Determinants and Effects of Competition

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

2022

vorgelegt von
Felix Montag

Referentin:	Prof. Dr. Monika Schnitzer
Koreferent:	Prof. Ali Yurukoglu, PhD
Promotionsabschlussberatung:	20. Juli 2022

Tag der mündlichen Prüfung: 28. Juni 2022

Namen der Berichterstatter: Monika Schnitzer, Ali Yurukoglu, Fabian Waldinger

Für meine Familie.

Acknowledgements

First and foremost, I would like to thank my supervisor Monika Schnitzer for her continuous support over the past years. I received my first insights into research as a research assistant in her group during my undergraduate studies. After returning to her group for the PhD, her advice, encouragements, and relentless support for my research plans were key to the completion of this dissertation. Ali Yurukoglu played a crucial part in my development over the past years. His mentorship allowed me to acquire the tools for research in empirical industrial organization and sharpen research ideas. Together with Fabian Waldinger and Ying Fan, their unwavering support during the academic job market mean that I can carry on as a researcher in the future.

I am very grateful to Martin Watzinger and Anna Gumpert, who have in large parts taught me how to do research. Although active in different fields, their guidance, support, willingness to discuss and give feedback on early stage ideas and sharing their insights were instrumental in many moments during my studies. Alessandra Allocca joined the group towards the end of my studies but has been key to their successful completion through her advice in research, sharing her experience of the job market, and many discussions about economics and beyond.

My colleagues at the Munich Graduate School of Economics, at the Evidence-Based Economics Program, and especially at the Seminar for Comparative Economics were always helpful and made the years of graduate school much nicer. Markus Nagler, Henrike Steimer, and Matthias Wilhelm deserve special mention for their friendship, support, and for making it a joy to come to work, even at the beginning of the PhD, where research can often be frustrating. I would also like to thank Ines Steinbach, Julia Zimmermann, Dagmar Ehrhardt, and the student assistants at the Seminar for Comparative Economics for their great help.

My co-authors Robin Mamrak, Alina Sagimuldina, Maximilian Schader, Monika Schnitzer, Christoph Winter, and Ali Yurukoglu 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.

While writing this dissertation I was lucky to spend extended periods abroad. I would like to thank Ali Yurukoglu for making possible my visit Stanford GSB. It is during that I learned the technical skills needed for my research today and developed the ideas for my job market paper. The time at Stanford also sparked our collaboration since, which I am very grateful for.

I would also like to thank the German Science Foundation and the EliteNetwork of Bavaria for their financial support. This dissertation also benefited from numerous comments and discussions at conferences and seminars.

Without my family, their support, and the opportunities they gave me, I would not be where I am today. Thank you. Most of all, I thank Josi. It is her patience, support, and encouragement that allowed me to write this dissertation. Without it, none of this would be possible.

Felix Montag
July 2022

Contents

List of Tables

List of Figures

Preface

"To keep our country moving, we have to bring fair competition back to this economy. That's why today, I'm signing an Executive Order promoting competition. It'll lower prices, increase wages, and take another critical step toward an economy that works for everyone."

Joseph R. Biden Jr., 2021

The determinants and effects of competition are front and center in today's political debate. Competition policy has advanced from a niche topic among experts to the center of political attention, as demonstrated by the Executive Order of the President of the United States of America. A fast growing literature finds a rise in concentration (Covarrubias, Gutiérrez, and Philippon, 2019), a rise in markups (De Loecker, Eeckhout, and Unger, 2020; Loecker and Eeckhout, 2021), as well as a rise in firm profits (e.g., Covarrubias et al., 2019 or De Loecker et al., 2020). Although there is an ongoing debate about whether these trends are also present in product markets (see, for example, Benkard, Yurukoglu, and Zhang, 2021), some ascribe these trends to technological change (D. Autor, Dorn, Katz, Patterson, and Van Reenen, 2020) or a decline in antitrust enforcement (Covarrubias et al., 2019).

A rise in concentration and market power is made responsible for phenomena such as the decline in the labor share (D. Autor et al., 2020), rising inequality (Baker and Salop, 2015), missing innovation (Watzinger, Fackler, Nagler, and Schnitzer, 2020), or the productivity growth slowdown (Olmstead-Rumsey, 2020). Whilst the final jury is still out on most of these issues, there is a broad consensus that promoting competition is beneficial. Many of the societal challenges in the years ahead require an efficient use of scarce resources and so fostering competition should be seen as a means to reaching this end.

This dissertation sheds light on three aspects relating to the determinants and effects of competition. The first two chapters strive to inform the design of policies with the intent to foster competition. The final chapter is concerned with how the price sensitivity of consumers determines the intensity of product market competition and how this affects the pass-through of commodity taxes.

The first chapter relates to the literature on merger policy, which is a tool used to keep markets competitive and constrain market power before it arises. In the first chapter I study how foreign entry changes the consumer welfare and domestic employment effects of a product market merger. In this context, foreign entry includes a foreign alternative buyer for the acquisition target, as well as post-merger product entry by foreign competitors.

The key methodological innovation of the chapter is that I set up a model of demand and supply, which features multi-product firms that endogenously choose their product portfolios, set prices, and hire workers. The model features manufacturers and consumers. Manufacturers choose their product portfolios and prices. Consumers make purchase decisions. The model is set up as a two-stage game. At the beginning of the game, each manufacturer is endowed with a set of potential products that it is technologically capable of producing. Each product is associated with an exogenous set of characteristics, a production location, and a marginal cost of production. In the first stage, each firm chooses which potential product to introduce into the market, at a per product fixed and sunk entry cost. In the second stage, firms set prices and consumers make purchases. Finally, the number of manufacturing jobs is determined. This is linear in the quantities of the product market equilibrium. Whether a job is created domestically or abroad depends on the exogenous production location for each product.

Combining hand-collected data on production locations with granular product market data, I estimate the structural parameters of this model to study the acquisition of Maytag by Whirlpool in the U.S. appliance industry. Using these estimates, I compare the consumer welfare and domestic employment effects of an acquisition of Maytag by Whirlpool with an alternative acquisition by the Chinese Haier that had no prior presence in the U.S. market. I find that an acquisition by Whirlpool is always worse for consumers but preserves more U.S. manufacturing jobs. It mildly increases the incentive for rivals to introduce new products and leads to an important decrease in the product portfolio for the merging parties. However, although foreign entry was at the heart of the merger clearance decision, I show that endogenous portfolio adjustments lead to more consumer harm overall. Finally, I estimate the job value for which the domestic employment effects offset consumer welfare losses. Overall, I cannot exclude that the employment effects are sufficient to offset the consumer welfare losses.

I contribute to the empirical literature on product market mergers by extending the analysis to incorporate the effects of mergers on consumer welfare and domestic employment when these mergers involve offshoring. This differs to the nascent literature on labor market power in merger analysis (e.g., Prager and Schmitt, 2021, Shapiro, 2019 or Marinescu and Hovenkamp, 2019). In my case, there is no overlap between the merging parties in local labor markets and thus also no change in labor market power.

The chapter also contributes novel evidence to how endogenous product portfolio choices change the consumer welfare effects of mergers. I find that even for an actual merger that was marginally cleared because of an entry defense, endogenous portfolio adjustments increase the harm to consumers. This is because foreign entry is mostly independent of the merger, whereas the merger leads to fewer products offered by the merging parties. Consistent with other evidence in the literature (e.g., Fan and Yang, 2020, Fan and Yang, 2021 or Caradonna, Miller, and Sheu, 2021) this suggests that product entry in retail product mergers only has a limited constraining effect on the merging parties.

The second chapter, which is based on joint work with Alina Sagimuldina and Christoph Winter, relates to the literature on policy tools to make markets more competitive after when there are significant impediments to competition. Our focus lies on mandatory price disclosure policies in markets where consumers have imperfect information about prices. The aim is to understand what determines the price effect of mandatory price disclosure, with an emphasis on the importance of how many consumers are well informed about prices already before the policy.

We contribute to the theoretical literature on mandatory price disclosure by providing a theoretical analysis of the question in the context of the canonical Varian (1980) model. On the supply side, it features 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 this price information always reaches a fixed number of consumers, irrespective of whether these are *shoppers* or *non-shoppers*. This yields a prediction where the magnitude of the price effect of mandatory price disclosure monotonically decreases in the ex ante share of *shoppers*.

We test the predictions in the context of the introduction of mandatory price disclosure 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, station-level price changes for Germany and France before and after the introduction of mandatory price disclosure in Germany. 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. We use a difference-in-differences design, where fuel stations in Germany are the treatment group and fuel stations in France the control group, to estimate the price effect of mandatory price disclosure for each fuel type.

We find that mandatory price disclosure decreases prices for all fuels but that this decrease is larger for gasoline than for diesel. The difference in treatment effects is particularly strong in the five months after the policy change. Thereafter, the treatment effect stabilizes at a lower level. Finally, we show that follow-on information shocks in the form of local radio reports about fuel prices can further increase the treatment effect.

The empirical analysis contributes to the literature on price transparency policies by shedding light on a novel mechanism by which mandatory price disclosure affects prices. In this context, our analysis highlights the importance of the share of consumers informed about prices before mandatory price disclosure. This complements the existing empirical literature that explores other mechanisms such as tacit collusion (Albæk, Møllgaard, and Overgaard, 1997; Luco, 2019) or that price comparison can increase the credibility of advertising (Ater and Rigbi, 2019).

Overall, the analysis suggests that mandatory price disclosure is most effective in markets where few consumers are well-informed before its introduction. Although the treatment effect weakens over time, policymakers can increase its strength again through complementary information campaigns.

The third chapter, which is based on joint work with Alina Sagimuldina and Monika Schnitzer, moves away from policies with the aim of fostering competition and instead goes a step further by studying the determinants of competition and its impact on 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. We therefore ask whether tax policy works when consumers have imperfect price information.

The theoretical analysis sheds light on how commodity tax pass-through is determined in a setting with consumers that are imperfectly informed about prices. Similar to the second chapter, we place the analysis in the framework of a Stahl (1989) model (an extension of Varian, 1980), where there are perfectly informed *shoppers* and uninformed *non-shoppers*. We extend this model to include commodity taxes and analyze how their pass-through varies with the share of *shoppers*. The more *shoppers* there are, the higher is the average price sensitivity of consumers, because a random consumer is more likely to react to a firm that is undercutting its rival.

We contribute to the theoretical literature by introducing a novel notion of price sensitivity to the 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 with perfect information is that the higher the price elasticity of aggregate demand, the lower the pass-through rate (e.g., Weyl and Fabinger, 2013).

In the second part of the third chapter, we extend the empirical literature on the determinants of pass-through (e.g., Nakamura and Zerom, 2010, Miravete, Seim, and Thurk, 2018 and Hollenbeck and Uetake, 2021). We exploit the introduction of a carbon price and a temporary reduction of the value-added tax in the German retail fuel market to study the predictions of the theoretical model. We use a difference-in-differences design and the universe of station-level prices in Germany and France. To account for differences in consumer information, we compare the pass-through of these tax changes for fuel types whose consumers differ in their level of price information. As predicted by the theory, we find that gasoline stations pass through taxes more to consumers that have a high sensitivity to prices. At the same time, theoretically and empirically, there is a hump-shaped relationship between the number of competitors and the pass-through rate.

This has wide-ranging implications for optimal tax policy, since it is at odds with predictions on how the price elasticity of demand affects pass-through in markets with perfect information. To illicit the intended behavioral responses to Pigouvian taxes therefore requires accounting for imperfect consumer information.

In summary, this dissertation offers new insights into driving forces behind the determinants and effects of competition. The mechanisms that are uncovered and the tools that are developed may hopefully contribute to designing economic policies that foster competition and help make efficient use of scarce resources to tackle the societal challenges ahead.

Chapter 1

Mergers, Foreign Entry, and Jobs: Evidence from the U.S. Appliance Industry

1.1 Introduction

Foreign competition can make markets more competitive and benefit consumers (e.g. Bai and Stumpner, 2019). It can also lead to the offshoring of jobs and harm domestic workers (e.g. D. H. Autor, Dorn, and Hanson, 2013). Traditional merger control overlooks the latter and narrowly focuses on competition. Voters however may care about overall domestic welfare. This can create a disconnect between voters' preferences and the objective of merger control. Proponents of industrial policy therefore argue that domestic employment effects should be considered in merger control.[1]

In this paper, I ask how foreign entry alters the consumer welfare and domestic employment effects of a merger between domestic competitors. In this context, foreign entry includes an alternative foreign buyer, as well as post-merger product entry by foreign competitors. Using a structural model of demand and supply, I analyze how a product market merger affects rival product entry, consumer welfare, and domestic employment. To account for the effects of product entry and exit on consumers and employment, I embed a consumer demand model into an endogenous product choice model, where the demand for domestic labor depends on produc-

[1]Many jurisdictions incorporate public interest considerations into merger control (see OECD, 2016). In Germany and South Africa, these include employment. There are no public interest considerations in merger control in the European Union and the United States.

tion locations and equilibrium quantities in the product market. I use this model to study the acquisition of Maytag by Whirlpool in the United States' market for clothes washers. I estimate the parameters of the model and simulate the consumer welfare and employment effects of two acquisitions: The observed acquisition of Maytag by Whirlpool, as well as a hypothetical acquisition by the alternative buyer at the time, Haier, which had no prior presence in the U.S. market. I provide descriptive evidence around the time of the actual merger to corroborate the predictions of the structural model.

Several findings emerge from the comparison of the two acquisitions: First, around the time of the acquisition of Maytag by Whirlpool, markups increase, but prices do not. Second, post-merger (foreign) rival product entry is mostly independent of who acquires Maytag. Third, a Whirlpool acquisition always leads to the removal of more merging party products than an acquisition by Haier. Fourth, a Whirlpool acquisition is always substantially worse for consumers. Fifth, a Whirlpool acquisition leads to less offshoring and a smaller decrease in U.S. manufacturing jobs. This effect is partially offset by a larger gain in market shares by foreign competitors after a Whirlpool acquisition. Sixth, I calculate how much each additional job maintained by the Whirlpool acquisition (relative to the acquisition by Haier) must be worth to counteract the larger decrease in consumer welfare due to the Whirlpool acquisition. Comparing this to the estimated local labor market effects of new multinational jobs by Setzler and Tintelnot (2021), I cannot exclude the possibility that a Whirlpool acquisition leads to higher domestic welfare. Seventh, welfare effects are unequally distributed. Relative consumer welfare losses mildly decrease in household income. Employment effects are concentrated in a few local labor markets.

The 2006 acquisition of Maytag by Whirlpool is a landmark case in which the Department of Justice (DoJ) unconditionally cleared the merger between the two largest U.S. laundry product manufacturers. Prior to the merger, the Chinese appliance manufacturer Haier made an offer to acquire Maytag. Since Haier did not have a prior U.S. market presence, this acquisition would not have decreased competition. However, Haier planned to relocate Maytag's production to its existing manufacturing plants in China (Goodman and White, 2005).[2] Since there are no public interest exceptions in U.S. merger control, the employment effects should not play a role in the decision. Instead, the DoJ argued that competition would remain unharmed by a Whirlpool

[2]Lacetera and Sydnor (2015) show that there is no inherent limitation to maintaining high-quality production after relocating production. This is consistent with frequent production relocations throughout the sample period.

acquisition as any attempt to raise prices would lead to imports by foreign competitors. This was heavily disputed (see Baker and Shapiro, 2008a).

For the empirical analysis, I construct a comprehensive data set of the U.S. residential laundry market between 2005 and 2015. The core of the product market data comes from *TraQline*, a representative survey of approximately $600,000$ U.S. households per year. On the production side, I hand-collect product-level data on the locations of plants manufacturing for the U.S. market. These location data serve three purposes: First, they allow constructing a production cost shifter that can be used as an instrumental variable for prices in the demand estimation. Second, they allow simulating the effects of different counterfactual scenarios on the number of U.S. manufacturing jobs. Third, they enable a data-driven approach to estimating the marginal cost efficiencies from offshoring.[3]

I descriptively document several trends around the time of Whirlpool's acquisition of Maytag. First, concentration strongly increases for clothes washers and dryers. Second, after controlling for product characteristics, prices of clothes washers and dryers by the merging parties do not increase compared to freestanding ranges by other brands, where there was only a small pre-merger overlap.[4] Third, while LG and Samsung introduce new clothes washers and dryers after the merger, this is also true for freestanding ranges. This suggests that product entry could be at least partially independent of the merger. Finally, I use a county-level difference-in-differences (DiD) design that shows that the closure of Maytag plants and of its headquarter (HQ) increases unemployment, decreases employment, and decreases average wages of the employed.[5]

Several questions remain: Does merger-independent entry reduce prices in the absence of the merger? Is overall entry sufficient to prevent the merging parties from increasing prices? If an acquisition by Whirlpool harmed consumers, could this harm be offset by benefits to U.S. workers? Answering these questions requires a model. As the descriptive trends for clothes

[3]This is similar in spirit to Miller and Weinberg (2017), who estimate how the Miller/Coors merger produced marginal cost efficiencies through a reduction in shipping distance.

[4]Ashenfelter, Hosken, and Weinberg (2013) study the price effects of the Maytag acquisition using the same empirical design with other data. They also do not find any price increases for clothes washers, however they find price increases of 14 percent for newly introduced Whirlpool dryers. I discuss how these differences in results could be related to the different data sources in Section 1.2.

[5]This is in line with recent evidence showing that the presence of multinational firms affects the wages of workers at other firms (see Card, Cardoso, Heining, and Kline, 2018, Alfaro-Ureña, Manelici, and Vasquez, 2021, or Setzler and Tintelnot, 2021). Furthermore, Jacobson, LaLonde, and Sullivan (1993) show that workers separating from distressed firms suffer long-term earnings losses and that these depend on local labor market conditions.

washers and dryers are very similar, I will focus on washers from hereon after.

The model features manufacturers and consumers. Manufacturers choose their product portfolios and prices. Consumers make purchase decisions. The model is set up as a two-stage game. At the beginning of the game, each manufacturer is endowed with a set of potential products that it is technologically capable of producing. Each product is associated with an exogenous set of characteristics, a production location, and a marginal cost of production. In the first stage, each firm chooses which potential product to introduce into the market, at a per product fixed and sunk entry cost.[6] Next, marginal cost and demand shocks are realized. In the second stage, firms set prices and consumers make purchases. I model consumer demand using a static random coefficients discrete choice model, where the price sensitivity of consumers depends on income and some consumers have an unobserved taste for front-loading clothes washers. Finally, the number of manufacturing jobs is determined. This is linear in the quantities of the product market equilibrium.[7] Whether a job is created domestically or abroad depends on the exogenous production location for each product.

On the demand side, the estimation is in the spirit of S. Berry, Levinsohn, and Pakes (2004). Informally, the non-linear demand parameters are identified by the correlation between household income and purchase prices and the correlation between the characteristics of the first and second choice products. I construct a cost shifter based on the production location of each product and the real exchange rate (RER) between the production location and the U.S. This cost shifter is then used as an instrumental variable for price, which is exogenous to product-level demand conditions (see Goldberg and Verboven, 2001 or Grieco, Murry, and Yurukoglu, 2021). The granularity of the data allows identifying rich substitution patterns and thus capture the closeness in competition between products.

On the supply side, I estimate the product-level marginal costs that rationalize the data assuming differentiated Bertrand-Nash competition (see Nevo, 2001). A growing literature is concerned with estimating bounds on the fixed costs of introducing a new product into the market using moment inequalities (see Pakes, Porter, Ho, and Ishii, 2015). Intuitively, the fixed and sunk cost of adding a product that was introduced to the market can at most be the expected variable

[6]Since I only observe product-level entry but no firm-level entry around the time of the merger, I focus on endogenous product choices and abstract away from firm entry.

[7]Wages are determined outside the model. They affect the demand for manufacturing workers through their effect on marginal costs and the product market equilibrium.

profit of the product. Similarly, the fixed and sunk cost of adding a product that is part of the set of potential products but is not introduced to the market must be at least as high as the expected variable profit of that product. Methodologically, the estimation of fixed cost bounds is closest to Eizenberg (2014). Finally, I combine evidence on the number of clothes washers that a manufacturing worker produces per year with the hand-collected product-level plant locations to estimate how different product market equilibria affect the demand for domestic manufacturing workers.

I encounter several empirical challenges. A first challenge is to identify the set of potential products that multi-product firms can introduce.[8] Studying an unconditionally cleared merger allows me to overcome this challenge. For rivals, the incentives to introduce new products are greatest after the Whirlpool acquisition. Thus, any rival product not observed after this acquisition is unlikely to be introduced after a Haier acquisition. In contrast, for the merging parties I observe any product that was removed because of the merger in their pre-merger product portfolio.[9] Draganska, Mazzeo, and Seim (2009) and Fan and Yang (2021) exploit cross-sectional variation in market structure to estimate the set of potential products. This is infeasible in my setting, since I study product portfolio choices at the national level.[10] To make the analysis usable in merger control, where the post-merger outcome cannot be observed, I describe how data available to competition authorities (but not to researchers) pre-merger can be used to estimate the set of potential products and simplify the fixed cost estimation.

A second empirical challenge is the multiplicity of equilibria when simulating counterfactual entry. Due to the large number of products, computing all potential equilibria is computationally infeasible. Instead, I follow a literature (e.g. Lee and Pakes, 2009, Wollmann, 2018 or Fan and Yang, 2020) that uses heuristic learning algorithms to determine equilibrium entry. Each player optimizes her portfolio sequentially, taking the choices of rivals as given. I iterate through players until there is no profitable one-step deviation. I exploit two institutional fea-

[8] An earlier literature on endogenous product entry focuses on single-product firms with discrete product types (e.g. Mazzeo, 2002 or Seim, 2006).

[9] I do not observe products that the merging parties do not carry pre-merger, do not introduce post-merger, but would introduce in the absence of the merger. However, these products are probably less important for firm profits and consumer welfare, since firms chose not to introduce them post-merger.

[10] Eizenberg (2014) analyzes a market without cross-sectional variation in entry. He estimates the set of potential products based on existing product lines and technologies. This works in his context, as he studies how the removal of a frontier technology affects the presence of older products. This is not a viable strategy to study the introduction of new products.

tures for the entry algorithm: First, since firms do not choose product portfolios after the merger from scratch, I initialize the entry algorithm at the pre-merger equilibrium. Second, I increase the computational tractability of the entry game by assuming that firms optimize their product portfolio brand-by-brand, whilst taking into account the effects on the profits of other brands of the same firm. Since firms segment products targeting different consumer groups by brand, this additional restriction should not have a strong impact on equilibrium entry.

The key methodological innovation of this paper is to propose a model to analyze the trade-off between the effects on consumer welfare and employment of a product market merger and estimate its structural parameters. This analysis differs to the nascent literature on labor market power in merger analysis (e.g. Prager and Schmitt, 2021, Shapiro, 2019 or Marinescu and Hovenkamp, 2019). In my case, there is no overlap between the merging parties in local labor markets and thus also no change in labor market power.[11] Instead, I ask how the identity and restructuring plans of different potential acquirers and product market rivals affects U.S. employment.[12]

The empirical results shed light on the interaction between the consumer welfare and employment effects of a product market merger. Without efficiencies, an acquisition of Maytag by Whirlpool leads to a decrease in consumer welfare between 6.6 and 10.1 percent compared to a Haier acquisition. However, it also leads to the maintenance of 1,021 to 1,507 additional U.S. manufacturing jobs. Decomposing the employment effect into a relocation and a reallocation effect shows that foreign competition is a double-edged sword. The relocation of Maytag jobs after a Haier acquisition is greater than after a Whirlpool acquisition, since the latter only partially offshored Maytag's production. Although the presence of competitors reduces the post-merger harm to consumers, the reallocation of market shares to competitors producing abroad also mildly decreases the employment benefits of a Whirlpool acquisition.

The estimates show that domestic employment benefits of a Whirlpool acquisition as compared to an acquisition by Haier could plausibly offset losses to consumer welfare. Considering clothes washers only and without efficiencies, I show that an annual average job value of between \$135,000 and \$316,000 is necessary to offset consumer welfare losses. Using a back-of-

[11] Maytag and Whirlpool do not operate plants in the same local labor markets pre-merger.

[12] Wollmann (2018) estimates how output changes with and without the 2009 automobile bailout affect employment. He assumes that all products are always produced in the United States.

the-envelope calculation, I find that this value is on average below $80,000 for other appliance categories. In comparison, Setzler and Tintelnot (2021) find that the total wage bill in a local labor market increases by around $113,000 per year for each additional job created by a foreign multinational firm. This does not include any other benefits of employment, which further increase the value of a job. Given these estimates, I cannot reject that the domestic employment effects overturned the consumer welfare effects of a comparison between these two acquisitions. These findings relate to a literature that quantifies the trade-off between consumer welfare and employment of trade liberalization (see Jaravel and Sager, 2020) and restrictions (see Hufbauer and Lowry, 2012 or Flaaen, Hortaçsu, and Tintelnot, 2020). Among these estimates, I find the lowest job values necessary to offset consumer welfare changes.

Finally, I contribute novel evidence to how endogenous product portfolio choices change the consumer welfare effects of mergers.[13] I find that even for an actual merger that was marginally cleared because of an entry defense, endogenous portfolio adjustments increase the harm to consumers. This is because foreign entry is mostly independent of the merger, whereas the merger leads to fewer products offered by the merging parties. Existing studies mostly consider hypothetical changes in concentration and find mixed results. Fan and Yang (2020) find that endogenous product adjustments exacerbate negative consumer welfare effects, whereas Wollmann (2018) finds the opposite. Fan and Yang (2021) show that product portfolio adjustments exacerbate negative merger effects in small markets and reduce consumer harm in larger markets. Under certain conditions, Caradonna et al. (2021) show that without marginal cost efficiencies product portfolio adjustments can never be profitable for the parties and also fully offset consumer harm. I find that marginal cost efficiencies also limit the strength of an entry defense, since they reduce the incentives for rivals to add new products.[14]

The remainder of the paper is structured as follows: The next section discusses the details of the case and describes the data. Section 1.3 presents the descriptive evidence, Section 1.4 outlines the industry model, Section 1.5 sketches the estimation strategy, Section 1.6 presents the results, Section 1.7 describes the welfare effects, Section 1.8 discusses simplifying estimation with

[13] A related literature (e.g., Werden and Froeb, 1998; S. Li, Mazur, Park, Roberts, Sweeting, and Zhang, Forthcoming; and Ciliberto, Murry, and Tamer, 2021) studies mergers and static entry for single-product firms. Garrido (2020) studies dynamic product entry decisions by multi-product firms assuming nested logit demand. Fan (2013) studies product repositioning after mergers. Several papers study the effect of mergers on entry and product variety for radio stations (e.g., S. T. Berry and Waldfogel, 2001; Sweeting, 2010; and Jeziorski, 2015).

[14] Cabral (2003) shows this theoretically for single-product firms.

proprietary data, and Section 1.9 concludes.

1.2 Institutional Setting and Data

1.2.1 The acquisition of Maytag by Whirlpool

Prior to its acquisition by Whirlpool, Maytag had been struggling financially for several years. Although the company had already cut costs by reducing its workforce by 20 percent, in 2004 it continued to struggle with cost pressure, a further decline in revenues and posted a net loss (Maytag, 2005). In May 2005, the management of Maytag agreed to be bought by a group of private investors for $1.13 billion (Barboza, 2005). In June 2005, the Chinese household appliance manufacturer Haier made a competing bid of $1.3 billion. One month later, Maytag's biggest manufacturing rival in the U.S. appliance market, Whirlpool, outbid Haier with an offer of $1.4 billion. A few days later, Haier withdrew its bid and in March 2006 Whirlpool acquired Maytag after an unconditional merger clearance by the Department of Justice.

Haier's bid came at a time when the Chinese government pushed its large companies to make foreign acquisitions to get access to foreign markets for its manufactured goods, particularly in the European Union and the United States.[15] Since Chinese acquirers were met with resistance, these acquisitions often targeted well-known brand names slipping into decline. This made the acquisition itself easier and also helped overcome the resistance of consumers towards Chinese brands in the product market.[16] With its weak financial performance and its strong brand portfolio, Maytag perfectly fit the bill. Haier, who previously had negligible sales in the U.S. appliance market, planned to use Maytag's brands, repair network and distribution channels, whilst offshoring production to Haier's existing plants in China (Goodman and White, 2005).

Against this backdrop, Whirlpool's bid for Maytag could be seen as fending off a foreign takeover. The main caveat, however, was that Whirlpool and Maytag were close competitors in the product market for several major appliance categories. In its investigation of the acquisition, the DoJ focused on residential clothes washers and dryers. For the manufacturing

[15] This was part of China's "Go Out Policy", promoting Chinese investments abroad (Goodman and White, 2005).

[16] A famous example is the 2005 acquisition of I.B.M.'s personal computer division by Lenovo.

of laundry products, this was a merger from four to three, where Whirlpool and Maytag were the largest and second largest manufacturers in the U.S. market. With its Kenmore brand Sears was another large brand owner in the laundry market; they however did not manufacture any appliances themselves but purchased them from original equipment manufacturers (OEMs) instead. For instance, all clothes washers sold under the Kenmore brand in 2005 were produced by Whirlpool. The DoJ concluded that despite the high market shares of the merging parties, they would not be successful in raising prices because "LG, Samsung, and other foreign manufacturers could increase their imports into the U.S." (Department of Justice, 2006). It therefore unconditionally cleared the acquisition. Baker and Shapiro (2008a) called this decision "[...] a highly visible instance of underenforcement" and Baker and Shapiro (2008b) described it as "fueling the perception that the Justice Department has adopted a very lax merger enforcement policy [...]". They conclude that in this case the DoJ was willing to accept entry and expansion arguments in a highly concentrated merger case, although entrants had thus far only achieved relatively low market shares.

1.2.2 The data

To analyze the implications of the Maytag acquisition by Whirlpool, I construct a comprehensive data set on the U.S. market for residential laundry products between 2005 and 2015.

Sales, products, and households

The centerpiece of the data comes from *TraQline*. This is a data set well-known across the appliance industry and is used by major retailers and all of the major brands in the industry as a source for market insights.[17] In every quarter, a representative sample of around 150,000 U.S. households is asked about appliance purchases. The survey is a repeated cross-section and in total around 600,000 households are surveyed every year. The data spans the years 2005 until 2015. For each respondent, *TraQline* records the number of appliances bought, the price,

[17] The only other comparable source of data on volume and value sales in the appliance industry is a, now discontinued, retailer panel by the NPD Group, which was the basis of the analysis by Ashenfelter et al. (2013). To the best of my knowledge, the key difference between the data sets is that the retailer panel does not include any sales from Sears, which, at the time, was the largest U.S. retailer for household appliances and accounted for an important share of Maytag and Whirlpool sales.

a detailed set of product characteristics (e.g. the brand or whether a product is Energy Star certified), other brands that the household considered buying, the retailer at which the appliance was bought, as well as a detailed set of household demographics. The data includes information for clothes washers and dryers, as well as for freestanding ranges.

Although *TraQline* records detailed characteristic information, respondents are not asked to provide the exact model specification of the appliance they purchased. I therefore use brand, retailer and key characteristics information to aggregate appliance purchases into products. Most brand owners use different brands to cluster their product offering according to the consumers that they target.[18] Thus, the brand of a product already captures much of the variation in, otherwise unobserved, product quality. Certain key product characteristics need to be reported by all survey respondents. For clothes washers, this includes whether a clothes washer is a regular top-loader (with an agitator), a high-efficiency top-loader (without an agitator) or a front-loader. Finally, I further refine the product definition by using information on the retailer at which the product is sold. Different retailers serve different customers. If a brand and key characteristics combination (e.g. a Whirlpool high-efficiency top-loading washing machine) is sold at both, a higher-end retailer such as Sears, and a lower-end retailer such as Best Buy, these products may still slightly differ in other characteristics.[19] To capture all of these sources in observed and unobserved characteristics variation, I define a product as a brand, retailer and key characteristics combination.

Other characteristics only need to be reported by a random subsample of respondents. This is to reduce the burden on respondents. Households that are selected to answer the more detailed characteristics questions do not have the possibility to opt-out, ruling out any selection problems. For clothes washers, these more detailed characteristics include whether it has a child lockout, the number of special programs, whether it is a stacked pair or whether it has additional noise insulation. For each product, I calculate the average value of these characteristics among the subsample of respondents.

Although household demographics allow constructing different geographic markets within the

[18] In its 2007 Annual Report, Whirlpool describes what each of its brands represents and what type of consumers it targets. Amana, for example, is described as stylish and affordable, whereas KitchenAid should stand for quality and craftsmanship, Whirlpool for innovation and Maytag for reliability.

[19] For retailers, I distinguish between Best Buy, H. H. Gregg, Home Depot, Lowe's, Sears, and all others. The latter group pre-dominantly includes smaller, regional retailers. A further disaggregation within this group would lead to many products with very few sales and thus noisy estimates.

U.S., I decide to aggregate products at the national level, because product entry is determined for each major retailer at the national level. I also aggregate responses at the yearly level.

I enrich the *TraQline* product data set with two additional product characteristics: the brand repair rate and brand-level advertising expenditures.

The brand repair rates come from Consumer Reports, a nonprofit consumer organization that tests products across multiple categories and publishes a monthly magazine with test results by product category. Major appliances have long been an important product category for Consumer Reports. Between 2005 and 2015, clothes washers were featured at least once a year. Each report included an overview of brand-level repair rates. This data is based on responses to the Annual Product Reliability Survey conducted by the Consumer Reports National Research Center for more than $100,000$ clothes washers. I digitize this information to create a measure of the perceived product reliability of a brand in a particular year.

Annual information on advertising expenditures comes from Kantar AdSpender between 2005 and 2015. This is a database that includes information on the annual advertising expenditure of a brand by product and media channel. I use the total advertising expenditure of a brand across media channels to capture variation in brand reputation over time. Benkard et al. (2021) use this data set to track brand ownership over time.

The *TraQline* data set only includes household demographics for respondents that purchase an appliance but not for those that do not. To identify how household income affects the sensitivity to prices in the demand estimation, I also need data on the unconditional distribution of income among the population of households (not only of those who purchased an appliance). For this, I draw a random sample of households from the IPUMS Current Population Survey (CPS). This data set includes rich demographic information for a representative household sample for every year in the analysis period.

1.2.3 Production locations and an instrumental variable for price

On the supply side, the core of the data consists of a hand-collected data set containing the locations of plants manufacturing clothes washers for the U.S. market at the product level. This data set serves three purposes. First, it allows constructing a product-level instrumental variable for prices based on differences in the production costs. Second, the product-level plant locations allow simulating how the number of U.S. clothes washer manufacturing jobs changes between counterfactual scenarios. Third, it enables a data-driven approach to estimate marginal cost efficiencies coming from offshoring and the resulting changes in production costs.

Figure 1.1 shows the plant locations of major clothes washer manufacturers for the U.S. market in 2005. To construct the panel of production locations, I collect production locations for all manufacturers with a market share of more than 3 percent in any year between 2005 and 2015. These are Electrolux, General Electric, LG, Maytag, Samsung, and Whirlpool. Whenever possible, I collect information on the exact plant location (e.g. Newton, Iowa). For the purpose of the analysis in this paper however, it is sufficient to know in which country a product is produced.

Figure 1.1: Clothes washer plants manufacturing for the U.S. market, 2005

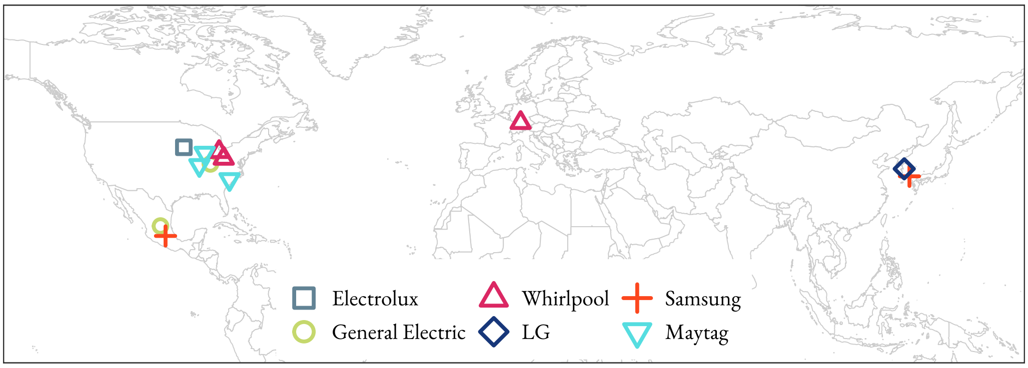

Notes: The map shows all plants manufacturing clothes washers for the U.S. market in 2005 by manufacturers with a market share of more than 3 percent in any year in the sample. The Appendix includes a map for 2007 in Figure A.5, for 2009 in Figure A.6 and for 2011 in Figure A.7.

For LG and Samsung, the production locations before 2012 are mostly based on the investigation by the U.S. International Trade Commission (USITC) into imports of large residential clothes washers from Mexico and South Korea. For 2012 until 2015, production locations for LG and Samsung are based on firm-level clothes washer imports based on the PIERS data set, which uses bill of landing documents and is reported in Flaaen et al. (2020).

For Electrolux, Maytag and Whirlpool, the bulk of the information on manufacturing plant locations is based on information in their annual reports. Since General Electric is not primarily an appliance manufacturer, its annual report does not contain information on appliance plant locations. I therefore base plant locations on a combination of documents from the USITC investigation and news reports. Finally, to make sure that plants produce clothes washers for the U.S. market, I check plant locations against import data split by top-loading and front-loading clothes washer at the country-level from the USITC.

Occasionally, a product is produced in multiple countries for the U.S. market (e.g. in 2008 Whirlpool front-loaders are produced in Mexico and Germany). In this case, I use the same sources as described above to construct weights on the share of the product produced in each production location. I summarize plant weights in Table A.1.

To explain the need for an instrumental variable for price and how I construct one, let us briefly jump ahead to the estimation of clothes washer demand as part of the structural model. As is well-known in the literature on demand estimation, there can be unobserved demand shocks that simultaneously affect prices and quantities. Simply regressing quantities on prices would therefore lead to biased estimates. To get an unbiased estimate of the reaction of quantities to price changes, I need an instrument for price that is unrelated to unobserved demand shocks (exogeneity) and has a sufficiently strong effect on prices (relevance).

An ideal instrument is a variable that captures differences in product-level marginal costs and is unrelated to demand. I use the product-level weighted average real exchange rate (RER) between the U.S. and the countries in which the production of the product is located. This is also used by Grieco et al. (2021) to estimate demand for automobiles. The RER comes from the Penn World Table. Product-level plant weights are constructed as described above.

I use the RER based on consumption expenditures. This is calculated by dividing the consumption of households at nominal prices by the the same consumption using the U.S. price level in 2005 and then multiplying this by the nominal exchange rate between the local currency and the U.S. dollar (Feenstra, Inklaar, and Timmer, 2015). It therefore consists of differences in the relative price levels and serves as a proxy for the local wage level, as well as fluctuations in the nominal exchange rate.

Figure 1.2 shows the evolution of the average RER over time and illustrates the source of the variation. The left panel plots the average RER of all production locations for a particular manufacturer. The average RER is based on the country-level RER of different plant locations of a manufacturer for a product in a particular year, weights that capture which share of a product is produced by a particular plant, and weights based on the sales volume of different products sold by a manufacturer. Although this masks within-manufacturer variation in the RER, already at this level there is significant variation. In the right panel, I further disentangle the average RER for Whirlpool and Maytag products.[20] This shows that there is additional variation in the RER below the manufacturer level, because the same manufacturer produces different products in different countries. For example, whereas all Maytag and Whirlpool top-loaders are produced in the U.S., over the sample period Maytag front-loaders were produced in the U.S. and Mexico and Whirlpool front-loaders in the U.S., Mexico and Germany.

Figure 1.2: Average real exchange rate over time

(a) RER by manufacturer **(b)** RER by product of merging parties

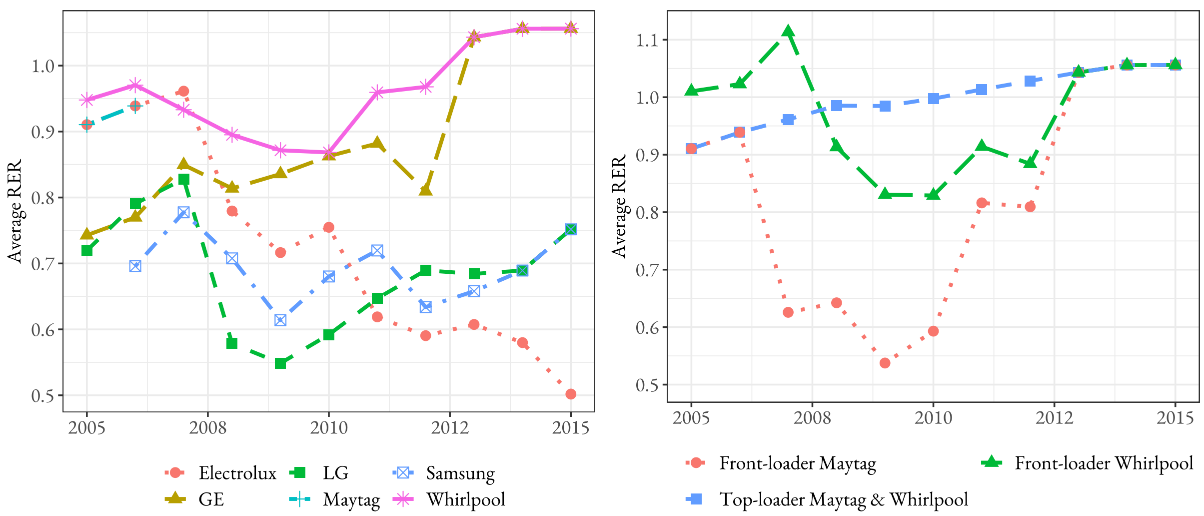

Notes: The left panel plots the average real exchange rate of all production locations by manufacturer over time. It includes the RER for all manufacturers with a market share of at least 3 percent in any year in the sample. The right panel plots the average RER of all production locations by product of the merging parties. The average RER is based on the plant locations in a particular year, the plant weights and the country-level RER. In the right panel, Maytag includes all products marketed under the brands owned by Maytag pre-acquisition (i.e. Admiral, Amana, MagicChef and Maytag) and Whirlpool includes all other brands owned by Whirlpool.

The large variation in the RER over time is also consistent with anecdotal evidence about the importance of the local cost of production for appliance manufacturers. One of the principal reasons why Maytag was struggling financially pre-merger was that its production costs were too high, in parts due to its lack of international production.[21] In a similar spirit, Electrolux launched its global cost-cutting program in 2004, with the aim to offshore more than half of its

[20] Maytag includes all products marketed under the brands owned by Maytag pre-acquisition (i.e. Admiral, Amana, MagicChef and Maytag) and Whirlpool includes all other brands owned by Whirlpool.

[21] This was highlighted throughout Maytag's 2004 annual report, as for example in the following: "Globalization of manufacturing is allowing companies to reduce costs by reaching around the world farther, faster and cheaper than ever before. It's no longer a trend we can watch with interest but a reality to which we are responding" (Maytag, 2005; p. 3).

production to low-cost countries by 2009 (Electrolux, 2007).[22] Both firms exclusively served the U.S. clothes washer market from the U.S. until 2007. This highlights the importance of production locations for costs and competitiveness in the appliance industry and also describes the source of variation in the cost measure: Changes in the RER between the U.S. and a particular production location over time, as well as changes in the production locations.

1.2.4 Labor market data

Finally, I use data on local labor markets from the U.S. Bureau of Labor Statistics (BLS). To analyze the labor market effects of plant closures, I am particularly interested in local wage and employment data. These come from the Quarterly Census of Employment and Wages (QCEW) and the Local Area Unemployment Statistics (LAUS).

The QCEW collects quarterly employment and wage data at the county level as reported by employers. I use the quarterly wages per employee, disaggregated by county and industry. These wages include total compensation, bonuses, stock options, severance payments, the cash value of meals and lodging, tips, and other gratuities. I annualize these wages for ease of interpretability.

The LAUS aggregates data from state-level workforce agencies. It includes monthly information on the number of employed and unemployed individuals for every U.S. county.

1.3 Descriptive Evidence

Before diving into the theoretical model, I document descriptive trends around the Maytag acquisition. To this end, I study the evolution of concentration, prices, product entry, and U.S. appliance manufacturing employment around the time of the acquisition.

[22] By the end of the sample period, Electrolux had lost most of its share of the U.S. laundry market and served its remaining customers from low-cost countries.

1.3.1 Changes in concentration

Table 1.1 shows the evolution of brand owner shares around the time of the Maytag acquisition by Whirlpool. Prior to the merger, Whirlpool and Maytag were the largest and third largest brand owners for laundry products in the U.S. market. Since Sears does not manufacture any appliances itself, Whirlpool and Maytag were also the largest and second largest laundry product manufacturers.[23] In contrast, Haier had no significant market shares in either product market.

The largest rival manufacturers of clothes washers and clothes dryers before the merger were General Electric and Electrolux. LG started gaining market shares, whereas Samsung was not yet present in the U.S. laundry market in 2005. It did, however, already have existing relationships with retailers, since it sold other products (e.g. consumer electronics) at these retailers.

The pre-merger Herfindahl-Hirschman Index (HHI) and the increase in the HHI because of the merger based on pre-merger market shares indicate that the transaction led to a strong increase in concentration.[24] According to the U.S. horizontal merger guidelines, the acquisition therefore potentially raises significant competitive concerns.[25]

Finally the evolution of market shares from just after the merger in 2007 to 2009 show that although some rivals gained market shares and the HHI gradually declined (as compared to the post-merger HHI based on pre-merger market shares), the increase in concentration due to the merger remains substantial and persistent.

[23] One approach could be to count all sales of Sears products towards the respective manufacturer. This would not be an appropriate reflection of market power, however, as Sears could switch supplier if faced with a large increase in prices. Indeed, although Whirlpool manufactured all Sears clothes washers prior to the merger, Sears switched to LG as a supplier of front-loading clothes washers in 2008. This shows that switching suppliers is not only a theoretical possibility and suggests that separately analyzing brand owners is more appropriate.

[24] The HHI is calculated as the sum of squared market shares using whole percentages (i.e. 1 to 100).

[25] The U.S. horizontal merger guidelines identify mergers with a pre-merger HHI between $1,500$ and $2,500$ and an increase in the HHI by more than 100 as potentially raising significant competitive concerns.

Table 1.1: Volume share by brand owner (%)

	Clothes washers			Clothes dryers		
	2005	2007	2009	2005	2007	2009
Whirlpool	25	44	42	27	44	42
Maytag	23			21		
Sears	25	20	18	25	21	19
General Electric	14	17	16	15	17	16
Electrolux	7	6	6	7	6	5
LG	3	7	10	2	6	10
Samsung	0	1	5	0	1	5
HHI	2,048	2,729	2,506	2,072	2,784	2,507
ΔHHI	1,149			1,124		

Notes: The table shows the market share in terms of volume sales by brand owners for clothes washers and clothes dryers pre-merger (2005) and post-merger (2007 and 2009). The HHI is calculated as the sum of squared market shares using whole percentages. The increase in the HHI is based on pre-merger market shares.

1.3.2 Evolution of prices

I next turn to the descriptive evolution of prices around the time of the acquisition. Ashenfelter et al. (2013) compare the evolution of Maytag and Whirlpool product prices for appliance categories with a large increase in concentration to categories with low increases in concentration. Since I use a different data source, I repeat the descriptive price analysis. In particular, the *NPD* data used by Ashenfelter et al. (2013) only includes product sales at a subset of retailers (e.g. omitting sales at Sears), which could lead to systematically different results.

As a comparison appliance category, I use freestanding ranges.[26] This is an appropriate control group if, in the absence of the merger, prices would have evolved similarly in the treatment and the control groups. Since I cannot observe the price evolution of laundry products without the merger directly, I use two indirect ways of assessing this assumption. First, I verify the parallel trends assumption prior to the acquisition. Second, I assess whether other market trends, in particular product entry, are likely to have affected the treatment and control groups similarly, had the merger not occurred.

The analysis starts in the first quarter of 2005 and ends in the last quarter of 2008. Each observation is a product in a particular quarter. To analyze the evolution of prices conditional

[26] Ashenfelter et al. (2013) use ranges, cooktops, ovens and freezers as comparison categories.

on product characteristics in the treatment and control groups, I estimate the parameters of the following model for each appliance category separately

$$log(p_{it}) = \beta x_{it} + \gamma_t + \epsilon_{it}\ , \tag{1.1}$$

where $log(p_{it})$ is the logarithm of price for product i at time t, x_{it} is a vector of product characteristics and γ_t are quarter $\times$ year fixed effects. For clothes washers and dryers, I only include products by Whirlpool and Maytag. For freestanding ranges, I only include products not produced by Whirlpool and Maytag. Instead of product fixed effects, I control for a rich set of characteristics, including the brand and the retailer. This has the advantage of not absorbing merger-specific price changes into the fixed effects for products present only before or only after the merger.[27]

Figure 1.3: Change in the average log price conditional on product characteristics

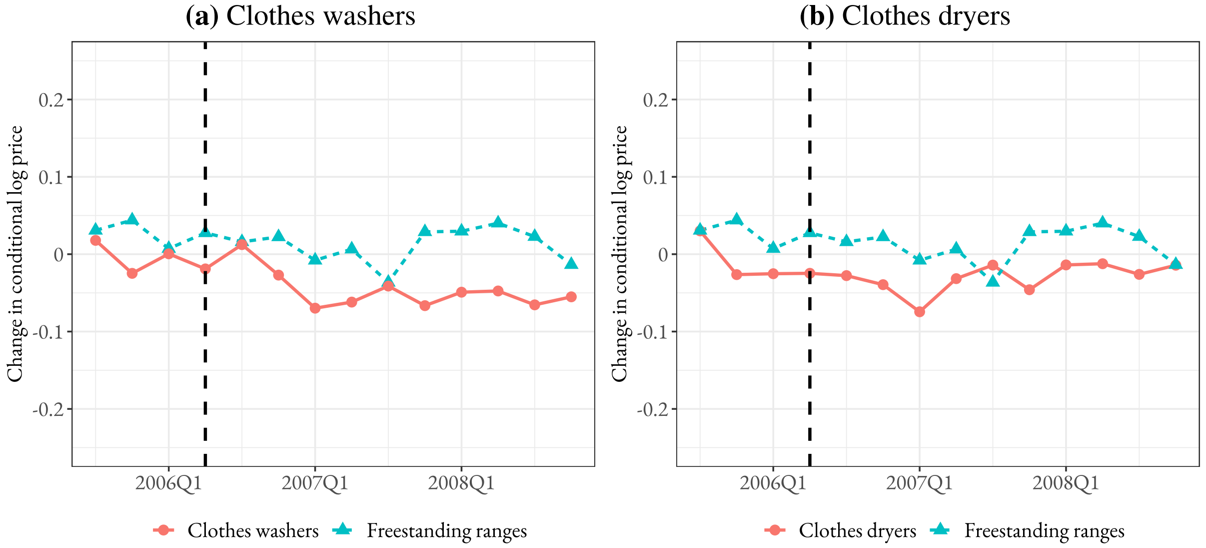

Notes: The solid red line shows the characteristics adjusted log price of Maytag and Whirlpool clothes washers and clothes dryers. The dashed blue line shows the characteristics adjusted log price of competitor freestanding ranges. The vertical line corresponds to the date of the merger, 30 March 2006.

Figure 1.3 plots the quarterly fixed effects γ_t for clothes washers and dryers over time, as well as for freestanding ranges as a control group. The time fixed effects evolve mostly horizontally for freestanding ranges, indicating that there are no important price increases over the observation period. For clothes washers by the merging parties, these are mildly decreasing over time. For

[27] Ashenfelter et al. (2013) show that controlling for product characteristics instead of product fixed effects yields to similar overall time trends, suggesting that there are no important additional unobserved product quality differences.

Table 1.2: Reduced form price effects of the Maytag acquisition

	Washers vs. ranges			Dryers vs. ranges		
	(1)	(2)	(3)	(4)	(5)	(6)
Merging parties × post	-0.030			-0.017		
	[-0.076, 0.016]			[-0.081, 0.046]		
Maytag × post		-0.049**	-0.026		-0.043	-0.015
		[-0.097, -0.001]	[-0.070, 0.018]		[-0.097, 0.011]	[-0.063, 0.032]
Whirlpool × post		-0.016	-0.006		0.007	0.028
		[-0.077, 0.045]	[-0.036, 0.023]		[-0.048, 0.062]	[-0.018, 0.075]
Characteristics controls	Yes	Yes	Yes	Yes	Yes	Yes
Quarter × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Brand fixed effects	Yes	Yes	No	Yes	Yes	No
Product fixed effects	No	No	Yes	No	No	Yes
Observations	3599	3599	3280	4088	4088	3739

Notes: Columns (1) to (3) compare the logarithm of prices for clothes washers and freestanding ranges. Columns (4) to (6) compare the logarithm of prices for clothes dryers and freestanding ranges. Differences in observations in columns (3) and (6) as compared to preceding columns are due to the iterative dropping of singleton observations when clustering standard errors. 95% confidence intervals are reported in parentheses. Standard errors are clustered at the brand level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

dryers, this decrease is less pronounced than for washers. Overall, the descriptive evidence suggests that there are no price increases for either clothes washers or dryers throughout the observation period.

I next estimate the price effects around the time of the merger separately for Maytag and Whirlpool products, using freestanding ranges as a control group. To do this, I estimate the parameters of the following model for washers (treatment) and freestanding ranges (control) and for dryers (treatment) and freestanding ranges (control)

$$log(p_{it}) = \alpha_1 \text{Maytag}_{it} \times \text{post}_t + \alpha_2 \text{Whirlpool}_{it} \times \text{post}_t + \beta x_{it} + \tau_i + \gamma_t + \epsilon_{it} \,. \tag{1.2}$$

The parameters of interest are α_1, which captures the average price increase for Maytag products and α_2, which captures the average price increase for Whirlpool products.

Table 1.2 includes the estimates of the reduced form effects of the Maytag acquisition on the logarithm of prices. Columns (1) and (4) include estimates from a regression where I pool Maytag and Whirlpool products together and estimate a joint price effect. These results suggest that there is no large price increase for clothes washers or dryers. Based on the 95% confidence intervals, I reject price increases of more than 1.6 percent for clothes washers and 4.6 percent for dryers.

In Columns (2) and (5), I disaggregate this by Maytag and Whirlpool products. Based on the

95% confidence intervals, I reject large price increases for Maytag products in both categories. For Whirlpool products, the point estimates are just below (washers) and just above (dryers) zero, however, the width of the confidence intervals do not allow me to reject price changes of between −7.7 and +4.5 percent for clothes washers and −4.8 and +6.2 percent for clothes dryers. In Columns (3) and (6) I repeat the previous analysis, however swapping brand fixed effects for more granular product fixed effects. This leads to a smaller price decrease for merging party products after the merger, but decreases are still found for Maytag clothes washers and dryers and Whirlpool washers.

A causal interpretation of these results could lead to two conclusions: First, the acquisition of Maytag by Whirlpool at most mildly increased prices for laundry products. Second, the acquisition similarly affected clothes washers and dryers.

Irrespective of whether these findings are causal, they are only partially in agreement with the findings by Ashenfelter et al. (2013). In line with their results, I do not find any reduced form evidence for clothes washer price increases around the time of the acquisition. In contrast to their results, I also do not find any reduced form evidence for large price increases for dryers.[28] Given the very similar evolution of market shares and prices for washers and dryers, it seems plausible to expect similar price effects of the merger for both categories. Although I cannot verify this claim, as I do not have access to the *NPD* data used by Ashenfelter et al. (2013), selection in how *NPD* recorded sales could be responsible for the different results.

In any event, the estimated price effects from the reduced form regressions should be interpreted with great caution. As previously described, a causal interpretation of these results requires that prices for laundry products would evolve similarly to prices for freestanding ranges in the absence of the merger. As also noted by Ashenfelter et al. (2013), product entry by LG and Samsung in the market for clothes washers may confound the reduced form estimates of the price effects of the merger. These entries may or may not be related to the merger. Similar market trends may or may not be present for clothes dryers and freestanding ranges.[29]

[28]Compared to ranges, they find an increase in prices for Maytag dryers newly introduced after the merger of 3 percent and of 14 percent for Whirlpool dryers newly introduced after the merger. They also find that the acquisition did not change prices of old Maytag dryers and reduced prices of old Whirlpool dryers by 6 percent. Unfortunately, the data does not allow me to identify when a product was first introduced to the market and so I cannot make this additional decomposition.

[29]Using more or different appliance categories in the control group does not necessarily alleviate the problem, since it remains difficult to establish that the control markets would have developed like the treatment markets in

Finally, the regression analysis does not treat products differently depending on their relative importance in the marketplace (i.e. their market share). Thus, if price changes are not homogeneous across all products, the estimated price changes may strongly be influenced by many products with relatively low market shares. If these are products that most consumers do not consider in any case, this may not be the most informative estimate to assess the price effects experienced by consumers.

1.3.3 Product entry

Rival product entry could affect the estimated price effects of the merger in two distinct ways: First, if the merger leads to merger-specific product entry, this can increase competition and decrease prices compared to a situation without merger-specific entry. Second, if there is merger-independent product entry by rivals into the residential laundry market around the time of the merger, this could also increase competition and reduce prices.

I therefore assess whether product entry by LG and Samsung occurred in the U.S. laundry market and whether this was different to entry patterns for freestanding ranges.

Figure 1.4 shows the evolution of the retailer presence by LG and Samsung for clothes washers, dryers, and freestanding ranges. Since I distinguish between five major retailers and "other retailers", the sum of retailers carrying LG and Samsung appliances can at most be twelve. Two trends emerge: First, the number of retailers carrying LG and Samsung laundry products increases around the time of the merger. By 2008, all major retailers carry LG and Samsung clothes washers and dryers. Second, there is also a strong and persistent increase in the number of retailers carrying LG and Samsung freestanding ranges. Growth is stronger, as it starts from a very low level, however full retailer coverage is only temporarily reached in 2009.

These results suggest that product entry occurred but was not necessarily merger-specific. Indeed, if we believe that merger-independent entry for laundry products is similar to the observed product entry for freestanding ranges, we would expect to observe product entry by LG and Samsung also in the absence of the Whirlpool acquisition.

the absence of the acquisition.

Figure 1.4: Retailer presence LG and Samsung by product category

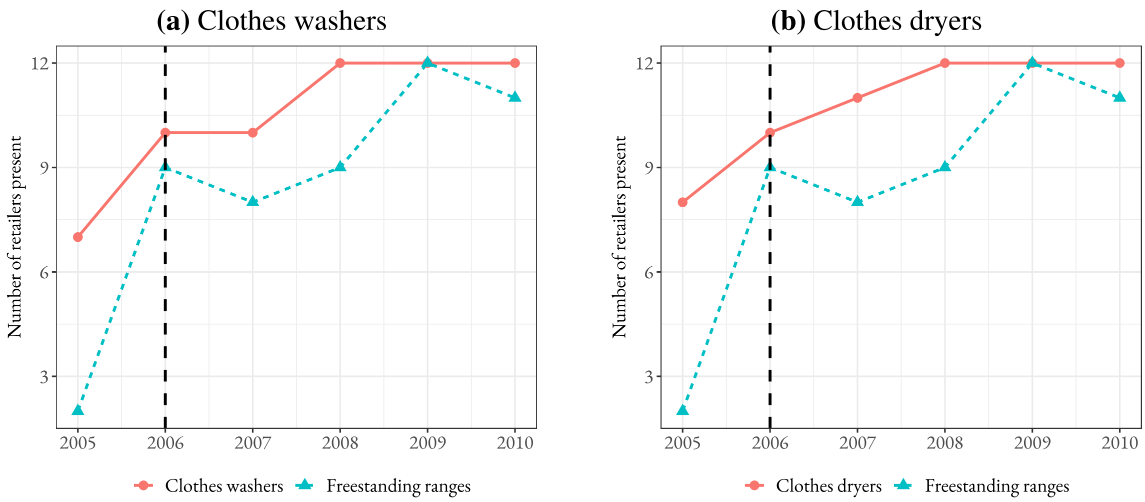

Notes: The solid red lines show the sum of retailers that carry clothes washers (left) or dryers (right) by LG and Samsung summed together. The dashed blue line shows the sum of retailers that carry freestanding ranges by LG and Samsung.

1.3.4 Labor market effects of plant closures

The analysis so far focused on the product market effects of the acquisition. Different acquisitions may also entail different changes to employment. For those to enter the overall welfare effects, appliance manufacturing jobs need to matter for local labor markets. In the following, I assess how Maytag plant closures by Whirlpool post-acquisition affected employment, unemployment, and wages of the employed in affected counties.

Although Whirlpool maintained some of Maytag's manufacturing plants (e.g. in Amana, Iowa, or Cleveland, Tennessee), shortly after the acquisition it shut down appliance manufacturing plants in Searcy, Arkansas (700 manufacturing jobs) and Herrin, Illinois (1,000 manufacturing jobs), as well as manufacturing and headquarter operations in Newton, Iowa (1,000 manufacturing and 1,800 corporate jobs). At the same time, Whirlpool announced adding 1,500 jobs at two existing plants in Ohio.

Figure 1.5 plots the number of employed persons and the unemployment rate in Jasper County, Iowa compared to the mean across other counties in Iowa.[30] Operations began shutting down in Jasper County on 31 December 2006, with manufacturing continuing until 31 October 2007. Already from the descriptive analysis it becomes clear that employment decreased persistently and unemployment shot up around the time of the plant closure and the shut down of corporate

[30]Jasper County is the county in which Newton is located.

operations. The increase in unemployment appears to be persistent and still present at the end of 2008.

Figure 1.5: Labor market effects of the plant and HQ closures in Jasper County, Iowa

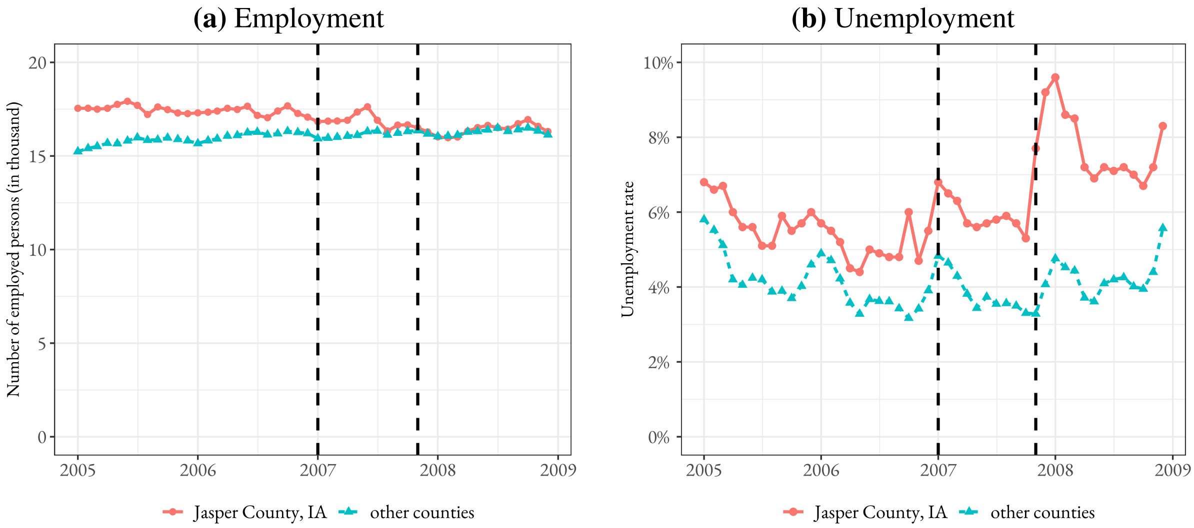

Notes: The solid red lines show the evolution of the total number of employed persons and the unemployment rate in percentages in Jasper County, Iowa (county of the Maytag plant and corporate offices in Newton), respectively. The dashed blue lines show the average number of employed persons and the unemployment rate by county for all other counties in Iowa, respectively. The vertical dashed lines indicate the shut down of operations on 31 December 2006 and 31 October 2007.

I next investigate whether this effect is also present when there are no corporate operations and only a plant closes. Figures 1.6 and 1.7 plot the evolution of employment and unemployment in White County, Arkansas, and Williamson County, Illinois compared to other counties in Arkansas and Illinois, respectively.[31] In both counties, the Maytag appliance manufacturing plants were shut down on the 31 December 2006. In both counties, there is an increase in unemployment in the year after the plant closure. The effect appears more pronounced and persistent for Williamson County. In both counties, the difference in unemployment appears to fade away after a year.

To quantitatively asses the local labor market effects of changes in Maytag employment, I estimate the parameters of the following regression model

$$\text{outcome}_{it} = \alpha_1 \mathbb{1}\,(\text{year}_t = 2007) \times \Delta\text{jobs}_i + \alpha_2 \mathbb{1}\,(\text{year}_t = 2008) \times \Delta\text{jobs}_i + \tau_i + \gamma_t + \epsilon_{it}\,, \quad (1.3)$$

where outcome_{it} is the number of employed persons, unemployed persons, or the average wage

[31] White County is where Searcy is located and Williamson County is where Herrin is located. I omit Cook County (Chicago) from the control group for Illinois.

Figure 1.6: Labor market effects of the plant closure in White County, AR

(a) Employment **(b)** Unemployment

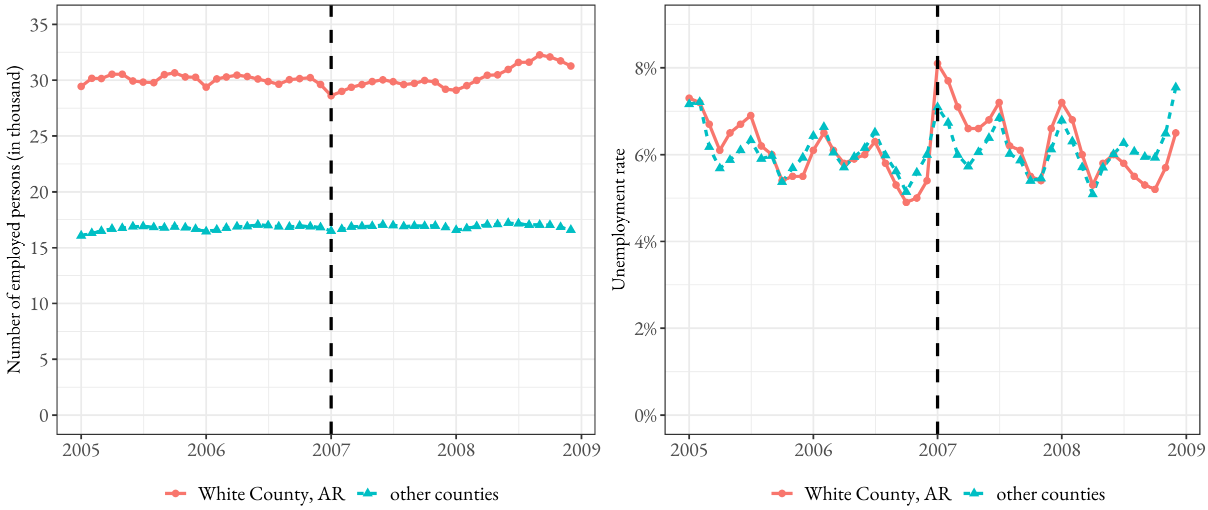

Notes: The solid red lines show the evolution of the total number of employed persons and the unemployment rate in percentages in White County, Arkansas (county of the Maytag plant in Searcy), respectively. The dashed blue lines show the average number of employed persons and the unemployment rate by county for all other counties in Arkansas, respectively. The vertical dashed lines indicate the shut down of operations on 31 December 2006.

Figure 1.7: Labor market effects of the plant closure in Williamson County, IL

(a) Employment **(b)** Unemployment

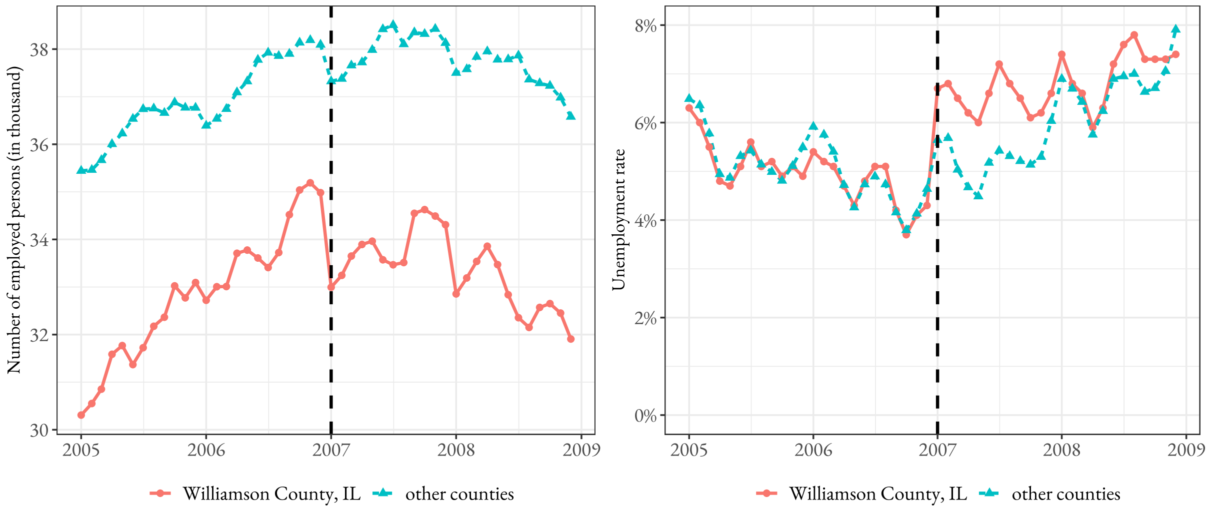

Notes: The solid red lines show the evolution of the total number of employed persons and the unemployment rate in percentages in Williamson County, Illinois (county of the Maytag plant in Herrin), respectively. The dashed blue lines show the average number of employed persons and the unemployment rate by county for all other counties in Illinois (with the exception of Cook County which contains Chicago), respectively. The vertical dashed lines indicate the shut down of operations on 31 December 2006.

Table 1.3: Reduced form labor market effects of plant and HQ closures

	Unemployment (persons)		Employment (persons)		Wages ($)	
	(1)	(2)	(3)	(4)	(5)	(6)
Plant & HQ closure × $\mathbb{1}$ (year = 2007)	163***		-1140***		-2472***	
	[151, 176]		[-1343,-937]		[-2666, -2278]	
Plant & HQ closure × $\mathbb{1}$ (year = 2008)	291***		-1716***		-6508***	
	[263, 319]		[-1983,-1449]		[-6740, -6275]	
Plant closure × $\mathbb{1}$ (year = 2007)		257		-288*		-329
		[-189, 704]		[-597,21]		[-1590, 931]
Plant closure × $\mathbb{1}$ (year = 2008)		8		-336**		-400
		[-545, 561]		[-639,-33]		[-1815, 1014]
County fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,752	8,352	4,752	8,448	1,584	2,816
Mean outcome in treated counties	1,130	1,123	11,840	13,815	34,404	25,524

Notes: Columns (1) and (2) compare the absolute number of unemployed persons in treated counties to all other counties in the same state. Columns (3) and (4) compare the absolute number of employed persons in treated counties to all other counties in the same state. Columns (5) and (6) compare the average annualized gross wage of employed persons in treated counties to all other counties in the same state. Columns (1), (3) and (5) compare Jasper County (county of Newton) to all other counties in Iowa. Columns (2), (4) and (6) compare White County (Searcy) and Williamson County (Herrin) to all other counties in Arkansas and Illinois. Cook County (county of Chicago), is omitted from any analyses involving Illinois. 95% confidence intervals are reported in parentheses. Standard errors are clustered at the county level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

of employed persons in a particular county i and time period t, Δjobs_i is an indicator variable equal to one if a particular county is affected by job cuts or newly created jobs by the merging parties, τ_i are county fixed effects and γ_t are time fixed effects.

I group counties into three different treatment groups and estimate separate regressions for each. The first treatment group is Jasper County, in which there was a shut down of manufacturing and corporate operations. The second treatment group are White County and Williamson County, in which only manufacturing plants were shut down. The third group is Marion County and Sandusky County, where Whirlpool created new jobs.

Table 1.3 summarizes the regression estimates for the elimination of jobs. Column (1) reports the effects on unemployment in Jasper County. I find that there is a statistically and economically significant increase in unemployment. The effect is persistent throughout the observation period, but is small in magnitude (around 300 persons in 2008) compared to the number of Maytag jobs lost ($1,000$ manufacturing and $1,800$ corporate jobs). This however only tells part of the story, as it masks other shifts into non-employment, such as early retirements, exits into education, as well as out-migration. The results in Column (3) show that the number of employed persons in Jasper County as compared to before the closing of operations declined by around $1,700$. Finally, Column (5) shows the effect on annualized average wages of employed persons. Again, there are large and statistically significant decreases in average wages.

Table 1.4: Reduced form labor market effects of new jobs

	Unemployment (persons)	Employment (persons)	Wages ($)
	(1)	(2)	(3)
New jobs × $\mathbb{1}$ (year = 2007)	-33	358	-88
	[-178,112]	[-151,867]	[-412,237]
New jobs × $\mathbb{1}$ (year = 2008)	-230**	656	-271
	[-458,-2]	[-169,1480]	[-1299,758]
County fixed effects	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes
Observations	4,224	4,224	1,408
Mean outcome in treated counties	2,067	27,006	32,452

Notes: Column (1) compares the absolute number of unemployed persons in Marion County (Marion) and Sandusky County (Clyde) to all other counties in Ohio. Column (2) compares the absolute number of employed persons in Marion County and Sandusky County to all other counties in Ohio. Column (3) compares the average annualized gross wage of employed persons in Marion County and Sandusky County to all other counties in Marion County and Sandusky County. 95% confidence intervals are reported in parentheses. Standard errors are clustered at the county level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

In Appendix Tables A.4 and A.6, I estimate wage effects separately for the manufacturing industry and all private sector jobs outside the manufacturing industry. Although wage decreases are much larger for the manufacturing industry, suggesting that many well-paying Maytag jobs were eliminated, I also find significant decrease in wages in other industries. This suggests that Maytag's presence also led to other well-paying jobs and exerted positive wage pressure on the labor market.

Columns (2) and (4) include the effects on unemployment and employment of only shutting down plants without the effects of closing the HQ. There is an economically meaningful but statistically noisy increase in unemployment. This effect however appears to only be transitory and disappears after a year. There is a more robust and persistent decrease in employment in affected counties of around 300 persons. Since the affected plants in the two treatment counties employed 700 and $1,000$ persons respectively, this suggests that around a third to half of the jobs were permanently lost and led to out-migration or other shifts into non-employment, beyond unemployment. The results in Column (6) and in Appendix Tables A.4 and A.6 suggest that there is a significant decrease in manufacturing wages but no effect on wages in other industries.

Table 1.4 shows the effect of relocating $1,500$ new jobs to two existing Whirlpool plants in two different counties in Ohio. On average, this is equivalent to 750 new jobs per affected county. The results in Columns (1) and (2) suggest that these new jobs led to a significant reduction in unemployment in 2008 and an increase in employment. Additional results in Appendix Tables A.5 and A.7 show that this effect is completely driven by an increase in employment in the manufacturing industry and accompanied by a modest decrease in employment in other industries. Wages do not increase, suggesting that these new jobs do not lead to positive wage pressure on the local labor market.

1.4 The Model

Three observations emerge from the preceding analysis. First, entry played a crucial role in the product market and so understanding entry is necessary to assess the effects of the merger. Second, although there were many changes in the product portfolios of existing firms, there was no entry by a new firm. The focus thus lies on endogenous portfolio choices and I abstract from firm-level entry. Third, there are frictions in local labor markets and important differences in the production locations of manufacturers. Where products are produced and by whom affects the welfare effects of the merger.

The model features manufacturers and consumers. Manufacturers choose their product portfolios and prices. Consumers make purchase decisions. The model proceeds in two stages. In the first stage, firms are endowed with a set of potential products that they are technologically capable of producing and their production locations. They observe product-level shocks to entry costs and decide which products to offer. At this stage, firms do not observe transitory demand and marginal cost shocks and only form expectations about these shocks. In the second stage, demand and marginal cost shocks realize and are observed by firms, upon which they set prices. Finally, households observe the products on offer and their characteristics, including prices, and make their purchase decisions. The number of domestic production jobs depends on equilibrium quantities in the product market and the location of production.

I solve this game backwards by searching for the Subgame Perfect Nash Equilibria (SPNE) of the game.[32] To estimate the parameters of the game, I require the existence of a SPNE but not its uniqueness.

[32]Whenever cost or demand shocks are observed by market participants, they remain unobserved by the econometrician.

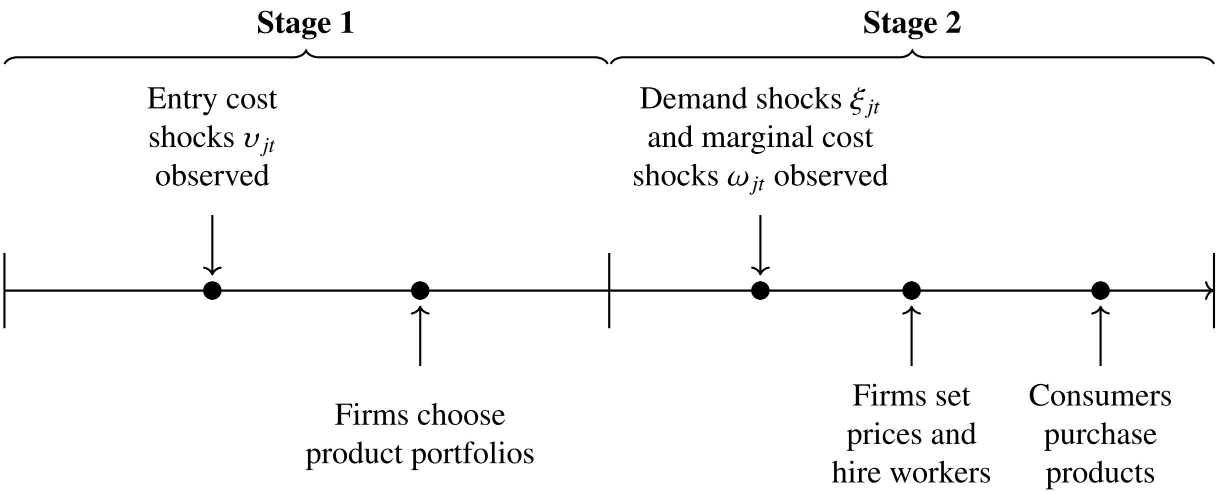

Figure 1.8: Model timeline

1.4.1 Demand model

Demand is a household-level discrete choice between different clothes washers. The demand model is based on the empirical discrete choice demand literature following S. Berry, Levinsohn, and Pakes (1995) and S. Berry, Levinsohn, et al. (2004). Every year, a household makes the choice between different clothes washers on offer in the market as well as not purchasing a clothes washer, i.e. choosing the outside good. This can be thought of as keeping the clothes washer already owned by the household or not owning a clothes washer at all (e.g. using a laundromat).

The utility of household i from buying clothes washer j in year t can be written as

$$u_{ijt} = x_{jt}\beta + \sigma^{FL}\nu_{it}^{FL}x_{jt}^{FL} - (\alpha + \kappa_{\alpha}\min(\$400\text{k}, z_{it}))p_{jt} + \xi_{jt} + \epsilon_{ijt}\,. \tag{1.4}$$

The vector x_{jt} includes non-price product characteristics, such as whether a clothes washer can be loaded from the front, whether it is Energy Star certified, or the number of special programs it includes. It also includes indicator variables for the brand and retailer at which the clothes washer was purchased, as well as year fixed effects and brand time trends.[33] p_{jt} is the price of a

[33]The full list of product characteristics are the price, the brand repair rate, the total advertising expenditure at the brand level, as well as indicator variables for whether a clothes washer is a front-loader, a Korean front-loader, a front-loader by Fisher & Paykel, a high-end European front-loader (i.e. Asko, Bosch, or Miele), has an agitator, is part of a stacked pair, has a stainless steel exterior, has a white exterior, is Energy Star certified, has additional noise insulation, has a child lockout. Finally, it includes retailer, brand and year fixed effects, as well as linear brand time trends.

clothes washer j at time t. I denote the set of products among which households can choose at time t as J_t.

Average tastes for price and non-price characteristics are captured by α and β respectively. x_{jt}^{FL} is an indicator variable for whether a particular clothes washer is a front-loader. ν_{it}^{FL} is an i.i.d. draw from a standard normal distribution and represents a household-specific unobserved taste shock for front-loaders. z_{it} is the income of household i at time t. Household incomes are capped at \$400,000, as this avoids positive price coefficients for households with very high incomes which can arise when income enters the price coefficient linearly. Incomes beyond this threshold have negligible effects on the estimated demand parameters in practice. σ^{FL} measures the dispersion in taste for front-loaders between households. κ_α captures how the sensitivity to prices varies with household income.

The remaining part of the utility function consists of an unobservable component constant across households, ξ_{jt}, as well as an idiosyncratic household-specific unobservable, ϵ_{ijt}. ξ_{jt} includes any remaining quality differences not captured by the product characteristics and fixed effects, as well as transitory demand shocks that vary between products but are common across households. Finally, ϵ_{ijt} is an i.i.d. draw from a type I extreme value (Gumbel) distribution.

To simplify notation, I separate utility into the mean utility δ_{jt} and the household-specific deviation $\mu_{ijt} + \epsilon_{ijt}$. The mean utility includes all utility components that are constant across households. I also define a vector $\theta = (\theta_1, \theta_2)$ which contains all the parameters of the demand model. Let $\theta_1 = (\alpha, \beta)$ contain all linear parameters of the model and $\theta_2 = (\sigma, \kappa)$ all nonlinear parameters. Since I can only identify utilities up to an affine transformation, I normalize the mean utility of the outside good to zero and so the utility of a household for the outside good reduces to ϵ_{i0t}.

The distributional assumptions on the household-specific unobservable allow deriving the familiar logit choice probabilities from this specification. By integrating over the joint distribution of household demographics $P_D(D)$ and the joint distribution of unobserved taste shocks $P_\nu(\nu)$, the model-predicted market share of product j in market t becomes

$$s_{jt} = \int \frac{exp(\delta_{jt} + \mu_{ijt})}{1 + \sum_{k \in J_t} exp(\delta_{kt} + \mu_{ikt})} P_D(D) P_\nu(\nu)\,. \tag{1.5}$$

1.4.2 Second stage: pricing

In the second stage, firms observe demand and marginal cost shocks and subsequently set prices. Each firm f chooses prices for the set of products it offers, J_{ft}, to maximize its variable profits, given by

$$VP_{ft} = \sum_{j \in J_{ft}} (p_{jt} - mc_{jt}) s_{jt} M_t \,, \tag{1.6}$$

where p_{jt} denotes the price of j at t, mc_{jt} its marginal costs and M_t denotes the total market size. Firms set prices by taking first-order conditions of the variable profit function with respect to the vector of prices for the products they are offering. For each product j, the equilibrium price must satisfy

$$p_{jt} = mc_{jt} - [(\nabla_p s \bullet \Lambda)^{-1} s]_{jt} \,, \tag{1.7}$$

where Λ is the ownership matrix and $\nabla_p s$ is the matrix of partial derivatives of market shares with respect to prices.[34]

Marginal costs can be decomposed into several components. In particular, the inverse hyperbolic sine of marginal costs depends on product- and market-specific components in the following way

$$arcsinh(mc_{jt}) = [x_{jt}, ic_{jt}]\gamma + \omega_{jt} \,, \tag{1.8}$$

where ic_{jt} is a vector of input costs, γ captures how product characteristics and input costs affect marginal costs and ω_{jt} is a transitory product-level unobserved marginal cost shock that is realized and observed by the firms in the pricing stage.[35]

1.4.3 First stage: entry

In the first stage, firms decide which products to offer. At the outset each firm is endowed with a set of potential products it can offer in market t, $\mathcal{J}_{ft}$. This can be thought of as the set of products that it is technologically capable of producing. It includes products that it sells already at a different retailer or in a different market and minor adjustments to existing products which

[34] The ownership matrix contains information on whether two products are offered by the same firm and so cross-price effects matter for the optimal pricing decision of firm f.

[35] The inverse hyperbolic sine is a transformation that approximates the natural logarithm. Its advantage is that zero is part of its definition area and it returns real numbers for negative inputs. See Bellemare and Wichman (2020) for more details.

it could perform in the short-term. It does not include products for which a firm would need to develop entirely new capabilities (e.g. launching its first front-loading clothes washer).

Introducing a product into the market comes at a fixed and sunk cost. This includes costs related to the final development of a product (e.g. a particular front-loader model), marketing or retailer investments. Empirically, I analyze markets at the yearly level. At the same time, Ashenfelter et al. (2013) show that the volumes of particular clothes washer models rapidly decline after twelve months. It therefore seems plausible that the fixed and sunk cost of introducing a product at a retailer in a particular year is independent of the product portfolio in previous years, since particular models are usually not kept on shelf for longer.

The fixed cost of introducing a new product can be decomposed into a brand- and market-specific component F_{bt} and a mean-zero idiosyncratic product- and market-specific fixed cost shock υ_{jt}. Thus $F_{jt} = F_{bt} + \upsilon_{jt}$ and $E[\upsilon_{jt}|j \in \mathcal{J}_{ft}] = 0$. Before deciding on its product portfolio, a firm observes the fixed cost shocks related to all products it could potentially add. It does not yet however observe the second stage marginal cost and demand shocks which I summarize as $e_{jt} = (\xi_{jt}, \omega_{jt})$. Instead, it chooses a product portfolio by trading off expected variable profits and the sum of fixed costs of different products. More specifically, it solves the following maximization problem:

$$\max_{J_{ft} \subseteq \mathcal{J}_{ft}} \{\Pi = E[VP(p)|J_{ft}] - \sum_{j \in J_{ft}} F_{jt}\}. \tag{1.9}$$

Since choosing an optimal product portfolio is a discrete choice, the first order conditions of this profit maximization only hold with inequality.

1.4.4 Demand for domestic workers

Let us now turn to the employment side. The aim of this exercise is to model how the number of U.S. manufacturing jobs changes if we hold production locations and the production technology fixed. I therefore do not model demand and supply in the labor market itself. That is not to say that the number of U.S. clothes washer manufacturing jobs would not change if, for example, wages increased. This would be reflected in the marginal costs of a clothes washer and thus affect equilibrium prices and quantities in the product market.

I assume that firms make longer-term decisions on where to produce which products outside of the model. The share of each product that is produced in the U.S. is therefore exogenously given. Similarly, the production technology $G(\cdot)$ is fixed and the number of manufacturing workers required is linear in the number of clothes washers. The demand for domestic clothes washer manufacturing workers by firm f therefore is

$$LD_{ft} = \sum_{j \in J_{ft}} G(q_{jt}) \times \text{domestic}_{jt} \,. \tag{1.10}$$

1.5 Estimation

In this section, I describe how to estimate the parameters of the model. As for the model description I proceed in reverse-order, beginning with the demand parameters.

1.5.1 Demand

The estimation of the demand parameters is similar to S. Berry, Levinsohn, et al. (2004). In a first step, I estimate the non-linear parameters of the utility function, σ^{FL} and κ_α. I identify these parameters by matching simulated moments to their analogues in the data. Informally, we can think of the data moments as identifying the structural parameters of their simulated equivalent.

The first data moment is based on the correlation between the clothes washer bought being a front-loader and the average share of front-loaders among the second-choice brand. Respondents to the *TraQline* survey are only asked which other brands they considered buying but not which exact model. Some brands carry both front-loaders and top-loaders. However, the share of front-loading clothes washers differs greatly between brands. Furthermore, the correlation between whether the first choice is a front-loader and the share of front-loaders among the second choice brand is important, with a correlation coefficient of 0.4. This suggests that there is a strong unobserved taste for front-loaders among some households, which can affect substitution patterns.

The second data moment is based on the correlation between the household income and the price of a clothes washer bought. Figure 1.9 shows the correlation between the household income and price. On average, the higher the income of a household, the higher the price of a clothes washer bought. This suggests that high income households are less sensitive to prices.

Figure 1.9: Correlation of average purchaser household income and price by product

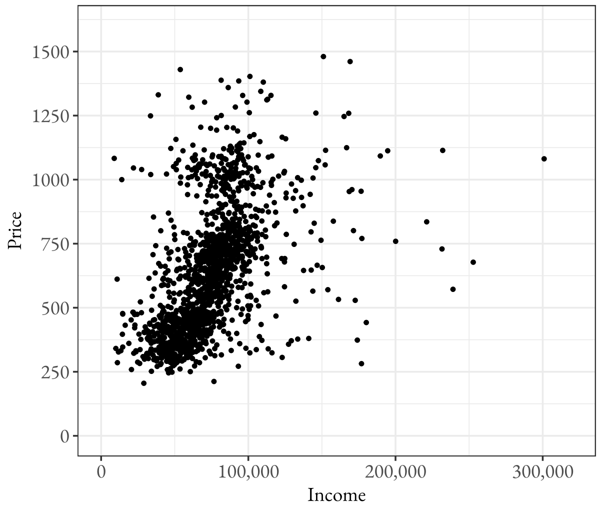

Notes: The plot shows the average annual income of households purchasing a particular clothes washer on the x-axis and the average price of that clothes washer on the y-axis. Each point is a product in a particular year. The correlation coefficient between the average income of households purchasing a particular clothes washer and its average price is 0.5.

To estimate the linear parameters of the utility function I first need to estimate the vector of mean utilities, δ. I estimate these by matching simulated market shares for each product to observed market shares. Before estimating the linear utility parameters α and β, I need to introduce a further assumption:

Assumption 1.1. $E[e_{jt}|X_{jt}, F_{jt}] = 0$ *for each* $j \in \mathcal{J}_t$.

This means that the second stage demand and marginal cost shocks are independent of the non-price product characteristics and the fixed costs of introducing a product. As explained by Eizenberg (2014), this is slightly stronger than the assumption that e_{jt} is realized after products are chosen, since it also means that firms cannot predict e_{jt}. This assumption nevertheless seems reasonable, as firms may still predict future costs and demand as they relate to observable characteristics, which I can control for. It only means that firms cannot predict unobservable transitory marginal cost and demand shocks.

Since prices can be adjusted frequently, they are likely correlated with ξ_{jt}. As explained in Section 1.2, I use an instrumental variable based on the production location and the real exchange rate, which affects costs but is otherwise unrelated to demand.

For the linear and non-linear demand parameters, standard errors are clustered at the brand level using the residual bootstrap. First, I estimate the linear and non-linear demand parameters using the original sample. Second, I compute the empirical distribution of demand residuals for every brand. Third, I re-sample demand residuals for every product from the empirical distribution of demand residuals for the respective brand, creating bootstrapped samples. Fourth, I re-estimate the linear and non-linear demand parameters for 100 bootstrapped samples. Finally, I compute the standard error of the parameter estimates using the bootstrapped samples.

1.5.2 Marginal costs and efficiencies

I compute marginal costs for each product by inverting the first order conditions of each firm's profit maximization problem. Under the model assumptions described above, the data are rationalized by a unique marginal cost and markup for each product.

Next, I estimate the parameters of the second-stage supply-side in Equation 1.8. This estimation serves two purposes. First, it allows me to residualize the inverse hyperbolic sine of marginal costs by the effect of product characteristics and input costs. I can hence split marginal costs into a part that is known to firms when making product entry decisions and the unobserved marginal cost shocks ω.

Second, it provides a data-driven estimate of the marginal cost efficiencies due to changes in the exchange rate and the cost level in different production locations. Since Equation 1.8 contains the relationship between the real exchange rate and marginal costs, it also allows me to estimate how marginal costs change if the RER changes. By using the post-merger production locations for Maytag products under the two acquisition scenarios (acquisition by Whirlpool and acquisition by Haier), I can estimate how these changes in production locations would affect the RER for Maytag products and their marginal costs.

For the Whirlpool acquisition, I use the observed post-merger production locations for Maytag products by Whirlpool in 2007. Marginal cost efficiencies in this case come from the relocation of front-loader production to Mexico. For the Haier acquisition, I use the publicly discussed relocation plans of Maytag's production to China.

1.5.3 Fixed cost bounds

The entry model in Section 1.4.3 only provides inequality conditions for profitable entry. It is hence not possibly to point identify entry costs in this setting. I therefore resort to partial identification and seek to estimate bounds on the identified set of fixed entry costs for every brand.

To estimate bounds on the fixed costs of adding a product, I need to determine the set of potential products of each firm. I refer to all products that a firm could have added as the potential products, to the potential products that it actually added as the active products and to the potential products that it chose not to add as the inactive products. Recall that the set of potential products of firm f in market t is denoted as $\mathcal{J}_{ft}$ and the set of active products as J_{ft}. I denote the set of inactive products of firm f as $\tilde{J}_{ft}$.

The set of active products are those products that we observe in the data. Before determining the set of inactive products, it is worth remembering that the goal is to estimate the fixed costs of adding or removing a product that is part of the set of products a firm is technologically capable of producing. Thus, if a firm does not have any front-loading washing machines among its active products, I do not consider that it could have added a front-loading washing machine in that particular year. Instead, I exploit the fact that I can distinguish sales at the retailer level and that appliance brand owners introduce different products at different retailers. For any active product (e.g. a front-loader by KitchenAid sold at Sears), all versions of the product that I do not observe in the data (e.g. a front-loader by KitchenAid sold at Best Buy, H. H. Gregg, Home Depot or Lowe's) is an inactive product. I therefore capture the fixed costs related to marketing, getting retail floor space for an additional product or customizing the product for the clientele of a particular retailer, but not of developing new technologies. This is appropriate in this case, since I am interested in estimating how the incentives to make portfolio adjustments change for existing players with already developed product portfolios. Furthermore, the development of new technologies is most likely a multi-year process that does not need to pay off within a year.

The estimation of the bounds on fixed costs resembles the procedure described by Eizenberg (2014). If the product entry that I observe is a pure strategy SPNE, then no firm can profitably deviate unilaterally from this equilibrium. More specifically, this means that no firm can increase its expected profits by unilaterally adding inactive products or removing active products.

To estimate bounds on the fixed costs of adding a product, I exploit a subset of the equilibrium conditions, namely that no firm has a profitable one-step deviation.[36]

Let us denote the equilibrium product portfolio (i.e. the set of active products) of firm f at time t as J^*_{ft}. For each active product j that a firm chooses to introduce in equilibrium, an upper-bound on the fixed cost of introducing a product is the expected incremental profit of offering that product holding other products fixed. That is,

$$F_{jt} \leq E_e[VP_{ft}(J^*_{ft}) - VP_{ft}(J^*_{ft} - \mathbf{1}^j_{ft})] \equiv \bar{F}_{jt}\,, \tag{1.11}$$

where $\bar{F}_{jt}$ is the upper-bound on fixed costs of adding product j at time t.

For each inactive product, a lower-bound on the fixed cost of introducing a product is the expected incremental profit of offering that product holding other products fixed. That is,

$$F_{jt} \geq E_e[VP_{ft}(J^*_{ft} + \mathbf{1}^j_{ft}) - VP_{ft}(J^*_{ft})] \equiv \underline{F}_{jt}\,, \tag{1.12}$$

where $\underline{F}_{jt}$ is the lower-bound on fixed costs of adding product j at time t.

These two conditions allow estimating the upper-bound on fixed costs of active products and the lower-bound on fixed costs of inactive products. I estimate the expected incremental variable profits using 500 draws from the joint empirical distribution of the demand and marginal cost shocks e_{jt}. Ultimately, I am interested in bounds on the brand-level average fixed costs in market t, F_{bt}. Constructing the upper-bound on F_{bt} only based on active products and the lower-bound based on inactive products is inadmissible, since product portfolio decisions are not independent of υ_{jt}, i.e. $E[\upsilon_{jt}|j \in J_{ft}] \neq 0$. Recall, however, that $E[\upsilon_{jt}|j \in \mathcal{J}_{ft}] = 0$, which means that the product-level fixed cost shock has mean zero conditional on products being part of the set of potential products. This means that if we can estimate a lower-bound on the fixed costs of adding active products and an upper-bound on the fixed costs of adding inactive products, we can get an unbiased estimate of bounds on the set of brand-level average fixed costs F_{bt}.

To fill the missing bounds, I follow the approach proposed by Eizenberg (2014). The details of the estimation procedure are described in Appendix A.3.4.

[36]In principle, I could add further restrictions on fixed cost bounds due to the lack of profitable multi-step deviations. In practice, restrictions based on multi-step deviations may be difficult to use, since the additional inequalities would include idiosyncratic fixed cost shocks υ_{jt} for each product.

For inference, I estimate the confidence sets using the same bootstrapped samples as for the demand estimation. D. W. K. Andrews (2000) shows that the bootstrap is inconsistent if the parameter is on the boundary of the parameter space, which is defined by inequality conditions. I however use this procedure as a first-order approximation of the consistent confidence sets.

1.5.4 U.S. employment

Estimating the equilibrium number of U.S. clothes washer manufacturing jobs under different scenarios requires the overall number of employees necessary to manufacture clothes washers in each scenario, as well as the corresponding production locations.

Recall that I assume that the number of employees necessary for the manufacturing process is directly proportional to the number of clothes washers sold. To simplify estimation I also assume that the production technology is linear and constant across products and manufacturers. I use information on the number of employees and clothes washer production from annual reports and news articles to calibrate how many clothes washers a manufacturing worker produces per year on average. I combine this number with the equilibrium quantity of clothes washers sold for each product, to estimate how many manufacturing jobs are necessary globally.[37]

The second step is to estimate the share of clothes washers that are produced in the United States. As described in greater detail in Section 1.2.3, I construct a granular data set that contains product-level information on the production location of clothes washers produced for the U.S. market. The equilibrium number of U.S. clothes washer manufacturing jobs is therefore the share of global manufacturing jobs multiplied by the share of a product's U.S. production.

To estimate how U.S. employment differs between acquisitions of Maytag by Haier or Whirlpool, I assume that Haier would offshore all Maytag jobs to China; whereas I use the observed post-merger production locations by Whirlpool after its acquisition of Maytag. The latter is necessary to also account for the partial offshoring of former Maytag manufacturing jobs by Whirlpool. Without doing so, I would overestimate the number of jobs maintained by Whirlpool.

[37] I describe this calibration in more detail in the Appendix Section A.3.5.

1.6 Estimation Results

1.6.1 Demand

Table 1.5 includes the demand estimates. Column (1) reports the first-stage results, where I regress the endogenous price variable on the instrumental variable (IV) for price, which is the real exchange rate, and include full controls. The results indicate that an increase in the RER by a full unit leads to an increase in clothes washer prices by \$191. The F-statistic is approximately 23, suggesting that the IV is relevant.[38]

Column (2) includes the reduced form estimates after regressing the outcome variable (the average utility that consumers get from purchasing clothes washer j at time t, δ_{jt}) on the instrument. As expected, the higher the RER, the lower the purchasing utility for a consumer. In Columns (3) and (4), I report the price coefficient for the simple logit demand model using OLS and the IV, respectively. By accounting for the endogeneity of prices, the average product-level own-price elasticity of residual demand changes from -0.96 to -2.42. Finally, I report the price effects for the full mixed logit model using IV in Column (5). The results suggest that there are significantly heterogeneous but correlated preferences across households. As expected, households with a higher household income are less sensitive to prices. Furthermore, households that purchase front-loaders also have an above average unobserved preference for other front-loaders. Accounting for these effects, I estimate that the average own-price elasticity of residual demand for clothes washers further reduces to -3.26.[39]

1.6.2 Marginal cost

Figure 1.10 shows the product-level marginal costs and the Lerner Index for the full observation period. I find that average marginal costs are around \$410 and range between close to zero and around \$1,500. The average Lerner Index in the sample is 40 percent.

[38] I follow I. Andrews, Stock, and Sun (2019), who recommend reporting the effective first-stage F-statistic due to Olea and Pflueger (2013) in cases with a single endogenous regressor. This is equivalent to the Kleibergen-Paap F-statistic in just-identified cases. In the just-identified case with a single endogenous regressor, we can also compare the F-statistic to the J. Stock and Yogo (2005) critical values.

[39] These elasticity estimates are comparable in magnitude to results by S. Houde (2018), who finds short-term own-price elasticities of residual demand for refrigerators of between -5.41 and -4.15, depending on household income and using weekly data.

Table 1.5: Demand estimates

	(1)	(2)	(3)	(4)	(5)
	First-stage	Reduced form	Logit OLS	Logit IV	Mixed logit IV
Dependent variable:	Price	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$	
Linear parameter					
Real exchange rate	1.909***	-0.787**			
	(0.398)	(0.358)			
Price ('00 2012 $)			-0.164**	-0.412**	-0.614***
			(0.062)	(0.202)	(0.024)
Non-linear parameter					
Income effect κ_{α}					0.070***
					(0.011)
Unobserved taste σ^{FL}					2.425***
					(0.016)
Characteristics	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Retailer FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Brand FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Brand time trends	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Year FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Observations	1,590	1,590	1,586	1,590	1,590
Kleibergen-Paap F-statistic	22.979				
Avg. own-price elasticity			-0.964	-2.416	-3.258

Notes: Column (1) presents results for the first stage regression of prices on the real exchange rate. Column (2) includes reduced form estimates for the simple logit model. Column (3) reports demand estimates for the simple logit without instrumenting for price. Column (4) presents demand estimates for the simple logit model using the RER as an instrumental variable for price. Column (5) shows demand estimates for the full mixed logit model presented in Section 1.4 and using the RER as an instrumental variable for price. For the mixed logit IV model, κ_{α}, σ^{FL}, and $\hat{\delta}_{jt}$ are estimated using simulated method of moments. The remaining linear parameters are estimated using linear IV regression. Standard errors are clustered at the brand level. The own-price elasticity of residual demand is computed at the product level and the average is calculated by weighting products according to their sales volume. Estimates for non-price characteristics are reported in Table A.8. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 1.10: Product-level distribution of marginal costs and Lerner Index

(a) Marginal costs ($) **(b)** Lerner Index (%)

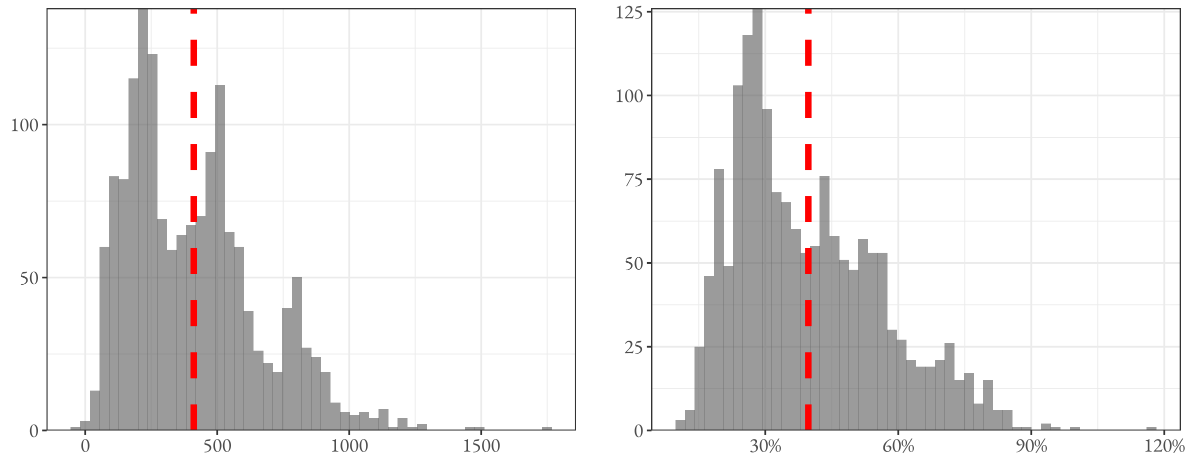

Notes: The histogram on the left depicts the distribution of product-level marginal cost (in $) estimates. The dashed red line indicates the average marginal costs. The histogram on the right shows the distribution of the product-level Lerner Index (markup over price). The dashed red line represents the average Lerner Index. Estimates are for the full observation period between 2005 and 2015.

Figure 1.11 shows the evolution of marginal costs and the Lerner Index by brand owner over time. After the removal of Maytag as a competitor marginal costs decreased. These decreases in marginal costs cannot all be merger-specific, as we would not expect to see marginal cost efficiencies for competitors as a result of the merger. At the same time profit margins increased.[40] This is also true for Whirlpool, although its Lerner Index after 2006 also includes Maytag, which had a significantly lower Lerner Index than Whirlpool pre-merger.

The inverse hyperbolic sine of marginal costs increases by 0.163 if the real exchange rate increases by one unit. Since the inverse hyperbolic sine is similar in form to the logarithmic function, we can approximately interpret this as meaning that marginal costs increase by 16.3 percent if the RER increases by one unit. Table 1.6 summarizes the estimated average marginal cost efficiencies due to offshoring between the two acquisition scenarios, disaggregated by Maytag's brands. These results suggest only modest marginal cost efficiencies after a Whirlpool acquisition and marginal cost efficiencies of up to 12.2 percent for some brands after a Haier acquisition.

This estimation is most likely too unfavorable to Whirlpool and too favorable to Haier. As

[40] Large movements in the Lerner Index for Samsung between 2006 and 2007 should be interpreted with caution, since these are based on relatively few Samsung products at the time.

Figure 1.11: Evolution of marginal cost and Lerner Index by brand owner

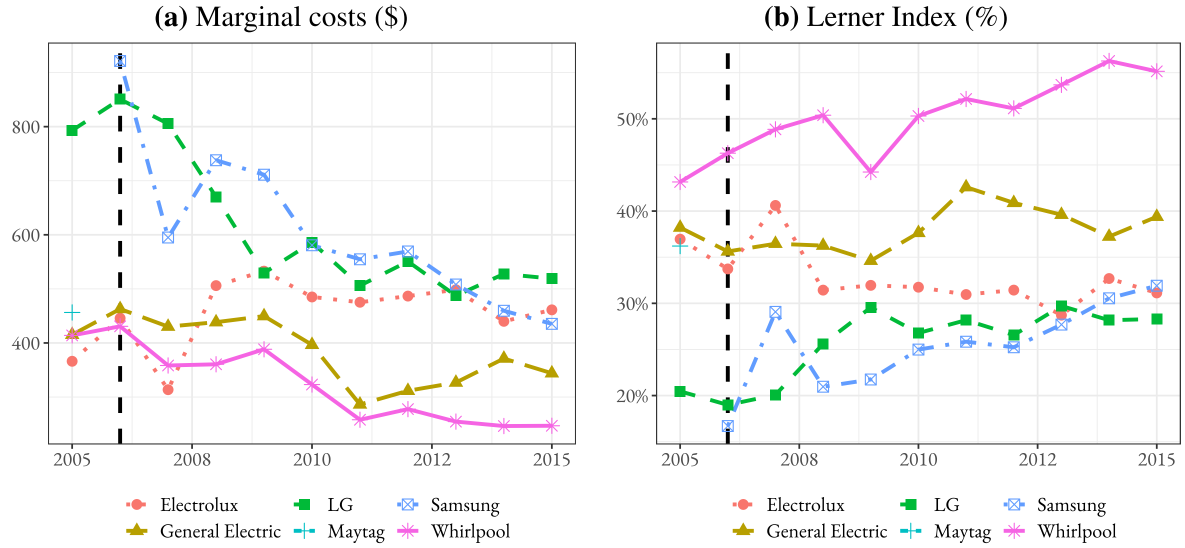

Notes: The plots show the evolution of marginal costs (left) and the Lerner Index (markup over price; right) by brand owner over time. The vertical line shows the time of the Maytag acquisition by Whirlpool. From 2006 onwards, Whirlpool also includes former Maytag products.

Table 1.6: Marginal cost efficiencies due to offshoring by brand (%)

Brand	Whirlpool acquisition	Haier acquisition
Admiral	0.0	12.2
Amana	1.4	11.7
Magic Chef	0.0	11.8
Maytag	2.1	10.7

Notes: The table includes estimated marginal cost efficiencies for Maytag products that arise due to offshoring in the two acquisition scenarios. Offshoring efficiencies are based on changes in the real exchange rate due to changes plant network between acquisition scenarios and the relationship between marginal costs and the RER.

previously discussed, Whirlpool made many changes to Maytag's production in the U.S. and cut costs without offshoring most of the production abroad. This would therefore lead to marginal cost efficiencies that are not reflected in the RER. In contrast, the increase in transportation costs and time-to-market which would arise from offshoring production to China is not factored in and is likely to diminish marginal cost efficiencies for Haier. In the alternative specification, I therefore set marginal cost efficiencies for both acquisitions to zero. This is likely too favorable to Whirlpool (relative to a Haier acquisition). As demonstrated by the general push towards offshoring at the time of the merger, the location-dependent production costs are too important for it to seem plausible that Whirlpool could completely offset Haier's advantage of producing in China by increasing the efficiency of Maytag's U.S. operations. These two marginal cost scenarios should therefore bound the true trade-off between the two acquisitions.

1.6.3 Fixed cost bounds

Finally, I estimate bounds on the fixed and sunk costs of product entry at the brand-level. Before interpreting these results, it is worth remembering that a product is defined as the combination of a brand, a retailer and major clothes washer characteristics (i.e. the distinction between front-loaders, regular top-loaders and high-efficiency top-loaders). Thus, the fixed cost sets that I estimate should be thought of as the cost of adding a product category (brand and major characteristic combination) at a particular retailer. In practice, this may be the more appropriate way of economically modeling product entry, since, for example, the marketing and sales costs of adding another slightly different Whirlpool front-loader at Sears are likely very low if Sears already offers a Whirlpool front-loader.

Table 1.7 describes the 95 percent confidence sets on the fixed costs of adding new products. As expected, I find that the range of plausible fixed costs to add products involves higher values for brands with large market shares (e.g. Maytag or Whirlpool) than brands with lower market shares (e.g. KitchenAid, Hotpoint or Westinghouse). This could be because the former are only offered at a retailer if this involves a full range of clothes washers within that product category, requiring more floor space as well as higher marketing expenditures.

Table 1.7: Brand-level fixed costs of adding a product ($M)

Brand owner	Brand	95 % confidence sets
Maytag	Admiral	[1.4, 1.8]
	Amana	[1.0, 2.8]
	Maytag	[5.9, 34.8]
Whirlpool	KitchenAid	[0.4, 1.0]
	Roper	[0.5, 3.4]
	Whirlpool	[3.9, 33.1]
General Electric	General Electric	[1.7, 20.5]
	Hotpoint	[0.3, 1.6]
Electrolux	Frigidaire	[1.8, 8.5]
	Westinghouse	[0.3, 1.3]
LG	LG	[1.5, 25.3]
Samsung	Samsung	[1.8, 10.3]

Notes: Brand-level fixed costs of adding or removing a product are based on all active and potential products in 2005 (pre-merger) and 2007 (post-merger). Brand owners listed in the table are based on pre-merger ownership of brands.

1.7 Welfare Effects of the Whirlpool Acquisition

In this section, I combine all of the estimation results so far, to compare the welfare effects of an acquisition of Maytag by Haier to the welfare effects of an acquisition by Whirlpool. In particular, I will focus on assessing how endogenous product portfolio choices, marginal cost efficiencies, and the inclusion of employment effects into the welfare assessment change the relative desirability of each acquisition. Since Haier had close to no presence in the U.S. laundry market prior to the merger, an acquisition by Haier without the marginal cost efficiencies is approximately equivalent to keeping a standalone Maytag in the product market.

1.7.1 Players and potential products

Endogenizing portfolio choices requires deciding who can add products, as well as estimating the set of potential products that players could add.

With endogenous product portfolio choices, I allow Electrolux, General Electric, LG, Maytag, Samsung, and Whirlpool to choose their product portfolios. These are all of the clothes washer manufacturers with a volume share of more than 3 percent. I do not endogenize portfolio decisions for Sears for two reasons: First, since Sears would not introduce Kenmore appliances to other major retailers, it could not react to the increase in prices or removal of products by the merging parties at Home Depot by introducing new Kenmore products at Home Depot. It also means that I do not observe any inactive products and cannot estimate the fixed cost of introducing new products for Sears. Second, since Sears does not manufacture products itself, it can only react by sourcing new products from existing manufacturers. The fact that Sears sourced all of its clothes washers from Whirlpool at the time of the merger might make it particularly difficult to react to price increases and portfolio changes of the merging parties by introducing new products in the short-run. To ease the computational burden, I fix the product portfolios of very small competitors. In practice this should only have minor effects on the welfare estimates.

The set of potential products of each player consist of the sum of their potential products in 2005 (pre-merger) and 2007 (post-merger).[41] Since I observe the acquisition scenario with the highest increase in market power in the data, these observed sets of potential products should be a good approximation of the actual set of potential products. This is because the higher the increase in market power, the lower the intensity of competition becomes and so the higher the incentives are for rivals to add new products. Thus, any product that was not added by rivals after the merger is also unlikely to have been added without the merger. Similarly, the incentives of adding new products was highest for the merging parties pre-merger. Thus, the pre-merger set of potential products should be a good proxy for the actual set of potential products of the merging parties.

Finally, I fix the products of players at smaller retailers that are not part of the five major retailers. This results in 135 potential products for the major manufacturers listed above and 69 exogenously active products (products of non-players and products of players at smaller retailers).

[41] Whenever different versions of the same product exist for 2005 and 2007, I choose the 2007 version of the product.

1.7.2 Portfolio choice algorithm

A well-known feature of product entry games is that there can be many potential equilibrium product portfolios. One way of identifying the set of potential equilibria is to estimate the expected variable profits for all possible product entry combinations and then check whether there are any combinations of product entry costs contained in the fixed cost confidence sets that make these product portfolios a SPNE of the entry game.[42] In this case, this is computationally infeasible at this time, since there are 2^{135} candidate equilibria. I instead leverage specificities of the case at hand to construct a heuristic portfolio choice algorithm. This algorithm is most closely related to the heuristic algorithm by Fan and Yang (2020).

First, I recognize that although firms incur the fixed cost of adding a potential product to their active portfolio every year, they do not start in a vacuum. More specifically, if there are multiple equilibria of the post-merger entry game it appears plausible to assume that equilibria closer in the product space to the pre-merger product portfolios are more likely to be realized. Thus, I initialize the portfolio choice algorithm at the pre-merger equilibrium.

After initializing the algorithm, there is an inner and an outer optimization loop to find a one-step equilibrium in portfolio choices. In the inner loop, a particular player computes both the expected change in firm-level profits of adding each inactive product separately to the brand's product portfolio, as well as the expected change in firm-level profits of removing each active product separately. If there is at least one profitable one-step deviation, the player implements this deviation and changes her product portfolio accordingly. I repeat this process until the player has no profitable one-step deviation left. In the outer loop, I repeat this process for each player. The pseudo-code in Appendix A.5 illustrates the steps of the portfolio choice algorithm.

In practice, I can considerably reduce the computational burden by optimizing product portfolios brand-by-brand instead of firm-by-firm. This requires computing fewer potential one-step deviations for every portfolio adjustment. Although I fully take into account how the introduction or removal of a product impacts the firm's expected profit (and not just that of the brand), the downside to this approach is that if products of two brands of the same firm are very close substitutes, the order of play could matter for which product enters. This is unlikely to play an

[42] This is the approach taken by Eizenberg (2014) in a setting where there are four brands and four product types. After adding some additional restrictions, he ends up with $2^9 = 512$ candidate equilibria.

important factor, as firms segment their products by brands and so products within a brand are much closer substitutes than between brands of the same firm.

Another way in which I reduce the computational burden is by only considering one-step deviations and disregarding multi-step deviations. This is necessary because checking for any multi-step deviations is also computationally infeasible in this case.[43] It could thus be that although there is no profitable one-step deviation, there nevertheless exists a profitable multi-step deviation. To assess whether this could be an important problem, it is helpful to consider when such a situation could arise. Since clothes washers are substitutes in the marketplace, if it is not profitable to add a particular clothes washer, it is also not profitable to add that and another potential clothes washer. The same logic applies to the removal of active clothes washers from the product portfolio. It is, however, possible that although adding a particular clothes washer is not profitable, it would be profitable to add the clothes washer and remove another washer from the product portfolio simultaneously. Similarly, it could be that it is profitable to add a clothes washer and remove two washers simultaneously. Overall however it may not be desirable to consider multi-step deviations with many different portfolio adjustments simultaneously, since it is more difficult to make many portfolio adjustments at the same time.

Finally, as I only set identify fixed costs, I repeatedly apply the portfolio choice algorithm for 50 different fixed cost draws for each product. I draw fixed costs from a uniform distribution, where the domain are the confidence sets of fixed costs for each brand. In all scenarios, I report 95 percent confidence sets for the welfare effects across fixed cost draws.

1.7.3 Product portfolio choices

I begin by comparing the product market effects of alternative acquisitions of Maytag by Haier or Whirlpool. Since Haier had close to no U.S. presence prior to the merger, in the product market, an acquisition by Haier without marginal cost efficiencies is approximately equivalent to keeping a standalone Maytag. Whenever I discuss the endogenous choice of the product portfolio, this means that all players can choose active products among the potential products. Since I do not observe realized demand and supply shocks for potential products, I estimate the

[43] To illustrate this point, brands have up to 15 potential products. Checking for all multi-step deviations would thus require checking up to $2^{15} = 32,768$ candidate deviations at each brand iteration.

Table 1.8: Number of products offered by each firm in different acquisition scenarios

	Endogenous portfolio adjustments					**No adjust.**	**Indep. adjust.**
Efficiencies:	No efficiencies		Offshoring				
Acquirer:	None / Haier	Whirlpool	None	Haier	Whirlpool		
Maytag	[21.9, 27.5]	[19.8, 25.4]	[21.5, 26.9]	[23.2, 28.2]	[19.7, 25.5]	21	23
Whirlpool	[27.8, 32.8]	[25.1, 30.6]	[27.5, 33.1]	[27.3, 32.1]	[25.5, 30.7]	27	27
LG + Samsung	[4.6, 10.3]	[5.0, 11.4]	[4.7, 10.7]	[4.6, 10.0]	[5.1, 11.1]	5	15
Electrolux + GE	[25.6, 34.2]	[27.2, 35.1]	[25.8, 34.2]	[24.8, 33.6]	[27.4, 35.4]	34	38
Total industry	[106, 117]	[103, 115]	[105, 117]	[105, 117]	[104, 114]	106	128

Notes: The first five columns include the 95% confidence sets on the number of products carried by each brand owner depending on who acquires Maytag and whether there are offshoring efficiencies. The final two columns show the number of products had there been no product portfolio adjustments (i.e. observed pre-merger portfolios in 2005) and had all portfolio adjustments been merger-independent, thus always leading to the same post-merger portfolios (i.e. observed post-merger portfolios in 2007). Confidence sets for the expected number of products offered with endogenous portfolio adjustments are based on 50 fixed cost draws for each potential product from a uniform distribution, where the domain are the confidence sets of brand-level fixed costs, and 500 demand and supply residual draws. Maytag includes all products marketed under the brands owned by Maytag pre-acquisition (i.e. Admiral, Amana, Magic Chef and Maytag).

expected welfare effects based on 500 demand and supply residual draws for each product.

Table 1.8 summarizes the number of products that firms choose to offer under different acquisition and marginal cost efficiency scenarios. As discussed in Section 1.5, I distinguish between a scenario in which there are no marginal cost efficiencies after either acquisition and a scenario where I credit Maytag products offshoring efficiencies, due to changes in the real exchange rate of the production locations that are chosen by the respective acquirer.[44] The final two columns show the observed number of products in 2005 and 2007. I consider the former to be the relevant product portfolios of firms under the assumption that there are no portfolio adjustments around the time of the merger. I consider the latter to be the relevant product portfolios of firms under the assumption that there are portfolio adjustments around the time of the merger, but that these adjustments are all independent of any acquisition.

Without marginal cost efficiencies, I find that a Maytag acquisition by Whirlpool leads to fewer products offered belonging to the pre-merger Maytag and Whirlpool. At the same time, for most fixed cost draws, there are slightly more products offered by rivals. Overall, there are fewer products offered.

Next, I compare the differences in endogenous portfolio adjustments between the two potential acquisitions to the differences in observed product portfolios between 2005 and 2007. This comparison is possible, because without efficiencies, in the product market an acquisition by

[44] An overview of the offshoring efficiencies for each acquisition scenario can be found in Table 1.6.

Haier is very similar to no acquisition at all (i.e. the pre-merger market in 2005). I observe a slight increase in the number of products offered by Maytag and a strong increase of products offered by rivals.[45] However I predict a decrease in the number of products offered by Maytag and Whirlpool and a smaller than observed increase of products offered by rivals. This suggests that the observed product portfolio adjustments are only partially driven by the merger and partially driven by other effects, such as an exogenous expansion in the set of potential products. A mere comparison of the market structure pre- and post-acquisition would underestimate the reduction in products by Maytag and Whirlpool and overestimate the expansion of the portfolio by rivals.[46]

1.7.4 Product market effects

Table 1.9 summarizes the product market effects of Maytag acquisitions by Haier and Whirlpool.[47] I begin by considering the effects of Maytag acquisitions using the pre-merger product portfolios and neglect any type of product portfolio adjustments. Without marginal cost efficiencies, prices after a Whirlpool acquisition increase by 2.7 percent and consumer welfare decreases by 4.9 percent. Total industry profits, as well as the profits of the merging parties, increase, however the increase in profits cannot offset the loss in consumer welfare. With offshoring efficiencies, an acquisition by Haier reduces average industry prices by 1 percent and increases consumer welfare by 3.1 percent, since there is close to no increase in market power. In contrast, with offshoring efficiencies an acquisition by Whirlpool increases prices by 2.6 percent and decreases consumer welfare by 4.3 percent.

[45] All products that are marketed under a brand owned by Maytag before the merger (i.e. Admiral, Amana, Magic Chef and Maytag) are denoted as "Maytag".

[46] This is in line with the descriptive evidence showing that there is also product entry by LG and Samsung in appliance categories unaffected by the Whirlpool acquisition.

[47] Since Haier has nearly no presence in the U.S. clothes washer market pre-merger, without efficiencies, an acquisition by Haier has no economically significant effect on the product market. Simulation results for Haier acquisitions are therefore relegated to Appendix Table A.9.

Table 1.9: Product market effects of Maytag acquisitions by Haier and Whirlpool

	No portfolio adjustments			**Merger-independent adjustments**			**Endogenous adjustments**		
Efficiencies:	**No efficiencies**	**Offshoring**		**No efficiencies**	**Offshoring**		**No efficiencies**	**Offshoring**	
Acquirer:	Whirlpool	Haier	Whirlpool	Whirlpool	Haier	Whirlpool	Whirlpool	Haier	Whirlpool
Average price	2.7% [1.7%, 3.7%]	-1.0% [-2.0%, 0.1%]	2.6% [1.5%, 3.6%]	2.9% [2.0%, 3.8%]	-0.9% [-1.0%, -0.8%]	2.7% [1.8%, 3.6%]	[1.8%, 5.1%]	[-1.9%, 0.2%]	[1.7%, 4.9%]
Consumer welfare	$-131M [$-206M, $-55M]	$83M [$12M, $155M]	$-116M [$-191M, $-42M]	$-156M [$-237M, $-75M]	$61M [$59M, $63M]	$-140M [$-221M, $-59M]	[$-300M, $-197M]	[$58M, $152M]	[$-302M, $-197M]
	-4.9% [-8.4%, -1.3%]	3.1% [0.3%, 5.9%]	-4.3% [-7.8%, -0.9%]	-4.9% [-8.0%, -1.8%]	1.9% [1.9%, 2.0%]	-4.4% [-7.5%, -1.3%]	[-10.0%, -6.6%]	[1.9%, 5.2%]	[-10.1%, -6.6%]
Industry profits	$66M [$27M, $105M]	$14M [$-38M, $65M]	$76M [$31M, $120M]	$81M [$41M, $121M]	$26M [$24M, $27M]	$85M [$46M, $125M]	[$76M, $133M]	[$18M, $74M]	[$78M, $132M]
							[5.6%, 10.3%]	[1.4%, 5.8%]	[5.8%, 10.2%]
Maytag + Whirlpool profits	$15M [$-23M, $54M]	$14M [$-72M, $99M]	$33M [$-11M, $77M]	$18M [$-18M, $54M]	$68M [$67M, $68M]	$30M [$-6M, $66M]	[$19M, $47M]	[$91M, $119M]	[$20M, $48M]
							[2.3%, 6.3%]	[-92.5%, 189.4%]	[2.4%, 6.2%]

Notes: The first three columns show the effects of a Whirlpool and Haier acquisition of Maytag without product portfolio adjustments. The next three columns show the same comparison for merger-independent portfolio adjustments and the final three columns for endogenous portfolio adjustments. Percentage changes in total profits can only be computed with endogenous adjustments, as this is the only scenario for which I compute fixed costs. Each adjustment scenario includes results without marginal cost efficiencies and with offshoring efficiencies. 95% confidence intervals for effects without product portfolio adjustments and with merger-independent portfolio adjustments are computed using 100 residual bootstrap draws. Confidence sets for the expected effects with endogenous portfolio adjustments are based on 50 fixed cost draws for each potential product from a uniform distribution, where the domain are the confidence sets of brand-level fixed costs, and 500 demand and supply residual draws. Maytag includes all products marketed under the brands owned by Maytag pre-acquisition (i.e. Admiral, Amana, Magic Chef and Maytag). Since Haier has nearly no presence in the U.S. clothes washer market pre-merger, without efficiencies, an acquisition by Haier has no economically significant effect on the product market. Simulation results for Haier acquisitions without efficiencies are therefore omitted for legibility. They can be found in Appendix Table A.9.

Next, I contrast this to the simulated merger effects if future product portfolio adjustments are fully independent of any acquisition.[48] Without marginal cost efficiencies, the predicted price increases of an acquisition by Whirlpool are modestly higher with merger-independent adjustments than without adjustments. Although this leads to a higher loss in consumer welfare in absolute terms, in both cases the Whirlpool acquisition leads to similar relative losses in consumer welfare. This is because the expansion of the overall product portfolio between 2005 and 2007 leads to an increase in the level of consumer welfare under either acquisition scenario. The comparison of both portfolio adjustment cases with offshoring efficiencies is similar.

With fully endogenous portfolio adjustments and no marginal cost efficiencies, the 95 percent confidence sets of the price effects of an acquisition by Whirlpool include modestly higher price increases than the 95 percent confidence intervals with fixed product portfolios. However, the decrease in consumer welfare is considerably higher than without endogenous portfolio adjustments. This is because I predict only a modest increase in the product portfolio by rivals and a larger decrease in the portfolio by Maytag and Whirlpool as a consequence of a Whirlpool acquisition.

If product portfolio adjustments are fully endogenous, offshoring efficiencies make both hypothetical acquisitions better for consumers. Since offshoring efficiencies are larger for Haier than for Whirlpool, this increases the discrepancy in the consumer welfare effects of the two acquisitions.

1.7.5 Trading off workers and consumers

Table 1.10 summarizes how the two potential acquisitions affect consumers and U.S. employment in different scenarios. Each figure is the difference between the effects of Maytag acquisitions by Whirlpool and Haier. As seen before, a Whirlpool acquisition is always worse for consumers than an acquisition by Haier, since it leads to a higher increase in market power, fewer offshoring efficiencies, only modestly more entry by rivals, and significantly fewer products by the merging parties.

Across all scenarios, an acquisition of Maytag by Whirlpool maintains significantly more jobs

[48] This could occur if there is an expansion of the set of potential products due to technological progress but that is unrelated to changes in market structure.

Table 1.10: Simulated effects of Maytag acquisitions by Whirlpool vs. Haier

	No portfolio adjustments		**Merger-independent adjustments**		**Endogenous adjustments**	
Efficiencies:	No efficiencies	Offshoring	No efficiencies	Offshoring	No efficiencies	Offshoring
Consumer welfare	$-130M [$-205M, $-55M]	$-200M [$-309M, $-90M]	$-155M [$-236M, $-74M]	$-201M [$-276M, $-127M]	[$-302M, $-197M]	[$-386M, $-275M]
	−4.8% [−8.3%, −1.3%]	−7.2% [−11.9%, −2.5%]	−4.9% [-8.0%, -1.8%]	−6.2% [-9.7%, -2.7%]	[−10.1%, −6.6%]	[-12.5%, -9.0%]
Domestic jobs maintained	1233 [677, 1790]	1202 [627, 1778]	1120 [644, 1597]	1201 [645, 1757]	[1021, 1507]	[1244, 1588]
	24.8% [2.6%, 47.1%]	24.2% [1.0%, 47.5%]	22.2% [5.9%, 38.4%]	24.2% [1.9%, 46.5%]	[24.0%, 37.0%]	[30.4%, 40.3%]
Offsetting job value	$106k [$19k, $192k]	$166k [$59k, $274k]	$138k [$63k, $214k]	$168k [$81k, $255k]	[$135k, $316k]	[$195k, $316k]

Notes: The first two columns show the effects of a Whirlpool vs. Haier acquisition of Maytag without product portfolio adjustments. The next two columns show the same comparison for merger-independent portfolio adjustments and the final two columns for endogenous portfolio adjustments. Each adjustment scenario includes results without marginal cost efficiencies and with offshoring efficiencies. 95% confidence intervals for effects without product portfolio adjustments and with merger-independent portfolio adjustments are computed using 100 residual bootstrap draws. Confidence sets for the expected effects with endogenous portfolio adjustments are based on 50 fixed cost draws for each potential product from a uniform distribution, where the domain are the confidence sets of brand-level fixed costs, and 500 demand and supply residual draws. Maytag includes all products marketed under the brands owned by Maytag pre-acquisition (i.e. Admiral, Amana, Magic Chef and Maytag).

in the U.S. than an acquisition by Haier. In Appendix A.5.3, I decompose this difference in U.S. employment into a relocation effect (due to different plant relocations between acquirers) and a reallocation effect (due to different market share reallocations after each acquisition). Whereas the former leads to more U.S. jobs after a Whirlpool acquisition, the latter pushes in the opposite direction, as foreign competitors gain more market shares after an acquisition by Whirlpool. I find that in this case, most of the change in employment can be explained by the relocation effect.

I next combine the product market and employment effects and estimate the average job value necessary such that the domestic jobs saved by a Whirlpool acquisition offset the losses in consumer welfare. Overall, this offsetting value is lowest without portfolio adjustments and marginal cost efficiencies and increases for merger-independent and then endogenous adjustments. It increases when accounting for marginal cost efficiencies, since this increases the consumer welfare losses of a Whirlpool compared to a Haier acquisition and makes Maytag products produced by Haier in China relatively more attractive. With endogenous portfolio adjustments and without marginal cost efficiencies, the average offsetting value of each additional job is between $135,000 and $316,000 per year.

To gauge whether an acquisition of Maytag by Whirlpool is better for U.S. welfare than an alternative acquisition by Haier requires estimating the consumer welfare and employment effects of the acquisitions for other product markets. Other product markets that Maytag was

Table 1.11: Comparing offsetting job value to other estimates in the literature

	Offsetting job value	Δ Total wage bill	Notes
Maytag acquisition by Whirlpool vs. Haier			
Clothes washers without efficiencies	[$135k, $316k]		Necessary value of a U.S. clothes washer job to offset losses in consumer welfare
Clothes washers with offshoring efficiencies	[$195k, $316k]		Necessary value of a U.S. clothes washer job to offset losses in consumer welfare
Other household appliances	[$50k, $79k]		Based on back-of-the-envelope calculation described in Appendix Section A.5.4
Other estimates			
Hufbauer and Lowry (2012)	$900k		Estimate that 2011 U.S. safeguard tariffs on tire imports from China saved 1,200 jobs and cost consumers $1.1 bn
Flaaen, Hortacsu and Tintelnot (2020)	$817k		Estimate that 2018 U.S. global safeguard tariffs on clothes washers created 1,800 jobs and cost consumers $1.6 bn
Jaravel and Sager (2020)	[$288k, $478k]		Estimate price and employment effects of U.S. trade liberalization with China
Setzler and Tintelnot (2021)		$113k	Estimate increase in the total wage bill (from direct and indirect effects) in the local labor market per additional foreign MNE job

active in are clothes dryers, dishwashers, ranges, and refrigerators. Without additional product market data, I cannot get a precise estimate of the consumer welfare and employment effects of the merger for these markets. Instead, I make a rough back-of-the-envelope approximation of the order of magnitude of these effects based on the market size and the change in the HHI in these markets compared to clothes washers, as well as reported overall Maytag U.S. employment pre-merger. The details of the estimation can be found in Appendix Section A.5.4. I find that without marginal cost efficiencies, for appliance categories that are not clothes washers, the average value of a job necessary to offset consumer welfare losses of a Whirlpool acquisition is between $50,000 and $79,000 per year. Unsurprisingly, this is significantly lower than for clothes washers, since in all other appliance categories (with the exception of clothes dryers) the increase in the HHI is much lower.

Table 1.11 summarizes the job value for which employment effects offset consumer harm and compares these to other estimates from the literature. To offset the consumer harm from trade restrictions, Hufbauer and Lowry (2012) and Flaaen et al. (2020) estimate that annual job values of between $800,000 and $900,000 are necessary. In comparison, the necessary job values to make a Whirlpool acquisition more attractive than an acquisition by Haier are modest. Results by Jaravel and Sager (2020) for trade liberalization show that trade with China increased U.S. consumer surplus by about $400,000 per displaced job. These are closer in magnitude to the results in this paper.

To determine which acquisition leads to higher U.S. welfare requires determining the value of a job. This is difficult, since even if I knew the wage for each job, it would be insufficient to infer its direct (for the worker that holds the manufacturing job) and indirect (for other workers in the economy) effects. Since I do not have the necessary variation and data to estimate these effects, I compare the offsetting job values to estimates by Setzler and Tintelnot (2021) of the direct and indirect local labor market effects of a job created by a foreign multinational firm.

They find that an additional foreign multinational job increases the total wage bill in a local labor market by $\$113,000$ per year. This includes wages for workers coming from non-employment, as well as the direct effect of the foreign multinational wage premium on employees previously employed at domestic firms.[49] It also includes wage gains for employed workers at domestic firms, as well as wages at newly created domestic jobs.[50] I apply these estimates to all appliance manufacturers, irrespective of their nationality, since Setzler and Tintelnot (2021) find similar wage premia for foreign and domestic multinationals and all of the appliance manufacturers fall in either of the two categories. I do not need to distinguish between local and national employment effects, since I consider each manufacturing job not created domestically to be created abroad.

There are many other positive effects related to an increase in the availability of jobs, that go beyond an increase in wages. Bearing this in mind, I consider the increase in the total wage bill by $\$113,000$ per year as a lower-bound estimate of the value of a U.S. appliance manufacturing job to the U.S. economy. This is at the lower end of the necessary job values to offset losses in consumer welfare for clothes washers with endogenous portfolio adjustments. Since this offsetting value is lower for other appliance categories, I cannot exclude that the sum of consumer welfare and domestic worker income is higher after an acquisition of Maytag by Whirlpool than after an acquisition by Haier. Overall, the aggregate gains to domestic workers from additional jobs are of similar magnitude as the losses in consumer welfare.

[49]They estimate that, on average, 87% of foreign multinational employees are previously employed by a domestic firm and that the multinational wage premium is 7%. They estimate that the average earning of a full-time worker is $\$62,600$ at a domestic firm and $\$75,200$ at a foreign multinational.

[50]They estimate a wage increase of 0.15% for workers employed at domestic firms for each one percentage point increase in the share of workers employed at foreign multinationals. Finally, they estimate a total local job multiplier of 1.50, which means 0.50 indirect jobs for every direct job created.

1.7.6 Unequal distribution of welfare effects

So far, the analysis focused on how consumers and workers overall are affected by the two alternative acquisitions. However, not all households need to be affected similarly by the acquisitions on the product market, as well as the employment side.

Figure 1.12 shows the simulated percentage change in consumer welfare between Maytag acquisitions by Whirlpool and Haier, depending on the annual income of a household. Without marginal cost efficiencies, the relative decrease in consumer welfare is highest for households with the lowest household income and the loss in consumer welfare modestly decreases the higher the income.[51] This relationship is steeper with fully endogenous portfolio adjustments than with exogenous adjustments or no adjustments at all.

With offshoring efficiencies, the relationship between household income and the consumer welfare losses after a Whirlpool acquisition of Maytag compared to a Haier acquisition is similar with exogenous or no portfolio adjustments. With endogenous portfolio adjustments, the loss in consumer welfare first increases in household income and then decreases again. This is because with offshoring efficiencies Haier introduces more new products of the former Maytag brand families post-merger. These products are mainly purchased by medium-income households. Foregoing these new Maytag products after an acquisition by Whirlpool thus exacerbates the losses in consumer welfare of a Whirlpool acquisition for households who otherwise would have considered buying these products.

The employment effects of the alternative potential acquisitions are also geographically highly unequally distributed. For the U.S. economy as a whole, 1,000 additional clothes washer manufacturing jobs do not have any significant effect on employment or wages. As I showed in Section 1.3, however, this is different for local labor markets. There, the closure of a manufacturing plants can decrease wages and employment at the county-level even two years after the plant closure. As illustrated in Figure 1.1, clothes washer manufacturing plants are concentrated in a few counties in Illinois, Iowa, Michigan, Ohio, and South Carolina. Thus, whereas most local labor markets are unaffected by the potential acquisitions, some are strongly affected. Even considering average employment effects at the level of the local labor market masks important heterogeneities in employment effects. Whereas some workers may only mildly be affected by

[51]The median annual household income for the simulated households is $55,000.

Figure 1.12: Consumer welfare change by household income: Whirlpool vs. Haier

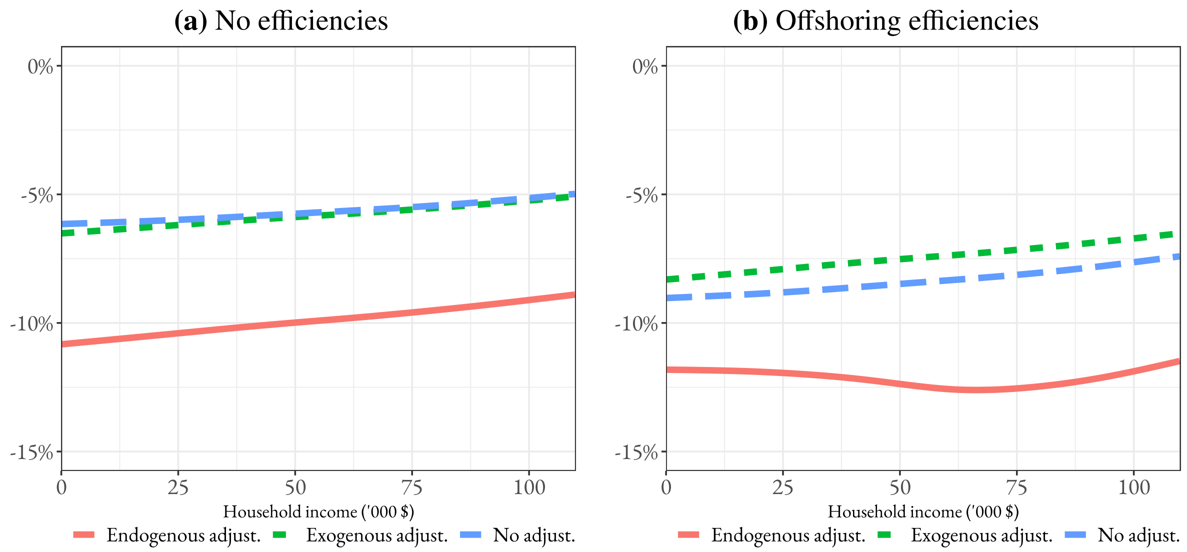

Notes: Both graphs show the simulated percentage change in consumer welfare between a Maytag acquisition by Whirlpool and Haier according to household income. The consumer welfare changes are shown for three adjustment scenarios: No product portfolio adjustments, exogenous adjustments, and fully endogenous adjustments. Simulations are based on 1,000 household draws per market. For expositional simplicity, the graphs only show households with an annual income of less than 110,000$, covering 80% of drawn households. In the left panel, no marginal cost efficiencies are credited. In the right panel, offshoring efficiencies are credited to Maytag products. Changes in absolute terms are reported in the Appendix, in Figure A.10.

a plant closure, others lose their job and their livelihood.

As I showed above, there is some, albeit modest, heterogeneity in the difference in consumer welfare effects between the two potential acquisitions. These are also unlikely to differ greatly geographically. In contrast, there are large heterogeneities in the geographic distribution of employment effects. This has implications for optimal policy. If household preferences are such that households have diminishing marginal utility of income and employment effects are not concentrated among the very wealthy, then an acquisition by Whirlpool may be domestic welfare improving, even if the increase in the total domestic wage bill as compared to a Haier acquisition is lower than the relative consumer welfare loss. Furthermore, other non-wage considerations related to job loss, such as mental or physical health, can improve the domestic welfare effects of a Whirlpool as compared to a Haier acquisition.

Finally, political considerations cannot be neglected completely. Whereas a loss in consumer welfare in the clothes washer market by $20 is unlikely to affect how voters cast their ballot, direct and indirect employment effects can. Thus, facilitating an acquisition of Maytag by Whirlpool as opposed to Haier can be politically more attractive.

1.8 Simplifying Estimation with Proprietary Data

There are two aspects of the analysis discussed so far which make it less attractive to implement in ex ante merger control. First, estimating fixed cost bounds is costly in terms of programming and computational time. Only identifying bounds on the fixed cost of product entry also reduces the precision of the estimation. Second, using observed post-merger entry to estimate the set of potential products is infeasible in ex ante merger analysis. Luckily, better data usually more easily available to competition authorities than to researchers can solve these problems.

To avoid the estimation of fixed costs, competition authorities can use market surveys and internal documents from market participants to form estimates of the costs of adding new products. Such a calibration exercise can simplify the overall analysis and potentially also lead to tighter fixed cost bounds.

A similar approach can be taken to estimate the set of potential products and their characteristics. Instead of having to observe what products are added by market participants after a merger, competition authorities can rely on market surveys and internal documents to identify the potential product pipeline.

This does not mean, however, that it is possible to completely forego the methods described in this paper to predict post-merger product portfolios. Estimating which products enter conditional on demand, a set of potential products, and associated marginal costs and fixed costs of adding a product remains necessary. Although using an algorithm to predict which products enter after a hypothetical merger cannot be substituted with market surveys or internal documents, these, as well as industry-specific knowledge, can help refine the heuristic entry algorithm. This can improve the reliability of portfolio choice predictions in practice.

1.9 Conclusion

This paper proposes a model to analyze the consumer welfare and employment effects of different potential product market mergers in the presence of foreign entry and describes how to estimate its structural parameters. To account for how mergers change the incentives to introduce new products, I allow firms to endogenously adjust their product portfolios. To account for employment effects, I model how the equilibrium in the product market affects the number of (domestic) employees required to manufacture the product. Structurally estimating the parameters of the model is not only possible for ex post merger evaluations, but also for prospective merger analysis. To facilitate its use in merger control, I describe how data that could be requested by competition authorities can be used to reduce the estimation burden and increase the precision of the estimates.

I apply the model to analyze the acquisition of Maytag by Whirlpool in the U.S. household appliance industry. Although the Department of Justice cleared the merger on the grounds that foreign product entry would sufficiently constrain the merging parties, the merger decreased consumer welfare by more than 5 percent. This is true when comparing it to no merger, as well as an alternative acquisition of Maytag by Haier. With endogenous product portfolio adjustments, rivals modestly increase the number of products they offer after a Whirlpool acquisition and the merging parties moderately decrease the number of products. Overall, endogenous portfolio adjustments increase the consumer harm of a Whirlpool acquisition. Rival product entry is therefore an insufficient constraint on the merging parties even in this landmark case where an entry defense was at the heart of the clearance decision.

By estimating the domestic employment effects of the two potential acquisitions, I investigate whether these could have made an acquisition by Whirlpool more desirable in terms of U.S. domestic welfare. I find that a Whirlpool acquisition leads to the preservation of more U.S. manufacturing jobs. I calculate the average value of a job necessary for domestic employment effects to offset the losses in consumer welfare. A comparison to estimates by Setzler and Tintelnot (2021) on the direct and indirect local labor market effects of new jobs by multinational firms leads to the conclusion that the aggregate gains to domestic workers from additional jobs is of similar magnitude as the consumer welfare losses. Overall, I cannot exclude that an acquisition of Maytag by Whirlpool leads to higher domestic welfare than an alternative acquisition by Haier.

This has important implications for policy. Since the employment effects of a product market merger can be of first order importance, these should not be ignored in merger analysis. Blocking acquisitions that could lead to the offshoring of jobs or allowing anti-competitive mergers that could lead to the preservation of jobs compared to an alternative acquisition may still not be optimal.

Instead, the framework laid out in this paper could be used to identify mergers in which employment effects are of first order importance. Whilst the merger decision could still be taken based on the consumer welfare standard, this would identify cases where there may be a need for complementary labor market policies.

Chapter 2

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

2.1 Introduction

Mandatory price disclosure (MPD) policies are increasingly used to make markets more competitive.[1] Studies estimating the local effect of mandatory price disclosure on prices find mixed results on their effect.[2] So far, there is limited evidence about why mandatory price disclosure sometimes lowers prices and sometimes does not.

In this paper, we ask what determines the price effect of mandatory price disclosure. 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, station-level price changes for Germany and France before and after the introduction of MPD in Germany. Second, MPD was introduced simultaneously for diesel and gasoline. On average, consumers buying gasoline are less informed about prices

This chapter is based on joint work with Alina Sagimuldina and Christoph Winter.

[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 (2019), who find that mandatory price disclosure decreased prices at Israeli supermarkets.

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.

Several findings emerge from our analysis: Theoretically we show that the more uninformed consumers there are prior to the introduction of MPD, the smaller 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 2.4 percent and diesel prices by around 1.3 percent. The difference in treatment effects is particularly strong in the five months after the introduction of MPD. Thereafter, the treatment effect stabilizes at around 1.5 percent for gasoline and 1.0 percent for diesel. We consistently confirm the empirical results using alternative information shocks and identification strategies. Overall this suggests that MPD is most effective in markets where few consumers are well-informed before its introduction.

The theoretical analysis is set in the context of the canonical model of 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. 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.[3] 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 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.

[3] 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.

In Germany, cars with diesel engines are driven by consumers that drive on average twice as many kilometers per year as gasoline buyers.[4] 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. 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 paper 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 disclosure 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 and 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 monotonous 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

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

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 (2019) 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. 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.[5] 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 paper relates to an extensive empirical literature that analyzes pricing decisions for retail fuel. 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-Dengler, Schutz, Weiss, and Yontcheva, 2020, or Byrne and de Roos, Forthcoming). In contrast, J.-F. 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. Eckert (2013) provides an overview of the earlier literature on pricing in fuel markets.

The remainder of this paper 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.

[5]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.

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 v. 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.[6] 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

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

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.

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*.[7] 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

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

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

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

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

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

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

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* $\frac{\partial F(p)}{\partial \Delta\phi \partial \phi_0} < 0$.

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 an ex ante higher 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.

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

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

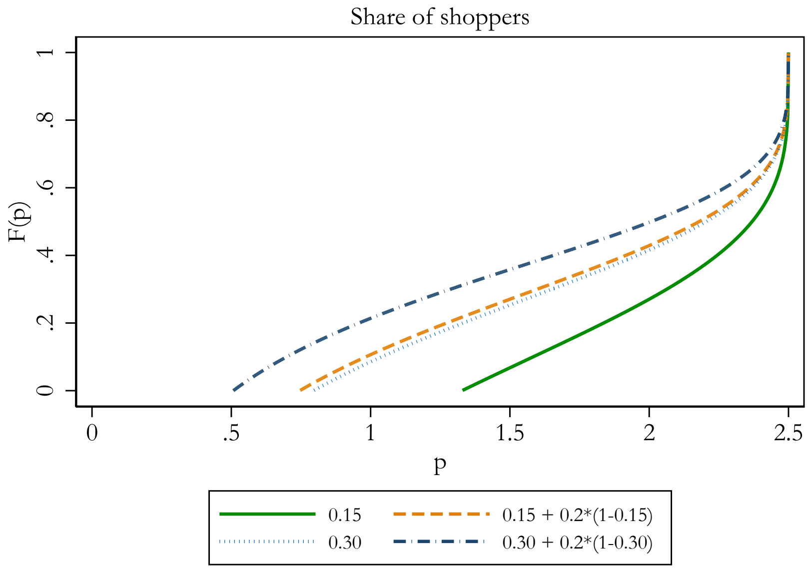

Note: The Figure shows simulation results of how the distribution from which sellers draw prices in the symmetric Nash equilibrium changes if the information shock Δ_ϕ hits the market. Parameter values: $\upsilon = 2.5$, $N = 5$, $\phi_{01} = 0.15$, $\phi_{02} = 0.30$ and $\Delta_\phi = 0.20$. The solid 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.

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, as well as on diesel, which were correspondingly used in 56 and 29 percent of passenger vehicles with combustion engines in Germany in 2013.[8]

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 on fuel costs later on. Diesel motorists are thus likely to have 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. Let us now see why 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,

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

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 may also receive a lump-sum or a per mile fuel allowance or are self-employed, which creates additional incentives to save on fuel costs. Thus, that many diesel vehicles are used for commercial purpose does not invalidate our observation that diesel motorists are on average more price sensitive than gasoline drivers.[9]

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, there were significant search costs for consumers. To find the prices of all potential sellers, a motorist would need to drive to all stations.[10]

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

[9] 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.

[10] 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.

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.[11] 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 firm name, 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.[12] 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.

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

[11] 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.

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

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.[13] 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.

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 $500m^3$ of fuel per year have to report all price changes to a government agency similar to the MTU in Germany. Regular

[13]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	13,783	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	1.59	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	0.10	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.

checks are carried out and fines imposed on fuel stations that do not comply with this rule. To the best of our knowledge, France is the only other European country for which station-level fuel prices are available during this period.[14] The French government makes all price information since 2007 publicly available on a government website.[15] 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.[16]

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

[14] Austria introduced mandatory price disclosure in 2011, however only publishes the five lowest prices in a local market. In addition, daily average prices at the state level are available for Austria.

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

[16] 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.

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 more rigorous 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 description of the industry 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.[17] 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.

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

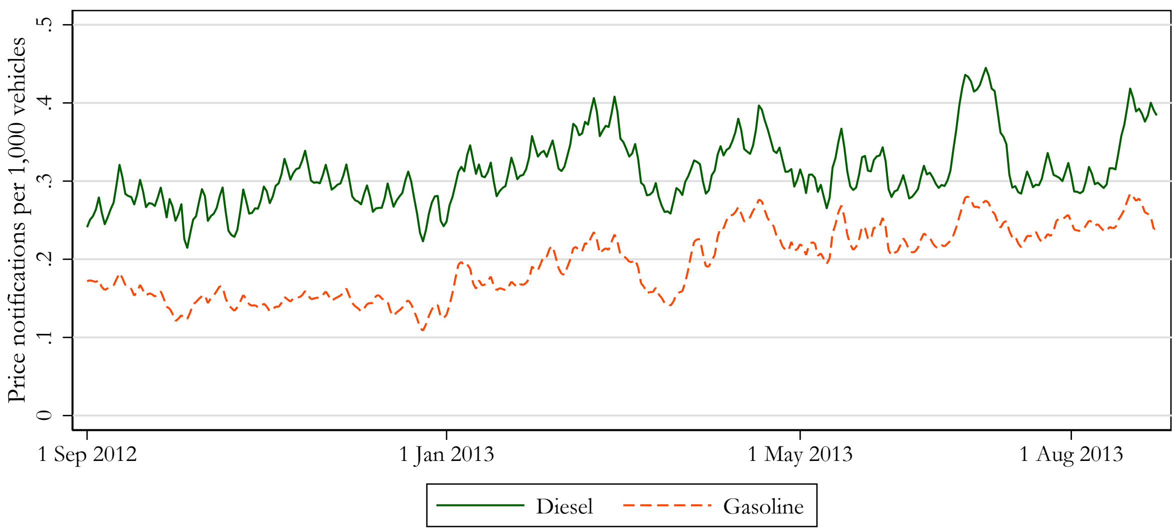

Notes: The Figure shows the daily number of self-reported price notifications by fuel type to a major German smartphone app per 1,000 diesel or gasoline vehicles in circulation. The data is available from September 2012 to August 2013. The solid line corresponds to the notification intensity for diesel.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

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

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.

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

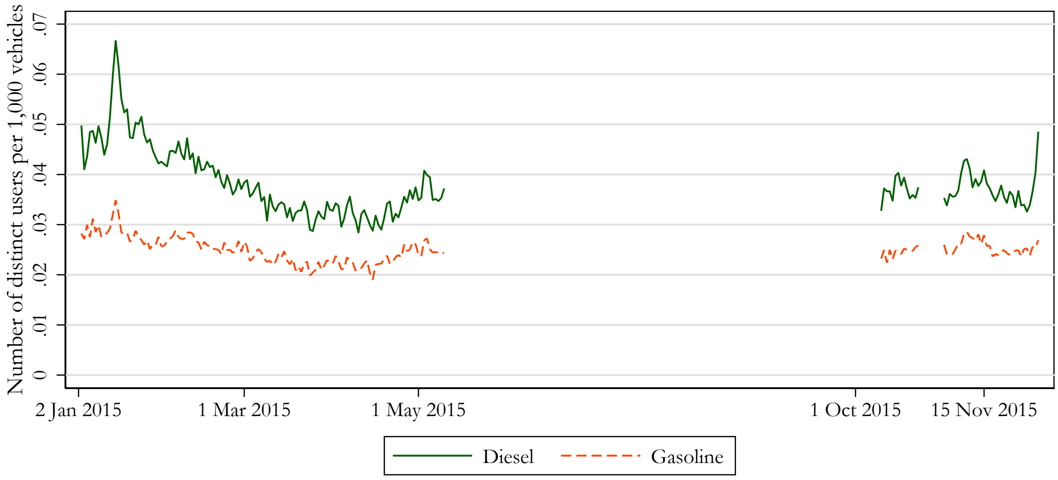

Notes: The Figure shows the daily number of distinct users who search for diesel or gasoline price in Germany in 2015, per 1,000 diesel or gasoline vehicles in circulation.

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.

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

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

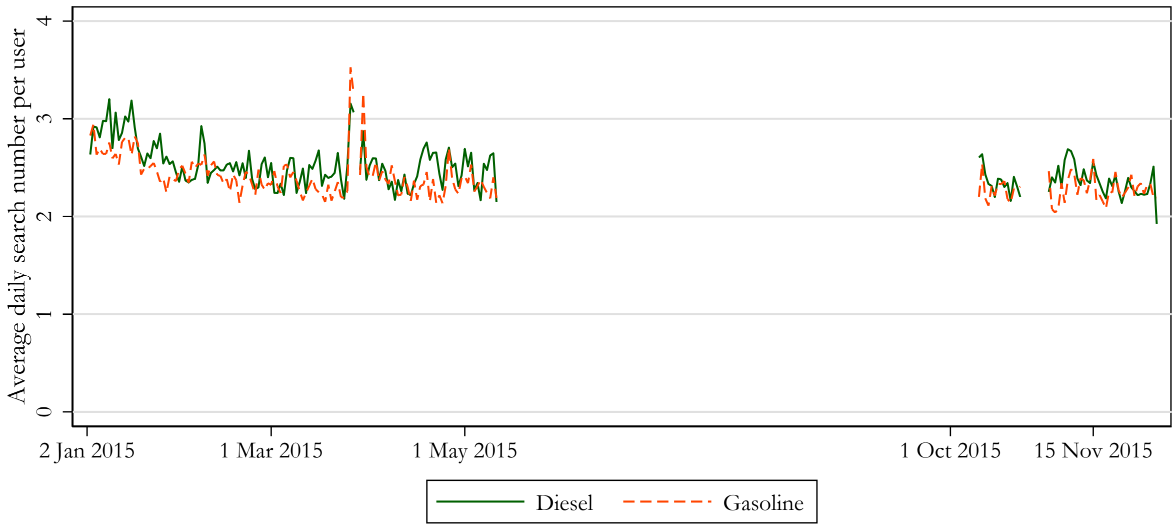

Notes: The Figure shows the daily number of price searches by fuel type at a major German smartphone app per 1,000 diesel or gasoline vehicles in circulation. The data is available for January to May and October to December 2015. The solid line corresponds to the search intensity for diesel. The dashed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

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

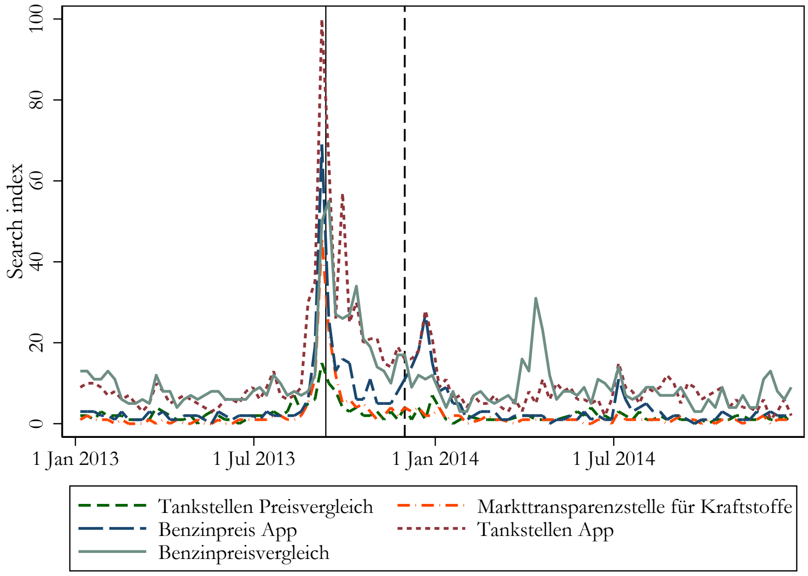

Notes: The figure shows the evolution of Google searches in Germany between 1 January 2013 and 31 December 2014 for MPD-related keywords. Searches are indexed such that 100 corresponds to the moment in time and keyword with the highest number of searches during the observation period. The search terms 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

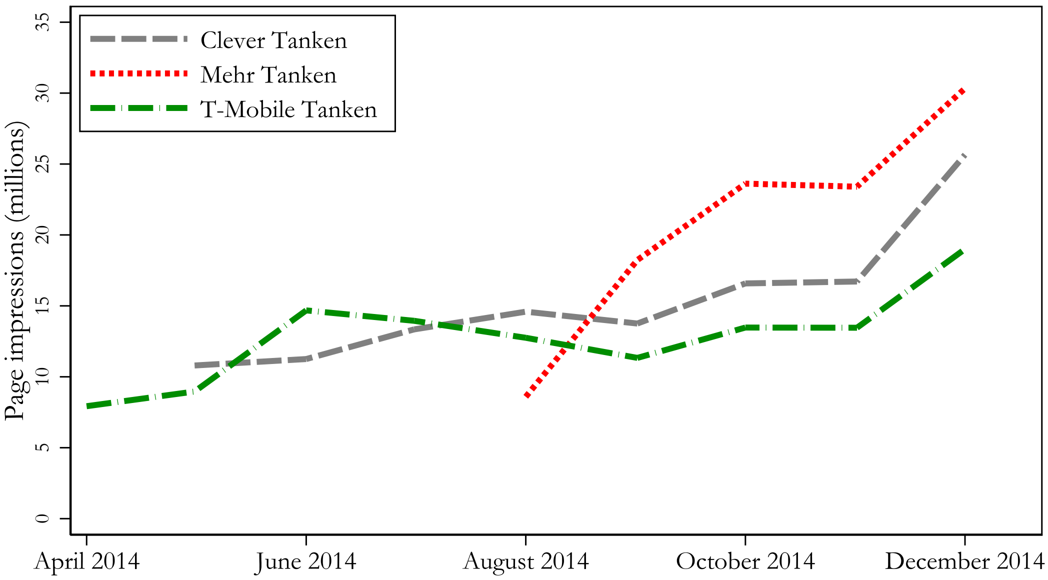

Notes: The Figure shows the evolution of monthly page impressions for three popular mobile price comparison applications. Each line begins when data for the particular app becomes available and ends at the end of our sample period, in December 2014.

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 evolve 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 price 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

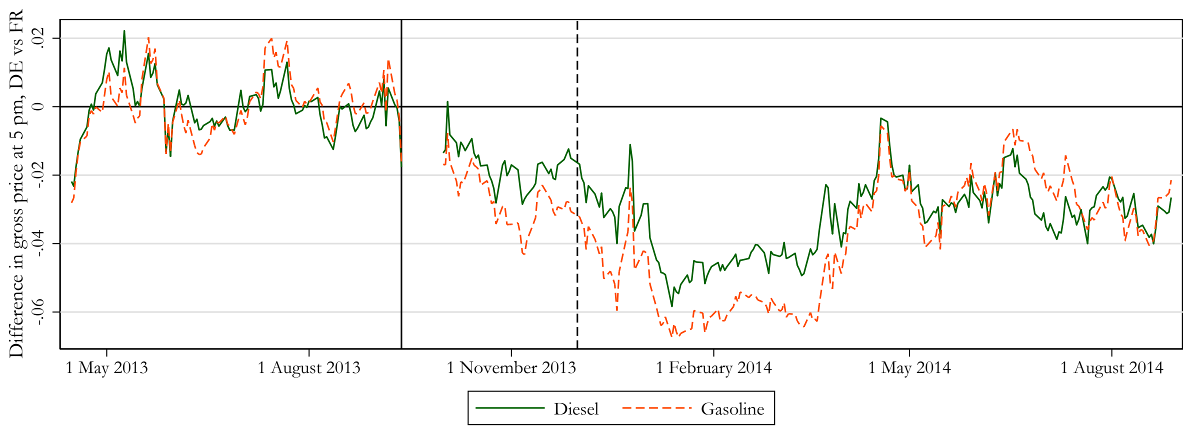

Notes: The solid line shows the evolution of the difference in daily diesel prices between Germany and France, demeaned by the corresponding average difference prior to MPD. The dashed line shows the evolution of the analogous difference in gasoline prices. The vertical solid line marks the beginning 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 may be 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.[18] Standard errors are computed via the jackknife method.[19]

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.[20]

Estimation of the treatment effect with SDID requires a balanced panel. We therefore estimate the effect on weekly average fuel prices and restrict our sample to fuel stations in Germany and France that have no missing weekly price observations.[21] This is the case for 52% 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.

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

[18]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.

[19]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.

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

[21]We employ weekly average fuel prices since a high share of stations in Germany have at least one missing daily fuel price observation during the time period used in the estimation of the treatment effect.

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. To the best of our knowledge, France is the only other country for which station-level fuel prices and retail margins are available for most stations for the full observation period.

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.[22]

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, we treat local monopolists in Germany as the control group and all other German stations as the treatment group. In the second, we use country-level weekly fuel prices for all countries

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

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.[23] 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:

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

where Y_{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.

[23] 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

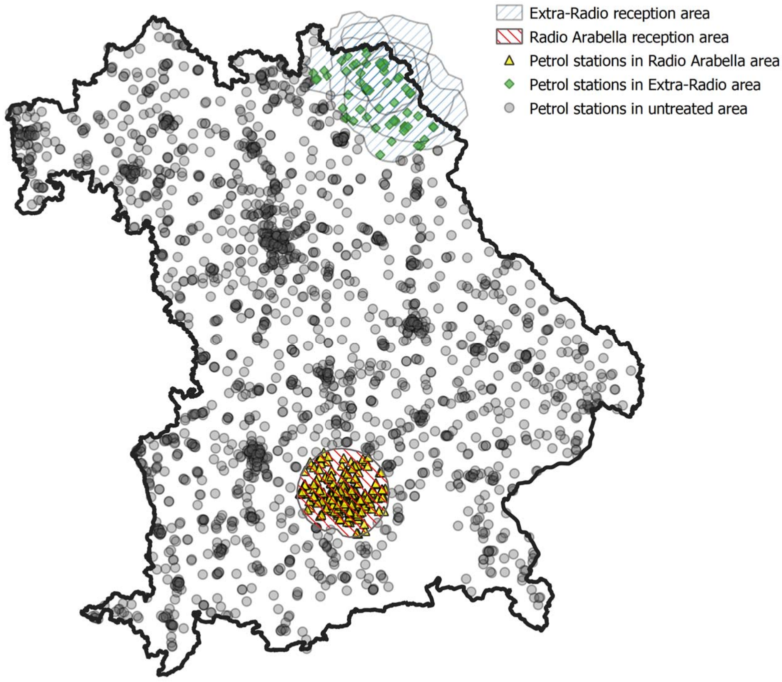

We can thus exclude that fuel stations in the control group are affected by reports of radio stations we have not surveyed. We restrict our analysis to the period October 2013 until September 2014, which is the twelve months after the beginning of the test phase of the MTU.

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.

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

	Gasoline	Diesel	Gasoline	Diesel
	(1)	(2)	(3)	(4)
MPD	-0.024***	-0.013***	-0.027***	-0.021***
	(0.002)	(0.003)	(0.002)	(0.001)
95% Confidence interval	[-0.028, -0.020]	[-0.018, -0.007]	[-0.030, -0.024]	[-0.023, -0.018]
Week FE	*Yes*	*Yes*	*Yes*	*Yes*
Station FE	*Yes*	*Yes*	*Yes*	*Yes*
Observations	666,106	783,951	52,969	58,408

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

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.[24]

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.

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

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

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.

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

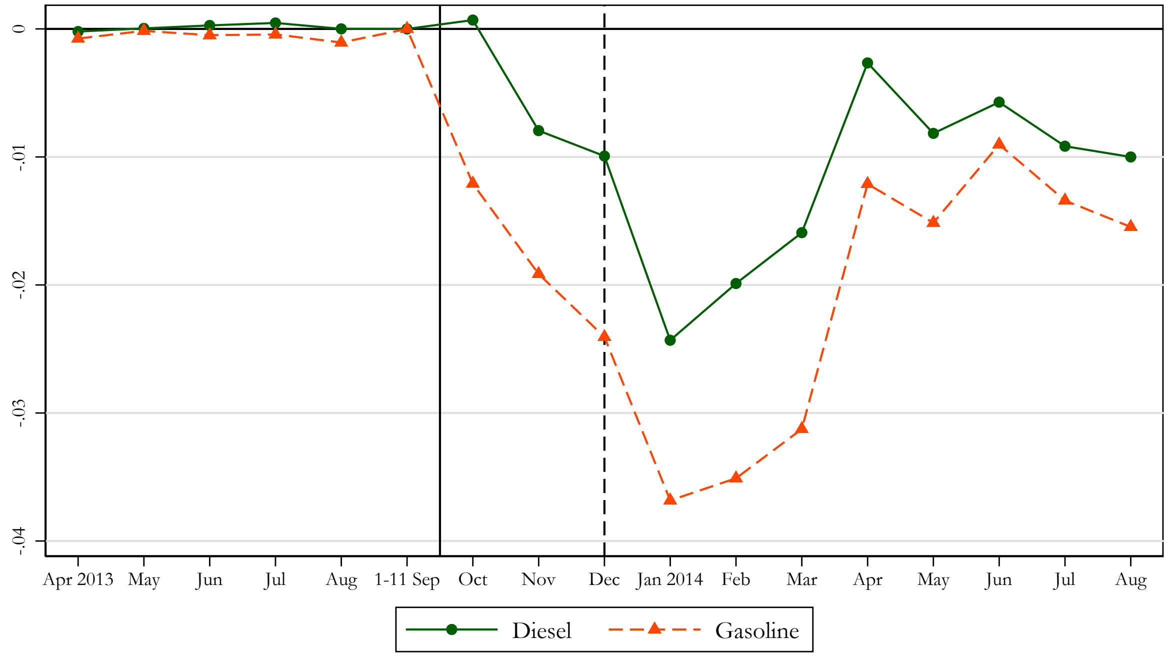

Notes: The Figure shows time-varying treatment effects of MPD on log weekly prices for gasoline and diesel between April 2013 and August 2014. The vertical solid line marks the beginning of the MTU test phase. The vertical dashed line marks the beginning of the MTU full-scale phase.

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***	0.001	-0.005***
	(0.003)	(0.0004)	(0.002)	(0.0004)
Date FE	*Yes*	*Yes*	*Yes*	*Yes*
Station FE	*Yes*	*Yes*	*Yes*	*Yes*
Observations	354,794	458,162	360,063	463,277
Adjusted R^2	0.426	0.425	0.306	0.305

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 paper, 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. In particular, 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 around 1.0 percent for diesel and 1.5 percent for gasoline. 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.

In this paper, we propose a new theoretical framework to analyze commodity tax pass-through

This chapter is based on joint work with Alina Sagimuldina 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 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, 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.

[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 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 paper 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, 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

[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.

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 (2019) 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.

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

[6]Some studies focus on particular industries, such as energy markets (see, e.g., Fabra and Reguant, 2014, Kopczuk et al., 2016, J. Li and J. H. 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 C. T. 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 (2019) 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.

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 J.-F. 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 the paper 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 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

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

[10] These include Chandra and Tappata (2011), Byrne and Roos (2017), Byrne and de Roos (Forthcoming) 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.

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 υ 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

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

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 υ.

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

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

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

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

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

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

[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

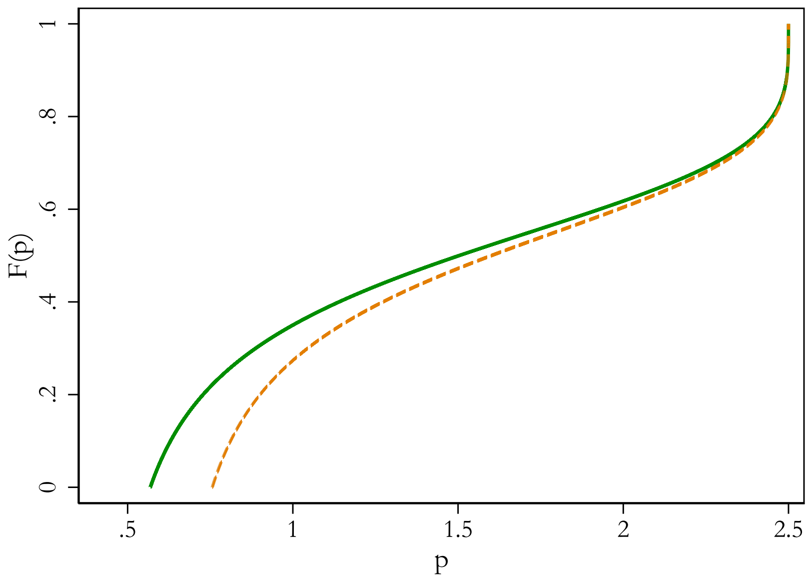

Note: The Figure shows simulation results of how the distribution from which sellers draw prices in the symmetric Nash equilibrium changes if the ad-valorem tax increases from τ to $\hat{\tau}$. The solid line corresponds to the distribution under τ. The dashed line corresponds to the distribution under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

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

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

price corresponds to the valuation of the good, v.

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 $\phi \to 1$, the pass-through rate $\rho_\tau \to 1$.*

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 $E[p]$

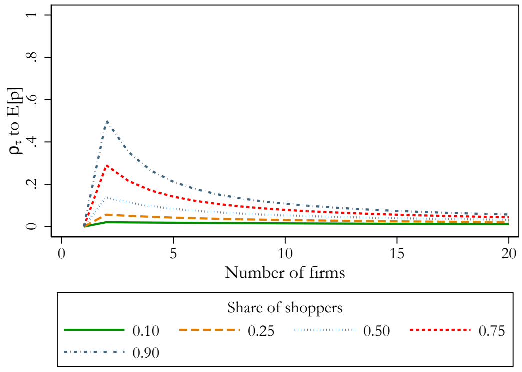

Parameter values: $\upsilon = 2.5$, $c = 0.4$, $\tau = 0.2$ and $\hat{\tau} = 0.22$.

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

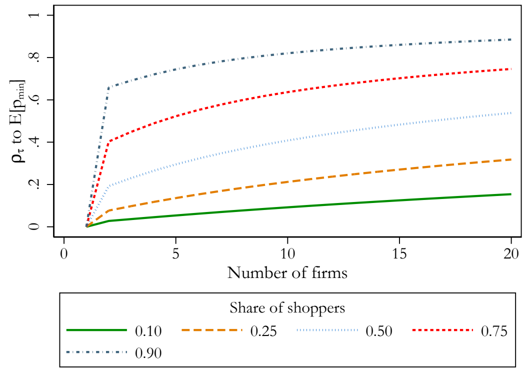

Parameter values: $\upsilon = 2.5$, $c = 0.4$, $\tau = 0.2$ and $\hat{\tau} = 0.22$.

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 decreases.

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

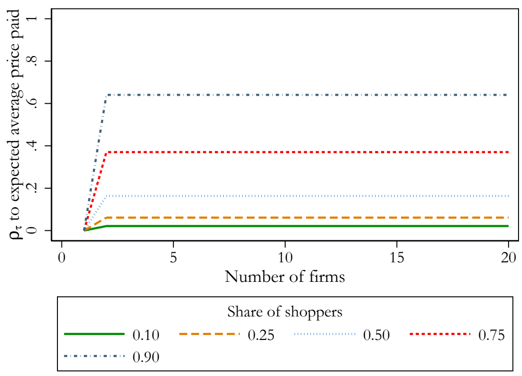

Parameter values: $\upsilon = 2.5$, $c = 0.4$, $\tau = 0.2$ and $\hat{\tau} = 0.22$.

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

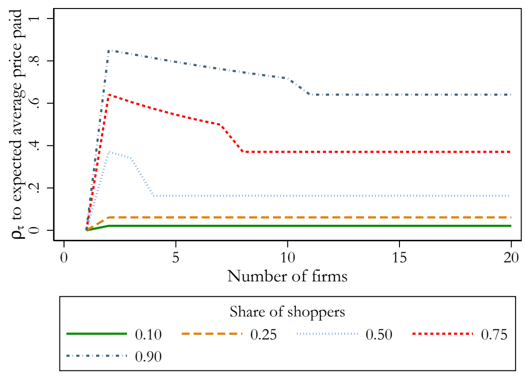

Parameter values: $\upsilon = 2.5$, $s = 0.75$, $c = 0.4$, $\tau = 0.2$ and $\hat{\tau} = 0.22$.

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 with combustion engines and gasoline accounts for the remaining 56 percent.[16] Substituting

[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.

[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.

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 passenger 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

[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)

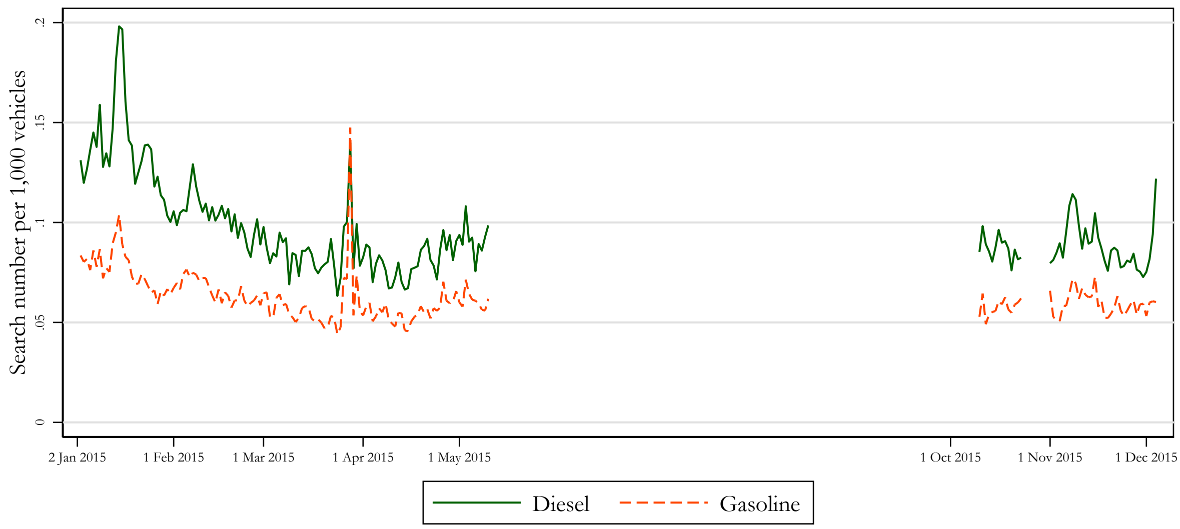

Notes: The Figure shows the daily number of price searches by fuel type on a major German smartphone app per 1,000 diesel or gasoline vehicles in circulation. The data is available for January to May and October to December 2015. The solid line corresponds to the search intensity for diesel. The dashed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

(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 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

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

[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.

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 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*.

[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/`.

[24] The full survey results can be found at `https://www.adac.de/news/umfrage-e10-tanken/`.

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

[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).

[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`.

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 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

[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.

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]

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

[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.

[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.

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.

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

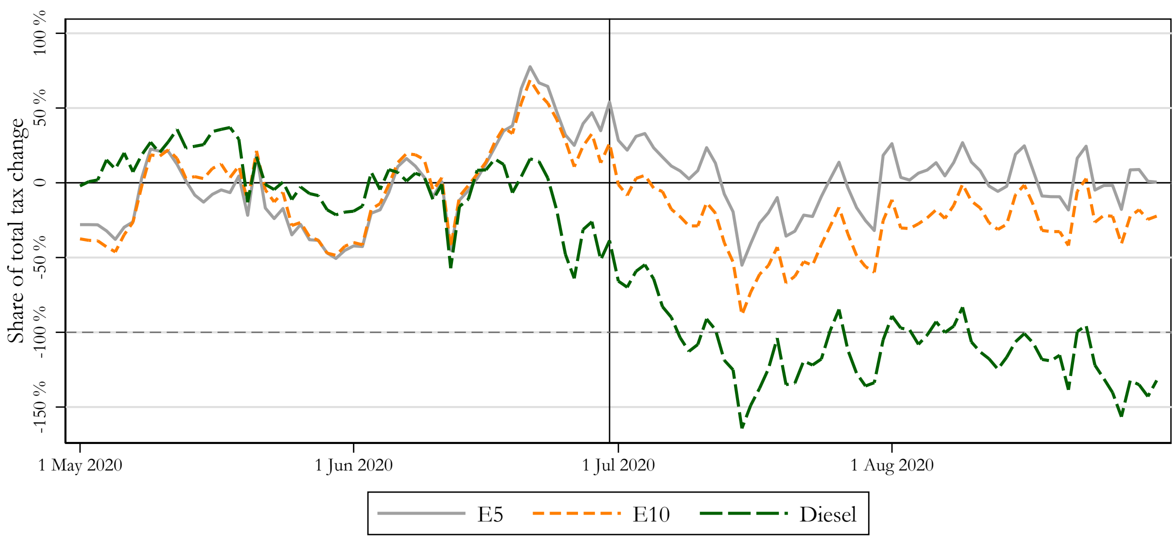

Notes: The solid line shows the nonparametric estimate of the daily average pass-through rate to prices for *E5*. The short-dashed and long-dashed lines show analogous estimates for *E10* and diesel, respectively. To estimate pass-through, we first subtract the average pre-period (1 May until 30 June 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

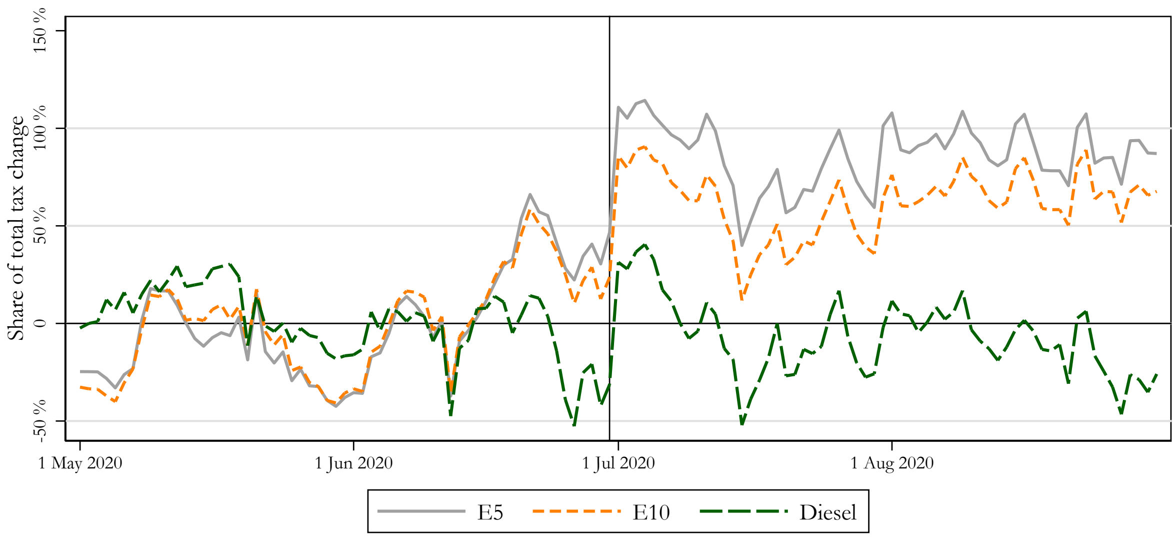

Notes: The solid line shows the nonparametric estimate of the daily average pass-through rate to retail margins for *E5*. The short-dashed and long-dashed lines show analogous estimates for *E10* and diesel, respectively. To estimate pass-through, we first subtract the average pre-period (1 May until 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.

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

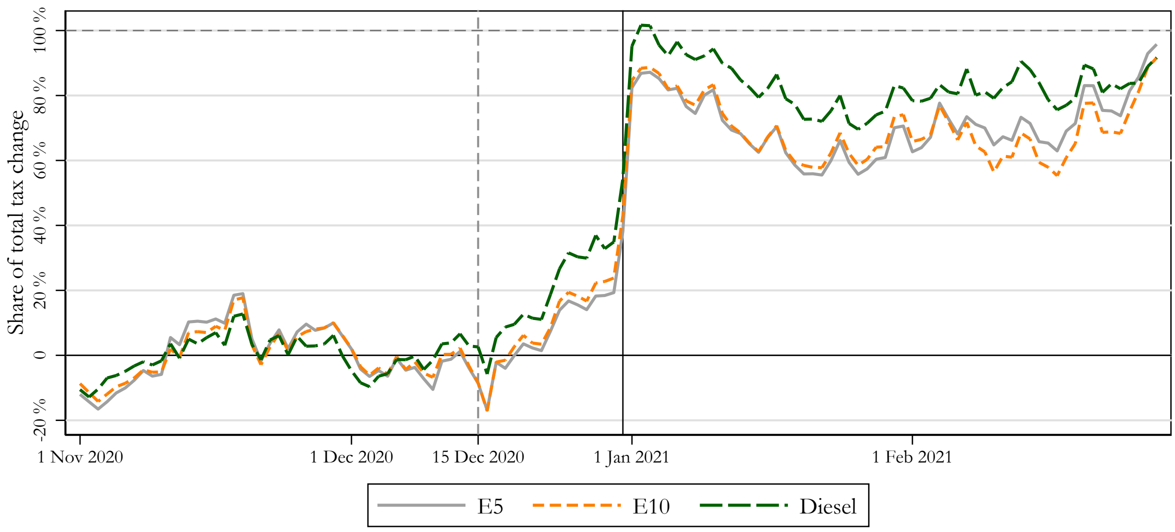

Notes: The solid line shows the nonparametric estimate of the daily average pass-through rate to prices for *E5*. The short-dashed and long-dashed lines show analogous estimates for *E10* and diesel, respectively. To estimate pass-through, we first subtract the average pre-period (1 November until 15 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.

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 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 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

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

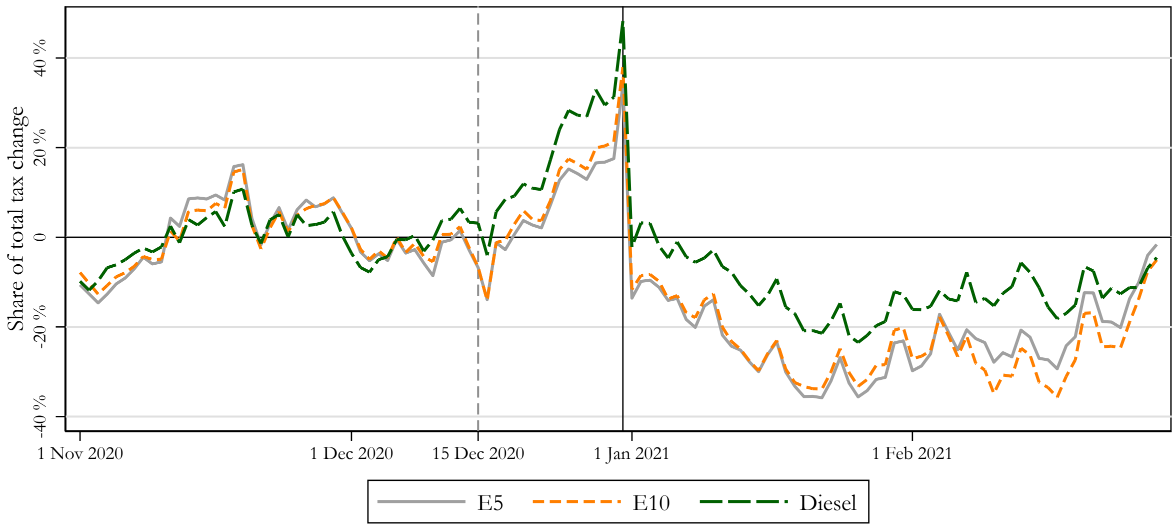

Notes: The solid line shows the nonparametric estimate of the daily average pass-through rate to retail margins for *E5*. The short-dashed and long-dashed lines show analogous estimates for *E10* and diesel, respectively. To estimate pass-through, we first subtract the average pre-period (1 November 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.

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 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

[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.

[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.

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:

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

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

[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.

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

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 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.

[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.

[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.

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 heterogeneities 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

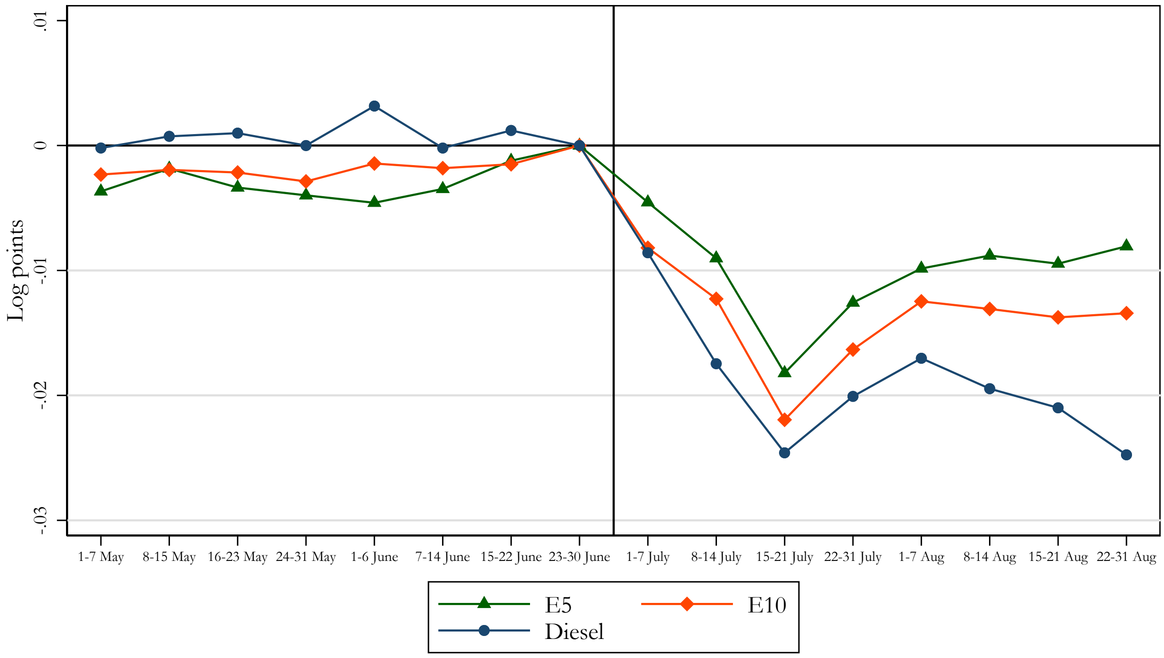

Notes: The graph shows the time-varying effect of the tax decrease on the log prices for *E5*, *E10* and diesel. The analysis period goes from 1 May until 31 August 2020. For the time-varying treatment effects, we estimate the model in Equation 3.3, weighted by the SDID unit and time weights. The vertical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

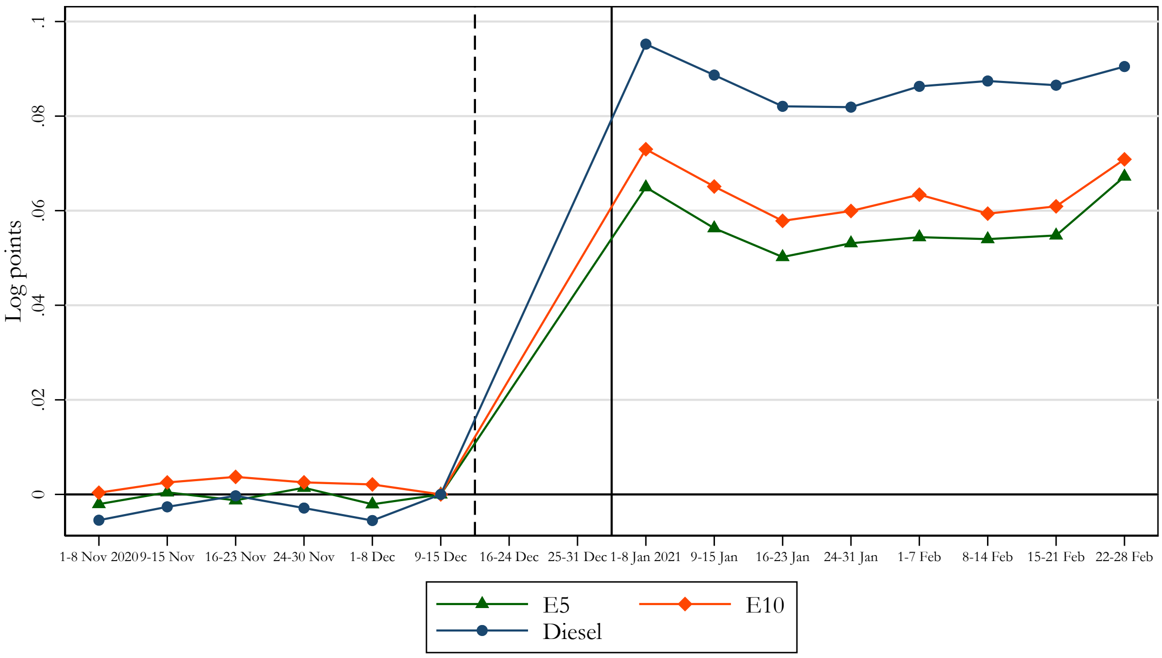

Notes: The graph shows the time-varying effect of the tax increase on log prices for *E5*, *E10* and diesel. The pre-treatment period goes from 1 November until 15 December 2020 and the post-treatment period from 1 January to 28 February 2021. For the time-varying treatment effects, we estimate the model 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.

Prior to the tax increase, the trends in the logarithm of fuel prices are similar between France

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

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]

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]

[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.

[43]In Appendix C.4, we report the results when we additionally control for regional mobility for retail and

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 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

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 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.

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*.

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

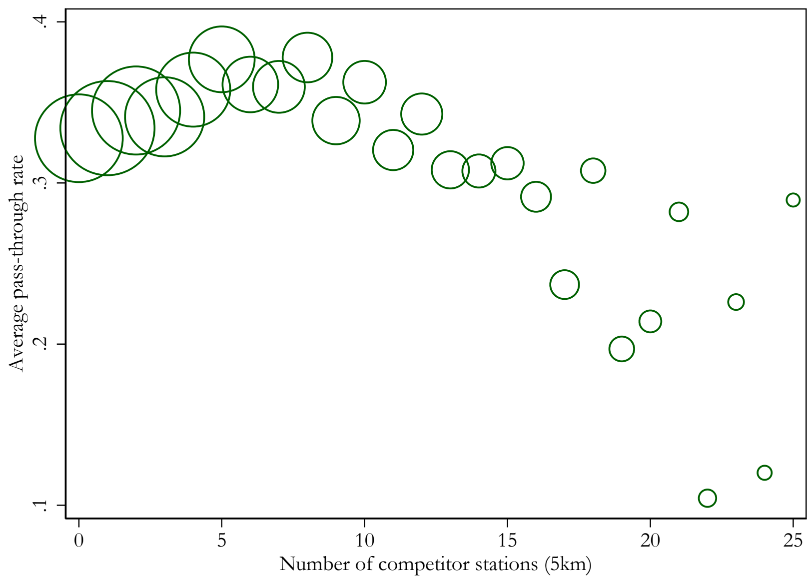

Notes: Each circle plots the average pass-through rate for a group of stations with a particular number of competitors within a 5 km catchment area. The number of competitor stations is trimmed at the top percentile.

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

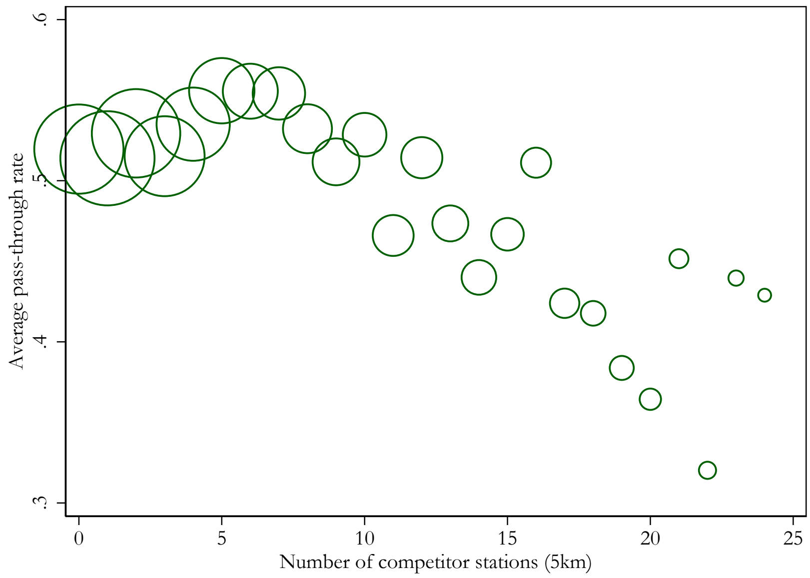

Notes: Each circle plots the average pass-through rate for a group of stations with a particular number of competitors within a 5 km catchment area. The number of competitor stations is trimmed at the top percentile.

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 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

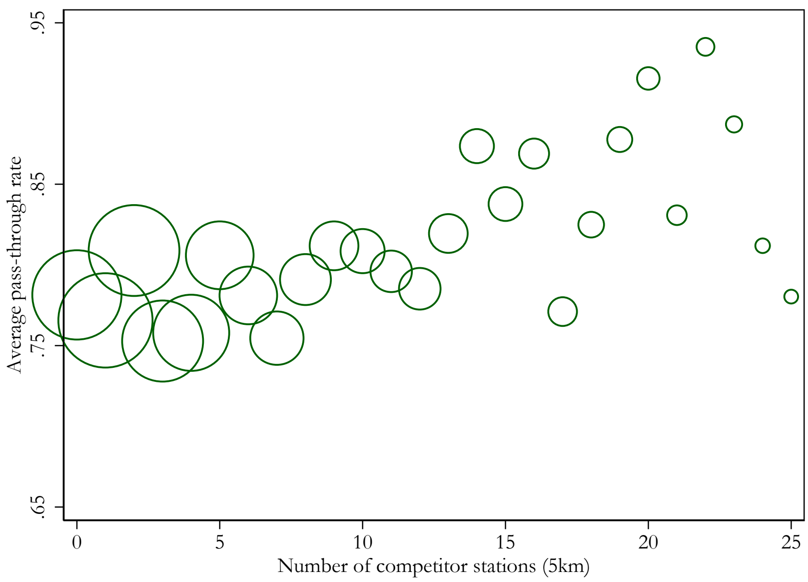

Notes: Each circle plots the average pass-through rate for a group of stations with a particular number of competitors within a 5 km catchment area. The number of competitor stations is trimmed at the top percentile.

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 paper, 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: Details on data set construction

A.1.1 Product market data set

In this section, we add further details on the construction of the product market data set.

Product data. As described in Section 1.2, for clothes washers, a product is defined as the combination of a brand, a retailer and whether the clothes washer is a front-loader, a regular top-loader (with an agitator) or a high-efficiency top-loader (without an agitator). For clothes washers, these are the key differentiating characteristics between products.

Figure A.1 illustrates the difference between a front-loader and a top-loader. Whether the former can be loaded from the front, the latter is loaded from the top. The former can therefore be stacked (i.e. a front-loading dryer can be placed on top of a front-loading washing machine), is more water and energy efficient, cleans better, and is usually more expensive than top-loaders. The latter can never be stacked, however, for top-loaders, there is an important distinction related to whether they have an agitator, which is illustrated in Figure A.2. top-loaders without an agitator are also called high-efficiency top-loaders. In all respects but stacking, they are in between regular top-loaders and front-loaders.[1]

Within a market (here, national at the yearly-level), I group responses that are the same along

[1]See, for example, McCabe (2016) for a detailed comparison of the different clothes washer types.

Figure A.1: Difference between a front-loader and a top-loader

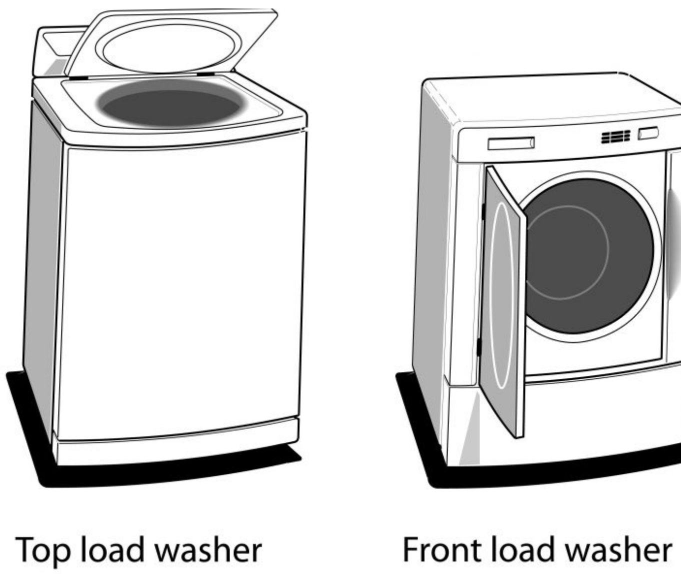

Figure A.2: top-loaders with and without agitator

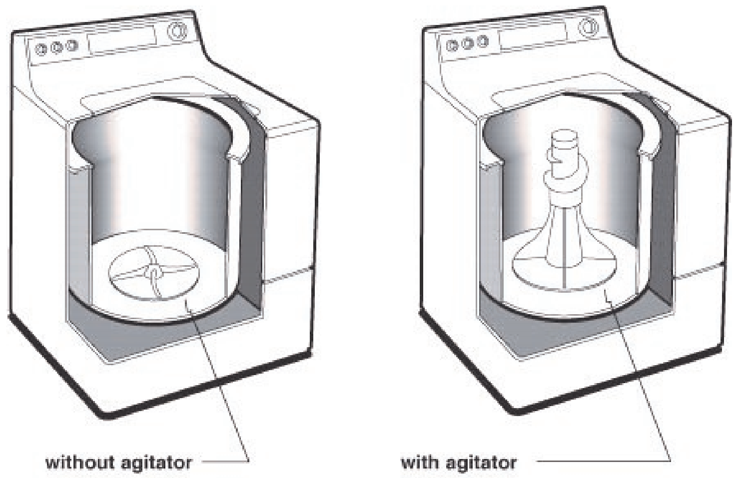

these three dimensions.[2] Doing so, I end up with $2,939$ products between 2005 and 2015. Using this product definition, many products are often very small and based only on a single responding household. Some responses also do not contain information on the brand. I therefore drop all products whose brand response is "Other Brands" or "Store Brand/Generic", as well as all products with a volume share of the clothes washer market of less than 0.01 percent. This results in a final product data set with $1,590$ products. Throughout the years, the remaining products account for between 97.3 and 99.0 percent of the volume share of all clothes washer sales in the *TraQline* data. Dropping very rare products should therefore not bias the estimation results.

For other characteristics, which are only available for a random subset of *TraQline* respondents, I calculate the within-group average of responses for that characteristic. These include whether a clothes washer is part of a stacked pair, whether its exterior is made of stainless steel, is white, or of a different color, whether it is Energy Star certified, has additional noise insulation or a child lockout, as well as the number of special programs it has.

Household income. Whereas the CPS data includes the exact income of the sampled households, the *TraQline* data only includes an income range for each household. To estimate how the price sensitivity of households depends on household income using a single parameter only, I need an exact income for each household. For this, I randomly draw a household income for each respondent based on the empirical distribution of household incomes and the income range that the household falls into.

This involves the following steps:

1. Compute the mid-point of the non-overlapping household income buckets for each response.

2. For each year, fit a log-normal income distribution to the observed household-level income range mid-points.

3. Draw $1,000,000$ incomes from the fitted log-normal income distribution.

[2]For 2006, I classify Maytag products as belonging to Whirlpool also for the first quarter, where the acquisition was not yet carried out. This is to avoid artificially inflating the number of clothes washer products in that year. Also, since merger talks were public since mid-2005, it seems unlikely that Maytag and Whirlpool would still compete heavily in the first quarter of 2006.

4. Allocate each income draw to a particular income bucket.

5. For each household, sample with replacement an income from the set of incomes that correspond to its income bucket.

A.1.2 Plant locations and plant location weights

Plant locations. Constructing the data set on plants manufacturing clothes washers for the U.S. market involves three steps: First, I use information from various sources, such as annual reports, news articles or the United States International Trade Commission's (USITC) anti-dumping hearing transcripts to identify the location of clothes washer plants by the major manufacturers. In many cases, this is insufficient to know whether a plant produces clothes washers for the U.S. or for another market. Second, I use information on the general imports of front-loader and top-loader clothes washers to the U.S. split by source country over time. I use this data to eliminate any plants that cannot plausibly produce substantive volumes for the U.S. market. Finally, I use this data to verify that there are production plants that can plausibly be responsible for the imported volumes for each country from which the U.S. imports substantial numbers of clothes washers.

Figure A.3 shows the evolution of annual imports of front-loaders and top-loaders into the U.S., split by source country. Across the sample period more than half of the front-loaders sold in the U.S. are imported. In 2005, Germany is the largest exporter of front-loaders into the United States. These are not produced by a German manufacturer, but by Whirlpool in its plant in Schorndorf, which was closed in 2012. Until 2012, LG and Samsung imported many of its front-loaders from South Korea and, like other manufacturers such as General Electric or Whirlpool, also from Mexico. After the imposition of anti-dumping duties on large residential clothes washers from Mexico and South Korea in 2012, imports from both countries declined and LG and Samsung moved their production to China (see Flaaen et al., 2020 for an in-depth discussion). In contrast, no country exported more than $50,000$ top-loaders to the U.S. until 2011, aside from a temporary spike in top-loader imports from Mexico in 2006 and 2007. Thereafter, LG and Samsung begin increasing their sales of top-loaders in the U.S. and import most of these from China.

Figure A.3: Clothes washer imports to the United States by source country

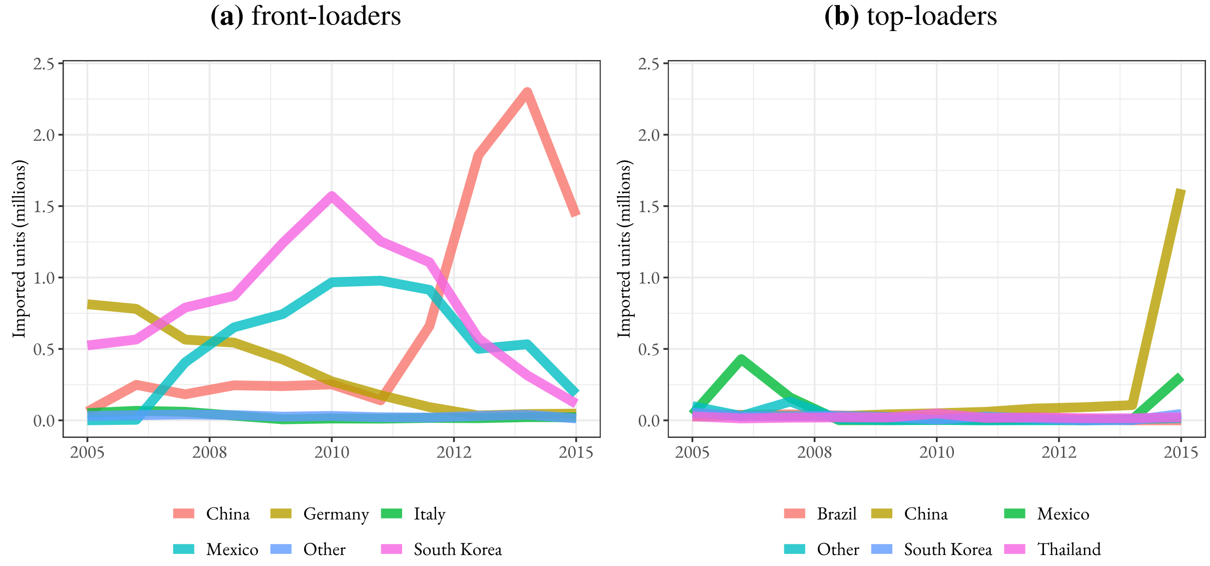

Notes: The left panel plots the annual general imports in terms of volume of front-loader washing machines (HS8450110080, HS8450200080, HS8450200090) imported into the U.S. by source country. The right panel plots the annual general imports in terms of volume of top-loader washing machines (HS8450110040, HS8450200040) imported into the U.S. by source country. The graphs include the top fix importing countries for each product class and groups all other importing countries into "Other". The data comes from the United States International Trade Commission.

For reference, according to Appliance Portrait (2006), 9.3 million clothes washers were sold across the U.S. in 2005. Of those, according to the *TraQline* data, around one-third are front-loaders and the rest top-loaders. The share of front-loaders gradually increased to over 40 percent in 2010 and then decreased again to around 25 percent in 2015. This suggests that although substantial amounts of front-loaders were imported into the U.S. throughout the sample period, most top-loaders were produced domestically.

By combining the clothes washer plant locations of major manufacturers with the USITC import data, I can identify which plants manufacture clothes washers for the U.S. market. Figures A.4, A.5, A.6, and A.7 show the locations of clothes washer plants for all manufacturers that have a volume share of more than 3 percent of the U.S. clothes washer market in any year in the sample.

Plant location weights. Finally, Table A.1 summarizes the plant location weights used to calculate the average real exchange rate for each product. Based on the plant locations, the aggregate USITC import data shown above, and the firm-level clothes washer imports for 2012 until 2015 based on PIERS bill of landing data and reported in Flaaen et al. (2020), these are

Figure A.4: Clothes washer plant locations 2005

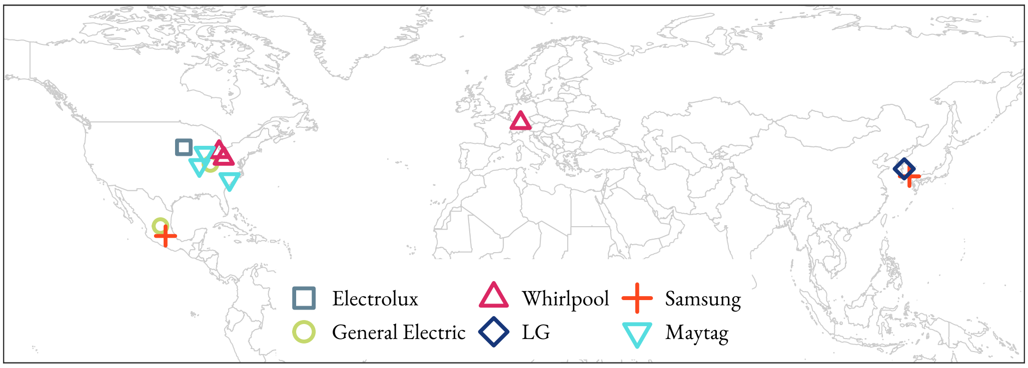

Notes: The map shows all plants manufacturing clothes washers for the U.S. market in 2005 by manufacturers with a market share of more than 3 percent in any year in the sample.

Figure A.5: Clothes washer plants manufacturing for the U.S. market, 2007

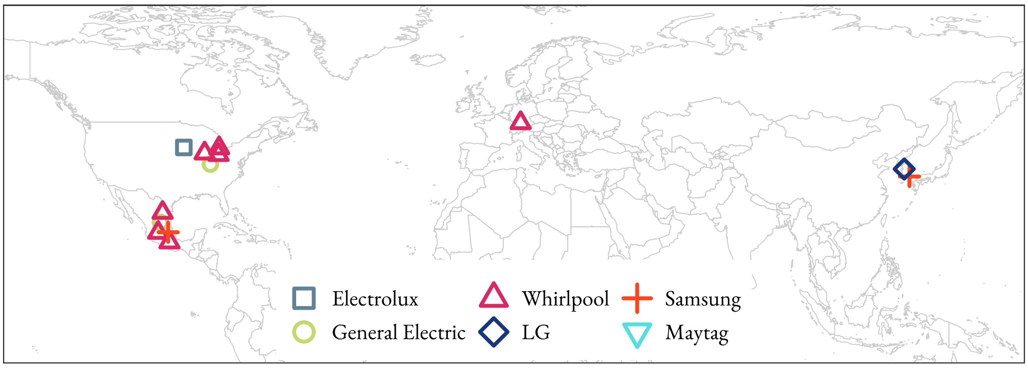

Notes: The map shows all plants manufacturing clothes washers for the U.S. market in 2007 by manufacturers with a market share of more than 3 percent in any year in the sample.

Figure A.6: Clothes washer plants manufacturing for the U.S. market, 2009

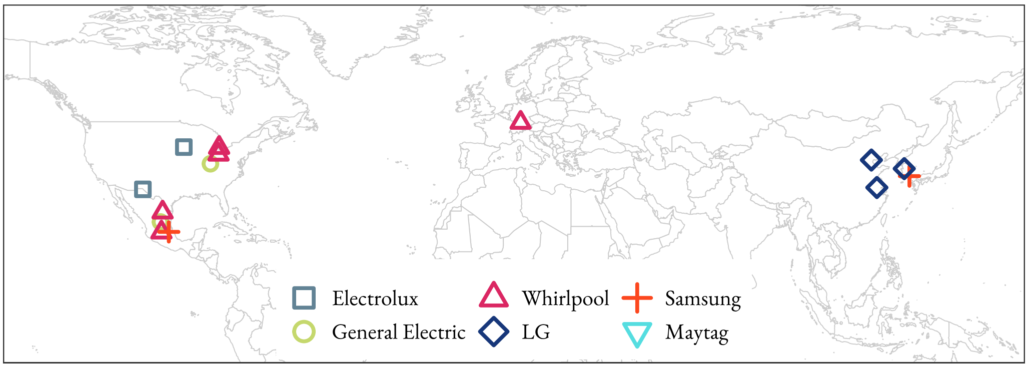

Notes: The map shows all plants manufacturing clothes washers for the U.S. market in 2009 by manufacturers with a market share of more than 3 percent in any year in the sample.

Figure A.7: Clothes washer plants manufacturing for the U.S. market, 2011

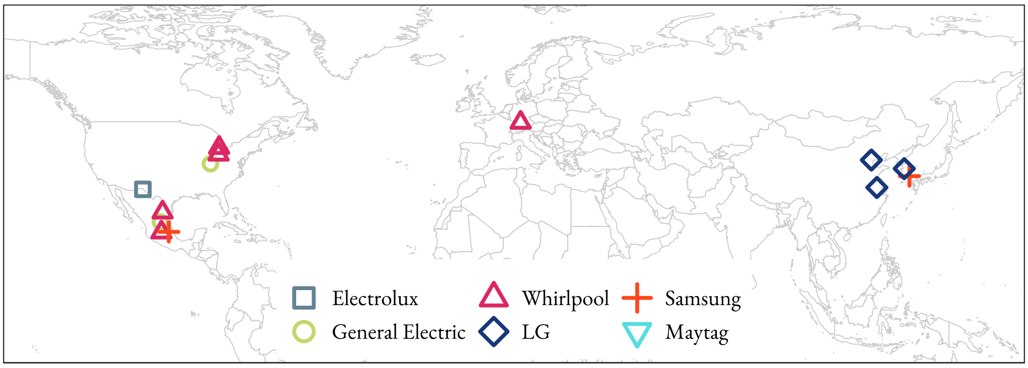

Notes: The map shows all plants manufacturing clothes washers for the U.S. market in 2011 by manufacturers with a market share of more than 3 percent in any year in the sample.

best estimates of which share of a product is sourced from which country in a particular year.

Table A.1: Plant location weights

Owner	Brand	Product	Years	China	Germany	Mexico	South Korea	USA
Electrolux	All brands	Front Loader	2005-2007					1
Electrolux	All brands	Front Loader	2008-2015			1		
Electrolux	All brands	Top Loader	2005-2010					1
Electrolux	All brands	Top Loader	2011-2015			1		
General Electric	All brands	Front Loader	2005-2012			1		
General Electric	All brands	Front Loader	2013-2015					1
General Electric	All brands	Top Loader	2005-2015					1
Whirlpool	All brands	Front Loader	2005		1			
Whirlpool	All other WP brands	Front Loader	2006-2007		1			
Whirlpool	All other WP brands	Front Loader	2008-2010		0.5	0.5		
Whirlpool	All brands	Front Loader	2011		0.33	0.33		0.33
Whirlpool	All brands	Front Loader	2012-2015					1
Whirlpool	Admiral, Amana, Maytag	Front Loader	2006-2010			1		
Whirlpool	Admiral, Amana, Maytag	Front Loader	2010			0.5		0.5
Whirlpool	All brands	Top Loader	2005-2015					1
LG	All brands	Front Loader	2005-2012				1	
LG	All brands	Front Loader	2013	0.67			0.33	
LG	All brands	Front Loader	2014-2015	1				
LG	All brands	Top Loader	2005-2007				1	
LG	All brands	Top Loader	2008-2015	1				
Samsung	All brands	Front Loader	2005-2011				1	
Samsung	All brands	Front Loader	2012	0.5			0.5	
Samsung	All brands	Front Loader	2013-2015	1				
Samsung	All brands	Top Loader	2005-2011			1		
Samsung	All brands	Top Loader	2012-2015	1				
Maytag	All brands	Front Loader	2005-2006					1
Maytag	All brands	Top Loader	2005-2006					1

A.2 Appendix to Section 1.3: Further descriptive evidence

A.2.1 Labor market effects

Figure A.8: Labor market effects of new jobs in Marion County, OH

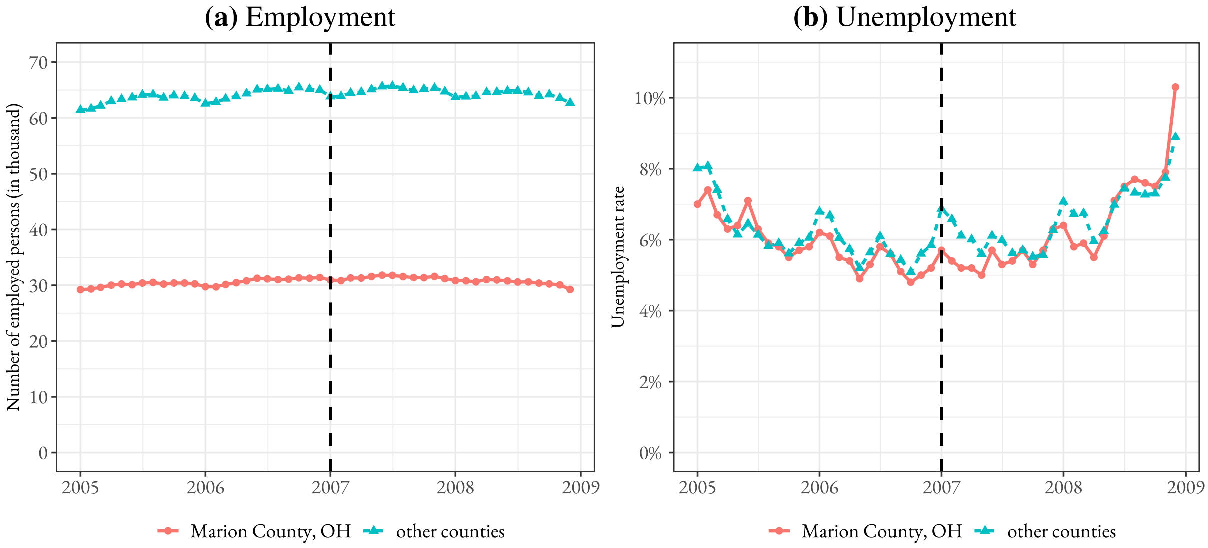

Notes: The solid red lines show the evolution of the total number of employed persons and the unemployment rate in percentages in Marion County, Ohio, respectively. The dashed blue lines show the average number of employed persons and the unemployment rate by county for all other counties in Ohio, respectively. The vertical dashed lines indicate the shut down of operations on 31 December 2006.

Figure A.9: Labor market effects of new jobs in Sandusky County, OH

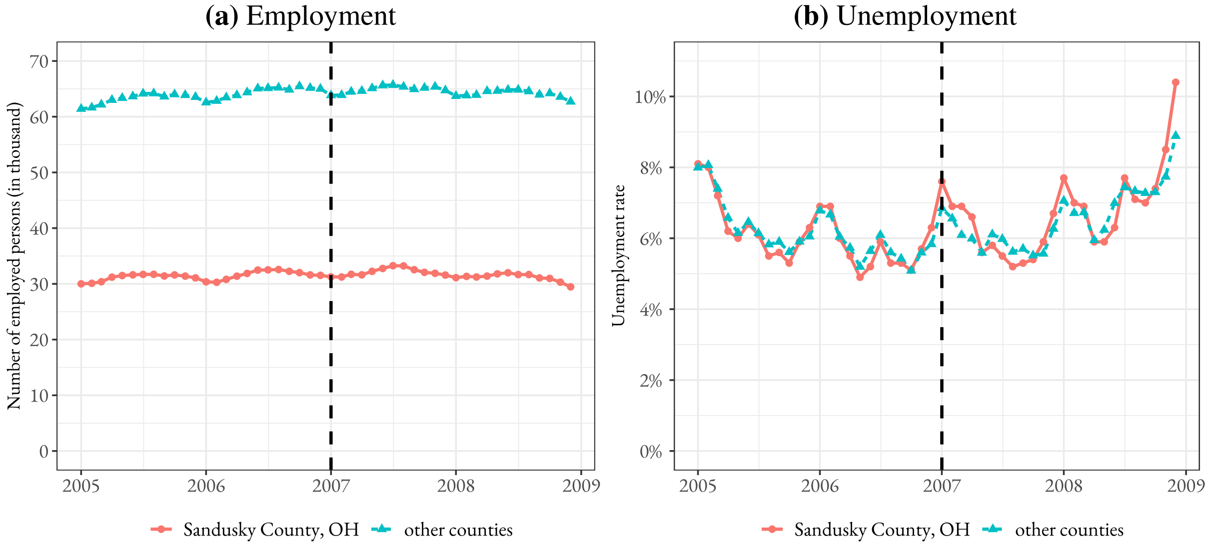

Notes: The solid red lines show the evolution of the total number of employed persons and the unemployment rate in percentages in Sandusky County, Ohio, respectively. The dashed blue lines show the average number of employed persons and the unemployment rate by county for all other counties in Ohio, respectively. The vertical dashed lines indicate the shut down of operations on 31 December 2006.

Table A.2: Reduced form labor market effects of plant and HQ closures (private sector)

	Employment (persons)		Wages ($)	
	(1)	(2)	(3)	(4)
Plant & HQ closure × $\mathbb{1}$ (year = 2007)	-1099***		-3385***	
	[-1275,-923]		[-3619,-3150]	
Plant & HQ closure × $\mathbb{1}$ (year = 2008)	-1621***		-8788***	
	[-1844,-1397]		[-9071,-8505]	
Plant closure × $\mathbb{1}$ (year = 2007)		-299**		-384
		[-532,-66]		[-1208,439]
Plant closure × $\mathbb{1}$ (year = 2008)		-323**		-523
		[-642,-5]		[-1373,327]
County fixed effects	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
Observations	4,752	8,448	1,584	2,804
Mean outcome in treated counties	9,328	11,151	34,022	23,274

Notes: Columns (1) and (2) compare the absolute number of private sector employees in treated counties to all other counties in the same state. Columns (3) and (4) compare the average annualized gross wage of private sector employees in treated counties to all other counties in the same state. Columns (1) and (3) compare Jasper County (county of Newton) to all other counties in Iowa. Columns (2) and (4) compare White County (Searcy) and Williamson County (Herrin) to all other counties in Arkansas and Illinois. Cook County (county of Chicago), is omitted from any analyses involving Illinois. 95% confidence intervals are reported in parentheses. Standard errors are clustered at the county level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.3: Reduced form labor market effects of new jobs (private sector)

	Employment (persons)	Wages ($)
	(1)	(2)
New jobs × $\mathbb{1}$ (year = 2007)	434*	-180
	[-51,918]	[-414,54]
New jobs × $\mathbb{1}$ (year = 2008)	724*	-492
	[-71,1518]	[-1196,211]
County fixed effects	Yes	Yes
Time fixed effects	Yes	Yes
Observations	4,224	1,396
Mean outcome in treated counties	22,373	31,604

Notes: Column (1) compares the absolute number of private sector employees in Marion County (Marion) and Sandusky County (Clyde) to all other counties in Ohio. Column (2) compares the average annualized gross wage of private sector employees in Marion County and Sandusky County to all other counties in Marion County and Sandusky County. 95% confidence intervals are reported in parentheses. Standard errors are clustered at the county level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.4: Reduced form labor market effects of plant and HQ closures (private sector excl. manufacturing)

	Employment (persons)		Wages ($)	
	(1)	(2)	(3)	(4)
Plant & HQ closure × $\mathbb{1}$ (year = 2007)	-389***		-2096***	
	[-554,-223]		[-2311,-1881]	
Plant & HQ closure × $\mathbb{1}$ (year = 2008)	-360***		-1774***	
	[-557,-162]		[-2023,-1524]	
Plant closure × $\mathbb{1}$ (year = 2007)		-11		37
		[-550,529]		[-518,591]
Plant closure × $\mathbb{1}$ (year = 2008)		-28		54
		[-424,368]		[-836,943]
County fixed effects	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
Observations	4,752	8,448	1,584	2,816
Mean outcome in treated counties	7,144	9,660	26,431	22,599

Notes: Columns (1) and (2) compare the absolute number of private sector employees excluding manufacturing in treated counties to all other counties in the same state. Columns (3) and (4) compare the average annualized gross wage of private sector employees excluding manufacturing in treated counties to all other counties in the same state. Columns (1) and (3) compare Jasper County (county of Newton) to all other counties in Iowa. Columns (2) and (4) compare White County (Searcy) and Williamson County (Herrin) to all other counties in Arkansas and Illinois. Cook County (county of Chicago), is omitted from any analyses involving Illinois. 95% confidence intervals are reported in parentheses. Standard errors are clustered at the county level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.5: Reduced form labor market effects of new jobs (private sector excl. manufacturing)

	Employment (persons)	Wages ($)
	(1)	(2)
New jobs × $\mathbb{1}$ (year = 2007)	-227	-430*
	[-650,196]	[-923,63]
New jobs × $\mathbb{1}$ (year = 2008)	-96	-441**
	[-648,455]	[-806,-77]
County fixed effects	Yes	Yes
Time fixed effects	Yes	Yes
Observations	4,224	1,408
Mean outcome in treated counties	22,373	31,604

Notes: Column (1) compares the absolute number of private sector employees excluding manufacturing in Marion County (Marion) and Sandusky County (Clyde) to all other counties in Ohio. Column (2) compares the average annualized gross wage of private sector employees excluding manufacturing in Marion County and Sandusky County to all other counties in Marion County and Sandusky County. 95% confidence intervals are reported in parentheses. Standard errors are clustered at the county level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.6: Reduced form labor market effects of plant and HQ closures (manufacturing sector)

	Employment (persons)		Wages ($)	
	(1)	(2)	(3)	(4)
Plant & HQ closure × $\mathbb{1}$ (year = 2007)	-710***		-802***	
	[-754,-666]		[-1259,-344]	
Plant & HQ closure × $\mathbb{1}$ (year = 2008)	-1261***		-22255***	
	[-1324,-1198]		[-22858,-21652]	
Plant closure × $\mathbb{1}$ (year = 2007)		-288		-1080
		[-758,181]		[-4105,1945]
Plant closure × $\mathbb{1}$ (year = 2008)		-296		-2638***
		[-824,233]		[-3870,-1406]
County fixed effects	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
Observations	4,752	8,448	1,572	2,659
Mean outcome in treated counties	2,185	1,491	58,368	27,603

Notes: Columns (1) and (2) compare the absolute number of employees in the manufacturing industry in treated counties to all other counties in the same state. Columns (3) and (4) compare the average annualized gross wage of employees in the manufacturing industry in treated counties to all other counties in the same state. Columns (1) and (3) compare Jasper County (county of Newton) to all other counties in Iowa. Columns (2) and (4) compare White County (Searcy) and Williamson County (Herrin) to all other counties in Arkansas and Illinois. Cook County (county of Chicago), is omitted from any analyses involving Illinois. 95% confidence intervals are reported in parentheses. Standard errors are clustered at the county level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.7: Reduced form labor market effects of new jobs (manufacturing sector)

	Employment (persons)	Wages ($)
	(1)	(2)
New jobs × $\mathbb{1}$ (year = 2007)	661***	-232
	[377,944]	[-1640,1177]
New jobs × $\mathbb{1}$ (year = 2008)	820***	-965
	[507,1133]	[-3074,1144]
County fixed effects	Yes	Yes
Time fixed effects	Yes	Yes
Observations	4,224	1,403
Mean outcome in treated counties	7,938	41,344

Notes: Column (1) compares the absolute number of employees in the manufacturing industry in Marion County (Marion) and Sandusky County (Clyde) to all other counties in Ohio. Column (2) compares the average annualized gross wage of number of employees in the manufacturing industry in Marion County and Sandusky County to all other counties in Marion County and Sandusky County. 95% confidence intervals are reported in parentheses. Standard errors are clustered at the county level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.3 Appendix to Section 1.5: Details on the estimation procedures

A.3.1 Details on estimating product characteristics for potential products

Potential products are all products that brand owners added to the market (active products), as well as all products that they could have added but did not (inactive products). Estimating the former is easy, since we can simply observe these in the market. Estimating the latter is more complicated.

The focus of the analysis in this paper lies on the decision of firms to add or remove products that they are technologically already capable of making. For example, if a firm does not carry front-loading washing machines, these will also not be part of its potential products. If, for example, Maytag sells regular top-loading washing machines under its Amana brand at Best Buy and Lowe's, but not at other major retailers, Amana regular top-loaders at other major retailers are potential products.[3]

[3]Major retailers are Best Buy, H. H. Gregg, Home Depot, Lowe's and Sears.

Product characteristics can mildly vary between retailers. That is, Amana top-loaders sold at Best Buy might modestly differ in their characteristics compared to Amana top-loaders sold at Lowe's. In the example, Amana regular top-loaders at Sears are an inactive product. To determine the exact product characteristics of this inactive product, I need to decide whether to attribute it the characteristics of the Amana regular top-loader sold at Best Buy or at Lowe's.

Whenever a particular combination of brand and key characteristic exists at two or more retailers, I use the following ordering of "closest" retailers to match other product characteristics:

- **Sears**: Home Depot, Lowe's, Best Buy, H. H. Gregg, Others
- **Home Depot**: Lowe's, Sears, Best Buy, H. H. Gregg, Others
- **Lowe's**: Home Depot, Best Buy, Sears, H. H. Gregg, Others
- **Best Buy**: Lowe's, H. H. Gregg, Home Depot, Sears, Others
- **H.H. Gregg**: Best Buy, Lowe's, Home Depot, Sears, Others

A.3.2 Details on the demand estimation

The estimation of the demand parameters follows S. Berry, Levinsohn, et al. (2004) and proceeds in two steps. First, I search for estimates $\hat{\kappa}_\alpha$ and $\hat{\sigma}_{FL}$ (jointly denoted by $\hat{\theta}_2$ of the non-linear parameters, as well as of the vector of mean utilities δ. Next, I estimate $\hat{\beta}$ for the vector of linear demand parameters. Wherever possible, I implement the best practices described by C. Conlon and Gortmaker (2020). For notational simplicity, I omit the time subscript t in this section. The details of the technical implementation should thus be seen as valid for a single market t and then repeated and averaged over markets.

The estimation of the non-linear parameters and the mean utilities proceeds in two iterative steps: In the inner loop, I search for the mean utilities given a guess of the non-linear parameters. In the outer loop, I search for the non-linear parameters that minimize the objective function, solving the inner loop at each step.

The first set of moments equates the observed market shares in the data with the simulated market shares from the demand model. To get an estimate $\hat{\delta}$ of the mean utilities, I proceed as follows: First, as described by S. Berry (1994), I invert the market share function $s_j(\delta_j;\theta)$ to obtain $\delta_j(s^n_j, s_j(\delta_j;\theta))$, where s^n_j denotes the market shares observed in the data and $s_j(\delta_j;\theta)$ denotes the simulated market shares implied by the model and the parameter vector θ.[4] Second, I use the fixed-point formulation due to S. Berry, Levinsohn, et al. (1995) to estimate $\hat{\delta}_j$. I use the SQUAREM described in Reynaerts, Varadha, and Nash (2012) to accelerate the convergence of the fixed-point iterations. As this is not guaranteed to converge, whenever convergence fails, I revert to the contraction mapping in S. Berry, Levinsohn, et al. (1995) which has guaranteed convergence. Finally, I speed up the inversion of market shares by using the reformulation of the contraction mapping in terms of consumer-specific choice probabilities for the outside option, described by Brunner, Heiss, Romahn, and Weiser (2020).

To estimate the market shares implied by the estimate $\hat{\theta}$ of the parameter vector, the model and the data, I need to solve the integral in Equation 1.5. As is standard in the literature, I approximate this integral using Monte Carlo simulations by drawing household demographics and unobserved taste shocks from the joint empirical distribution for 1000 households. Household demographics come from the CPS. I draw unobserved taste shocks from a standard normal distribution, using scrambled Halton draws (see Owen, 2017).

The second set of moments fits the covariance between the price of the first-choice clothes washer and the average income of households purchasing the product. I compute the moment as follows

$$\sum_j \frac{n_j}{n} p_j \left\{ \left(\frac{1}{n_j} \sum_{i s.t. y^1_i = j} z_i \right) - E\left[z | y^1 = j, \theta \right] \right\}, \tag{A.1}$$

where J continues to denote a product, n denotes the total number of households, n_j denotes the number of households buying good j, y^1_i denotes the first choice product of household i, p_j continues to denote the price of product j, and z_i the income of household i.

The third set of moments fits the covariance between whether the first-choice clothes washer is a front-loader and the share of front-loaders among products of the second choice brand. In contrast to S. Berry, Levinsohn, et al. (2004), I do not observe the exact second-choice product

[4] Note, that $s_j(\delta_j;\theta)$ also depends on the product and household characteristics, which I omitted to simplify notation.

but only the second-choice brand. In particular, I use the following moment condition

$$\sum_j \left(\frac{n_j}{n} x_j^{FL} \sum_{b' \neq b_j} x_{b'}^{FL} \left\{ \frac{n_{jb'}}{n_j} - E\left[\mathbb{1}\left(b^2 = b' | y^1 = j, \theta \right) \right] \right\} \right), \tag{A.2}$$

where b denotes a brand, b_j denotes the brand of product j, b^2 denotes the brand of the second choice, x_j^{FL} indicates whether product j is a front-loader and $x_{b'}^{FL}$ denotes the volume-weighted share of front-loaders among products sold of brand b.

The objective function that I minimize in the outer loop to estimate $\hat{\theta}_2$ consists of the moments in Equations A.1 and A.2. Since there are two nonlinear parameters and two moment conditions, the parameters are just-identified and we estimate $\hat{\theta}_2$ using the method of simulated moments. I therefore estimate

$$\hat{\theta}_{2,MSM} = \text{argmin}\, \hat{m}(\theta_2)' \, \hat{m}(\theta_2) \; . \tag{A.3}$$

Solving the minimization problem above does not only allow recovering the nonlinear parameters of the demand model, but also the mean utilities $\hat{\delta}$. In the final step, I estimate the linear parameters of the demand model using the following specification:

$$\hat{\delta}_j = x_j \beta - \alpha p_j + \xi_j \, . \tag{A.4}$$

As explained in Section 1.5, I assume that the non-price product characteristics are independent of unobserved quality differences ξ_j, whereas the price can be correlated with these unobserved differences. To solve the endogeneity problem, I use an instrumental variables estimator, where the product-level real exchange rate serves as a cost shifting instrumental variable for price, as described in Section 1.2.

Market size and share of the outside good

To compute the total market size, I assume that every seventh household is a potential purchaser of a clothes washer in a particular year. According to Consumer Reports, in 2009 the average life expectancy of a clothes washer was ten years. Many households will consider buying a clothes washer already before the end of the life expectancy of their washer, e.g. to get a new washer with novel features. Some households will consider new washing machines for multiple years. Households that recently purchased a washer are unlikely to be on the lookout for a new one immediately. It therefore seems plausible that the true market size is somewhere between a fifth and a tenth of the number of households. The results are robust to alternative market size assumptions.

To compute firm profits, consumer welfare and estimate entry cost bounds in Dollar terms for the U.S. population, I need to scale the estimates by the number of households that are in the market for clothes washers in a particular year. There are two alternative estimation methods: We can take the total number of U.S. households in a particular year and assume that the market size is one seventh of these households. Alternatively, I can use estimates of the annual total clothes washers shipped as reported by Appliance Portrait and divide this by the share of the inside good. Both methods yield similar results for the years around the merger date and so I assume that the total market size in the U.S. is around 15 million households.

A.3.3 Speeding up the computation of expected profits

Both, the estimation of fixed costs, and the heuristic entry algorithm require computing the expected profits of firms for many different product portfolios. This is computationally costly and since it has to be repeated many times, speeding up this process is crucial. In the following, I briefly describe the key elements that helped speed up the computations for this paper.[5]

Computing equilibrium prices. Each draw of the second-stage marginal cost and demand shocks e_{jt} requires re-estimating the equilibrium price vector for all active products. Since I use 500 draws of e_{jt} to approximate the expected variable profits for a single product portfolio, it is also necessary to re-compute equilibrium prices 500 times for each product portfolio. Speeding up this process is therefore crucial. Furthermore, not all methods to re-compute equilibrium prices necessarily converge.

Morrow and Skerlos (2011) compare different numerical methods to re-compute equilibrium prices using the Nash-Bertrand first order conditions. They find that applying Newton methods to this problem is reliable but slow. On the other hand, they show that fixed point iteration on

[5] As noted in the Online Appendix to Wollmann (2018), implementing the computations in Julia has significant speed advantages, as it can handle loop commands at comparable speed to "vectorized" code in Matlab.

the BLP-markup equation need not converge and is slow. Instead, they propose a reformulated markup equation, the ζ-markup, which is fast and reliable. I therefore compute equilibrium prices by using fixed point iteration on the ζ-markup equation.

Drawing e_{jt}. The heuristic algorithm to choose product portfolios requires comparing the expected profits of the current product portfolio to the expected profits of any product portfolio that is within a one-step change of the current product portfolio. This involves revisiting the same product portfolios many times.

An important feature of the heuristic portfolio choice algorithm is to use the same e_{jt} draws for the same product when computing the expected profits of different product portfolios. In terms of economics, this is desirable because there is no good reason for why a firm should form its expectation about demand and cost shocks for a product differently based on what other products are in the market. In terms of computations, this is desirable because it means that I only need to compute expected profits of all firms for a given set of product portfolios once. Every time that the algorithm re-visits the particular set of product portfolios, I can re-use the memorized expected profits and do not need to re-compute equilibrium prices and expected profits.

A.3.4 Details on the fixed cost estimation

I follow the approach proposed by Eizenberg (2014) and fill the missing bounds by adding two further assumptions.

Assumption 3.1. $\sup_{j\in\mathcal{J}_{bt}} F_{jt} = F_{bt}^U < \infty$ *and* $\inf_{j\in\mathcal{J}_{bt}} F_{jt} = F_{bt}^L > -\infty$ *(bounded support)*

Assumption 3.1 states that the fixed costs associated with introducing a new product have a bounded support. This assumption does not need to be fulfilled in all contexts. If F_{jt} is the cost of developing a new breakthrough technology, it could be that no money in the world makes the necessary invention possible. Since I consider F_{jt} to be the cost of introducing a product at a new retailer and developing new products interior to a firm's technological capability frontier, it seems plausible that there exists an upper-bound to the necessary fixed costs. At the same time, the cost of developing and introducing a new product in this context should never be negative and so the existence of a lower-bound of the fixed cost support, F_b^L, is an innocuous assumption.

Assumption 3.2. $[F_b^L, F_b^U] \subset$ *supp(expected change in variable profit due to the elimination or addition of a single product of brand b).*

Assumption 3.2 adds further restrictions on the support of F_{jt}. For each brand b, the support of the fixed costs of introducing any potential product is contained within the support of expected changes in variable profits of firm f if any potential product of brand b is introduced. The intuition behind this assumption is quite simple. If fixed costs of introducing different potential products of a particular brand come from the same distribution and there exists a blockbuster product that increases expected variable products of the firm so much, that it would always be introduced, then I observe this product as an active product in the data and the expected change in variable profit of adding this product must be higher than the fixed cost of introducing any potential product. Similarly, if there exists a product that has such a small impact on the expected change in variable profit, such that it would never be introduced, then I will always observe this product as an inactive product and the expected change in variable profit of adding this product must be lower than the fixed cost of introducing any potential product.

With these additional assumptions, I can fill the missing upper- and lower-bounds on the fixed costs of potential products. I fill the missing lower-bound on fixed costs for active products by using the minimum change in firm-level expected variable profits among inactive products of the same brand. I fill the missing upper-bound on fixed costs for inactive products by using the maximum change in firm-level expected variable profits among active products of the same brand. The product-level bounds on fixed costs for active and inactive products are defined as

$$L_{jt}(\theta) = \begin{cases} VP_{bt}^L(\theta) & j \in J_{bt} \\ \underline{F}_{jt}(\theta) & j \in \tilde{J}_{bt} \end{cases} \qquad U_{jt}(\theta) = \begin{cases} \bar{F}_{jt}(\theta) & j \in J_{bt} \\ VP_{bt}^U(\theta) & j \in \tilde{J}_{bt} \end{cases}.$$

Since $E[\upsilon_{jt} | j \in \mathcal{J}_{ft}] = 0$, and with estimates on the upper- and lower-bound on fixed costs for all $j \in \mathcal{J}_{ft}$, I can now apply an unconditional expectation, such that

$$E[L_{jt}(\theta)] \leq F_{bt} \leq E[U_{jt}(\theta)] \quad \forall j \in \mathcal{J}_{bt}\,. \tag{A.5}$$

To estimate the set in A.5, I replace the true parameter vector θ by the first stage estimator $\hat{\theta}$ and estimate the change in firm-level variable profits of removing any active product and adding any inactive product in the data. I use $\min_{j \in \tilde{J}_{bt}}\{\underline{F}_{jt}(\hat{\theta})\}$ as an estimator for $VP^{L}_{bt}(\theta)$ and $\max_{j \in J_{bt}}\{\bar{F}_{jt}(\hat{\theta})\}$ as an estimator for $VP^{U}_{bt}(\theta)$.

Finally, I compute the within brand and market sample average across $L_{jt}(\hat{\theta})$ and $U_{jt}(\hat{\theta})$, to estimate bounds on the set of brand- and market-level fixed costs. This estimation procedure produces unbiased estimates and overall leads to wide and conservative fixed cost bounds.

A.3.5 Details on the employment calibration

To simulate the employment effects of the different hypothetical acquisitions, I need an estimate of how many clothes washers a manufacturing worker produces on average per year. Since I do not have systematic data on employment by manufacturer and appliance category, I calibrate the number of clothes washers produced by manufacturing workers based on different sources.

In 2005, Maytag produced clothes washers and dryers in Newton, Iowa (1,000 manufacturing jobs) and Herrin, Illinois (1,000 manufacturing jobs) and dryers in Searcy, Arkansas (700 manufacturing jobs).[6] In addition, there was a small plant manufacturing clothes washers and dryers in Florence, South Carolina (60 manufacturing jobs).[7] According to Appliance Portrait (2006), Maytag shipped 1.75 million clothes washers and 1.6 million dryers in 2005. On average, these are around 1,200 clothes washers and dryers per manufacturing worker per year.

In 2011, the Whirlpool plant manufacturing front-loading clothes washers in Schorndorf, Germany, had 500 employees and produced 200,000 clothes washers.[8] This amounts to 400 clothes washers per manufacturing worker per year.

To simplify matters, I assume that the number of employees necessary to produce clothes washers linearly increases in the number of clothes washers and that this technology is constant over time and across manufacturers, products, and production locations. With richer data and depending on the institutional context, all of these assumptions can be relaxed.

[6] See https://www.nbcnews.com/id/wbna12718867.

[7] See https://www.twice.com/news/maytag-close-florence-laundry-facility-27876.

[8] See https://www.stuttgarter-zeitung.de/inhalt.bauknecht-in-schorndorf-konzern-gibt-den-standort-auf.2559fd28-6719-48b9-a055-5956c7f61c03.html.

Based on the evidence described above, I calibrate that a manufacturing worker produces on average around $1,000$ clothes washers per year. Among clothes dryers, top-loading washers, and front-loading washers, the first are the simplest products to produce and the last the most complex. It therefore seems plausible that the estimate for Whirlpool front-loaders is an overall underestimate of the number of clothes washers produced by worker and the estimate based on Maytag washers and dryers an overestimate. Either way, choosing a relatively high number of clothes washers per manufacturing worker is a conservative approach, since it likely underestimates the employment effects of either acquisition.

A.4 Appendix to Section 1.6: Further results of the structural estimation

A.4.1 Demand estimation

Table A.8: Detailed estimates of linear demand parameters

	(1)	(2)	(3)	(4)	(5)
	First-stage	Reduced form	Logit OLS	Logit IV	Mixed logit IV
Dependent variable:	Price	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$	$\hat{\delta}_{jt}$
Real exchange rate	1.909***	-0.787**			
	(0.398)	(0.358)			
Price ('00 2012 $)			-0.164**	-0.412**	-0.614***
			(0.062)	(0.202)	(0.024)
Front-loader	0.174	0.267	0.358	0.339	-0.686***
	(0.205)	(0.267)	(0.244)	(0.215)	(0.019)
Korean	-0.563***	1.746***	1.569***	1.514***	1.483***
front loader	(0.179)	(0.353)	(0.349)	(0.348)	(0.011)
Fisher & Paykel	-4.506***	-0.624	-1.455***	-2.481***	-3.116***
front-loader	(0.331)	(0.412)	(0.480)	(0.859)	(0.093)
European high-end	0.071	1.235***	1.192***	1.264*	1.272***
front-loader	(1.311)	(0.314)	(0.438)	(0.715)	(0.032)
Agitator	-2.510***	0.952***	0.540**	-0.083	-0.449***
	(0.276)	(0.270)	(0.252)	(0.532)	(0.060)
Stacked pair	0.493*	-0.225	-0.147	-0.022	0.027**
	(0.280)	(0.149)	(0.149)	(0.202)	(0.011)
Stainless steel	0.481	-0.052	0.009	0.146	0.180***
exterior	(0.603)	(0.247)	(0.270)	(0.362)	(0.011)
White exterior	-0.289	0.677***	0.624***	0.558***	0.506***
	(0.360)	(0.130)	(0.101)	(0.131)	(0.009)
Energy Star	0.023	0.089	0.092	0.099	0.110***
	(0.182)	(0.126)	(0.126)	(0.138)	(0.004)
Extra noise	0.395*	0.248**	0.312**	0.411**	0.466***
insulation	(0.207)	(0.125)	(0.120)	(0.162)	(0.010)
Number of special	0.009	0.050	0.052	0.054	0.052***
programs	(0.058)	(0.035)	(0.039)	(0.047)	(0.001)
Child lockout	-0.073	0.204	0.200	0.174	0.174***
	(0.164)	(0.172)	(0.167)	(0.171)	(0.005)
Repair rate	-2.397	2.048	1.627	1.060	0.733***
	(3.156)	(3.272)	(2.957)	(2.793)	(0.129)
Total advertising	-0.006	0.004	0.003	0.001	0.001***
expenditure	(0.005)	(0.002)	(0.002)	(0.002)	(0.0002)
Retailer FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Brand FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Brand time trends	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Year FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Observations	1,590	1,590	1,586	1,590	1,590
Kleibergen-Paap F-statistic	22.979				
Avg. own-price elasticity			-0.964	-2.416	-3.258

Notes: Column (1) presents results for the first stage regression of prices on the real exchange rate. Column (2) includes reduced form estimates for the simple logit model. Column (3) reports demand estimates for the simple logit without instrumenting for price. Column (4) presents demand estimates for the simple logit model using the RER as an instrumental variable for price. Column (5) shows demand estimates for the full mixed logit model presented in Section 1.4 and using the RER as an instrumental variable for price. Standard errors are clustered at the brand level. The own-price elasticity of residual demand is computed at the product level and the average is calculated by weighting products according to their sales volume. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.5 Appendix to Section 1.7: Details on the welfare effects

A.5.1 Entry algorithm

The pseudo-code in Algorithm 1 illustrates this portfolio choice algorithm described in Section 1.7.

Algorithm 1: Heuristic portfolio choice algorithm

```
adjust_B = 1;
J^* ← pre-merger portfolio of all firms;
while adjust_B = 1 do                    /* until no adjustment for any brand */
    adjust_B = 0;
    for b ∈ B do                                   /* iterate through brands */
        adjust_b = 1;
        while adjust_b = 1 do              /* until no adjustment for brand b */
            ΔE[Π_remove] = [];
            ΔE[Π_add] = [];
            for j ∈ J_b do                    /* iterate through active products */
                append(ΔE[Π_remove], E_e[Π_f(J^*_f − 1^j_b)]);
            end
            for j ∈ J̃_b do                  /* iterate through inactive products */
                append(ΔE[Π_add], E_e[Π_f(J^*_f 1^j_b)]);
            end
            if max(max(ΔE[Π_remove]), max(ΔE[Π_add])) > 0 then
                if max(ΔE[Π_remove]) > max(ΔE[Π_add]) then
                    k = findmax( ΔE[Π_remove] );
                    J^*_f ← J^*_f − 1^k_b;
                    adjust_B = 1;
                else
                    k = findmax( ΔE[Π_add] );
                    J^*_f ← J^*_f + 1^k_b;
                    adjust_B = 1;
                end
                adjust_b = 0;
            end
        end
    end
end
```

A.5.2 Detailed product market effects

Table A.9: Product market effects of Maytag acquisitions by Haier and Whirlpool

	No portfolio adjustments				**Merger-independent adjustments**				**Endogenous adjustments**			
Efficiencies:	**No efficiencies**		**Offshoring**		**No efficiencies**		**Offshoring**		**No efficiencies**		**Offshoring**	
Acquirer:	Haier	Whirlpool	Haier	Whirlpool	Haier	Whirlpool	Haier	Whirlpool	Haier	Whirlpool	Haier	Whirlpool
Prices												
Total industry	0.0%	2.7%	-0.1%	2.6%	0.1%	2.9%	-0.9%	2.7%	[-0.9%, 1.1%]	[1.8%, 5.1%]	[-1.9%, 0.2%]	[1.7%, 4.9%]
	[-0.0%, 0.3%]	[1.7%, 3.7%]	[-2.0%, 0.1%]	[1.5%, 3.6%]	[0.0%, 0.2%]	[2.0%, 3.8%]	[-1.0%, -0.8%]	[1.8%, 3.6%]				
Maytag	0.0%	5.8%	-5.9%	4.7%	0.0%	8.8%	-5.2%	7.5%	[-2.1%, 2.7%]	[3.6%, 11.5%]	[-6.2%, 1.6%]	[3.2%, 11.3%]
	[-0.1%, 0.2%]	[2.8%, 8.8%]	[-10.8%, 1.0%]	[1.6%, 7.8%]	[-0.0%, 0.1%]	[5.6%, 12.0%]	[-5.3%, -5.1%]	[4.3%, 10.6%]				
Whirlpool	0.0%	5.9%	0.0%	6.2%	0.0%	6.1%	-0.2%	6.1%	[-3.0%, 3.0%]	[1.1%, 9.8%]	[-3.1%, 1.9%]	[1.0%, 10.0%]
	[-0.0%, 0.0%]	[3.0%, 8.9%]	[-0.2%, 0.2%]	[3.2%, 9.2%]	[-0.0%, 0.0%]	[3.3%, 8.9%]	[-0.2%, -0.2%]	[3.2%, 8.9%]				
Consumer welfare												
All consumers	$-0.7M	$-131M	$83M	$-116M	$-0.9M	$-156M	$61M	$-140M	[$-32M, $33M]	[$-300M, $-197M]	[$58M, $152M]	[$-302M, $-197M]
	[$-5M, $3M]	[$-206M, $-55M]	[$12M, $155M]	[$-191M, $-42M]	[$-3M, $1M]	[$-237M, $-75M]	[$59M, $63M]	[$-221M, $-59M]				
	-0.0%	-4.9%	3.1%	-4.3%	-0.0%	-4.9%	1.9%	-4.4%	[-1.1%, 1.1%]	[-10.0%, -6.6%]	[1.9%, 5.2%]	[-10.1%, -6.6%]
	[-0.2%, 0.1%]	[-8.4%, -1.3%]	[0.3%, 5.9%]	[-7.8%, -0.9%]	[-0.0%, 0.0%]	[-8.0%, -1.8%]	[1.9%, 2.0%]	[-7.5%, -1.3%]				
Variable profits												
Total industry	$0.8M	$66M	$14M	$76M	$0.6M	$81M	$26M	$85M	[$-18M, $23M]	[$30M, $97M]	[$42M, $91M]	[$24M, $97M]
	[$-2M, $4]	[$27M, $105M]	[$-38M, $65M]	[$31M, $120M]	[$-1M, $-2M]	[$41M, $121M]	[$24M, $27M]	[$46M, $125M]				
	0.0%	3.7%	0.8%	4.2%	0.0%	4.1%	1.3%	4.3%	[-0.9%, 1.2%]	[1.5%, 4.9%]	[2.1%, 4.7%]	[1.2%, 4.9%]
	[-0.1%, 0.2%]	[1.2%, 6.1%]	[-2.3%, 3.8%]	[1.5%, 6.9%]	[-0.0%, 0.1%]	[2.0%, 6.2%]	[1.2%, 1.4%]	[2.2%, 6.5%]				
Maytag + Whirlpool	$0.8M	$15M	$14M	$33M	$0.02M	$18M	$68M	$30M	[$-29.1M, $35.3M]	[$-67M, $24M]	[$90M, $186M]	[$-69M, $21M]
	[$-1M, $2M]	[$-23M, $54M]	[$-72M, $99M]	[$-11M, $77M]	[$-1M, $1M]	[$-18M, $54M]	[$67M, $68M]	[$-6M, $66M]				
	0.1%	1.7%	3.2%	3.7%	0.0%	1.8%	17.6%	3.0%	[-5.6%, 7.0%]	[-5.4%, 1.8%]	[15.9%, 38.8%]	[-5.5%, 1.6%]
	[-0.0%, 0.5%]	[-1.9%, 5.3%]	[-16.7%, 23.1%]	[-0.4%, 7.8%]	[-0.1%, 0.1%]	[-1.2%, 4.8%]	[17.5%, 17.7%]	[0.0%, 6.0%]				
Total profits												
Total industry	$0.8M	$66M	$14M	$76M	$0.6M	$81M	$26M	$85M	[$-19M, $19M]	[$76M, $133M]	[$18M, $74M]	[$78M, $132M]
	[$-2M, $4]	[$27M, $105M]	[$-38M, $65M]	[$31M, $120M]	[$-1M, $-2M]	[$41M, $121M]	[$24M, $27M]	[$46M, $125M]				
									[-1.5%, 1.4%]	[5.6%, 10.3%]	[1.4%, 5.8%]	[5.8%, 10.2%]
Maytag + Whirlpool	$0.8M	$15M	$14M	$33M	$0.02M	$18M	$68M	$30M	[$-5.8M, $6.4M]	[$19M, $47M]	[$91M, $119M]	[$20M, $48M]
	[$-1M, $2M]	[$-23M, $54M]	[$-72M, $99M]	[$-11M, $77M]	[$-1M, $1M]	[$-18M, $54M]	[$67M, $68M]	[$-6M, $66M]				
									[-12.8%, 11.8%]	[2.3%, 6.3%]	[-92.5%, 189.4%]	[2.4%, 6.2%]

Notes: The first four columns show the effects of a Whirlpool and Haier acquisition of Maytag without product portfolio adjustments. The next four columns show the same comparison for merger-independent portfolio adjustments and the final four columns for endogenous portfolio adjustments. Each adjustment scenario includes results without marginal cost efficiencies and with offshoring efficiencies. 95% confidence intervals for effects without product portfolio adjustments and with merger-independent portfolio adjustments are computed using 100 residual bootstrap draws. Confidence sets for the expected effects with endogenous portfolio adjustments are based on 50 fixed cost draws for each potential product from a uniform distribution, where the domain are the confidence sets of brand-level fixed costs, and 500 demand and supply residual draws. Maytag includes all products marketed under the brands owned by Maytag pre-acquisition (i.e. Admiral, Amana, Magic Chef and Maytag).

Figure A.10: Consumer welfare change by household income

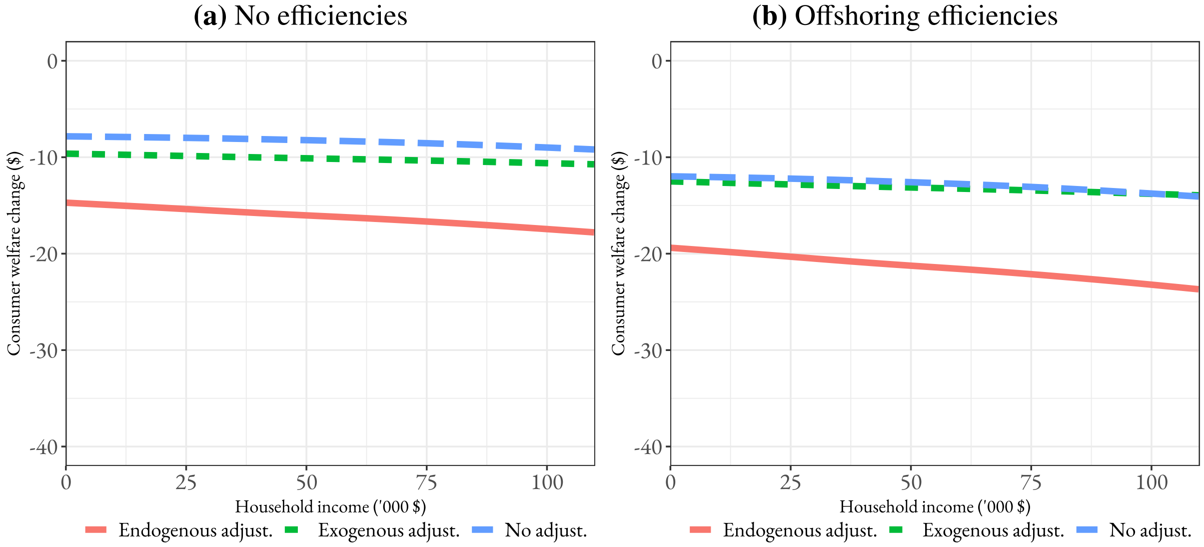

Notes: Both graphs show the absolute change in consumer welfare between a Maytag acquisition by Whirlpool and Haier according to household income. The consumer welfare changes are shown for three adjustment scenarios: No product portfolio adjustments, exogenous adjustments, and fully endogenous adjustments. Simulations are based on $1,000$ household draws per market. For expositional simplicity, the graphs only show households with an annual income of less than 110,000$, covering 80% of drawn households. In the left panel, no marginal cost efficiencies are credited. In the right panel, offshoring efficiencies are credited to Maytag products.

A.5.3 Employment effect decomposition

Figure A.11: Decomposition of employment effects

(a) No efficiencies **(b)** Offshoring efficiencies

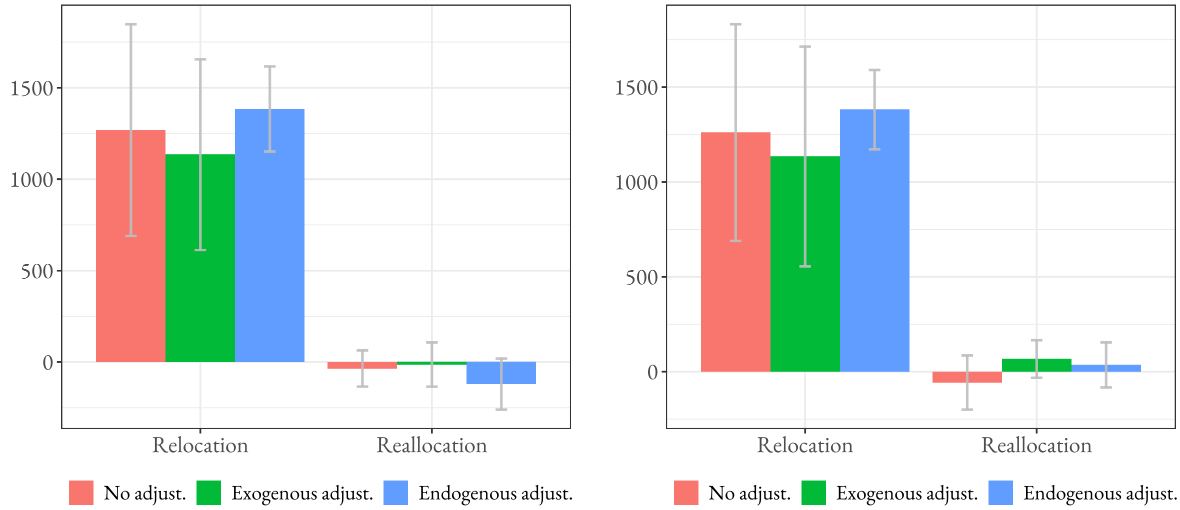

Notes: The histograms decompose the employment effects in Table 1.10 into a relocation and a reallocation effect. The former is related to differences in U.S. employment between Maytag acquisitions by Whirlpool and Haier because of different plant relocation plants. The latter is related to differences in U.S. employment between the acquisitions due to different reallocations of market shares between the merging parties, competitors, and the outside good.

A.5.4 Estimating consumer welfare and employment effects for other household appliances

In this subsection, I outline a very rough estimation of the consumer welfare and employment effects of alternative Maytag acquisitions by Whirlpool vs. Haier for other household appliances. This exercise is based on very strong assumptions and rough approximations. The estimated magnitudes should therefore be interpreted with caution. The aim of this exercise is merely to provide an idea about the direction in which the offsetting job value would shift, had I incorporated all affected appliance categories in the estimation.

The first difficulty is to estimate the consumer welfare effects of the hypothetical acquisitions in other product categories. In a first step, I identify which product categories to consider.

Table A.10 reports volume shares by manufacturer pre-merger for all appliance categories where Maytag was active and overlapped with either Haier or Whirlpool.[9] The data is based on a consensus survey among market participants conducted by the industry journal *Appliance*. Since products are attributed to who manufactures a product, not who markets and sells it, there are

[9] Maytag also sold floor care products, however, neither Haier, nor Whirlpool were active in that market.

Table A.10: Volume share by manufacturer from *Appliance* in 2005 (%)

	Clothes dryers	Clothes washers	Dishwashers	Ranges	Refrigerators
Whirlpool	56	51	34	19	25
Maytag	20	19	17	13	11
General Electric	13	17	25	45	29
Electrolux	9	9	18	16	25
Bosch			3		
Haier					2
Other	2	4	3	4	8
HHI	3,769	3,348	2,412	2,762	2,281
ΔHHI	2,191	1,938	1,156	467	550

Notes: The data is based on a consensus survey of market participants by the industry journal *Appliance*. Appliances are attributed to manufacturers, not brand owners, leading to discrepancies as compared to the *TraQline* data. For example, all clothes washers marketed by Sears are counted to Whirlpool, since they are manufactured by Whirlpool in 2005. Refrigerators only includes standard refrigerators and omits compact and built-in under-the-counter refrigerators. These constitute a different market.

some important discrepancies. In particular, products sold by Sears are manufactured by competitors. For example, all clothes washers marketed by Sears in 2005 were manufactured by Whirlpool.

As expected, market concentration and the overlap between Whirlpool and Maytag is largest for laundry products. The increase in the HHI after an acquisition by Whirlpool is around 2,000 for laundry products, 1,000 for dishwashers, and 500 for ranges and refrigerators. Haier and Maytag only overlap in the market for refrigerators, however, also there the overlap is unlikely to lead to any price effects.

To get a rough approximation of the loss in consumer welfare, I now make several strong assumptions. Since Nocke and Whinston (2020) show that emphasizing the change in HHI is more important for merger screening than the post-merger HHI, I focus on the increase in the HHI in what follows.

First, I assume that the percentage change in consumer welfare is proportional to the change in the HHI and that this relationship is the same across product categories. With endogenous portfolio adjustments and without efficiencies, the 95 percent confidence set of consumer welfare losses after a Whirlpool compared to a Haier acquisition for clothes washers ranges from 5.9 to 9.5 percent. Based on this, I assume that consumer welfare decreases between 6.7 and 10.4 percent for clothes dryers, 3.5 and 5.5 percent for dishwashers, 1.4 and 2.2 percent for ranges, and 1.7 and 2.6 percent for refrigerators.

Second, we need an estimate of pre-merger consumer welfare for each appliance category. I assume that pre-merger consumer welfare is proportional to the number of appliances of a particular appliance type sold in 2005. Based on the estimates from *Appliance*, there were 8.2 million clothes dryers, 9.2 million clothes washers, 7.4 million dishwashers, 10.0 million ranges, and 11.1 million standard sized refrigerators sold in 2005.

Taken together, I find a total decrease in consumer welfare of a Maytag acquisition by Whirlpool compared to Haier of between $600 million (clothes washers: 194 million $; clothes dryers: $195 million; dishwashers: $93 million; ranges: $51 million; refrigerators: $67 million) and $936 million (clothes washers: $304 million; clothes dryers: $303 million; dishwashers: $146 million; ranges: $79 million; refrigerators: $104 million).

The second difficulty is to estimate the employment effects of the hypothetical acquisitions in other product categories. As seen for the analysis of clothes washers, employment effects stem from three sources: the relocation of Maytag production, the reallocation of inside good market shares between competitors, and changes in the overall share of the inside good (i.e. the market size). For simplicity, I only estimate the difference in the relocation effect on employment for other appliance products. Since the overlap for ranges and refrigerators is low, changes in employment due to the reallocation of market shares, as well as changes in the market size are going to be low. For clothes dryers and dishwashers, omitting these effects leads to an overestimate of the difference in employment effects between the acquisitions, since the reallocation to competitors producing abroad and the decrease in market size only occur after an acquisition by Whirlpool. The decomposition of employment effects for clothes washers showed, however, that these latter two effects are significantly smaller in magnitude compared to the relocation effect.

According to Maytag (2005), the company had 16,900 employees in its home appliance division, of which 85 percent worked in the United States. According to news reports, Whirlpool cut 4,500 U.S. jobs at Maytag plants after its acquisition and created 1,500 new U.S. jobs at existing Whirlpool plants. Haier is assumed to offshore all Maytag jobs after an acquisition. Thus, Haier would reduce U.S. appliance manufacturing by 11,365 additional jobs compared to an acquisition by Whirlpool.

The number of additional U.S. jobs cut at Maytag because of production relocation after a Haier acquisition in home appliance categories that are not clothes washers are thus the sum of 11,365 and the model prediction of Maytag clothes washer jobs relocated by Whirlpool post-merger. The 95 percent confidence set for the latter is between 495 and 829. Overall, the number of additional Maytag jobs relocated by Haier in other product categories is between 11,860 and 12,194.

Taking the estimated product market and employment effects together, this results in an offsetting job value of between \$49,139 and \$78,907 for all other household appliances.

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

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

The expected profit of a firm i in equilibrium is

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

The expected price is

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

The expected minimum price is

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

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

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 partial 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 partial 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 begin by studying how the marginal effect of the share of *shoppers* changes with the ex ante share of *shoppers* for the minimum price. From the proof of Lemma 2.1, we know that the partial derivative of the minimum price with respect to the share of *shoppers* ϕ is strictly negative. For $0 < \phi < 1$, a second partial derivative of the minimum price with respect to the share of *shoppers* ϕ is strictly positive:

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

This means that the marginal effect of the share of *shoppers* on the minimum price decreases in magnitude with the ex ante share of *shoppers*.

Next, we study how the marginal effect of the share of *shoppers* varies with the ex ante share of *shoppers* for the equilibrium density function. We start by taking a second partial derivative of the equilibrium density function with respect to the share of *shoppers* ϕ:

$$\frac{\partial^2 F(p)}{\partial \phi^2} = k\left[\frac{N-2}{(N-1)(1-\phi)} - 2\right], \tag{B.11}$$

where

$$k = \frac{1}{N(N-1)\phi^3}\frac{\upsilon - p}{p}\left[\frac{\upsilon - p}{Np}\frac{1-\phi}{\phi}\right]^{\frac{1}{N-1}-1} > 0\,.$$

In Equation B.11, the component k is strictly positive since in equilibrium firms draw prices from the interval $[\underline{p}, p_r]$ and the share *shoppers* is such that $0 < \phi < 1$. From the proof of Lemma 2.1, we also know that the partial derivative of the equilibrium density function with respect to the share of *shoppers* ϕ is strictly positive.

Then, whether the marginal effect of ϕ on prices in equilibrium is stronger for ex ante higher share of *shoppers* depends on the sign of the expression in parenthesis in Equation B.11:

$$\frac{\partial^2 F(p)}{\partial \phi^2} \gtrless 0 \quad \text{iff} \quad \frac{N-2}{(N-1)(1-\phi)} - 2 \gtrless 0\,.$$

After rearranging and simplifying the above expression, the condition becomes

$$\frac{\partial^2 F(p)}{\partial \phi^2} \gtrless 0 \quad \text{iff} \quad \phi \gtrless \frac{N}{2(N-1)}\,.$$

Thus, when $0 < \phi < 1 \quad \frac{\partial^2 F(p)}{\partial \phi^2} < 0 \quad$ if $\phi \leq \frac{N}{2(N-1)}$ and $\frac{\partial F(p)}{\partial \phi} > 0$.

In other words, as long as the share of *shoppers* does not exceed $\frac{N}{2(N-1)}$, the marginal effect of the share of *shoppers* ϕ on the equilibrium density function will be stronger for consumer groups that are on average less informed ex ante.

□

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

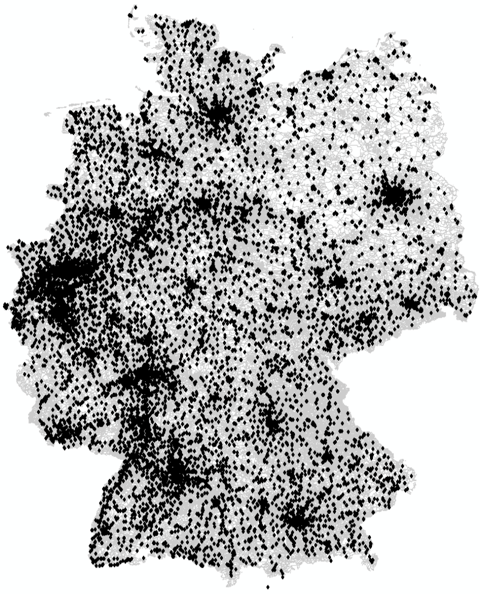

Note: The Figure shows the geographic distribution of fuel stations in Germany.

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 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

[2]This is consistent with findings by Haucap, Heimeshoff, Kehder, Odenkirchen, and Thorwarth (2017) for Germany and Gautier and Saout (2015) for France.

Table B.1: Share of stations in percent by brand

	Pre-MTU	Post-MTU
Aral	20.1	18.1
Shell	14.2	14.2
Esso	5.7	5.4
Total	7.0	4.7
Jet	5.0	4.6
Orlen	4.7	4.2
Agip	2.0	3.1
Hem	3.0	2.8
OMV	2.6	2.3
Non-integrated	35.8	40.6

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.

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] 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.

[3]The daily fuelling patterns are described in detail in Figure B.6 in Appendix B.2.1.

[4]This is consistent with pricing patterns in the data described in Figure B.7 in Appendix B.2.1.

[5]There are numerous newspaper articles on intertemporal price dispersion during our observation period, which suggest that consumers are aware of these patterns.

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

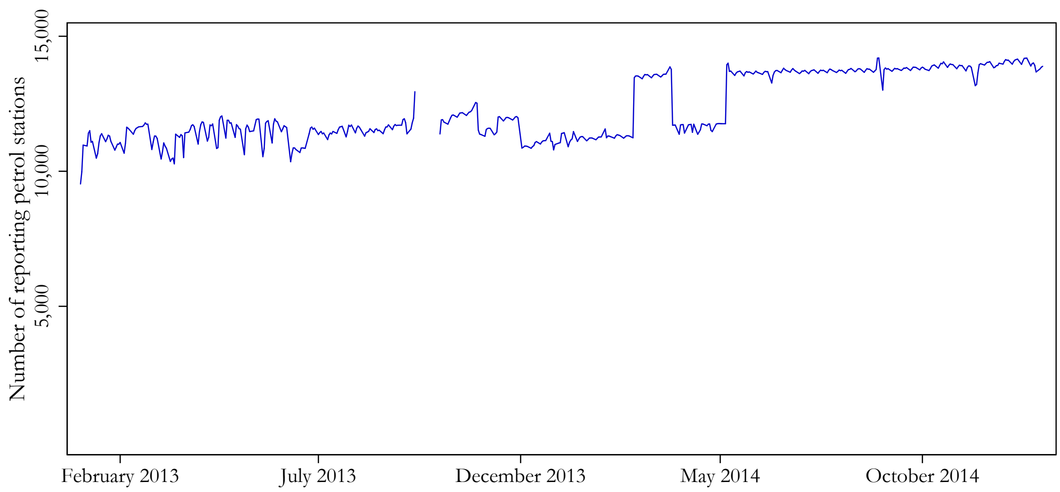

Notes: The Figure shows the average daily number of fuel stations with a positive price report at 5 pm in Germany in our sample.

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.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.

[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: Number of daily price changes

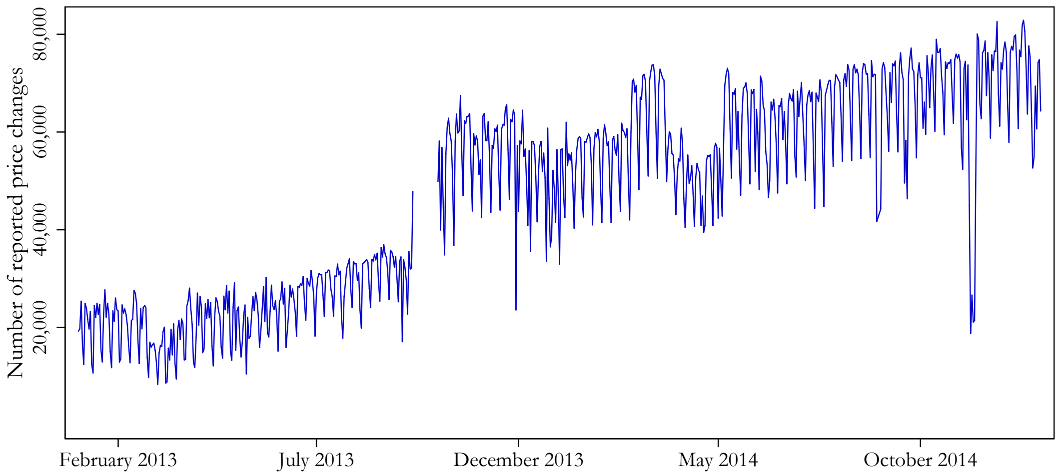

Notes: The Figure shows the average daily number of price changes in Germany in our data. In the pre-MTU period consecutive reports of the same price are not considered a price change.

Figure B.4: Notification patterns over the day

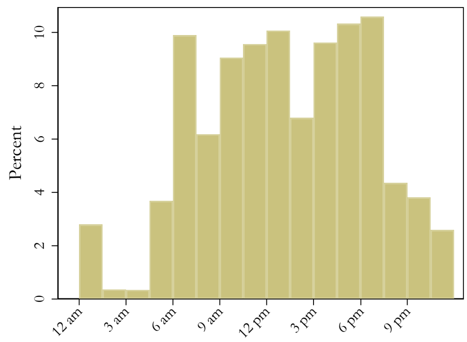

(a) Pre-MTU

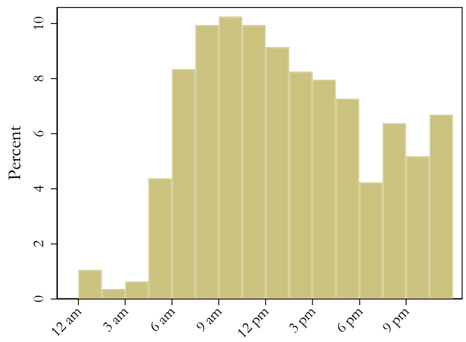

(b) Post-MTU

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.6 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.

Figure B.6: Daily fuelling patterns

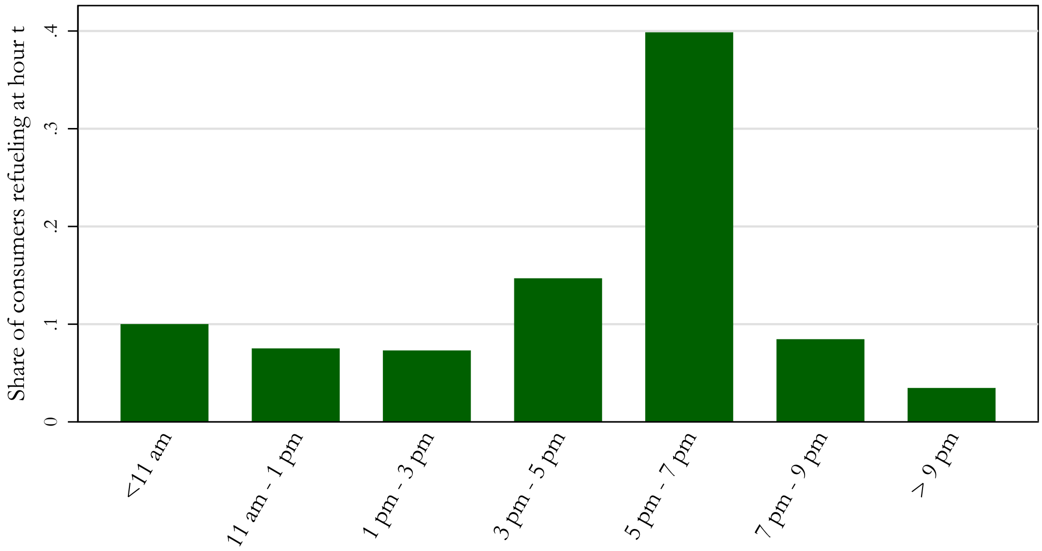

Notes: The Figure shows the average fuelling patterns by German motorists over the day. Data is based on a representative survey among drivers commissioned by the German Federal Ministry of Economic Affairs.

The fuelling patterns are also consistent with price patterns reported in Figure B.7. Whereas *E5* 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.7: Daily price patterns

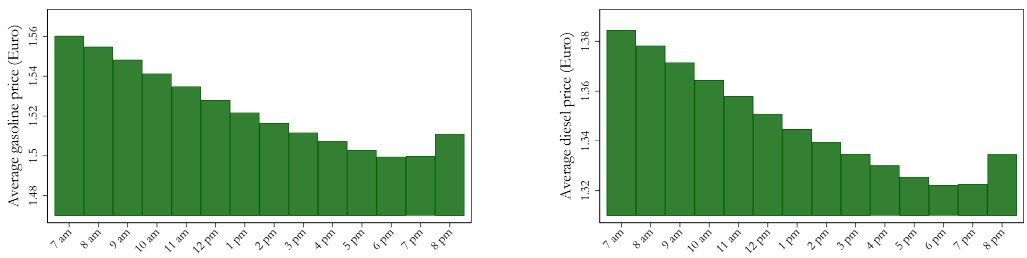

(a) Gasoline **(b)** Diesel

Notes: Panel (a) shows the average *E5* price for every hour between 7 am and 20 pm in Germany between 2013 and 2014. Panel (b) shows the average diesel price for every hour between 7 am and 20 pm in Germany between 2013 and 2014.

B.2.2 Distribution of fuel stations by SDID unit weights in France

Figure B.9 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 *E5* 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.9: Geographic distribution of fuel stations by SDID unit weights, France

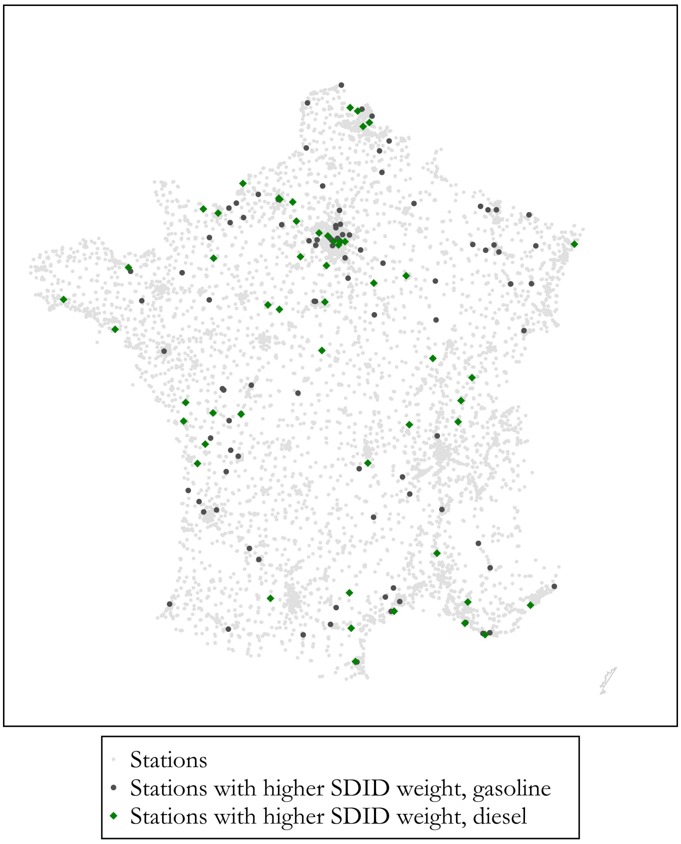

Notes: The Figure shows the geographic distribution of fuel stations in France. Stations that receive a disproportionally high unit weight in the SDID estimation are highlighted.

B.3 Appendix to Section 2.6: Results

In this Section we provide further empirical evidence on the average effect of the MTU on *E5* 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 MTU_{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 the mandatory price disclosure led to the decline by 2.93% to 3.01% for *E5* price and 2.33% to 2.75% for diesel price. The effects are economically and statistically significant, and, similarly to the results estimated via SDID, remain larger for *E5*.

Table B.2: Effect of MPD on the logarithm of gross prices

	Gasoline	Diesel	Gasoline	Diesel
	(1)	(2)	(3)	(4)
MPD	-0.029***	-0.023***	-0.030***	-0.028***
	(0.0002)	(0.0002)	(0.001)	(0.001)
Date FE	*Yes*	*Yes*	*Yes*	*Yes*
Station FE	*Yes*	*Yes*	*Yes*	*Yes*
Observations	4,235,999	4,830,137	374,344	402,985
Adjusted R^2	0.825	0.797	0.806	0.730

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.10 illustrates the identification strategy for the Donut-SDID analysis graphically. 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. Fuel stations more than 100 kilometers away from the border could be subject to very different market conditions and are thus also not considered. Each point in Figure B.10 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 *E5* price.

Figure B.10: Fuel stations 20 to 100 kilometers from the Franco-German border

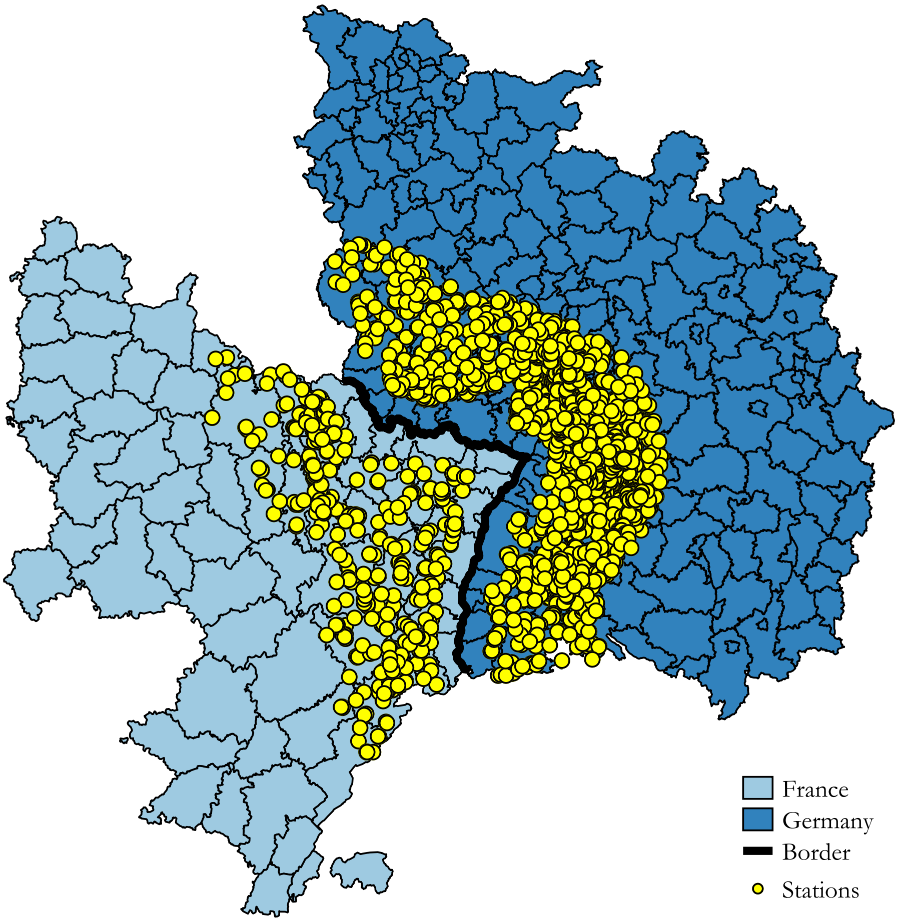

Notes: The thick, solid line represents the Franco-German border. Each point on the right of the border represents a fuel station in Germany, which is 20 to 100 kilometers away from the border. Each point on the left side of the border represents a fuel station in France, which is 20 to 100 kilometers 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	-0.032***	-0.023***	-0.028***	-0.020***	-0.026***	-0.020***
	(0.002)	(0.001)	(0.002)	(0.002)	(0.002)	(0.001)
95% CI	[-0.036, -0.029]	[-0.026, -0.020]	[-0.032, -0.024]	[-0.022, -0.017]	[-0.030, -0.022]	[-0.022, -0.017]
Week FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Station FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Observations	10,094	12,005	22,001	25,725	39,690	44,590

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 *E5* 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 2.62% for *E5* and 1.48% for diesel. When the sample is restricted to the Donut-SDID, the corresponding estimates indicate a decline of 1.51% for *E5* and 1.4% 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	-0.026***	-0.015***	-0.015***	-0.014***
	(0.008)	(0.007)	(0.020)	(0.005)
95% Confidence interval	[-0.043, -0.010]	[-0.029, -0.001]	[-0.055, 0.025]	[-0.025, -0.003]
Germany × crude oil price	*Yes*	*Yes*	*Yes*	*Yes*
Week FE	*Yes*	*Yes*	*Yes*	*Yes*
Station FE	*Yes*	*Yes*	*Yes*	*Yes*
Observations	666,106	783,951	52,969	58,408

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 3.25 and 1.44 Eurocent for *E5* 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 3.37 Eurocent for *E5* and 2.29 Eurocent for diesel. The effect of the MTU introduction is statistically and economically significant, and is larger for *E5* fuel.

Table B.5: Effect of MPD on retail margins

	Gasoline	Diesel	Gasoline	Diesel
	(1)	(2)	(3)	(4)
MPD	-3.245***	-1.441***	-3.366***	-2.288***
	(0.283)	(0.179)	(0.198)	(0.127)
95% Confidence interval	[-3.800, -2.690]	[-1.791, -1.091]	[-3.754, -2.978]	[-2.537, -2.039]
Week FE	*Yes*	*Yes*	*Yes*	*Yes*
Station FE	*Yes*	*Yes*	*Yes*	*Yes*
Observations	666,106	783,951	52,969	58,408
Mean retail margin	8.39	10.82	10.88	11.23

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 *E5* 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 0.07 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)	(0.0004)	(0.0003)	(0.0004)	(0.0004)
Date FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Station FE	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Observations	2,233,878	1,355,754	1,109,831	2,252,605	1,367,029	1,119,258
Share local monopolists	64.3%	42.3%	29.4%	64.3%	42.3%	29.4%
Adjusted R^2	0.814	0.817	0.817	0.665	0.672	0.672

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 *E5* 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 0.13 to 0.15 Eurocent 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.

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***
	(0.006)	(0.006)	(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	1258	1258	1258	1258
Adjusted R^2	0.868	0.836	0.879	0.859

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

B.3.6 Difference-in-differences analysis: European countries as a control

To test the validity of France as a counterfactual, we additionally estimate the effect of the MTU introduction on fuel prices in Germany using other 26 European countries as a control group.[7] To do so, we use information on country-level weekly average net *E5* 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 log net *E5* and diesel price estimated via difference-in-differences and using other European countries as a control. The estimation is based on data between 15 April 2013 and 31 March 2014, and we control for date and country fixed effects in all columns. In Columns (3) and (4), we additionally include the crude oil price at the port of Rotterdam interacted with country indicators into estimation, which allows for differential pass-through of oil cost shocks across countries.

Table B.7 shows that the MTU introduction led to a decline of 3% to 3.3% in the net price of *E5* and 1.5% to 1.8% in the net price of diesel, when the effects are estimated using the other European countries as a control. The ranking of the effects with respect to the fuel types and their magnitude remain robust to using this alternative control group.

[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*

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

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

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

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 occur, 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 ful-

[1] For a more detailed proof, see Varian (1980).

filled, 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 $[\underline{p}, p_r]$

$$F(p_i) = 1 - (\frac{p_r - p_i}{p_i - c(1+\tau)}\frac{1-\phi}{N\phi})^{\frac{1}{N-1}}\,. \tag{C.9}$$

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

$$E[\pi_i] = E[\pi(p_r)] = (\frac{p_r}{1+\tau} - c)\frac{1-\phi}{N}M\,. \tag{C.10}$$

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

$$E[p] = \int_{\underline{p}}^{p_r} pf(p)dp = p_r - \int_{\underline{p}}^{p_r} F(p)dp\,, \tag{C.11}$$

after integrating by parts. We can then insert the equilibrium price distribution and simplify the expression, which yields

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

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

$$E[p_{min}] = \int_{\underline{p}}^{p_r} pf_{min}(p)dp = \int_{\underline{p}}^{p_r} p\frac{p_r - p}{p - c(1+\tau)}\frac{1-\phi}{N\phi} f(p)dp\,. \tag{C.14}$$

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 $\hat{\underline{p}} > \underline{p}$ 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.

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

Figure C.1: Marginal cost pass-through to the equilibrium pricing strategy

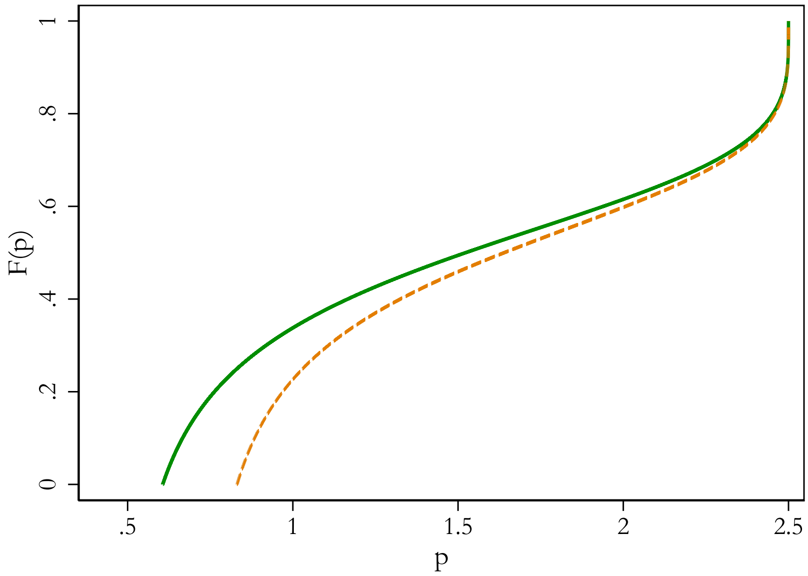

Note: The Figure shows simulation results of how the distribution from which sellers draw prices in the symmetric Nash equilibrium changes if marginal costs increase from c to $\hat{c}$. Parameter values: $\upsilon = 2.5$, $s = 0.75$, $\tau = 0.2$, $c = 0.4$, and $\hat{c} = 0.6$.

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.

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. $\rho_{c,\underline{p}} \to 1 + \tau$.*

[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 $\hat{c} - c$.

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 $E[p]$

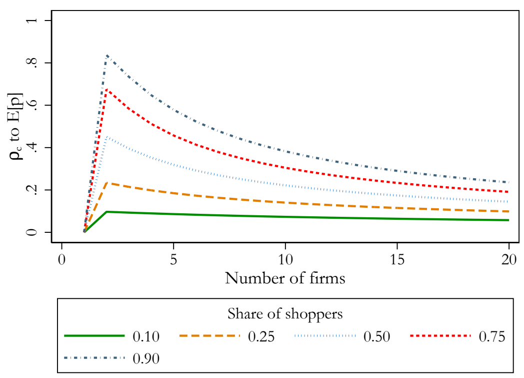

Parameter values: $\upsilon = 2.5$, $\tau = 0.2$, $c = 0.4$ and $\hat{c} = 0.44$.

Figure C.3: Pass-through of c to $E[p_{min}]$

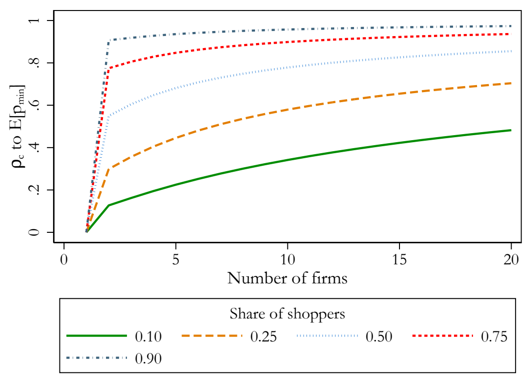

Parameter values: $\upsilon = 2.5$, $\tau = 0.2$, $c = 0.4$ and $\hat{c} = 0.44$.

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 decreases.

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

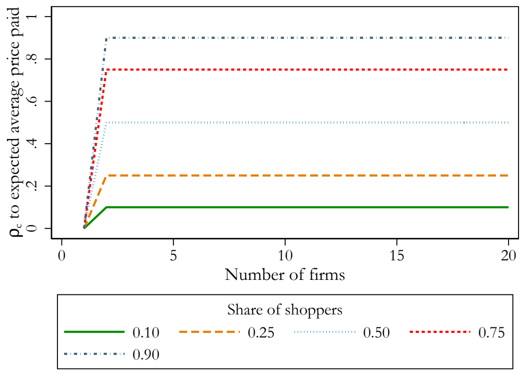

Parameter values: $\upsilon = 2.5$, $\tau = 0.2$, $c = 0.4$ and $\hat{c} = 0.44$.

Figure C.5: ρ_c to $E[\phi p_{min} + (1-\phi)p]$, p_r endogenous

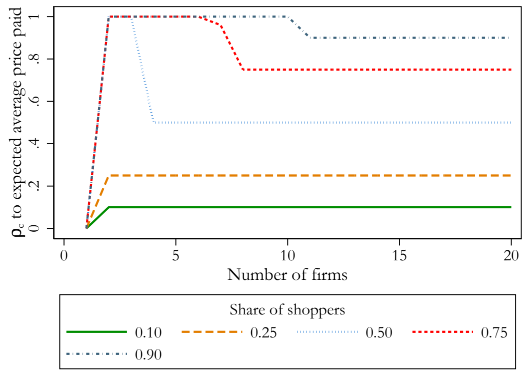

Parameter values: $\upsilon = 2.5$, $s = 0.75$, $\tau = 0.2$, $c = 0.4$ and $\hat{c} = 0.44$.

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

$$\frac{\partial \underline{p}}{\partial \tau} = c(1 + \frac{1-\phi}{\phi N})^{-1} > 0\,.$$

Thus, with $0 < \phi < 1$, pass-through of τ to the minimum element of the support of the equilibrium pricing strategy is strictly positive.

Next, we assess the pass-through of the ad-valorem tax to $F(p)$ if $0 < \phi < 1$. Taking the first

[3] $\underline{p}$ is not defined for $\phi = 0$ or $\phi = 1$.

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

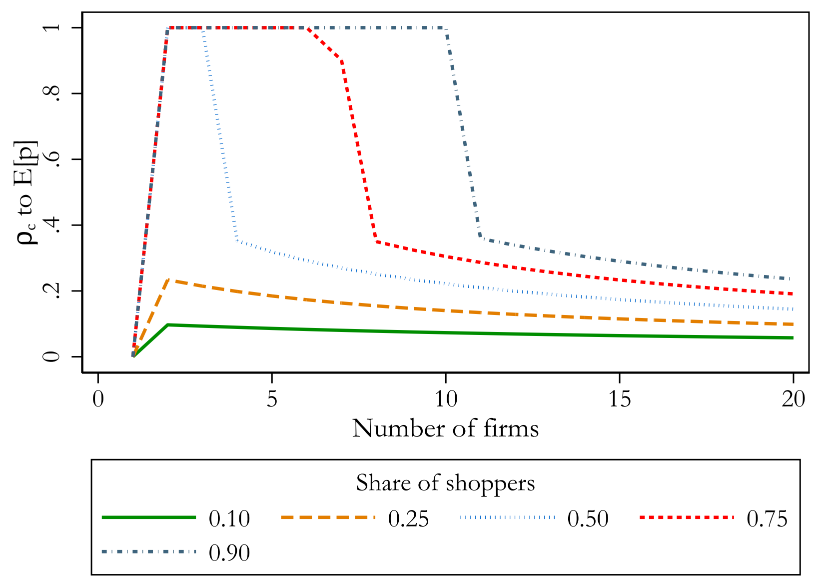

Note: The Figure shows simulation results of how the pass-through of marginal costs to the expected price varies with the share of shoppers and the number of active sellers. We fix the following parameter values for these simulations: $\upsilon = 2.5$, $s = 0.75$, $\tau = 0.2$, $c = 0.4$ and $\hat{c} = 0.44$.

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

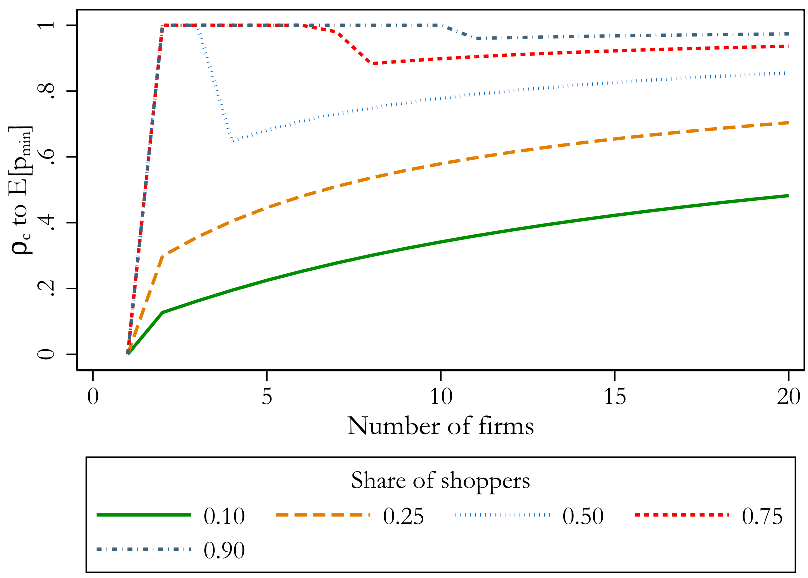

Note: The Figure shows simulation results of how the pass-through of marginal costs to the expected minimum price varies with the share of shoppers and the number of active sellers. We fix the following parameter values for these simulations: $\upsilon = 2.5$, $s = 0.75$, $\tau = 0.2$, $c = 0.4$ and $\hat{c} = 0.44$.

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

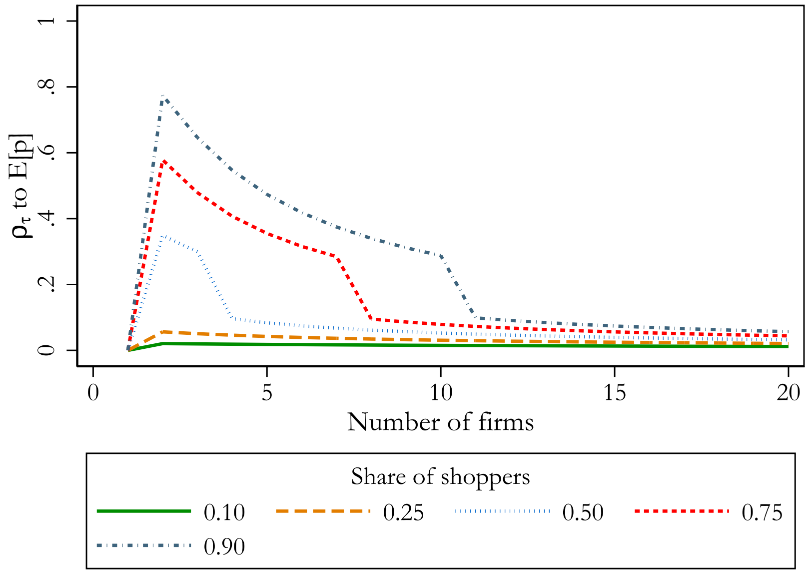

Note: The Figure shows simulation results of how the pass-through of an ad-valorem tax to the expected price varies with the share of shoppers and the number of active sellers. We fix the following parameter values for these simulations: $\upsilon = 2.5$, $s = 0.75$, $c = 0.4$, $\tau = 0.2$ and $\hat{\tau} = 0.22$.

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

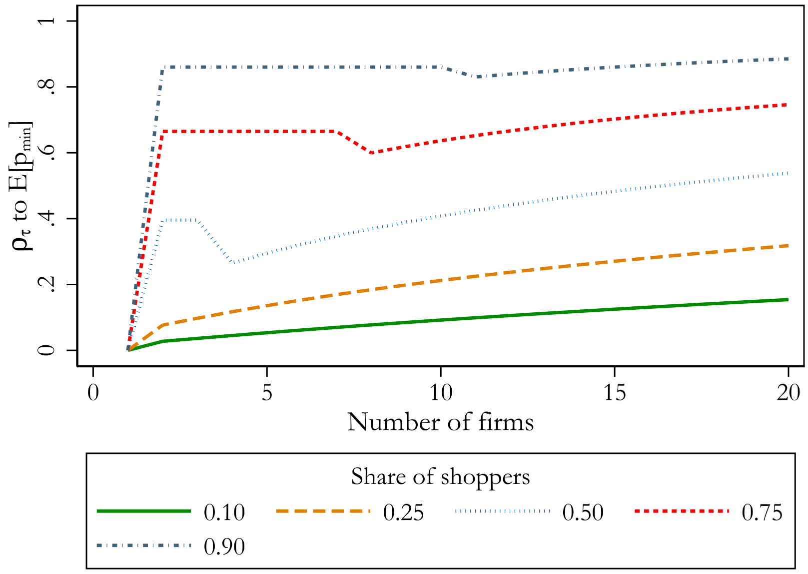

Note: The Figure shows simulation results of how the pass-through of an ad-valorem tax to the expected minimum price varies with the share of shoppers and the number of active sellers. We fix the following parameter values for these simulations: $\upsilon = 2.5$, $s = 0.75$, $c = 0.4$, $\tau = 0.2$ and $\hat{\tau} = 0.22$.

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)

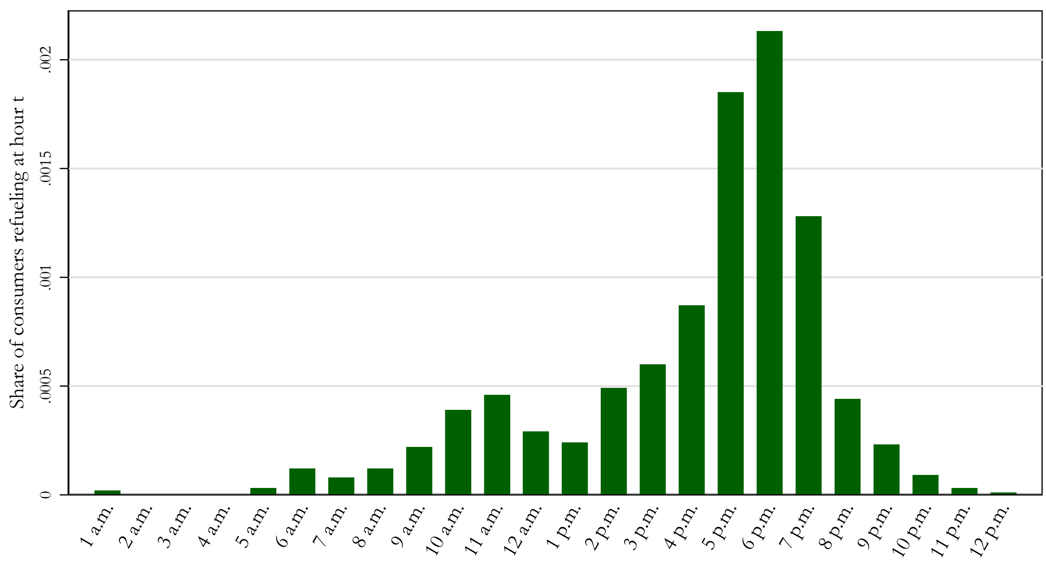

Notes: The Figure shows shares of drivers in Germany who fuel at a given hour of a day. Data is based on a representative survey of motorists in Germany, commissioned by the German Federal Ministry of Economic Affairs.

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)

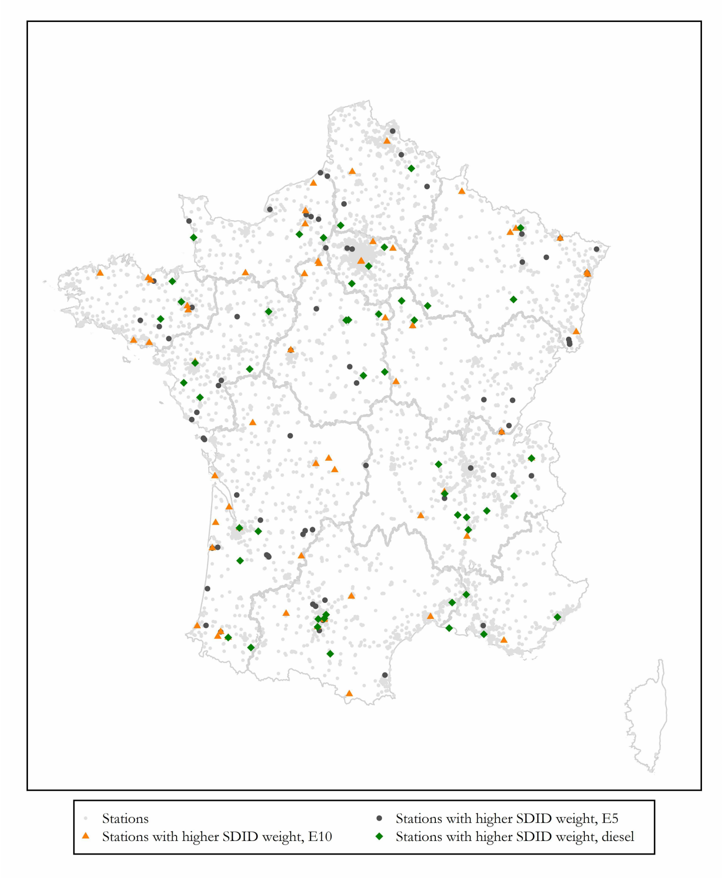

Notes: The Figure shows the geographic distribution of fuel stations in France for the analysis of the tax decrease. Stations with a disproportionally high unit weight in the SDID pass-through estimation for *E5*, *E10* or diesel are highlighted.

Figure C.12: France: distribution of fuel stations by SDID unit weights (tax increase)

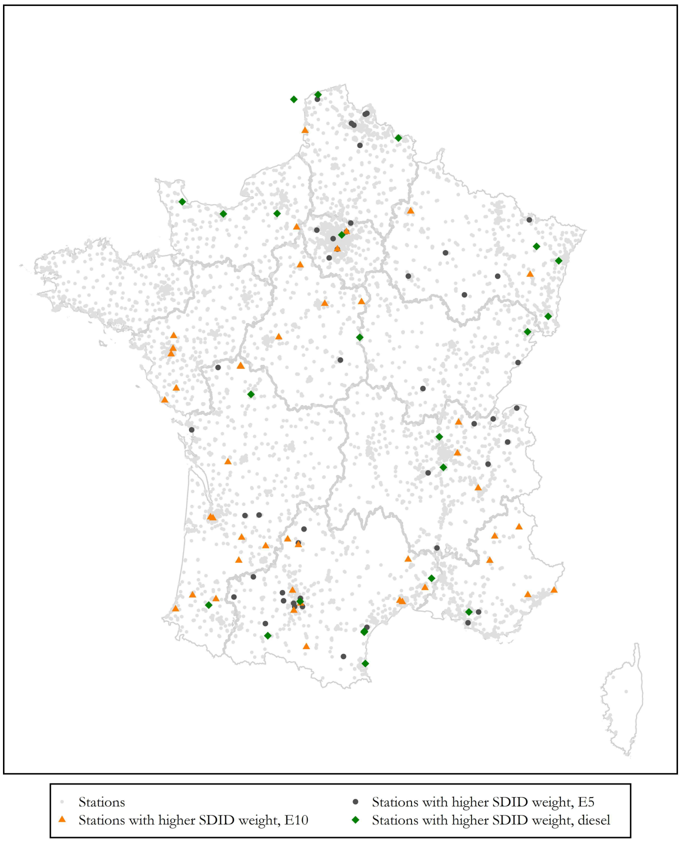

Notes: The Figure shows the geographic distribution of fuel stations in France for the analysis of the tax increase. Stations with a disproportionally high unit weight in the SDID pass-through estimation for *E5*, *E10* or diesel are highlighted.

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 × 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

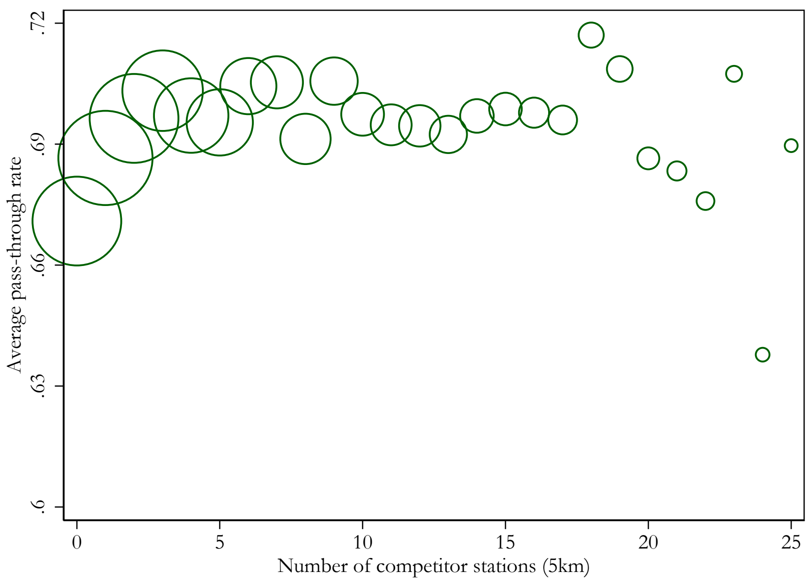

Notes: Each circle plots the average pass-through rate for a group of stations with a particular number of competitors within 5 km catchment area. The number of competitor stations is trimmed at the top percentile.

Figure C.14: Average pass-through by number of competitor stations, E10

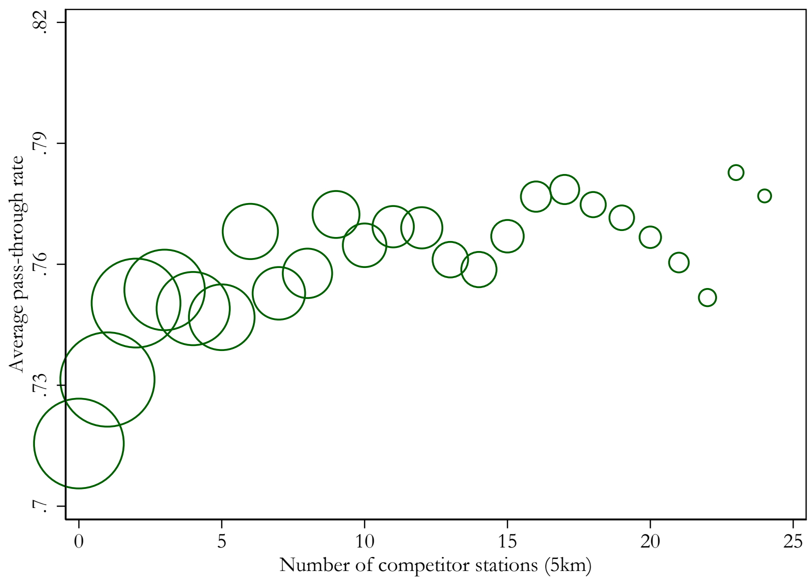

Notes: Each circle plots the average pass-through rate for a group of stations with a particular number of competitors within 5 km catchment area. The number of competitor stations is trimmed at the top percentile.

Figure C.15: Average pass-through by number of competitor stations, diesel

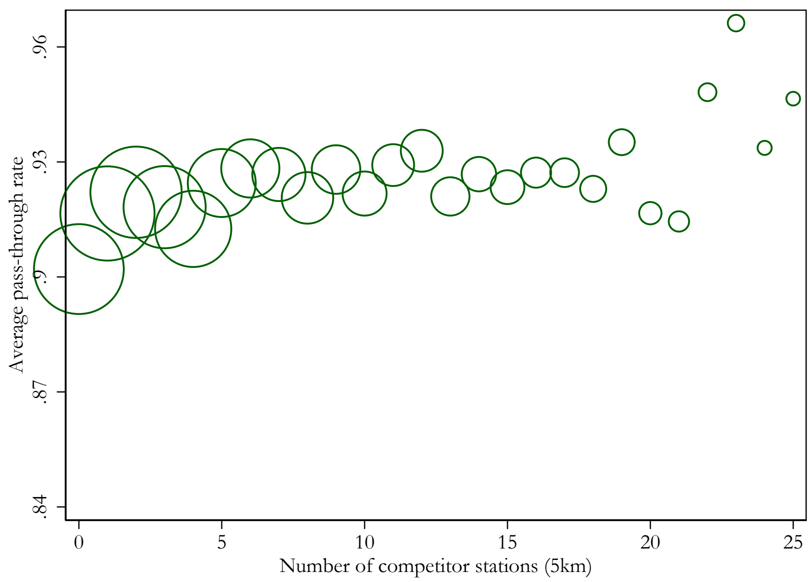

Notes: Each circle plots the average pass-through rate for a group of stations with a particular number of competitors within 5 km catchment area. The number of competitor stations is trimmed at the top percentile.

Bibliography

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.

Alfaro-Ureña, A., Manelici, I., and Vasquez, J. P. (2021). "The Effects of Multinationals on Workers: Evidence from Costa Rican Microdata." *mimeo*.

Andrews, D. W. K. (2000). "Inconsistency of the Bootstrap when a Parameter is on the Boundary of the Parameter Space." *Econometrica*, 68 (2): pp. 399–405.

Andrews, I., Stock, J., and Sun, L. (2019). "Weak Instruments in IV Regression: Theory and Practice." *Annual Review of Economics*, 11: pp. 727–753.

Appliance Portrait (2006). 29^{th} *Annual Portrait of the U.S. Appliance Industry*. Dana Chase.

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.

Ashenfelter, O. C., Hosken, D. S., and Weinberg, M. C. (2013). "The Price Effects of a Large Merger of Manufacturers: A Case Study of Maytag-Whirlpool." *American Economic Journal: Economic Policy*, 5 (1): pp. 239–61.

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. (2019). "Price Transparency, Media and Informative Advertising." *mimeo*.

Autor, D., Dorn, D., Katz, L. F., Patterson, C., and Van Reenen, J. (Feb. 2020). "The Fall of the Labor Share and the Rise of Superstar Firms." *The Quarterly Journal of Economics*, 135 (2): pp. 645–709.

Autor, D. H., Dorn, D., and Hanson, G. H. (2013). "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review*, 103 (6): pp. 2121–2168.

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.

Bai, L. and Stumpner, S. (2019). “Estimating US Consumer Gains from Chinese Imports.” *American Economic Review: Insights*, 1 (2): pp. 209–224.

Baker, J. B. and Salop, S. C. (2015). “Antitrust, Competition Policy, and Inequality.” *Georgetown Law Journal*, 104 (1-28).

Baker, J. B. and Shapiro, C. (2008a). “Detecting and Reversing the Decline in Horizontal Merger Enforcement.” *Antitrust*, 22 (3): pp. 29–36.

– (2008b). “Reinvigorating Horizontal Merger Enforcement.” In: ed. by R. Pitofsky. Oxford University Press.

Barboza, D. (2005). “From China, A New Bid for Maytag and Status.” *New York Times*. URL: `https://www.nytimes.com/2005/06/22/business/worldbusiness/from-china-a-new-bid-for-maytag-and-status.html`.

Bellemare, M. F. and Wichman, C. J. (2020). “Elasticities and the Inverse Hyperbolic Sine Transformation.” *Oxford Bulletin of Economics and Statistics*, 82 (1): pp. 50–61.

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.

Benkard, C. L., Yurukoglu, A., and Zhang, A. L. (2021). “Concentration in Product Markets.” *NBER Working Paper 28745*.

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.

Berry, S. (1994). “Estimating Discrete-Choice Models of Product Differentiation.” *The RAND Journal of Economics*, 25 (2): pp. 242–262.

Berry, S., Levinsohn, J., and Pakes, A. (1995). “Automobile Prices in Market Equilibrium.” *Econometrica*, 63 (4): pp. 841–890.

– (2004). “Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market.” *Journal of Political Economy*, 112 (1): pp. 68–105.

Berry, S. T. and Waldfogel, J. (2001). “Do Mergers Increase Product Variety? Evidence from Radio Broadcasting.” *The Quarterly Journal of Economics*, 116 (3): pp. 1009–1025.

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.

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.

Brunner, D., Heiss, F., Romahn, A., and Weiser, C. (2020). “Reliable Estimation of Random Coefficient Logit Demand Models.” *mimeo*.

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. (2019). “Learning to Coordinate: A Study in Retail Gasoline.” *American Economic Review*, 109 (2): pp. 591–619.

– (Forthcoming). “Startup Search Costs.” *American Economic Journal: Microeconomics*.

Byrne, D. P. and Roos, N. de (2017). “Consumer Search in Retail Gasoline Markets.” *The Journal of Industrial Economics*, 65 (1): pp. 183–193.

Cabral, L. M. (2003). “Horizontal mergers with free-entry: why cost efficiencies may be a weak defense and asset sales a poor remedy.” *International Journal of Industrial Organization*, 21 (5): pp. 607–623.

Caradonna, P., Miller, N. H., and Sheu, G. (2021). “Mergers, Entry, and Consumer Welfare.” *Georgetown McDonough School of Business Research Paper 3537135*.

Carbonnier, C. (2007). “Who pays sales taxes? Evidence from French VAT reforms, 1987–1999.” *Journal of Public Economics*, 91 (5): pp. 1219–1229.

Card, D., Cardoso, A. R., Heining, J., and Kline, P. (2018). “Firms and Labor Market Inequality: Evidence and Some Theory.” *Journal of Labor Economics*, 36 (1): pp. 13–17.

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.

Ciliberto, F., Murry, C., and Tamer, E. (2021). "Market Structure and Competition in Airline Markets." *Journal of Political Economy*, 129 (11): pp. 2995–3038.

Conlon, C. and Gortmaker, J. (2020). "Best practices for differentiated products demand estimation with PyBLP." *The RAND Journal of Economics*, 51 (4): pp. 1108–1161.

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.

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.

– (Forthcoming). "Managing Households' Expectations with Unconventional Policies." *Review of Financial Studies*.

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.

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.

Department of Justice (2006). *Statement on the Closing of its Investigation of Whirlpool's Acquisition of Maytag*. Washington, DC: Author. URL: https://www.justice.gov/archive/atr/public/press_releases/2006/215326.htm.

Diamond, P. A. (1971). "A Model of Price Adjustment." *Journal of Economic Theory*, 3: pp. 156–168.

Draganska, M., Mazzeo, M. J., and Seim, K. (2009). "Beyond plain vanilla: Modeling joint product assortment and pricing decisions." *Quantitative Marketing and Economics*, 7 (2): pp. 105–146.

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.

Eizenberg, A. (2014). "Upstream Innovation and Product Variety in the U.S. Home PC Market." *The Review of Economic Studies*, 81 (3): pp. 1003–1045.

Electrolux (2007). *Annual Report 2006*. URL: `https://www.electroluxgroup.com/en/annual-reports-27641/?decade=2000-2009`.

Fabra, N. and Reguant, M. (2014). "Pass-Through of Emissions Costs in Electricity Markets." *American Economic Review*, 104 (9): pp. 2872–2899.

Fan, Y. (2013). "Ownership Consolidation and Product Characteristics: A Study of the US Daily Newspaper Market." *American Economic Review*, 103 (5): pp. 1598–1628.

Fan, Y. and Yang, C. (2020). "Competition, Product Proliferation, and Welfare: A Study of the US Smartphone Market." *American Economic Journal: Microeconomics*, 12 (2): pp. 99–134.

– (2021). "Estimating Discrete Games with Many Firms and Many Decisions: An Application to Merger and Product Variety." *mimeo*.

Feenstra, R. C., Inklaar, R., and Timmer, M. P. (2015). "The Next Generation of the Penn World Table." *American Economic Review*, 105 (10): pp. 3150–82.

Flaaen, A., Hortaçsu, A., and Tintelnot, F. (2020). "The Production Relocation and Price Effects of US Trade Policy: The Case of Washing Machines." *American Economic Review*, 110 (7): pp. 2103–27.

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.

Garrido, F. (2020). "Mergers Between Multi-Product Firms With Endogenous Variety: Theory and an Application to the RTE-Cereal Industry." *mimeo*.

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. (2019). "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*.

Goldberg, P. K. and Verboven, F. (2001). "The Evolution of Price Dispersion in the European Car Market." *The Review of Economic Studies*, 68 (4): pp. 811–848.

Goodman, P. and White, B. (2005). "Haier Withdraws Maytag Bid." *Washington Post*. URL: `https://www.washingtonpost.com/archive/business/2005/07/20/haier-withdraws-maytag-bid/f05586dd-097d-41b8-8cee-99c3bd17a72b/`.

Grieco, P. L., Murry, C., and Yurukoglu, A. (2021). "The Evolution of Market Power in the Us Auto Industry." *NBER Working Paper 29013*.

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.

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*.

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.

Houde, J.-F. (2012). "Spatial Differentiation and Vertical Mergers in Retail Markets for Gasoline." *American Economic Review*, 102 (5): pp. 2147–2182.

Houde, S. (2018). "How consumers respond to product certification and the value of energy information." *The RAND Journal of Economics*, 49 (2): pp. 453–477.

Hufbauer, G. and Lowry, S. (2012). "US Tire Tariffs: Saving Few Jobs at High Cost." *Peterson Institute for International Economics*.

Jacob, M., Müller, M., and Wulff, T. (2021). "Do Consumers Pay the Corporate Tax?" *TRR 266 Accounting for Transparency Working Paper Series No. 15*.

Jacobson, L. S., LaLonde, R. J., and Sullivan, D. G. (1993). “Earnings Losses of Displaced Workers.” *American Economic Review*, 83 (4): pp. 685–709.

Jaravel, X. and Sager, E. (2020). “What are the Price Effects of Trade? Evidence from the U.S. and Implications for Quantitative Trade Models.” *mimeo*.

Jeziorski, P. (2015). “Empirical Model of Dynamic Merger Enforcement–Choosing Ownership Caps in US Radio.” *mimeo*.

Johnson, R. N. (2002). “Search Costs, Lags and Prices at the Pump.” *Review of Industrial Organization*, 20 (1): pp. 33–50.

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.

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.

Lacetera, N. and Sydnor, J. (2015). “Would You Buy a Honda Made in the United States? The Impact of Production Location on Manufacturing Quality.” *The Review of Economics and Statistics*, 97 (4): pp. 855–876.

Lee, R. S. and Pakes, A. (2009). “Multiple Equilibria and Selection by Learning in an Applied Setting.” *Economics Letters*, 104 (1): pp. 13–16.

Lemus, J. and Luco, F. (2021). “Price Leadership and Uncertainty about Future Costs.” *The Journal of Industrial Economics*, 69 (2): pp. 305–337.

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.

Li, S., Mazur, J., Park, Y., Roberts, J., Sweeting, A., and Zhang, J. (Forthcoming). "Repositioning and Market Power After Airline Mergers." *The RAND Journal of Economics*.

Loecker, J. D. and Eeckhout, J. (2021). "Global Market Power." *mimeo*.

Luco, F. (2019). "Who Benefits from Information Disclosure? The Case of Retail Gasoline." *American Economic Journal: Microeconomics*, 11 (2): pp. 277–305.

Marinescu, I. and Hovenkamp, H. (2019). "Anticompetitive Mergers in Labor Markets." *Indiana Law Journal*, 94 (3): p. 1031.

Martin, S. (2020). "Market Transparency and Consumer Search - Evidence from the German Retail Gasoline Market." *Dice Discussion Paper, No. 350*.

Maytag (2005). *2004 Annual Report*. Newton, IA: Author.

Mazzeo, M. J. (2002). "Product Choice and Oligopoly Market Structure." *The RAND Journal of Economics*, 33 (2): pp. 221–242.

McCabe, L. (2016). "Should You Get a Front-Load or Top-Load Washing Machine?" *New York Times*. URL: https://www.nytimes.com/wirecutter/blog/should-you-get-a-front-or-top-load-washing-machine/.

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.

Miller, N. H. and Weinberg, M. C. (2017). "Understanding the Price Effects of the MillerCoors Joint Venture." *Econometrica*, 85 (6): pp. 1763–1791.

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. and Winter, C. (2020). "Price Transparency Against Market Power." *mimeo*.

Morrow, W. R. and Skerlos, S. J. (2011). "Fixed-Point Approaches to Computing Bertrand-Nash Equilibrium Prices Under Mixed-Logit Demand." *Operations Research*, 59 (2): pp. 328–345.

Nakamura, E. and Zerom, D. (2010). "Accounting for Incomplete Pass-Through." *The Review of Economic Studies*, 77 (3): pp. 1192–1230.

Nevo, A. (2001). "Measuring Market Power in the Ready-to-Eat Cereal Industry." *Econometrica*, 69 (2): pp. 307–342.

Nocke, V. and Whinston, M. (2020). "Concentration Screens for Horizontal Mergers." *NBER Working Paper 27533*.

OECD (2016). "Public Interest Considerations in Merger Control." *DAF/COMP/WP3 /WD(2016)30*.

Olea, J. L. M. and Pflueger, C. (2013). "A Robust Test for Weak Instruments." *Journal of Business & Economic Statistics*, 31 (3): pp. 358–369.

Olmstead-Rumsey, J. (2020). "Market Concentration and the Productivity Slowdown." *mimeo*.

Owen, A. B. (2017). *A randomized Halton algorithm in R*.

Pakes, A., Porter, J., Ho, K., and Ishii, J. (2015). "Moment Inequalities and Their Application." *Econometrica*, 83 (1): pp. 315–334.

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.

Prager, E. and Schmitt, M. (2021). "Employer Consolidation and Wages: Evidence from Hospitals." *American Economic Review*, 111 (2): pp. 397–427.

Reynaerts, J., Varadha, R., and Nash, J. (2012). "Enhancing the Convergence Properties of the BLP (1995) Contraction Mapping." *KU Leuven Working Papers of VIVES - Research Centre for Regional Economics 482522*.

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.

Seim, K. (2006). "An Empirical Model of Firm Entry with Endogenous Product-Type Choices." *The RAND Journal of Economics*, 37 (3): pp. 619–640.

Setzler, B. and Tintelnot, F. (2021). "The Effects of Foreign Multinationals on Workers and Firms in the United States." *The Quarterly Journal of Economics*, 136 (3): pp. 1943–1991.

Shapiro, C. (2019). "Protecting Competition in the American Economy: Merger Control, Tech Titans, Labor Markets." *Journal of Economic Perspectives*, 33 (3): pp. 69–93.

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.

Stock, J. and Yogo, M. (2005). *Testing for Weak Instruments in Linear IV Regression; Identification and Inference for Econometric Models*. Ed. by D. W. K. Andrews. New York: Cambridge University Press, pp. 80–108.

Sweeting, A. (2010). "The Effects of Mergers on Product Positioning: Evidence From the Music Radio Industry." *The RAND Journal of Economics*, 41 (2): pp. 372–397.

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.

Werden, G. J. and Froeb, L. M. (1998). "The Entry-Inducing Effects of Horizontal Mergers: An Exploratory Analysis." *The Journal of Industrial Economics*, 46 (4): pp. 525–543.

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.

Wollmann, T. G. (2018). "Trucks without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles." *American Economic Review*, 108 (6): pp. 1364–1406.

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 21.07.2022

Felix Montag