

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

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

October 14, 2026

Abstract

Imperfect consumer price information gives sellers market power. We study how this shapes tax pass-through. In a consumer search model, pass-through rises with the share of consumers who know all prices and is higher for the lowest price in a market, paid by informed consumers, than for the average price, paid by uninformed consumers. Pass-through to the average price is nonmonotonic in the number of sellers, whereas the pass-through gap between the two prices widens with it. Five tax changes in German and French fuel markets support these predictions, comparing fuel types with differently informed buyers within the same stations.

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

*Montag: New York University, Stern School of Business (email: felix.montag@nyu.edu); Mamrak: Charles River Associates (email: robin@mamrak.de); Sagimuldina: Charles River Associates (email: asagimuldina@crai.com); Schnitzer: Ludwig-Maximilians-University Munich, Department of Economics (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, Simon Martin, Eugenio Miravete, Andreas Peichl, Chris Snyder, Meredith Startz, Dario Tortarolo, Thomas Wollmann, and participants at many conferences and seminars 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 common cost shocks, such as taxes, is fundamental for the design of public policies. 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.¹ 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 extend a theoretical consumer search model where some consumers know all prices and others need to search for prices sequentially to study tax pass-through. Using this model, we show how the relationship between competition and pass-through differs between models with fully informed or imperfectly informed consumers. We test the theoretical predictions using multiple tax changes in the retail fuel market. A key feature of this setting is that we can estimate how tax pass-through differs between products and local markets that differ according to the primitives of our model.

We make three central contributions that we show theoretically and test empirically. First, we show that the higher the share of consumers that are well informed about prices, the higher is pass-through of tax increases and decreases. Second, pass-through to the price paid by well-informed consumers is higher than to the price paid by uninformed consumers in the same market. Finally, when some consumers are uninformed about prices, the relationship between the number of sellers and pass-through to the average price is nonmonotonic, whereas the difference in pass-through between the prices paid by well-informed and by uninformed consumers widens with the number of sellers.

Taken together, the pass-through patterns we document are difficult to reconcile with models that do not feature imperfect consumer information about prices. The retail fuel market is an ideal setting for two reasons. First, within a fuel type, products are homogeneous, which eliminates other sources of imperfect information. Second, retail fuel has a high degree of price transparency relative to other products. A finding of effects here suggests that they are likely present, and potentially larger, in markets with less price transparency. Despite this high degree of price transparency, retail fuel markets exhibit significant price dispersion that is difficult to predict for consumers. Some consumers use price comparison apps and are well informed about prices. Others are not.

¹Adachi and Fabinger (2022) generalize this framework to allow for richer governmental intervention, and Kroft et al. (2024) allow for free entry and love-of-variety preferences.

For the theoretical framework, we adapt the Stahl (1989) consumer search model to study tax pass-through. Well-informed consumers (“shoppers”) know all prices and buy from the seller offering the lowest price; uninformed consumers (“nonshoppers”) search sequentially and, in equilibrium, pay the first price they draw. Firms set prices using mixed strategies, so equilibrium prices are dispersed. From an *ex ante* perspective, shoppers pay the expected minimum price, whereas nonshoppers pay the expected posted price. The degree of market power depends on the number of sellers and the share of well-informed consumers, as a larger share incentivizes firms to compete on prices. As the number of sellers grows, however, it becomes increasingly unlikely for any one of them to attract the shoppers, so firms increasingly price for their captive nonshoppers, near the reservation price, where pass-through is low. This is the source of the nonmonotonic relationship between the number of sellers and pass-through to the average price. Consumer information thus shapes pass-through through the elasticity of each seller’s residual demand, holding aggregate demand fixed. In Section 7, we contrast these results to different full-information benchmarks, in which pass-through can only depend on demand curvature and seller conduct and show that these cannot explain our empirical findings.

Our empirical setting is the German retail fuel market, for which we observe daily prices at the station level. We examine a temporary decrease in the value-added tax (VAT) in July 2020 and its reversal, combined with the introduction of a carbon emissions price, in January 2021, comparing prices of stations in Germany with those in France in a difference-in-differences (DID) design. Three French tax changes in 2022/23 provide additional evidence on the difference in pass-through to the minimum and the average posted price.²

The setting disentangles the two channels of competition in the model. Fuel stations sell all three fuel types, and the fuel type consumers purchase is correlated with their incentive to become informed about prices: there is strong evidence that diesel drivers are better informed than gasoline drivers, and that drivers fueling with *E10* are better informed than those fueling with *E5*.³ We use search data from a price comparison app to support this ranking. Comparing fuel types within the same station therefore varies the share of well-informed consumers while holding the station network fixed, and comparing markets with different numbers of sellers varies competition while holding the consumer type fixed.

Our first finding is that tax pass-through is higher for fuel types with a larger share of well-informed consumers. For the January 2021 tax increase, estimated pass-through is 70% for *E5*, 75% for *E10*, and 88% for diesel and the ordering is the same for the tax decrease.

²We cannot study differences in pass-through between fuel types in France, as the French tax changes occurred in 2022/23, when different fuel types were hit differently by global energy shocks.

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

While most of the empirical literature on tax pass-through and market power has treated consumers as fully informed about prices (Nakamura and Zerom, 2010; Miravete, Seim, and Thurk, 2018; Hollenbeck and Uetake, 2021), we show that the price information consumers have is a quantitatively important determinant of pass-through. Relatedly, Duso and Szücs (2017) find higher cost pass-through for electricity tariffs that consumers actively need to choose than for default tariffs, and Kosonen (2015) finds that Finnish hairdressers pass on VAT decreases more for advertised services.

Two concurrent papers study closely related mechanisms. Westphal (2025) finds that sellers pass through cost changes that consumers are aware of more than changes they are unaware of. A key difference to our setting is that this variation concerns awareness of the cost change, whereas we study widely publicized tax changes and vary how well consumers know prices. Fischer, Martin, and Schmidt-Dengler (2026) estimate a consumer search model on pre-pandemic German retail fuel data and, simulating the 2020 VAT cut, find that better-informed consumers face higher effective pass-through, benefiting high-income regions disproportionately. The two approaches are complementary: their structural model can evaluate counterfactual tax designs, whereas we estimate pass-through quasi-experimentally from multiple observed tax changes across fuel types and market sizes.⁴

Our second 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 holds for every fuel type in each of the three tax increases we study, in Germany and in France, and it survives a placebo test based on single-station monopoly markets, in which the two prices coincide by construction. Tax incidence thus differs across consumer groups within the same market. This extends the literature on the distributional effects of tax pass-through, in which Harju et al. (2022) find lower fuel tax pass-through in high-income areas and Conlon, Rao, and Wang (2024) show that the sin tax burden is concentrated among few households. We add a within-market margin, between informed and uninformed consumers buying at the same stations.⁵

Our third finding concerns the number of sellers and has two parts. First, pass-through to the average price shows no systematic increase beyond duopoly, in line with the predicted nonmonotonic profile. Second, in the tax increases, the pass-through gap between the minimum and the average price widens with the number of sellers. The empirical record on competition and pass-through is mixed, with evidence that tax pass-through increases (e.g.,

⁴Their value-of-information measure, the difference between the average and the minimum price in a market, is the level counterpart of the pass-through gap we study.

⁵Our 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-shaped relationship between household income and consumer search.

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. The nonmonotonic profile can rationalize these patterns, as estimated relationships depend on the range of market sizes in the data. Relatedly, Fischer, Martin, and Schmidt-Dengler (2025) find that consumers with easy access to price information gain the most from entry in the German retail fuel market, the price-level counterpart of our finding that entry widens this gap.

That more sellers can soften price competition when consumers are imperfectly informed is a classical insight (Rosenthal, 1980; Stahl, 1989; Janssen and Moraga-González, 2004). These papers concern price levels; we extend the insight to tax pass-through and decompose competition into two channels with different pass-through implications, which our empirical design tests separately. In concurrent theoretical work, Garrod et al. (2026) analyze cost pass-through in the Varian (1980) model under general demand and also contrast the two notions of competition, and Albrecht and Whitmeyer (2026) characterize cost pass-through by quantile in a framework where consumers differ in their consideration sets. Relative to these analyses, our model endogenizes the reservation price through sequential search, delivers the pass-through profile in the number of sellers in closed form, and is taken to the data.

Within the retail fuel market, work on tax pass-through has focused on supply-side determinants (e.g., Marion and Muehlegger, 2011), whereas a separate literature documents that imperfect price information matters for pricing in these markets.⁶ We connect the two: differences in consumer price information across fuel types, and in the number of sellers across markets, shape how taxes are passed through into fuel prices.

Our findings matter for the design of commodity taxes. The lower the share of well-informed consumers, the lower is the pass-through of such taxes. Well-informed consumers, who pay prices close to the market minimum, bear a larger share of the tax on the price they pay than uninformed consumers, so incidence also differs within a market. The number of sellers, in contrast, is a poor guide to pass-through when market power derives from imperfect information. These implications extend beyond retail fuel: markets with both well-informed and uninformed consumers are widespread, and models of competition with imperfect information explain, for instance, price differences between online and brick-and-mortar stores (Baye, Morgan, and Scholten, 2006).

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.

⁶See, e.g., Chandra and Tappata (2011), Byrne and Roos (2017), Luco (2019), Pennerstorfer et al. (2020), Martin (2024), and Montag, Sagimuldina, and Winter (2026).

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

2.1 Data

Our dataset includes real-time price changes for almost all fuel stations in Germany and France, along with various station characteristics. German stations are mandated to report price changes to the Market Transparency Unit at the Federal Cartel Office.⁷ Similarly, in France, a government agency requires stations to report price changes.⁸ We construct a daily price for each station as the average of its hourly prices, weighted by the share of drivers who refuel at each hour according to a survey of motorists. 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 nonoverlapping local markets using a hierarchical clustering algorithm based on driving time, as in previous studies (e.g., Carranza, Clark, and Houde, 2015, Lemus and Luco, 2021, and Assad et al., 2024). The idea underlying this approach is to find clusters of stations that are naturally separated from each other. We prune the clustering tree at the 85th percentile of the distribution of inconsistency coefficients and classify stations with no competitor within a 10-minute drive as monopoly markets, a cutoff the European Commission also uses in merger assessments. The details of our clustering method are explained in Appendix A.2. Table 1 shows that we assign the 14,648 German stations to 3,408 local markets with an average size of 4.3 stations. In France, there are 9,075 fuel stations assigned to 2,781 local markets.

We measure the number of competitors in a local market by the number of competing price setters. Two stations whose prices are set by the same entity count as one price setter. For Germany, we have two data sources that allow us to establish a common price setter across stations. First, the station dataset contains information on the brand of a fuel

⁷TankerKönig, a price comparison website, provides access to these data. See https://dev.azure.com/tankerkoenig/_git/tankerkoenig-data.

⁸See <https://www.prix-carburants.gouv.fr/rubrique/opendata/>. In France, fuel stations selling less than 500 m³ of fuel per year are exempt from reporting price changes.

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	2.5 min	n/a
Median driving distance to closest competitor	1.4 km	n/a
B. Prices (in Euro), <i>E5</i>		
Mean gross price	1.41	1.53
Mean price net of taxes and duties	0.53	0.58
C. Prices (in Euro), <i>E10</i>		
Mean gross price	1.39	1.49
Mean price net of taxes and duties	0.51	0.57
D. Prices (in Euro), 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	3,408	2,781
Mean no. of stations in market	4.30	3.26
Mean no. of competitors	3.67	n/a
F. Local monopolies		
Share of monopoly markets	16%	n/a
Median driving time to closest competitor	10.5 min	n/a
Median driving distance to closest competitor	9.5 km	n/a
G. Prices (in Euro), <i>E5</i>		
Mean average posted price	1.42	1.52
Mean minimum price	1.37	1.51
H. Prices (in Euro), <i>E10</i>		
Mean average posted price	1.40	1.49
Mean minimum price	1.35	1.46
I. Prices (in Euro), diesel		
Mean average posted price	1.25	1.45
Mean minimum price	1.21	1.42

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. Nonoverlapping markets are defined with a hierarchical clustering algorithm, as explained in Appendix A.2. A competitor is defined as a competing price setter (available only for Germany). Panels G to I exclude monopoly markets, defined as markets with a single price setter in Germany and as markets with a single station in France (where ownership information is unavailable).

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.⁹ Moreover, some firms operate fuel stations under different brands or even without brands. “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). On average, there are 3.7 different price setters per market in Germany. Furthermore, 16% of markets contain only a single price setter and are thus monopoly markets. The median monopoly station is 10.5 minutes away from its closest competitor, whereas the median driving time to the closest competitor across all stations is only 2.5 minutes. For France, our dataset does not contain information on fuel station brand or ownership.

Finally, we use data on search queries in 2015 from a major German smartphone app that enables users to compare fuel prices across stations. For each fuel price search, 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 search intensity between consumers who 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%.¹⁰ Diesel and gasoline can be considered separate markets in the short term because of the high cost of substitution on both the demand and supply sides.

Standard gasoline (called *Super*) has an octane rating of 95 and is 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.¹¹ Although *E5* and *E10* are not taxed differently, *E10* is typically 4–6 eurocents cheaper in Germany because of a minimum biofuels quota. Since *E5* and *E10* have the same octane rating, quality differences between the two fuel types are minimal.

⁹We conducted several interviews with market participants. All our interviewees confirmed that prices are set at the headquarters level for both large integrated conglomerates and 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->.

¹⁰Based 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.

¹¹In addition, there are other types of *Super* with an octane rating of 98, but their market share in Germany is only approximately 5%.

2.3 Price dispersion

The summary statistics in Table 1 reveal substantial differences between the average posted and the minimum price 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 the intraday price cycle or product differentiation, as well as random and unpredictable price changes.

Part of this variation is an intraday price cycle. The average price of *E10* in Germany is more than 10 eurocents higher at approximately 7:30 am than at approximately 10 pm, and on average there are 14 daily price changes for *E10* at German fuel stations in 2019 (Appendix Figure A5). For this reason, we weight daily station prices by the time at which drivers fuel and compare stations at a fixed hour in Section 5.¹²

To identify the random and unpredictable price dispersion for consumers, we focus on 5 pm and calculate the absolute price deviation of fuel stations from the mean price in their local market for all nonmonopolistic stations, regressing daily 5 pm prices on market $\times$ date fixed effects and computing the resulting residuals. However, some stations may always deviate from the mean price in the same way because of product differentiation. They may, for example, be in a particularly attractive location or offer better amenities. To isolate the nonconstant part of the deviation from the market mean, we further include station fixed effects. The remaining price variation is difficult for consumers to predict.

Table 2 decomposes the observed price dispersion into its predictable and unpredictable components. On average, the absolute price deviation from the market mean is 1.6 eurocents for *E5* and diesel, and 1.7 eurocents for *E10*. The mean absolute deviation from the mean after we include 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 approximately 2.9 eurocents for all fuel types, which is substantial.

Stylized Fact 1. *A substantial share of price dispersion is unpredictable to consumers.*

¹²The model in Section 3 does not feature price cycles, but cycles are not inconsistent with it. In experimental markets based on the Burdett and Judd (1983) and Stahl (1989) search models, Cason, Friedman, and Hopkins (2021) observe price cycles whose prices, pooled over time, are well described by the mixed strategy equilibrium distribution.

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

Table 2: Within-market price residual (in Euro), 5 pm, 2019

	Stations	Markets		
	Mean abs. deviation	$p_{25}-p_{75}$	$p_{10}-p_{90}$	Range
<i>A. E5</i>				
Market $\times$ date FE	.0163	.0273	.0416	.0425
Market $\times$ date FE and station FE	.0104	.0179	.0278	.0285
<i>B. E10</i>				
Market $\times$ date FE	.0173	.0293	.0450	.0458
Market $\times$ date FE and station FE	.0105	.0182	.0281	.0288
<i>C. Diesel</i>				
Market $\times$ date FE	.0161	.0272	.0414	.0424
Market $\times$ date FE and station FE	.0103	.0177	.0278	.0285
No. of stations/markets	14,138	2,898	2,898	2,898
Observations	3,505,247	756,378	756,378	756,378

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. The table also shows the distribution of this absolute deviation after we include 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.

2.4 Consumer information

Because stations must report price changes immediately, smartphone apps provide consumers with real-time price information. These apps provide perfect information on prices to users, whereas nonusers can discover prices only by driving from station to station.

Stylized Fact 2. *Some consumers are well informed about prices (app users), whereas others need to search for prices sequentially.*

The degree of knowledge that consumers have about prices often correlates with the type of fuel 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, fuel costs are lower, making them a fixed-cost investment in a lower marginal cost of driving. Drivers who select diesel engines therefore have a higher incentive to save on fuel costs and are more likely to use price comparison apps.

Our search data support this. Normalizing the number of distinct app users searching for each fuel type by the number of registered vehicles, we find that the number of users who search for diesel prices is approximately 50% higher than the number of users who search

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

for gasoline prices. This is in line with the hypothesis that, on average, the share of diesel drivers who are well informed about prices is higher than the corresponding share of gasoline car drivers. More details are presented in Appendix A.3.

Montag, Sagimuldina, and Winter (2026) show that price comparison apps are widely used in Germany and that their introduction led to an appreciable decrease in average prices for diesel and gasoline. Comprehensive data on usage and market shares do not exist, but the Federal Cartel Office reports that more than 40 information providers actively distribute price information through apps and websites,¹⁵ and Google Maps has included real-time fuel prices in its navigation since October 2020.¹⁶ Our search data predate the tax changes by five or more years. We rely on them only for the ranking of search intensity across fuel types. This ranking is driven by persistent structural factors, such as mileage patterns, and is unlikely to have reversed. Moreover, our main empirical comparisons exploit variation across fuel types within the same station, which holds the local market and its app coverage fixed.

Commercial vehicles often run on diesel, and drivers who do not pay for their own fuel may be less price sensitive. This concern is limited. At least 66% of diesel passenger cars are owned by private individuals: as of January 1, 2020, there were 15.1 million diesel passenger vehicles but only 5.2 million commercial passenger vehicles of any engine type (Kraftfahrt-Bundesamt, 2021). Many commercial drivers, such as the self-employed or those receiving a fuel allowance, also have an incentive to reduce fuel costs. If anything, price-insensitive commercial diesel drivers would work against finding the highest pass-through for diesel.

In addition to differences between diesel and gasoline, there are differences in price sensitivity between buyers of *E5* and *E10* in Germany. These differences 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 further segment consumers according to how well informed they are.¹⁷ Despite *E5* being more expensive, the majority of gasoline drivers in Germany purchase *E5*. According to the German Automobile Association (ADAC), *E10* is approximately 1.5% less efficient than *E5*.¹⁸ However, this efficiency gap accounts for only a fraction of the observed difference in prices between *E5* and *E10* since *E10* is usually 4–6 eurocents cheaper. A 2020 ADAC survey suggests that the difference in price sensitivity between *E5* and *E10* is due to preferences and a lack of information. Among respondents

¹⁵See: https://www.bundeskartellamt.de/DE/Aufgaben/MarkttransparenzstelleFuerKraftstoffe/TankApps/tankapps_node.html (accessed on April 2, 2025).

¹⁶See: <https://t3n.de/news/google-maps-suche-guenstigen-1327331/> (accessed on April 2, 2025).

¹⁷Although biofuels can pose a significant threat to the engine of a vehicle not compatible with *E10*, approximately 90% of gasoline-run vehicles, including all vehicles produced after 2012, are compatible. A complete list of compatible vehicles can be found at <https://www.dat.de/e10/>.

¹⁸See <https://www.adac.de/verkehr/tanken-kraftstoff-antrieb/benzin-und-diesel/e10-tanken/>.

fueling with *E10*, the most cited reason is lower prices (72%), followed by environmental concerns (37%). Among those who do not fuel with *E10*, the most cited reasons are technical concerns (51%) and uncertainty about cost and benefits (23%).¹⁹

This evidence is consistent with more buyers of *E10* than of *E5* choosing to become informed about prices in Germany. Our search data provide further supporting evidence 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*. France saw no comparable *E10* controversy, and French gasoline drivers predominantly buy *E10*.

Stylized Fact 3. *Search intensity differs by fuel type. In Germany, it is higher for diesel than for gasoline and for E10 than for E5.*

3 Theoretical Model

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

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 divide into two groups. A fraction ϕ of fully informed shoppers know the prices of all sellers and buy from the lowest-price seller. The remaining fraction $1 - \phi$ of uninformed nonshoppers draw a first price for free, know the distribution of prices, and can sequentially search at an incremental search cost s until they find a price weakly below their reservation price p_r .

On the supply side, an exogenous number N of sellers, indexed by i , produce at a constant marginal cost of c .²⁰ Sales are subject to an ad valorem tax τ .

Sellers first choose their price, and consumers then make search and purchase decisions. The equilibrium concept is a (perfect Bayesian) Nash equilibrium, in which sellers play

¹⁹The full survey results can be found at <https://www.adac.de/news/umfrage-e10-tanken/>.

²⁰We treat the number of sellers as exogenous because our empirical study concerns a short-term tax adjustment during which entry seems unlikely. Entry is straightforward to endogenize. Suppose symmetric potential entrants pay a fixed and sunk cost F before prices are set, and let the reservation price equal the valuation. By Equation (B3) in the appendix, the expected profit per entrant is then strictly decreasing in the number of entrants, so the equilibrium N^* is the largest integer at which it covers F . We focus on $N^* \geq 2$, since consumer informedness matters only when at least two sellers are active.

mutual best responses and nonshoppers search optimally given correct beliefs about the price distribution.

Prices refer to the gross price paid by consumers. Sellers bear the statutory incidence of the tax and (partially) “pass through” its cost to consumers. Two derived quantities govern everything that follows. The first is the tax-inclusive marginal cost

$$\kappa \equiv c(1 + \tau),$$

the lowest gross price at which a seller breaks even. Taxes and costs affect equilibrium prices only through κ . The second is the competition index (see Baye, Morgan, and Scholten, 2006)

$$\Lambda \equiv \frac{N\phi}{1 - \phi}.$$

Its role is most transparent in a seller’s residual demand. A seller charging price p always serves its captive share $(1 - \phi)/N$ of nonshoppers. With some probability w , it has the lowest price in the market and additionally wins all shoppers. Its expected demand is therefore $M \frac{1-\phi}{N} (1 + \Lambda w)$. The index Λ measures the demand that is contested in the price competition relative to the captive base a seller serves regardless, and with it the price elasticity of the seller’s residual demand. As $\Lambda \rightarrow 0$, each seller prices as a monopolist over its captive consumers. At a fixed number of sellers, as $\phi \rightarrow 1$ and hence $\Lambda \rightarrow \infty$, pricing approaches the Bertrand outcome. A larger number of sellers also raises Λ , but it lowers each seller’s probability of winning the contested demand. This tension is the source of the nonmonotonic effects of the number of sellers established below.

Let $dm/d\kappa$ denote the response, in levels, of a price statistic m of the equilibrium price distribution, a quantile, the expected price, or the expected minimum price, to a change in the tax-inclusive cost. The pass-through rate of the ad valorem tax is

$$\rho_{\tau}^m = \frac{\partial m}{\partial \tau} \cdot \frac{1 + \tau}{m} = \frac{\kappa}{m} \cdot \frac{dm}{d\kappa}.$$

We write ρ_{τ}^{avg} and ρ_{τ}^{min} for the rates to the expected price and the expected minimum price and drop the superscript when no confusion arises. Since all cost and tax shifters act through κ , the mechanisms are identical for a per unit tax.

3.2 Equilibrium price distribution

While there exists no pure strategy Nash equilibrium in prices, there is a unique symmetric mixed strategy equilibrium in which all sellers draw a price from an interval $[\underline{p}, p_r]$.²¹ The reservation price is $p_r = \min\{E[p] + s, v\}$. Because no seller ever charges above p_r , nonshoppers buy at the first price they draw and do not search sequentially in equilibrium. If searching is sufficiently cheap, it equals the expected price at the next draw plus the search cost s . With relatively high search costs, it is simply the valuation v .

Two features of the equilibrium drive every result. First, all equilibrium prices lie between the cost floor κ and the reservation price p_r , and competition compresses them toward cost. The expected price, which is the average price paid by nonshoppers, is

$$E[p] = \kappa + (p_r - \kappa) (1 - T_1(\Lambda, N)),$$

where $T_1 \in (0, 1)$ is the fraction of the gap $p_r - \kappa$ that competition closes in expectation. Equivalently, each equilibrium price is a weighted average of the cost floor κ and the reservation price p_r , with the weight on cost equal to the contested share of the seller's demand at that price. T_1 is the expected cost weight in the price a nonshopper pays. It depends only on the competition index Λ and the number of sellers N . Lemma B2(i) in Appendix B.2 gives it in closed form. When the reservation price does not respond to the tax, only the cost anchor moves, and T_1 is exactly $dE[p]/d\kappa$.

Second, in equilibrium, the expected minimum price, which is the average price paid by shoppers, is tied to the expected price in the following way:²²

$$E[p_{min}] = \kappa + \frac{1 - \phi}{\phi} (p_r - E[p]). \quad (1)$$

Sellers earn their entire rent from nonshoppers. Industry profit equals what sellers would earn by charging the reservation price and serving only their captive consumers. The margin that shoppers pay over cost is therefore pinned down by the rent that sellers fail to extract from nonshoppers, $p_r - E[p]$, scaled by the ratio of nonshoppers to shoppers (Lemma B2(ii) in the appendix).

²¹Lemma B1 in Appendix B.1 derives the equilibrium price distribution and the support floor $\underline{p}$ in closed form.

²²The expected minimum price refers to the average price paid by shoppers when the game is repeated across time or space. At a given time and location, there is one minimum price and N prices.

The two branches of the reservation price partition the model into two regimes. By the structure of the equilibrium, the expected gain from one more search is $p_r - E[p] = (p_r - \kappa) T_1$, so the consistency condition $p_r = E[p] + s$ is linear in p_r and solves to

$$p_r = \min \left\{ \kappa + \frac{s}{T_1(\Lambda, N)}, v \right\}$$

(Lemma B2(iii) in the appendix). When $s > (v - \kappa) T_1$, nonshoppers never gain from an additional search, and the reservation price is anchored at the valuation, $p_r = v$. This is the pricing model of Varian (1980), which we call the no-search regime. When $s \leq (v - \kappa) T_1$, the reservation price falls below the valuation and responds to costs. This is the Stahl (1989) case, which we call the sequential-search regime. Search still does not occur in equilibrium, but the option to continue searching now disciplines the reservation price.²³

The two regimes differ in how the price distribution responds to the tax-inclusive cost. In the no-search regime, a higher κ raises the floor of the distribution while its top stays at the valuation. The distribution is compressed from below, and the expected price rises by $T_1 < 1$ per euro of cost. In the sequential-search regime, $p_r - \kappa = s/T_1$ does not depend on κ , so the entire price distribution translates one-for-one with the tax-inclusive cost (Lemma B2(iv) in the appendix). Every quantile, the expected price, and the expected minimum price rise one-for-one with κ , and the rate to any such price m is $\kappa/m < 1$.

3.3 Effect of price sensitivity on the pass-through rate

An increase in the ad valorem tax τ raises every quantile $q < 1$ of the equilibrium price distribution, and with it the expected price and the expected minimum price. The pass-through rates established below are strictly between zero and one. We first analyze how the pass-through rate varies with the price sensitivity of consumers.²⁴

Proposition 1. *Fix $N \geq 2$ and $s > 0$. 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, $\rho_\tau = 1$. For $0 < \phi < 1$, the pass-through rate of the ad valorem tax to every quantile $q < 1$ of the price distribution, and to the expected price and the expected minimum price, is strictly increasing in ϕ , with $\rho_\tau \rightarrow 0$ as $\phi \rightarrow 0$ and $\rho_\tau \rightarrow 1$ as $\phi \rightarrow 1$.*

²³At the boundary $s = (v - \kappa) T_1$ the two branches meet, and the propositions assign the boundary itself to the sequential-search branch. A tax increase raises κ and moves the market into the no-search branch, whereas a tax decrease moves it into the sequential-search branch, so pass-through at the boundary is the one-sided derivative in the direction of the tax change.

²⁴The proofs of all propositions can be found in Appendix B.3.

The extreme cases are familiar benchmarks. With no shoppers, the equilibrium is degenerate at the monopoly price, which is unaffected by the tax, and pass-through is zero (Diamond, 1971). With only shoppers, it is the classical Bertrand outcome at $\kappa = c(1 + \tau)$, and pass-through is complete. In between, a larger share of shoppers raises the competition index Λ and pushes price levels toward cost. In the no-search regime, it also raises $dm/d\kappa$ for every price statistic m . Both of these effects raise the pass-through rate. This holds true for every quantile below the top of the support and therefore also for the expected price and the expected minimum price, which are fixed-weight averages of those quantiles.

At a fixed number of sellers, a larger informed share raises the price elasticity of each seller's residual demand, so imperfect price information is itself a source of market power.

Pass-through also differs between the average and the minimum price.

Proposition 2. *With $0 < \phi < 1$, $N \geq 2$, and $s > 0$, the pass-through rate of the ad valorem tax τ to the expected minimum price is strictly higher than the rate to the expected price, $\rho_{\tau}^{min} > \rho_{\tau}^{avg}$, in both regimes. In the no-search regime, the same ordering holds for $dE[p_{min}]/d\kappa$ and $dE[p]/d\kappa$.*

A tax increase raises the cost floor κ . In the no-search regime, this compresses the support of the price distribution from below. The minimum price is closer to the floor and rises by more in euro terms than the average price, and it is lower to begin with. Both forces push toward a higher rate. In the sequential-search regime, both prices rise by the same euro amount, and the ordering comes entirely from the lower base of the minimum price.

3.4 Effect of the number of sellers on the pass-through rate

The number of active sellers is a second dimension of competition, often more salient in empirical applications.

Proposition 3. *With $0 < \phi < 1$ and $s > 0$, pass-through of the ad valorem tax τ to the expected price is zero under monopoly ($N = 1$), strictly positive for every finite $N \geq 2$, strictly decreasing in N on $N \geq 2$, and converges to zero as $N \rightarrow \infty$. Pass-through to the expected price therefore peaks at $N = 2$.*

The intuition combines two monopoly benchmarks. Facing inelastic unit demand, the monopolist charges consumers their valuation and does not adjust the price when taxes change. As $N \rightarrow \infty$, the expected price converges back to the monopoly price. Each seller wins the shoppers with vanishing probability and increasingly prices for its captive nonshoppers near the reservation price. The decline is strict and monotonic after $N = 2$, so the peak occurs exactly at duopoly. Entry leaves each seller's expected contested demand

unchanged relative to its captive base. Shoppers always go to some seller, so a seller wins them with probability $1/N$, and its captive base shrinks at the same rate. What entry changes is the dispersion of prices. A smaller captive base lowers the profit a seller receives by pricing at the reservation price, so it sacrifices less by undercutting and the contest for shoppers turns fiercer. The bottom of the price distribution falls. At the same time each seller wins that contest less often, so it spends more of the time pricing for its captives alone, near the reservation price. The distribution stretches in both directions. Pass-through at the cost-anchored bottom cannot exceed one, while over the growing mass of near-monopoly prices it falls toward zero, so the average declines. In full-information models with log-concave demand of constant curvature, by contrast, the rate of pass-through to $E[p]$ rises monotonically in the number of sellers.

Section 3.5 traces the profile numerically under downward-sloping demand and imperfect comparability. The interior peak survives, and imperfect comparability shifts it beyond duopoly, whereas the hump disappears under log-convex demand and constant-elasticity demand leaves a rate of exactly one at every N .

Propositions 1 and 3 together imply that the effect of more intense competition on pass-through depends on its source. Holding the number of sellers fixed, a larger share of informed consumers raises Λ and raises pass-through to every moment of the price distribution. Entry also raises Λ , yet beyond duopoly it lowers pass-through to the average price. Two markets can therefore become more competitive by the same amount, as measured by Λ , with opposite consequences for pass-through. The number of sellers is not a sufficient statistic for the intensity of competition.

Proposition 4. *With $0 < \phi < 1$ and $s > 0$, the pass-through rate of τ to the expected minimum price is strictly increasing in N in the no-search regime, with $\rho_\tau^{min} \rightarrow 1$ as $N \rightarrow \infty$, and exactly constant in N in the sequential-search regime.*

In the no-search regime, the expected minimum price converges to the cost floor κ as the number of sellers increases, leaving sellers less margin to absorb a tax increase. Both the level price change (the numerator of ρ_τ) and the compression of the price toward κ (the denominator) push the rate toward full pass-through. In the sequential-search regime, two effects cancel exactly. More sellers intensify the contest for shoppers, which pushes the minimum price down, while the higher expected price raises the reservation price of nonshoppers, which relaxes competition and pushes the minimum price up. Equation (1) forces the two to offset, because the margin shoppers pay is tied to the search margin of nonshoppers, which is constant in the sequential-search regime: $E[p_{min}] = \kappa + s(1 - \phi)/\phi$, independent of the number of sellers.

Proposition 4 is a statement within each regime. Because T_1 falls with N , entry can move a market from the sequential-search to the no-search regime, and across this switch the rate to the expected minimum price can fall, since its level coefficient drops from one to $T_N < 1$ (Appendix B.3). Corollary 1 below is unaffected by the switch.

Propositions 3 and 4 combine into a prediction that does not require knowing the regime.

Corollary 1. *With $0 < \phi < 1$ and $s > 0$, the difference between the two rates, $\rho_\tau^{min} - \rho_\tau^{avg}$, is strictly increasing in N on $N \geq 2$, within each regime and across the switch between them.*

In the no-search regime, entry moves the two rates in opposite directions. Pass-through to the price that informed consumers pay rises, while pass-through to the price that uninformed consumers pay on average falls. In the sequential-search regime, the rate to the minimum price is flat and the rate to the average price falls, so the gap, which we call the pass-through gap, widens there as well. Combining the two moments with Equation (1) shows, in addition, that the number of sellers drops out entirely from the average price paid in the market.

Corollary 2. *Let $\bar{p} = \phi E[p_{min}] + (1 - \phi) E[p]$ denote the average transacted price. For every $N \geq 2$, $0 < \phi < 1$, and $s > 0$, the average transacted price responds to any cost or tax shifter by $d\bar{p}/d\kappa = \phi$ in the no-search regime and $d\bar{p}/d\kappa = 1$ in the sequential-search regime, independent of the number of sellers in either case.*

Equation (1) implies $\bar{p} = \phi \kappa + (1 - \phi) p_r$. The average transacted price is a ϕ -weighted average of the cost floor and the reservation price. In the no-search regime, the reservation price is fixed at the valuation, so an increase in the tax-inclusive cost κ raises the average transacted price by ϕ times the change in κ , for any number of sellers. In the sequential-search regime, the reservation price translates with cost and $d\bar{p}/d\kappa = 1$ regardless of ϕ .

In our empirical setting, a price setter can operate several stations within a local market. Appendix B.4 extends the model to multi-station sellers and shows that the propositions carry over, with the competition index based on stations. The comparative statics in the number of sellers require a constant number of stations per seller. Two additions matter for the empirical analysis. Only a seller's cheapest station competes for the informed demand, so the market minimum across stations is the minimum across price setters, and the number of price setters is the empirical counterpart of N . Noncontesting stations price at the reservation price and dilute the station average, so at a given number of sellers, a denser station network widens the pass-through gap.

3.5 Extensions

Two assumptions give the model its sharpest edges: consumers have inelastic unit demand, and shoppers compare every seller in the market. Retail fuel markets have neither feature exactly. Aggregate fuel demand slopes downward, and stations differ in location and visibility, so even informed consumers compare the sellers relevant to their route rather than all of them. Appendix B.5 relaxes the two assumptions and traces the resulting pass-through profiles numerically. Here we summarize how they change the predictions.

With downward-sloping demand, the upper anchor of the price distribution is no longer the fixed valuation but the monopoly price, which itself moves with the tax-inclusive cost. Pass-through at the two extremes of the profile, monopoly and the many-sellers limit, then equals the positive monopoly rate rather than zero. The hump of Proposition 3 survives with its peak at duopoly, but the post-peak decline is compressed. It runs toward the monopoly rate, and its depth is governed by the curvature of demand.²⁵

With imperfect comparability, modeled as each shopper comparing each rival only with some probability, the contest for the informed demand builds gradually rather than arriving in full with the first competitor. The peak of the profile moves beyond duopoly, the more so the lower the comparability, and the profile flattens as it stretches.

4 Policy Changes and Descriptive Evidence

We analyze several tax changes in the German and French retail fuel markets between 2020 and 2023 to test the predictions of Section 3. This section describes the tax changes and presents descriptive evidence on their pass-through.

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 eurocents per liter was levied on gasoline and 47.04 eurocents per liter on diesel in Germany. In France, the lump-sum fuel tax varied by region, ranging from 67 to 70 eurocents per liter for gasoline (and approximately 61 eurocents per liter for diesel). In addition, Germany and France have a VAT of 19% and 20%, respectively, levied on the tax-inclusive price of fuel.

²⁵The hump requires the monopoly pass-through rate to lie below the pass-through of the cost-anchored low prices, which holds for log-concave demand (Weyl and Fabinger, 2013). With log-convex demand, the hump disappears, and with constant-elasticity demand the rate equals one for every N (Appendix B.5). Garrod et al. (2026) analyze cost pass-through in the no-search regime under general demand and characterize how the shift of the price distribution depends on whether demand is log-concave or log-convex.

The first tax change was a temporary reduction of the German VAT from 19% to 16% between July and December 2020, part of the fiscal response to the COVID-19 pandemic. 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 euros per emitted tonne of CO₂ on oil, gas, and fuel. For *E5* and *E10*, this translates into a per unit tax of 6.00 eurocents per liter (7.14 eurocents including VAT). For diesel, the per unit tax is 6.69 eurocents per liter (7.96 eurocents including VAT).

We cannot separately identify the pass-through of the simultaneous VAT increase and the introduction of the carbon price on 1 January 2021, so we estimate their joint pass-through. This is unproblematic for the theoretical predictions, since in the model both taxes affect prices only through the tax-inclusive marginal cost.

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, the fuel tax on gasoline and diesel fell by 18 eurocents per liter. This rebate increased to 30 eurocents between 1 September and 15 November 2022. It then dropped to 10 eurocents between 16 November and 31 December 2022, before being phased out on 1 January 2023.

Germany also implemented a temporary tax rebate, which we exclude from our analysis because of the intense public scrutiny and the 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.²⁶


4.2 Descriptive evidence on heterogeneous pass-through

Before turning to the econometric analysis, we descriptively study the pass-through of the 2020/21 tax changes in Germany by comparing fuel prices over time in Germany and France. We can thus see whether pass-through differs across fuel types, for which the evidence in Section 2 suggests different shares of informed consumers.

Panel A of Figure 1 displays nonparametric estimates of the pass-through rate by fuel type for the 2020 temporary VAT reduction in Germany. 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 pretrends. The pass-through rate is 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 approximately two weeks.

²⁶In Appendix C, we present additional descriptive evidence showing that the 2022 tax changes in Germany are not suitable for our analysis.

Figure 1: 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 nonparametric estimate of the daily average pass-through rate to prices for *E5*. The short-dashed and long-dashed lines show analogous estimates for *E10* and diesel, respectively. To estimate pass-through, we demean the daily average price in each country by its average pre-period price (1 May to 30 June 2020 for the tax decrease, Panel A; 1 November to 15 December 2020 for the tax increase, Panel B), take the difference between Germany and France, and divide it by the difference under full pass-through. Because we divide by the absolute size of the tax change, full pass-through corresponds to -100% for the tax decrease and $+100\%$ for the tax increase, whereas the pass-through rates reported in the text and tables are defined so that full pass-through equals $+100\%$ for both tax changes. For the tax decrease, full pass-through would correspond to a price drop of 2.52% . For the tax increase, full pass-through would correspond to a price increase of 2.59% from the VAT increase, plus the newly introduced carbon price of 7.14 eurocents per liter (including VAT) for *E5* and *E10* and 7.96 eurocents per liter for diesel, which we convert into a percentage using the average German price in the week from 9 to 15 December 2020. The vertical solid line marks the starting date of the tax change.

Panel B of Figure 1 presents nonparametric estimates of the pass-through rate by fuel type for the subsequent tax increase in winter 2020/21. Unlike for the tax decrease, there appear to be anticipatory effects in the pass-through of the tax increase in the last two weeks of December. Therefore, we 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 afterward. Again, pass-through is highest for diesel, consistent with the theoretical predictions. Differences in pass-through between *E5* and *E10* appear less pronounced.

5 Empirical Strategy

5.1 Difference-in-differences strategy

To estimate the average rate of pass-through of the tax changes to 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.²⁷ Our baseline specification is the following DID model:

$$Y_{it} = \beta \text{Tax}_{it} + \gamma X_{it} + \alpha_i + \pi_t + \epsilon_{it}, \quad (2)$$

where Y_{it} is the logarithm of the weighted average price of gasoline or diesel at a fuel station i at date t . We compute this daily average price by weighting a station's price at each hour between 6 am and 10 pm, using data on hourly fueling patterns (see Appendix A). Tax_{it} equals one for fuel stations in Germany from 1 July 2020 onward for the tax reduction and from 1 January 2021 onward for the subsequent tax increase. X_{it} is a vector of controls. Finally, α_i and π_t are fuel station and date fixed effects. For the VAT component, the pass-through rate implied by β is the model's pass-through rate ρ_τ from Section 3. We report pass-through rates as β divided by $\ln(1 + f)$, where f is the proportional price change that full pass-through of the tax change implies. A rate of 100% therefore corresponds to full pass-through.

5.2 Stations in neighboring country as a control group

Two assumptions must hold 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 France before and after the tax change, other than the policy change itself. Second, there

²⁷Conversely, 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.

should be no spillover effects from the tax changes onto the fuel market in the neighboring country. Both are likely to be satisfied for the 2020/21 tax changes, but less so for 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. Moreover, we restrict our analysis to relatively narrow time windows around the reforms, which limits the scope for transitory shocks that differently affect German and French stations in 2020/21.

The most obvious threats to identification concern transitory demand and supply shocks. On the demand side, 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 on 15 June 2020. On the supply side, fuel stations in Germany and France procure most of their refined oil from similar sources, and the European single market harmonizes the movement of goods between the two countries.

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. 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 control group.

Focusing on tax changes in opposite directions further reduces concerns about confounders. A transitory shock that inflates the estimated pass-through for one fuel type during the July 2020 decrease would have to reverse sign in January 2021 to produce the same ranking for the increase. Finding consistent differences across fuel types for both tax changes therefore suggests that we estimate actual differences in pass-through.

For 2022/23, the picture is different. First, multiple tax changes 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. This may lead to spillover effects and break 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 to the global oil market in 2022.²⁸ We therefore do not interpret the magnitude of pass-through for the 2022 tax changes or the differences between fuel types. These tax changes remain informative

²⁸Appendix C presents supporting evidence. The German and French diesel and gasoline markets started diverging right after the start of Russia’s invasion and before any tax change was announced, and margins in France increased immediately after Germany’s large fuel tax cut on 1 June 2022, suggesting spillover effects.

about the difference in pass-through to the average posted price and the minimum price within a given fuel type.

5.3 Testing the propositions empirically

To test Proposition 1, we estimate the baseline DID model in Equation (2) 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.²⁹ According to the theoretical model and the specificities of the German industry, we expect pass-through to be highest for diesel and lowest for *E5*. Since Proposition 1 concerns every quantile of the distribution of equilibrium prices, this should hold for any moment of the price distribution.

To test Proposition 2, we estimate tax pass-through to the average price and the minimum price in a market for the 2020/21 tax changes and the 2022/23 French tax changes. We now include the French tax changes because we compare pass-through to the average price and the minimum price *within* each fuel type. To remain as close as possible to the expected price and the expected minimum price in the theoretical model, we take the cross-station average and the cross-station minimum of the prices posted at a fixed hour, 5 pm, within nonoverlapping geographic markets in Germany and France. Fixing the hour makes this a comparison across sellers at a point in time, as in the model, rather than one that also reflects the intraday price cycle. We then estimate the following triple DID variation of our baseline model:

$$Y_{jts} = \beta_1 \text{Tax}_{jt} + \beta_2 \text{Tax}_{jt} \cdot \text{Min}_s + \alpha_{js} + \pi_{ts} + \epsilon_{jts}, \quad (3)$$

where Y_{jts} is the logarithm of the 5 pm average or minimum price of gasoline or diesel in market j at date t , as indicated by the price type s , and Min_s equals one for the minimum price. Tax_{jt} equals one for markets affected by the tax change at date t , and α_{js} and π_{ts} are market $\times$ price type and date $\times$ price type fixed effects. Divided by the full-pass-through benchmark, β_1 yields the pass-through rate to the average posted price, and β_2 yields the difference in the pass-through rate between the minimum price and the average posted price in a local market. Because this benchmark carries the sign of the tax change, Proposition 2 predicts a positive rate difference for all fuel types and all tax changes, that is, $\beta_2 > 0$ for the tax increases and $\beta_2 < 0$ for the tax decreases.

To test Proposition 3, we estimate tax pass-through for the 2020/21 tax changes at the station level. We first estimate a pass-through rate for every station in Germany for each fuel type. For each station, we re-estimate Equation (2) adding an interaction between the

²⁹ Additionally, in Appendix D.4, we run a pooled regression with all three fuel types. This allows a formal test of whether pass-through differs significantly across fuel types.

treatment indicator and an indicator for that station, and take the sum of the coefficients on the treatment indicator and on the interaction as the station’s estimate. We then group stations by the number of competing price setters in a local nonoverlapping market and nonparametrically regress the station-level pass-through rate on indicators for these groups. The multi-station extension of the model (Section 3 and Appendix B.4) shows that the predictions extend to price setters operating several stations. This lets us compare pass-through rates *within* fuel type, holding an important source of variation in price sensitivity fixed. A perfect-information model with log-concave demand of constant curvature predicts that the average pass-through rate increases monotonically in the number of competitors. Proposition 3 instead predicts a nonmonotonic relationship. As we do not observe brands or price setters for France, we can test Proposition 3 only for German tax changes.

To test Corollary 1, we combine the two previous exercises. Proposition 4 predicts that pass-through to the minimum price rises with the number of sellers in the no-search regime and is flat in the sequential-search regime. This prediction is difficult to test, as we do not observe the regime a market is in. The corollary combines it with Proposition 3 into a prediction that holds in both regimes: the difference between pass-through to the minimum and to the average price, the pass-through gap, widens in the number of sellers.

We estimate the triple DID model in Equation (3) with the treatment interactions specified separately for each market-size bin, the markets with a given number of competing price setters, with monopoly markets split by whether the monopolist operates one or several stations, so that β_2 delivers the pass-through gap for each bin. Because the minimum and average price coincide by construction in single-station markets, the bin of single-station monopoly markets serves as a placebo, and we report each bin’s gap relative to it. Estimating all bins jointly, for each fuel type and tax change, makes the standard errors of this correction valid despite the shared control group. French markets enter as a single control group that is not split by market size. Their before-after change in the gap is therefore common to all bins and drops out of the placebo correction, so the corrected profile is the change in the German gap by market size relative to the placebo bin.

Appendix D.12 estimates the same system on pseudo reform dates inside the pre-periods. We estimate the gap on the 5 pm data, for the German tax changes with markets grouped by the number of competing price setters and, in the appendix, for the French tax changes with markets grouped by the number of stations.

5.4 Robustness checks

We verify that our empirical results are robust to alternative specifications. The details are in the appendix.

Our results survive the inclusion of additional controls (Appendix D.1): demand-related shocks proxied by the Google Mobility Report for the baseline pass-through rates, and demographic characteristics such as income per capita and population density for the pass-through by number of sellers. They also survive restricting the sample to a balanced panel of stations with a price observation every day (Appendix D.2), which rules out that temporary closures drive the findings.³⁰ Our preferred specification accounts for anticipatory effects in winter but not in summer, following the descriptive evidence in Figure 1; the results are robust to changing this assumption (Appendix D.3). A synthetic difference-in-differences strategy (Arkhangelsky et al., 2021), which reweights control stations to match pre-treatment price trends and thereby addresses parallel-trends concerns directly, delivers the same conclusions, except that the difference between *E5* and *E10* in the tax increase is no longer statistically significant (Appendix D.5).

For the minimum-price analyses, we address the concern that the market minimum reflects short-lived outlier prices by re-estimating the gap with the price at the 10th percentile of the within-market price distribution (Appendix D.6). This leaves the differential pass-through to the minimum price intact. A placebo on pseudo reform dates inside the pre-periods produces gaps well below the estimated ones for the tax increase (Appendix D.12). Appendix D.7 reports the pass-through gap for the French tax changes by the number of sellers. Finally, Appendix D.8 repeats the analysis with overlapping markets centered on each station, defined by a 5 km driving-distance radius, with the same conclusions.

6 Results

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 (2).

Columns (1) to (3) show that the tax decrease lowered prices for all fuel types. Under full pass-through, we expect prices for each fuel product to decrease by approximately

³⁰Entry and exit of stations have a minimal impact in our setting. In the sample period for the analysis of the tax decrease (increase), only 0.34% (0.37%) of the German stations that reported at least one price newly entered the market, while 0.23% (0.33%) exited.

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% [42%, 49%]	93% [91%, 96%]	70% [69%, 71%]	75% [74%, 75%]	88% [87%, 88%]
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 (2). 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, using data from 1 May to 31 August 2020. Columns (4) to (6) present average treatment effect estimates of the German VAT increase and introduction of a carbon price on 1 January 2021 on the same outcomes, using data from 1 November to 15 December 2020 and from 1 January to 28 February 2021. Standard errors clustered at the market level are shown in parentheses. The implied pass-through rates, β divided by the log-point price change that full pass-through implies, are reported with their 95% confidence intervals in brackets.

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

2.52%.³¹ We estimate that 28% of the tax decrease is passed on to consumers fueling *E5* while the pass-through rates for *E10* and diesel are 46% and 93%, 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% for diesel.³² We find a joint pass-through rate of the tax increases of 70% for *E5*, 75% for *E10*, and 88% for diesel.

This comparison isolates the first of the two channels of competition set out in Section 3. It varies the share of informed consumers across fuel types while holding the station network, and hence the number of sellers, fixed. Our findings are consistent with Proposition 1, which establishes 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*. The pooled regression in Appendix D.4 confirms that these differences are statistically significant,

³¹With a decrease in the VAT rate from 19% before the VAT decrease to 16% after the policy change, this is $\frac{1.16-1.19}{1.19} \cdot 100 \approx -2.52\%$.

³²Under full pass-through, a change in the VAT rate from 16% to 19% would increase the fuel price by $\frac{1.19-1.16}{1.16} \cdot 100 \approx 2.59\%$. 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 2: Pass-through to market-level minimum vs. average prices



Notes: The figure shows the difference in the market-level pass-through rate between the minimum price and the average posted price. The average posted price is the mean of the 5 pm prices of all stations in a nonoverlapping market on a given day, and the minimum price is the lowest of them. Markers show the difference implied by the DID estimate β_2 in Equation (3), estimated separately for each tax change and fuel type, with 95% confidence intervals based on standard errors clustered at the market level. We use data for the two months before and after each tax change, except that we exclude the second half of December 2020 for the German tax increase (anticipatory effects), and the post-treatment period on 31 December 2022 for the French tax increase on 16 November 2022, and use 16 November to 31 December 2022 as the pre-treatment period for the French tax increase on 1 January 2023. Solid (hollow) symbols indicate point estimates that are statistically significant (insignificant) at the 5% level. Crosses mark the uncorrected placebo estimates from single-station monopoly markets.

even with $\text{station} \times \text{date}$ fixed effects. Since almost all German fuel stations sell all three fuel types, the differences in pass-through are very unlikely to be driven by supply-side factors.

6.2 Pass-through to the average and minimum price

Figure 2 shows the difference in the pass-through rate between the minimum price and the average posted price within nonoverlapping geographic markets.³³

Pass-through to the minimum price is significantly higher than to the average posted price for every fuel type in each of the three tax increases. Single-station monopoly markets provide a placebo test of the design. Their minimum and average price coincide by construction, so an estimated gap in these markets can only reflect differential trends in the control group’s price dispersion. Figure 2 marks these placebo estimates. For the three increases, they are close to zero, whereas for the two decreases they are of the same order as the estimated gaps themselves, indicating that the control group’s dispersion shifted over

³³Market-level minimum and average prices evolved similarly in Germany and France before the tax changes. Repeating the nonparametric estimation of Section 4 with market-level prices yields results very similar to Figure 1 (Appendix C.1).

the COVID-affected decrease windows. We therefore do not read the decrease-episode gaps as informative about the prediction.

The increase episodes therefore support Proposition 2. The results suggest that informed consumers, who typically buy fuel at prices closer to the within-market minimum, bear a larger share of the tax on the price they pay than uninformed consumers, who buy at the average posted price (Appendix D.11).

6.3 Number of sellers and tax pass-through

Turning to the second channel, the number of sellers, Figure 3 shows the relationship between the average pass-through rate and the number of competing price setters of a focal station for the 2020/21 German tax changes and the three fuel types.

Panels A, C, and E show the tax decrease in summer 2020. For *E5* and *E10*, the average pass-through rate is relatively low for local monopolists and significantly higher in markets with two competing price setters, but it does not increase any further beyond two. The fitted lines point to a mildly hump-shaped relationship, in line with the peak at duopoly and subsequent decline that Proposition 3 establishes. For diesel, pass-through differs little between monopolistic and competitive markets, but again it does not monotonically increase in the number of competitors.


Panels B, D, and F repeat the analysis for the tax increase in winter 2020/21 and yield similar relationships for all fuel types. For *E5* and *E10*, pass-through is again low for monopolists and then flat or mildly decreasing beyond two sellers. For diesel, it increases mildly up to about five competing price setters before leveling off, the flattened and right-shifted profile that imperfect comparability produces.

Our findings differ from Genakos and Pagliero (2022), who find that pass-through monotonically increases in the number of sellers. Notably, theirs is a setting with infrequent price changes, where imperfect information likely plays less of a role.

In Appendix D.1 we show that our findings remain unchanged when we control for demographic characteristics, such as county-level income per capita or population density.

In Appendix D.9, we formally test for nonmonotonicity using the U-test of Lind and Mehlum (2010), which tests a monotone or U-shaped relationship against an inverted-U shape. The estimated coefficients satisfy the necessary conditions for an inverted U for every fuel type and tax change. With standard errors clustered at the market level, the test rejects a monotone relationship at the 10% level for *E5* and *E10* in the tax increase, and at the 5% level for *E5* once we add demographic controls. In the remaining cases the hump is not statistically distinguishable from a monotone profile.

Figure 3: Average pass-through by number of competing price setters



Notes: The figure shows how the rate of pass-through to the average price varies with the number of competing price setters in a market. Panels A, C, and E show the German VAT decrease on 1 July 2020, and Panels B, D, and F the German VAT increase and introduction of a carbon price on 1 January 2021, for $E5$, $E10$, and diesel, respectively. Each circle plots the average pass-through rate for stations with a given number of competing price setters within a nonoverlapping local market, with circle size proportional to the number of such stations. The 95% confidence intervals come from a nonparametric regression of station-level pass-through rates on market-size indicators, with monopoly markets the omitted category and standard errors clustered at the market level. The solid and dashed lines show fractional polynomial and quadratic fits. Stations in markets above the 97.5th percentile of the number of competing price setters are excluded.

While the evidence cannot definitively rule out a monotonically increasing relationship, it is more consistent with a hump-shaped profile. The findings are therefore consistent with Proposition 3, which establishes that pass-through to the expected price is nonmonotonic in the number of sellers.³⁴ The finer features of the estimated profiles, an interior peak at a moderate number of sellers, a mild decline beyond it, and a level well above zero throughout, line up with the extended model of Section 3.5.

6.4 Number of sellers and the pass-through gap

We now turn to the number of sellers and the minimum price. Corollary 1 predicts that the pass-through gap widens in the number of sellers in both regimes of the model.


Figure 4 plots the pass-through gap by the number of competing price setters within a nonoverlapping market for the two German tax changes and the three fuel types. Single-station monopoly markets provide the placebo bin discussed in Section 6.2. As in Section 5, we estimate all market-size bins jointly and report each bin’s gap relative to the placebo bin.

For the tax increase in Panel B, the pattern matches the prediction. The placebo bin is close to zero, and the corrected gap widens with the number of competing price setters for all three fuel types, from about +2 percentage points in duopoly markets to between +3 and +6 percentage points in the largest markets, all significant at the 5% level. The estimates track the uniform-level-change benchmark and exceed it moderately for diesel and *E10*. Both features are consistent with the model. The uniform-level-change benchmark is the gap that arises when every price moves by the same euro amount, as in the sequential-search regime, and the compression of the minimum price toward the cost floor in the no-search regime adds to it. A specification in euro levels, in which a uniform level change implies a gap of exactly zero, isolates this compression. It is statistically significant in most bins but only about one percentage point for diesel and *E10* (Appendix D.11). A pseudo reform date in the autumn of 2019 produces corrected gaps within half a percentage point of zero, significantly below the actual gaps in every bin (Appendix D.12).

For the tax decrease in Panel A, the design fails its placebo: single-station monopoly markets show gaps of five to ten percentage points before correction, reflecting differential

³⁴The empirical content we test is the nonmonotone shape, not the proposition’s two point predictions, which are knife-edge properties of the model: with inelastic unit demand for a homogeneous good, a monopolist absorbs the tax fully, and the peak falls exactly at duopoly. Empirical monopoly markets face residual competition from stations beyond the market boundary and horizontally differentiated demand, so their pass-through is low but not zero in Figure 3, and travel costs can shift the peak beyond two price setters. Section 3.5 and Appendix B.5 make both points formal: with downward-sloping demand, the profile starts at and declines toward the positive monopoly rate, and with imperfect comparability, the peak moves beyond duopoly.

Figure 4: Pass-through gap by number of sellers



Notes: The figure shows how the placebo-corrected difference in the market-level pass-through rate between the minimum price and the average posted price varies with the number of competing price setters in a nonoverlapping market. All market-size bins are estimated jointly, and each bin's gap is reported relative to the bin of single-station monopoly markets, shown at zero at $N = 1$. This correction removes differential trends in the control group's price dispersion. The second marker at $N = 1$ shows multi-station monopoly markets, for which the extended model also predicts a zero gap (Lemma B5(i)). Panel A depicts the German VAT decrease on 1 July 2020 and Panel B the German VAT increase and introduction of a carbon price on 1 January 2021. In each panel, the circles, diamonds, and triangles plot the corrected gap for diesel, E10, and E5, respectively, grouped by the number of competing price setters (1 to 5, and 6 or more). Vertical spikes show 95% confidence intervals based on standard errors clustered at the market level. Solid (hollow) markers indicate estimates that are statistically significant (insignificant) at the 5% level. The dashed lines show the uniform-level-change benchmark, the gap implied by the pre-reform within-market price dispersion if every price in a market moved by the same euro amount along the observed price-level path.

trends in the French control group's price dispersion over the COVID-affected summer of 2020. The corrected gaps are negative throughout for the gasoline products and close to zero for diesel except in the largest markets. About half of this negative profile is the uniform-level-change benchmark, whose sign flips because German price levels recovered over the window; the remainder reflects a compression of within-market price dispersion in large German markets as conditions normalized. The decrease episode's gasoline price levels are also sensitive to the crude-oil control (Appendix D.10). We therefore read the decrease episode as uninformative about Corollary 1.

The three French tax changes corroborate the increase result. Appendix D.7 reports the 5 pm by seller count analysis for France, where markets are grouped by the number of stations because price-setter identity is not observed. Since a price setter can operate several stations, the station count bounds the number of sellers from above, so the French profile is a coarser proxy for the model's N than the German grouping. For both French tax increases, the gap widens monotonically in the number of stations for all three fuel types, from close to zero in single-station markets, which double as the placebo bin in the French grouping, to between +9 and +33 percentage points in the largest markets. The French tax decrease, like its German counterpart, is small and mixed.

Taken together, the three tax increases, in Germany and in France, support Corollary 1. The more sellers a market has, the more pass-through to the price informed consumers pay exceeds pass-through to the average price paid by uninformed consumers, so entry widens the difference in tax incidence between the two groups within a market.

7 Discussion

Our analysis shows that a model of imperfect consumer price information accounts jointly for the empirical patterns. This section asks whether alternative explanations do as well. We first consider full-information models and then alternative ways of modeling imperfect information.

7.1 Alternative hypotheses with full information

Before turning to specific alternatives, we note the benchmark most of the tax pass-through literature uses for imperfect competition. The conduct parameter approach of Weyl and Fabinger (2013) nests most full-information oligopoly models with symmetric sellers, including symmetrically differentiated Nash-in-prices, a natural description of retail fuel competition.³⁵ With constant marginal costs and conduct invariant to quantities, assumptions that Genakos and Pagliero (2022) argue are reasonable for retail fuel, pass-through in this benchmark depends only on the conduct parameter and the curvature of demand.

Demand-side differences across fuel types. Under full information, differences in pass-through between fuel types could arise from differences in the curvature of aggregate demand, in its price elasticity, or in consumers' price sensitivity. Each of these channels may explain part of the cross-fuel ordering. None explains the within-market patterns. With symmetric sellers and full information, there is no within-market price dispersion and no gap between pass-through to the minimum and the average price. In that benchmark with log-concave demand of constant curvature, more competitors raise pass-through monotonically, whereas we find no systematic increase beyond duopoly. Demand curvature is also not an alternative to our mechanism. Section 3.5 incorporates downward-sloping demand into the model, and the interior peak and the decline toward the monopoly rate survive, whereas under full information with log-concave demand of constant curvature the profile rises monotonically.

Vertical differentiation. Some stations may simply be better than others, for example because of their location. Such stations face a lower price elasticity of residual de-

³⁵The approach goes back to Bresnahan (1989) and Genesove and Mullin (1998).

mand even under perfect information and therefore have lower pass-through. Combined with a reason for differences across fuel types, this may explain some of the empirical patterns. However, this would not explain the nonmonotonic relationship between the number of competitors and pass-through. Most importantly, vertical differentiation cannot explain within-market price dispersion after controlling for market $\times$ date and station fixed effects, as documented in Table 2. Appendix Figure A6 illustrates for one market how station pricing produces undercutting, intraday price cycles, and rank reversals, patterns vertical differentiation cannot generate.

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

7.2 Alternative hypotheses with imperfect information

Rockets and feathers. Tappata (2009) endogenizes the decision to become informed in a Varian (1980) framework with cost uncertainty and shows that pass-through can be asymmetric because search, and with it price dispersion, is high when costs and prices are low. The mechanism therefore requires price dispersion to be higher when prices are low. In Appendix A.4, we show that search intensity and price dispersion are uncorrelated or mildly positively correlated with the price level in our data. The premise of the model does not apply in our setting.

Reference price dependence. In the duopoly search model of Lewis (2011), consumers search when a drawn price exceeds a reference price formed from past prices, so cost increases that consumers do not observe trigger search and are passed through faster than decreases. A publicized tax change moves the reference price itself and switches this search response off. Applied to our setting, the model thus predicts no differential pass-through of the tax changes, and any variation across fuel types would again have to come from differences in the share of price-comparing consumers, our mechanism.

Unobservable input costs. Janssen, Pichler, and Weidenholzer (2011) and Janssen and Shelegia (2015) extend the Stahl (1989) model to input costs that consumers do not observe, so nonshoppers cannot condition their behavior on the realized cost, which raises price levels. For cost shocks that consumers do observe, these models reduce to the Stahl (1989) benchmark underlying our model. Since the tax changes we study were broadly publicized, with known timing, direction, and magnitude, the unobservable-cost extension adds no predictions in our setting. Janssen and Shelegia (2020) study pass-through with imperfect information and differentiated products.

Although one can model imperfect information in the retail fuel market differently, none of the obvious alternatives better explains our empirical results, and we cannot formally test between models. Taken together, however, the findings are difficult to reconcile with models that do not feature imperfect consumer information, and our model provides a parsimonious joint account of them.

8 Conclusion

In this paper, we show that imperfect price information helps explain heterogeneity in tax pass-through. When consumers do not know all prices, firms have market power, and the extent of that market power determines how much of a tax reaches consumers.

Three results stand out and set our analysis apart from one with perfect information: First, the more well-informed consumers there are, the higher is tax pass-through. Second, taxes are passed through more to the price paid by well-informed consumers than to the price paid by uninformed consumers, a result we establish with the three tax increases we study. Third, pass-through to the average price does not rise systematically with the number of sellers, whereas the gap between pass-through to the minimum and to the average price widens with it.

These results bear on policies that work through prices. Unconventional fiscal policy relies on consumers expecting a pre-announced tax increase to reach the prices they pay (D’Acunto, Hoang, and Weber, 2018, 2022), and a Pigouvian tax changes behavior only to the extent that it is passed through. Both therefore pass through more fully the larger the share of well-informed consumers, and both reach informed consumers more than uninformed ones. The same holds for any common cost shock. Better-informed consumers lower price levels but raise the share of a cost change that reaches prices.

The gap between the average and the minimum price in a market is observable wherever price panels exist, and its response to a common cost shock is a testable implication of how much consumer information constrains sellers. The mechanism we study is therefore testable in any market where consumers differ in how well they know prices, well beyond retail fuel.

References

- Adachi, Takanori, and Michal Fabinger. 2022. “Pass-Through, Welfare, and Incidence Under Imperfect Competition.” *Journal of Public Economics* 211:104589.
- Albrecht, Brian C., and Mark Whitmeyer. 2026. “Pass-through with Price Dispersion.” *arXiv preprint arXiv:2601.17964*.
- 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. 2024. “Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market.” *Journal of Political Economy* 132 (3): 723–771.
- Baye, Michael R., Dan Kovenock, and Casper G. de Vries. 1992. “It Takes Two to Tango: Equilibria in a Model of Sales.” *Games and Economic Behavior* 4 (4): 493–510.
- 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, 323–375. Elsevier.
- Bresnahan, Timothy F. 1989. “Empirical Studies of Industries with Market Power.” Chap. 17 in *Handbook of Industrial Organization*, 2:1011–1057. Elsevier.
- Burdett, Kenneth, and Kenneth L. Judd. 1983. “Equilibrium Price Dispersion.” *Econometrica* 51 (4): 955–969.
- 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. 2021. “An Experimental Investigation of Price Dispersion and Cycles.” *Journal of Political Economy* 129 (3): 789–841.
- 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.
- Conlon, Christopher, Nirupama Rao, and Yinan Wang. 2024. “Who Pays Sin Taxes? Understanding the Overlapping Burdens of Corrective Taxes.” *Review of Economics and Statistics* 106 (6): 1719–1729.

- 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.
- Dana, James D., Jr. 1999. “Equilibrium Price Dispersion Under Demand Uncertainty: The Roles of Costly Capacity and Market Structure.” *RAND Journal of Economics* 30 (4): 632–660.
- Diamond, Peter A. 1971. “A Model of Price Adjustment.” *Journal of Economic Theory* 3 (2): 156–168.
- Duso, Tomaso, and Florian Szücs. 2017. “Market Power and Heterogeneous Pass-Through in German Electricity Retail.” *European Economic Review* 98:354–372.
- Fischer, Kai, Simon Martin, and Philipp Schmidt-Dengler. 2025. “The Heterogeneous Effects of Entry on Prices.” *Journal of the European Economic Association*, <https://doi.org/10.1093/jeea/jvaf061>.
- . 2026. “Indirect Taxation in Consumer Search Markets: The Case of Retail Fuel.” *Working Paper*.
- Garrod, Luke, Ruochen Li, Antonio Russo, and Chris M. Wilson. 2026. “Understanding Cost Pass-Through when Prices are Dispersed.” *Economic Journal*, <https://doi.org/10.1093/ej/ueag097>.
- 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.
- 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.
- 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.
- Huber, Stephan, and Christoph Rust. 2016. “Calculate Travel Time and Distance with OpenStreetMap 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.
- Janssen, Maarten C. W., and José Luis Moraga-González. 2004. “Strategic Pricing, Consumer Search and the Number of Firms.” *Review of Economic Studies* 71 (4): 1089–1118.
- 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.pdf?__blob=publicationFile&v=1.
- Kroft, Kory, Jean-William P. Laliberté, René Leal-Vizcaíno, and Matthew J. Notowidigdo. 2024. “Efficiency and Incidence of Taxation with Free Entry and Love-of-Variety Preferences.” *American Economic Journal: Economic Policy* 16 (2): 300–334.
- Lemus, Jorge, and Fernando Luco. 2021. “Price Leadership and Uncertainty about Future Costs.” *Journal of Industrial Economics* 69 (2): 305–337.
- 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.
- 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 (9–10): 1202–1212.
- Martin, Simon. 2024. “Market Transparency and Consumer Search – Evidence from the German Retail Gasoline Market.” *RAND Journal of Economics* 55 (4): 573–602.

- McAfee, R. Preston. 1995. “Multiproduct Equilibrium Price Dispersion.” *Journal of Economic Theory* 67 (1): 83–105.
- 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. 2026. “When Does Mandatory Price Disclosure Lower Prices? Evidence from the German Fuel Market.” *Journal of Political Economy Microeconomics* 4 (1): 106–146.
- Morgan, John, Henrik Orzen, and Martin Sefton. 2006. “An Experimental Study of Price Dispersion.” *Games and Economic Behavior* 54 (1): 134–158.
- 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.
- Rosenthal, Robert W. 1980. “A Model in Which an Increase in the Number of Sellers Leads to a Higher Price.” *Econometrica* 48 (6): 1575–1579.
- Rothschild, Michael, and Joseph E. Stiglitz. 1970. “Increasing Risk: I. A Definition.” *Journal of Economic Theory* 2 (3): 225–243.
- Stahl, Dale O. 1989. “Oligopolistic Pricing with Sequential Consumer Search.” *American Economic Review* 79 (4): 700–712.
- 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.
- Westphal, Ryan M. 2025. “What You Don’t Know Can’t Pass Through: Consumer Beliefs and Pass-through Rates.” *Working Paper*.
- 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


A.1 Construction of 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 changes, along with the precise time and date stamp for every change. On average, in 2019, fuel stations in Germany changed 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 reweight the hourly shares to produce weights for the hours between 6 am and 10 pm.

In Germany, taxes and duties consist of VAT, a lump-sum energy tax, and a fee for oil storage. The lump-sum energy tax is 65.45 eurocents per liter for *E5* and *E10* gasoline, and 47.04 eurocents per liter for diesel. The fee for oil storage is 0.27 eurocents per liter for *E5*

Figure A1: Daily fueling patterns (Germany)



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

and *E10* and 0.30 eurocents per liter for diesel.³⁶ 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 eurocents per liter, depending on the metropolitan region and fuel type.³⁷ In addition, the French VAT rate on retail fuel is 20%.

We impose a few restrictions on which fuel stations we include in our analysis. In Germany, we drop stations located on highways (i.e., “Autobahn”) because prices at these stations are typically approximately 20 to 30 eurocents higher than those at regular fuel stations. We identify highway stations based on their address and manual checks. In France, we focus only on stations in mainland France (i.e., we exclude stations on the island of Corsica and overseas). As in Germany, we exclude French stations located on motorways, which the French price data flag directly.

A.2 Nonoverlapping markets

To group fuel stations into nonoverlapping local markets, we use an agglomerative hierarchical clustering algorithm based on the driving time between stations. This approach follows Carranza, Clark, and Houde (2015), Lemus and Luco (2021), and Assad et al. (2024).

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.³⁸ Eventually, all stations are combined into a single cluster.

The “height” of each link, that is, the average driving time needed to join two clusters, is informative about natural separations between groups of stations. We compute an inconsistency coefficient for each link, which captures the height of the current link relative to the heights of previous links.³⁹ If the driving time required to link two clusters is much higher


³⁶See <https://www.avd.de/kraftstoff/staatlicher-anteil-an-den-kraftstoffkosten/>.

³⁷See <http://www.financespubliques.fr/glossaire/terme/TICPE/>.

³⁸See Appendix C in Lemus and Luco (2021) for an example and illustration.

³⁹See http://cda.psych.uiuc.edu/multivariate_fall_2012/matlab_help/cluster_analysis.pdf for additional details.

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 nonoverlapping markets by means of 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.


than the driving time between stations *within* these clusters, the two clusters are likely to represent separate local markets, and the inconsistency coefficient is high.

Finally, based on the clustering tree and the inconsistency coefficients for each link, we group stations into nonoverlapping markets. We do this 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 of this distribution. This threshold is in line with choices in prior literature on retail fuel markets, which uses the 80th (Assad et al., 2024) or 90th percentile (Lemus and Luco, 2021). We verify that our results do not hinge on the specific choice of the threshold.

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' drive as monopoly markets, without including them in the clustering procedure.⁴⁰

⁴⁰See, e.g., the European Commission's assessment of the takeover of OMV stations by EG Group in the German market (Case M.10134 – EG Group/OMV Germany Business).

Figure A3: Distribution of market sizes in Germany



Notes: The figure shows the distribution of the number of stations per nonoverlapping local market in Germany, with markets defined by the hierarchical clustering algorithm.


Figure A2 illustrates the outcome of the clustering algorithm, showing the nonoverlapping markets for the city of Munich. Figure A3 presents the distribution of market sizes in Germany. The median market consists of four stations, and the 90th percentile is at eight stations.

A.3 Search intensity by fuel type

We use data on search queries in 2015 from a major German price comparison smartphone app to document the ordering of search intensity across the three fuel types. This ordering underlies the test of Proposition 1 in Section 6, which requires that the share of well-informed consumers is higher for diesel than for *E10* and higher for *E10* than for *E5*.

Figure A4 plots the three series on a common scale. The ratio of distinct searchers to vehicles in circulation is approximately 50% higher for diesel than for gasoline, and, within gasoline, it is substantially higher for *E10* than for *E5*. The full ordering diesel > *E10* > *E5* matches the ordering of estimated pass-through rates in Section 6.

Figure A4: Search intensity by fuel type in Germany




Notes: The figure shows the daily number of distinct users searching for fuel prices by fuel type on a major German smart-phone app, per 1,000 vehicles in circulation and on a common scale. The solid, dashed, and short-dashed lines correspond to diesel, *E10*, and *E5*, respectively. The diesel series is normalized by the number of diesel vehicles. The *E5* and *E10* series are normalized by the number of gasoline vehicles and adjusted for the relative market shares of *E5* and *E10* within the gasoline market. We report distinct users rather than total searches to adjust for the higher mileage of diesel drivers. The data are available for January to mid-May and mid-October to early December 2015.

A.4 Additional evidence on search and price dispersion

Figure A5 shows the average price of *E10* in Germany in 2019 at different times of the day. Figure A6 presents price cycles at a more disaggregated level than these 2019 averages, showing that these pricing patterns do not merely result from averaging over time and across stations. We zoom in on one local market in the city of Munich (market no. 21 in the map in Figure A2) and present the stations' raw prices on four consecutive Mondays in the fall of 2019. First, on each of the four days, the stations' pricing follows a similar pattern, in line with that shown in Figure A5. 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 A6, 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 days, so a substantial amount of price variation is unpredictable to consumers, consistent with the mixed strategy equilibrium of our model.

In Table 2, we analyze price dispersion more systematically by computing within-market price residuals for 5 pm prices. Panel A of Figure A7 graphically illustrates the distribution of these residuals. Panel B shows the corresponding residuals when we addi-

Figure A5: 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.


tionally include station fixed effects to absorb any time-invariant price differences across stations. These residuals correspond to the price variation that remains difficult for consumers to predict. Consistent with Table 2, Figure A7 shows that this dispersion is substantial.

As outlined in Section 7, the endogenous-search model of Tappata (2009) predicts that consumers search more, and price dispersion is higher, when prices are low.

Figure A8 shows the average number of searches per app user in 2015 for *E10*, along with the development of the gross price of *E10*. Search intensity and the price level are almost entirely uncorrelated.


Figure A9 depicts the relationship between price dispersion and the price level for *E10* in 2019. The figure corresponds to the specification with only market $\times$ date fixed effects in the last column of Panel B in Table 2. If anything, Figure A9 points to a positive correlation between price dispersion and the price level, implying that gains from search are higher when prices are higher, at odds with Tappata (2009).

Figure A6: Daily price cycles for *E10* on selected Mondays in one local market




Notes: The figure shows prices of *E10* for five stations in one local market in the city of Munich over the course of the day, on four Mondays in the fall of 2019. Fuel prices are updated at five-minute intervals.

Figure A7: Within-market price residuals, 5 pm, 2019




Notes: The figure shows the distribution of the deviation of a fuel station's price from the average price in the same market on the same day (within-market residuals), at 5 pm for *E10* and for all stations that are not local monopolists, using data for all weekdays in 2019. The fixed effects included are indicated in the panel titles.

Figure A8: 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 are available for January to mid-May and mid-October to early December 2015.

Figure A9: 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. We compute price dispersion as the difference between the maximum and the minimum price within a local market, using prices at 5 pm.

B Appendix to Section 3: Theoretical Model

This appendix contains the lemmas behind the propositions of Section 3 (Lemmas B1–B3), the proofs of the propositions, and the extensions: multi-station sellers (Lemmas B4 and B5), downward-sloping demand, and imperfect comparability.

B.1 Equilibrium price distribution

Lemma B1. *Fix $0 < \phi < 1$ and $N \geq 2$, and let $p_r \in (\kappa, v]$ be the reservation price of nonshoppers. There is a unique symmetric mixed strategy Nash equilibrium, in which all sellers draw a price from the interval $[\underline{p}, p_r]$ according to the distribution $F(p_i)$ defined by the equal-profit condition*

$$(p_i - \kappa) (1 + \Lambda w(p_i)) = p_r - \kappa, \quad w(p_i) \equiv (1 - F(p_i))^{N-1}, \quad (\text{B1})$$

Solving for F and for the support floor,

$$F(p_i) = 1 - \left(\frac{p_r - p_i}{\Lambda (p_i - \kappa)} \right)^{\frac{1}{N-1}}, \quad \underline{p} = \kappa + \frac{p_r - \kappa}{1 + \Lambda}. \quad (\text{B2})$$

The expected profits of a seller are

$$E[\pi_i] = \left(\frac{p_r}{1 + \tau} - c \right) \frac{1 - \phi}{N} M. \quad (\text{B3})$$

Proof. *No pure strategy equilibrium and no mass points.* The undercutting arguments are standard (Varian, 1980). The only model-specific element is that the relevant cost floor is the tax-inclusive marginal cost κ rather than the untaxed marginal cost.

Support. No consumer buys at a price above p_r , and a seller that does not win the shoppers optimally charges its captive nonshoppers their reservation price, so the maximum of the support is p_r . Because no seller prices above p_r , nonshoppers buy at the first price they draw, and there is no sequential search in equilibrium.

Equal-profit condition. A seller charging p_i serves its captive share $(1 - \phi)/N$ of nonshoppers and wins the shoppers with probability $w(p_i) = (1 - F(p_i))^{N-1}$; ties have probability zero. Its expected profit is the margin $(p_i - \kappa)/(1 + \tau)$ times its expected demand $M \frac{1-\phi}{N} (1 + \Lambda w(p_i))$, and equilibrium requires this profit to be constant on the support. Evaluating the constant at p_r , where $w(p_r) = 0$, gives Equations (B1) and (B3). Solving Equation (B1) for F gives Equation (B2), and setting $w = 1$ gives $\underline{p}$, the price that wins the

shoppers with certainty. The equal-profit condition thus uniquely identifies the symmetric equilibrium for a given reservation price.⁴¹ $\square$

B.2 Structure and comparative statics of the equilibrium

Lemma B2 (structure of the equilibrium). *Fix $N \geq 2$, $0 < \phi < 1$, and $s > 0$, and define*

$$T_1(\Lambda, N) \equiv \int_0^1 \frac{\Lambda u^{N-1}}{1 + \Lambda u^{N-1}} du, \quad T_N(\Lambda, N) \equiv \int_0^1 \frac{\Lambda u^{N-1}}{1 + \Lambda u^{N-1}} N u^{N-1} du. \quad (\text{B4})$$

The symmetric equilibrium has the following properties.

(i) *For any reservation price $p_r \in (\kappa, v]$, the q -th quantile of F is $p(q) = \kappa + (p_r - \kappa) \frac{1}{1 + \Lambda w_q}$ with $w_q = (1 - q)^{N-1}$. The moments are $E[p] = \kappa + (p_r - \kappa)(1 - T_1)$ and $E[p_{min}] = \kappa + (p_r - \kappa)(1 - T_N)$, so that $p_r - E[p] = (p_r - \kappa) T_1$. Moreover, $0 < T_1 < T_N < 1$.*

(ii) *For any $p_r \in (\kappa, v]$,*

$$E[p_{min}] = \kappa + \frac{1 - \phi}{\phi} (p_r - E[p]), \quad \text{equivalently} \quad T_N = 1 - \frac{N}{\Lambda} T_1. \quad (\text{B5})$$

(iii) *The reservation price is*

$$p_r = \min \left\{ \kappa + \frac{s}{T_1(\Lambda, N)}, v \right\}. \quad (\text{B6})$$

(iv) *In the sequential-search regime, $p(q) - \kappa = s/[T_1(1 + \Lambda w_q)]$ does not depend on κ : every quantile, the expected price, and the expected minimum price rise one-for-one with κ ; $E[p] = \kappa + s(1 - T_1)/T_1$; $E[p_{min}] = \kappa + s(1 - \phi)/\phi$; and the pass-through rate of the ad valorem tax to any such price m is $\rho_\tau^m = \kappa/m < 1$.*

Proof. (i) Setting $F(p(q)) = q$ in Equation (B1) gives $w = (1 - q)^{N-1} \equiv w_q$ and the quantile formula; $q = 0$ gives $\underline{p}$. For a random draw p from F , the variable $u \equiv 1 - F(p)$ is uniform on $[0, 1]$, and the equal-profit condition writes the draw as $p = \kappa + (p_r - \kappa)/(1 + \Lambda u^{N-1})$. Taking expectations gives $E[p] = \kappa + (p_r - \kappa)(1 - T_1)$, hence $p_r - E[p] = (p_r - \kappa) T_1$. The lowest of N independent prices corresponds to the largest of the N draws of u , with density Nu^{N-1} on $[0, 1]$, so $E[p_{min}] = \kappa + (p_r - \kappa)(1 - T_N)$. Write $\psi(x) \equiv x/(1 + x)$, which lies strictly between zero and one for $x > 0$. By Equation (B4), T_1 and T_N are the expectations of $\psi(\Lambda u^{N-1})$ under the uniform density and under the density Nu^{N-1} on $(0, 1)$, respectively,

⁴¹Uniqueness refers to the symmetric equilibrium. Pricing games of this type also admit asymmetric mixed strategy Nash equilibria (Baye, Kovenock, and de Vries, 1992).

so $0 < T_1 < 1$ and $0 < T_N < 1$. The largest of N draws first-order stochastically dominates a single draw and $\psi(\Lambda u^{N-1})$ is strictly increasing in u , so $T_N > T_1$.

(ii) We compute total expected industry profit in two ways. By the equal-profit condition (Equation (B3)), total industry profit is $NE[\pi_i] = (p_r - \kappa)(1 - \phi)M/(1 + \tau)$. Summing over consumers instead, shoppers, of mass ϕM , buy at the realized minimum price and nonshoppers, of mass $(1 - \phi)M$, at a single random draw from F , with expected margins $(E[p_{min}] - \kappa)/(1 + \tau)$ and $(E[p] - \kappa)/(1 + \tau)$ per unit. Equating the two expressions and multiplying by $(1 + \tau)/M$ yields

$$\phi(E[p_{min}] - \kappa) + (1 - \phi)(E[p] - \kappa) = (1 - \phi)(p_r - \kappa). \quad (\text{B7})$$

Rearranging gives the first form. Substituting part (i) and $(1 - \phi)/\phi = N/\Lambda$ gives $\phi(1 - T_N) = (1 - \phi)T_1$, i.e., $T_N = 1 - \frac{N}{\Lambda}T_1$.

(iii) By part (i), the branch $p_r = E[p] + s$ is the linear equation $s = (p_r - \kappa)T_1$, with unique solution $\kappa + s/T_1$, and this is the reservation price whenever it does not exceed v . Otherwise $E[p] + s > v$ at $p_r = v$, and $p_r = v$. This gives Equation (B6).

(iv) By part (iii), $p_r - \kappa = s/T_1$ in the sequential-search regime, and T_1 depends only on (Λ, N) . Substituting into the quantile representation in part (i), $p(q) - \kappa = s/[T_1(1 + \Lambda w_q)]$ does not depend on κ , so every quantile, and with it the expected price and the expected minimum price, shifts one-for-one with κ . In particular, $E[p] = p_r - s = \kappa + s(1 - T_1)/T_1$ and, by part (ii) with $p_r - E[p] = s$, $E[p_{min}] = \kappa + s(1 - \phi)/\phi$. For any such price $m > \kappa$, the pass-through rate is $\rho_\tau^m = \frac{\kappa}{m} \cdot \frac{dm}{d\kappa} = \frac{\kappa}{m} < 1$. $\square$

The number of sellers moves both arguments of the pass-through coefficients. It enters directly and through the competition index $\Lambda(N) = N\phi/(1 - \phi)$.⁴²

Lemma B3 (comparative statics of the pass-through coefficients). *Fix $0 < \phi < 1$ and write $\Lambda(N) = N\phi/(1 - \phi)$.*

- (i) *At fixed $N \geq 2$, T_1 , T_N , and $\Lambda w_q/(1 + \Lambda w_q)$ (for $q < 1$) are strictly increasing in ϕ .*
- (ii) *$T_1(\Lambda(N), N)$ is strictly decreasing in N on the integers $N \geq 2$.*
- (iii) *$T_N(\Lambda(N), N)$ is strictly increasing in N on the integers $N \geq 2$.*

⁴²The price-level counterparts of parts (ii) and (iii) of Lemma B3, entry raising the expected price and lowering the expected minimum price at a fixed reservation price, have antecedents in Rosenthal (1980) and Morgan, Orzen, and Sefton (2006). In the sequential-search model, Stahl (1989) obtains monotone comparative statics only under an auxiliary hypothesis (his Hypothesis HN).

(iv) *Limits: as $N \rightarrow \infty$ along $\Lambda = \Lambda(N)$, $T_1 \rightarrow 0$ and $T_N \rightarrow 1$. At fixed N , as $\phi \rightarrow 1$, T_1 , T_N , and $\Lambda w_q/(1 + \Lambda w_q)$ (for $q < 1$) converge to one; as $\phi \rightarrow 0$, all three converge to zero.*

Proof. (i) At fixed N , Λ is strictly increasing in ϕ , and $\Lambda w/(1 + \Lambda w)$ is strictly increasing in Λ at every $w > 0$. Both integrands in Equation (B4) therefore increase pointwise on $(0, 1)$, and so do T_1 , T_N , and the quantile coefficient.

(ii) Let U be uniform on $[0, 1]$ and define $X_N \equiv \Lambda(N)U^{N-1}$, the ratio of contested to captive demand at a random draw from the equilibrium price distribution. By Equation (B4), $T_1 = E[\psi(X_N)]$, where ψ , as in the proof of Lemma B2(i), maps this ratio into the cost weight and is strictly concave. First, $E[X_N] = \Lambda(N)/N = \phi/(1 - \phi)$ for every N : entry leaves the average contested demand unchanged and only redistributes it across the price distribution. Second, the distribution functions $G_N(x) = (x/\Lambda(N))^{1/(N-1)}$ on $[0, \Lambda(N)]$ cross exactly once. The difference $\ln G_{N+1}(x) - \ln G_N(x)$ is linear and strictly decreasing in $\ln x$, so the two cross at most once where both are interior; since $G_{N+1} > G_N$ near zero and $G_{N+1}(\Lambda(N)) < 1 = G_N(\Lambda(N))$, they cross exactly once, from above. Equal means and a single crossing from above make X_{N+1} a mean-preserving spread of X_N , and the strict concavity of ψ gives $T_1(N+1) = E[\psi(X_{N+1})] < E[\psi(X_N)] = T_1(N)$ (Rothschild and Stiglitz, 1970). Entry spreads the contested demand around a constant mean, and the concavity of the cost weight turns the spread into lower pass-through.

(iii) By Lemma B2(ii), $T_N = 1 - \frac{1-\phi}{\phi} T_1(\Lambda(N), N)$, and $N/\Lambda(N) = (1 - \phi)/\phi$ does not depend on N , so part (ii) implies that T_N is strictly increasing in N . Economically, total industry rents equal what sellers would earn by charging their captive nonshoppers the reservation price, independent of the number of sellers. At a given reservation price, entry raises the rents extracted from nonshoppers (part (ii)), so the rents extracted from shoppers fall by exactly the same amount, and T_N rises.

(iv) For the limit in N , fix $u_0 \in (0, 1)$ and split the integral defining T_1 at u_0 : $T_1 \leq \Lambda(N)u_0^{N-1} + (1 - u_0)$. Since $\Lambda(N)$ grows linearly while u_0^{N-1} vanishes geometrically, $\limsup_N T_1 \leq 1 - u_0$ for every $u_0 < 1$, so $T_1 \rightarrow 0$, and part (iii) gives $T_N \rightarrow 1$. At fixed N , as $\phi \rightarrow 1$, the integrand in Equation (B4) rises to one pointwise on $(0, 1]$, so $T_1 \rightarrow 1$ and $T_N \geq T_1 \rightarrow 1$ by monotone convergence, and $\Lambda w_q/(1 + \Lambda w_q) \rightarrow 1$ for $w_q > 0$. As $\phi \rightarrow 0$, all three coefficients are bounded above by $\Lambda \rightarrow 0$. $\square$

B.3 Proofs of the propositions

In the no-search regime, $p_r = v$ is fixed, and each object of interest has the form

$$m(T) = \kappa + (v - \kappa)(1 - T), \quad T \in \left\{ \frac{\Lambda w_q}{1 + \Lambda w_q}, T_1, T_N \right\}, \quad (\text{B8})$$

by Lemma B2(i). Only the cost anchor moves with κ , so $dm/d\kappa = T$ and the pass-through rate is $\rho_\tau = \kappa T/m(T)$. Since

$$\frac{\partial \rho_\tau}{\partial T} = \kappa \frac{m(T) + T(v - \kappa)}{m(T)^2} = \frac{\kappa v}{m(T)^2} > 0, \quad (\text{B9})$$

every monotonicity statement for a coefficient T transfers to the corresponding rate. In the sequential-search regime, $dm/d\kappa = 1$ and the rate to a quantile, the expected price, or the expected minimum price m is $\rho_\tau^m = \kappa/m$ (Lemma B2(iv)), so the monotonicity of a rate in a parameter is the reverse of the monotonicity of the price level m .

The two regimes meet where $s = (v - \kappa)T_1(\Lambda, N)$. Price levels are continuous there, because $p_r = \min\{\kappa + s/T_1, v\}$ is continuous, whereas the rate to such a price m equals κ/m on the sequential-search side and $\kappa T/m$ with $T < 1$ on the no-search side. Because T_1 is strictly increasing in ϕ and strictly decreasing in N (Lemma B3(i)–(ii)), the switch is unique along either dimension. The sequential-search regime is active for ϕ above a threshold at fixed N , and for N below a threshold at fixed ϕ . Crossing the switch preserves the monotonicity in ϕ of Proposition 1 and the monotonicity in N of Proposition 3, but not that of Proposition 4; Corollary 1 holds across the switch, as its proof below shows. As ϕ rises past the threshold, the rate to every quantile and to both expected prices jumps up, from $\kappa T/m$ to κ/m . As N crosses the threshold, write $\hat{N}$ for the first integer in the no-search regime and $T_1(N)$ for $T_1(\Lambda(N), N)$. The condition $s \leq (v - \kappa)T_1(\hat{N} - 1)$ gives $E[p](\hat{N} - 1) \leq \kappa + (v - \kappa)(1 - T_1(\hat{N} - 1))$, so

$$\rho_\tau^{avg}(\hat{N} - 1) \geq \frac{\kappa}{\kappa + (v - \kappa)(1 - T_1(\hat{N} - 1))} > \frac{\kappa T_1(\hat{N})}{\kappa + (v - \kappa)(1 - T_1(\hat{N}))} = \rho_\tau^{avg}(\hat{N}),$$

using $T_1(\hat{N}) < 1$ and $T_1(\hat{N}) < T_1(\hat{N} - 1)$: the rate to the expected price also falls across the switch. The rate to the expected minimum price, by contrast, can fall across the switch, because its level coefficient drops discretely from one to $T_N(\hat{N}) < 1$, whereas the price level varies continuously across the regime boundary. For example, with $\phi = 1/2$, $\kappa = 1$, $v = 2$, and $s = 0.4$, the market is in the sequential-search regime at $N = 2$ and in the no-search

regime at $N = 3$, and ρ_τ^{min} falls from 0.71 to 0.43. Proposition 4 is therefore a statement within each regime.

Proof of Proposition 1. If $\phi = 0$, the equilibrium is degenerate at the monopoly price, $\underline{p} = p_r = v$. Sellers extract the entire valuation, fully absorb the tax, and $\rho_\tau = 0$. If $\phi = 1$, it is degenerate at the competitive price, $\underline{p} = p_r = \kappa$. Sellers earn zero profits, pass the tax through fully, and $\rho_\tau = 1$.

No-search regime. For $0 < \phi < 1$ with $p_r = v$, the coefficients $dm/d\kappa$ are $\Lambda w_q/(1 + \Lambda w_q)$ for $p(q)$ with $q \in [0, 1)$, T_1 for $E[p]$, and T_N for $E[p_{min}]$. Each is strictly increasing in ϕ by Lemma B3(i), and each rate is strictly increasing in its coefficient by Equation (B9).

Sequential-search regime. Each price level is κ plus a term strictly decreasing in ϕ : $s(1 - \phi)/\phi$ for $E[p_{min}]$; $s(1 - T_1)/T_1$ for $E[p]$, because T_1 is strictly increasing in ϕ (Lemma B3(i)); and $s/[T_1(1 + \Lambda w_q)]$ for $p(q)$, because T_1 and Λ are both strictly increasing in ϕ . Hence every rate κ/m is strictly increasing in ϕ .

The rate jumps up at the unique switch in ϕ , so every pass-through rate is strictly increasing in ϕ on $(0, 1)$. The sequential-search condition fails for ϕ small enough, so the limit $\phi \rightarrow 0$ runs through the no-search regime, where the coefficients converge to zero and the price levels to v (Lemma B3(iv)), and every rate converges to zero. As $\phi \rightarrow 1$, every price level converges to κ in either regime, and every rate to one. $\square$

Proof of Proposition 2. In the no-search regime, $E[p] = m(T_1)$ and $E[p_{min}] = m(T_N)$ with $T_1 < T_N$ (Lemma B2(i)), so $E[p_{min}] < E[p]$ and $dE[p_{min}]/d\kappa = T_N > T_1 = dE[p]/d\kappa$; the rate ordering follows from Equation (B9). In the sequential-search regime, both moments rise one-for-one with κ , and $E[p_{min}] = \kappa + s(1 - \phi)/\phi < \kappa + s(1 - T_1)/T_1 = E[p]$, because $T_1 < \phi$ is equivalent to $T_N > T_1$ by Lemma B2(ii). Hence $\rho_\tau^{min} = \kappa/E[p_{min}] > \kappa/E[p] = \rho_\tau^{avg}$ in both regimes. $\square$

Proof of Proposition 3. If $N = 1$, the monopolist facing unit demand charges $p = v$: a higher price loses all demand, and a lower price adds none. The monopoly price is independent of τ , and pass-through is zero.

No-search regime. For $2 \leq N < \infty$, $\rho_\tau^{avg} = \kappa T_1/E[p]$ is strictly positive, since $T_1 > 0$ (Lemma B2(i)). By Lemma B3(ii), T_1 is strictly decreasing in N along $\Lambda = \Lambda(N)$, so ρ_τ^{avg} is strictly decreasing in N by Equation (B9). By Lemma B3(iv), $T_1 \rightarrow 0$ and $E[p] \rightarrow v$ as $N \rightarrow \infty$, so $\rho_\tau^{avg} \rightarrow 0$.

Sequential-search regime. $E[p] = \kappa + s(1 - T_1)/T_1$ is strictly increasing in N , because T_1 is strictly decreasing in N (Lemma B3(ii)), so $\rho_\tau^{avg} = \kappa/E[p]$ is strictly decreasing in N and remains strictly positive.

The rate falls across the unique switch in N , so pass-through to the expected price is strictly positive and strictly decreasing on $N \geq 2$ and converges to zero. With zero pass-through at $N = 1$, it peaks at $N = 2$. $\square$

Proof of Proposition 4. No-search regime. $\rho_\tau^{min} = \kappa T_N / E[p_{min}]$. By Lemma B3(iii), T_N is strictly increasing in N along $\Lambda = \Lambda(N)$, so ρ_τ^{min} is strictly increasing in N by Equation (B9). By Lemma B3(iv), $T_N \rightarrow 1$ and $E[p_{min}] \rightarrow \kappa$ as $N \rightarrow \infty$, so $\rho_\tau^{min} \rightarrow 1$.

Sequential-search regime. $E[p_{min}] = \kappa + s(1 - \phi)/\phi$ does not depend on N (Lemma B2(iv)), so $\rho_\tau^{min} = \kappa / E[p_{min}]$ is constant in N . $\square$

Proof of Corollary 1. Write $\Delta(N) \equiv \rho_\tau^{min} - \rho_\tau^{avg}$ and $\delta \equiv v - \kappa$.

Within each regime. In the no-search regime, ρ_τ^{min} is strictly increasing and ρ_τ^{avg} strictly decreasing in N (Propositions 4 and 3), so Δ is strictly increasing. In the sequential-search regime, ρ_τ^{min} is constant and ρ_τ^{avg} strictly decreasing in N (Proposition 3), so Δ is again strictly increasing.

Across the switch. Let $\hat{N}$ be the first integer in the no-search regime, so that $\delta T_1(\hat{N}) < s \leq \delta T_1(\hat{N} - 1)$. If $\hat{N} = 2$, every $N \geq 2$ lies in the no-search regime and the previous paragraph applies, so let $\hat{N} \geq 3$. Two auxiliary objects. For any $N \geq 2$ and any $p_r > \kappa$, let $H(N, p_r) \equiv \kappa / E[p_{min}] - \kappa / E[p]$ denote the gap that the levels of Lemma B2(i) would produce under full pass-through (i.e., every price moving one-for-one with the tax-inclusive cost). It equals $\Delta(N)$ where the sequential-search regime is active and $p_r = \kappa + s / T_1(N)$. For any $N \geq 2$, let $\Delta^{ns}(N; \delta) \equiv \kappa T_N / E[p_{min}] - \kappa T_1 / E[p]$ denote the no-search formula at $p_r = v$, written as a function of δ at fixed κ , ϕ , and s . It equals $\Delta(N)$ where the no-search regime is active.

First, at fixed s , $H(N, \kappa + s / T_1(N))$ is strictly increasing in N whether or not the sequential-search regime is active at N : at $p_r = \kappa + s / T_1(N)$, $E[p_{min}] = \kappa + s(1 - \phi)/\phi$ does not depend on N while $E[p] = \kappa + s(1 - T_1)/T_1$ strictly increases in N , by the sequential-search paragraphs of the proofs of Propositions 3 and 4.

Second, $\Delta^{ns}(N; \delta)$ is strictly increasing in N on $N \geq 2$ by Lemma B3(ii)–(iii) and Equation (B9), and $\Delta^{ns}(N; \delta) - H(N, v) = \kappa[(1 - T_1)/E[p] - (1 - T_N)/E[p_{min}]] > 0$, because $y \mapsto y/(\kappa + \delta y)$ is strictly increasing and $1 - T_1 > 1 - T_N$.

Third, by Equation (B8), $\Delta^{ns}(\hat{N}; \delta) = \kappa(T_N - T_1)(\kappa + \delta) / [(\kappa + \delta(1 - T_N))(\kappa + \delta(1 - T_1))]$ with $T_1 = T_1(\hat{N})$ and $T_N = T_N(\hat{N})$. Its derivative in δ has the sign of $\kappa^2(T_1 + T_N - 1) - 2\kappa\delta(1 - T_1)(1 - T_N) - \delta^2(1 - T_1)(1 - T_N)$, which is strictly decreasing in δ . Hence $\Delta^{ns}(\hat{N}; \delta)$ is first increasing and then decreasing in δ , so on any interval it is bounded below by the smaller of its values at the two endpoints of the closure.

Fixing s , the switch condition confines δ to $[s/T_1(\hat{N} - 1), s/T_1(\hat{N})]$, and $\Delta(\hat{N} - 1)$ does not depend on δ . At $\delta = s/T_1(\hat{N} - 1)$, the reservation price at $\hat{N} - 1$ equals v , so $\Delta(\hat{N} - 1) = H(\hat{N} - 1, v) < \Delta^{ns}(\hat{N} - 1; \delta) < \Delta^{ns}(\hat{N}; \delta)$ by the second fact. (Under the opposite boundary convention, $\Delta(\hat{N} - 1) = \Delta^{ns}(\hat{N} - 1; \delta) < \Delta^{ns}(\hat{N}; \delta)$, so the comparison is unaffected.) At $\delta = s/T_1(\hat{N})$, $H(\hat{N}, v) = H(\hat{N}, \kappa + s/T_1(\hat{N})) > \Delta(\hat{N} - 1)$ by the first fact, and $\Delta^{ns}(\hat{N}; \delta) > H(\hat{N}, v)$ by the second. Since $\Delta(\hat{N}) = \Delta^{ns}(\hat{N}; \delta)$ exceeds $\Delta(\hat{N} - 1)$ at both endpoints of the closure and is bounded below by the smaller endpoint value in between, $\Delta(\hat{N}) > \Delta(\hat{N} - 1)$ on the whole interval. $\square$

Proof of Corollary 2. By Equation (B7),

$$\bar{p} \equiv \phi E[p_{min}] + (1 - \phi) E[p] = \kappa + (1 - \phi) (p_r - \kappa) = \phi \kappa + (1 - \phi) p_r.$$

In the no-search regime, $p_r = v$ is independent of κ , so $d\bar{p}/d\kappa = \phi$, which does not depend on N . In the sequential-search regime, $p_r = \kappa + s/T_1$ shifts one-for-one with κ , so $d\bar{p}/d\kappa = \phi + (1 - \phi) = 1$. $\square$

B.4 Multi-station sellers

Setup. An exogenous number N of sellers, indexed by f , operate stations: seller f chooses a price for each of its $S_f \geq 1$ stations, and $S = \sum_f S_f$ is the total number of stations. All other primitives are unchanged. The relevant competition index is now based on stations,

$$\Lambda_S \equiv \frac{S\phi}{1 - \phi} = \frac{S}{N} \Lambda \geq \Lambda,$$

because a low price defends the captive demand of a single station while contesting the same mass of shoppers as before.

Lemma B4 (equilibrium with multi-station sellers). *Fix $0 < \phi < 1$, $N \geq 2$, and $p_r \in (\kappa, v]$.*

- (i) *Contesting stations. In any best response, a seller prices at most one station strictly below p_r and all remaining stations at p_r .*
- (ii) *Equilibrium. There is a symmetric mixed strategy equilibrium, unique up to the assignment of prices to a seller's stations, in which every seller, regardless of its station count, draws one price from the distribution in Equation (B2) with Λ_S in place of Λ , posts it at one station (its contesting station), and posts p_r at its remaining stations. Expected profits are $E[\pi_f] = \left(\frac{p_r}{1+\tau} - c\right) \frac{1-\phi}{S} S_f M$.*

(iii) Market-level prices. Let p_{min} denote the lowest and $\bar{p}^{st}$ the average of the S station prices. Then p_{min} is the minimum of the N contesting-station draws,

$$E[p_{min}] = \kappa + (p_r - \kappa) \left(1 - T_N(\Lambda_S, N)\right), \quad E[\bar{p}^{st}] = \kappa + (p_r - \kappa) \left(1 - \frac{N}{S} T_1(\Lambda_S, N)\right), \quad (\text{B10})$$

and the relationship of Lemma B2(ii) holds with the station average in place of the expected price,

$$E[p_{min}] = \kappa + \frac{1 - \phi}{\phi} (p_r - E[\bar{p}^{st}]). \quad (\text{B11})$$

Proof. (i) Shoppers buy at the market-wide lowest price and nonshoppers buy at their drawn station if and only if its price is at most p_r , so the shopper demand a seller attracts depends on its price vector only through its lowest price m_f . Hold rival prices fixed. A station priced above p_r sells nothing and earns strictly more at p_r . A station priced in (m_f, p_r) serves only its captive nonshoppers, and raising it to p_r strictly raises its margin without changing m_f . If two or more stations are priced at $m_f < p_r$, raising all but one to p_r leaves m_f unchanged, redirects the seller's shopper demand to its remaining lowest-price station, and strictly raises the captive margin at the stations moved up. Any best response therefore prices at most one station strictly below p_r .

(ii) By part (i), a seller's problem reduces to choosing the price $m_f \leq p_r$ of its contesting station. Its expected profit is the payoff of the game in Lemma B1 with the captive share $(1 - \phi)/N$ replaced by $(1 - \phi)/S$, that is, with Λ replaced by Λ_S , plus the captive rents of its remaining $S_f - 1$ stations at p_r , which do not depend on m_f . The reduced game is therefore symmetric across sellers regardless of the ownership distribution $(S_1, \dots, S_N)$, and Lemma B1 applies verbatim. In the unique symmetric equilibrium of the reduced game, every seller draws its contesting price from the atomless distribution in Equation (B2) with Λ_S in place of Λ , so ties in the lowest price have probability zero. Evaluating the constant contesting profit at $m_f = p_r$ and adding the captive rents of the remaining stations gives $E[\pi_f]$.

(iii) Contesting prices lie strictly below p_r almost surely and all remaining stations price at p_r , so the lowest station price equals the lowest of the N contesting draws, and the calculation in the proof of Lemma B2(i) with Λ_S gives $E[p_{min}]$ in Equation (B10). The station average places weight $1/S$ on each contesting draw and $(S - N)/S$ on p_r , which rearranges to $E[\bar{p}^{st}]$ in Equation (B10). For Equation (B11), we compute total expected industry profit in two ways, as in the proof of Lemma B2(ii). Summing $E[\pi_f]$ over sellers gives $(p_r - \kappa)(1 - \phi)M/(1 + \tau)$. Summing over consumers, shoppers generate an expected margin of $(E[p_{min}] - \kappa)/(1 + \tau)$ per unit, and nonshoppers, who buy at a station drawn

uniformly at random, generate $(E[\bar{p}^{st}] - \kappa)/(1 + \tau)$ per unit. Equating the two expressions and rearranging gives Equation (B11). $\square$

Lemma B5 (pass-through with multi-station sellers). *Let $p_r = v$ and $0 < \phi < 1$.*

- (i) Monopoly. *If $N = 1$, every station is priced at v for any $S_1 \geq 1$. The market minimum and the station average coincide, and pass-through of the ad valorem tax to both is zero.*
- (ii) Minimum versus average. *For $N \geq 2$, $dE[p_{min}]/d\kappa = T_N(\Lambda_S, N) > \frac{N}{S} T_1(\Lambda_S, N) = dE[\bar{p}^{st}]/d\kappa$, and the same strict ordering holds for the pass-through rates: Proposition 2 carries over.*
- (iii) Informed share and average transacted price. *At fixed (N, S) , the pass-through rates to the market minimum and the station average are strictly increasing in ϕ , as in Proposition 1. The average transacted price satisfies $\bar{p} = \phi \kappa + (1 - \phi) v$ and $d\bar{p}/d\kappa = \phi$ for every N, S , and ownership distribution: Corollary 2 is unchanged.*
- (iv) Stations per seller. *At a fixed $N \geq 2$, an increase in the number of stations S strictly raises level and rate pass-through to the market minimum, strictly lowers both to the station average, and strictly widens the gap. In particular, since the station-average coefficient at $S = N$ equals $T_1(\Lambda, N)$, the gap is strictly wider than with single-station sellers whenever $S > N$.*

Proof. By Lemma B4(iii), both market-level moments have the form $m(T) = \kappa + (v - \kappa)(1 - T)$ of Equation (B8), with coefficients $T = T_N(\Lambda_S, N)$ for the market minimum and $T = \frac{N}{S} T_1(\Lambda_S, N)$ for the station average, and Equation (B9) transfers every monotonicity statement for a level coefficient to the corresponding pass-through rate.

(i) With a single seller there is no contest. Shoppers buy at the seller's lowest station price whatever its level, and nonshoppers buy at any price up to the valuation, so raising any station price to v leaves demand unchanged and strictly raises the margin. Pricing every station at v is the unique optimum, and the valuation does not depend on κ .

(ii) By Lemma B2(i), $T_N(\Lambda_S, N) > T_1(\Lambda_S, N) \geq \frac{N}{S} T_1(\Lambda_S, N)$.

(iii) At fixed (N, S) , Λ_S is strictly increasing in ϕ , so both coefficients are strictly increasing in ϕ by the pointwise argument in the proof of Lemma B3(i). For the average transacted price, rearranging Equation (B11) with $p_r = v$ gives $\bar{p} = \phi E[p_{min}] + (1 - \phi) E[\bar{p}^{st}] = \phi \kappa + (1 - \phi) v$, as in the proof of Corollary 2.

(iv) At fixed N , Λ_S is strictly increasing in S , so $T_N(\Lambda_S, N)$ is strictly increasing in S by the pointwise argument in the proof of Lemma B3(i). For the station average, treat S as continuous and write $A(S) \equiv \frac{N}{S} T_1(\Lambda_S, N)$. Differentiating under the integral sign,

$$\Lambda \frac{\partial T_1}{\partial \Lambda} = \int_0^1 \frac{\Lambda u^{N-1}}{(1 + \Lambda u^{N-1})^2} du < T_1,$$

because the integrand lies strictly below that of T_1 , so $A'(S) = \frac{N}{S^2} [\Lambda_S \partial T_1 / \partial \Lambda - T_1] < 0$, and pass-through to the station average strictly falls in S , with monotonicity on the integers following. Using $\frac{N}{\Lambda_S} = \frac{N}{S} \frac{1-\phi}{\phi}$, the gap is $T_N - \frac{N}{S} T_1 = 1 - \frac{N}{S\phi} T_1(\Lambda_S, N) = 1 - A(S)/\phi$, so it strictly widens in S , and at $S = N$ the coefficient $A(N) = T_1(\Lambda, N)$ recovers the single-station model. $\square$

The comparative statics in which we vary N carry over at a constant station ratio. Along a sequence of markets in which N increases with S/N constant, $\Lambda_S = aN$ with a constant in N , and $E[aN U^{N-1}] = a$ is again constant in N , so the mean-preserving-spread argument of Lemma B3(ii) and the limit argument of Lemma B3(iv) apply verbatim with $\Lambda(N)$ replaced by aN . Pass-through to the station average therefore peaks at duopoly and declines to zero, the rate to the market minimum rises to one, and the gap between the level coefficients, $1 - \frac{N}{S\phi} T_1(\Lambda_S, N)$, strictly widens. Since the minimum coefficient rises, the average coefficient falls, and each rate is strictly increasing in its coefficient (Equation (B9)), the gap between the rates widens as well: Propositions 3 and 4 and Corollary 1 carry over. The across-the-switch argument in the proof of Corollary 1 transfers verbatim, with Λ_S in place of Λ and $\frac{N}{S} T_1(\Lambda_S, N)$ in place of T_1 .

Three implications map to the empirical exhibits. First, only a seller's cheapest station competes for the informed demand, so the market minimum across stations coincides with the minimum across price setters. The minimum price should therefore be compared across markets by the number of price setters, not by the number of stations. Second, the station average mixes contesting and noncontesting stations. At a given number of sellers, markets in which sellers operate more stations have lower pass-through to the average posted price and a wider pass-through gap (part (iv) of Lemma B5). If stations per seller rise with the number of sellers, the widening of the empirical gap across market-size bins combines the contest channel of Proposition 4 with this composition channel. Table B1 shows that they do, but modestly: the mean number of stations per price setter rises from 1.06 in monopoly markets to 1.22 in markets with six or more sellers, while the share of markets in which some seller operates more than one station rises from 5% to 69%. Third, a monopoly market has no contest for any station network (part (i)). The equilibrium's within-seller price structure

Table B1: Stations per price setter by number of sellers

	Number of competing price setters					
	1	2	3	4	5	6+
<i>A. Tax decrease (1 July 2020)</i>						
Diesel	1.06 (5%)	1.15 (24%)	1.17 (36%)	1.13 (36%)	1.15 (47%)	1.21 (68%)
E10	1.07 (6%)	1.15 (25%)	1.16 (34%)	1.14 (37%)	1.17 (51%)	1.22 (69%)
E5	1.06 (5%)	1.15 (23%)	1.16 (35%)	1.14 (38%)	1.16 (48%)	1.22 (69%)
<i>B. Tax increase (1 January 2021)</i>						
Diesel	1.06 (5%)	1.15 (23%)	1.17 (37%)	1.14 (38%)	1.15 (46%)	1.22 (68%)
E10	1.07 (6%)	1.15 (25%)	1.16 (35%)	1.14 (38%)	1.16 (51%)	1.23 (69%)
E5	1.06 (5%)	1.15 (23%)	1.17 (36%)	1.14 (39%)	1.15 (47%)	1.22 (69%)

Notes: The table reports, by number of competing price setters, the mean number of stations per price setter in German nonoverlapping markets, with the share of markets in which at least one price setter operates more than one station in parentheses. Panels A and B use the market samples of the tax decrease and the tax increase, respectively.

is stark, one contesting station and the rest at the reservation price, but station-level cost or demand heterogeneity would smooth it without restoring a contest at $N = 1$. A nonzero pass-through gap in monopoly markets therefore reflects within-seller price dispersion from sources outside the model, not the mechanism of Proposition 4. The placebo bin in Section 6.4 rests on this prediction.⁴³

The sequential-search regime also carries over. A nonshopper's next search samples a random station, so the consistency condition becomes $p_r = E[\bar{p}^{st}] + s$, which solves to $p_r = \kappa + s S/[N T_1(\Lambda_S, N)]$ by Equation (B10). As in Lemma B2(iv), the price distribution translates one-for-one with the tax-inclusive cost, and Equation (B11) gives $E[p_{min}] = \kappa + s(1 - \phi)/\phi$, independent of the number of sellers and of the station network.

⁴³McAfee (1995) and Dana (1999) derive within-firm price dispersion from multiproduct clearinghouse pricing and from costly capacity with demand uncertainty, respectively. Station-specific pricing by chains is well documented in our setting (Assad et al., 2024).

B.5 Extensions: downward-sloping demand and imperfect comparability

Throughout, we work in the no-search regime. The equilibrium characterizations and the two endpoint results are analytic. The peak locations and profile shapes are computed from the closed-form quantile representations at the parameter values reported in the notes to Figure B1.

Downward-sloping demand. Replace unit demand with a per-consumer demand function $D(p)$, common to both consumer groups, such that $\pi(p) \equiv (p - \kappa) D(p)$ is strictly quasiconcave with an interior maximizer, the monopoly price $\hat{p}(\kappa)$. The derivation of Lemma B1 applies with the margin $p - \kappa$ replaced by $\pi(p)$, because demand multiplies both terms of a seller's expected sales. The equal-profit condition becomes


$$\pi(p) (1 + \Lambda w(p)) = \pi(\hat{p}), \quad w(p) = (1 - F(p))^{N-1},$$

with the top of the support at $\hat{p}$, the price a seller charges when it concedes the contest. Unlike the valuation, this upper anchor moves with the tax-inclusive cost, at the monopoly pass-through rate that the curvature of demand determines (Weyl and Fabinger, 2013). Two endpoint results follow analytically. Under monopoly the seller prices at $\hat{p}$ and passes through at the monopoly rate. As $N \rightarrow \infty$, $\Lambda(N)u^{N-1} \rightarrow 0$ for any fixed $u < 1$ as in Lemma B3(iv), so $\pi(p(u)) \rightarrow \pi(\hat{p})$ pointwise and the price distribution concentrates at $\hat{p}$. Pass-through in N therefore starts at the monopoly rate and returns to it.

Panel A of Figure B1 traces the pass-through rate to the expected price for linear demand, $D(p) = 1 - p$, next to the unit-demand baseline, with the valuation set equal to the linear-demand monopoly price at the initial tax so that the two variants share the same monopoly benchmark price. The hump of Proposition 3 survives with its peak at duopoly, and the post-peak decline is compressed toward the positive monopoly rate. For log-concave demand, the monopoly rate lies below the pass-through of the cost-anchored low prices, and the profile is a hump. For log-convex demand, the ordering reverses, and the computed profile of $dE[p]/d\kappa$ inverts to a trough at duopoly. In the constant-elasticity case, the rate nonetheless stays exactly at one for every N , because without a valuation the equilibrium has no price scale other than κ and is homogeneous of degree one in it.

Imperfect comparability. Let each shopper observe each seller independently with probability $\eta \in (0, 1]$ and buy from the cheapest observed seller. Here $\eta = 1$ is the baseline model, and a shopper who observes no seller does not buy, an event that carries no pricing incentives. A seller charging price p wins a given shopper with probability $\eta (1 - \eta F(p))^{N-1}$, so the equal-profit condition holds with the demand composition $(1 - \phi)/N + \phi \eta (1 - \eta F(p))^{N-1}$,

Figure B1: Pass-through and the number of sellers in the extended models



Notes: The figure shows the pass-through rate of the ad valorem tax τ to the expected price as a function of the number of sellers in the no-search regime, computed from the equilibrium price distribution for a tax change from $\tau = 0.2$ to $\hat{\tau} = 0.22$. Panel A compares the unit-demand baseline with linear captive demand $D(p) = 1 - p$ at full comparability, $\eta = 1$. Panel B holds captive demand linear and varies the probability η with which a shopper compares each rival. Parameter values: $c = 0.4$, $\phi = 0.25$; in the unit-demand variant, $v = 0.74$, the linear-demand monopoly price at the initial tax.

and the quantile representation again solves in closed form. The highest price in the support now carries contested demand, because a seller is the only one a shopper observes with probability $\eta(1 - \eta)^{N-1} > 0$. The contest also builds gradually. A shopper compares $\eta(N - 1)$ rivals in expectation, so the dilution force behind Proposition 3 sets in only once $\eta(N - 1)$ is large.

Panel B of Figure B1 combines the two extensions, tracing the profile for linear demand as comparability falls. The computed peak moves from duopoly at $\eta = 1$ to about six sellers at $\eta = 0.5$ and about thirteen at $\eta = 0.25$, and the profile flattens as it stretches. Across the plotted variants, all with log-concave demand, an interior peak and an eventual decline toward the monopoly rate are invariant. With log-convex demand the computed level profile inverts, and in the constant-elasticity case the rate is exactly one at every N , as noted above. The decline separates the information mechanism from oligopoly models with fully informed consumers and log-concave demand of constant curvature, in which pass-through to the average price increases monotonically in the number of sellers.

C Appendix to Section 4: Descriptive Evidence


This appendix presents additional descriptive evidence on tax pass-through and supports the caution in Section 5 regarding the 2022/23 tax changes.

C.1 Additional descriptive evidence

Section 6.2 studies pass-through to the minimum and the average posted price within local markets, based on market-level rather than station-level prices. Figures C1 and C2 therefore repeat the nonparametric estimation of Figure 1 using the market-level average and minimum price, respectively.


Figures C1 and C2 show results very similar to Figure 1. Market-level average and minimum prices evolve similarly across fuel types before the tax changes, and the post-change ordering of pass-through rates matches Figure 1. The nonparametric evidence is thus robust to using market-level prices and different moments of the price distribution.

Figure C1: Price change as share of total tax change (market-level average prices)



Notes: The figure depicts the price change as a share of the total tax change for the tax decrease in July 2020 (Panel A) and the tax increase in January 2021 (Panel B). Unlike in Figure 1, a price observation here is the average posted price in a nonoverlapping local market on a given day. Construction and pre-periods follow the note to Figure 1. The solid, short-dashed, and long-dashed lines show the estimates for *E5*, *E10*, and diesel, respectively. For the tax decrease, full pass-through corresponds to a price drop of 2.52%, plotted as -100% . Using average absolute prices from 24 June until 30 June 2020, this translates to a decrease of 3.26 eurocents for *E5*, 3.18 eurocents for *E10*, and 2.75 eurocents for diesel. For the tax increase, full pass-through corresponds to a price increase of 2.59% from the VAT increase, plus the newly introduced carbon price. Using absolute prices from 9 December until 15 December 2020, this translates to an increase of 10.40 eurocents for *E5*, 10.27 eurocents for *E10*, and 10.78 eurocents for diesel. The vertical solid line marks the starting date of the tax change.

Figure C2: Price change as share of total tax change (market-level minimum prices)



Notes: The figure depicts the price change as a share of the total tax change for the tax decrease in July 2020 (Panel A) and the tax increase in January 2021 (Panel B). Unlike in Figure 1, a price observation here is the minimum price in a nonoverlapping local market on a given day. Construction and pre-periods follow the note to Figure 1. The solid, short-dashed, and long-dashed lines show the estimates for *E5*, *E10*, and diesel, respectively. For the tax decrease, full pass-through corresponds to a price drop of 2.52%, plotted as -100% . Using average absolute prices from 24 June until 30 June 2020, this translates to a decrease of 3.13 eurocents for *E5*, 3.04 eurocents for *E10*, and 2.62 eurocents for diesel. For the tax increase, full pass-through corresponds to a price increase of 2.59% from the VAT increase, plus the newly introduced carbon price. Using absolute prices from 9 December until 15 December 2020, this translates to an increase of 10.27 eurocents for *E5*, 10.14 eurocents for *E10*, and 10.66 eurocents for diesel. The vertical solid line marks the starting date of the tax change.

C.2 The 2022/23 tax changes

Diesel and gasoline prices diverged before the first 2022 tax change, so the 2022/23 episodes cannot be used to compare pass-through *across* fuel types. In addition, the German tax cut of June 2022 appears to have spilled over to French margins.

Figure C3 shows the evolution of gross fuel prices in Germany and France in February and March of 2022. Two patterns emerge. First, diesel and gasoline prices diverged in March 2022 in both countries, reflecting the differential demand shocks from Russia’s invasion of Ukraine on 24 February 2022. Second, the shock hit the two countries unevenly. Relative to 1 February 2022, diesel (gasoline) prices increased by up to 43% (29%) in Germany but only by up to 31% (19%) in France.


On 1 April 2022 (i.e., right after the time window shown in Figure C3), France introduced a fuel tax rebate of 18 eurocents per liter on both diesel and gasoline. Because of 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 starting on 1 June 2022, which we exclude for the reasons given in Section 4. The size of the 2022 rebates also raises the spillover concern discussed in Section 5. They may have changed the opportunity cost of selling fuel across countries.

Figure C4 presents evidence of such potential spillover effects, showing an increase in French gross margins over crude and taxes immediately after the introduction of the German tax cut on 1 June 2022. To compute margins for *E5*, *E10*, and diesel, we subtract taxes, duties, and the estimated input cost of crude oil from the gross price.⁴⁴ The figure shows that margins at French stations increased by approximately 5 eurocents on the day when the German fuel rebate went into effect and continued increasing in subsequent weeks. Estimating pass-through of the German tax cut on 1 June 2022 with France as the control group is therefore problematic.


⁴⁴To compute these 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 approximately 19 gallons of gasoline, 12 gallons of diesel, or 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. Approximately 54% of the Brent oil price per barrel corresponds to the production of 19 gallons of gasoline, while approximately 34% corresponds to the production of 12 gallons of diesel. Finally, we transform these values into the approximate input cost per liter of gasoline and diesel.

Figure C3: Evolution of gross prices in early 2022



Notes: The figure shows the evolution of daily gross fuel prices in February and March of 2022. All prices are normalized to one at 1 February 2022. The solid, short-dashed, and long-dashed lines show prices for *E5*, *E10*, and diesel, respectively.

Figure C4: Margins in France around 1 June 2022



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

D Appendix to Section 6: Empirical Results

D.1 Robustness: Additional controls

We begin by presenting robustness checks where we include additional control variables into our regressions.

Table D1: 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*** (0.0008)	-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*** (0.0056)	0.0470*** (0.0053)	0.0790*** (0.0037)	0.0202*** (0.0037)	0.0755*** (0.0031)
Pass-through rate	32% [28%, 35%]	49% [46%, 52%]	92% [89%, 94%]	70% [69%, 71%]	73% [72%, 74%]	86% [85%, 86%]
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 (2), additionally controlling for regional mobility for retail and recreational purposes and to workplaces, using data from the Google Mobility Report. Sample periods, standard errors, and the reporting of implied pass-through rates follow Table 3.

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

D.1.1 Baseline pass-through estimation

In Table D1, we re-estimate the baseline model, additionally controlling for regional mobility for retail and recreational purposes and to workplaces, using data from the Google Mobility Report. The estimates are close to those in Table 3. Pass-through remains significantly higher for diesel than for gasoline and significantly higher for *E10* than for *E5*, consistent with Proposition 1.

D.1.2 Number of sellers and pass-through


We now assess the robustness of our results on the relationship between pass-through and the number of sellers in a market. One potential concern regarding our baseline analysis in Section 6.3 is that markets with different numbers of competing price setters may also be different in terms of other characteristics, which could explain different pass-through rates.

We address this potential concern by taking into account county-level demographic characteristics obtained from the German Federal Statistical Office. Specifically, we include income per capita, population density, and an indicator for whether the county is a city (i.e., “Kreisfreie Stadt”). We then run a nonparametric regression of the station-level pass-

through rate on indicators for markets with different numbers of competing price setters, as well as the demographic control variables.

The results in Figure D1 look very similar to those in Figure 3. For *E5* and *E10*, pass-through rises from monopoly to two competing price setters and does not increase further, with the fitted lines again pointing to a mildly hump-shaped profile. For diesel, the initial rise is much weaker, but the profile is still mildly hump-shaped, and the U-test in Table D7 rejects a monotone relationship for it too, at the 10% level in the tax increase. The nonmonotonic relationship between the number of sellers and pass-through is thus robust to controlling for demographic characteristics.

Figure D1: Average pass-through by number of competitors (with demographic controls)



Notes: The figure shows how the rate of pass-through to the average price varies with the number of competing price setters in a market. Unlike in Figure 3, the nonparametric regression additionally controls for county-level demographic characteristics (income per capita, population density, and a city indicator). Construction and reporting otherwise follow the note to Figure 3.

D.2 Robustness: Balanced sample

We now restrict the sample to stations in Germany and France with a price observation on every day of the sample period. For diesel, this retains 83% of German and 63% of French stations for the tax reduction, and 82% and 72% for the tax increase. The estimates in Table D2 are lower than those in Table 3 by no more than about five percentage points, except for *E5* in the tax decrease, where the rate falls from 28% to 14%. The ranking of pass-through rates across fuel types is unchanged.

Table D2: 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*** (0.0004)	0.0833*** (0.0003)
Pass-through rate	14% [10%, 17%]	42% [39%, 45%]	89% [86%, 91%]	68% [67%, 69%]	73% [73%, 74%]	88% [87%, 88%]
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 (2). Unlike in Table 3, the model is estimated on a balanced sample of stations with a price observation for every day. Sample periods, standard errors, and the reporting of implied pass-through rates follow Table 3.

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

D.3 Robustness: Anticipatory effects

In Table D3, we examine how our estimates of pass-through rates change when we alter the assumptions on anticipatory effects. In columns (1) to (3), we estimate the pass-through rate of the tax decrease when we drop the second half of June 2020 from the control period. In this case, the gap between pass-through rates among *E5*, *E10* and diesel widens, but the order of estimates in terms of magnitude remains the same. This is not our preferred estimation strategy since we do not think that there is sufficient evidence of anticipatory pass-through of the tax decrease in June 2020. We therefore treat these point estimates of the pass-through rate with caution. 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 when we include the second half of December 2020 in the control period. In this case, the point estimate of the pass-through rate decreases from 70% to 65% for *E5*, from 75% to 67% for *E10*, and from 88% to 73% for diesel. This is expected since Figure 1 reveals important anticipatory effects of the tax pass-through in the second half of December 2020. The ranking across fuel types is preserved, but the differences shrink by half or more. The gap between *E10* and *E5* falls from 5 to 2 percentage points, between diesel and *E10* from 13 to 6, and between diesel and *E5* from 18 to 8.

Table D3: 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%]	92% [89%, 96%]	65% [65%, 66%]	67% [66%, 68%]	73% [73%, 74%]
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 (2) with altered assumptions on anticipatory effects. Columns (1) to (3) estimate the German VAT reduction on 1 July 2020 using data from 1 May to 15 June 2020 and from 1 July to 31 August 2020. Columns (4) to (6) estimate the German VAT increase and introduction of a carbon price on 1 January 2021 using data from 1 November 2020 to 28 February 2021. Standard errors and the reporting of implied pass-through rates follow Table 3.

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

D.4 Robustness: Pooled regression

In Section 6, we estimate separate DID regressions for each fuel type. We conclude that pass-through is highest for diesel and lowest for *E5* for both the tax decrease and the tax increase, since the confidence intervals of the pass-through estimates do not overlap. An alternative approach to estimating the difference in pass-through across fuel types is by pooling all fuel types in one regression. We implement such a pooled regression in Table D4, controlling for date×fuel fixed effects and station×fuel fixed effects. Moreover, we control for an interaction term of the crude oil price with the fuel type and with an indicator for stations in Germany. This allows for differential pass-through of the crude oil price in Germany and France, and for each fuel type.

By construction, columns (1) and (3) replicate the baseline *E5* pass-through rates in Table 3. The interaction coefficients give the difference in pass-through for *E10* relative to *E5* and for diesel relative to *E10*, so the levels for *E10* and diesel are recovered by summing the coefficients.

One advantage of the pooled regression is that we can additionally control for date×station fixed effects. That is, we can estimate the difference in pass-through across fuel types by only exploiting price variation *within* stations on any given day. The resulting estimates are shown in columns (2) and (4) of Table D4. For the tax decrease, we estimate that the pass-through rate is 21 percentage points higher for *E10* than for *E5* and an additional 45 percentage points higher for diesel than for *E10*. Similarly, for the tax increase, we estimate that the pass-through rate is three percentage points higher for *E10* than for *E5* and an additional 14 percentage points higher for diesel than for *E10*. These differences across fuel types are very similar in magnitude to those in our baseline specification estimated without date×station fixed effects. The drawback of the pooled regression with date×station fixed effects is that we can only estimate differences in pass-through, as the absolute magnitude of the price change (i.e., the coefficient for *E5*) is absorbed by the fixed effects.

The estimates indicate that the difference in pass-through between all fuel types is statistically significant at the 1% level. This is true even when we control for date×station fixed effects. Therefore, the pooled regression further supports Proposition 1, which establishes that the pass-through rate is higher when there are more price-sensitive consumers.

Table D4: Effect of the tax change on log prices (pooled regression)

	Tax decrease		Tax increase	
	(1)	(2)	(3)	(4)
Tax change				
<i>E5</i>	-0.0071*** (0.0004)		0.0561*** (0.0004)	
<i>E10</i> vs. <i>E5</i>	-0.0045*** (0.0003)	-0.0053*** (0.0002)	0.0049*** (0.0002)	0.0038*** (0.0002)
Diesel vs. <i>E10</i>	-0.0122*** (0.0003)	-0.0115*** (0.0003)	0.0223*** (0.0003)	0.0226*** (0.0003)
Pass-through				
<i>E5</i>	28% [25%, 31%]		70% [69%, 71%]	
<i>E10</i> vs. <i>E5</i>	18pp [16pp, 20pp]	21pp [19pp, 22pp]	4pp [4pp, 5pp]	3pp [3pp, 3pp]
Diesel vs. <i>E10</i>	48pp [45pp, 51pp]	45pp [42pp, 48pp]	13pp [13pp, 14pp]	14pp [13pp, 14pp]
Date $\times$ fuel fixed effects	Yes	Yes	Yes	Yes
Station $\times$ fuel fixed effects	Yes	Yes	Yes	Yes
Date $\times$ station fixed effects	No	Yes	No	Yes
DE $\times$ fuel $\times$ oil price	Yes	Yes	Yes	Yes
Observations	7,161,112	7,101,790	6,112,024	6,064,209

Notes: The table presents DID estimates using a variation of the model in Equation (2) that pools all three fuel types. The base category is *E5*. The table reports the change in log prices for *E10* relative to *E5*, and for diesel relative to *E10*. Columns (1) and (2) cover the tax decrease and columns (3) and (4) the tax increase. Sample periods, standard errors, and the reporting of implied pass-through rates follow Table 3. For the tax increase, the full-pass-through benchmark differs across fuel types because of the carbon price, so the pass-through-rate differences in column (4) combine that column's interaction coefficients with the *E5* level from column (3) and the fuel-specific benchmarks. For the tax decrease, the benchmark is common to all fuel types, and column (2) needs no such anchor.

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

D.5 Robustness: Synthetic difference-in-differences analysis

Next, we repeat our baseline pass-through estimation but employ a synthetic difference-in-differences (SDID) approach instead of a standard DID. SDID reweights control units to match pre-treatment trends, in the spirit of synthetic control methods (Arkhangelsky et al., 2021).

We first 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. We use clustered bootstrapping with 300 replications and clustering at the station level to estimate standard errors. 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.

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 price and the average price, as SDID does not allow a triple interaction term.

D.5.1 Baseline pass-through estimation


Table D5 reports the SDID estimates. For the tax decrease, the pass-through rates are 23% for *E5*, 45% for *E10*, and 82% for diesel. For the tax increase, the joint pass-through rate of the VAT increase and the carbon price is 78% for *E5* and *E10* and 90% for diesel. Pass-through for *E5* and *E10* is here statistically indistinguishable, while the ranking of pass-through rates across fuel types and their magnitudes remain close to Table 3.

D.5.2 Number of sellers and pass-through

Figure D2 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 SDID (i.e., with additional weights) and a restricted balanced sample of stations. We estimate the time and unit weights only once and then use them to estimate each station-specific treatment effect.

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

Figure D2: Average pass-through by number of competitors (SDID)



Notes: The figure shows how the rate of pass-through to the average price varies with the number of competing price setters in a market. Unlike in Figure 3, pass-through rates are estimated with an SDID strategy on a restricted balanced sample. Construction and reporting otherwise follow the note to Figure 3.

Table D5: 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%, 48%]	82% [78%, 85%]	78% [77%, 80%]	78% [77%, 79%]	90% [90%, 91%]
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. Standard errors obtained via clustered bootstrap with 300 replications are shown in parentheses. Sample periods and the reporting of implied pass-through rates follow Table 3.

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

D.6 Robustness: Price at 10th percentile

A potential concern about the market minimum is that it may reflect outlier prices that are valid for only a short time. Figure D3 therefore replaces the minimum with the price at the 10th percentile of the distribution of all hourly prices within a market on a given day. These estimates are not placebo-corrected, so the decrease-episode bins carry the caveat from Section 6.4. For the tax increases, the informative episodes, the results are very similar to Figure 2. The differential pass-through to the low end of the price distribution persists, so the pass-through gaps are not driven by outlier prices.

Figure D3: Pass-through to market-level minimum (10th percentile) vs. average prices




Notes: The figure shows the difference in the market-level pass-through rate between the minimum price and the average posted price. The average posted price is the average daily price in a nonoverlapping market, weighting the price at every full hour between 6 am and 10 pm equally. The 10th percentile of all hourly prices in the market on a given day replaces the minimum price. Markers show the difference implied by the DID estimate β_2 in Equation (3), estimated separately for each tax change and fuel type, with 99% confidence intervals based on standard errors clustered at the market level. Estimation windows follow the note to Figure 2. Solid (hollow) symbols indicate point estimates that are statistically significant (insignificant) at the 1% level.

D.7 The pass-through gap by number of sellers: French tax changes

Figure D4 repeats the placebo-corrected 5 pm analysis of Section 6 by the number of sellers for the three French tax changes. As discussed in Section 6.4, markets in France are grouped by the number of stations rather than the number of competing price setters. The two French tax increases (Panels B and C) show a clean, monotone widening of the corrected gap in the number of stations for all three fuel types, from the placebo bin at zero to between +9 and +33 percentage points in the largest markets, in line with Corollary 1. The French tax decrease of April 2022 (Panel A) is small and mixed: the gap widens monotonically for *E5* but is nonmonotonic for diesel and *E10*, with all estimates within a few percentage points of zero.

Figure D4: Pass-through gap by number of stations (5 pm, France)



Notes: The figure shows how the placebo-corrected difference in the market-level pass-through rate between the cross-station minimum price and the average posted price at 5 pm varies with the number of stations in a nonoverlapping market. All market-size bins are estimated jointly and reported relative to the placebo bin of single-station markets, in which the minimum and average price coincide by construction (shown at zero at $N = 1$). Panels A, B, and C depict the French tax decrease on 1 April 2022, the French tax increase on 16 November 2022, and the French tax increase on 1 January 2023, respectively. Because station-level price-setter identity is not available for France, markets are grouped by the number of stations (1, 2, 3, and 4 or more) rather than by the number of competing price setters as in the German analysis. In each panel, the circles, diamonds, and triangles plot the corrected gap for diesel, E10, and E5, respectively. Vertical spikes show 95% confidence intervals based on standard errors clustered at the market level. Solid (hollow) markers indicate estimates that are statistically significant (insignificant) at the 5% level. For the tax increase on 16 November 2022, we use the period until 31 December 2022 as the post-treatment period; for the tax increase on 1 January 2023, we use 16 November until 31 December 2022 as the pre-treatment period.

D.8 Robustness: Overlapping markets


The main analysis groups stations into nonoverlapping markets by hierarchical clustering (Appendix A.2). An alternative is to center a market on each station and include all competitors within a 5 km driving distance, which yields overlapping markets, one per station (Table 1). Luco (2019) argues that station-centered markets are preferable when the researcher does not observe the universe of stations; since we observe the universe in both countries, we keep the clustering definition as the baseline and use the overlapping definition as a robustness check.

D.8.1 Pass-through to the average and minimum price

In Figure D5, we re-estimate the difference in pass-through between the minimum price and the average posted price for all the tax changes in our sample period. Unlike in Figure 2, we now use overlapping markets, and prices are measured over the whole day rather than at 5 pm.

When using overlapping markets, the conclusions are unchanged. These estimates are not placebo-corrected, so the decrease-episode bins carry the caveat from Section 6.4. For the tax increases, the informative episodes, pass-through rates to the minimum price remain

Figure D5: Pass-through to market-level minimum vs. average prices (overlapping markets)



Notes: The figure shows the difference in the market-level pass-through rate between the minimum price and the average posted price. The average posted price is the average daily price in a local market, weighting the price at every full hour between 6 am and 10 pm equally, and the minimum price is the minimum in the market at any point during the day. Markets are overlapping, centered on each station and including all competing price setters within a driving distance of 5 km. Markers show the difference implied by the DID estimate β_2 in Equation (3), estimated separately for each tax change and fuel type, with 99% confidence intervals based on standard errors clustered at the market level. Estimation windows follow the note to Figure 2. Solid (hollow) symbols indicate point estimates that are statistically significant (insignificant) at the 1% level.


significantly higher than those to the average posted price in all nine cases in Figure D5. For the tax decrease, the picture is mixed, as in the main text.

In overlapping markets, the prices of stations in areas with a higher density of fuel stations receive a higher weight because they belong to multiple markets, and a single price notification can be the minimum price in several nearby markets at once. The estimates in Figure D5 may therefore be overly sensitive to individual, typically independent, low-price stations.

D.8.2 Number of sellers and pass-through

Figure D6 repeats the analysis of Figure 3 for the 2020/21 German tax changes using overlapping markets. The inverse U-shaped relationship between the average pass-through rate and the number of sellers is pronounced for all three fuel types in the tax decrease, with the highest pass-through rates in markets with seven to eight competing price setters, and for *E5* and *E10* in the tax increase, with a peak around ten to eleven sellers. The exception is diesel in the tax increase, where pass-through is continuously increasing in the number of sellers, consistent with its exception in the formal test of Table D8.

Figure D6: Average pass-through by number of competitors (overlapping markets)



Notes: The figure shows how the rate of pass-through to the average price varies with the number of competing price setters in a market. Unlike in Figure 3, markets are overlapping, centered on each station and including all competing price setters within a driving distance of 5 km. Standard errors are clustered by the station's nonoverlapping market. Construction and reporting otherwise follow the note to Figure 3.

D.9 Formal test of nonmonotonicity

The results in Figure 3 indicate that pass-through to the average price is not monotonically increasing in the number of sellers. In this section, we formally test for nonmonotonicity by applying the U-test of 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.

To this end, we regress the station-level pass-through rate on the number of competing price setters, its square, and a constant. A positive coefficient on the linear term, a negative coefficient on the quadratic term, and an extreme point within the data range are necessary but not sufficient conditions for an inverse U-shape, so Lind and Mehlum (2010) propose

Table D6: U-test of nonmonotonicity

	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	0.0116 (0.0235)	0.0189 (0.0297)	0.0290 (0.0231)	0.0089 (0.0057)	0.0086 (0.0058)	0.0059 (0.0048)
(No. of competing price setters) ²	-0.0019 (0.0024)	-0.0023 (0.0032)	-0.0034 (0.0024)	-0.0011* (0.0006)	-0.0010* (0.0006)	-0.0004 (0.0005)
Constant	0.2505*** (0.0496)	0.4017*** (0.0587)	0.8764*** (0.0477)	0.6820*** (0.0126)	0.7236*** (0.0127)	0.8596*** (0.0101)
Inverse-U test						
<i>t</i> -value	0.4103	0.6116	1.2057	1.4605	1.4142	0.5212
<i>p</i> -value	0.3408	0.2704	0.1140	0.0721	0.0787	0.3011
Extreme point	3.0455	4.1154	4.3252	4.2032	4.4541	6.8152
Lower bound	1	1	1	1	1	1
Slope at lower bound	0.0078	0.0143	0.0223	0.0068	0.0066	0.0050
<i>t</i> -value	0.4103	0.6116	1.2057	1.4605	1.4142	1.3185
Upper bound	10	10	10	10	10	10
Slope at upper bound	-0.0264	-0.0271	-0.0380	-0.0123	-0.0107	-0.0028
<i>t</i> -value	-1.0227	-0.7425	-1.4395	-2.0106	-1.7524	-0.5212
Observations	14,000	13,454	14,267	14,034	13,485	14,312

Notes: The table shows the estimates from our regressing the station-level pass-through rate on the number of competing price setters in a local nonoverlapping market, its square, and a constant. Standard errors clustered at the market level are in parentheses. Columns (1) to (3) present estimates for the German VAT reduction on 1 July 2020, while columns (4) to (6) present estimates for the German VAT increase and introduction of a carbon price on 1 January 2021. Stations in markets above the 97.5th percentile of the number of competing price setters are excluded. The bottom part of the table shows the application of the U-test of Lind and Mehlum (2010).

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

a formal test that also takes into account the slopes at the lower and upper bounds of the data range.

We present the baseline estimates in Table D6, following the exposition in Pennerstorfer et al. (2020); Table D7 adds demographic characteristics to the regression, which addresses the concern that markets with different numbers of competing price setters may differ in other characteristics. In both cases, the estimated coefficients satisfy all three necessary conditions for each fuel type and tax change. Standard errors are clustered at the market level, since the station-level rates within a market share the market’s price shock. In the baseline specification, the inverse-U test rejects the null hypothesis of a monotonic or U-shaped relationship at the 10% level for *E5* and *E10* in the tax increase, with p -values of 0.07 and 0.08; in the other cases the p -values range from 0.11 to 0.34. With demographic controls, the test rejects at the 5% level for *E5* in the tax increase, and the p -value for *E10* is 0.05. Table D8 repeats the test with the overlapping markets of Appendix D.8; the necessary conditions hold, the test rejects at the 1% level for all three fuel types in the tax decrease, and the p -values for *E5* and *E10* in the tax increase are 0.12 and 0.14. For diesel in the tax increase, the estimated slope does not turn negative within the data range and the estimates suggest an increasing relationship between the number of competitors and average pass-through.

Table D7: U-test of nonmonotonicity (with demographic 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)
No. of competing price setters	0.0208 (0.0221)	0.0312 (0.0285)	0.0290 (0.0229)	0.0098* (0.0057)	0.0098* (0.0058)	0.0052 (0.0047)
(No. of competing price setters) ²	-0.0022 (0.0022)	-0.0029 (0.0031)	-0.0031 (0.0024)	-0.0011* (0.0006)	-0.0010* (0.0006)	-0.0004 (0.0005)
Income per capita	-0.0000*** (0.0000)	-0.0000*** (0.0000)	-0.0000*** (0.0000)	-0.0000*** (0.0000)	-0.0000*** (0.0000)	0.0000*** (0.0000)
Population density	0.1027*** (0.0230)	0.1019*** (0.0239)	0.1390*** (0.0199)	0.0089 (0.0062)	0.0084 (0.0063)	-0.0104** (0.0049)
City (dummy)	-0.1926*** (0.0525)	-0.1777*** (0.0546)	-0.1925*** (0.0463)	0.0243* (0.0143)	0.0265* (0.0145)	0.0261** (0.0111)
Constant	1.3067*** (0.0809)	1.4784*** (0.0854)	1.5100*** (0.0762)	0.8607*** (0.0211)	0.9051*** (0.0215)	0.7876*** (0.0165)
Inverse-U test						
<i>t</i> -value	0.9222	0.7490	1.2449	1.6531	1.6407	0.5708
<i>p</i> -value	0.1782	0.2269	0.1066	0.0492	0.0505	0.2841
Extreme point	4.7850	5.4610	4.6803	4.4356	4.6676	6.3859
Lower bound	1	1	1	1	1	1
Slope at lower bound	0.0165	0.0255	0.0228	0.0076	0.0077	0.0043
<i>t</i> -value	0.9222	1.1319	1.2449	1.6531	1.6407	1.1551
Upper bound	10	10	10	10	10	10
Slope at upper bound	-0.0227	-0.0259	-0.0330	-0.0123	-0.0111	-0.0029
<i>t</i> -value	-0.9544	-0.7490	-1.2677	-2.0243	-1.8125	-0.5708
Observations	13,987	13,442	14,253	14,021	13,473	14,298

Notes: As Table D6, with demographic characteristics added to the regression.

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

Table D8: U-test of nonmonotonicity (overlapping markets)

	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	0.0407*** (0.0133)	0.0463*** (0.0147)	0.0295*** (0.0112)	0.0159*** (0.0034)	0.0168*** (0.0036)	0.0033 (0.0026)
(No. of competing price setters) ²	-0.0034*** (0.0012)	-0.0038*** (0.0013)	-0.0025*** (0.0010)	-0.0008*** (0.0003)	-0.0009*** (0.0003)	-0.0000 (0.0002)
Constant	0.1669*** (0.0301)	0.3249*** (0.0324)	0.8659*** (0.0253)	0.6458*** (0.0081)	0.6851*** (0.0085)	0.8599*** (0.0061)
Inverse-U test						
<i>t</i> -value	2.6874	2.4988	2.4167	1.1801	1.0920	n/a
<i>p</i> -value	0.0036	0.0063	0.0079	0.1190	0.1374	n/a
Extreme point	5.9368	6.1198	5.8528	9.4039	9.5173	1.1e+03
Lower bound	1	1	1	1	1	1
Slope at lower bound	0.0338	0.0387	0.0245	0.0142	0.0151	0.0033
<i>t</i> -value	3.0370	3.1771	2.5978	4.9936	4.9808	1.5250
Upper bound	12	12	12	12	12	12
Slope at upper bound	-0.0415	-0.0445	-0.0310	-0.0044	-0.0044	0.0033
<i>t</i> -value	-2.6874	-2.4988	-2.4167	-1.1801	-1.0920	1.0845
Observations	14,088	13,522	14,355	14,122	13,552	14,396

Notes: As Table D6, with overlapping markets centered on each station and including all competing price setters within a driving distance of 5 km.

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

D.10 Robustness: The crude-oil control and the decrease window

The baseline station-level specification in Equation (2) controls for the crude oil price interacted with a Germany indicator, whereas the market-level triple DID in Equation (3) includes no controls. Table D9 shows that this difference only matters for the gasoline decrease. Adding the crude-oil control to the market-level regressions raises the average-price pass-through of the tax cut from -0.08 to 0.24 for *E5* and from 0.16 to 0.41 for *E10*, close to the station-level baseline in Table 3, while diesel and all estimates for the tax increase move little. The pass-through gap coefficients are unchanged throughout, because the crude-oil control is a level control and does not affect dispersion.

The sensitivity of the estimates to including a control for the crude oil price is specific to the time window around the tax decrease, because crude oil prices recovered steeply from their COVID trough around this time. Table D10 narrows the estimation window around 1 July 2020. As the window shrinks from the full sample to two weeks on each side of the

Table D9: Market-level pass-through with and without the crude-oil control

	Average price (β_1)		Min.-price gap (β_2)	
	no control	crude control	no control	crude control
<i>A. Tax decrease (1 July 2020)</i>				
Diesel	0.88 (0.01)	0.88 (0.01)	0.08 (0.01)	0.08 (0.01)
E10	0.16 (0.01)	0.41 (0.02)	0.01 (0.01)	0.01 (0.01)
E5	-0.08 (0.02)	0.24 (0.02)	0.02 (0.01)	0.02 (0.01)
<i>B. Tax increase (1 January 2021)</i>				
Diesel	0.92 (0.00)	0.87 (0.00)	0.03 (0.00)	0.03 (0.00)
E10	0.76 (0.00)	0.74 (0.00)	0.03 (0.00)	0.03 (0.00)
E5	0.77 (0.00)	0.70 (0.01)	0.03 (0.00)	0.03 (0.00)

Notes: The table reports market-level 5 pm pass-through rates from the aggregate triple DID in Equation (3), estimated without controls and with the crude oil price interacted with a Germany indicator, the control used in the station-level baseline. β_1 is the pass-through rate to the average posted price and β_2 the difference in the pass-through rate between the minimum and the average price, both divided by the full-pass-through benchmark. Standard errors, clustered at the market level and divided by the same benchmark, are in parentheses.

reform, the gap between the station-level gasoline estimates with and without the crude-oil control narrows from 31 percentage points for *E10* and 37 for *E5* to 13 for both, whereas the diesel gap widens from 8 to 16. Gasoline pass-through in the narrow window is substantially positive under either specification. The ranking of pass-through rates across fuels is the same in every window and under either specification.

Table D10: Station-level pass-through of the tax decrease by estimation window

	Full window		± 4 weeks		± 2 weeks	
	no control	crude	no control	crude	no control	crude
Diesel	0.86 (0.01)	0.93 (0.01)	0.83 (0.01)	0.90 (0.01)	0.67 (0.01)	0.82 (0.01)
E10	0.15 (0.02)	0.46 (0.02)	0.45 (0.01)	0.73 (0.01)	0.49 (0.01)	0.63 (0.01)
E5	-0.09 (0.02)	0.28 (0.02)	0.24 (0.01)	0.58 (0.01)	0.35 (0.01)	0.48 (0.01)

Notes: The table reports station-level pass-through rates of the 1 July 2020 tax decrease from the DID model in Equation (2), with and without the crude-oil control, for three estimation windows: the full baseline window (1 May to 31 August 2020), four weeks on each side of the reform, and two weeks on each side. Standard errors, clustered at the market level and divided by the full-pass-through benchmark, are in parentheses.

D.11 Robustness: Level pass-through gaps

The log specification in Equation (3) measures proportional responses, so if every price in the market changes by the same level, this produces a positive pass-through gap because the minimum is the smaller base. Estimating price responses in levels removes this concern. Table D11 reports the placebo-corrected level gaps by market size. What remains is the differential level response of the minimum price, the compression toward cost of the no-search regime. For the tax increase, the compression is statistically significant in most bins but small, about one percentage point of the full-pass-through amount for diesel and *E10* and about half a percentage point for *E5* in markets with two to four sellers, where it turns significantly negative in the largest markets. For the tax decrease, the gaps turn negative in the largest markets, the level counterpart of the dispersion compression documented in Section 6.4.

Table D11: Placebo-corrected level pass-through gaps by market size

	Number of competing price setters				
	2	3	4	5	6+
<i>A. Tax decrease (1 July 2020)</i>					
Diesel	-0.001 (0.007)	-0.004 (0.006)	0.012 (0.008)	0.001 (0.009)	-0.022*** (0.008)
E10	-0.016** (0.007)	-0.019** (0.008)	-0.013* (0.007)	-0.034*** (0.010)	-0.041*** (0.012)
E5	-0.014** (0.006)	-0.026*** (0.007)	-0.017** (0.008)	-0.034*** (0.009)	-0.058*** (0.010)
<i>B. Tax increase (1 January 2021)</i>					
Diesel	0.006*** (0.002)	0.010*** (0.002)	0.015*** (0.002)	0.012*** (0.003)	0.011*** (0.003)
E10	0.009*** (0.003)	0.012*** (0.002)	0.008*** (0.003)	0.009*** (0.003)	0.004 (0.003)
E5	0.004** (0.002)	0.005** (0.002)	0.005* (0.002)	0.001 (0.003)	-0.008** (0.003)

Notes: The table reports placebo-corrected differences in pass-through between the minimum and the average posted price from the triple DID in Equation (3) estimated in levels rather than logarithms, by the number of competing price setters in a market. Coefficients are divided by the euro amount that full pass-through of each tax change would imply, so entries are comparable to pass-through rates; standard errors, clustered at the market level and divided by the same amount, are in parentheses. All market-size bins are estimated jointly and reported relative to the placebo bin of single-station monopoly markets. Under a uniform level response of all prices, every entry would be zero.

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

D.12 Robustness: Placebo reform dates in 2019

We probe the placebo-corrected design of Figure 4 by estimating the same triple DID on pseudo reform dates in 2019, a year without a fuel tax change in either country. The first is 1 July 2019, estimated on 1 May to 31 August 2019, the calendar window of the tax decrease. The second is 1 November 2019, estimated on 1 September to 31 December 2019. The corrected gaps at the pseudo dates are scaled by the true episodes' full-pass-through benchmarks for comparability with Figure 4. At the autumn pseudo date they lie within half a percentage point of zero in every market-size bin for all three fuel types. At the summer pseudo date they are small for gasoline, between -1.4 and $+1.5$ percentage points, and negative for diesel, reaching -3.3 percentage points, so within-market dispersion of diesel prices compressed over the summer of 2019 as well.

Table D12 tests whether the actual corrected gaps differ from these placebo gaps. We stack the actual and the pseudo-episode samples, estimate the triple DID jointly with stack-specific fixed effects and treatment terms, and report the difference between the actual and the pseudo corrected gap for each bin. For the tax increase, the actual gap exceeds the

Table D12: Actual versus 2019 placebo corrected gaps: stacked estimates

	Number of competing price setters					Joint
	2	3	4	5	6+	<i>p</i> -value
<i>A. Tax decrease (1 July 2020) vs. pseudo decrease (1 July 2019)</i>						
Diesel	0.017*	-0.002	0.019*	0.019	-0.017	0.104
	(0.010)	(0.009)	(0.011)	(0.013)	(0.013)	
E10	-0.024**	-0.051***	-0.050***	-0.073***	-0.104***	<0.001
	(0.010)	(0.011)	(0.010)	(0.015)	(0.016)	
E5	-0.015*	-0.049***	-0.043***	-0.061***	-0.104***	<0.001
	(0.008)	(0.009)	(0.010)	(0.013)	(0.013)	
<i>B. Tax increase (1 January 2021) vs. pseudo increase (1 November 2019)</i>						
Diesel	0.025***	0.036***	0.050***	0.053***	0.063***	<0.001
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	
E10	0.025***	0.031***	0.035***	0.042***	0.046***	<0.001
	(0.003)	(0.003)	(0.003)	(0.004)	(0.004)	
E5	0.021***	0.025***	0.034***	0.035***	0.035***	<0.001
	(0.003)	(0.003)	(0.003)	(0.004)	(0.004)	

Notes: The table reports, for each market-size bin, the difference between the placebo-corrected pass-through gap of the actual tax change (Figure 4) and that of a pseudo reform date in 2019: 1 July 2019, estimated on 1 May to 31 August 2019 (Panel A), and 1 November 2019, estimated on 1 September to 31 December 2019 (Panel B). No tax changed on either date. Both gaps are estimated in one regression that stacks the actual and the pseudo-episode samples with stack-specific fixed effects and treatment terms, so the standard errors of the differences, clustered at the market level, come from the joint fit. Coefficients are divided by the true episodes' full-pass-through benchmarks as in Figure 4. The last column reports the *p*-value of a Wald test that the differences are zero in all bins.

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

placebo gap in every bin, by 2 to 6 percentage points, and the difference is larger in larger markets. For the tax decrease, the differences are negative for gasoline and a Wald test rejects equality of the actual and placebo gaps, in line with the failed placebo of that episode in Section 6, whereas for diesel the actual and placebo gaps are statistically indistinguishable, with a joint *p*-value of 0.10.