

Montag, Felix; Mamrak, Robin; Sagimuldina, Alina; Schnitzer, Monika

Working Paper

Imperfect price information, market power, and tax pass-through

New Working Paper Series, No. 337

Provided in Cooperation with:

George J. Stigler Center for the Study of the Economy and the State, The University of Chicago Booth School of Business

Suggested Citation: Montag, Felix; Mamrak, Robin; Sagimuldina, Alina; Schnitzer, Monika (2023) : Imperfect price information, market power, and tax pass-through, New Working Paper Series, No. 337, University of Chicago Booth School of Business, Stigler Center for the Study of the Economy and the State, Chicago, IL

This Version is available at:

<https://hdl.handle.net/10419/279798>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Stigler Center

for the Study of the Economy and the State

Imperfect Price Information, Market Power, and Tax Pass-Through

Felix Montag

Tuck School of Business

Robin Mamrak

Ludwig-Maximilians-University Munich

Alina Sagimuldina

Ludwig-Maximilians-University Munich

Monika Schnitzer

Ludwig-Maximilians-University Munich

November 2023

New Working Paper Series No. #337

Stigler Center for the Study of the Economy and the State
University of Chicago Booth School of Business
5807 S Woodlawn Ave
Chicago, IL 60637

Imperfect Price Information, Market Power, and Tax Pass-Through

Felix Montag, Robin Mamrak, Alina Sagimuldina, and Monika Schnitzer*

October 18, 2023

Abstract

Pass-through determines how consumers respond to taxes. We investigate the impact of imperfect price information on pass-through of commodity taxes. Our theoretical model predicts that the pass-through rate increases with the share of well-informed consumers. Pass-through is higher for the minimum price, paid by well-informed consumers, than for the average price, paid by uninformed consumers. Moreover, pass-through to the average price is non-monotonic with respect to the number of sellers. An empirical analysis of multiple recent tax changes in the German and French retail fuel markets confirms our theoretical predictions. Our results have implications for tax policy and shed light on the relative effectiveness of Pigouvian taxes versus regulation.

Keywords: pass-through, taxes, imperfect information, competition

*Montag: Tuck School of Business at Dartmouth and Stigler Center, Booth School of Business at the University of Chicago (email: felix.montag@tuck.dartmouth.edu); Mamrak: Department of Economics, Ludwig-Maximilians-University Munich (email: robin.mamrak@econ.lmu.de); Sagimuldina: Department of Economics, Ludwig-Maximilians-University Munich (email: alina.sagimuldina@econ.lmu.de); Schnitzer: Department of Economics, Ludwig-Maximilians-University Munich (email: schnitzer@econ.lmu.de); We thank Takanori Adachi, David Agrawal, Heski Bar-Isaac, Pierre Boyer, Thiess Büttner, David Byrne, Clemens Fuest, Tuomas Kosonen, Kory Kroft, Maarten Janssen, Erzo Luttmer, Erin Mansur, Simon Martin, Eugenio Miravete, Andreas Peichl, Markus Reisinger, Dominik Sachs, Chris Snyder, Meredith Startz, Dario Tortarolo, Thomas Wollmann, and participants of the 2021 CESifo Area Conference on Public Economics, the 2021 Spring Meeting of Young Economists, the 8th CRC Retreat, the 2021 EARIE Conference, the 2021 Congress of the Verein für Socialpolitik, the 2021 IIPF Annual Congress, the Consumer Search Digital Seminar Series, the Franco-German Fiscal Policy Seminar, the Conseil des Prélèvements Obligatoires for helpful comments and suggestions. We thank Christoph Winter for sharing the search data with us. Financial support by the German Research Foundation (DFG) through CRC/TRR190 (Project number 280092119) is gratefully acknowledged. Mamrak and Montag thank the International Doctoral Program “Evidence-Based Economics” of the Elite Network of Bavaria for financial support. Sagimuldina thanks the Research Training Group “GRK 1928 – Microeconomic Determinants of Labor Productivity” of the German Research Foundation for financial support.

1 Introduction

Understanding how sellers pass through taxes is fundamental for the design of optimal tax policies. When firms have market power, pass-through can affect the impact of Pigouvian taxes, the effectiveness of unconventional fiscal policy, and the distributional consequences of commodity taxes. Competitive conduct is a key determinant of pass-through. Weyl and Fabinger (2013) present a unified theoretical framework to study pass-through under imperfect competition, where competition is captured by a conduct parameter. A key assumption of this framework is that consumers possess complete price information. No such framework exists when there is imperfect information.

In this paper, we investigate how market power resulting from imperfect price information affects commodity tax pass-through. We show that imperfect consumer information about prices best explains the pass-through behavior of firms observed in our empirical application. Imperfect information is a common feature in most markets, affecting consumers' sensitivity to price differences. Empirically, we study retail fuel products, which have a high degree of price transparency and homogeneity as compared to other products. Nevertheless, they exhibit significant price dispersion, consistent with imperfectly informed consumers. We use a theoretical consumer search model in which some consumers know all prices and others have to search for prices sequentially. We adapt this model to derive how pass-through of a common cost shock, such as a tax change, depends on the share of well-informed consumers and the number of sellers. We demonstrate that when market power comes from imperfectly informed consumers, this affects the relationship between market concentration and pass-through in a way that cannot be captured by the conduct parameter approach.

Our theoretical analysis has key implications for the analysis of pass-through. First, we show that the share of well-informed consumers is positively related to the average price sensitivity of consumers (i.e., the price elasticity of residual demand) and the pass-through rate. Second, since firms play mixed strategies, the pass-through to the price paid by well-informed consumers is higher than the pass-through to the price paid by uninformed consumers in the same market. Third, we find that when market power is derived from imperfect information, the relationship between the number of competitors and the pass-through to the average price is non-monotonic. This suggests that the number of competitors can be a poor predictor of pass-through. Fourth, the full information conduct parameter approach (see Weyl and Fabinger, 2013) cannot nest models where market power derives from imperfect information. Its results on the determinants of pass-through thus do not extend to imperfect information.

To empirically test the model predictions, we analyze the pass-through of multiple recent tax changes in the German and French retail fuel markets, using detailed price data

at the station level. A unique aspect of the setting is that we can separately study fuel products that differ in how well their users are informed about prices. We find that tax pass-through is higher for fuel types with a larger share of well-informed consumers. We also show that pass-through to the minimum price is higher than to the average posted price for nearly all of the tax changes that we study. Finally, we find a non-monotonic relationship between the number of competitors and pass-through to the average price in a local market. Together, the different findings strongly suggest that imperfect consumer information about prices is necessary to explain the observed pass-through patterns.

These results are widely applicable beyond the retail fuel market, because markets with both well-informed and uninformed consumers are widespread across the economy. For instance, models of competition with imperfect information are used to explain price differences between online and brick-and-mortar stores (Baye, Morgan, and Scholten, 2006). In such settings, understanding commodity tax pass-through requires a stronger emphasis on information, rather than the number of competitors.

Our findings are crucial to assess the impact of tax policy. The lower the share of well-informed consumers, the lower is the pass-through of commodity taxes. This can make Pigouvian taxes less effective, as there will be a smaller output response from consumers, compared to a setting with full information.¹ Since pass-through differs between well-informed and uninformed consumers, output reactions across consumer groups will be different, which has distributional implications. Pigouvian taxes may therefore induce stronger quantity reactions by well-informed consumers as compared to uninformed consumers. This affects the relative benefits of Pigouvian taxes versus regulation. Similarly, if few consumers are well informed about prices, this lowers tax pass-through and limits the possibility to stimulate the economy using unconventional fiscal policy.

In the theoretical analysis, we modify the Stahl (1989) model to examine pass-through. This model features a homogeneous good, as well as fully informed shoppers and uninformed non-shoppers who can search for prices sequentially. The degree of market power depends on the number of competitors and the share of well-informed consumers, as a higher share incentivizes firms to compete on prices.

The equilibrium of the model is characterized by a distribution of prices, because firms set prices using mixed strategies. Well-informed shoppers always buy from the seller offering the lowest price. Uninformed consumers do not search in equilibrium and instead pay the first price they draw. From an ex ante perspective, informed shoppers pay the expected minimum price, whereas uninformed non-shoppers pay the expected price.

¹Although output can also be reduced with market power, Conlon and Rao (2023) demonstrate that limiting competition to address negative externalities results in much higher welfare costs than using taxes.

The model offers several predictions regarding how competition affects pass-through. First, the higher the proportion of well-informed consumers, the greater the pass-through rate to all prices. Second, pass-through to the minimum price is higher than to the average price in a market. Third, the pass-through rate to the expected price first increases and eventually declines as the number of sellers increases. This is because above a certain threshold of competitors, it becomes increasingly unlikely for a particular firm to attract shoppers. Consequently, firms are more likely to charge a higher price and only cater to uninformed non-shoppers. With imperfect price information, a larger number of sellers does not necessarily lead to lower average prices.

Our theoretical analysis of tax pass-through differs from traditional analyses, as we consider market power derived from imperfect information. Many studies in the empirical literature on tax pass-through assume perfectly competitive markets (e.g., Chetty, Looney, and Kroft, 2009 or Chetty, 2009). In contrast, a growing theoretical literature considers how firms with market power pass through taxes (e.g., Sumner, 1981, Bulow and Pfleiderer, 1983, Stern, 1987, and Hamilton, 1999), with Weyl and Fabinger (2013) providing a general model to capture the intensity of competition.² All of these models assume that consumers are fully informed about prices.

Some studies depart from the full information assumption. Many of these models assume that consumers are aware of posted net prices but that the tax component applied at checkout is less salient (e.g., Chetty, Looney, and Kroft, 2009 or Kroft et al., 2023). Similarly, Busse, Silva-Risso, and Zettelmeyer (2006) analyze differences in the pass-through of promotions by auto manufacturers that differ in how salient they are to consumers. Approaches that rely on differences in salience between the net price and taxes cannot explain findings in our context, where the gross price including taxes is posted.

Different models lead to distinct theoretical predictions, which can be empirically falsified. We discuss alternative hypotheses of what could explain the empirical results in full information, as well as imperfect information frameworks. We conclude that none of these hypotheses can explain the empirical findings jointly as well as the Stahl (1989) model.

We test the predictions empirically using multiple tax changes in the German and French fuel markets. This industry is ideal for this study because the fuel type consumers purchase is highly correlated with their incentive to be informed about prices. In Germany, there is strong evidence suggesting that diesel drivers are better informed about prices than gasoline drivers.³ Moreover, drivers fueling *E5* gasoline are less informed about prices than

²Adachi and Fabinger (2022) generalize this to allow for richer governmental intervention and Kroft et al. (2021) allow for free entry and love-of-variety preferences.

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

those fueling *E10*. We use search data from a price comparison smartphone app to confirm these hypotheses. Using differences between fuel types, we can test the predictions about the relationship between pass-through and the share of well-informed consumers. Since fuel stations sell all three fuel types, we can disentangle different components of market power. We can test how pass-through varies across consumer groups with different levels of information, while holding the station network fixed. Similarly, we can test how pass-through varies across stations with different numbers of competitors, while holding the consumer type fixed.

Our empirical analysis examines the impact of a temporary decrease in the value-added tax (VAT) and the introduction of a carbon emissions price in Germany in 2020/21. We estimate pass-through rates separately for each fuel type by comparing daily prices of German stations with those in France, using a difference-in-differences (DID) design. We test the robustness of our findings by analyzing three French tax changes in 2022/23.

The first empirical finding is that tax pass-through is higher for fuel types with a higher share of well-informed consumers. The empirical literature on tax pass-through and market power (see, for example, Miravete, Seim, and Thurk, 2018, Hollenbeck and Uetake, 2021, or Nakamura and Zerom, 2010) has so far ignored imperfect information. Most closely related to our mechanism, Duso and Szücs (2017) find higher cost pass-through for electricity tariffs that consumers actively need to choose than for default tariffs. Similarly, Kosonen (2015) find that Finnish hairdressers pass on VAT decreases more for advertised services. Our results also relate to Eizenberg, Lach, and Oren-Yiftach (2021), who find that spatial frictions and differences in the sensitivity to lower prices between different neighborhoods leads to spatial differences in market power and price levels.

The second empirical finding is that pass-through to the minimum price, paid by well-informed consumers, is higher than pass-through to the average posted price, paid by uninformed consumers. This extends the literature on the distributional implications of tax pass-through. For example, Harju et al. (2022) find lower fuel tax pass-through in high-income areas. Conlon, Rao, and Wang (forthcoming) show that the sin tax burden is concentrated among few households that exhibit similar purchasing patterns. Our understanding of who searches is restricted to differences between consumers buying different fuel types. Byrne and Martin (2021) review the literature and document an inverse-U relationship between household income and consumer search.

The third empirical finding is a non-monotonic relationship between the number of sellers and the pass-through rate to the expected price, or average price. The existing empirical literature is mixed, with evidence that tax pass-through increases (e.g., Genakos and Pagliero, 2022), decreases (e.g., Miller, Osborne, and Sheu, 2017 and Hindriks and Serse, 2019) or is uncorrelated (e.g., Kopczuk et al., 2016) with the number of competitors.

We find higher pass-through for tax increases than tax decreases. Yet, as we never observe a symmetric tax increase and decrease by the same magnitude, this is not evidence of asymmetric tax pass-through beyond what can be explained by standard economic models. In contrast, Benzarti et al. (2020) find higher pass-through for tax increases than decreases when tax changes are symmetric.⁴ There are a number of studies estimating the pass-through of different tax changes, but none of them studies its relation with market power.⁵

Finally, by emphasizing how imperfect consumer information affects tax pass-through in oligopolistic markets, we contribute to a literature on the supply side determinants of tax pass-through in fuel markets (e.g., Marion and Muehlegger, 2011 or Fischer, Martin, and Schmidt-Dengler, 2023).⁶ An earlier literature studies fuel market cost pass-through using error correction models.⁷ Relatedly, Borenstein, Cameron, and Gilbert (1997) conclude that asymmetric pass-through may be explained by tacit collusion or imperfect information. Deltas and Polemis (2020) show that results from error correction models strongly depend on the research design and data features.

The remainder of the paper is structured as follows: Section 2 describes the data and derives stylized facts about the fuel market. Section 3 outlines the theoretical model. Section 4 introduces the tax changes and provides descriptive evidence. Section 5 presents the empirical strategy. Section 6 discusses the estimation results. Section 7 contrasts the empirical results with alternative hypotheses and Section 8 concludes.

2 Consumer Information in the Retail Fuel Market

We begin by describing the data and highlight the key features of the retail fuel markets in Germany and France.

2.1 Data

Our comprehensive dataset includes real-time price changes for almost all fuel stations in Germany and France, along with various station characteristics. German stations are man-

⁴When search is high, Heim (2021) finds pass-through is high for cost decreases and low for cost increases.

⁵These include single industry studies (e.g., Fabra and Reguant, 2014, Li and Stock, 2019, Ganapati, Shapiro, and Walker, 2020, Büttner and Madzharova, 2021, Dubois, Griffith, and O’Connell, 2020, Harding, Leibtag, and Lovenheim, 2012, or Conlon and Rao, 2020), and cross-industry studies (e.g., Benedek et al., 2020).

⁶Imperfect information is known to play an important role for pricing more generally in these markets (see, e.g., Chandra and Tappata, 2011, Byrne and Roos, 2017, Luco, 2019, Pennerstorfer et al., 2020, Martin, forthcoming, or Montag, Sagimuldina, and Winter, 2023).

⁷For a review of the literature, see Eckert (2013).

dated to report price changes to the Market Transparency Unit at the Federal Cartel Office.⁸ Similarly, in France, a government agency requires stations to report price changes, providing researchers access to this data.⁹ We construct daily weighted average prices for each station, using the time of price changes. See Appendix A.1 for details on our dataset construction.

We analyze data from January 2019 to February 2023. We calculate summary statistics for 2019 to capture pre-intervention markets, as all tax changes occurred between 2020 and 2023. The top panel of Table 1 presents station-level summary statistics.

To analyze local price dispersion and competitive dynamics, we group fuel stations into non-overlapping markets using a hierarchical clustering algorithm based on driving time, as done in previous studies (e.g., Carranza, Clark, and Houde, 2015, Luco, 2019, or Assad et al., 2020). The idea underlying this approach is to find clusters of stations that are naturally separated from each other. The details of our clustering method are explained in Appendix A.2. Table 1 shows that we assign the 14,648 German stations to 3,479 local markets with an average size of 4.2 stations. In France, there are 9,075 fuel stations assigned to 2,769 local markets. France has fewer markets and fewer stations per market, which is most likely related to its lower population and lower population density.

Ultimately, we are interested in the number of competitors in a local market. We measure the number of competitors by the number of competing price setters. That is, if there are two stations for which the same entity sets prices, we want to treat it as a single price setter. For Germany, we have two data sources that allow us to establish a common price setter between stations. First, the station dataset contains information on the brand of a fuel station. Prices at stations belonging to a brand of the vertically integrated fuel producers (e.g., Aral or Shell) are set centrally by the brand’s headquarters, irrespective of whether the station is owned by the fuel producer or by a third-party owner-operator.¹⁰ Moreover, some firms operate fuel stations under different brands or even unbranded stations. “Wer-zu-wem” is a database that contains ownership information for many such stations and allows us to group brands together with a common price setter (e.g., the brand Elan belonging to TotalEnergies). Ultimately, we compute the number of competing price setters in a local market. On average, there are 3.6 different price setters per market in Germany. Furthermore, 16% of markets only contain a single price setter and are thus monopoly markets. The median monopoly station is 10.6 minutes away from its closest competitor,

⁸TankerKönig, a price comparison website, provides access to this data.

⁹See <https://www.prix-carburants.gouv.fr/rubrique/opendata/>. In France, fuel stations selling less than 500 m^3 of fuel per year are exempt from reporting price changes.

¹⁰We conducted several interviews with market participants. All our interviewees confirmed that prices are set at the headquarter level both for large integrated conglomerates as well as for most small- and medium-sized station operators. See, e.g., for Shell: <https://support.shell.de/hc/de/articles/360010715077-Wer-bestimmt-die-Kraftstoffpreise-an-den-Shell-Stationen->.

Table 1: Summary statistics

	Germany	France
<i>Station level</i>		
A. Station characteristics		
Number of stations	14,648	9,075
Median driving time to closest competitor	2.4min	n/a
Median driving distance to closest competitor	1.4km	n/a
B. Prices, <i>E5</i>		
Mean gross price	1.41	1.53
Mean price net of taxes and duties	0.53	0.58
C. Prices, <i>E10</i>		
Mean gross price	1.39	1.49
Mean price net of taxes and duties	0.51	0.57
D. Prices, Diesel		
Mean gross price	1.25	1.45
Mean price net of taxes and duties	0.57	0.60
<i>Market level</i>		
E. Market characteristics		
Number of markets	3,479	2,769
Mean no. of stations in market	4.21	3.28
Mean no. of competitors	3.60	n/a
F. Local monopolies		
Share of monopoly markets	16%	n/a
Median driving time to closest competitor	10.6min	n/a
Median driving distance to closest competitor	9.4km	n/a
G. Prices, <i>E5</i>		
Mean average posted price	1.42	1.53
Mean minimum price	1.38	1.51
H. Prices, <i>E5</i>		
Mean average posted price	1.40	1.48
Mean minimum price	1.35	1.46
I. Prices, Diesel		
Mean average posted price	1.25	1.45
Mean minimum price	1.21	1.43

Notes: The Table shows summary statistics for 2019 (i.e., before all tax changes). The top panel presents data at the station level, whereas the bottom panel presents data at the market level. Non-overlapping markets are defined using a hierarchical clustering algorithm, as explained in Appendix A.2. A competitor is defined as a competing price setter (available only for Germany).

whereas the median driving time to the closest competitor across all stations is only 2.4 minutes. For France, our dataset does not contain information on the brand or ownership of a fuel station.

Finally, we leverage data on search queries in 2015 from a major German smartphone app that enables users to compare fuel prices across stations. Anytime a user searches for fuel prices, the dataset contains the searcher’s location, a time stamp, a unique user ID, and the fuel type searched. This allows us to document intensive and extensive margin differences in the search intensity between consumers that search for different fuel types.

2.2 Fuel types

Diesel and gasoline are the two primary fuel types for passenger vehicles with combustion engines. In Germany, diesel accounts for 43% of the volume share and gasoline accounts for the remaining 57%.¹¹ The high costs of substitution between these two types on the demand and supply sides means they can be considered as separate markets in the short term.

Gasoline can be classified according to its octane rating and ethanol share. Standard gasoline (called *Super*) has an octane rating of 95 and can be further distinguished by its ethanol share. Gasoline with a 5% share of ethanol is referred to as *E5*, while *E10* has a 10% ethanol share.¹² While *E5* and *E10* are not taxed differently, *E10* is typically 4-6 Eurocent cheaper in Germany due to a minimum biofuels quota. Since *E5* and *E10* have the same octane rating, there are no quality differences between the fuel types.

2.3 Price dispersion


The summary statistics in Table 1 reveal substantial price dispersion within local markets on a particular day for fuel stations selling the same products. To understand the sources of this variation, we decompose it into components related to intertemporal differences in demand or product differentiation, and random, unpredictable price changes.

Figure 1 illustrates that the average price of *E10* in Germany at different times of the day varies significantly, with prices at around 7:30 am being more than 10 Eurocent higher than prices at around 10 pm. On average, there are 14 daily price changes for *E10* at German fuel stations in 2019. As noted by Holt, Igami, and Scheidegger (2023), these price cycles are different to other countries (e.g., Australia). They are not cost-driven, as costs can be assumed not to vary within a day. Instead, the price cycles serve two purposes: first,

¹¹Based on 2019 data from the German Ministry of Transportation’s *Verkehr in Zahlen 2020/2021*. Truck diesel prices are not included as they are not reported to the Market Transparency Unit.

¹²In addition, there are other types of *Super* with an octane rating of 98, but their market share in Germany is only around 5%.

Figure 1: Average daily price cycles for *E10* in Germany, 2019



Notes: The Figure shows average prices of *E10* in 2019 across fuel stations in Germany at different times of the day. Fuel prices are updated in five-minute intervals.

high prices in the morning and lower prices in the evening are consistent with intertemporal price discrimination, where prices are high when drivers have little time to search for better prices. Second, frequent price changes during the day make it difficult for drivers to learn which station is the cheapest at a particular point in time, making it more likely for sellers to be able to sell at a price higher than the minimum price in the market.

Although the theoretical predictions in Section 3 do not feature price cycles, these are not inconsistent with the model. Cason, Friedman, and Hopkins (2020) show that in experimental markets that have similar features to the Burdett and Judd (1983) or Stahl (1989) search models, the observed pricing patterns feature price cycles. After collapsing prices over time, the mixed strategy equilibrium distribution in the theoretical search models does well and explaining the observed price distribution.

To identify the random and unpredictable price dispersion for consumers, we narrow our focus to a particular time of day, 5 pm, and calculate the absolute price deviation of fuel stations from the mean price in their local market for all non-monopolistic stations. We do this by regressing daily 5 pm prices on market \times date fixed effects and taking the resulting residuals. However, some stations may always deviate from the mean price in the same way due to product differentiation. They may, for example, be in a particularly attractive location or offer better amenities. To isolate the non-constant part of the deviation from the

Table 2: Within market price residual, 5 pm, 2019

	Stations	Markets		
	Mean abs. deviation	p_{25-p75}	p_{10-p90}	Range
<i>A. E5</i>				
Market \times date FE	.0162	.0271	.0408	.0417
Market \times date FE and station FE	.0104	.0177	.0272	.0279
<i>B. E10</i>				
Market \times date FE	.0173	.0291	.0439	.0449
Market \times date FE and station FE	.0105	.0181	.0275	.0281
<i>C. Diesel</i>				
Market \times date FE	.0161	.0269	.0407	.0416
Market \times date FE and station FE	.0103	.0175	.0272	.0278
No. of stations/markets	14,140	2,971	2,971	2,971
Observations	3,507,612	775,431	775,431	775,431

Notes: The Table shows the distribution of the average absolute deviation of a fuel station's price from the average price in the same market on the same day at 5 pm for each fuel type and for all stations that are not local monopolists. We use data for all weekdays in 2019. It also shows the distribution of this absolute deviation after controlling for station fixed effects. The mean absolute deviation shows the average across all fuel stations. We compute the different range measures by calculating the range for each individual market on a particular day and then averaging across days and markets.

market mean, we further control for station fixed effects. The remaining price variation is unpredictable even to the most sophisticated consumers.

Table 2 decomposes the observed price dispersion into predictable and unpredictable components. On average, the absolute price deviation from the market mean is 1.6 Eurocent for *E5* and diesel and 1.7 Eurocent for *E10*. The mean absolute deviation from the mean after controlling for station fixed effects, which is the unpredictable component, remains above 1.0 Eurocent for all fuel types. In Appendix A.4, we present these within market price residuals graphically.¹³ Furthermore, the average difference between the cheapest and the most expensive fuel station in a local market in terms of the unpredictable component is around 2.8 Eurocent for all fuel types, which is substantial.

Stylized Fact 1. *There is a substantial share of price dispersion that is random and unpredictable to consumers.*

2.4 Consumer information

Fuel stations in Germany and France are required to immediately report price changes, enabling real-time price information to be available to consumers via smartphone apps. These

¹³We also show price cycles at a more disaggregated level by zooming in on one local market and individual days.

apps provide perfect information on prices to users, whereas non-users can only discover prices by driving from station to station.

Stylized Fact 2. *Some consumers know all prices (app users), whereas others need to search for prices sequentially.*

How well informed consumers are about prices often correlates with the fuel type they purchase. Frequent drivers often prefer diesel cars. Accordingly, diesel passenger vehicles drive 19,200 kilometers per year on average, compared to 10,800 kilometers for gasoline passenger vehicles.¹⁴ Although diesel cars are more expensive to buy, the cost of fuel is usually lower, making it a fixed-cost investment to lower the marginal cost of driving. Therefore, drivers who select diesel engines have a higher incentive to save on fuel costs, so they are more likely to use price comparison apps.

To assess differences in search intensity across fuel types, we use data on search queries in 2015 from a major German price comparison smartphone app. Normalizing the number of app users by the number of registered vehicles, we find that the distinct number of users who search for diesel prices is around 50% higher than the number of users who search for gasoline prices. This is in line with the hypothesis that, on average, the share of diesel drivers that are well informed about prices is higher than the share for drivers of gasoline cars. Further details are presented in Appendix A.3.

Commercial vehicles often run on diesel. If drivers of commercial vehicles do not pay for their own fuel, they may be less sensitive to prices. It is therefore worth discussing why commercial vehicles are not a concern for our analysis. First, as of 1 January 2020, there were 15.1 million passenger vehicles with a diesel engine, but, including those with a gasoline engine, there were only 5.2 million commercial passenger vehicles (Kraftfahrt-Bundesamt, 2021). Hence, at least 66% of passenger cars with a diesel engine are owned by private individuals. Second, some commercial drivers, such as those receiving a fuel allowance or those that are self-employed, also have an incentive to reduce fuel costs. Therefore, the fact that many commercial vehicles run on diesel does not undermine our finding that drivers of diesel vehicles are, on average, more price sensitive than drivers of gasoline vehicles.

In addition to differences between diesel and gasoline, there are differences in price sensitivity between buyers of *E5* and *E10* in Germany. These are likely driven by unwarranted concerns about potential damage to the engine caused by biofuels, which arose around the introduction of *E10* in 2011 and help us to further segment consumers according to how

¹⁴Based on 2019 data from the German Ministry of Transportation’s *Verkehr in Zahlen 2020/2021*.

well informed they are.¹⁵ Despite being more expensive, the majority of gasoline drivers in Germany purchase *E5*. According to the German Automobile Association (ADAC), *E10* is around 1.5% less efficient than *E5*.¹⁶ However, this accounts only for a fraction of the observed difference in prices between *E5* and *E10*, since *E10* is usually 4-6 Eurocent cheaper. A survey conducted by the ADAC in 2020 suggests that the difference in price sensitivity between *E5* and *E10* is due to preferences and a lack of information. The survey found that among respondents fueling *E10*, the most cited reason for doing so is lower prices (72%), followed by environmental concerns (37%). Among those not fueling *E10*, the most cited reasons are technical concerns (51%) and uncertainty about the cost and benefits (23%).¹⁷ Since the octane rating is the same for *E5* and *E10*, there are no quality differences between the two fuel types.

This evidence strongly suggests that, among drivers of gasoline cars, more buyers of *E10* choose to become informed about prices in Germany. Again, we confirm this hypothesis with our search data in Appendix A.3. In particular, we find that, adjusted for the relative market shares of *E5* and *E10* within the gasoline market, search intensity is substantially higher among consumers buying *E10* than those purchasing *E5*. In France, in contrast, no such controversy regarding *E10* existed. Therefore, drivers of gasoline vehicles in France predominantly buy *E10*.

Stylized Fact 3. *The share of well-informed consumers (app users) differs between fuel types. In Germany, it is higher for diesel than for gasoline and it is higher for E10 than E5.*

3 Theoretical Model

Motivated by the stylized facts in Section 2, we adapt the Stahl (1989) model to analyze the determinants of pass-through in a setting where firms sell a homogeneous good to consumers who are either fully informed or can search for lower prices. The model generates testable predictions tailored to the empirical setting.

3.1 Setup

On the demand side, there is a mass M of consumers, each with the same valuation v and inelastic unit demand for a homogeneous good. Consumers in the market can be divided

¹⁵Although biofuels can pose a significant threat to the engine of a vehicle that is not compatible with *E10*, around 90% of gasoline-run vehicles, including all vehicles produced after 2012, are compatible. A full list of compatible vehicles can be found at <https://www.dat.de/e10/>.

¹⁶See <https://www.adac.de/verkehr/tanken-kraftstoff-antrieb/benzin-und-diesel/e10-tanken/>.

¹⁷The full survey results can be found at <https://www.adac.de/news/umfrage-e10-tanken/>.

into two groups: fully informed shoppers, who know the prices of all sellers and always buy from the lowest-price seller, and non-shoppers, who draw a first price for free, know the distribution of prices, and can decide to sequentially search for prices at an incremental search cost s until they find a price that is weakly below their reservation price p_r . The model assumes that a fraction ϕ of consumers is fully informed shoppers and the remaining fraction $1 - \phi$ consists of non-shoppers.

On the supply side, there is an exogenous number of sellers denoted by N , which produce at a constant marginal cost of c .¹⁸ Sellers are indexed by i . Sales are subject to an ad valorem tax τ .

Sellers first choose their price and consumers then make search and purchase decisions. We search for the subgame perfect Nash equilibrium of the game via backward induction.

Before we proceed, we introduce some additional notation. Whenever mentioning prices, we refer to the gross price paid by consumers. We assume that sellers bear the initial incidence of the tax and then (partially) “pass through” the cost of the tax to consumers. It is well-established in the theoretical literature that equilibrium prices are equivalent regardless of whether the initial tax incidence is on buyers or sellers. We denote the pass-through rate of marginal costs as $\rho_c = \frac{\partial p}{\partial c}$. 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 denoted as

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

We focus on the determinants of the pass-through rate of the ad valorem tax. In Appendix B.3, we show that the main mechanisms are the same for a per unit tax.

The model differs from traditional models of pass-through in its notion of price sensitivity. While many models capture consumers’ sensitivity to price changes through the price elasticity of aggregate demand or the closeness of substitution of differentiated products, we consider the share of shoppers ϕ and the incremental search cost of non-shoppers s as the primary determinants of price sensitivity. A larger share of shoppers results in more consumers purchasing from the lowest-price seller, thus reducing the expected profit of setting a price above the market minimum. Similarly, lower search costs for non-shoppers incentivize them to search for lower prices, leading to lower reservation prices and prices overall. In Section 7, we discuss why the price elasticity of demand or product differentiation in a full information model cannot explain the empirical findings.

¹⁸We endogenize entry in Appendix B.2.

3.2 Equilibrium price distribution

In the following, we characterize the equilibrium while the analysis of the model is relegated to Appendix B.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 that a seller charges. 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

$$p_r = \min \{E[p] + s, v\} .$$

If searching sequentially is sufficiently cheap, the reservation price of non-shoppers is the sum of the expected price at the next draw and the search cost s . With relatively high search costs, the reservation price of non-shoppers is simply the valuation of the good v and the model boils down to the well-known Varian (1980) setting.

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 profits of a seller are

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

In equilibrium, non-shoppers buy at the first price they draw, making the expected price equal to the average price paid by non-shoppers. On the other hand, shoppers buy from the lowest-price seller, resulting in the expected minimum price being equal to the average price paid by shoppers.¹⁹

¹⁹The average minimum price refers to the average price paid by shoppers if this game is often repeated across time or space. At a given time and location, there is only one minimum price and N prices.

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

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

3.3 Pass-through of an ad valorem tax

To analyze how ad valorem taxes are passed through to consumers, we first examine the impact of an increase in the ad valorem tax τ on the equilibrium pricing strategy. We assume that the search cost s is sufficiently high so that the reservation price p_r is equal to the consumers' valuation v , simplifying the framework to a Varian (1980) setting. In Section 3.5, we relax this assumption and use numerical examples to show that our results hold even when search costs are low, the Stahl (1989) setting with sequential search.²⁰

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.²¹

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

When the share of shoppers is strictly positive, increasing the ad valorem tax τ leads to a shift in the support of prices from which sellers draw in equilibrium towards higher prices. Additionally, for each price on this support, the likelihood of a drawn price being lower than that price decreases with an increase in the ad valorem tax rate to $\hat{\tau}$. As the share of shoppers converges to zero, the Nash equilibrium converges towards a degenerate distribution at the monopoly price, the classical result by Diamond (1971). The monopoly price corresponds to the valuation of the good, v .

Since the minimum element of the support of prices and the density function monotonically 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.

²⁰An alternative simplification would be setting $N = 2$, which is less desirable to study the effect of competition.

²¹The proof of Proposition 1 and all following Propositions can be found in Appendix B.4.

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

Let us begin by examining the two extreme cases. If there are no shoppers, the Nash equilibrium is a degenerate distribution at the monopoly price, which is unaffected by the ad valorem tax, and pass-through is zero. However, if the share of shoppers approaches one, the Nash equilibrium approaches the classical Bertrand equilibrium, where the Nash equilibrium is a degenerate distribution at $c(1 + \tau)$, and there is full pass-through.

As the share of shoppers ϕ increases from zero to one, the pass-through rate of the ad valorem tax to the lower bound of the equilibrium price strategy strictly increases. Furthermore, the rate at which an increase in the tax from τ to $\hat{\tau}$ reduces the probability of drawing a price below a certain price p , i.e., from $F(p)$ to $\hat{F}(p)$, also strictly increases as the share of shoppers increases. Therefore, the pass-through rate increases with the share of shoppers, and it reaches full pass-through as the share of shoppers approaches one.

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

Besides the share of informed consumers, the number of active sellers is also an important dimension of competition, often more salient in empirical applications.

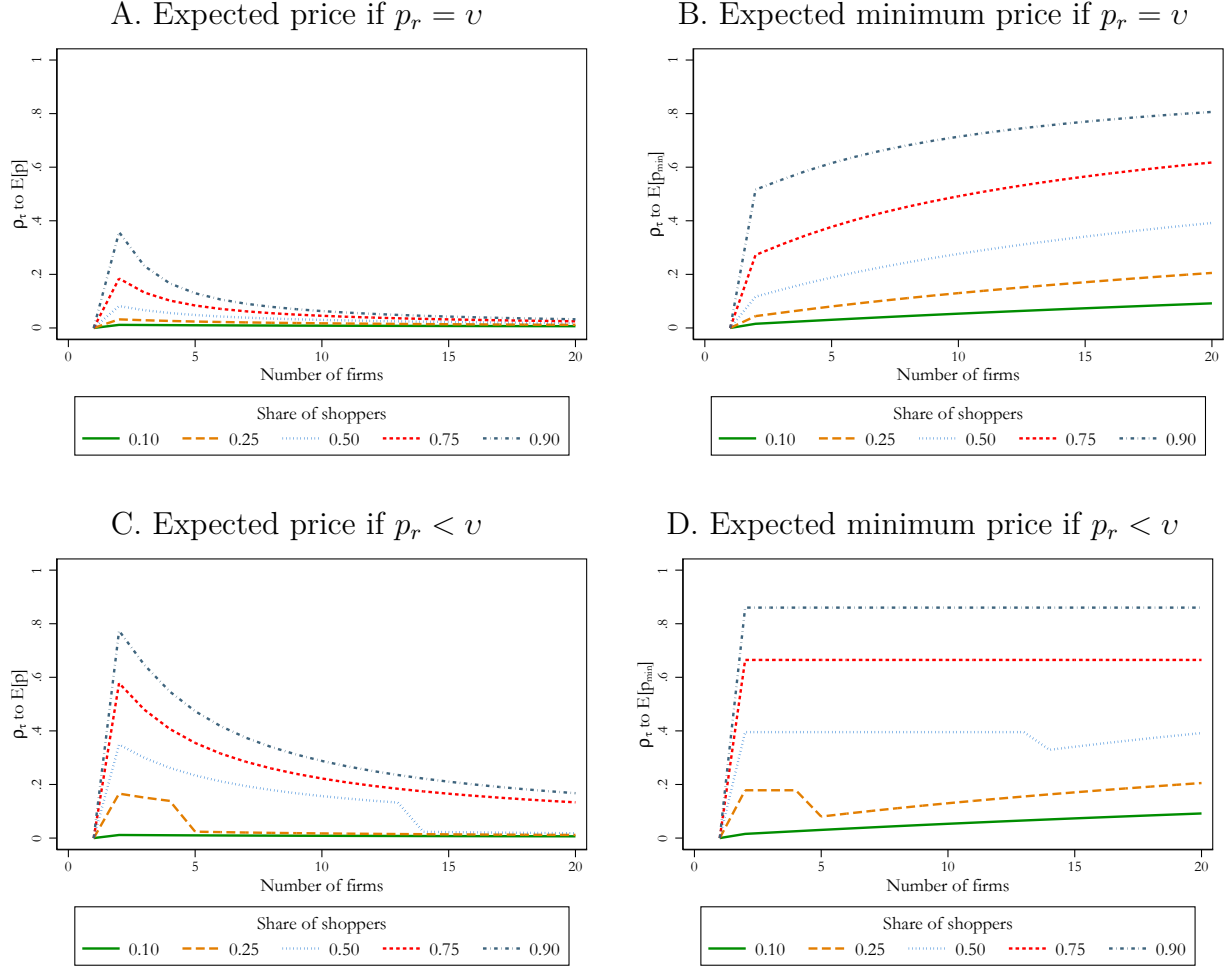
Proposition 3. *With $0 < \phi < 1$, as $N \rightarrow \infty$ the pass-through of τ to the minimum element of the equilibrium price support converges to full pass-through, i.e., $\rho_{\tau,p} \rightarrow 1$.*

With more sellers, competition for shoppers becomes more intense, leading to convergence of the minimum price that sellers consider charging in the symmetric Nash equilibrium towards $c(1 + \tau)$. As a result, the pass-through rate of the ad valorem tax to \underline{p} increases.

Showing how an increase in N affects the pass-through rate to $F(p)$, $E[p]$ and $E[p_{min}]$ analytically is more difficult. Instead, we numerically simulate how a change in the tax affects $E[p]$ and $E[p_{min}]$ for a given set of parameters and varying the number of sellers, N .

We show the numerical results for a particular choice of parameter values in Figure 2. Panels A and B illustrate how pass-through of an ad valorem tax to $E[p]$ and $E[p_{min}]$ varies with the number of sellers in a Varian (1980) setting, where the sequential search cost of non-shoppers s is so high that their reservation price is equal to their valuation of the good, i.e., $p_r = v$. Panels C and D show how pass-through of an ad valorem tax to $E[p]$ and $E[p_{min}]$

Figure 2: Numbers of sellers and tax pass-through



Notes: The Figure shows simulation results of how the pass-through rate of the ad valorem tax τ varies with the number of sellers. Panel A and B respectively show how the pass-through rate to the expected price, $E[p]$, and to the expected minimum price, $E[p_{min}]$, vary with the number of sellers if the reservation price is exogenous. Panel C and D show the same if the reservation price of non-shoppers, p_r , is endogenous. In all panels, the different lines correspond to different values of the share of shoppers, ϕ . Parameter values: $v = 4.5$, $c = 0.4$, $\tau = 0.2$, $\hat{\tau} = 0.22$, $s = \infty$ (without sequential search) and $s = 0.75$ (with sequential search).

varies with the number of sellers in a Stahl (1989) setting, where the sequential search cost of non-shoppers is sufficiently low such that their reservation price depends on the price they expect to draw if they were to search, i.e., $p_r < v$.

There are two key results from this numerical exercise. If at least some consumers are not perfectly informed, whatever parameter values we choose, there is always a non-monotonic relationship between the number of sellers and the pass-through rate to $E[p]$. This is not the case in models with perfect information, where the pass-through rate to $E[p]$ monotonically increases in the number of sellers.

For pass-through to $E[p_{min}]$, the results are more nuanced. In a Varian (1980) setting more sellers decrease $E[p_{min}]$. The more $E[p_{min}]$ converges to marginal costs, the lower the margin that sellers could use to absorb a tax increase. Thus, pass-through to $E[p_{min}]$ monotonically increases in N . In a Stahl (1989) setting there is an additional countervailing effect. If the reservation price is endogenous, this is a function of the expected price. When the number of sellers increases, $E[p]$ increases and p_r increases. This decreases the incentive for sellers to set lower prices, increasing the expected minimum price and decreasing pass-through to $E[p_{min}]$. Depending on the relative strength of these two effects, when the reservation price is endogenous, pass-through to the expected minimum price can increase or decrease in the number of sellers.

Since there is no clear prediction about the relationship between the number of sellers and the expected minimum price, the key testable implication of the numerical exercise is that if there is imperfect information, the relationship between the number of sellers and pass-through to $E[p]$ is non-monotonic. Although we cannot prove that this is always true, our numerical results, combined with the following analogy to Stahl (1989) give us confidence that this is true for any parameter value.

In a simpler setting without taxes and or marginal costs, but for a wider class of demand functions, Stahl (1989) shows how the equilibrium price distribution behaves when there is an increase in the number of sellers. For the special case of inelastic unit demand, pass-through is inversely related to price. The higher the equilibrium price, the lower is pass-through. Although with taxes and marginal costs, both of which are necessary to analyze pass-through, it becomes intractable to prove the relationship between the number of sellers and pass-through, we can still learn something about the relationship between the number of sellers and pass-through from the relationship between the number of sellers and equilibrium prices.

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 $N \rightarrow \infty$. As N increases from one to two, the equilibrium price distribution shifts from a degenerate distribution at the monopoly price to a more competitive distribution that includes prices below the monopoly price. Intuitively, to get equilibrium prices below the monopoly price requires more than one seller. The more sellers there are in a market, the less likely it becomes for each individual seller to have the lowest price and attract shoppers. Each seller increases the likelihood of charging the reservation price of non-shoppers and foregoing the possibility to sell to shoppers. Accordingly, the expected price first decreases and then increases as N increases. Similarly, we expect the pass-through rate of ad valorem taxes to $E[p]$ to increase and then decrease as N goes to infinity.

3.6 Deriving empirically testable predictions

Our empirical setting deviates from the theoretical model in several ways. Incorporating these features into the model is difficult, so we qualitatively discuss how they affect the testable predictions. In Section 7, we discuss how well alternative hypotheses can explain the empirical results.

In the theoretical model, players have expectations about the average price and the minimum price in a market. These are the expected price and the expected minimum price, respectively. In the empirical application, we do not observe these expectations. Instead, we observe many different local markets. The sample equivalents to these theoretical objects are therefore the average price and the minimum price in a market.

The first prediction is based on Propositions 1 and 2. Since close to all stations in Germany sell all three fuel types under consideration, station-level product differentiation should not affect the relative pass-through between fuel types. Propositions 1 and 2 deal with the full distribution of equilibrium prices. Prediction 1 should therefore hold for any moment of the price distribution.

Prediction 1. *Pass-through is higher when the share of well-informed consumers is higher. In Germany, we expect pass-through to be highest for diesel and lowest for E5.*

In the model, consumers purchase at most one unit of the good and consumers cannot choose to buy different quantities in response to different prices. If the price elasticity of aggregate demand is highest for diesel and lowest for E5, this works in the opposite direction to the effect of consumer information. Finding higher pass-through for diesel would thus further underscore the importance of imperfect information. If the price elasticity of aggregate demand is lowest for diesel and highest for E5, the price elasticity of aggregate demand would be an observationally equivalent explanation with respect to Prediction 1. Finding support for Predictions 2 and 3 continues to require imperfect information, however.

The final two predictions are based on results from the numerical exercise. Figure 2 shows that pass-through to the minimum price, paid by well-informed consumers, is predicted to be higher than pass-through to the average price, paid by uninformed consumers. This holds for any $N \geq 2$ and remains the case with horizontal differentiation.

Prediction 2. *Pass-through to the minimum price is higher than pass-through to the average price.*

There is some degree of horizontal product differentiation caused by different locations of fuel stations. Travelling between stations comes at a cost, and the degree of substitution between stations decreases with travel time. Horizontal differentiation, through the distance

between stations, increases a station’s market power. The closer the competing stations are to each other, the lower their market power becomes. In contrast to perfectly homogeneous fuel stations, having only two rivaling stations may not be enough to achieve perfect competition, even with full information. With imperfect information, this effect works in the opposite direction to the increase in the pass-through to the average price observed when $N > 2$. Hence, the pass-through peak may occur at a higher number of competitors than $N = 2$.

Prediction 3. *The relationship between the number of competitors and pass-through to the average posted price is non-monotonic.*

4 Policy Changes and Descriptive Evidence

We analyze multiple tax changes in the German and French retail fuel markets from 2020 to 2023, to verify whether pass-through can be explained by competition under imperfect consumer information. We first provide an overview of the tax changes and then present descriptive evidence on the pass-through of these interventions.

4.1 Tax changes in the retail fuel market

Taxes account for the largest share of fuel prices in Germany and France. In 2019, a lump-sum energy tax of 65.45 Eurocent per liter was levied on gasoline and 47.04 Eurocent per liter on diesel in Germany. In France, the lump-sum fuel tax varied by region, ranging from 67 to 70 Eurocent per liter for gasoline (and around 61 Eurocent per liter for diesel). In addition, Germany and France have a value-added tax of 19% and 20%, respectively, that is levied on fuel tax-inclusive price of fuel.

The retail fuel markets in Germany and France experienced several tax changes between 2020 and 2023. These changes include a temporary VAT reduction in Germany to combat the economic impact of the Covid-19 pandemic, the introduction of a carbon tax in the German fuel market, and temporary reductions in the energy tax in both countries in 2022/23 to address price increases resulting from the Russian invasion of Ukraine.

The first tax change was a temporary reduction of the value-added tax in Germany from 19% to 16% between July and December 2020. On 1 January 2021, at the same time as the VAT was raised back to 19%, the German Federal Government also introduced a carbon price of 25 Euro per emitted tonne of CO₂ on oil, gas, and fuel. For *E5* and *E10*, this translates into a per unit tax of 6.00 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).

We cannot separately identify the pass-through of the simultaneous increase in the VAT and the introduction of the carbon emissions price in Germany on 1 January 2021. Instead, we jointly estimate their pass-through rate. This does not raise concerns regarding the theoretical predictions, since we show that the mechanisms that determine pass-through of an ad valorem tax and a per unit tax are the same.

In 2022, several tax changes occurred as a response to the Russian invasion of Ukraine and the resulting surge in energy prices. In France, between 1 April and 31 August 2022, there was a decrease in the fuel tax on gasoline and diesel of 18 Eurocent per liter. This rebate increased to 30 Eurocent between 1 September and 15 November 2022. It then dropped to 10 Eurocent between 16 November and 31 December 2022, before being completely phased out on 1 January 2023. Instead, the government introduced a lump-sum transfer to poorer households depending on the use of their car to commute to work.

Germany also implemented a temporary tax rebate, but this tax change is not studied in our analysis due to intense public scrutiny and a concurrent market investigation by the Federal Cartel Office. To appease the public, the vertically integrated oligopolists heavily advertised that they would pass through the tax change fully and quickly.²²

4.2 Descriptive evidence on heterogeneous pass-through

Before turning to the econometric analysis, we descriptively study the pass-through of the 2020 temporary VAT reduction in Germany, by comparing fuel price trends in Germany and France. This allows us to observe whether the pass-through differs between markets with a higher share of informed consumers (diesel) and those with fewer informed consumers (*E5*).

Panel A of Figure 3 displays non-parametric estimates of the VAT pass-through rate by fuel type in Germany during the 2020 temporary VAT reduction. Prices before the tax reduction evolve similarly for the three fuel types, suggesting that post-reduction differences in pass-through rates are not driven by pre-trends. Pass-through rates are highest for diesel and lowest for *E5*, consistent with the theoretical prediction that pass-through is higher when more price-sensitive consumers are present in the market. Pass-through is relatively fast and stabilizes after about two weeks.

Panel B of Figure 3 presents non-parametric estimates of the pass-through rate by fuel type for the winter 2020/21 tax increase. Unlike in the case of the tax decrease, there are anticipatory effects in passing through the tax increases in the last two weeks of December. We therefore drop the second half of December 2020 from the econometric analysis. The sharp increase in the implied pass-through rate around 1 January 2021 stabilizes afterwards. Pass-

²²In Appendix C, we present additional descriptive evidence showing that the 2022 tax changes in Germany are not suitable for our analysis.

Figure 3: Price change as share of total tax change



Notes: The Figure depicts the price change as a share of the total tax change for the tax decrease in July 2020 and the tax increase in January 2021 in panels A and B, respectively. The solid line shows the non-parametric 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 price in Germany (France) from the daily average price in Germany (France). The pre-period is from 1 May until 30 June 2020 for the tax decrease (panel A) and from 1 November until 15 December 2020 for the tax increase (panel B). Next, we compute the difference between demeaned average prices in Germany and France. Finally, we divide this difference by the difference under full pass-through. For the tax decrease, full pass-through would correspond to a price drop by 2.52%. Using average absolute prices from 24 June until 30 June (i.e., in the week prior to the tax change), this translates to a price decrease by 3.24 Eurocent for *E5*, 3.15 Eurocent for *E10*, and 2.72 Eurocent for diesel under full pass-through. For the tax increase, full pass-through would correspond to a price increase by 2.59% due to the VAT increase, plus the newly introduced carbon price. Using absolute prices in the week from 9 December until 15 December 2020 (i.e., in the week prior to the appearance of anticipatory effects), this translates to a price increase by 10.37 Eurocent for *E5*, 10.24 Eurocent for *E10*, and 10.75 Eurocent for diesel under full pass-through. The vertical solid line marks the starting date of the tax change. The horizontal dashed line indicates full pass-through.

through is highest for diesel, which is consistent with the theoretical predictions. Differences in pass-through between *E5* and *E10* appear less pronounced.

5 Empirical Strategy

Next, we estimate pass-through rates of the different tax changes separately by fuel type using a difference-in-differences (DID) strategy.

5.1 Difference-in-differences strategy

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. We use French fuel prices as the control group to estimate pass-through of the 2020/21 tax changes in Germany.²³ The treatment effect is the change in the difference between average fuel prices in Germany and France between pre- and post-treatment periods.

For our baseline pass-through estimation, we estimate the following DID model:

$$Y_{it} = \beta \text{Tax}_{it} + \gamma X_{it} + \alpha_i + \pi_t + \epsilon_{it}, \quad (1)$$

where Y_{it} is the logarithm of the weighted average price of gasoline or diesel at a fuel station i at date t . 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 onward. For the analysis of the subsequent tax increase, these are fuel stations in Germany from 1 January 2021 onward. For the analysis of the French tax changes in 2022/23, Tax_{it} equals one for French stations in the respective post-treatment periods. X_{it} is a vector of controls, which in the baseline specification only includes an interaction term of the crude oil price with an indicator for stations in Germany. This allows for differential pass-through of the crude oil price in Germany and France. Finally, the variables α_i and π_t correspond to fuel station and date fixed effects, respectively. We cluster standard errors at the market level.

5.2 Stations in neighboring country as a control group

Two assumptions must be met to identify the impact of the tax changes on fuel prices. First, there should be no temporary shocks that differentially affect fuel stations in Germany and

²³Conversely, we use German stations as the control group to estimate pass-through of the 2022/23 tax changes in France. For simplicity, we explain the DID strategy for the baseline tax changes in Germany.

France before and after the tax change, other than the policy change itself. Second, there should be no spillover effects from the tax changes onto the fuel market in the neighboring country. Both of these assumptions are likely to have been satisfied for the tax changes in 2020/21, but this is less likely for the tax changes in 2022/23.

Station fixed effects account for time-invariant differences between fuel stations in Germany and France, while date fixed effects control for transitory shocks that identically affect German and French stations. The two countries are similar in their geographic location, size, and wealth, and we restrict our analysis to relatively narrow time windows around the reforms, which should alleviate concerns about transitory shocks differently affecting German and French fuel stations in 2020/21.

To strengthen our claim that the effects are not influenced by transitory shocks, we consider the most obvious threats to identification. Public and school holidays in Germany and France are highly correlated, and travel restrictions due to the Covid-19 pandemic were lifted simultaneously in both countries, and the rest of the Schengen Area, from 15 June 2020. As most holidaymakers within Europe typically travel across several EU countries, and France and Germany are popular travel destinations in close proximity, it is likely that demand shocks affected fuel stations in both countries similarly.

Transitory supply shocks should also affect German and French fuel stations in a similar way. Due to their geographic proximity, fuel stations in Germany and France procure most of their refined oil from similar sources. 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 concerning the fuel market were implemented in Germany and France during our analysis period other than the tax changes discussed in Section 4. In general, there is no fuel price-setting regulation in Germany and France, and both countries have mandatory disclosure of fuel prices, which reaffirms our choice of France as a suitable control group.

Focusing on tax changes in opposite directions reduces concerns about confounding factors driving our results. If we find similar heterogeneities in pass-through for the tax increase in January 2021 as for the tax decrease in July 2020, a transitory shock confounding our estimates in July 2020 would also have to be present in January 2021 but in the opposite direction. For instance, if we overestimated the diesel pass-through rate in July 2020 because of a positive demand shock in France, then overestimating pass-through for diesel in January 2021 would require France to experience a negative demand shock. This scenario is unlikely, and finding consistent heterogeneities between the two tax changes suggests we are robustly estimating actual differences in pass-through.

In 2022/23, the picture is different. First, there are multiple tax changes that occurred in Germany and France, sometimes simultaneously, making it impossible to identify these separately. Second, the tax changes are so large that they may change the opportunity cost of selling fuel in the other country, leading to spillover effects and breaking the stable unit treatment value assumption (SUTVA). Third, gasoline and diesel markets were hit differently by Russia’s invasion of Ukraine, as diesel is a close substitute for heating oil whereas gasoline is not. Fourth, Germany and France were affected differently by the shocks on the global oil market in 2022.²⁴ As a consequence, we should refrain from interpreting the magnitude of pass-through for the 2022 tax changes, as well as differences between fuel types. Instead, analyzing these tax changes can be helpful to understand the difference in pass-through to the average posted price and the minimum price within a given fuel type.

5.3 Testing the theoretical predictions empirically

The aim of our empirical exercise is to test the theoretical predictions in Section 3 empirically. To test Prediction 1, we estimate the baseline DID model in Equation (1) separately for the three different fuel types for each of the 2020/21 tax changes in Germany. We then compare the pass-through rates across fuel types. According to the theoretical model and the specificities of the industry, we expect pass-through to be highest for diesel and lowest for *E5*. This should be the case for the tax decrease, as well as the increase.

To test Prediction 2, we estimate tax pass-through to the average price and the minimum price in a market for the 2020/21 tax changes, as well as the 2022/23 French tax changes. To remain as close as possible to the expected price and the expected minimum price in the theoretical model, we compute pass-through rates for the average posted price and the minimum price within non-overlapping geographic markets in Germany and France. We then estimate the following triple-DID variation of our baseline model:

$$Y_{jts} = \beta_1 \text{Tax}_{jt} + \beta_2 \text{Tax}_{jt} \cdot \text{Min}_s + \alpha_{js} + \pi_{ts} + \epsilon_{jts}, \quad (2)$$

where Y_{jts} is the logarithm of the price of gasoline or diesel in market j at date t . Note that the regression is now at the market level (instead of the station level). The market-level price is either the average posted price or the minimum price, as indicated by the “price type” s . Accordingly, Min_s is an indicator that equals one when the price type is the minimum price. As before, Tax_{jt} is a dummy variable that equals one for markets affected by the tax change

²⁴We present evidence in support of these points in Appendix C. We show that the diesel and gasoline markets in Germany and France started developing differently right after the start of Russia’s invasion and before any tax change was announced. We also show that margins increased in France immediately after Germany introduced a large fuel tax cut on 1 June 2022, suggesting spillover effects.

at date t . The variables α_{js} and π_{ts} correspond to market \times price type and date \times price type fixed effects. In this model, β_1 allows us to derive the average pass-through rate to the average posted price (i.e., when $\text{Min}_s = 0$). In contrast, β_2 allows us to derive the difference in the pass-through rate between the minimum price and average posted price in a local market. As the theoretical model predicts pass-through to the minimum price to be higher than to the average price, we expect $\beta_2 > 0$ when estimating the parameters of the model in Equation (2). This prediction holds for all fuel types and for all tax changes.

To test Prediction 3, we estimate tax pass-through for the 2020/21 tax changes at the station level. An important feature of our setting is that we can do this comparison within fuel type and thus 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 and fuel type, we estimate the model in Equation (1) adding an interaction term between the treatment period and the station fixed effect. 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 competing price setters in a market and calculate the average pass-through rate for each group. In a perfect information model, we expect this relationship to be monotonically increasing. Instead, our model predicts a hump-shaped relationship between the number of competitors and average tax pass-through.

5.4 Robustness checks

We run several additional analyses to verify that our empirical results are robust to alternative model specifications.

In Appendix D.1, we present DID estimates of the baseline pass-through rates when we include additional control variables into our regression model. In particular, we directly account for demand-related shocks by including regional information on the daily mobility to work and to retail and recreational places from the Google Mobility Report. Our results are robust to including these control variables.

In Appendix D.2, we further show that our results are robust to estimating the baseline DID model on a restricted (balanced) sample of fuel stations for which we have a price observation on every day. This allows us to rule out that our findings are driven by the temporary closure (e.g., on weekends or holidays) of some fuel stations.

Based on the descriptive evidence in Figure 3, our preferred specification is to account for anticipatory effects in winter (tax increase) but not in summer (tax decrease). In Appendix D.3, we show that our main empirical findings are robust to changing this as-

sumption. In Appendix B.5, we also provide a brief theoretical discussion for the emergence of anticipatory price increases before a tax increase and a tax decrease.

In Appendix D.4, we employ a synthetic difference-in-differences (SDID) strategy (Arkhangelsky et al., 2021). SDID is a variation of DID that aims to match pre-treatment trends between the treatment and control groups using weights. Our results are robust to using this alternative empirical strategy.

Finally, in Appendix D.5, we assess the robustness of our results regarding the differential pass-through to the average and the minimum price. One potential concern is that the minimum price in a local market may reflect outlier prices that are only valid for a short period of time. We address this concern by replacing the minimum price with the price at the 10th percentile of the distribution of all hourly prices across all stations in a market on a given day. We find that our results on the difference in pass-through between the minimum and the average price are robust to this specification and hence are not driven by outliers.

6 Results

The following section presents the results from the empirical estimation.

6.1 Consumer information and tax pass-through

Table 3 presents the estimated average treatment effect of the 2020/21 tax changes on fuel prices for *E5*, *E10*, and diesel, using the DID model in Equation (1). The outcome variable in all columns is the logarithm of price, including taxes and duties, for a given station and date.

Columns (1) to (3) show the effect of the tax decrease. The tax reduction caused prices for all fuel types to decrease. Under full pass-through, we expect prices for each fuel product to decrease by about 2.52%.²⁵ We estimate that 95% of the tax decrease is passed on to diesel consumers, while the pass-through rates for *E10* and *E5* is 46% and 28%, respectively.

Columns (4) to (6) show that the tax increase raised prices for all fuel products. Under full pass-through, we expect an increase in prices by 8.31% for *E5*, 8.54% for *E10*, and 9.97%

²⁵With a decrease in the VAT rate from 19% before the VAT decrease to 16% after the policy change, this is $\frac{1.16-1.19}{1.19} * 100 \approx -2.52\%$.

Table 3: Effect of the tax change on log prices

	Tax decrease			Tax increase		
	<i>E5</i>	<i>E10</i>	Diesel	<i>E5</i>	<i>E10</i>	Diesel
	(1)	(2)	(3)	(4)	(5)	(6)
Tax change	-0.0071*** (0.0004)	-0.0116*** (0.0004)	-0.0239*** (0.0004)	0.0561*** (0.0004)	0.0611*** (0.0003)	0.0834*** (0.0003)
Pass-through rate	28% [25%, 31%]	46% [43%, 49%]	95% [92%, 97%]	68% [67%, 68%]	72% [71%, 72%]	84% [83%, 84%]
Date fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Station fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
DE \times oil price	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,133,377	2,324,131	2,703,604	1,804,703	1,988,649	2,318,672

Notes: The Table presents DID estimates using the model in Equation (1). Columns (1) to (3) present average treatment effect estimates of the German VAT reduction on 1 July 2020 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 CO₂ emissions tax on 1 January 2021 on *E5*, *E10*, and diesel log prices, respectively. Columns (4) to (6) use data from 1 November to 15 December 2020 for the pre-treatment period, and from 1 January to 28 February 2021 for the post-treatment period. Standard errors clustered at the market level are shown in parentheses. We also compute the pass-through rates corresponding to the point estimates and report their 95% confidence intervals in brackets.

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

for diesel.²⁶ We find a joint pass-through rate of the tax increases of 68% for *E5*, 72% for *E10*, and 84% for diesel.


Overall, our findings are consistent with Prediction 1 that the pass-through rate is higher when there are more price sensitive consumers. For both tax changes, pass-through is significantly higher for diesel than for gasoline. Within gasoline, pass-through is significantly higher for *E10* than for *E5*. Thus, the ordering of all the point estimates in Table 3 is consistent with our prediction. Since we observe that almost all fuel stations in Germany sell all three types of fuel, the differences in the pass-through rates cannot be explained by supply-side factors. In Section 7, we discuss why differences in the price elasticity of aggregate demand between fuel types are unlikely to explain the results.

6.2 Pass-through to the average and minimum price

Figure 4 shows the difference in the pass-through rate between the minimum and the average price for different tax changes in the German and French retail fuel markets between July 2020 and January 2023. As previously noted, the 2022 tax changes are inadequate for

²⁶Under full pass-through, a change in the VAT rate from 16% to 19% would increase the fuel price by $\frac{1.19-1.16}{1.16} * 100 \approx 2.59\%$. To estimate by what percentage the fuel price would increase if the carbon emissions price was fully passed through, we divide the gross price per liter on carbon emissions for each fuel type by the average fuel price in Germany in the week from 9 December until 15 December 2020 (i.e., before we start seeing anticipatory effects).

Figure 4: Pass-through to market-level minimum vs. average prices



Notes: The Figure shows the difference in the market-level pass-through rate (in percentage points) between the minimum price and the average posted price. The average posted price is the average daily price within a non-overlapping market by weighting the price at every full hour of the day between 6 am and 10 pm equally. The minimum price is the minimum price within a non-overlapping market at any point of time during the day. The Figure depicts the pass-through rates implied by the DID estimate β_2 in Equation (2) along with 99% confidence intervals, based on standard errors clustered at the market level. Regressions are estimates separately for each tax change and fuel type. For most tax changes, we use data for the two months before and after every tax change. Exceptions include the German tax increase on 1 January 2021, where we exclude the second half of December to account for anticipatory effects. For the tax increase in France on 16 November 2022, we only use the period until 31 December 2022 as post-treatment period. Similarly, for the French tax increase on 1 January 2023, we use the period from 16 November until 31 December 2022 as pre-treatment period. Solid (hollow) symbols indicate statistically significant (insignificant) point estimates at the 1% level.

measuring the relative and absolute pass-through between different fuel types. Nonetheless, they allow us to compare pass-through to the minimum and average posted prices within a particular fuel type.

For the tax increases, pass-through rates to the minimum price are always significantly higher than those to the average posted price. This supports our theoretical prediction. For the tax decrease, this is true in two cases, while there are three cases where the difference in pass-through rates is statistically indistinguishable from zero and one case where pass-through to the average posted price is slightly higher. Overall, we observe a statistically significantly higher pass-through rate to the minimum price than to the average price in 11 out of the 15 cases depicted in Figure 4.

The findings presented in Figure 4 support Prediction 2 that the pass-through to the expected minimum price is higher than to the expected price. Informed consumers, who typically buy fuel at prices closer to the within-market minimum, bear more of the cost of a tax increase (and gain more from a tax cut) than uninformed consumers, who buy fuel at the average posted price.

6.3 Number of sellers and tax pass-through

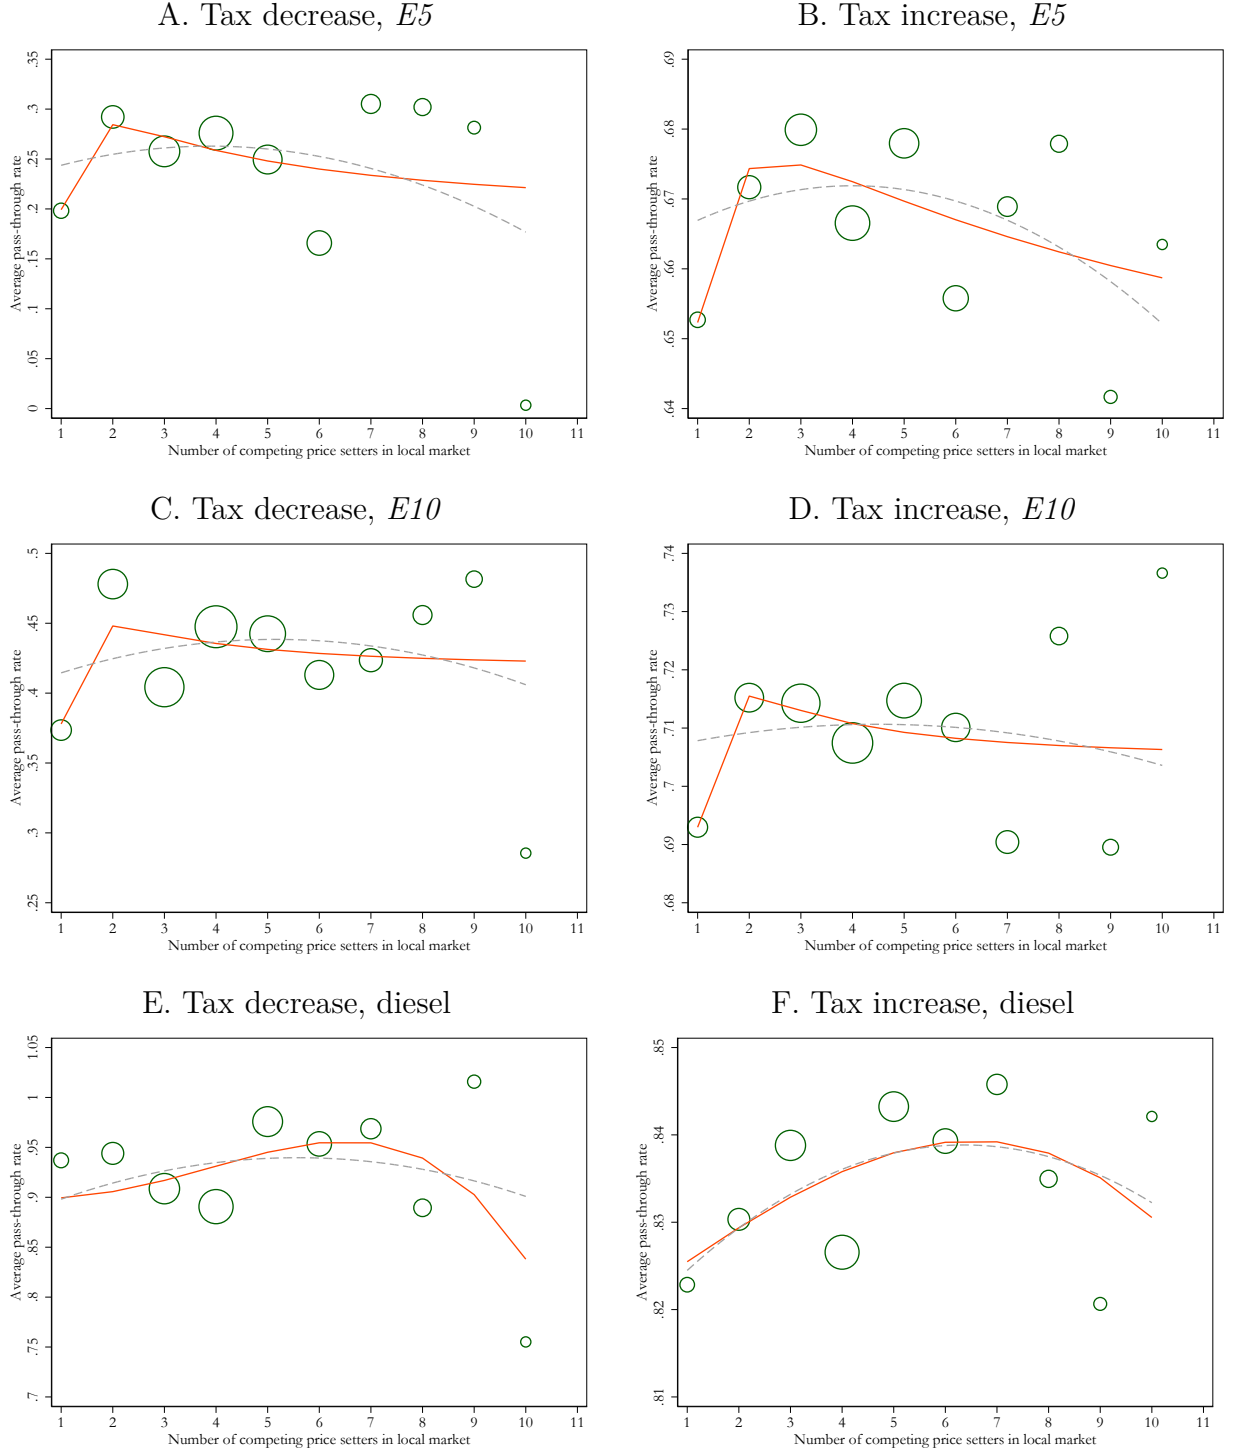
Finally, we study how the pass-through rate varies with the number of sellers in the market. Figure 5 shows the relationship between the pass-through rate and the number of competitors of a focal station for the 2020/21 German tax changes and the three fuel types. Each circle corresponds to the average pass-through rate for stations with a particular number of competing price setters within a non-overlapping local market. The size of a circle is proportional to the total number of stations with a given number of competitors. We also plot the curves of a fractional polynomial fit as well as a quadratic fit.

Panels A, C, and E depict the pass-through rates for the tax decrease in summer 2020 for *E5*, *E10*, and diesel, respectively. Panel A shows that the average pass-through rate for *E5* is relatively low for local monopolists. It is higher for markets with two competing price setters and then tends to decrease in the number of competitors. We observe a similar non-monotonic relationship between the number of sellers and the average pass-through rate for *E10*. This pattern looks strikingly similar to the numerically simulated patterns in Figure 2. For diesel, the relationship is flatter for smaller markets, but then the average pass-through rate also declines with the number of competitors.

In panels B, D, and F of Figure 5, we repeat this analysis for the tax increase in winter 2020/21. For all fuel types, we find similar relationships as for the tax decrease. For *E5* and *E10*, pass-through is again relatively low for local monopolists and tends to decrease in the number of competitors when there are at least two competing price setters. For diesel, the pass-through rate is mildly increasing up to around six or seven competing price setters and then decreases in the number of sellers.

Overall, the results in Figure 5 indicate that pass-through to the average price is not monotonically increasing in the number of sellers, in line with our theoretical prediction under imperfect information. The fractional polynomial fits for *E5* and *E10* closely resemble our simulations in Figure 2 with a peak at $N = 2$. For diesel, the relationship between the number of sellers and the average pass-through rate has an inverted-U shape with a peak at a higher number of sellers than in the case of *E5* and *E10*. The different pattern between diesel and gasoline may suggest that if pass-through is already higher on average, the number

Figure 5: Average pass-through by number of competitors



Notes: The Figure shows how the pass-through rate to the average price varies with the number of competing price setters in a market. Panels A, C, and E depict the pass-through rates for the German VAT decrease on 1 July 2020 for $E5$, $E10$, and diesel, respectively. Panels B, D, and F depict the pass-through rates for the German VAT increase and introduction of a carbon price on 1 January 2021 for $E5$, $E10$, and diesel, respectively. In every panel, each circle plots the average pass-through rate for a group of stations with a particular number of competing price setters within a non-overlapping local market. The size of a circle is proportional to the total number of stations with a given number of competitors. The solid line shows a fractional polynomial fit. The dashed line shows a quadratic fit. The number of competitor stations is trimmed at the 97.5th percentile.

of sellers may have less of an impact on pass-through rates than if pass-through is at a lower level.

In Appendix D.6, we formally test for non-monotonicity by applying the U-test by Lind and Mehlum (2010). This allows us to test the null hypothesis of a monotone or U-shaped relationship against the alternative hypothesis of an inverted-U shape. We reject the null hypothesis at the 10% level for both tax changes for diesel and *E5*. For *E10*, the estimates still speak in favor of a hump-shaped relationship, but we cannot formally reject that the pass-through rate is increasing in the number of competitors. Overall, we conclude that our empirical findings are consistent with Prediction 3 that the relationship between the number of sellers and pass-through to the expected price is non-monotonic.

7 Alternative Hypotheses

While our analysis shows that the imperfect consumer information model proposed by Stahl (1989) effectively accounts for the relationship between competition and tax pass-through that we observe empirically, this does not prove that there is no observationally equivalent alternative explanation for the empirical findings. In the following, we collect alternative candidate explanations and discuss how well they can explain the empirical findings.

7.1 Alternative hypotheses with full information

Before discussing alternative hypotheses, it is worth briefly reviewing the full information conduct parameter approach, which is a widely used approach of modeling imperfect competition in the tax pass-through literature. A natural way of modeling competition in the retail fuel market with full information is as symmetrically differentiated Nash-in-prices. Sellers have a symmetric cost structure, offer a homogeneous good, and are located in different places, with pricing as their primary decision variable. This is a special case of the analysis presented in Weyl and Fabinger (2013).

Most of the literature that uses pass-through as a sufficient statistic for welfare results assumes that markets are perfectly competitive (Weyl and Fabinger, 2013).²⁷ To extend this analysis to oligopolistic markets, Weyl and Fabinger (2013) use the conduct parameter approach, which encompasses most models of oligopolistic competition with symmetric sellers and perfect information.²⁸ Additionally, Weyl and Fabinger (2013) extend their analysis to

²⁷Sumner (1981) notes that the price elasticity of residual demand is a critical factor in tax pass-through for oligopolistic markets. Bulow and Pfleiderer (1983) demonstrate that the degree of pass-through depends on the functional form of demand, and by measuring it, the curvature of demand can be determined.

²⁸This was first introduced by Bresnahan (1989) and Genesove and Mullin (1998).

cases of asymmetric competition, including homogeneous product oligopoly, differentiated Nash-in-prices, and monopolistic competition with perfect information.

The conduct parameter approach features a parameter θ , which varies between 0 for perfect competition, 1 for monopoly, and $1/N$ for Cournot competition with N symmetric competitors. For full information models with symmetric competitors, including symmetrically differentiated Nash-in-prices, pass-through of a per-unit tax can be expressed as

$$\rho = \frac{1}{1 + \frac{\theta}{\epsilon_\theta} + \frac{\epsilon_D - \theta}{\epsilon_S} + \frac{\theta}{\epsilon_{ms}}} , \quad (3)$$

where ϵ_D is the price elasticity of aggregate demand, ϵ_S is the price elasticity of supply, and ϵ_{ms} is the price elasticity of marginal surplus, i.e., the curvature of demand.

Genakos and Pagliero (2022) argue that for the retail fuel market it is reasonable to assume that marginal cost are constant and that the conduct parameter does not vary with quantities. In case of the former, $\frac{\epsilon_D - \theta}{\epsilon_S} = 0$. In case of the latter, $\frac{\theta}{\epsilon_\theta} = 0$. In case of symmetric competition with constant marginal costs and conduct invariant to quantities, pass-through therefore simplifies to

$$\rho = \frac{1}{1 + \frac{\theta}{\epsilon_{ms}}} . \quad (4)$$

The relationship between pass-through and competition is unclear since the curvature of demand may either increase or decrease with competition.

Next, we discuss alternative hypotheses of what could cause differences in pass-through rates for retail fuel, which do not rely on imperfectly informed consumers. Some of these are nested by the conduct parameter approach, whereas others are not. We contrast their predictions to the empirical evidence.

The curvature of demand. Analyzing the pass-through equation used by Genakos and Pagliero (2022), the simplest conclusion would be that differences in pass-through between different fuel types are based on differences in the curvature of aggregate demand. Although this *may* explain differences in pass-through between fuel types, it explains neither differences in pass-through between the average and the minimum price in a market nor a non-monotonic relationship between the number of competitors and pass-through. In fact, holding the curvature of demand fixed and assuming that θ decreases in the number of competitors, pass-through should increase monotonically with the number of competitors. Furthermore, with symmetric firms and full information, there should not be any within-market price dispersion.

Vertical differentiation. One reason for within-market differences in prices and pass-through that is unrelated to imperfect information could be vertical differentiation.

One could argue that some stations are just better than others (e.g., because of their location). These stations may face a lower price elasticity of residual demand even under perfect information and therefore also have lower pass-through. Together with a mechanism of why we should expect differences in pass-through between fuel types, this *may* explain some of the empirical patterns. However, this would not explain the non-monotonic relationship between the number of competitors and pass-through, as well as the existence of random price dispersion. Most importantly, a story of vertical differentiation would require that the ranking of stations by price within a market is mostly unchanged. This is at odds with what we observe empirically.

The price elasticity of aggregate demand. Similar to the curvature of demand, we can begin by analyzing it through the lens of the symmetric competition conduct parameter approach. If we relax the assumption that marginal cost is constant and allow for the price elasticity of supply to be positive but below infinity, all other things being equal, a higher price elasticity of aggregate demand leads to a lower pass-through rate. If the price elasticity of aggregate demand is lowest for diesel and highest for *E5*, this *may* explain differences in pass-through between fuel types without requiring imperfect information. However, like with the curvature of demand, it does not explain differences in pass-through between the average and the minimum price, random price dispersion, or a non-monotonic relationship between the number of competitors and pass-through.

Price sensitivity unrelated to information. An alternative approach would be to let consumers differ in how strongly they react to lower prices (i.e., their price sensitivity) whilst assuming that they are perfectly informed. Even with perfect information, consumers may only react to a lower price if it is worth their time (e.g., because the lower-price station is further away). Purchasers of different fuel types may differ in their price sensitivity, which may affect competitive conduct. This could be modeled, for example, as the price coefficient in a logit demand model. Combined with firms competing in prices, the higher the price sensitivity of consumers, the higher is the intensity of competition and the higher is tax pass-through. Without imperfect information, however, frequent changes in the ranking of stations by price, random price dispersion, and a non-monotonic relationship between the number of competitors and pass-through remain to be explained.

Overall, although the empirical findings can often partially be explained by alternative hypotheses that do not rely on imperfect information, the results in their entirety are inconsistent with any of these alternative hypotheses.

7.2 Alternative hypotheses with imperfect information

Another set of alternative hypotheses relies on modeling competition under imperfect information differently. We briefly discuss what we consider to be the most obvious alternative modeling choices.

Rockets and feathers. Tappata (2009) proposes a dynamic model to explain the “rockets and feathers” phenomenon – prices rising faster than they fall – via consumer search and cost uncertainty. In this model, atomistic consumers have unit demand and value the good at v . Consumers have the option to purchase access to an information clearinghouse, which allows them to observe all market prices, or they can choose to draw a single price at random. Some consumers have zero access costs and are always perfectly informed, while others draw access costs from a continuous distribution. This model is a variant of Varian (1980) in which the decision to become informed is endogenized.

Marginal costs in this model can be high or low, and they follow a first-order Markov process. Firms use mixed strategies to set prices. When production costs are high, the gap between marginal cost and v is narrow, resulting in low price dispersion and limited search gains. Conversely, if production costs are low, the gap between marginal cost and v is large, leading to high price dispersion and greater search gains. As cost decreases are possible only when marginal costs are high, they occur during periods of low search, resulting in slow pass-through. In contrast, cost increases are quickly passed through when marginal costs are low, prices are low, and search is high.

For this mechanism to explain the empirical findings, price dispersion, a measure for the gains from search, needs to be higher when prices are low. However, in Appendix A.4, we show that this is not the case. Since the basic premise of the model does not apply in our empirical application, modeling competition under imperfect information in this way is unlikely to be better at explaining the empirical findings.

Reference price dependence. In a duopoly market where costs are unobservable to consumers, Lewis (2011) finds that whether consumers search depends on how the first price they draw compares to a reference price (e.g., the previous period’s price). A positive cost shock increases the probability of the first price exceeding the reference price, inducing more search and higher pass-through. Conversely, a negative cost shock reduces search and lowers pass-through.

The model’s main finding is that pass-through is faster for cost increases and slower for cost decreases. However, there are several drawbacks to analyzing our empirical application through the lens of this model. First, since it is a duopoly model, it does not allow analyzing the relationship between pass-through and the number of competitors. Second, the consumer search protocol is suboptimal, as it is unrelated to the *actual* gains from search, and only

applicable if cost shocks – tax changes in our application – are unobservable and a surprise to consumers. Finally, unlike the German fuel market, where price cycles occur intra-day and are unrelated to cost, price cycles result from cost shocks, similar to Tappata (2009).

Unobservable input costs. Whereas we treat input costs (i.e., mainly the price of oil and tax rates in our case) as observable to consumers, an alternative would be to treat these costs as unobservable. Janssen, Pichler, and Weidenholzer (2011) extend the Stahl (1989) model to include unobservable input prices, while assuming an exogenous share of shoppers. Janssen and Shelegia (2015) extend this to a vertical market, where an upstream manufacturer sets the input price. They find that lower sequential search costs for non-shoppers lead to less elastic upstream demand, incentivizing the manufacturer to reduce downstream retailer profits and resulting in higher prices compared to a vertically integrated monopolist.²⁹

Our empirical application differs from this setting in several ways. First, vertical integration is prevalent in the industry. Second, upstream prices, represented by the oil price, are more transparent than in other industries. Even if the oil price is not observed by consumers on a daily basis, the cost shocks analyzed in this paper (i.e., the sizeable tax changes) were broadly publicized and salient for consumers. Also the timing, direction, and magnitude of the tax changes were well-known. It therefore seems unlikely that treating input costs as unobservable would increase the explanatory power of the model.

Although one can make the case for modeling imperfect information in the retail fuel market differently, none of the most obvious alternative ways of modeling competition appear better suited to our empirical setting. Therefore, we conclude that our empirical findings in their entirety are best explained by the imperfect consumer information model by Stahl (1989).

8 Conclusion

In this paper, we highlight the role of imperfect information in explaining heterogeneities in tax pass-through. We show that when consumers do not know all prices, firms have market power and this affects tax pass-through. While our approach imposes more structure on conduct and demand than Weyl and Fabinger (2013), this allows us to be more flexible in modeling consumer information.

²⁹Janssen, Pichler, and Weidenholzer (2011) and Janssen and Shelegia (2015) both find that in the presence of unobservable input costs or upstream prices, downstream prices are higher. However, neither has predictions about pass-through. Janssen and Shelegia (2020) study pass-through with imperfect information and differentiated products.

Three results stand out and set this apart from an analysis with perfect information: first, the more well-informed consumers there are, the higher is tax pass-through. Second, taxes (and tax cuts) are passed through more to the price paid by well-informed consumers than to the price paid by uninformed consumers. Third, there is no monotonic relationship between the number of sellers and pass-through.

The results of this study have important implications for policy. When markets are imperfectly competitive, pass-through affects the effectiveness of unconventional fiscal policies (see D’Acunto, Hoang, and Weber, 2018 or D’Acunto, Hoang, and Weber, 2022), how strongly prices react to Pigouvian taxes, and the distributional consequences of such policies. These considerations, as well as uncertainty about the level of consumer information, can affect the relative benefits of regulation versus Pigouvian taxes or subsidies.

We shed light on a novel explanation of what determines tax pass-through, relevant to any market where it is costly for consumers to learn about prices. Future research should aim to develop tools to estimate the parameters of demand and supply models with imperfectly informed consumers. This would allow estimating how a proposed tax or subsidy would affect prices and quantities, even in the presence of firms with market power and consumers that are imperfectly informed about prices.

References

- Adachi, Takanori, and Michal Fabinger. 2022. “Pass-Through, Welfare, and Incidence Under Imperfect Competition.” *Journal of Public Economics* 211:1045–1089.
- Arkhangelsky, Dmitry, Susan Athey, David A. Hirshberg, Guido W. Imbens, and Stefan Wager. 2021. “Synthetic Difference-in-Differences.” *American Economic Review* 111 (12): 4088–4118.
- Assad, Stephanie, Robert Clark, Daniel Ershov, and Lei Xu. 2020. “Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market.” *CESifo Working Paper No. 8521*.
- Baye, Michael R., John Morgan, and Patrick Scholten. 2006. “Information, Search, and Price Dispersion.” Chap. 6 in *Handbook on Economics and Information Systems*, edited by Terrence Hendershott, 323–376. Elsevier.
- Benedek, Dora, Ruud A. De Mooij, Michael Keen, and Philippe Wingender. 2020. “Varieties of VAT Pass Through.” *International Tax and Public Finance* 27 (4): 890–930.
- Benzarti, Youssef, Dorian Carloni, Jarkko Harju, and Tuomas Kosonen. 2020. “What Goes Up May Not Come Down: Asymmetric Incidence of Value-Added Taxes.” *Journal of Political Economy* 128 (12): 4438–4474.
- Borenstein, Severin, A. Colin Cameron, and Richard Gilbert. 1997. “Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?” *Quarterly Journal of Economics* 112 (1): 305–339.
- Bresnahan, Timothy F. 1989. “Empirical Studies of Industries with Market Power.” Chap. 17 in *Handbook of Industrial Organization*, 2:1011–1057. Elsevier.
- Bulow, Jeremy I., and Paul Pfleiderer. 1983. “A Note on the Effect of Cost Changes on Prices.” *Journal of Political Economy* 91 (1): 182–185.
- Burdett, Kenneth, and Kenneth L. Judd. 1983. “Equilibrium Price Dispersion.” *Econometrica* 51 (4): 955–969.
- Busse, Meghan, Jorge Silva-Risso, and Florian Zettelmeyer. 2006. “\$1,000 Cash Back: The Pass-Through of Auto Manufacturer Promotions.” *American Economic Review* 96 (4): 1253–1270.
- Büttner, Thiess, and Boryana Madzharova. 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): 103–134.
- Byrne, David P., and Leslie A. Martin. 2021. “Consumer Search and Income Inequality.” *International Journal of Industrial Organization* 79:102716.

- Byrne, David P., and Nicolas de Roos. 2017. "Consumer Search in Retail Gasoline Markets." *Journal of Industrial Economics* 65 (1): 183–193.
- Carranza, Juan Esteban, Robert Clark, and Jean-François Houde. 2015. "Price Controls and Market Structure: Evidence from Gasoline Retail Markets." *Journal of Industrial Economics* 63 (1): 152–198.
- Cason, Timothy N., Daniel Friedman, and Ed Hopkins. 2020. "Discrete Prices and the Incidence and Efficiency of Excise Taxes." *Journal of Political Economy* 129 (3): 790–841.
- Chandra, Ambarish, and Mariano Tappata. 2011. "Consumer Search and Dynamic Price Dispersion: An Application to Gasoline Markets." *RAND Journal of Economics* 42 (4): 681–704.
- Chetty, Raj. 2009. "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods." *Annual Review of Economics* 1 (1): 451–488.
- Chetty, Raj, Adam Looney, and Kory Kroft. 2009. "Salience and Taxation: Theory and Evidence." *American Economic Review* 99 (4): 1145–1177.
- Conlon, Christopher, and Nirupama Rao. 2023. "The Cost of Curbing Externalities with Market Power: Alcohol Regulations and Tax Alternatives." *NBER Working Paper No. 30896*.
- Conlon, Christopher, Nirupama Rao, and Yinan Wang. Forthcoming. "Who Pays Sin Taxes? Understanding the Overlapping Burdens of Corrective Taxes." *Review of Economics and Statistics*.
- Conlon, Christopher T., and Nirupama L. Rao. 2020. "Discrete Prices and the Incidence and Efficiency of Excise Taxes." *American Economic Journal: Economic Policy* 12 (4): 111–143.
- D’Acunto, Francesco, Daniel Hoang, and Michael Weber. 2018. "Unconventional Fiscal Policy." *AEA Papers and Proceedings* 108:519–523.
- . 2022. "Managing Households’ Expectations with Unconventional Policies." *Review of Financial Studies* 35 (4): 1597–1642.
- Deltas, George, and Michael Polemis. 2020. "Estimating Retail Gasoline Price Dynamics: The Effects of Sample Characteristics and Research Design." *Energy Economics* 92:1049–1076.
- Diamond, Peter A. 1971. "A Model of Price Adjustment." *Journal of Economic Theory* 3:156–168.
- Dubois, Pierre, Rachel Griffith, and Martin O’Connell. 2020. "How Well Targeted Are Soda Taxes?" *American Economic Review* 110 (11): 3661–3704.

- Duso, Tomaso, and Florian Szücs. 2017. “Market Power and Heterogeneous Pass-Through in German Electricity Retail.” *European Economic Review* 98:354–372.
- Eckert, Andrew. 2013. “Empirical Studies of Gasoline Retailing: A Guide to the Literature.” *Journal of Economic Surveys* 27 (1): 140–166.
- Eizenberg, Alon, Saul Lach, and Merav Oren-Yiftach. 2021. “Retail Prices in a City.” *American Economic Journal: Economic Policy* 13 (2): 175–206.
- Fabra, Natalia, and Mar Reguant. 2014. “Pass-Through of Emissions Costs in Electricity Markets.” *American Economic Review* 104 (9): 2872–2899.
- Fischer, Kai, Simon Martin, and Philipp Schmidt-Dengler. 2023. “The Heterogeneous Effects of Entry on Prices.” *CEPR Discussion Paper 18297*.
- Ganapati, Sharat, Joseph S. Shapiro, and Reed Walker. 2020. “Energy Cost Pass-Through in US Manufacturing: Estimates and Implications for Carbon Taxes.” *American Economic Journal: Applied Economics* 12 (2): 303–342.
- Genakos, Christos, and Mario Pagliero. 2022. “Competition and Pass-Through: Evidence from Isolated Markets.” *American Economic Journal: Applied Economics* 14 (4): 35–57.
- Genesove, David, and Wallace P. Mullin. 1998. “Testing Static Oligopoly Models: Conduct and Cost in the Sugar Industry, 1890–1914.” *RAND Journal of Economics* 29 (2): 355–377.
- Hamilton, Stephen F. 1999. “Tax Incidence under Oligopoly: A Comparison of Policy Approaches.” *Journal of Public Economics* 71 (2): 233–245.
- Harding, Matthew, Ephraim Leibtag, and Michael F. Lovenheim. 2012. “The Heterogeneous Geographic and Socioeconomic Incidence of Cigarette Taxes: Evidence from Nielsen Homescan Data.” *American Economic Journal: Economic Policy* 4 (4): 169–198.
- Harju, Jarkko, Tuomas Kosonen, Marita Laukkanen, and Kimmo Palanne. 2022. “The Heterogeneous Incidence of Fuel Carbon Taxes: Evidence from Station-Level Data.” *Journal of Environmental Economics and Management* 112:102607.
- Heim, Sven. 2021. “Asymmetric Cost Pass-Through and Consumer Search: Empirical Evidence from Online Platforms.” *Quantitative Marketing and Economics* 19 (2): 227–260.
- Hindriks, Jean, and Valerio Serse. 2019. “Heterogeneity in the Tax Pass-Through to Spirit Retail Prices: Evidence From Belgium.” *Journal of Public Economics* 176:142–160.
- Hollenbeck, Brett, and Kosuke Uetake. 2021. “Taxation and Market Power in the Legal Marijuana Industry.” *RAND Journal of Economics* 52 (3): 559–595.
- Holt, Timothy, Mitsuru Igami, and Simon Scheidegger. 2023. “Detecting Edgeworth Cycles.” *SSRN Electronic Journal*.

- Huber, Stephan, and Christoph Rust. 2016. "Calculate Travel Time and Distance with Open-StreetMap Data Using the Open Source Routing Machine (OSRM)." *Stata Journal* 16 (2): 416–423.
- Janssen, Maarten, Paul Pichler, and Simon Weidenholzer. 2011. "Oligopolistic Markets with Sequential Search and Production Cost Uncertainty." *RAND Journal of Economics* 42 (3): 444–470.
- Janssen, Maarten, and Sandro Shelegia. 2015. "Consumer Search and Double Marginalization." *American Economic Review* 105 (6): 1683–1710.
- . 2020. "Beliefs and Consumer Search in a Vertical Industry." *Journal of the European Economic Association* 18 (5): 2359–2393.
- Johnson, Ronald N. 2002. "Search Costs, Lags and Prices at the Pump." *Review of Industrial Organization* 20 (1): 33–50.
- Kopczuk, Wojciech, Justin Marion, Erich Muehlegger, and Joel Slemrod. 2016. "Does Tax-Collection Invariance Hold? Evasion and the Pass-Through of State Diesel Taxes." *American Economic Journal: Economic Policy* 8 (2): 251–286.
- Kosonen, Tuomas. 2015. "More and Cheaper Haircuts after VAT Cut? On the Efficiency and Incidence of Service Sector Consumption Taxes." *Journal of Public Economics* 131:87–100.
- Kraftfahrt-Bundesamt. 2021. "Der Fahrzeugbestand im Überblick am 1. Januar 2020." https://www.kba.de/DE/Statistik/Fahrzeuge/Bestand/Jahresbilanz__Bestand/2020/2020__b_ueberblick_.pdf?__blob=publicationFile&v=1.
- Kroft, Kory, Jean-William P Laliberté, René Leal-Vizcaino, and Matthew J Notowidigdo. 2023. "Salience and Taxation with Imperfect Competition." *Review of Economic Studies*.
- . 2021. "Efficiency and Incidence of Taxation with Free Entry and Love-of-Variety Preferences." *NBER Working Paper No. 28838*.
- Lewis, Matthew S. 2011. "Asymmetric Price Adjustment and Consumer Search: An Examination of the Retail Gasoline Market." *Journal of Economics & Management Strategy* 20 (2): 409–449.
- Li, Jing, and James H. Stock. 2019. "Cost Pass-Through to Higher Ethanol Blends at the Pump: Evidence from Minnesota Gas Station Data." *Journal of Environmental Economics and Management* 93:1–19.
- Lind, Jo Thori, and Halvor Mehlum. 2010. "With or Without U? The Appropriate Test for a U-Shaped Relationship." *Oxford Bulletin of Economics and Statistics* 72 (1): 109–118.
- Luco, Fernando. 2019. "Who Benefits from Information Disclosure? The Case of Retail Gasoline." *American Economic Journal: Microeconomics* 11 (2): 277–305.

- Marion, Justin, and Erich Muehlegger. 2011. "Fuel tax incidence and supply conditions." *Journal of Public Economics* 95:1202–1212.
- Martin, Simon. Forthcoming. "Market Transparency and Consumer Search - Evidence from the German Retail Gasoline Market." *RAND Journal of Economics*.
- Miller, Nathan H., Matthew Osborne, and Gloria Sheu. 2017. "Pass-through in a Concentrated Industry: Empirical Evidence and Regulatory Implications." *RAND Journal of Economics* 48 (1): 69–93.
- Miravete, Eugenio J., Katja Seim, and Jeff Thurk. 2018. "Market Power and the Laffer Curve." *Econometrica* 86 (5): 1651–1687.
- Montag, Felix, Alina Sagimuldina, and Christoph Winter. 2023. "Whom to Inform About Prices? Evidence from the German Gasoline Market." *Mimeo*.
- Nakamura, Emi, and Dawit Zerom. 2010. "Accounting for Incomplete Pass-Through." *Review of Economic Studies* 77 (3): 1192–1230.
- Pennerstorfer, Dieter, Philipp Schmidt-Dengler, Nicolas Schutz, Christoph Weiss, and Biliana Yontcheva. 2020. "Information and Price Dispersion: Theory and Evidence." *International Economic Review* 61 (2): 871–899.
- Stahl, Dale O. 1989. "Oligopolistic Pricing with Sequential Consumer Search." *American Economic Review* 79 (4): 700–712.
- Stern, Nicholas. 1987. "The Effects of Taxation, Price Control and Government Contracts in Oligopoly and Monopolistic Competition." *Journal of Public Economics* 32 (2): 133–158.
- Sumner, Daniel A. 1981. "Measurement of Monopoly Behavior: An Application to the Cigarette Industry." *Journal of Political Economy* 89 (5): 1010–1019.
- Tappata, Mariano. 2009. "Rockets and Feathers: Understanding Asymmetric Pricing." *RAND Journal of Economics* 40 (4): 673–687.
- Varian, Hal R. 1980. "A Model of Sales." *American Economic Review* 70 (4): 651–659.
- Weyl, E. Glen, and Michal Fabinger. 2013. "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition." *Journal of Political Economy* 121 (3): 528–583.

Appendix

A Appendix to Section 2: Data, Prices, and Search

In this appendix, we provide additional details on the construction of our price dataset and the construction of non-overlapping local markets. We also present supplementary descriptive evidence on search and price dispersion in the retail fuel market.

A.1 Construction of the price dataset

We construct the station-level price panel for Germany and France as follows. For each fuel station in our dataset, we observe a price every time it is changed, along with a precise time and date stamp for every change. On average, in 2019, fuel stations in Germany changed fuel prices 14 times per day, whereas there was typically one price change per day at French fuel stations. Based on the distribution of price changes, we construct hourly fuel prices from 6 am until 10 pm for every fuel station in Germany and France.

In the next step, we compute daily weighted average prices from the hourly distribution of price changes. To construct the weights, we use data on hourly fueling patterns reported in a representative survey among drivers for the German Federal Ministry of Economic Affairs. Figure A1 shows the share of motorists in Germany who fuel at a particular time of day. We further re-weight the hourly shares to produce weights for the hours between 6 am and 10 pm.


In Table 1, we also compute prices net of taxes and duties for both Germany and France. 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 65.45 Eurocent per liter for *E5* and *E10* gasoline, and 47.04 Eurocent per liter for diesel. The fee for oil storage is 0.27 Eurocent per liter for *E5* and *E10*, and 0.30 Eurocent per liter for diesel.³⁰ Before the temporary VAT reduction in 2020, the German VAT rate on retail fuel was 19%. In mainland France, fuel products are subject to a lump-sum tax of 60 to 70 Eurocent per liter, depending on the metropolitan region and fuel type.³¹ In addition, the French VAT rate on retail fuel is 20%.

We make a few restrictions to the fuel stations that we include in our analysis. In Germany, we drop stations located on highways (i.e., “Autobahn”), because these stations are typically around 20 to 30 Eurocent more expensive than regular fuel stations. We identify

³⁰See <https://www.avd.de/kraftstoff/staatlicher-anteil-an-den-kraftstoffkosten/>.

³¹See <http://www.financespubliques.fr/glossaire/terme/TICPE/>.

Figure A1: 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.

highway stations based on their address as well as manual checks. In France, we only focus on stations in mainland France (i.e., excluding stations on the island of Corsica and overseas).

A.2 Non-overlapping markets

To group fuel stations into non-overlapping local markets, we use an agglomerative hierarchical clustering algorithm based on the driving time between stations. This approach follows Carranza, Clark, and Houde (2015), Luco (2019), and Assad et al. (2020).

In the first step, we compute the driving time between all pairs of fuel stations in each country. To do this, we use the `osrmtime` Stata package by Huber and Rust (2016), which relies on OpenStreetMap data using the Open Source Routing Machine (OSRM).

Next, we implement the hierarchical clustering algorithm separately for stations in Germany and France. The algorithm begins with each station in a separate cluster. Then, iteratively, the algorithm combines the closest two clusters into a larger cluster and records the additional driving time required to link the clusters. We use average linkage, implying that two clusters are linked based on the average driving time between the stations in the two clusters. As this procedure moves on, the algorithm builds a clustering tree that indicates which clusters have been linked at which iteration and how much additional driving time is required to link two clusters.³² Eventually, all stations are combined into a single cluster.

³²See Appendix C in Luco (2019) for an example and illustration.

The objective of the clustering exercise is to find clusters of stations that are naturally separated from each other. The “height” of each link (i.e., the average driving time needed to link one cluster and another) is informative about such natural separations. Formally, we compute an inconsistency coefficient for each link, which captures the height of the current link relative to the heights of previous links.³³ A high inconsistency coefficient indicates that two clusters are far apart from each other (i.e., there is an inconsistency when linking the two clusters). The idea underlying this inconsistency measure is twofold. First, two clusters linked at a low additional driving time are more likely to belong to one local geographic market than two clusters linked at a much higher additional driving time. This is true irrespective of whether the original clusters are individual stations or groups of stations linked in a previous iteration. Second, if the driving time required to link two clusters is similar to the driving time required to link clusters (or individual stations) in previous iterations, then there is unlikely to be a natural border between this group of stations. In contrast, if the driving time required to link two clusters is much higher than the time needed to drive from station to station *within* these two clusters, then the two clusters are likely to represent separate local markets.

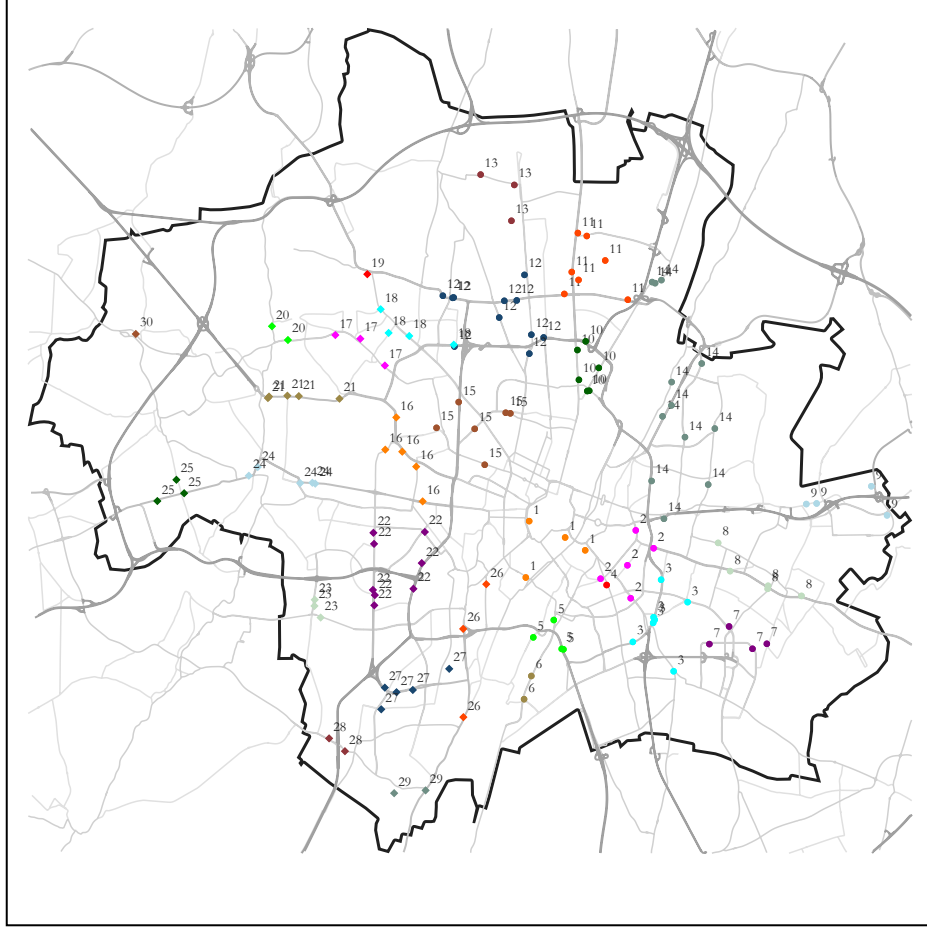
Finally, based on the clustering tree and the inconsistency coefficients for each link, we group stations into non-overlapping markets. This is done by pruning the tree at a selected threshold in the distribution of the inconsistency coefficients. We choose to prune the clustering tree at the 85th percentile in the distribution of the inconsistency coefficients. This threshold is in line with prior literature on retail fuel markets that used the 80th percentile (Assad et al., 2020) or 90th percentile (Luco, 2019). We verified that our results do not hinge on the specific choice of the threshold.

As pointed out by Luco (2019) and Assad et al. (2020), the advantage of this agglomerative hierarchical clustering approach is that researchers do not need to specify the number or size of markets. Instead, in the entire procedure outlined above, we only decide where to prune the clustering tree. One potential drawback of the approach is that the clustering algorithm tends to group stations in rural areas together, although they may be far away from each other in terms of absolute driving time. Therefore, we make one more explicit choice and define all fuel stations with no competitor within 10 minutes as monopoly markets, without including them in the clustering procedure.³⁴

³³See http://cda.psych.uiuc.edu/multivariate_fall_2012/matlab_help/cluster_analysis.pdf for additional details.

³⁴The European Commission also frequently uses market definitions based on 10-minute driving time. An example in the Commission’s assessment of the recent takeover of OMV stations by EG Group in the German market (see Case M.10134 – EG Group / OMV Germany Business).


Figure A2: Example of local markets in the city of Munich



Notes: The Figure shows a map of Munich where all fuel stations (indicated by solid squares or diamonds) are grouped into non-overlapping markets, using the agglomerative hierarchical clustering algorithm. The solid black line indicates the city boundary. The gray lines represent the road network, with thicker and darker lines indicating larger roads.

To illustrate the outcome of applying the clustering algorithm, the map in Figure A2 shows how we group stations into non-overlapping markets, using the city of Munich as an example. The Figure highlights that nearby stations are usually grouped together into one local market. It also shows that the algorithm often identifies “natural” clusters of stations. In Figure A3, we also present the distribution of market sizes in Germany. The median market consists of four stations and the 90th percentile is at seven stations. That is, the vast majority of the local markets defined by the clustering algorithm has a very reasonable size.

Figure A3: Distribution of market sizes in Germany



Notes: The Figure shows the distribution of market sizes in Germany when applying the hierarchical clustering algorithm to identify non-overlapping local markets. The histogram depicts the number of markets (on the y-axis) by the number of stations per market (on the x-axis).

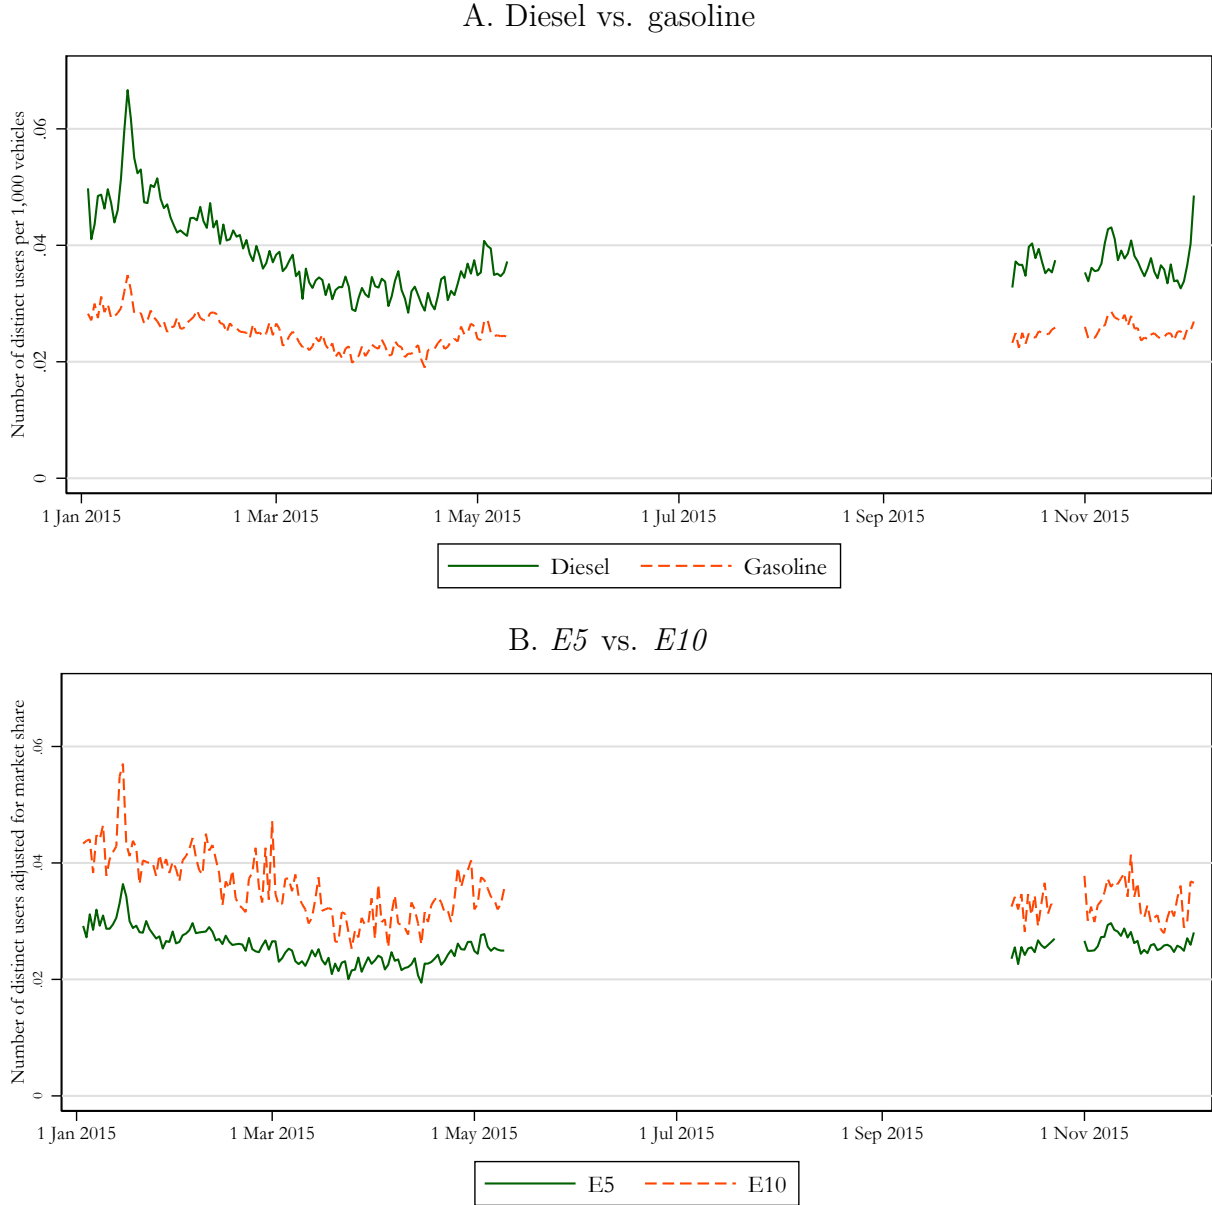
A.3 Search intensity by fuel type

In this section, we use data on search queries in 2015 from a major German price comparison smartphone app to confirm that the share of well-informed consumers is higher for diesel than for gasoline and higher for *E10* than *E5*.

Panel A of Figure A4 shows the daily number of distinct users searching for fuel prices by fuel type. Normalizing the number of users by the number of registered vehicles, we see that the ratio of searchers to the number of vehicles in circulation is around 50% higher for diesel than for gasoline. We report the number of distinct searchers rather than the total number of searches to adjust for the higher mileage of diesel drivers.

Panel B of Figure A4 shows the number of distinct searchers for *E5* and *E10*, divided by the number of gasoline vehicles in circulation and adjusted for the relative market shares of *E5* and *E10* within the gasoline market. This shows that the search intensity is substantially higher among consumers buying *E10* than those purchasing *E5*.

Figure A4: Consumer search patterns in Germany

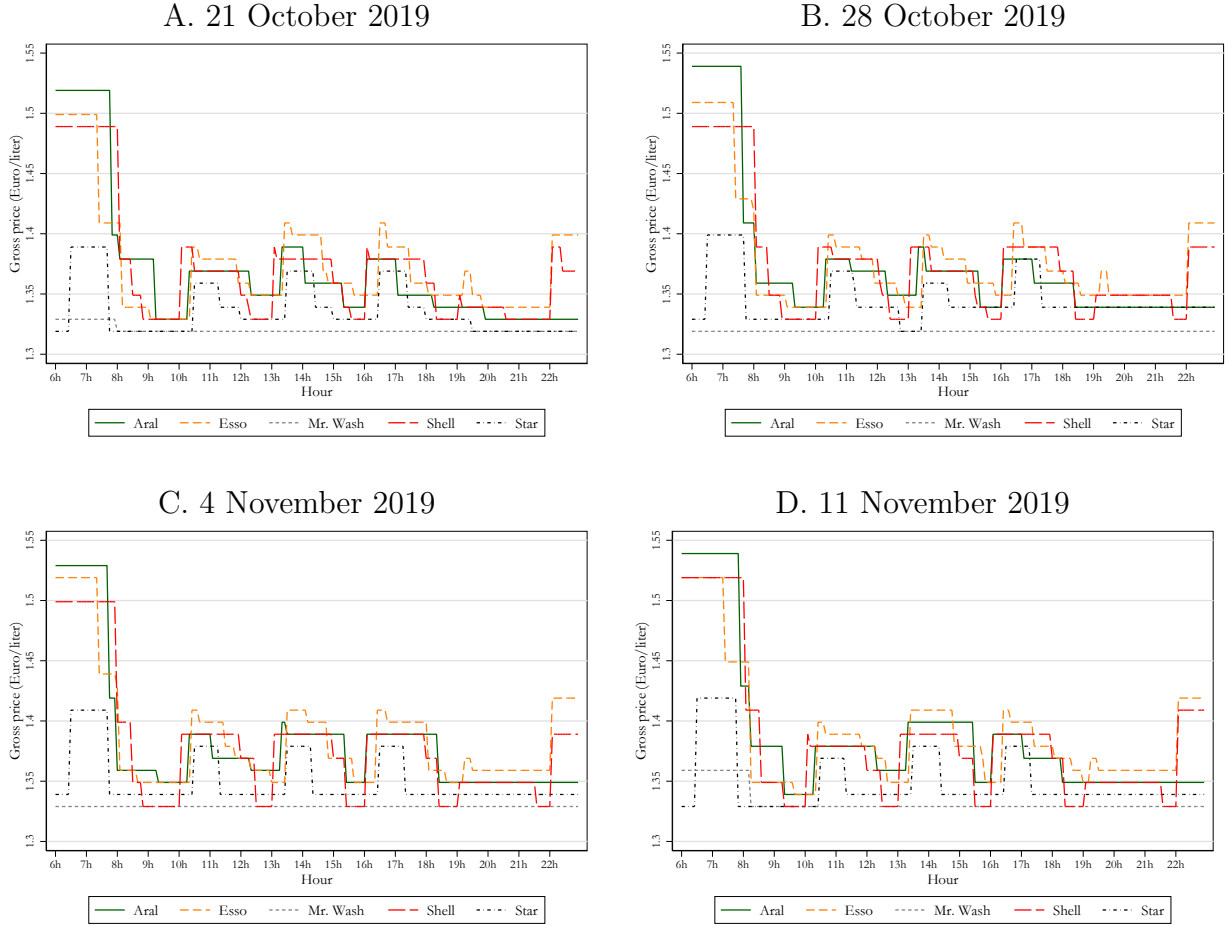


Notes: The Figure shows the daily number of searchers by fuel type on a major German smartphone app. The data is available for January to mid-May and mid-October to early December 2015. Panel A shows the number of distinct users who search for diesel vs. gasoline prices per 1,000 diesel or gasoline vehicles in circulation. The solid line corresponds to the search intensity for diesel, whereas the dashed line corresponds to the search intensity for gasoline. Panel B shows the number of distinct users who search for E5 vs. E10 per 1,000 gasoline vehicles in circulation and adjusted for the relative market shares of E5 and E10. The solid line corresponds to the search intensity for E5, whereas the dashed line corresponds to the search intensity for E10.

A.4 Additional evidence on search and price dispersion

In this section, we present additional evidence on search intensity and price dispersion. Figure 1 shows average daily price cycles for E10 in Germany in 2019. We now present price cycles at a more disaggregated level to show that these pricing patterns do not merely


Figure A5: Daily price cycles for *E10* on selected Mondays in one local market



Notes: The Figure shows prices of *E10* for five different stations in one local market in the city of Munich at different times of a specific day. Fuel prices are updated in five-minute intervals. Panels A, B, C, and D depict prices on 21 October, 28 October, 4 November, and 11 November (all Mondays), respectively.

result from averaging over time and across stations. Therefore, in Figure A5, we zoom in on one local market in the city of Munich (see market no. 24 in the map in Figure A2) and present the stations' raw prices on four consecutive Mondays in the fall of 2019. Several things are noteworthy in the Figure. First, on each of the four days, the stations' pricing follows a similar pattern, which is in line with that shown in Figure 1. Price increases typically occur at the same time, whereas the timing of price decreases is more idiosyncratic. Second, there are persistent differences in the average price level across stations, consistent with some degree of product differentiation (e.g., due to station amenities). For example, in Figure A5, Mr. Wash typically sets the lowest price, as this is the only station that does not belong to a vertically integrated brand. Finally, even at a particular time, the order of the stations' prices may vary across different days. This indicates that there is a substantial

Figure A6: Within market price residuals, 5 pm, 2019



Notes: The Table shows the distribution of the deviation of a fuel station's price from the average price in the same market (i.e., within market residuals) on the same day, at 5 pm for *E10* and for all stations that are not local monopolists. We use data for all weekdays in 2019. Panel A shows residuals when only controlling for market \times date fixed effects. In panel B, we additionally control for station fixed effects.

amount of price variation that is unpredictable to consumers, which is consistent with the mixed strategy equilibrium in our theoretical model.


In Table 2, we analyzed price dispersion more systematically by computing within market price residuals for 5 pm prices. Panel A of Figure A6 graphically illustrates the distribution of these residuals. Panel B shows the corresponding residuals when we additionally control for station fixed effects to absorb any time invariant price differences across stations. These residuals correspond to variation in prices that is unpredictable even to the most sophisticated consumers. Consistent with our stylized fact and the numbers in Table 2, Figure A6 shows that this unpredictable price dispersion is substantial.

Next, we investigate whether and how search and price dispersion are correlated with the absolute price level. As outlined in Section 7, the model with endogenous search by Tappata (2009) predicts that consumers search more when prices are low. Similarly, the model predicts that price dispersion (i.e., a measure of the gains from search) is high when prices are low.

Figure A7 shows the average number of searches per app user in 2015 for *E10*, along with the development of the gross price of *E10*. As can be seen in the Figure, search intensity and the price level are almost entirely uncorrelated. That is, in our empirical application, there is no evidence that consumers change the intensive margin of their search behaviour in response to changes in the price level.


Figure A8 depicts the relationship between price dispersion and the price level for *E10* in 2019. Price dispersion is computed as the difference between the maximum and the

Figure A7: Search per user and price level of *E10*, 2015



Notes: The Figure shows the development of daily search intensity (on the left axis, solid line) and the daily weighted average fuel price (on the right axis, dashed line) for *E10* in Germany in 2015. Search intensity is measured by the average number of searches per distinct user on a major German smartphone app. The search data is available for January to mid-May and mid-October to early December 2015.

Figure A8: Price dispersion and price level for *E10*, 2019



Notes: The Figure shows the development of daily price dispersion (on the left axis, solid line) and the daily weighted average fuel price (on the right axis, dashed line) for *E10* in Germany in 2019. Price dispersion is computed as the difference between the maximum and the minimum price within a local market, using prices at 5 pm.

minimum price within a local market, using daily prices at 5 pm. Thus, it corresponds to the specification with only market \times date fixed effects in the last column of panel B in Table 2. If anything, Figure A8 points to a positive correlation between price dispersion and the price

level. This would imply that gains from search are higher when prices are higher, which is at odds with the predictions by Tappata (2009).

B Appendix to Section 3: Theoretical Model

This appendix complements the theoretical model in Section 3. Here, we formally solve the model, prove our propositions, and consider extensions such as endogenous entry or pass-through of marginal costs.

B.1 Equilibrium price distribution

Lemma 1. *There is no pure strategy Nash equilibrium in prices if $N \geq 2$.*

Proof. Suppose that all N sellers choose 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. \square

Lemma 2. *There are no mass points in the equilibrium pricing strategies.*

Proof. Suppose that any price is played with positive probability. This means 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.³⁵ \square

Lemma 3. *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*

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

³⁵For a more detailed proof, see Varian (1980).

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

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

The expected profits 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

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

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 fulfilled, non-shoppers with a particular first draw of p search as long as the expected gain of searching is greater than s . Thus, search occurs so long as

$$s < p - \int_{\underline{p}}^{p_{max}} pf(p)dp. \quad (\text{B1})$$

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 $p_r \leq v$ must therefore satisfy

$$H(p_r; \phi, N, s) \equiv p_r - \int_{\underline{p}}^{p_r} pf(p)dp - s = 0. \quad (\text{B2})$$

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 v , 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

$$p_r = \begin{cases} E[p] + s & \text{if } E[p] + s < v \\ v & \text{otherwise} \end{cases}. \quad (\text{B3})$$

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

$$E[\pi(\underline{p})] = E[\pi(p_r)]. \quad (\text{B4})$$

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

$$\left(\frac{\underline{p}}{1+\tau} - c\right) \left(\phi + \frac{1-\phi}{N}\right) M = \left(\frac{p_r}{1+\tau} - c\right) \frac{1-\phi}{N} M. \quad (\text{B5})$$

We can simplify this expression and re-arrange it to yield an expression for the lowest price sellers may draw in equilibrium

$$\underline{p} = \frac{p_r}{\frac{\phi N}{1-\phi} + 1} + c \frac{1+\tau}{1 + \frac{1-\phi}{\phi N}}. \quad (\text{B6})$$

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

$$E[\pi(p_i)] = E[\pi(p_r)] \quad \forall \quad p_i \in [\underline{p}, p_r]. \quad (\text{B7})$$

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

$$\left(\frac{p_i}{1+\tau} - c\right) \left(\phi + \frac{1-\phi}{N}\right) (1 - F(p_i))^{N-1} M + \left(\frac{p_i}{1+\tau} - c\right) \left(\frac{1-\phi}{N}\right) (1 - (1 - F(p_i))^{N-1}) M = \left(\frac{p_r}{1+\tau} - c\right) \frac{1-\phi}{N} M. \quad (\text{B8})$$

We can solve this equation for the equilibrium density function according to which each seller i draws its prices from the support $[\underline{p}, p_r]$:

$$F(p_i) = 1 - \left(\frac{p_r - p_i}{p_i - c(1+\tau)} \frac{1-\phi}{N\phi}\right)^{\frac{1}{N-1}}. \quad (\text{B9})$$

For a given number of entrants N and a given set of exogenous parameters, Equations (B3), (B6), and (B9) uniquely identify the symmetric mixed strategy Nash equilibrium in prices.

We can derive the expected profit of each seller i in this equilibrium. Since the expected profit of each seller in the symmetric equilibrium is the same for any price chosen with positive probability, the expected profit of seller i drawing a price from the equilibrium price distribution is

$$E[\pi_i] = E[\pi(p_r)] = \left(\frac{p_r}{1+\tau} - c\right) \frac{1-\phi}{N} M. \quad (\text{B10})$$

Finally, we can derive the expected prices paid by non-shoppers and shoppers, namely the expected price and the expected minimum price. The expected price is

$$E[p] = \int_{\underline{p}}^{p_r} p f(p) dp = p_r - \int_{\underline{p}}^{p_r} F(p) dp, \quad (\text{B11})$$

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

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

To derive the expected minimum price we begin by setting up the probability density function of the minimum price. This can be written as

$$f_{\min}(p) = N(1 - F(p))^{N-1} f(p). \quad (\text{B12})$$

After inserting $F(p)$ and simplifying the expression, this yields

$$f_{min}(p) = \frac{p_r - p}{p - c(1 + \tau)} \frac{1 - \phi}{\phi} f(p). \quad (\text{B13})$$

The expected minimum price is then

$$E[p_{min}] = \int_p^{p_r} p f_{min}(p) dp = \int_p^{p_r} p \frac{p_r - p}{p - c(1 + \tau)} \frac{1 - \phi}{N\phi} f(p) dp. \quad (\text{B14})$$

After adding and subtracting $c(1 + \tau)$ in the numerator of the first fraction and further simplifications, we get that

$$E[p_{min}] = \frac{1 - \phi}{\phi} \left[\int_p^{p_r} p \frac{p_r - c(1 + \tau)}{p - c(1 + \tau)} f(p) dp - E[p] \right].$$

Finally, we can use integration by parts and rearrange terms to get the following expression for the expected minimum price:

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

□

B.2 Endogenous entry

To consider endogenous entry, we assume that 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 .

In this case, the game proceeds in two stages. In the first stage, firms decide whether to enter the market. 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.

In the main analysis, we 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 there need to be at least two sellers active in the market for the informedness of consumers to matter.

Lemma 4. *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*

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

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 $1, \dots, N$ 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 (B10). 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

$$\left(\frac{p_r}{1+\tau} - c\right) \frac{1-\phi}{F} M - 1 < N^* \leq \left(\frac{p_r}{1+\tau} - c\right) \frac{1-\phi}{F} M. \quad (\text{B15})$$

□

B.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 4. *With $0 < \phi < 1$, for any $\hat{c} > c$ the minimum element of the support of the equilibrium pricing strategy $\hat{p} > \underline{p}$ and the Nash equilibrium pricing strategy with c first-order stochastically dominates (FOSD) the pricing strategy with \hat{c} , i.e. $\hat{F}(p) \leq F(p) \quad \forall p$.*

Analogous to the explanation for ad valorem taxes, this means that if the share of shoppers is strictly positive, an increase in c leads to a shift in the support of the prices from which sellers draw in equilibrium towards higher prices. Furthermore, for each price on the

equilibrium pricing support, the likelihood that a drawn price is below said price decreases if marginal costs increase from c to \hat{c} .

As for the pass-through of ad valorem taxes, the pass-through of marginal costs converges to zero as the share of shoppers converges to zero. Since the minimum element of the support of prices and the density function monotonically 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 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 $\phi \rightarrow 1$, the pass-through rate $\rho_c \rightarrow 1 + \tau$.*

We can begin by looking at the cases when there are no shoppers and when there are only shoppers. If there are no shoppers, all sellers choose the monopoly price and pass-through of marginal costs is zero. If all consumers are shoppers, there is full pass-through of marginal costs or per unit taxes.³⁶

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.

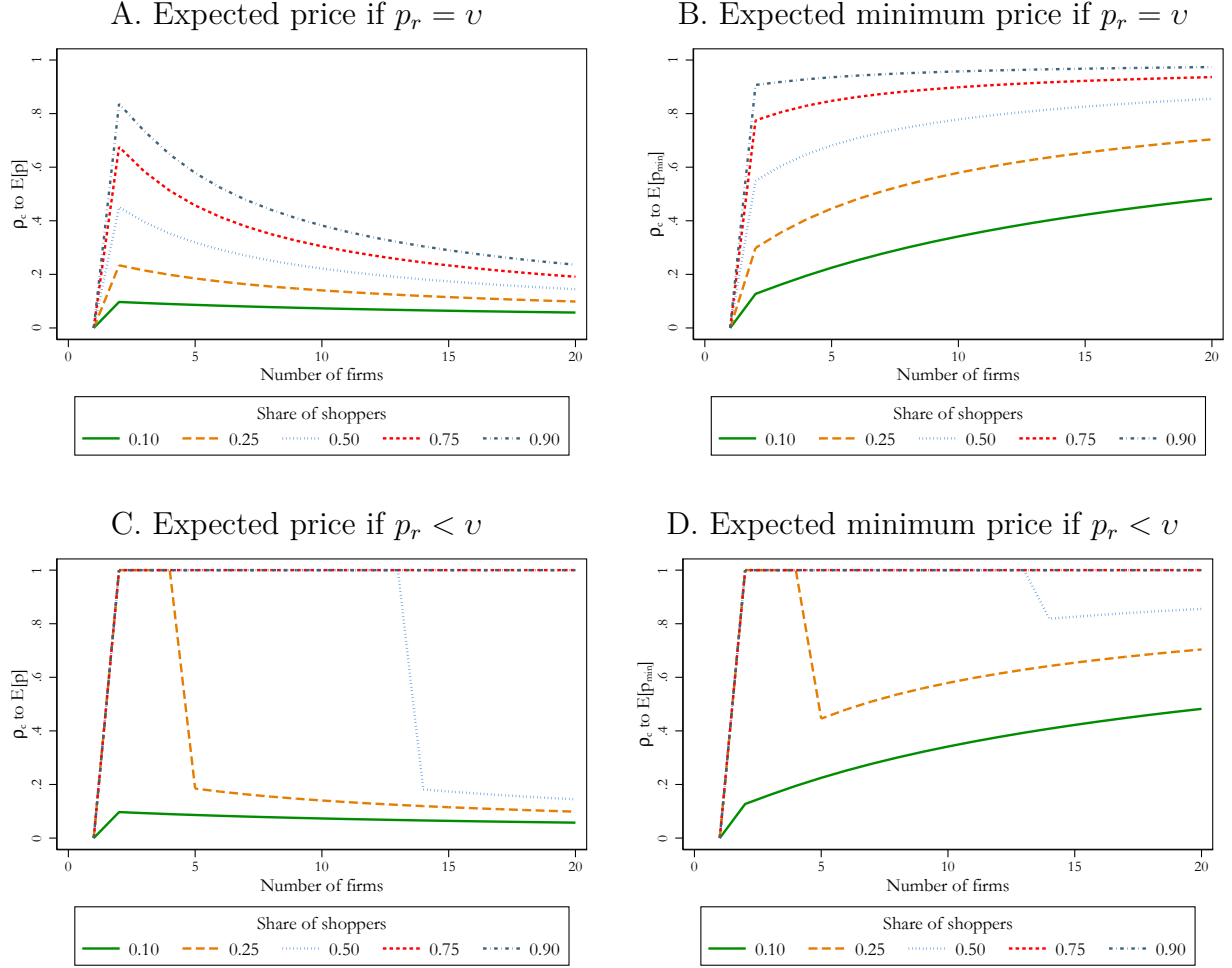
Let us now consider how pass-through of marginal costs varies with the number of active sellers. As we will see, all of our results and intuitions with respect to ad valorem tax pass-through extend to marginal costs.

Proposition 6. *With $0 < \phi < 1$, as $N \rightarrow \infty$ the pass-through of c to the minimum element of the equilibrium price support converges to full pass-through, i.e., $\rho_{c,p} \rightarrow 1 + \tau$.*

As the number of sellers increases, competition for shoppers becomes fiercer and the pass-through rate of marginal costs to \underline{p} increases. Furthermore, we also expect pass-through of marginal costs to $E[p]$ to first increase and then decrease, whereas pass-through to $E[p_{min}]$ should always increase as $N \rightarrow \infty$. The same reasoning as laid out for ad valorem taxes applies.

³⁶Although 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 $\hat{c} - c$.

Figure B1: Numbers of sellers and marginal cost pass-through



Notes: The Figure shows simulation results of how the pass-through rate of the ad valorem tax τ varies with the number of sellers. Panel A and B respectively show how the pass-through rate to the expected price, $E[p]$, and to the expected minimum price, $E[p_{min}]$, vary with the number of sellers if the reservation price is exogenous. Panel C and D show the same if the reservation price of non-shoppers, p_r , is endogenous. In all panels, the different lines correspond to different values of the share of shoppers, ϕ . Parameter values: $v = 4.5$, $\tau = 0.2$, $c = 0.4$, $\hat{c} = 0.44$, $s = \infty$ (without sequential search) and $s = 0.75$ (with sequential search).

The results from the numerical simulation in Figure B1 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. This is the case with and without sequential search (Panel A vs. Panel C). Pass-through to the expected minimum price always increases in the number of sellers if there is no sequential search (Panel B). If non-shoppers can search sequentially, the pass-through to the expected minimum price can either be monotonically increasing in the number of sellers or there can be a non-monotonic relationship between the number of sellers and the pass-through rate.

B.4 Proof of Propositions

Proof of Proposition 1. First, we assess the pass-through of τ to \underline{p} if $0 < \phi < 1$.³⁷ 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 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 $\hat{\tau} > \tau$ $\hat{F}(p) \leq F(p) \quad \forall p \in [\underline{p}, p_r]$. \square

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 $\underline{p} = p_r = v$. 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:

$$\frac{\partial^2 F(p)}{\partial \tau \partial \phi} = (\frac{1}{N-1})^2 (\frac{p_r - p}{p - c(1 + \tau)})^{\frac{1}{N-1}} \frac{c}{p - c(1 + \tau)} (\frac{1 - \phi}{\phi N})^{\frac{1}{N-1}-1} \frac{1}{\phi^2 N} > 0.$$

³⁷ \underline{p} is not defined for $\phi = 0$ or $\phi = 1$.

Thus, for higher ϕ , an increase from τ to $\hat{\tau}$ decreases the probability that prices are below a certain p more strongly. \square

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 $N \rightarrow \infty$. We find that

$$\lim_{N \rightarrow \infty} \rho_{\tau, \underline{p}} = \lim_{N \rightarrow \infty} \frac{\partial \underline{p}}{\partial \tau} \cdot \frac{1 + \tau}{\underline{p}} = \frac{c(1 + \tau)}{c(1 + \tau)} = 1.$$

Thus, with $N \rightarrow \infty$, pass-through of a value-added tax to the minimum element of the support of the equilibrium pricing strategy converges to full pass-through. \square

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

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

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

$$\frac{\partial F(p)}{\partial c} = - \left(\frac{1 - \phi}{\phi N}\right)^{\frac{1}{N-1}} \frac{1}{N-1} \left(\frac{p_r - p}{p - c(1 + \tau)}\right)^{\frac{1}{N-1}} \frac{1 + \tau}{p - c(1 + \tau)} < 0.$$

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

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 $\underline{p} = p_r = v$. 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.³⁸

³⁸Although 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 $\hat{c} - c$.

Finally, we study the case where $0 < \phi < 1$. Let us begin by analyzing how the pass-through rate changes with ϕ

$$\frac{\partial^2 p}{\partial c \partial \phi} = (1 + \tau) \left(1 + \frac{1 - \phi}{\phi N}\right)^{-2} \frac{1}{\phi^2 N} > 0.$$

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

$$\frac{\partial^2 F(p)}{\partial c \partial \phi} = \left(\frac{1}{N-1}\right)^2 \left(\frac{p_r - p}{p - c(1 + \tau)}\right)^{\frac{1}{N-1}} \frac{1 + \tau}{p - c(1 + \tau)} \left(\frac{1 - \phi}{\phi N}\right)^{\frac{1}{N-1} - 1} \frac{1}{\phi^2 N} > 0.$$

Thus, for higher ϕ , an increase from c to \hat{c} decreases the probability that prices are below a certain p more strongly. \square

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 $N \rightarrow \infty$. We find that

$$\lim_{N \rightarrow \infty} \rho_{c,p} = \lim_{N \rightarrow \infty} \rho_{c,p} (1 + \tau) \left(1 + \frac{1 - \phi}{\phi N}\right)^{-1} = 1 + \tau.$$

Thus, with $N \rightarrow \infty$, pass-through of marginal costs to the minimum element of the support of the equilibrium pricing strategy converges to full pass-through. \square

B.5 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 they 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 they have the option to wait also in the absence of a tax change. 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 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 Appendix to Section 4: Descriptive Evidence

In our main empirical analysis, we focus on a temporary VAT reduction in Germany in the second half of 2020. The 2020/21 tax changes are the only recent policy shifts that allow us to study differences in pass-through across fuel types. Thus, they allow us to test our first theoretical prediction that pass-through increases in the price sensitivity of consumers. In addition, in Section 6.2, we use tax changes in France in 2022/23 to study differences in pass-through to the average posted price and the minimum price *within* a given fuel type.

In this appendix, we present additional descriptive evidence suggesting that we should be cautious with using the 2022/23 tax changes for other analyses. First, we argue that comparisons across fuel types are problematic for the 2022/23 tax changes. Second, we show that there may have been spillover effects to France of the German tax cut in June 2022.

Figure C1 shows the development of gross fuel prices in Germany and France in February and March of 2022. Panels A and B present German and French prices, respectively. All prices are normalized to one on 1 February 2022. Two findings emerge from this Figure. First, there was a divergence between diesel and gasoline prices in March 2022 in both Germany and France. This is because gasoline and diesel markets were hit differently by Russia's invasion of Ukraine on 24 February 2022. As diesel is a close substitute for heating

Figure C1: Evolution of gross prices in early 2022



Notes: The Figure shows the evolution of daily gross fuel prices in February and March of 2022. Panels A and B present German and French prices, respectively. All prices are normalized to one on 1 February 2022. The solid line shows prices for *E5*. The short-dashed and long-dashed lines show prices for *E10* and diesel, respectively.


oil, demand for diesel increased relatively more than that for gasoline. As a consequence, diesel prices also increased disproportionately. Second, Figure C1 shows that Germany and France were affected differently by the shocks on the global oil market. While diesel (gasoline) prices increased by up to 43% (29%) in Germany relative to 1 February 2022, they only increased by up to 31% (19%) in France. That is, fuel prices increased much more in Germany than in France, following Russia’s invasion of Ukraine.

On 1 April 2022 (i.e., right after the time window shown in Figure C1), France introduced a fuel tax rebate of 18 Eurocent per liter on both diesel and gasoline. Due to the divergence in diesel and gasoline prices prior to the French tax cut, we cannot use this tax change to compare pass-through *across* fuel types.

Similarly, Germany implemented a temporary tax rebate on diesel and gasoline starting on 1 June 2022. As discussed in Section 2, we do not analyze this tax change because of intense public scrutiny and a concurrent market investigation by the Federal Cartel Office. An additional concern regarding the 2022 tax changes is that they were so large that they may have changed the opportunity cost of selling fuel across countries. That is, there may have been spillover effects from Germany to France (or vice versa), which would violate the stable unit treatment value assumption (SUTVA) underlying our empirical approach.

Figure C2 presents evidence of such potential spillover effects, showing an increase in French retail margins immediately after the introduction of the German tax cut on 1 June 2022. To compute margins for *E5*, *E10*, and diesel, we subtract taxes and duties as well as

Figure C2: Margins in France around 1 June 2022



Notes: The Figure shows the evolution of daily retail margins at French stations in May and June of 2022. The solid line shows margins for *E5*. The short-dashed and long-dashed lines show margins for *E10* and diesel, respectively. The vertical solid line marks the starting date of the German tax rebate on 1 June 2022.

the estimated input cost of crude oil from the gross price.³⁹ The Figure shows that margins at French stations increased by approximately 5 Eurocent on the day when the German fuel rebate went into effect and continued increasing in subsequent weeks. Therefore, estimating pass-through of the Germany tax cut on 1 June 2022 with France as the control group is problematic and not done in this paper.

D Appendix to Section 6: Empirical Results

In this appendix, we present additional results and several robustness checks for our empirical findings in Section 6.

³⁹To compute retail margins, we obtain daily data on the Brent price of crude oil at the port of Rotterdam from the US Energy Information Administration. On average, one barrel (42 gallons) of crude oil is refined into around 19 gallons of gasoline, 12 gallons of diesel, and 13 gallons of other products (e.g., jet fuel). See <https://www.eia.gov/energyexplained/oil-and-petroleum-products/refining-crude-oil.php>. Assuming that among the other products only jet fuel is of high value, we split the price of one barrel into the cost of producing gasoline, diesel, and jet fuel to compute the share of the Brent price that corresponds to a particular fuel product. Around 54% of the Brent oil price per barrel corresponds to the production of 19 gallons of gasoline, while around 34% corresponds to the production of 12 gallons of diesel. Finally, we then transform these values into the approximate input cost per liter of gasoline and diesel.

Table E1: Effect of the tax change on log prices (with controls)

	Tax decrease			Tax increase		
	<i>E5</i>	<i>E10</i>	Diesel	<i>E5</i>	<i>E10</i>	Diesel
	(1)	(2)	(3)	(4)	(5)	(6)
Tax change	-0.0080*** (0.0004)	-0.0125*** (0.0004)	-0.0234*** (0.0004)	0.0557*** (0.0004)	0.0600*** (0.0004)	0.0814*** (0.0003)
Retail & recreation	0.0025*** (0.0008)	0.0032*** (0.0006)	0.0046*** (0.0005)	-0.0018*** (0.0006)	-0.0043*** (0.0005)	-0.0052*** (0.0005)
Workplaces	0.0130*** (0.0009)	0.0118*** (0.0007)	-0.0018** (0.0007)	0.0000 (0.0008)	-0.0001 (0.0007)	-0.0039*** (0.0006)
DE \times oil price	0.2282*** (0.0066)	0.1931*** (0.0055)	0.0470*** (0.0052)	0.0790*** (0.0036)	0.0202*** (0.0036)	0.0755*** (0.0031)
Pass-through rate	32% [29%, 35%]	49% [46%, 53%]	93% [90%, 96%]	67% [66%, 68%]	70% [69%, 71%]	82% [81%, 82%]
Date fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Station fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,128,241	2,318,268	2,694,252	1,804,493	1,988,459	2,318,185

Notes: The Table presents DID estimates using the model in Equation (1), but we now additionally control for regional mobility for retail and recreational purposes and to workplaces, using data from the Google Mobility Report. Columns (1) to (3) present average treatment effect estimates of the German VAT reduction on 1 July 2020 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 CO₂ emissions tax on 1 January 2021 on *E5*, *E10*, and diesel log prices, respectively. Columns (4) to (6) use data from 1 November to 15 December 2020 for the pre-treatment period, and from 1 January to 28 February 2021 for the post-treatment period. Standard errors clustered at the market level are shown in parentheses. We also compute the pass-through rates corresponding to the point estimates and report their 95% confidence intervals in brackets.

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

D.1 Robustness: Additional controls

In Table E1, we report results on the effect of the tax change on *E5*, *E10*, and diesel prices when we additionally control for regional mobility for retail and recreational purposes and to workplaces. To this end, we use data from the Google Mobility Report. Overall, the point estimates of the pass-through rates are similar to our main estimation results in Table 3.

The results in columns (1) to (3) show that the average price for *E5* decreased by 0.80% after the tax reduction in July 2020, whilst average prices for *E10* and diesel decreased by 1.25% and 2.34%, respectively. This implies pass-through rates of 32% for *E5*, 49% for *E10*, and 93% for diesel. The results in columns (4) to (6) show that, following the subsequent tax increase, the average price of *E5* increased by about 5.57%, whereas *E10* and diesel prices increased by about 6.00% and 8.14%, respectively. We compute a joint pass-through rate of the VAT increase and the carbon emissions price of 67% for *E5*, 70% for *E10*, and 82% for diesel.

Overall, the estimates in Table E1 are close to our baseline estimates without controls. Therefore, they show that including controls in our regression model does not affect our main results. In particular, pass-through is still significantly higher for diesel than for gasoline and it is significantly higher for *E10* than *E5*, consistent with Prediction 1.

D.2 Robustness: Balanced sample

In our baseline approach, we estimate the DID model in Equation (1) on an unbalanced panel of fuel stations. That is, we do not observe a price for every station on every day. This could be, for example, because some stations are closed on weekends or holidays or due to permanent station closures or new openings. To ensure that this does not drive our results, we now restrict the sample to fuel stations in Germany and France for which we have a price observation on every day in our sample period. For diesel, for example, this is the case for 83% of fuel stations in Germany and 62% in France for the analysis of the tax reduction, and for 81% of stations in Germany and 72% in France for the analysis of the tax increase.

Columns (1) to (3) show the effect of the tax decrease. We estimate that 90% of the tax decrease is passed on to diesel consumers, while the pass-through rates for *E10* and *E5* is 43% and 14%, respectively. Columns (4) to (6) show the estimates for the tax increase in January 2021. We find pass-through rates of 66% for *E5*, 70% for *E10*, and 84% for

Table E2: Effect of the tax change on log prices (balanced sample)

	Tax decrease			Tax increase		
	<i>E5</i>	<i>E10</i>	Diesel	<i>E5</i>	<i>E10</i>	Diesel
	(1)	(2)	(3)	(4)	(5)	(6)
Tax change	-0.0035*** (0.0004)	-0.0107*** (0.0004)	-0.0226*** (0.0004)	0.0544*** (0.0004)	0.0601*** (0.0003)	0.0833*** (0.0003)
Pass-through rate	14% [11%, 17%]	43% [39%, 46%]	90% [87%, 93%]	66% [65%, 66%]	70% [70%, 71%]	84% [83%, 84%]
Date fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Station fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
DE \times oil price	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,734,669	1,967,631	2,174,886	1,465,464	1,691,664	1,922,544

Notes: The Table presents DID estimates using the model in Equation (1). Unlike in Table 3, we now estimate the model on a balanced sample of stations for which we have a price observation on every day. Columns (1) to (3) present average treatment effect estimates of the German VAT reduction on 1 July 2020 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 CO₂ emissions tax on 1 January 2021 on *E5*, *E10*, and diesel log prices, respectively. Columns (4) to (6) use data from 1 November to 15 December 2020 for the pre-treatment period, and from 1 January to 28 February 2021 for the post-treatment period. Standard errors clustered at the market level are shown in parentheses. We also compute the pass-through rates corresponding to the point estimates and report their 95% confidence intervals in brackets.

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

Table E3: Effect of the tax change on log prices (accounting for anticipatory effects)

	Tax decrease			Tax increase		
	<i>E5</i>	<i>E10</i>	Diesel	<i>E5</i>	<i>E10</i>	Diesel
	(1)	(2)	(3)	(4)	(5)	(6)
Tax change	0.0006 (0.0005)	-0.0058*** (0.0005)	-0.0235*** (0.0005)	0.0523*** (0.0003)	0.0548*** (0.0003)	0.0696*** (0.0003)
Pass-through rate	-2% [-7%, 2%]	23% [19%, 27%]	93% [90%, 97%]	63% [62%, 64%]	64% [64%, 65%]	70% [69%, 70%]
Date fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Station fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
DE \times oil price	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,869,158	2,037,695	2,369,148	2,084,101	2,295,092	2,677,702

Notes: The Table presents DID estimates using the model in Equation (1). Columns (1) to (3) present average treatment effect estimates of the German VAT reduction on 1 July 2020 on *E5*, *E10*, and diesel log prices, respectively. Columns (1) to (3) use data from 1 May to 15 June 2020 for the pre-treatment period, and from 1 July to 31 August 2020 for the post-treatment period. Columns (4) to (6) present average treatment effect estimates of the VAT increase and CO₂ emissions tax on 1 January 2021 on *E5*, *E10*, and diesel log prices, respectively. Columns (4) to (6) use data from 1 November to 28 February 2021. Standard errors clustered at the market level are shown in parentheses. We also compute the pass-through rates corresponding to the point estimates and report their 95% confidence intervals in brackets.

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

diesel. That is, most pass-through rates are slightly lower than in our baseline specification in Table 3 where we use the full unbalanced sample. Importantly, however, our main result regarding the heterogeneous pass-through across fuel types remains unchanged.

Overall, the estimates show that our results are robust to estimating the DID model on a restricted (balanced) sample of fuel stations for which we have a price observation on every day.

D.3 Robustness: Anticipatory effects

In Table E3, 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 result regarding 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 decreases from 68% to 63% for *E5*, from 72% to 64% for *E10*, and from 84% to 70% for diesel. This is expected, since we can graphically see important anticipatory effects of the tax pass-through in the second half of December 2020 (see Figure 3). Therefore, including this time period into the control period necessarily leads to an underestimate of the pass-through rate. Reassuringly, the difference between fuel types remains similar to our main results, although the difference in pass-through for *E5* and *E10* is very small. Although not accounting for anticipatory effects would slightly modify our estimates, the overall conclusions remain the same. Yet, 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.

D.4 Robustness: Synthetic difference-in-differences analysis

Next, we repeat our baseline pass-through estimation employing a synthetic differences-in-differences (SDID) approach instead of a standard DID. SDID is a variation of DID that aims to match pre-treatment trends between the treatment and control groups using weights. In this sense, SDID is similar to synthetic control methods (Arkhangelsky et al., 2021).

With the SDID approach, we estimate pass-through using a two-step procedure. First, we calculate unit and time weights that minimize the difference in pre-treatment trends between treated and control groups and the difference in outcomes between pre- and post-treatment periods for the control group. In the second step, we estimate a DID model similar to that in Equation (1) using the weights from the first step. We use clustered bootstrapping with 300 replications and clustering at the station level to estimate standard errors.

Formally, to estimate the average pass-through rate of the tax changes on fuel prices, SDID solves the following minimization problem:

$$(\hat{\beta}^{sdid}, \hat{\mu}, \hat{\alpha}, \hat{\pi}) = \arg \min_{\beta, \mu, \alpha, \pi} \left\{ \sum_{i=1}^N \sum_{t=1}^T (Y_{it} - \mu - \alpha_i - \pi_t - \text{Tax}_{it}\beta)^2 \hat{w}_i^{sdid} \hat{\lambda}_t^{sdid} \right\}, \quad (\text{E1})$$

where $\hat{\beta}^{sdid}$ is the estimated effect of the policy change, and \hat{w}_i^{sdid} and $\hat{\lambda}_t^{sdid}$ are the SDID unit and time weights, respectively. As in our baseline DID model in Equation (1), Y_{it} is the logarithm of the weighted average 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 . The variables α_i and π_t again correspond to fuel station and date fixed effects, respectively. Using SDID requires a balanced panel, but we have shown above in Appendix D.2 that this sample restriction by itself does not affect our baseline estimates.

Table E4: Effect of the tax change on log prices (SDID)

	Tax decrease			Tax increase		
	<i>E5</i>	<i>E10</i>	Diesel	<i>E5</i>	<i>E10</i>	Diesel
	(1)	(2)	(3)	(4)	(5)	(6)
Tax change	-0.0058*** (0.0007)	-0.0115*** (0.0004)	-0.0209*** (0.0005)	0.0625*** (0.0005)	0.0636*** (0.0004)	0.0859*** (0.0004)
Pass-through rate	23% [17%, 28%]	45% [42%, 49%]	83% [79%, 86%]	75% [74%, 77%]	75% [74%, 75%]	86% [85%, 87%]
Date fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Station fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,734,669	1,967,631	2,174,886	1,465,464	1,691,664	1,922,544

Notes: The Table presents SDID estimates using the model in Equation (E1). Columns (1) to (3) present average treatment effect estimates of the German VAT reduction on 1 July 2020 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 CO₂ emissions tax on 1 January 2021 on *E5*, *E10*, and diesel log prices, respectively. Columns (4) to (6) use data from 1 November to 15 December 2020 for the pre-treatment period, and from 1 January to 28 February 2021 for the post-treatment period. Standard errors obtained via clustered bootstrap with 300 replications are shown in parentheses. We also compute the pass-through rates corresponding to the point estimates and report their 95% confidence intervals in brackets.

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

Using the SDID approach, we first re-estimate the baseline pass-through rates of the German tax changes in 2020/21. In a second step, we also assess the robustness of our results regarding the relationship between the pass-through rate and the number of sellers in the market. We cannot use SDID to estimate the difference in the pass-through rates between the minimum and the average price, as SDID does not allow for a triple interaction term.

D.4.1 Baseline pass-through estimation

Table E4 shows the results of estimating the regression model presented in Equation (E1) for the analysis of the 2020/21 tax changes in Germany. Columns (1) to (3) show the effect of the tax decrease. For *E5*, the pass-through rate is 23%, while around 45% and 83% of the tax decrease is passed on to consumers who refuel with *E10* and diesel, respectively. Therefore, the ranking of pass-through rates with respect to fuel types is robust to using an SDID strategy.

Columns (4) to (6) show that the tax increase raised prices for all fuel products. We find a joint pass-through rate of the VAT increase and the introduction of a CO₂ emissions tax of 75% for *E5* and *E10*, and 86% for diesel. That is, for the tax increase, pass-through is again significantly higher for diesel than for gasoline, whereas pass-through rates for *E5* and *E10* are statistically indistinguishable.

Overall, the ranking of pass-through rates with respect to fuel types and their magnitude remain largely robust to using SDID instead of a simple DID approach.

D.4.2 Number of sellers and pass-through

Figure E1 shows the relationship between the pass-through rate and the number of competitors of a focal station when we estimate the station-level pass-through rates with an SDID (i.e., with additional weights) and using a restricted balanced sample of stations. We estimate the time and unit weights only once and then use them for estimating each station-specific treatment effect.

The results look very similar to our main results in Figure 5. Consistent with Prediction 3, we still find a non-monotonic relationship between the number of competing price setters in a local market and the average pass-through rate.

For *E5* and *E10*, pass-through is again relatively low for local monopolists for both the tax decrease in summer 2020 and the tax increase in winter 2020/21. With at least two competing price setters in a local market, the average pass-through tends to decrease in the number of sellers in the case of the tax decrease and appears constant in the case of the tax increase. As before, for diesel, the relationship between the number of sellers and pass-through has an inverted-U shape with a peak at a higher number of sellers than in the case of *E5* and *E10*.

In summary, our analysis shows that the non-monotonic relationship between the number of sellers and pass-through is robust to using an SDID strategy.

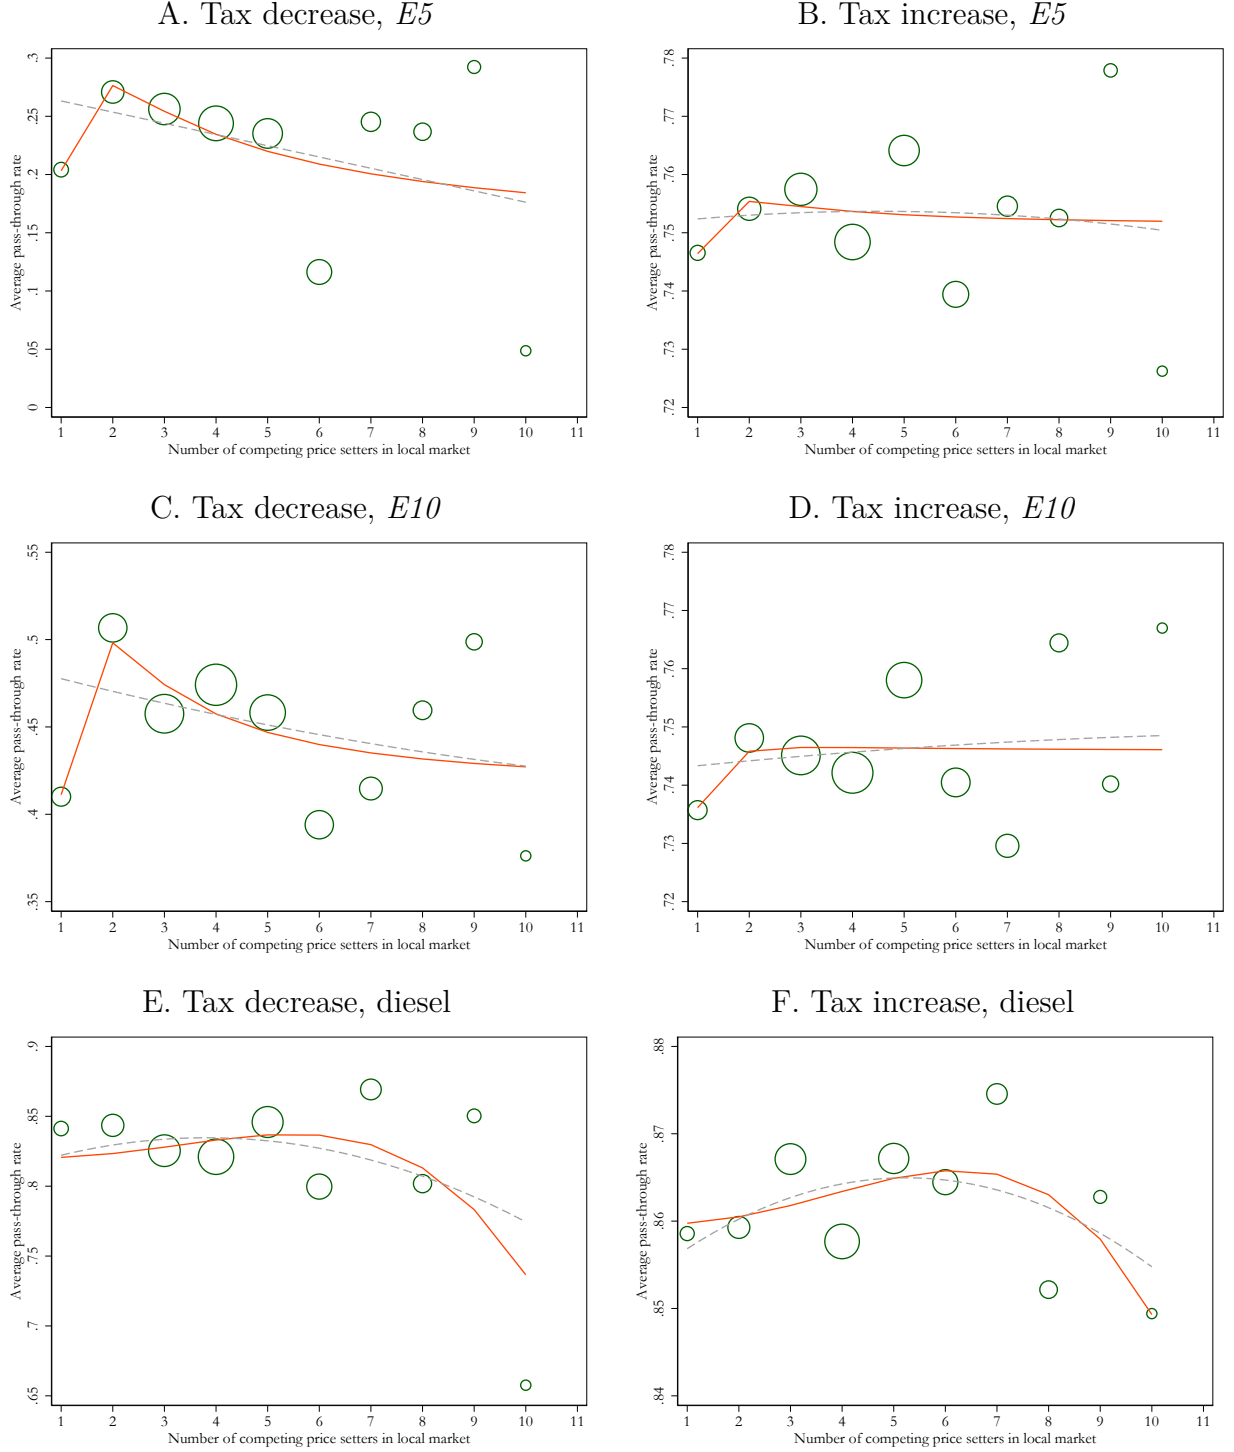
D.5 Robustness: Price at 10th percentile

In Figure 4, we show that pass-through to the minimum price in a local market is generally higher than pass-through to the average posted price. One potential concern regarding the minimum price is that it may reflect outlier prices that are only valid for a short period of time. In that case, the minimum price that we identify would not actually be relevant for consumers.

We address this point in Figure E2, where we re-estimate the difference in the pass-through rate between the minimum and the average price. This time, however, instead of the minimum price in a local market, we use the price at the 10th percentile of the distribution of all hourly prices across all stations in that market on a given day.


The results look very similar to our main specification in Figure 4. They indicate that our findings on the difference in pass-through between the minimum and the average price are not driven by outlier prices. If anything, the estimated difference between pass-through to the minimum price and the average price slightly becomes more pronounced for the tax decreases. Therefore, E2 provides even stronger evidence in support of Prediction 2 that

Figure E1: Average pass-through by number of competitors (SDID)



Notes: The Figure shows how the pass-through rate to the average price varies with the number of competing price setters in a market. Unlike in Figure 5, pass-through rates are estimated using a SDID strategy and a restricted balanced sample. Panels A, C, and E depict the pass-through rates for the German VAT decrease on 1 July 2020 for $E5$, $E10$, and diesel, respectively. Panels B, D, and F depict the pass-through rates for the German VAT increase and introduction of a carbon price on 1 January 2021 for $E5$, $E10$, and diesel, respectively. In every panel, each circle plots the average pass-through rate for a group of stations with a particular number of competing price setters within a non-overlapping local market. The size of a circle is proportional to the total number of stations with a given number of competitors. The solid line shows a fractional polynomial fit. The dashed line shows a quadratic fit. The number of competitor stations is trimmed at the 97.5th percentile.

Figure E2: Pass-through to market-level minimum (10th percentile) vs. average prices



Notes: The Figure shows the difference in the market-level pass-through rate (in percentage points) between the minimum price and the average posted price. The average posted price is the average daily price within a non-overlapping market by weighting the price at every full hour of the day between 6 am and 10 pm equally. Unlike in Figure 4, we now replace the minimum price with price at the 10th percentile of the distribution of all hourly prices within a non-overlapping market on a given day. The Figure depicts the pass-through rates implied by the DID estimate β_2 in Equation (2) along with 99% confidence intervals, based on standard errors clustered at the market level. Regressions are estimates separately for each tax change and fuel type. For most tax changes, we use data for the two months before and after every tax change. Exceptions include the German tax increase on 1 January 2021, where we exclude the second half of December to account for anticipatory effects. For the tax increase in France on 16 November 2022, we only use the period until 31 December 2022 as post-treatment period. Similarly, for the French tax increase on 1 January 2023, we use the period from 16 November until 31 December 2022 as pre-treatment period. Solid (hollow) symbols indicate statistically significant (insignificant) point estimates at the 1% level.

pass-through is higher to the price paid by well-informed consumers than to the price paid by uninformed consumers.

D.6 Formal test for non-monotonicity

The results in Figure 5 indicate that pass-through to the average price is not monotonically increasing in the number of sellers. In this section, we formally test for non-monotonicity by applying the U-test by Lind and Mehlum (2010). This allows us to test the null hypothesis of a monotone or U-shaped relationship against the alternative hypothesis of an inverted-U shape.

Table E5: U-test for non-monotonicity

	Tax decrease			Tax increase		
	<i>E5</i>	<i>E10</i>	Diesel	<i>E5</i>	<i>E10</i>	Diesel
	(1)	(2)	(3)	(4)	(5)	(6)
No. of competing price setters	0.0178 (0.0124)	0.0136 (0.0137)	0.0221** (0.0110)	0.0043 (0.0030)	0.0017 (0.0032)	0.0064*** (0.0023)
(No. of competing price setters) ²	-0.0023* (0.0012)	-0.0013 (0.0014)	-0.0020* (0.0010)	-0.0005** (0.0003)	-0.0002 (0.0003)	-0.0005** (0.0002)
Constant	0.2283*** (0.0294)	0.4030*** (0.0312)	0.8779*** (0.0262)	0.6632*** (0.0074)	0.7070*** (0.0079)	0.8186*** (0.0058)
Inverse-U test						
<i>t</i> -value	1.3043	0.8777	1.6463	1.3138	0.4782	1.7083
<i>p</i> -value	0.0961	0.1900	0.0499	0.0945	0.3163	0.0438
Extreme point	3.8900	5.1693	5.5786	3.9854	4.2498	6.3627
Lower bound	1	1	1	1	1	1
Slope at lower bound	0.0132	0.0110	0.0182	0.0032	0.0013	0.0054
<i>t</i> -value	1.3043	0.9926	2.0243	1.3138	0.4782	2.7812
Upper bound	10	10	10	10	10	10
Slope at upper bound	-0.0279	-0.0127	-0.0175	-0.0065	-0.0022	-0.0036
<i>t</i> -value	-2.3261	-0.8777	-1.6463	-2.4840	-0.7975	-1.7083
Observations	13,998	13,394	14,250	14,033	13,426	14,296

Notes: The Table shows the estimates from regressing the station-level pass-through rate on the number of competing price setters in a local non-overlapping market, its square, as well as a constant. Columns (1) to (3) present estimates for the German VAT reduction on 1 July 2020, while columns (4) to (6) present estimates for the VAT increase and CO₂ emissions tax on 1 January 2021. The bottom part of the Table shows the application of the U-test by Lind and Mehlum (2010).

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

To this end, we first regress the station-level pass-through rate on the number of competing price setters and its square. As the number of competing price setters in a local market is always positive, a necessary condition for a hump-shaped relationship is that the coefficient on the linear term is positive, that the coefficient on the quadratic term is negative, and that the extreme point lies within the data range. Lind and Mehlum (2010) argue, however, that these conditions are not sufficient to guarantee an inverse-U-shaped relationship. Therefore, they propose a formal statistical test for a U (or inverse-U) shape, taking into account the slope at the lower and upper bound of the data range.

We present our estimation results in Table E5, following the exhibition in Pennerstorfer et al. (2020). The Table shows that our estimated coefficients satisfy all three necessary conditions for each fuel type and tax changes, although not all estimates are statistically significantly different from zero. The *p*-value on the inverse-U test indicates that we reject the null hypothesis of a monotonic or U-shaped relationship at the 10% level for both tax

changes for diesel and *E5*. For *E10*, the pass-through rate is also increasing (decreasing) in the number of competing price setters for local monopolies (large markets), but the slopes at the bounds are imprecisely estimated. Therefore, we cannot formally reject a monotonic relationship between the number of competitors and the pass-through rate, but the estimates for *E10* still speak in favor of a hump-shaped relationship.

Overall, the U-test by Lind and Mehlum (2010) further supports Prediction 3 that the relationship between the number of sellers and pass-through to the expected price is non-monotonic.