the minimum element of the support of the equilibrium pricing strategy is strictly positive.

[3] $\underline{p}$ is not defined for $\phi = 0$ or $\phi = 1$.

Next, we assess the pass-through of the ad-valorem tax to $F(p)$ if $0 < \phi < 1$. Taking the first derivative with respect to τ, we find that

$$\frac{\partial F(p)}{\partial \tau} = -(\frac{1-\phi}{\phi N})^{\frac{1}{N-1}} \frac{1}{N-1} (\frac{p_r - p}{p - c(1+\tau)})^{\frac{1}{N-1}} \frac{c}{p - c(1+\tau)} < 0 \,.$$

Thus, with $0 < \phi < 1$, for any $\hat{\tau} > \tau$ $\hat{F}(p) \leq F(p) \quad \forall p \in [\underline{p}, p_r]$.

□

Proof of Proposition 2. Let us begin by examining the case where $\phi = 0$. In this case, the price equilibrium is a degenerate distribution at the monopoly price, with $\underline{p} = p_r = \upsilon$. An increase in τ is fully absorbed by sellers, since these already fully extract the entire valuation from consumers.

Next, we examine the case where $\phi = 1$. In this case, the price equilibrium is a degenerate distribution at the perfectly competitive price, with $\underline{p} = p_r = c(1+\tau)$. An increase in the ad-valorem tax τ is now fully passed through to consumers, as sellers already operate at zero profits and absorbing some of the marginal cost would mean that they would be making losses.

Finally, we study the case where $0 < \phi < 1$.

Let us begin by analyzing how the pass-through rate changes with ϕ

$$\frac{\partial^2 \underline{p}}{\partial \tau \partial \phi} = c(1 + \frac{1-\phi}{\phi N})^{-2} \frac{1}{\phi^2 N} > 0 \,.$$

Thus, with $0 < \phi < 1$, the pass-through of τ to the minimum element of the support of the equilibrium pricing strategy strictly increases in ϕ.

Next, we consider how the effect of an increase from τ to $\hat{\tau}$ on the cumulative density function of the pricing strategy changes if ϕ increases

$$\frac{\partial^2 F(p)}{\partial \tau \partial \phi} = (\frac{1}{N-1})^2 (\frac{p_r - p}{p - c(1+\tau)})^{\frac{1}{N-1}} \frac{c}{p - c(1+\tau)} (\frac{1-\phi}{\phi N})^{\frac{1}{N-1}-1} \frac{1}{\phi^2 N} > 0 \,.$$

Thus, for higher ϕ, an increase from τ to $\hat{\tau}$ decreases the probability that prices are below a certain p more strongly.

□

Proof of Proposition 3. To see how the pass-through rate of a value-added tax τ to the minimum element of the support varies with N, we study the limit to which the pass-through rate converges as $N \to \infty$. We find that

$$\lim_{N\to\infty} \rho_{\tau,\underline{p}} = \lim_{N\to\infty} \frac{\partial \underline{p}}{\partial \tau} \cdot \frac{1+\tau}{\underline{p}} = \frac{c(1+\tau)}{c(1+\tau)} = 1\,.$$

Thus, with $N \to \infty$, pass-through of a value-added tax to the minimum element of the support of the equilibrium pricing strategy converges to full pass-through.

□

Proof of Proposition 4. We begin by assessing the pass-through of marginal costs to $\underline{p}$ if $0 < \phi < 1$. Taking the first derivative with respect to c, we find that

$$\frac{\partial \underline{p}}{\partial c} = (1+\tau)(1 + \frac{1-\phi}{\phi N})^{-1} > 0\,.$$

Thus, with $0 < \phi < 1$, pass-through of marginal costs to the minimum element of the support of the equilibrium pricing strategy is strictly positive.

Next, we assess the pass-through of marginal costs to $F(p)$ if $0 < \phi < 1$. Taking the first derivative with respect to c, we find that

$$\frac{\partial F(p)}{\partial c} = -(\frac{1-\phi}{\phi N})^{\frac{1}{N-1}} \frac{1}{N-1} (\frac{p_r - p}{p - c(1+\tau)})^{\frac{1}{N-1}} \frac{1+\tau}{p - c(1+\tau)} < 0\,.$$

Thus, with $0 < \phi < 1$, for any $\hat{c} > c$, $\hat{F}(p) \leq F(p) \quad \forall p \in [\underline{p}, p_r]$.

□

Proof of Proposition 5. Again, we begin by examining the case where $\phi = 0$. In this case, the price equilibrium is a degenerate distribution at the monopoly price, with $\underline{p} = p_r = \upsilon$. An increase in marginal costs is fully absorbed by sellers, since these already fully extract the entire valuation from consumers.

Next, we examine the case where $\phi = 1$. In this case, the price equilibrium is a degenerate distribution at the perfectly competitive price, with $\underline{p} = p_r = c(1+\tau)$. An increase in c is now fully passed through to consumers.[4]

Finally, we study the case where $0 < \phi < 1$.

Let us begin by analyzing how the pass-through rate changes with ϕ

$$\frac{\partial^2 \underline{p}}{\partial c \partial \phi} = (1+\tau)(1+\frac{1-\phi}{\phi N})^{-2}\frac{1}{\phi^2 N} > 0\,.$$

Thus, with $0 < \phi < 1$, the pass-through of c to the minimum element of the support of the equilibrium pricing strategy strictly increases in ϕ.

Next, we consider how the effect of an increase from c to $\hat{c}$ on the cumulative density function of the pricing strategy changes if ϕ increases

$$\frac{\partial^2 F(p)}{\partial c \partial \phi} = (\frac{1}{N-1})^2(\frac{p_r - p}{p - c(1+\tau)})^{\frac{1}{N-1}}\frac{1+\tau}{p-c(1+\tau)}(\frac{1-\phi}{\phi N})^{\frac{1}{N-1}-1}\frac{1}{\phi^2 N} > 0\,.$$

Thus, for higher ϕ, an increase from c to $\hat{c}$ decreases the probability that prices are below a certain p more strongly.

□

Proof of Proposition 6. To see how the pass-through rate of marginal costs to the minimum element of the support varies with N, we study the limit to which the pass-through rate converges as $N \to \infty$. We find that

$$\lim_{N\to\infty} \rho_{c,\underline{p}} = \lim_{N\to\infty} \rho_{c,\underline{p}}(1+\tau)(1+\frac{1-\phi}{\phi N})^{-1} = 1+\tau\,.$$

Thus, with $N \to \infty$, pass-through of marginal costs to the minimum element of the support of the equilibrium pricing strategy converges to full pass-through.

□

[4] Although an increase in the marginal cost from c to $\hat{c}$ leads to an increase of $(\hat{c}-c)(1+\tau)$ to consumers, we would still classify this case as full pass-through (instead of over-shifting) since the producer price only increases by $\hat{c} - c$.

C.1.5 Allowing for sequentially searching non-shoppers

In this section, we simulate how the pass-through of marginal costs and ad-valorem taxes to the expected price and the expected minimum price vary with the share of shoppers and the number of sellers, if we allow non-shoppers to search sequentially. We find that the qualitative results remain unchanged to a situation where non-shoppers cannot search sequentially.

Marginal cost pass-through

Figure C.6: Marginal cost pass-through to the expected price

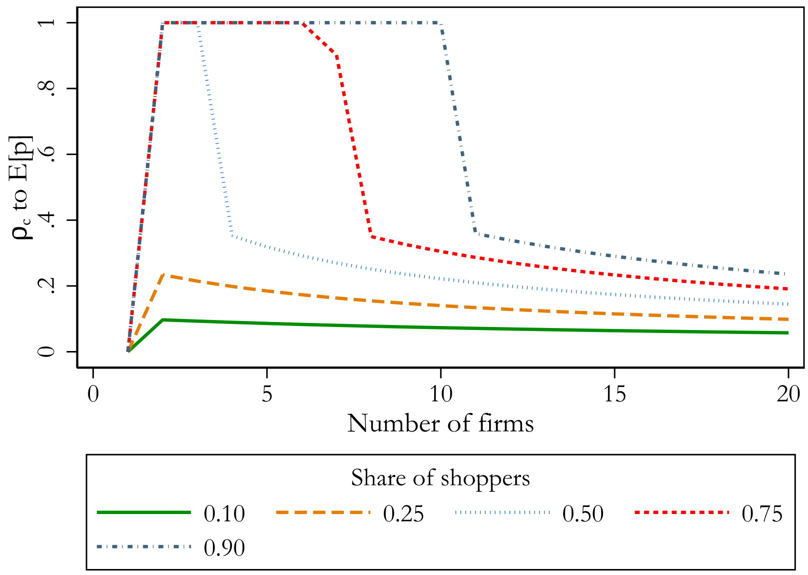

Note: The Figure shows simulation results of how the pass-through of marginal costs to the expected price varies with the share of shoppers and the number of active sellers. We fix the following parameter values for these simulations: $\upsilon = 2.5$, $s = 0.75$, $\tau = 0.2$, $c = 0.4$ and $\hat{c} = 0.44$.

The higher the share of shoppers, the higher is the pass-through rate of marginal costs to the expected price. For a given share of shoppers, marginal cost pass-through to the expected price first increases and then decreases in the number of sellers.

Figure C.7: Marginal cost pass-through to the expected minimum price

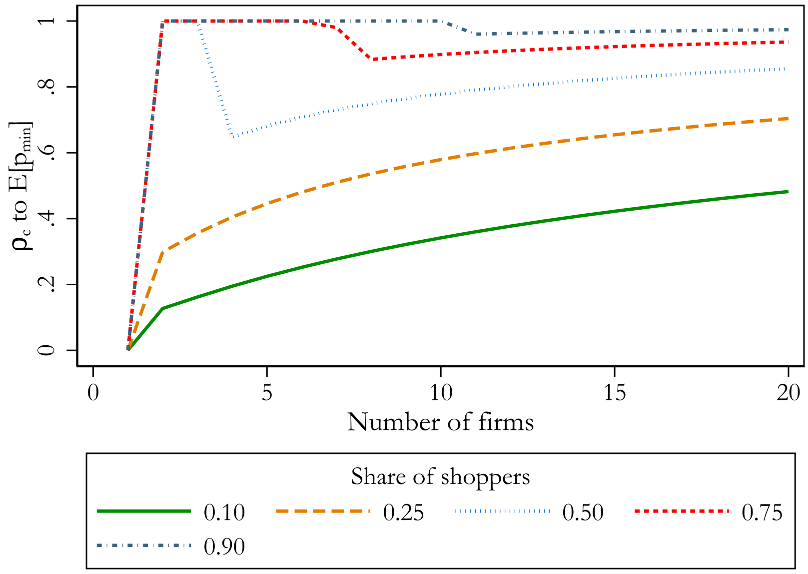

Note: The Figure shows simulation results of how the pass-through of marginal costs to the expected minimum price varies with the share of shoppers and the number of active sellers. We fix the following parameter values for these simulations: $\upsilon = 2.5$, $s = 0.75$, $\tau = 0.2$, $c = 0.4$ and $\hat{c} = 0.44$.

The higher the share of shoppers, the higher is the pass-through rate of marginal costs to the expected minimum price. For sufficiently low shares of shoppers and holding the share of shoppers fixed, marginal cost pass-through to the expected minimum price increases in the share of shoppers. This is as in the case without sequentially searching non-shoppers. For sufficiently high shares of shoppers, the pass-through rate first increases in the number of sellers, then decreases and then increases again. This is different to when we do not allow for sequentially searching non-shoppers.

Ad-valorem tax pass-through

Figure C.8: Ad-valorem tax pass-through to the expected price

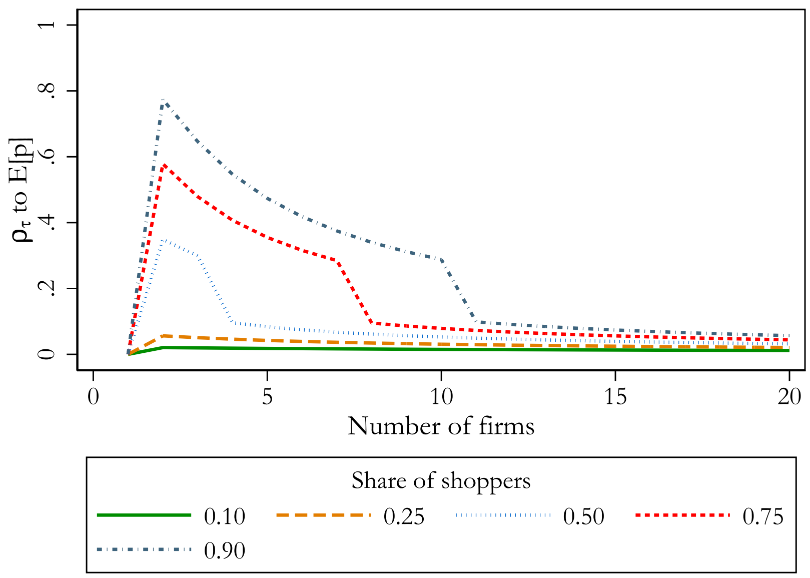

Note: The Figure shows simulation results of how the pass-through of an ad-valorem tax to the expected price varies with the share of shoppers and the number of active sellers. We fix the following parameter values for these simulations: $\upsilon = 2.5$, $s = 0.75$, $c = 0.4$, $\tau = 0.2$ and $\hat{\tau} = 0.22$.

The higher the share of shoppers, the higher is the pass-through rate of an ad-valorem tax to the expected price. For a given share of shoppers, ad-valorem tax pass-through to the expected price first increases and then decreases in the number of sellers.

Figure C.9: Ad-valorem tax pass-through to the expected minimum price

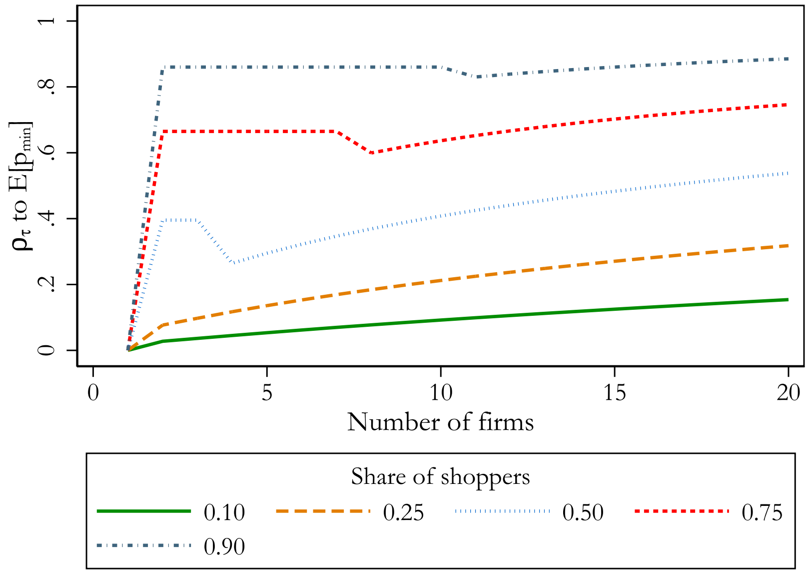

Note: The Figure shows simulation results of how the pass-through of an ad-valorem tax to the expected minimum price varies with the share of shoppers and the number of active sellers. We fix the following parameter values for these simulations: $\upsilon = 2.5$, $s = 0.75$, $c = 0.4$, $\tau = 0.2$ and $\hat{\tau} = 0.22$.

The higher the share of shoppers, the higher is the pass-through rate of an ad-valorem tax to the expected minimum price. For sufficiently low shares of shoppers and holding the share of shoppers fixed, ad-valorem tax pass-through to the expected minimum price increases in the share of shoppers. This is as in the case without sequentially searching non-shoppers. For sufficiently high shares of shoppers, the pass-through rate first increases in the number of sellers, then decreases and then increases again. This is different to when we do not allow for sequentially searching non-shoppers.

C.1.6 Dynamics and anticipatory effects

Since we analyze pass-through in a static model, we abstract from how expectations about future prices affect current price setting. Nevertheless, we briefly discuss how expectations may lead to anticipatory effects if extended to a dynamic framework. In particular, anticipatory price increases before a tax increase and a tax decrease are not at odds with the more long-term relationship between price sensitivity, competition, and pass-through that we focus on in this paper.

First, let us extend our model and consider a dynamic framework in which there are not only informed shoppers and uninformed non-shoppers, but within both groups also patient consumers (who could buy before or after the tax change) and impatient consumers (who cannot or do not want to wait).

Let us now consider how an anticipatory price increase could occur before a large pre-announced tax decrease. In this case, all patient consumers wait until the next period. Sellers cannot compete for patient consumers before the tax decrease and so are left with impatient consumers that do not have the option to wait. Within the group of shoppers and non-shoppers, patient consumers are more price sensitive since, also in the absence of a tax change, they have the option to wait. Before a large pre-announced tax decrease, the more price sensitive consumer groups within shoppers and non-shoppers drop out. Compared to a situation without a tax change, equilibrium prices therefore increase and quantities decrease.

Finally, let us consider how an anticipatory price increase could occur before a large pre-announced tax increase. In this case, the option of waiting for another period becomes worse for patient consumers. Therefore, patient consumers become more likely to accept a particular price draw before the tax increase than if there is no pre-announced tax change. For impatient consumers, nothing changes. Patient consumers therefore are willing to accept higher prices than without a large pre-announced tax increase and are more likely to buy in the current period, whereas impatient consumers behave just as they do without a pre-announced tax increase. Compared to a situation without a tax change, equilibrium prices therefore increase and quantities also increase.

C.2 Appendix to Section 3.4: Data and descriptive evidence

C.2.1 Data

Details on constructing the price and margin data set

We construct the price panel at fuel stations in France and Germany as follows. For each fuel station in our data set, we observe a fuel price every time it is changed along with a precise time and date stamp of a change. On average, fuel stations in Germany change fuel prices 15 times a day, whereas there is typically one price change a day at French fuel stations. Based on the distribution of price changes at German fuel stations, we construct hourly fuel prices from 6 am until 10 pm for each day between 1 May and 31 August 2020 and between 1 November 2020 and 28 February 2021. For France, we keep a fuel price at 5 pm for our empirical analysis since fuel prices do not change frequently over the day.

For German fuel stations, we compute daily weighted average prices from the hourly distribution of price changes that we observe. To construct the weights, we use the data on hourly fueling patterns reported in a representative survey among drivers for the German Federal Ministry of Economic Affairs. Figure C.10 shows shares of motorists in Germany who fuel at a given time period during a day. We further re-weight the hourly shares to produce weights for the hours between 6 am and 10 pm.

Figure C.10: Daily fueling patterns (Germany)

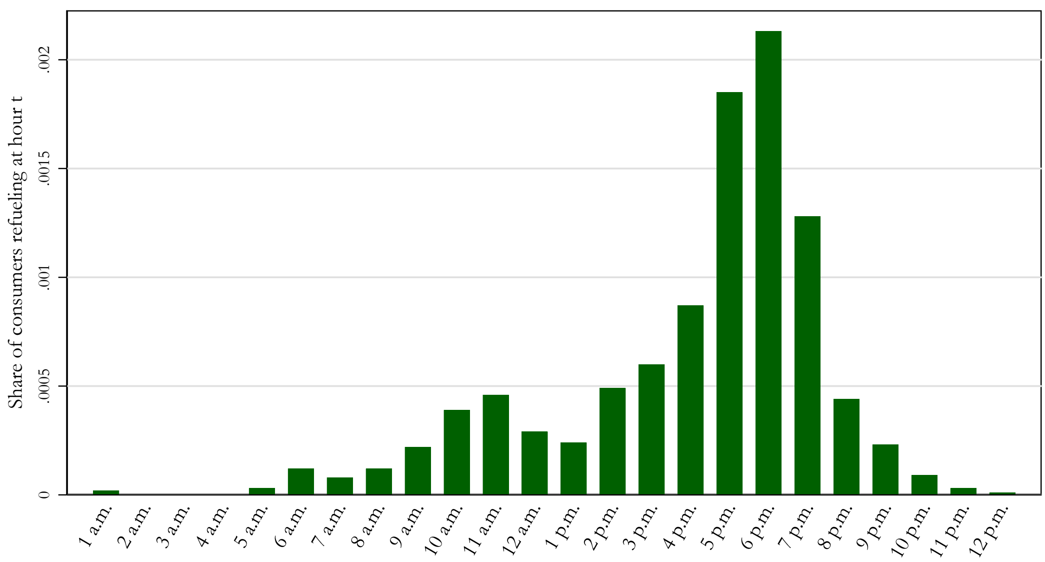

Notes: The Figure shows shares of drivers in Germany who fuel at a given hour of a day. Data is based on a representative survey of motorists in Germany, commissioned by the German Federal Ministry of Economic Affairs.

We also compute retail margins. To compute retail margins, we subtract taxes and duties in France and Germany, as well as an estimate of the input cost of crude oil.

In Germany, taxes and duties consist of the value-added tax, a lump-sum energy tax, and a fee for oil storage. The lump-sum energy tax is at 0.6545 Euro per liter for *E5* and *E10* gasoline, and at 0.4704 Euro per liter for diesel. The fee for oil storage is at 0.0027 Euro per liter for *E5* and *E10*, and at 0.0030 Euro per liter for diesel.[5] Before the VAT reduction, the VAT rate on retail fuel was 19 percent. Between 1 July 2020 and 31 December 2020, this was temporarily reduced to 16 percent. On 1 January 2021, the VAT rate was raised back to 19 percent. At the same time, the German Federal Government introduced a CO_2 price of 25 Euro per emitted tonne of CO_2 on oil, gas and fuel.[6]

In France, the VAT rate on retail fuel is 20 percent, with the exception of Corsica Island, where it is 13 percent. In addition to the VAT, fuel products in France are subject to a lump-sum tax of 0.60 to 0.70 Euro per liter, depending on the metropolitan region and fuel product type.[7]

We obtain daily data on the Brent price of crude oil at the port of Rotterdam from the US Energy Information Administration. A barrel (42 gallons) of crude oil is on average refined into around 19 gallons of gasoline, 12 gallons of diesel, and 13 gallons of other products, such as jet fuel, petroleum coke, and still gas. Among products different from gasoline and diesel, only jet fuel (of which around 4.3 gallons are refined from a barrel of crude oil) yields sizable commercial value.[8]

Assuming that among the other products only jet fuel is of high value, we split the price of a barrel into the cost of producing gasoline, diesel, and jet fuel to compute a share of the Brent price that corresponds to a particular fuel product. Around 54 percent of the Brent oil price per barrel corresponds to the production of 19 gallons of gasoline, and around 34 percent - to the production of 12 gallons of diesel, which we further transform into the input cost per liter of gasoline and diesel. We therefore compute retail margins of *E5*, *E10*, and diesel by subtracting taxes and duties, as well as the approximate input cost of crude oil from the observed fuel price.

Summary statistics for winter 2020/21

In Table C.1, we report summary statistics for the time window around the tax increase. Our analysis is based on the pre-treatment period of 1 November to 15 December 2020 and post-

[5] See https://www.avd.de/kraftstoff/staatlicher-anteil-an-den-krafstoffkosten/.

[6] For *E5* and *E10*, this translates into a per unit tax of 6 Eurocent per liter (7.14 Eurocent including VAT). For diesel, the per unit tax is 6.69 Eurocent per liter (7.96 Eurocent including VAT). Further details can be found in the "Brennstoff-Emissionshandelsgesetz" (2020 Fuel Emissions Trading Act).

[7] See http://www.financespubliques.fr/glossaire/terme/TICPE/.

[8] See https://www.eia.gov/energyexplained/oil-and-petroleum-products/refining-crude-oil.php.

Table C.1: Summary statistics

	Germany pre-treatment	Germany post-treatment	France pre-treatment	France post-treatment
A. Station characteristics				
Number of stations	14,554	14,491	8,832	9,146
Median comp. nr. (5km markets)	4	4	2	2
Share of local monopolists	13%	13%	19%	19%
B. Prices, *E5*				
Mean price	1.23	1.40	1.35	1.45
Mean price net of taxes and duties	.41	.46	.44	.52
Mean retail margin	.13	.11	.16	.17
C. Prices, *E10*				
Mean price	1.19	1.35	1.32	1.41
Mean price net of taxes and duties	.37	.42	.43	.51
Mean retail margin	.09	.07	.15	.15
D. Prices, diesel				
Mean price	1.05	1.24	1.23	1.33
Mean price net of taxes and duties	.43	.50	.42	.50
Mean retail margin	.16	.15	.14	.15
E. Mobility data				
Retail & recreation	-28.8%	-56.8%	-40.7%	-37.8%
Workplaces	-16.1%	-28.9%	-25.1%	-24%

Notes: "pre-treatment" and "post-treatment" refer to fuel stations in Germany and France before and after the increase of the VAT rate and introduction of carbon emissions tax, respectively. The pre-treatment phase goes from 1 November until 15 December 2020. The post-treatment phase goes from 1 January until 28 February 2021.

treatment period of 1 January to 28 February 2021. Table C.1 shows that the price level is generally higher in France than in Germany. Gross prices increase in France by around 9 to 10 Eurocent between pre- and post-tax increase. In Germany, gross prices increase by about 16 to 19 Eurocent, depending on the fuel type. At the same time, net prices in Germany increase between 5 and 7 Eurocent. This is smaller than in France and suggests that the increase in the VAT and the introduction of CO_2 tax were not completely passed on to consumers.

Table C.1 also shows mobility patterns in France and Germany. In both countries, visits to workplaces were around 16 to 29 percent lower in November 2020 to February 2021 compared to their pre-pandemic levels. At the same time, visits to retail and recreational facilities were around 40 percent lower in France and 29 to 57 percent lower in Germany than in the baseline period of 3 January to 6 February 2020.

Summary statistics using SDID weights

In Table C.2, we report summary statistics for the analysis of the tax decrease restricted to the balanced sample used in the SDID analysis. The analysis is based on the pre-treatment period of 1 May to 30 June 2020 and post-treatment period of 1 July to 31 August 2020. In the last two columns, we report summary statistics where we weigh fuel stations in the control group by the station weights they receive in the SDID analysis. In contrast to the summary statistics in Table 3.1, Table C.2 is based on the balanced panel which is required for the estimation with SDID. Due to the sample restriction, the total number of stations in Germany and France is lower than in Table 3.1.

Table C.2 shows that characteristics of the unweighted and weighted control groups are similar. As in the summary statistics based on the full sample in Table 3.1, relative increase in retail margins in Germany remains highest for *E5* and lowest for diesel when we restrict the sample to a balanced panel.

Table C.3 reports analogous summary statistics for the analysis of the tax increase. The last two columns correspond to the control group weighted by the weights in SDID. Table C.3 is based on the balanced panel used in the estimation by SDID, so the number of stations is lower than in Table C.1 that reports summary statistics for the full sample. Across unweighted and weighted control groups, price characteristics and mobility indicators are similar. As in the summary statistics based on the full sample, Table C.3 shows that relative decline in margins in Germany after the tax increase is lowest for diesel.

Table C.2: Summary statistics, tax decrease

	DE pre-change	DE post-change	FR pre-change	FR post-change	FR, weighted pre-change	FR, weighted post-change
A. Station characteristics						
Number of stations	12,171	12,171	5,523	5,523	5,523	5,523
Median comp. nr. (5km markets)	4	4	3	3	2	2
Share of local monopolists	11%	11%	15%	15%	16%	16%
B. Prices, *E5*						
Mean price	1.21	1.27	1.29	1.34	1.28	1.35
Mean price net of taxes and duties	.36	.44	.38	.43	.38	.43
Mean retail margin	.13	.16	.15	.15	.15	.16
C. Prices, *E10*						
Mean price	1.18	1.23	1.26	1.32	1.26	1.33
Mean price net of taxes and duties	.34	.40	.38	.43	.38	.43
Mean retail margin	.11	.13	.15	.15	.15	.16
D. Prices, diesel						
Mean price	1.05	1.07	1.19	1.24	1.20	1.24
Mean price net of taxes and duties	.41	.45	.38	.42	.39	.43
Mean retail margin	.18	.17	.15	.14	.16	.15
E. Mobility data						
Retail & recreation	-22.3%	-2.4%	-34.1%	1.4%	-34.2%	0%
Workplaces	-21.8%	-20.5%	-29.6%	-27.6%	-29.5%	-27.8%

Notes: DE (FR) "pre-change" and "post-change" refer to fuel stations in Germany (France) before and after the reduction of the VAT rate, respectively. The pre-VAT change phase goes from 1 May until 30 June 2020. The post-VAT change phase starts on 1 July 2020. All columns are based on the balanced panel, which is used in the estimation by SDID. Columns labeled with "FR, weighted" correspond to summary statistics on stations in France, when these are weighted by the SDID unit weights.

Table C.3: Summary statistics, tax increase

	DE pre-change	DE post-change	FR pre-change	FR post-change	FR, weighted pre-change	FR, weighted post-change
A. Station characteristics						
Number of stations	12,077	12,077	6,632	6,632	6,632	6,632
Median comp. nr. (5km markets)	4	4	3	3	3	3
Share of local monopolists	11%	11%	17%	17%	9%	9%
B. Prices, *E5*						
Mean price	1.24	1.40	1.34	1.44	1.36	1.46
Mean price net of taxes and duties	.41	.46	.43	.51	.44	.53
Mean retail margin	.13	.11	.15	.15	.17	.17
C. Prices, *E10*						
Mean price	1.19	1.35	1.32	1.41	1.32	1.41
Mean price net of taxes and duties	.37	.42	.43	.50	.43	.50
Mean retail margin	.09	.07	.15	.15	.15	.15
D. Prices, diesel						
Mean price	1.05	1.24	1.23	1.33	1.23	1.32
Mean price net of taxes and duties	.43	.50	.41	.50	.41	.49
Mean retail margin	.16	.15	.14	.14	.14	.14
E. Mobility data						
Retail & recreation	-28.9%	-56.8%	-41.8%	-38.7%	-41.5%	-38.3%
Workplaces	-16.1%	-28.8%	-26.4%	-24.9%	-25.8%	-24.7%

Notes: DE (FR) "pre-change" and "post-change" refer to fuel stations in Germany (France) before and after the increase of the VAT rate and introduction of carbon emissions tax, respectively. The pre-treatment phase goes from 1 November until 15 December 2020. The post-treatment phase goes from 1 January until 28 February 2021. All columns are based on the balanced panel, which is used in the estimation by SDID. Columns labeled with "FR, weighted" correspond to summary statistics on stations in France, when these are weighted by the SDID unit weights.

C.3 Appendix to Section 3.5: VAT Pass-through Estimation

C.3.1 Synthetic difference-in-differences

In the following, we give a brief overview of the SDID method developed by Arkhangelsky et al. (2021).

Consider a balanced panel with N units, T time periods, and outcomes denoted by Y_{it}. Units from 1 to N_{co} and time periods from 1 to T_{pre} are not exposed to the binary treatment $W_{it} \in \{0, 1\}$. Units from N_{tr} to N and time periods from T_{post} to T are exposed to the treatment. To compute the SDID estimator $\hat{\tau}^{sdid}$, the SDID method proceeds via the following algorithm:

1. Compute the regularization parameter according to Equation (C.17)
2. Compute the unit weights $\hat{w}_i^{sdid}$ solving the minimization problem in Equation (C.16)
3. Compute the time weights $\hat{\lambda}_t^{sdid}$ solving the minimization problem in Equation (C.18)
4. Compute the SDID estimator $\hat{\tau}^{sdid}$ by solving the following minimization problem:

$$(\hat{\tau}^{sdid}, \hat{\mu}, \hat{\alpha}, \hat{\beta}, \hat{\gamma}) = \underset{\tau,\mu,\alpha,\beta,\gamma}{\arg\min} \left\{ \sum_{i=1}^{N} \sum_{t=1}^{T} (Y_{it} - \mu - \alpha_i - \beta_t - X_{it}\gamma - W_{it}\tau)^2 \, \hat{w}_i^{sdid} \hat{\lambda}_t^{sdid} \right\}$$

where X_{it} is a vector of controls.[9]

In Steps 1 to 2, the unit weights are computed by solving

$$(\hat{w}_0, \hat{w}^{sdid}) = \underset{w_0 \in \mathbb{R}, w \in \Omega}{\arg\min}\, l_{unit}(w_0, w), \text{ where} \tag{C.16}$$

$$l_{unit}(w_0, w) = \sum_{t=1}^{T_{pre}} \left(w_0 + \sum_{i=1}^{N_{co}} w_i Y_{it} - \frac{1}{N_{tr}} \sum_{i=N_{co}+1}^{N} Y_{it} \right)^2 + \xi^2 T_{pre} \|w\|_2^2,$$

$$\Omega = \left\{ w \in \mathbb{R}_+^N : \sum_{i=1}^{N_{co}} w_i = 1, w_i = N_{tr}^{-1} \text{ for all } i = N_{co}+1, .., N \right\}.$$

[9] See Arkhangelsky et al. (2021) for further details.

ξ is the regularization parameter and w_0 is the intercept. The regularization parameter matches a one period change in the outcome for the control units in the pre-treatment period and is set to

$$\xi^2 = \frac{1}{N_{co}T_{pre}} \sum_{i=1}^{N_{co}} \sum_{t=1}^{T_{pre}} (\Delta_{it} - \bar{\Delta})^2, \text{ where} \tag{C.17}$$

$$\Delta_{it} = Y_{i,(t+1)} - Y_{it}, \text{ and } \bar{\Delta} = \frac{1}{N_{co}(T_{pre}-1)} \sum_{i=1}^{N_{co}} \sum_{t=1}^{T_{pre}-1} \Delta_{it}.$$

In Step 3, the time weights are computed by solving

$$(\hat{\lambda}_0, \hat{\lambda}^{sdid}) = \underset{\lambda_0 \in \mathbb{R}, \lambda \in \Lambda}{\arg\min}\, l_{time}(\lambda_0, \lambda), \text{ where} \tag{C.18}$$

$$l_{time}(\lambda_0, \lambda) = \sum_{i=1}^{N_{co}} \left(\lambda_0 + \sum_{t=1}^{T_{pre}} \lambda_t Y_{it} - \frac{1}{T_{post}} \sum_{t=T_{pre}+1}^{T} Y_{it} \right)^2,$$

$$\Lambda = \left\{ \lambda \in \mathbb{R}_+^T : \sum_{t=1}^{T_{pre}} \lambda_t = 1, \lambda_t = T_{post}^{-1} \text{ for all } t = T_{pre}+1, .., T \right\}.$$

C.4 Appendix to Section 3.6: Empirical Results

C.4.1 Geographical distribution of station weights in the SDID

Figures C.11 and C.12 show the geographical distribution of stations in France. In Figure C.11, we highlight stations that receive a disproportionally high weight in the SDID pass-through estimation of the tax decrease for *E5*, *E10* and diesel. Analogously, in Figure C.12 we highlight stations that receive a disproportionately high weight in the SDID pass-through estimation of the tax increase. The control stations with higher SDID weights are scattered throughout France and there does not appear to be any regional cluster that particularly influences the estimation results.

Figure C.11: France: distribution of fuel stations by SDID unit weights (tax decrease)

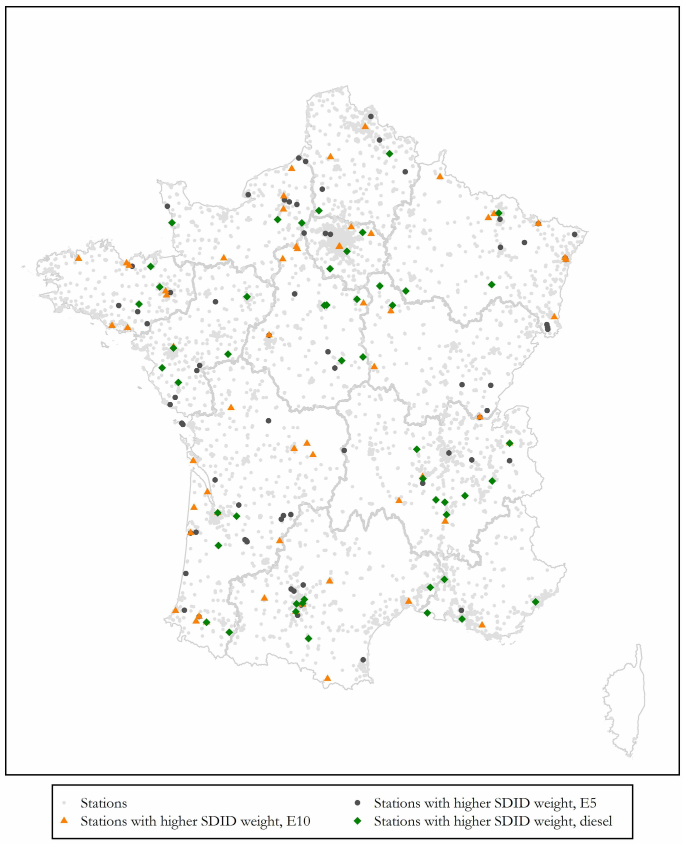

Notes: The Figure shows the geographic distribution of fuel stations in France for the analysis of the tax decrease. Stations with a disproportionally high unit weight in the SDID pass-through estimation for *E5*, *E10* or diesel are highlighted.

Figure C.12: France: distribution of fuel stations by SDID unit weights (tax increase)

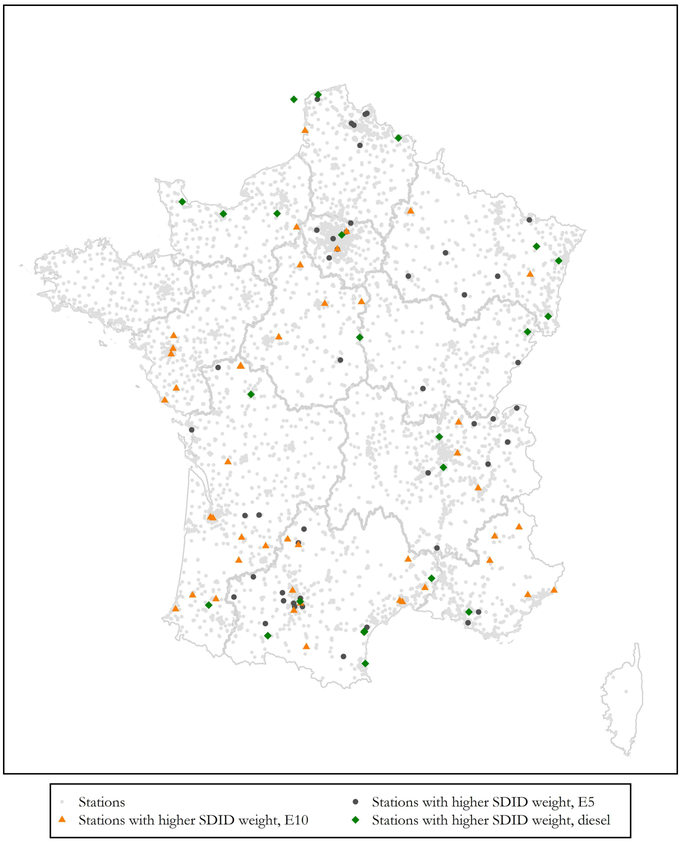

Notes: The Figure shows the geographic distribution of fuel stations in France for the analysis of the tax increase. Stations with a disproportionally high unit weight in the SDID pass-through estimation for *E5*, *E10* or diesel are highlighted.

C.4.2 Robustness: Pass-through estimation with additional controls

In Table C.4, we report results on the effect of the tax change on *E5*, *E10* and diesel prices when we control for regional mobility for retail and recreational purposes and to workplaces, and allow the changes in the crude oil price to differentially affect fuel prices in France and Germany. Overall, the point estimates of the pass-through rates are very similar (no deviation of more than 2 percentage points) to our main estimation results in Table 3.2.

The coefficients in Columns (1) to (3) correspond to the effect of the tax decrease on *E5*, *E10* and diesel prices, and the coefficients in Columns (4) to (6) correspond to the effect of the subsequent tax increase.

The results in Columns (1) to (3) show that the tax decrease led to a decline in prices of all fuel products, which is statistically significant at the 1 percent level and economically significant. The average price for *E5* decreases by 0.88 percent after the VAT reduction, whilst average prices for *E10* and diesel decrease by 1.27 and 2.01 percent, respectively.

Under full pass-through, we would expect prices for each fuel product to decrease by about 2.52 percent. An estimated decline of 2.01 percent in diesel prices is therefore relatively close to full pass-through. Around 80 percent of the tax decrease is passed on to consumers who buy diesel. For *E10*, the pass-through rate is 50 percent. Finally, we estimate that 35 percent of the tax decrease is passed through to consumers of *E5*.

The results in Columns (4) to (6) show that the subsequent tax increase led to an increase in prices of all fuel products. The average price of *E5* increased by about 5.8 percent, whereas *E10* and diesel prices increase by about 6.5 and 8.8 percent after the tax increase, respectively.

Next, we estimate the pass-through rate of the tax increase. Under full pass-through, we would expect an increase in prices by 8.15 percent for *E5*, 8.37 percent for *E10* and 9.66 percent for diesel. We find a joint pass-through rate of the VAT increase and the carbon emissions price of 71 percent for *E5*, 77 percent for *E10* and 91 percent for diesel. This is very close to the pass-through of 69 percent for *E5*, 75 percent for *E10* and 92 percent for diesel, estimated without the additional controls.

Table C.4: Effect of the tax change on log prices (percent)

	E5	E10	Diesel	E5	E10	Diesel
	(1)	(2)	(3)	(4)	(5)	(6)
Tax change	-.0088***	-.0127***	-.0201***	.0577***	.0647***	.0878***
	(.0012)	(.0012)	(.0013)	(.0015)	(.0016)	(.0014)
Pass-through rate	35%	50%	80%	71%	77%	91%
	[25%, 45%]	[41%, 60%]	[70%, 90%]	[67%, 74%]	[73%, 81%]	[88%, 94%]
Retail & recreation	Yes	Yes	Yes	Yes	Yes	Yes
Workplaces	Yes	Yes	Yes	Yes	Yes	Yes
DE × oil price	Yes	Yes	Yes	Yes	Yes	Yes
Date fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Station fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,736,145	1,968,984	2,176,362	1,485,120	1,712,984	1,945,736

Notes: Columns (1) to (3) present average treatment effect estimates of the VAT reduction on E5, E10, and diesel log prices, respectively. Columns (1) to (3) use data from 1 May to 31 August 2020. Columns (4) to (6) present average treatment effect estimates of the VAT increase and CO_2 emissions tax on E5, E10, and diesel log prices, respectively. Columns (4) to (6) use data from 1 November to 15 December 2020 for pre-treatment period, and from 1 January to 28 February 2021 for post-treatment period. 95% confidence intervals on pass-through rates are reported in parentheses. Standard errors are computed using the jackknife method and are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

C.4.3 Robustness: Anticipatory effects

In Table C.5, we estimate pass-through rates if we change the assumptions on anticipatory effects. In Columns (1) to (3), we estimate the pass-through rate of the tax decrease if we drop the second half of June 2020 from the control period. In this case, the gap between pass-through rates between *E5*, *E10* and diesel widens, but the order remains the same. This is not our preferred estimation strategy, since we do not think that there is sufficient evidence for an anticipatory pass-through of the tax decrease in June 2020. We would therefore treat the point estimates of the pass-through rate with caution. Reassuringly, however, our main results, which is the heterogeneity of pass-through with respect to the price sensitivity of consumers, does not change.

In Columns (4) to (6), we report the estimates of the pass-through rate for the tax increase if we include the second half of December 2020 into the control period. In this case, the point estimate of the pass-through rate for *E5* decreases from 69 percent to 65 percent, for *E10* from 75 to 65 percent and for diesel from 92 percent to 84 percent. This is expected, since we

Table C.5: Effect of the tax change on log prices (percent)

	E5	E10	Diesel	E5	E10	Diesel
	(1)	(2)	(3)	(4)	(5)	(6)
Tax change	.0037***	-.0051***	-.0223***	.0531***	.0544***	.0811***
	(.0014)	(.0018)	(.0009)	(.0040)	(.0031)	(.0029)
Pass-through rate	-15%	20%	88%	65%	65%	84%
	[-25%, -4%]	[7%, 34%]	[81%, 95%]	[56%, 75%]	[58%, 72%]	[78%, 90%]
Date fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Station fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,524,420	1,728,864	1,910,952	1,690,320	1,952,760	2,219,160

Notes: Columns (1) to (3) present average treatment effect estimates of the VAT reduction on E5, E10, and diesel log prices, respectively. Columns (1) to (3) use data from 1 May to 15 June for pre-treatment period, and 1 July to 31 August 2020 for post-treatment period. Columns (4) to (6) present average treatment effect estimates of the VAT increase and CO_2 emissions tax on E5, E10, and diesel log prices, respectively. Columns (4) to (6) use data from 1 November to 31 December 2020 for pre-treatment period, and from 1 January to 28 February 2021 for post-treatment period. 95% confidence intervals on pass-through rates are reported in parentheses. Standard errors are computed using the jackknife method and are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

can graphically see important anticipatory effects of the tax pass-through in the second half of December 2020 and so including this time period into the control period necessarily leads to an underestimate of the pass-through rate. The difference between gasoline and diesel remains similar to our main results. The difference between *E5* and *E10* disappears. Although not accounting for anticipatory effects would slightly modify the results, the overall conclusions remain the same. Overall, however, the important anticipatory effects that are obvious in the data lead us to believe that excluding the second half of December 2020 from the analysis is preferable.

C.4.4 Robustness: Difference-in-differences analysis

Using the SDID requires us to restrict our analysis to a balanced subsample of our data. To make sure that our main results are not driven by this sample restriction, we repeat the analysis by estimating the following DID using the full, unbalanced panel:

$$Y_{it} = \beta_0 + \beta_1 \text{Tax}_{it} + \alpha X_{it} + \mu_i + \gamma_t + \epsilon_{it}\,, \tag{C.19}$$

where Y_{it} is the logarithm of the price of gasoline or diesel at a fuel station i at date t, and Tax_{it} is a dummy variable that equals one for stations affected by the tax change at date t. As for the SDID specification, we also include results of a specification where we include a vector of controls, X_{it}, with regional mobility for retail and recreational purposes, mobility to work, and an interaction term of crude oil price with an indicator of stations in Germany. μ_i and γ_t correspond to fuel station and date fixed effects, respectively.

Table C.6 shows the results of estimating the regression model presented in Equation C.19 using the logarithm of price as an outcome variable for the analysis of the tax decrease. The coefficients in Columns (1) to (3) correspond to the effect of the tax decrease on *E5*, *E10* and diesel prices without mobility controls. Columns (4) to (6) show the effects when we control for mobility.

For *E5*, the pass-through rate is 31 percent, and around 49 and 93 percent of the tax decrease is passed on to consumers who refuel with *E10* and diesel, respectively. This is very close to the pass-through rates of 34, 52 and 79 percent for *E5*, *E10* and diesel, respectively, estimated using the SDID method for the balanced subsample. The ranking of pass-through rates with respect to fuel types and their magnitude therefore are robust to using this alternative specification.

Table C.6: Effect of the tax decrease on log prices (percent)

	E5	E10	Diesel	E5	E10	Diesel
Tax decrease	-.0069***	-.0115***	-.0237***	-.0079***	-.0123***	-.0233***
	(.0003)	(.0002)	(.0002)	(.0003)	(.0002)	(.0002)
Retail & recreation				.0016***	.0033***	.0039***
				(.0005)	(.0004)	(.0003)
Workplaces				.0131***	.0115***	-.0017***
				(.0004)	(.0004)	(.0003)
DE $\times$ oil price	.1952***	.1624***	.0394***	.2245***	.1919***	.0451***
	(.0053)	(.0033)	(.0030)	(.0053)	(.0033)	(.0031)
Pass-through rate	27%	46%	94%	31%	49%	93%
Date fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Station fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,150,748	2,332,890	2,725,295	2,149,177	2,329,576	2,721,105
Adjusted R^2	0.889	0.887	0.952	0.890	0.887	0.952
Mean price	1.24	1.21	1.06	1.24	1.21	1.06

Notes: Columns (1) to (3) present estimates without mobility control variables on E5, E10, and diesel log prices, respectively. Columns (4) to (6) present estimates on E5, E10, and diesel log prices from estimation with mobility controls. All columns use data from 1 May to 31 August 2020. Standard errors clustered at the fuel station level are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

We also estimate the effect of the tax increase with the DID specification in Equation C.19 using the full, unbalanced panel.

Table C.7 shows the results of estimating the regression model presented in Equation C.19 using the logarithm of price as an outcome variable for the analysis of tax increase. The coefficients in Columns (1) to (3) correspond to the effect of the VAT rate increase and the CO_2 tax on *E5*, *E10* and diesel prices without mobility controls. Columns (4) to (6) show the effects when we control for mobility. In all columns, we control for an interaction term of crude oil price with an indicator of stations in Germany.

For *E5*, the pass-through rate is 69 percent. For *E10* and diesel, the pass-through is 72 and 84 percent, respectively. This is close to the pass-through rates of 69, 75 and 92 percent for *E5*, *E10* and diesel, respectively, estimated using the SDID method for the balanced subsample. The ranking of pass-through rates with respect to fuel types and their magnitude remain robust to using this alternative specification.

Table C.7: Effect of the tax increase on log prices (percent)

	E5	E10	Diesel	E5	E10	Diesel
Tax increase	.0561***	.0610***	.0831***	.0560***	.0602***	.0813***
	(.0003)	(.0002)	(.0002)	(.0003)	(.0002)	(.0002)
Retail & recreation				-.0013**	-.0039***	-.0054***
				(.0006)	(.0004)	(.0003)
Workplaces				.0010**	.0004	-.0030***
				(.0004)	(.0004)	(.0003)
DE × oil price	.0801***	.0229***	.0807***	.0783***	.0193***	.0778***
	(.0035)	(.0026)	(.0019)	(.0032)	(.0025)	(.0019)
Pass-through rate	69%	73%	86%	69%	72%	84%
Date fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Station fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,808,265	1,985,213	2,322,408	1,807,129	1,982,431	2,318,890
Adjusted R^2	0.949	0.950	0.973	0.949	0.951	0.973
Mean price	1.33	1.28	1.15	1.33	1.28	1.15

Notes: Columns (1) to (3) present estimates without mobility control variables on E5, E10, and diesel log prices, respectively. Columns (4) to (6) present estimates on E5, E10, and diesel log prices from estimation with mobility controls. All columns use data from 1 November until 15 December 2020 for pre-treatment and from 1 January until 28 February 2021 for post-treatment. Standard errors clustered at the fuel station level are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

C.4.5 Number of sellers and tax pass-through for tax increase

Figures C.13 to C.15 show the relationship between the pass-through rate of the tax increase and the number of competitors of a focal station within 5 km catchment area for *E5*, *E10* and diesel. As for the tax decrease, there appears to be a mild hump-shamped relationship between the number of competitors and the pass-through rate for *E5*. For *E10* and diesel, we seem to only observe the upward-sloping part of the hump. Interestingly, as for the tax decrease, the hump-shaped relationship between the number of competitors and the pass-through rate appears to weaken for higher pass-through rates.

Figure C.13: Average pass-through by number of competitor stations, E5

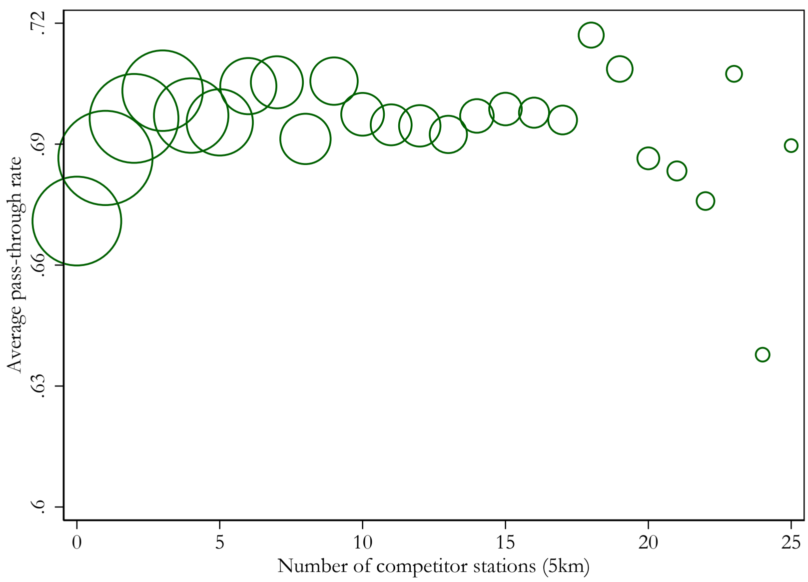

Notes: Each circle plots the average pass-through rate for a group of stations with a particular number of competitors within 5 km catchment area. The number of competitor stations is trimmed at the top percentile.

Figure C.14: Average pass-through by number of competitor stations, E10

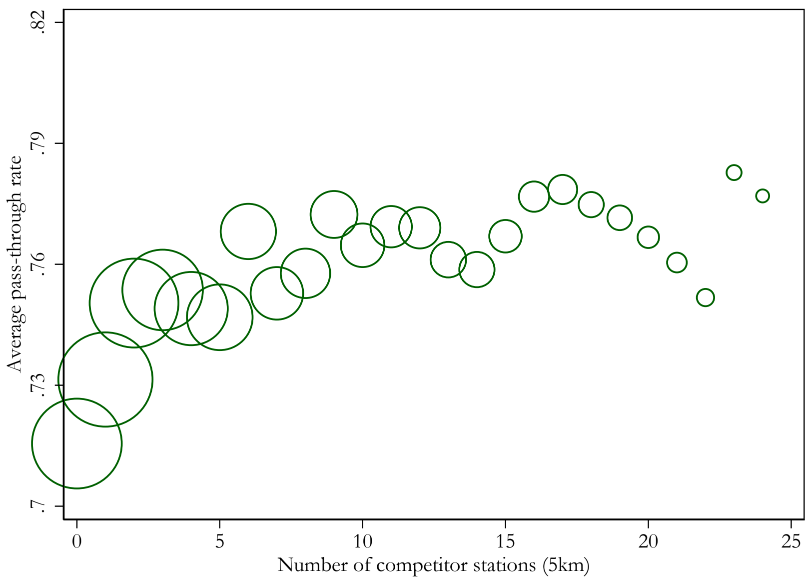

Notes: Each circle plots the average pass-through rate for a group of stations with a particular number of competitors within 5 km catchment area. The number of competitor stations is trimmed at the top percentile.

Figure C.15: Average pass-through by number of competitor stations, diesel

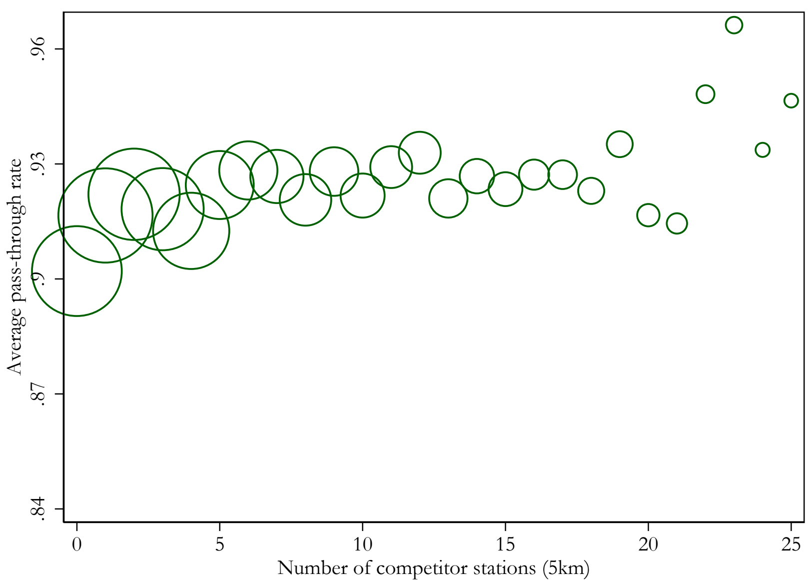

Notes: Each circle plots the average pass-through rate for a group of stations with a particular number of competitors within 5 km catchment area. The number of competitor stations is trimmed at the top percentile.

Bibliography

Acemoglu, D. (2010). "When Does Labor Scarcity Encourage Innovation?" *Journal of Political Economy*, 118 (6): pp. 1037–1078.

Acemoglu, D. and Restrepo, P. (2022). "Demographics and Automation." *The Review of Economic Studies*, 89 (1): pp. 1–44.

Acs, Z. J. and Audretsch, D. B. (1988). "Innovation in Large and Small Firms: An Empirical Analysis." *American Economic Review*, 78 (4): pp. 678–690.

Ahmadpoor, M. and Jones, B. F. (2017). "The Dual Frontier: Patented Inventions and Prior Scientific Advance." *Science*, 357 (6351): pp. 583–587.

Albæk, S., Møllgaard, P., and Overgaard, P. B. (1997). "Government-Assisted Oligopoly Coordination? A Concrete Case." *The Journal of Industrial Economics*, 45 (4): pp. 429–443.

Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., and Wager, S. (2021). "Synthetic Difference-in-Differences." *American Economic Review*, 111 (12): pp. 4088–4118.

Arora, A., Belenzon, S., and Suh, J. (2022). "Science and the Market for Technology." *Management Science*, 68 (10): pp. 7176–7201.

Assad, S., Clark, R., Ershov, D., and Xu, L. (2020). "Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market." *CESifo Working Paper No. 8521*.

Ater, I. and Rigbi, O. (2023). "Price Transparency, Media, and Informative Advertising." *American Economic Journal: Microeconomics*, 15 (1): pp. 1–29.

Autor, D., Dorn, D., Katz, L. F., Patterson, C., and Van Reenen, J. (2020). "The Fall of the Labor Share and the Rise of Superstar Firms." *The Quarterly Journal of Economics*, 135 (2): pp. 645–709.

Bachmeier, L. J. and Griffin, J. M. (2003). "New Evidence on Asymmetric Gasoline Price Responses." *The Review of Economics and Statistics*, 85 (3): pp. 772–776.

Baker, J. B. and Salop, S. C. (2015). “Antitrust, Competition Policy, and Inequality.” *Georgetown Law Journal*, 104 (1-28).

Barron, J. M., Taylor, B. A., and Umbeck, J. R. (2004). “Number of Sellers, Average Prices, and Price Dispersion.” *International Journal of Industrial Organization*, 22: pp. 1041–1066.

Benedek, D., De Mooij, R. A., Keen, M., and Wingender, P. (2019). “Varieties of VAT Pass Through.” *International Tax and Public Finance*, 27 (4): pp. 890–930.

Benzarti, Y. and Carloni, D. (2019). “Who Really Benefits from Consumption Tax Cuts? Evidence from a Large VAT Reform in France.” *American Economic Journal: Economic Policy*, 11 (1): pp. 38–63.

Benzarti, Y., Carloni, D., Harju, J., and Kosonen, T. (2020). “What Goes Up May Not Come Down: Asymmetric Incidence of Value-Added Taxes.” *Journal of Political Economy*, 128 (12): pp. 4438–4474.

Bertrand, J. (1883). “Théorie mathématique de la richesse sociale.” *Journal des Savants*: pp. 499–508.

Besley, T. J. and Rosen, H. S. (1999). “Sales Taxes and Prices: An Empirical Analysis.” *National Tax Journal*, 52 (2): pp. 157–78.

Blandhol, C., Mogstad, M., Nilsson, P., and Vestad, O. L. (2020). “Do Employees Benefit from Worker Representation on Corporate Boards?” *NBER Working Paper No. 28269*.

Borenstein, S., Cameron, A. C., and Gilbert, R. (1997). “Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?” *The Quarterly Journal of Economics*, 112 (1): pp. 305–339.

Bradley, D., Kim, I., and Tian, X. (2016). “Do Unions Affect Innovation?” *Management Science*, 63 (7): pp. 2251–2271.

Brown, Z. Y. (2019a). “An Empirical Model of Price Transparency and Markups in Health Care.” *mimeo*.

– (2019b). “Equilibrium Effects of Health Care Price Information.” *The Review of Economics and Statistics*, 101 (4): pp. 699–712.

Büttner, T. and Madzharova, B. (2021). “Unit Sales and Price Effects of Pre-announced Consumption Tax Reforms: Micro-level Evidence from European VAT.” *American Economic Journal: Economic Policy*, 13 (3): pp. 103–134.

Byrne, D. P. and de Roos, N. (2017). "Consumer Search in Retail Gasoline Markets." *The Journal of Industrial Economics*, 65 (1): pp. 183–193.

– (2019). "Learning to Coordinate: A Study in Retail Gasoline." *American Economic Review*, 109 (2): pp. 591–619.

– (2022). "Start-up Search Costs." *American Economic Journal: Microeconomics*, 14 (2): pp. 81–112.

Carbonnier, C. (2007). "Who pays sales taxes? Evidence from French VAT reforms, 1987–1999." *Journal of Public Economics*, 91 (5): pp. 1219–1229.

Chandra, A. and Tappata, M. (2011). "Consumer Search and Dynamic Price Dispersion: an Application to Gasoline Markets." *The RAND Journal of Economics*, 42 (4): pp. 681–704.

Chetty, R., Looney, A., and Kroft, K. (2009). "Salience and Taxation: Theory and Evidence." *American Economic Review*, 99 (4): pp. 1145–1177.

Conchon, A. (2011). "Board-Level Employee Representation Rights in Europe: Facts and Trends." *European Trade Union Institute Report 121*.

Conlon, C. T. and Rao, N. L. (2020). "Discrete Prices and the Incidence and Efficiency of Excise Taxes." *American Economic Journal: Economic Policy*, 12 (4): pp. 111–43.

Connolly, R. A., Hirsch, B. T., and Hirschey, M. (1986). "Union Rent Seeking, Intangible Capital, and Market Value of the Firm." *The Review of Economics and Statistics*, 68 (4): pp. 567–577.

Correia, S., Guimarães, P., and Zylkin, T. (2020). "Fast Poisson Estimation with High-Dimensional Fixed Effects." *The Stata Journal*, 20 (1).

Covarrubias, M., Gutiérrez, G., and Philippon, T. (2019). "From Good to Bad Concentration? U.S. Industries over the past 30 years." *NBER Working Paper No. 25983*.

D'Acunto, F., Hoang, D., and Weber, M. (2018). "Unconventional Fiscal Policy." *AEA Papers and Proceedings*, 108: pp. 519–523.

– (2022). "Managing Households' Expectations with Unconventional Policies." *The Review of Financial Studies*, 35 (4): pp. 1597–1642.

Danzer, A. M., Feuerbaum, C., and Gaessler, F. (2020). "Labor Supply and Automation Innovation." *IZA DP No. 13429*.

De Loecker, J. and Eeckhout, J. (2021). "Global Market Power." *mimeo*.

De Loecker, J., Eeckhout, J., and Unger, G. (2020). “The Rise of Market Power and the Macroeconomic Implications.” *The Quarterly Journal of Economics*, 135 (2): pp. 561–644.

Dechezleprêtre, A., Hémous, D., Olsen, M., and Zanella, C. (2022). “Induced Automation: Evidence from Firm-level Patent Data.” *mimeo*.

Deltas, G. (2008). “Retail Gasoline Price Dynamics and Local Market Power.” *The Journal of Industrial Economics*, 56 (3): pp. 613–628.

Deltas, G. and Polemis, M. (2020). “Estimating Retail Gasoline Price Dynamics: the Effects of Sample Characteristics and Research Design.” *Energy Economics*, 92: pp. 1049–76.

Diamond, P. A. (1971). “A Model of Price Adjustment.” *Journal of Economic Theory*, 3: pp. 156–168.

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

Dubois, P., Griffith, R., and O’Connell, M. (2020). “How Well Targeted Are Soda Taxes?” *American Economic Review*, 110 (11): pp. 3661–3704.

Duso, T. and Szücs, F. (2017). “Market Power and Heterogeneous Pass-Through in German Electricity Retail.” *European Economic Review*, 98: pp. 354–372.

Eckert, A. (2013). “Empirical Studies of Gasoline Retailing: A Guide to the Literature.” *Journal of Economic Surveys*, 27 (1): pp. 140–166.

Fabra, N. and Reguant, M. (2014). “Pass-Through of Emissions Costs in Electricity Markets.” *American Economic Review*, 104 (9): pp. 2872–2899.

FitzRoy, F. R. and Kraft, K. (1993). “Economic Effects of Codetermination.” *The Scandvinavian Journal of Economics*, 95 (3): pp. 365–375.

Freeman, R. B. and Lazear, E. P. (1995). *An Economic Analysis of Works Councils*. In “Works Councils: Consultation, Representation, Cooperation in Industrial Relations”, Joel Rogers and Wolfgang Streeck, eds. (Cambridge, MA: NBER Comparative Labor Market Series).

Freeman, R. B. and Medoff, J. L. (1985). “What Do Unions Do?” *Industrial and Labor Relations Review*, 38 (2): pp. 244–263.

Fuest, C., Neumeier, F., and Stöhlker, D. (2020). “The Pass-Through of Temporary VAT Rate Cuts Evidence from German Retail Prices.” *ifo Working Papers*, 341.

Ganapati, S., Shapiro, J. S., and Walker, R. (2020). "Energy Cost Pass-Through in US Manufacturing: Estimates and Implications for Carbon Taxes." *American Economic Journal: Applied Economics*, 12 (2): pp. 303–342.

Gautier, E. and Saout, R. L. (2015). "The Dynamics of Gasoline Prices: Evidence from Daily French Micro Data." *Journal of Money, Credit and Banking*, 47 (6): pp. 1063–1089.

Genakos, C. and Pagliero, M. (2022). "Competition and Pass-Through: Evidence from Isolated Markets." *American Economic Journal: Applied Economics*.

German Federal Cartel Office (2011). *Sektoruntersuchung Kraftstoffe*.

German Ministry for Economic Affairs and Energy (2018). *Bericht über die Ergebnisse der Arbeit der Markttransparenzstelle für Kraftstoffe und die hieraus gewonnenen Erfahrungen*.

Gold, M. (2011). "'Taken on Board': an Evaluation of the Influence of Employee Board-Level Representatives on Company Decision-Making across Europe." *European Journal of Industrial Relations*, 17 (1): pp. 41–56.

Gold, M., Kluge, N., and Conchon, A. (2010). *"In the Union and on the Board": Experiences of Board-Level Employee Representatives across Europe*. Brussels: European Trade Union Institute.

Gorton, G. and Schmid, F. A. (2004). "Capital, Labor, and the Firm: a Study of German Codetermination." *Journal of the European Economic Association*, 2 (5): pp. 863–905.

Green, E. J. and Porter, R. H. (1984). "Noncooperative Collusion under Imperfect Price Information." *Econometrica*, 52 (1): pp. 87–100.

Grout, P. (1984). "Investment and Wages in the Absence of Binding Contracts: a Nash Bargaining Approach." *Econometrica*, 52 (2): pp. 449–460.

Hamilton, S. F. (1999). "Tax Incidence under Oligopoly: a Comparison of Policy Approaches." *Journal of Public Economics*, 71 (2): pp. 233–245.

Harding, M., Leibtag, E., and Lovenheim, M. F. (2012). "The Heterogeneous Geographic and Socioeconomic Incidence of Cigarette Taxes: Evidence from Nielsen Homescan Data." *American Economic Journal: Economic Policy*, 4 (4): pp. 169–198.

Harhoff, D., Scherer, F. M., and Vopel, K. (2003). "Citations, Family Size, Opposition and the Value of Patent Rights." *Research Policy*, 32 (8): pp. 1343–1363.

Harju, J., Jäger, S., and Schoefer, B. (2021). "Voice at Work." *NBER Working Paper No. 28522*.

Haucap, J., Heimeshoff, U., Kehder, C., Odenkirchen, J., and Thorwarth, S. (2017). "Auswirkungen der Markttransparenzstelle für Kraftstoffe (MTS-K): Änderungen im Anbieter- und Nachfragerverhalten." *DICE Ordnungspolitische Perspektiven*.

Hirsch, B. T. and Link, A. N. (1987). "Labor Union Effects on Innovative Activity." *Journal of Labor Research*, 8: pp. 323–332.

Hollenbeck, B. and Uetake, K. (2021). "Taxation and Market Power in the Legal Marijuana Industry." *The RAND Journal of Economics*, 52 (3): pp. 559–595.

Hosken, D. S., McMillan, R. S., and Taylor, C. T. (2008). "Retail gasoline pricing: What do we know?" *International Journal of Industrial Organization*, 26: pp. 1425–1436.

Houde, J.-F. (2012). "Spatial Differentiation and Vertical Mergers in Retail Markets for Gasoline." *American Economic Review*, 102 (5): pp. 2147–2182.

I.M.U. (2021). *I.M.U. - Aufsichtsrätebefragung 2019/2021*. Mitbestimmungsreport Nr. 68.

Jacob, M., Müller, M., and Wulff, T. (2021). "Do Consumers Pay the Corporate Tax?" *TRR 266 Accounting for Transparency Working Paper Series No. 15*.

Jäger, S., Noy, S., and Schoefer, B. (2022a). "Codetermination and Power in the Workplace." *Journal of Law and Political Economy*, 3 (1).

– (2022b). "The German Model of Industrial Relations: Balancing Flexibility and Collective Action." *Journal of Economic Perspectives*, 36 (4): pp. 53–80.

Jäger, S., Schoefer, B., and Heining, J. (2021). "Labor in the Boardroom." *The Quarterly Journal of Economics*, 136 (2): pp. 669–725.

Jensen, M. C. and Meckling, W. H. (1979). "Rights and Production Functions: an Application to Labor-Managed Firms and Codetermination." *The Journal of Business*, 52 (4): pp. 469–506.

Johnson, R. N. (2002). "Search Costs, Lags and Prices at the Pump." *Review of Industrial Organization*, 20 (1): pp. 33–50.

Kim, E. H., Maug, E., and Schneider, C. (2018). "Labor Representation in Governance as an Insurance Mechanism." *Review of Finance*, 22 (4): pp. 1251–1289.

Klepper, S. (1996). "Entry, Exit, Growth, and Innovation over the Product Life Cycle." *American Economic Review*, 86 (3): pp. 562–583.

Kopczuk, W., Marion, J., Muehlegger, E., and Slemrod, J. (2016). "Does Tax-Collection Invariance Hold? Evasion and the Pass-Through of State Diesel Taxes." *American Economic Journal: Economic Policy*, 8 (2): pp. 251–286.

Kosonen, T. (2015). "More and Cheaper Haircuts after VAT Cut? On the Efficiency and Incidence of Service Sector Consumption Taxes." *Journal of Public Economics*, 131: pp. 87–100.

Kraft, K., Stank, J., and Dewenter, R. (2011). "Co-Determination and Innovation." *Cambridge Journal of Economics*, 35 (1): pp. 145–172.

Kraftfahrt-Bundesamt (2013). *Personenkraftwagen am 1. Januar 2013 nach ausgewählten Merkmalen*. URL: https://www.kba.de/DE/Statistik/Fahrzeuge/Bestand/Jahrebilanz_Bestand/2013/2013_b_jahresbilanz_tabellen.html?__lang=en&fromStatistic=3532350&fromStatistic=3532350&yearFilter=2013&yearFilter=2013.

– (2021). *Personenkraftwagen am 1. Januar 2021 nach ausgewählten Merkmalen*. URL: https://www.kba.de/DE/Statistik/Fahrzeuge/Bestand/Jahresbilanz/fz_b_jahresbilanz_archiv/2021/2021_b_barometer.html?nn=2597888.

Krieger, J. L., Schnitzer, M., and Watzinger, M. (2022). "Standing on the Shoulders of Science." *HBS Working Paper No. 21-128*.

Kühn, K.-U. and Vives, X. (1995). *Information Exchanges aong Firms and their Impact on Competition*. Vol. Office for Official Publications of the European Communities, Luxembourg.

Lemus, J. and Luco, F. (2021). "Price Leadership and Uncertainty about Future Costs." *The Journal of Industrial Economics*, 69 (2): pp. 305–337.

Lewis, M. (2008). "Price Dispersion and Competition with Differentiated Sellers." *The Journal of Industrial Economics*, 106 (3).

Li, J. and Stock, J. H. (2019). "Cost Pass-Through to Higher Ethanol Blends at the Pump: Evidence from Minnesota Gas Station Data." *Journal of Environmental Economics and Management*, 93: pp. 1–19.

Luco, F. (2019). "Who Benefits from Information Disclosure? The Case of Retail Gasoline." *American Economic Journal: Microeconomics*, 11 (2): pp. 277–305.

Mann, K. and Püttmann, L. (2021). "Benign Effects of Automation: New Evidence from Patent Texts." *The Review of Economics and Statistics*: pp. 1–45.

Martin, S. (2020). "Market Transparency and Consumer Search - Evidence from the German Retail Gasoline Market." *Dice Discussion Paper, No. 350*.

Marx, M. and Fuegi, A. (2020). “Reliance on Science: Worldwide Front-Page Patent Citations to Scientific Articles.” *Strategic Management Journal*, 41 (9): pp. 1572–1594.

McCrary, J. (2008). “Manipulation of the Running Variable in the Regression Discontinuity Design: a Density Test.” *Journal of Econometrics*, 142 (2): pp. 698–714.

Miller, N. H., Osborne, M., and Sheu, G. (2017). “Pass-through in a concentrated industry: empirical evidence and regulatory implications.” *The RAND Journal of Economics*, 48 (1): pp. 69–93.

Miravete, E. J., Seim, K., and Thurk, J. (2018). “Market Power and the Laffer Curve.” *Econometrica*, 86 (5): pp. 1651–1687.

Montag, F., Sagimuldina, A., and Schnitzer, M. (2021). “Does Tax Policy Work When Consumers Have Imperfect Price Information? Theory and Evidence.” *CESifo Working Paper No. 9138*.

Montag, F., Sagimuldina, A., and Winter, C. (2023). “Whom to Inform about Prices? Evidence from the German Fuel Market.” *mimeo*.

Montag, F. and Winter, C. (2020). “Price Transparency Against Market Power.” *mimeo*.

Nakamura, E. and Zerom, D. (2010). “Accounting for Incomplete Pass-Through.” *The Review of Economic Studies*, 77 (3): pp. 1192–1230.

Olmstead-Rumsey, J. (2022). “Market Concentration and the Productivity Slowdown.” *mimeo*.

Pennerstorfer, D., Schmidt-Dengler, P., Schutz, N., Weiss, C., and Yontcheva, B. (2020). “Information and Price Dispersion: Theory and Evidence.” *International Economic Review*, 61 (2): pp. 871–899.

Pew Research Center (2017). “Automation in Everyday Life.”

Poege, F., Harhoff, D., Gaessler, F., and Baruffaldi, S. (2019). “Science Quality and the Value of Inventions.” *Science Advances*, 5 (12): pp. 1–8.

Raiser, T., Veil, R., and Jacobs, M. (2015). *Mitbestimmungsgesetz und Drittelbeteiligungsgesetz*. Berlin: Walter de Gruyter.

Romer, P. M. (1986). “Increasing Returns and Long-Run Growth.” *Journal of Political Economy*, 94 (5): pp. 1002–1037.

Rossi, F. and Chintagunta, P. K. (2016). “Price Transparency and Retail Prices: Evidence from Fuel Price Signs in the Italian Highway System.” *Journal of Marketing Research*, 53 (3): pp. 407–423.

Silvia, S. J. (2013). *Holding the Shop Together: German Industrial Relations in the Postwar Era*. Cornell University Press.

Stahl, D. O. (1989). "Oligopolistic Pricing with Sequential Consumer Search." *American Economic Review*, 79 (4): pp. 700–712.

Stern, N. (1987). "The Effects of Taxation, Price Control and Government Contracts in Oligopoly and Monopolistic Competition." *Journal of Public Economics*, 32 (2): pp. 133–158.

Svejnar, J. (1981). "Relative Wage Effects of Unions, Dictatorship and Codetermination: Econometric Evidence from Germany." *The Review of Economics and Statistics*, 63 (2): pp. 188–197.

The Economist (2020). *Deutschland AG Rethinks Workers' Role in Management*. Last accessed 28 February 2023. URL: https://www.economist.com/business/2020/02/01/deutschland-ag-rethinks-workers-role-in-management.

– (2022). *The Growing Demand for More Vigorous Antitrust Action*. Last accessed 7 March 2023. URL: https://www.economist.com/special-report/2022/01/10/the-growing-demand-for-more-vigorous-antitrust-action.

Varian, H. R. (1980). "A Model of Sales." *American Economic Review*, 70 (4): pp. 651–659.

Verlinda, J. A. (2008). "Do Rockets Rise Faster and Feathers Fall Slower in an Atmosphere of Local Market Power? Evidence from the Retail Gasoline Market." *The Journal of Industrial Economics*, 56 (3): pp. 581–612.

Watzinger, M., Fackler, T. A., Nagler, M., and Schnitzer, M. (2020). "How Antitrust Enforcement Can Spur Innovation: Bell Labs and the 1956 Consent Decree." *American Economic Journal: Economic Policy*, 12 (4): pp. 328–359.

Weyl, E. G. and Fabinger, M. (2013). "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition." *Journal of Political Economy*, 121 (3): pp. 528–583.

Eidesstattliche Versicherung

Ich versichere hiermit eidesstattlich, dass ich die vorliegende Arbeit selbständig und ohne fremde Hilfe verfasst habe. Die aus fremden Quellen direkt oder indirekt übernommenen Gedanken sowie mir gegebene Anregungen sind als solche kenntlich gemacht. Die Arbeit wurde bisher keiner anderen Prüfungsbehörde vorgelegt und auch noch nicht veröffentlicht. Sofern ein Teil der Arbeit aus bereits veröffentlichten Papers besteht, habe ich dies ausdrücklich angegeben.

München, den 16.03.2023

Alina Sagimuldina