Product-Harm Crises and Intermediaries in the Used Car Market*

Masahiro Nishida†

November 5, 2026

Click Here for the Latest Version

Abstract

Product-harm crises, such as recalls and scandals, have been studied as negative information shocks about product quality. In used car markets, such a shock reaches households through dealers, which choose not only the retail price of an affected car but also whether to sell it to a household or to another dealer at wholesale, where the price of an unsold car feeds back into retail prices. This paper asks how dealers adjust the prices of affected cars, how household purchases respond, and how these responses would change if interdealer trade were restricted. Focusing on Mitsubishi Motors' 2016 fuel economy scandal in Japan as a demand shock, I use wholesale auction records and online retail listings to compare the affected models that Mitsubishi built for Nissan with unaffected cars before and after the disclosure. Prices of affected cars fell in both markets, at retail by less at Nissan-affiliated dealers than at independent dealers, whereas the probability of a sale to a household rose only briefly. To measure the long-run effect of the scandal and the consequences of restricting interdealer trade, I estimate an overlapping-generations model in which dealers of different types, each a local monopolist at retail, trade with one another at wholesale, sell to households, and resell unsold cars at wholesale. The estimated loss in households' valuation of affected cars is concentrated at independent dealers, so that the retail price decline at Nissan-affiliated dealers comes largely through the lower wholesale price. In counterfactuals that segment the wholesale market by dealer type or require dealers to sell their trade-ins at retail, the gap between the retail price responses of the two dealer types widens: the interdealer market pools the two losses into one wholesale price and lets dealers pass unwanted cars on rather than discount them.

*I am grateful to Kenneth Hendricks and Jean-François Houde for their continuous support and invaluable advice. I am also grateful to Karam Kang, Ashley Swanson, Yuta Toyama, and participants at the Wisconsin IO Lunch Workshop for their insightful discussion and feedback. I thank Naoki Aizawa, Satoshi Imahie, Taiju Kitano, Hiroyuki Kubota, Tarun Kushwaha, Davi Mittelstadt, Yohei Nishimura, Nicolas Schutz, Kazuma Takakura, Ana Sofia Teles, Makoto Watanabe, and participants at the JEMIOW×J-TREE Joint Workshop, Kyoto, Musashi, and Waseda for their helpful comments and suggestions. I gratefully acknowledge PROTO Corporation for providing their data, and Shigeyoshi Shimizu and the staff members at PROTO Corporation for sharing industry knowledge about the used car market in Japan. This paper used generative AI tools (ChatGPT and Claude) as assistants for data analysis and manuscript preparation. All opinions, findings, and errors are my own.

†Department of Economics, University of Wisconsin-Madison. mnishida@wisc.edu

1 Introduction

Product-harm crises, such as defect revelations, recalls, and corporate scandals, have long been studied as negative information shocks about product quality.[1] Automobiles are a leading example, because recalls are frequent and widely publicized.[2] In used car markets, however, such a shock does not reach households directly, because most household purchases of used cars go through dealers. A dealer that holds an affected car chooses not only its retail price but also which cars to put on its lot and, once a car is there, whether to sell it to a household or to another dealer in the interdealer wholesale market. The retail price that households observe is therefore only one of the dealer's choices.

These choices are linked through the interdealer market. A dealer buys cars at wholesale and sells there the cars that its households do not buy, so when a product-harm crisis lowers the demand for an affected car, the fall in its wholesale price feeds back into retail prices and into whether a dealer keeps the car or passes it on. The effect of such a crisis on the used car market therefore depends on how dealers trade with one another, and measuring it requires taking the interdealer market into account. I study these responses in the Japanese used car market, using as the demand shock the 2016 fuel economy scandal of Mitsubishi Motors, which had overstated the fuel economy ratings of some of its models. This paper asks two questions. First, how do dealers adjust the prices of affected cars, and how do household purchases respond? Second, how would these responses change if interdealer trade were restricted?

I answer the first question by comparing affected with unaffected cars before and after the disclosure of the scandal, at wholesale and at retail. Prices of affected cars fell in both markets, and at retail they fell less at the dealers affiliated with the affected brand than at independent dealers. The probability that an affected car sold to a household within 3 months rose for the cars listed in the months right after the disclosure, at the lower prices, but returned to its previous level in the medium run. This comparison has two limits. Because unaffected cars are substitutes for affected cars, the scandal moves their prices and sales too, so these estimates measure the effect on affected cars net of the effect on unaffected cars rather than the effect on affected cars alone. The comparison also speaks only to the short run and the medium run: a year or more after the disclosure, new model generations and other changes in the market overlay the scandal, so the difference between affected and unaffected cars no longer isolates it.

For the long-run effect and for the second question, I build an equilibrium model of dealers that trade cars with one another at wholesale, sell them to households, and resell unsold cars at wholesale, in which affected and unaffected cars are substitutes at wholesale. Its long-run effect is the difference between the stationary equilibria before and after the scandal. The model's estimates suggest that the loss in households' valuation of affected cars was far larger at independent dealers than at the dealers affiliated with the affected brand. In counterfactuals that restrict interdealer

[1] See, for example, Van Heerde, Helsen, and Dekimpe (2007), Cleeren, Dekimpe, and Van Heerde (2017), and Astvansh, Antia, and Tellis (2025).

[2] See, for example, Liu and Shankar (2015) and van den Bijgaart and Cerruti (2024) on automobile recalls.

trade, the gap between the retail price changes at the two types of dealers widens.

The setting is the Japanese used car market, where dealers play a central role: most used cars that households buy come from dealers. Dealers acquire their inventory from households, as trade-ins or outright purchases, and from other dealers, mainly at auctions open only to licensed dealers, and they differ in how much they rely on each source. Affiliated dealers hold an exclusive dealership contract with a manufacturer and sell its new cars as well as used cars, most of which carry their own brand because they arrive as trade-ins from new car customers. Independent dealers, by contrast, sell only used cars and rely more on auctions. The scandal concerns kei cars, the smallest regulatory class of passenger vehicles in Japan. On April 20, 2016, Mitsubishi disclosed that it had overstated the fuel economy ratings of four kei car models, two Mitsubishi-badged models and their physically identical twins, the Dayz and the Dayz Roox, which Mitsubishi built for Nissan. Retesting found the actual fuel economy 5 to 16 percent below the ratings. Because the models still met the safety standards, no recall was issued and the cars stayed legal to drive and resell, so the disclosure was a pure information shock to their perceived quality. I study the Nissan-badged models because the reputational damage fell mainly on Mitsubishi.

To observe both markets, I combine two datasets. The first consists of listing records from wholesale auto auctions, covering about 70 percent of all auction listings in Japan. Each record gives whether the vehicle sold, its transaction price if it did, its characteristics, and a condition score assigned by inspectors. The second consists of monthly snapshots of one of Japan's largest online used car listing platforms, on which dealers advertise their inventory to households. The snapshots record each listing's asking price, dealer, and vehicle characteristics. The novelty of the data is that they record where each car went when it left the dealer, such as to a household, an auction, or another dealer. Studies of dealer pricing that use online listings, such as those from Cars.com, observe when a listing disappears but not where the car went, and treat its removal as a sale to a household.[3] A dealer that fails to sell a car at retail can instead send it to the interdealer market, and dealers often do so. Because my data record the destination, I observe the price that the dealer set and whether a household bought the car separately, which is what the first question requires. The sample covers kei cars of the 2013 to 2016 vintages from August 2015 to December 2016, 9 months before the disclosure and 8 months after it, and compares the Nissan-badged affected models with the kei cars of six other manufacturers, excluding Mitsubishi-badged cars and Mitsubishi's dealers.

To answer the first question, I compare affected with unaffected cars before and after the disclosure, controlling for vehicle characteristics such as age and mileage, with fixed effects for narrowly defined models, months, and auction venues or dealers. At wholesale, prices of affected cars fall by 57.3 thousand JPY (6.8 percent) in the short run, the first 3 months after the disclosure, and by 33.6 thousand JPY (3.8 percent) in the medium run, August to December 2016. Their unsold rate rises by 7.8 percentage points in the short run but shows no change in the medium run, so the wholesale market adjusts through prices rather than quantities. Retail listing prices fall by

[3] See Huang, Luo, and Xia (2019) and Li, Murry, Tian, and Zhou (2024, 2025).

25.4 thousand JPY (2.3 percent) in the short run and by 17.3 thousand JPY (1.6 percent) in the medium run. The decline is smaller at Nissan-affiliated dealers, 20.9 thousand JPY (1.8 percent) and then 15.7 thousand JPY (1.5 percent), than at independent dealers, 42.9 thousand JPY (4.0 percent) and then 22.0 thousand JPY (2.0 percent). The probability that an affected car sells to a household within 3 months falls, although not significantly, for the cars listed in April 2016 at prices set before the disclosure, rises by 5.5 and 7.7 percentage points at Nissan-affiliated and independent dealers in the short run, and shows no significant change in the medium run. In the medium run, then, affected cars sold at lower prices but as often as before, which suggests that households' willingness to pay fell by about as much as retail prices.

These comparisons have their limits. Because the scandal affects unaffected cars as well, they measure the effect on affected cars net of the effect on unaffected cars rather than the scandal's own effect on affected cars, and because the market changes in many other ways over a longer horizon, they cannot measure its long-run effect. Nor do they reveal how the interdealer wholesale market shapes dealers' responses. I therefore build a model that accounts explicitly for the substitution between affected and unaffected cars in the wholesale market, use it to measure the long-run effect of the scandal, and quantify how the interdealer market shapes dealers' responses.

The model is an overlapping-generations model in which a continuum of dealers enters in every period, each dealer lives for at most 2 periods and holds at most one car, and cars are of two brands, affected and unaffected. Dealers are of three types: Nissan-affiliated dealers hold only affected cars, dealers affiliated with other brands hold only unaffected cars, and independent dealers hold either brand, so that for them the two brands are substitutes at wholesale. Each dealer is the monopolist in its own retail market, which has a single household, and dealers differ in the demand they face and in the cost of holding a car that does not sell. On entry, dealers trade in a competitive interdealer wholesale market, where the supply consists of the trade-ins with which some dealers enter and the cars left unsold by the previous cohort. A dealer that holds a car then posts a retail price once, and its household buys with a probability that decreases in the price. A dealer whose household does not buy sells the car at wholesale in the next period. The wholesale price is therefore both the cost of a car and the value of an unsold one, so a loss in households' valuation of affected cars lowers their wholesale price as well as their retail prices. In a stationary state, the shock is absorbed mainly by retail and wholesale prices, and sale probabilities change little.

I estimate the model in two stages. The first stage estimates the distributions of demand and holding costs across dealers by fitting the pre-scandal stationary equilibrium to pre-scandal retail and wholesale data: listing rates, the means and dispersions of listing prices, sale probabilities, the covariance between prices and sales, and wholesale prices. The second stage introduces the scandal as a loss in households' valuation of affected cars that differs between Nissan-affiliated and independent dealers. Assuming that the market reaches a new stationary equilibrium 2 years after the disclosure, I solve this post-scandal equilibrium and the transition path from the pre-scandal equilibrium to it, and I choose the two losses so that the wholesale price, retail prices, and sale probabilities over the first 8 months of the path match the reduced-form estimates that

compare affected with unaffected cars. The estimates attribute the loss mainly to the households of independent dealers, so that the retail price decline at Nissan-affiliated dealers comes largely through the wholesale price that both dealer types share.

Finally, to answer the second question, I use the estimated model to ask how dealers' responses would change if interdealer trade were restricted. In the baseline, every dealer can sell its trade-in at wholesale at once and buy any car that its type can hold at one price per brand. The first counterfactual scenario segments the wholesale market by dealer type, so that dealers trade only with dealers of their own type. It is motivated by industry practice: manufacturers run their own auction venues, encourage their affiliated dealers to trade there, and restrict the participation of independent dealers, in order to reallocate cars of their own brand within their dealer networks, although this segmentation is incomplete in practice. The second scenario requires a dealer that takes a car in trade to offer it at retail before it can sell it at wholesale. Under segmentation, each dealer type's loss stays in its own wholesale market: the wholesale price of affected cars falls more in the market of independent dealers and less in the manufacturers' networks, the retail decline at Nissan-affiliated dealers shrinks, and the decline at independent dealers deepens. Under mandatory retail of trade-ins, the trade-ins that independent dealers must list are cars that their households rarely buy, so the retail price of affected cars at independent dealers falls more and their sale probability falls, whereas sale probabilities otherwise change little. The interdealer market thus pools the losses of the two dealer types into one wholesale price and lets dealers pass unwanted cars on rather than discount them, which spreads the price declines across dealer types and keeps household purchases nearly unchanged.

Related Literature

This paper contributes to four literatures: on intermediaries in markets for durable assets, on dealers in over-the-counter markets, on product-harm crises, and on wholesale auto auctions. In markets for durable assets, dealers provide liquidity by holding inventory and matching with searching buyers, and the thickness of the market governs how easily assets are reallocated across users (Gavazza, 2011, 2016).[4] For used cars, Gavazza, Lizzeri, and Roketskiy (2014) develop an equilibrium model with trading frictions, Li, Murry, Tian, and Zhou (2024, 2025) show how search frictions shape dealers' inventories, prices, and the gains from intermediation, Huang, Luo, and Xia (2019) model dynamic pricing with learning about demand, and Murry and Zhou (2020) study dealer colocation under consumer search. These papers model the dealer's interaction with searching customers in detail but treat its procurement side in reduced form. I take the opposite modeling choice: this paper abstracts from search frictions and dynamic pricing at retail and models explicitly the interdealer wholesale market, whose price is both the cost of a car and the value of an unsold one and through which cars move across dealers.

Dealers that hold inventory and trade both with customers and with one another have been studied mainly in financial over-the-counter markets, where a dealer-to-customer segment coexists

[4] For a survey of frictions in product markets, including markets for durable goods, see Gavazza and Lizzeri (2021).

with an interdealer segment (Duffie, Gârleanu, and Pedersen, 2005; Hugonnier, Lester, and Weill, 2020). Dunne, Hau, and Moore (2015) and Colliard, Foucault, and Hoffmann (2021) show how dealers intermediate between the two segments and how their access to the interdealer market shapes their inventories and the prices that customers face. The retail and wholesale layers of the used car market parallel these two segments, and I study how a demand shock passes through each layer for a physical durable good whose dealers differ in the demand they face and in their cost of holding a car that does not sell.

A large marketing literature documents that recalls and scandals damage demand for the affected firm's products and brand (Van Heerde, Helsen, and Dekimpe, 2007; Liu and Shankar, 2015; Cleeren, Dekimpe, and Van Heerde, 2017; Astvansh, Antia, and Tellis, 2025). In used car markets, van den Bijgaart and Cerruti (2024) show that a recall raises trading in some segments and lowers it in others, as the information triggers sorting or adverse selection. The event closest to the Mitsubishi scandal is Volkswagen's 2015 emissions scandal, whose effects on used car markets have been measured in three settings: Strittmatter and Lechner (2020) document falling prices and rising listings of Volkswagen cars in Germany, Ater and Yoseph (2022) find that trading volumes in Israel fell only for private sellers, and Che, Katayama, and Lee (2023) measure spillovers within the Volkswagen group in eBay auctions. In the new car market, Bachmann et al. (2023) quantify the damage that the scandal did to the collective reputation of other German manufacturers. Whereas these studies measure how a crisis moves prices and trading volumes, this paper follows the affected cars through dealers, which set the prices that households see and choose whether to sell a car to a household or to another dealer.

Wholesale auto auctions are the mechanism through which used cars are reallocated across dealers: Overby and Forman (2015) and Subramanian and Overby (2017) show that electronic bidding widens the geographic reach of purchases, enables spatial arbitrage, and reduces price dispersion across locations. Other studies examine frictions within the auction, namely adverse selection and search (Genesove, 1993, 1995), information disclosure (Tadelis and Zettelmeyer, 2015), and the post-auction bargaining between seller and buyer, including its efficiency, the bargaining power of the two sides, and the mediators who run it (Larsen, 2021; Larsen and Zhang, 2026; Larsen, Lu, and Zhang, 2026). Rather than studying these frictions, I treat the auction as a competitive market inside an equilibrium model of dealers estimated around a large demand shock, so that the counterfactuals show how restricting interdealer trade changes the responses of wholesale prices, retail prices, and household purchases.

2 Institutional Background

2.1 Used Car Market in Japan

Consumers in Japan purchased an estimated 2.40 million used cars in 2015 (Recruit Carsensor Co., Ltd., 2018). Dealers play a central role in these purchases: 94% of used car purchases by Japanese consumers go through dealers, compared with 71% in the United States (Ministry of

Land, Infrastructure, Transport and Tourism, 2015).

Dealers acquire their inventory from two sources: direct purchases from households, either as trade-ins when a household buys a replacement vehicle or as outright buyouts, and purchases from other dealers. Interdealer trade takes place mainly at wholesale auto auctions, which are open only to licensed dealers.[5] The mix of sources varies with vehicle condition, as direct purchases from households account for a larger share of acquisitions among young, low-mileage vehicles in good condition. At an auction, the seller brings its vehicle to a physical venue and sets a reserve price, and buyers bid in an ascending-price format, either at the venue or online. Most venues hold an auction once a week. A vehicle whose highest bid falls short of the reserve price goes unsold, and such vehicles often stay at the venue to be relisted the following week (USS Co., Ltd., 2021, 2022).

Dealers that retail used cars fall into two types, affiliated and independent. Affiliated dealers sell a manufacturer's new cars under an exclusive dealership contract and also sell used cars; although their used cars may be of any brand, most carry the manufacturer's own brand (Section 3.4). Independent dealers specialize in used cars and range from small owner-operated shops to nationwide chains, some of which run repair shops or car accessory stores alongside their used car business. Because affiliated dealers take in trade-ins when they sell new cars, they have more opportunities to acquire relatively well-kept used cars directly from households, and they rely less on auctions for their inventory than independent dealers do.

The scandal studied in this paper concerns kei cars, the smallest regulatory class of passenger vehicles in Japan.[6] Kei cars are inexpensive to buy and to own because they carry lower taxes and insurance premiums than regular passenger vehicles, and their owners use them mainly for everyday trips such as shopping and commuting. They are especially common in small cities and rural areas, where public transportation is less convenient (Japan Automobile Manufacturers Association, 2016). In 2015, kei cars accounted for 35.9% of new passenger car sales (Japan Automobile Manufacturers Association, 2018) and 43.9% of used car purchases by consumers (Recruit Carsensor Co., Ltd., 2018).

2.2 The Mitsubishi Fuel Economy Scandal

On April 20, 2016, Mitsubishi Motors (hereafter Mitsubishi), a Japanese automaker, disclosed that it had overstated the official fuel economy ratings of four kei car models. Two of these models, the eK Wagon and the eK Space, were sold under Mitsubishi's own brand, and the other two, the Dayz and the Dayz Roox, were their physically identical twins that Mitsubishi developed and built for Nissan Motor (hereafter Nissan), which sold them under its own brand (Mitsubishi Motors Corporation, 2016).[7] The models were launched in 2013 and 2014, and by the end of

[5] Japan Used Car Dealers Association (2021) reports that auctions account for 59% of its member dealers' used car acquisitions, direct purchases from users, as trade-ins or buyouts, for 27%, and other channels for 14%.

[6] Kei cars are limited to a length of 3.40 meters, a width of 1.48 meters, a height of 2.00 meters, and an engine displacement of 0.66 liters, with caps on passenger and cargo capacity.

[7] Mitsubishi Motors, press release, April 20, 2016, `http://www.mitsubishi-motors.com/publish/pressrelease_jp/corporate/2016/news/detailg420.html`. All sources given in the footnotes of this subsection are in Japanese, and links to pages that are no longer online lead to copies saved by the Internet Archive.

March 2016, Mitsubishi had sold 157,000 Mitsubishi-badged units and built 468,000 Nissan-badged units. Throughout the paper, I call these four models the affected models and describe each as Mitsubishi-badged or Nissan-badged according to the brand under which it was sold.

Official fuel economy ratings in Japan come from tests that manufacturers conduct themselves under a procedure prescribed by the government, and the ratings are shown in catalogs and determine eligibility for tax reductions, which gives manufacturers an incentive to inflate them (Tanaka, 2020). Because the test is run under artificial conditions, on-road fuel economy falls short of the rating for every model, so consumers could not detect the misconduct from the gap they experienced (Aizawa, 2020). Mitsubishi had measured the running resistance that enters the test, the resistance from tires and air, with methods that deviated from the regulations and had selectively used favorable data points, inflating the ratings. Government retesting found the actual fuel economy of the four models 5 to 16 percent below their ratings, and the ratings were revised downward in June 2016.[8]

Both companies stopped selling the affected models on the day of the disclosure and resumed taking orders on July 1, 2016, after the ratings had been revised. In June 2016, each announced compensation of 100,000 yen (about 900 USD) per vehicle for customers who were using an affected vehicle at the time of the disclosure, so buyers who acquired one afterward were not eligible.[9] The Consumer Affairs Agency later found the fuel economy claims misleading and, in 2017, ordered both companies to pay administrative surcharges.[10]

The misconduct did not lead to a recall. Japanese regulations require recalls only for vehicles that fail to meet the safety standards, and because the affected models still met them, including the emission standards, the government's remedy was to correct the registered ratings, and the vehicles remained legal to drive and to resell. Fuel economy is nonetheless a key performance measure for kei cars, whose buyers value low running costs, so the disclosure lowered the perceived quality of the affected models without any change to the vehicles themselves.

After the disclosure, the government ordered every manufacturer to review its fuel economy testing, and in May 2016 Suzuki reported that it, too, had tested its models, including kei cars, with methods that deviated from the regulations. Retesting found that every Suzuki model exceeded its

[8]The misconduct came to light when Nissan, which was to lead the development of the next-generation models, measured the fuel economy of a current model with its own equipment in the fall of 2015 and found a large gap from the registered rating (Mitsubishi Motors Corporation, 2016). The retesting results are reported in Ministry of Land, Infrastructure, Transport and Tourism, press release, June 21, 2016, `https://www.mlit.go.jp/common/001149179.pdf`.

[9]Mitsubishi Motors, press release, June 23, 2016, `http://www.mitsubishi-motors.com/publish/pressrelease_jp/corporate/2016/news/detailg622.html`; Nissan, press release, June 23, 2016, `https://global.nissannews.com/ja-JP/releases/160623-02-j`.

[10]In January 2017, the agency issued cease-and-desist orders against both companies under the Act against Unjustifiable Premiums and Misleading Representations and ordered Mitsubishi to pay a surcharge of 485.07 million yen for its regular passenger models. The surcharges for the four affected models were small, 3.68 million yen for Mitsubishi and 3.17 million yen for Nissan, because the surcharge system covers only sales from April 1, 2016, and the amounts were reduced for self-reporting and for the compensation payments. Nissan's surcharge was revoked in December 2018 after an administrative appeal found that Nissan had not failed to exercise due care. Sources: Consumer Affairs Agency, press releases, January 27, 2017, and June 14, 2017, amended July 21, 2017; Administrative Appeal Review Board, report, October 31, 2018.

registered rating, so its ratings stood and its vehicles stayed on sale. Suzuki's disclosure therefore appears to have had little effect on the reputation of its kei cars, which the analysis includes among the unaffected cars.[11]

In the empirical analysis, I compare the Nissan-badged affected models with the kei cars of six unaffected manufacturers and leave out the Mitsubishi-badged models, for two reasons. First, the reputational damage appears to have fallen mainly on Mitsubishi: its stock price fell by 42% within 3 trading days of the disclosure, whereas Nissan's stock price did not fall.[12] Second, according to an industry expert, Mitsubishi instructed its affiliated dealers to stop selling the affected models, both new and used, whereas Nissan gave no such instruction. Prices and sales of the Nissan-badged models therefore reflect the demand shock without the manufacturer's own restrictions on dealer sales, providing a cleaner setting for studying how dealers respond to it.

3 Data and Descriptive Statistics

3.1 Wholesale Auto Auctions

My first dataset consists of listing records from the wholesale auto auctions described in Section 2.1, held at 83 of the roughly 120 auction venues in Japan (Ministry of Land, Infrastructure, Transport and Tourism, 2015) and covering about 70 percent of all auction listings in the country.[13] Each record represents one vehicle offered at auction, whether or not it was sold; a vehicle goes unsold when its highest bid fails to meet the reserve price. Seller identities are not observed, and neither are the number of bidders, the bidding process, or reserve prices.

For each listing, the dataset provides the auction date, the venue, the transaction price for sold vehicles, and the highest bid for unsold vehicles. Vehicle characteristics include mileage, the year of first registration, a repair-history indicator, the manufacturer, the model name, and the model code; repair history indicates that structural frame components have been repaired or replaced and excludes cosmetic repairs. Catalog specifications, such as the registered fuel economy rating, engine displacement, and the manufacturer's suggested retail price, are merged from manufacturer catalogs. A distinctive feature of the auction data is the condition score assigned by trained inspectors at each venue before the sale: inspectors evaluate exterior damage, interior wear, and mechanical condition and assign an overall grade on a scale from S (near-new) down to 1, without considering popularity or rarity. Appendix B describes the scores and how I group them.

[11] By the May 18, 2016 deadline, only Mitsubishi and Suzuki reported irregularities. Suzuki had tested 16 models with methods that deviated from the regulations. Mitsubishi had also used a non-compliant method for most of its regular passenger models, and retesting lowered the ratings of eight of them by about 4 percent on average. Sources: Suzuki Motor, press release, May 18, 2016, `https://www.suzuki.co.jp/release/d/2016/0518/`; Ministry of Land, Infrastructure, Transport and Tourism, press release, August 30, 2016, `https://www.mlit.go.jp/report/press/jidosha08_hh_002372.html`.

[12] The 42% decline is reported by Nihon Keizai Shimbun, April 22, 2016, `https://www.nikkei.com/article/DGXLASFL22HBC_S6A420C1000000/`. From the close on April 19 to the close on April 22, 2016, Nissan's stock price rose by 3.9%, compared with 3.3% for the TOPIX index (author's calculation from daily closing prices).

[13] According to the data provider.

3.2 Online Retail Platform

My second dataset comes from Goo-net, one of the largest online used car listing platforms in Japan, operated by PROTO Corporation. The platform is a dealer-to-household advertising marketplace: dealers nationwide post their inventory, and consumers browse listings and contact dealers. The data consist of monthly snapshots taken on the first day of each month, so individual vehicles can be tracked across months as they remain listed, forming a listing-month panel.

For each listing, the data record a unique listing identifier, the listing price at which the dealer advertises the vehicle (transaction prices are not observed), the date the vehicle was first posted, and vehicle characteristics comparable to those in the auction data. For dealers, I observe a unique dealer identifier, the dealer's name, and its location. Two further variables are central to the analysis. The removal date records when the dealer deleted the vehicle from its inventory management system, and the destination variable categorizes what happened to the vehicle after removal: sale to a consumer, sale at auction, sale to another dealer, transfer to another store within the same company, export, or scrap. Together they allow me to measure when a vehicle actually left the dealer's inventory and whether it reached a household.

Two limitations require attention. First, listings that begin and end between consecutive snapshots are not captured. Second, dealers purchase a limited number of listing slots and do not necessarily post their entire inventory, and both the removal date and the destination are manually entered by dealers, so they contain some missing values. Appendix C examines the destination margin and shows that the recorded destinations behave sensibly.

3.3 Sample Construction

I restrict both datasets to gasoline-powered kei passenger cars of the 2013 to 2016 vintages,[14] observed over the 17 months from August 2015 through December 2016, which include 9 months before the disclosure and 8 months after it. The affected group consists of the Nissan Dayz (model code B21W) and Dayz Roox (B21A), whose Mitsubishi-badged twins are excluded for the reasons given in Section 2.2. The unaffected group consists of the kei cars of the six other manufacturers: Daihatsu, Toyota, Subaru, Honda, Suzuki, and Mazda.[15] In both datasets, I drop vehicles with fewer than 500 kilometers on the odometer, most of which are registered but unused cars sold as substitutes for new cars, and vehicles with 100,000 kilometers or more.[16]

[14] A car's vintage is the year in which it was first registered as a new car.

[15] The other Mitsubishi- and Nissan-badged kei models, mostly the Nissan Moco built by Suzuki, are also excluded. In the sample period, they account for 3.3 percent of the kei cars of these vintages first listed on the platform and 3.5 percent of those listed at auction.

[16] Auction odometers are recorded in whole thousands of kilometers, whereas retail readings are either exact or rounded to the nearest 1,000 kilometers, and many independent and other-brand-affiliated dealers switched from rounded to exact readings in mid-2016. The mileage controls in both markets therefore use readings rounded to the nearest 1,000 kilometers, which removes this change in reporting precision and puts retail readings on the auction scale. On this scale, the mileage bins are 1,000 kilometers wide up to 29,000 kilometers, 5,000 kilometers wide from 30,000 up to 60,000 kilometers, and 10,000 kilometers wide in the 60,000s and the 70,000s, with a single bin for 80,000 kilometers and above.

Because the auction data record the date of each auction, I group auctions into monthly periods running from the 20th of one calendar month to the 19th of the next, so that the disclosure on April 20, 2016 opens a new period, and label each period by the calendar month in which it ends. The auction sample thus covers auctions held from July 20, 2015 through December 19, 2016 and contains 223,120 listings, of which 148,860 are sold with a positive price.

On the platform, the full entry sample consists of the first observed listing of each vehicle, excluding all listings of Mitsubishi's dealers, whether of affected or unaffected cars, and those of one nationwide used car chain that never records where its vehicles went. Of these first listings, 174,209 have a positive listing price after trimming the top and bottom 1 percent of listing prices within each vehicle group and listing month.

I classify each retail dealer by its name into one of three groups: Nissan's franchised dealers, the franchised dealers of the other manufacturers, and independent used car dealers. Because franchised dealers acquire most of their used inventory as trade-ins of their own brand (Section 2.1), the analysis compares, within each vehicle group, the franchised dealers of the car's own brand with independent dealers. For affected cars, these are Nissan's franchised dealers and independent dealers; for unaffected cars, they are the franchised dealers of the car's own brand, counting Toyota and Daihatsu as one brand because Toyota's kei cars are built by Daihatsu, and independent dealers. In both cases, I call the franchised dealers of the car's own brand the affiliated dealers, and Nissan-affiliated or other-brand-affiliated dealers where the brand matters. The design drops 10,657 of the 174,209 first listings, 6.1 percent: 1,741 unaffected cars at Nissan-affiliated dealers, 725 affected cars at other-brand-affiliated dealers, and 8,191 unaffected cars of another brand at other-brand-affiliated dealers.[17] The price analysis uses the remaining 163,552 first listings. To measure when and where new listings leave the dealer's inventory, I keep only the dealers that record an exit destination in every month from August 2015 through December 2016, which I call the reporting sample, and apply the same dealer-type rule to its listings.

3.4 Descriptive Statistics

Throughout the paper, prices are deflated by the consumer price index and expressed in thousand JPY at 2016 prices. Table 1 reports summary statistics for the wholesale auction sample by vehicle group and period. Affected and unaffected kei cars sell at similar average prices and go unsold at similar rates: before the disclosure, the mean transaction price is 898 thousand JPY for affected cars and 882 thousand JPY for unaffected cars, and the shares of listings that go unsold are 0.325 and 0.361. In both groups, mileage and vehicle age are higher after the disclosure than before, a mechanical consequence of restricting the sample to the 2013–2016 vintages, which age over the window.

Table 2 reports descriptive statistics for the first listings of the price analysis by vehicle group, dealer type, and period; the figures that follow are for the pre-disclosure period, and the post-

[17] Including these listings, with the affected cars at other-brand-affiliated dealers as a third dealer-type cell, changes the estimates of Section 4 by less than one standard error.

Table 1: Summary Statistics: Wholesale Auction Sample

	Pre				Post			
Variable	Mean	P25	P50	P75	Mean	P25	P50	P75
Panel A: Nissan-affected kei								
Transaction price (1K JPY)	898	755	897	1,051	860	693	846	1,037
Unsold (share)	0.325				0.316			
Mileage (1K km)	9.0	3.0	5.0	12.0	15.2	5.0	12.0	21.0
Vehicle age (years)	1.32	1.00	1.00	2.00	1.77	1.00	2.00	2.00
High condition score (share)	0.461				0.271			
Above-middle condition score (share)	0.796				0.754			
Repaired (share)	0.056				0.065			
Venues	72				74			
Observations: listings	10,676				9,890			
Observations: sold listings	7,208				6,768			
Panel B: Unaffected kei								
Transaction price (1K JPY)	882	692	868	1,066	880	678	866	1,074
Unsold (share)	0.361				0.301			
Mileage (1K km)	9.9	2.0	5.0	12.0	16.4	4.0	10.0	23.0
Vehicle age (years)	1.31	1.00	1.00	2.00	1.69	1.00	2.00	3.00
High condition score (share)	0.521				0.352			
Above-middle condition score (share)	0.859				0.789			
Repaired (share)	0.040				0.060			
Venues	76				76			
Observations: listings	111,792				90,762			
Observations: sold listings	71,475				63,409			

Notes: Mean and 25th, 50th, and 75th percentiles by vehicle group and period, Nissan-affected kei cars in Panel A and unaffected kei cars in Panel B; share rows report the mean only. Transaction price is in thousand JPY, CPI-deflated to 2016 prices, and is computed over sold listings only. Mileage is in thousand km. Vehicle age is the auction year minus the first-registration year. A high condition score is an auction condition score of 5 or above, and an above-middle score is 4 or above; both shares are over all listings, including those without a recorded score. Venues is the number of auction venues with at least one listing in the cell. The sample is the main auction sample: kei cars of vintage 2013–2016 with 500 km to under 100,000 km on the odometer, Nissan-affected kei cars and unaffected kei cars. Pre covers the auction months August 2015 to April 2016 and post May to December 2016. An auction month runs from the 20th of the previous calendar month to the 19th and is labelled by the month in which it ends, so the sample covers auctions from July 20, 2015 to December 19, 2016 and the post period opens on April 20, 2016.

disclosure columns show the same patterns. Most cars of both groups are listed by the affiliated dealers of their own brand: Nissan-affiliated dealers list 8,199 of the 10,202 affected cars, and other-brand-affiliated dealers list 51,343 of the 71,045 unaffected cars. Within each group, affiliated dealers list their cars at higher prices than independent dealers, 1,201 against 1,067 thousand JPY for affected cars and 1,138 against 1,073 thousand JPY for unaffected cars, or roughly 13 and 6 percent higher. Their cars are also about half a year younger on average and have lower mileage, 5.4 against 9.6 thousand km for affected cars and 6.3 against 11.7 for unaffected cars, and far fewer of them have a repair history: the repaired share is 0.007 against 0.170 for affected cars and 0.005 against 0.128 for unaffected cars.

Table 2: Summary Statistics: Retail First Listings by Vehicle Group and Dealer Type

	Affiliated								Independent							
	Pre				Post				Pre				Post			
Variable	Mean	P25	P50	P75	Mean	P25	P50	P75	Mean	P25	P50	P75	Mean	P25	P50	P75
Panel A: Nissan-affected vehicles (Affiliated = Nissan-affiliated dealers)																
Listing price (1K JPY)	1,201	1,077	1,198	1,334	1,177	1,020	1,188	1,326	1,067	899	1,050	1,213	1,043	891	1,001	1,192
Mileage (1K km)	5.4	2.0	3.0	6.0	9.5	3.0	5.0	12.0	9.6	3.0	6.0	13.0	15.2	5.1	12.0	21.5
Vehicle age (years)	0.97	0.00	1.00	1.00	1.34	1.00	1.00	2.00	1.45	1.00	1.00	2.00	1.83	1.00	2.00	3.00
Listing duration (months)	2.1	1.0	2.0	3.0	2.2	1.0	2.0	3.0	3.2	1.0	2.0	4.0	3.2	1.0	2.0	4.0
Repaired (share)	0.007				0.009				0.170				0.173			
Free warranty (share)	0.986				0.985				0.846				0.871			
Paid warranty (share)	0.005				0.000				0.126				0.072			
Observations	8,199				6,011				2,003				2,387			
Dealers	506				506				795				882			
Panel B: Unaffected vehicles (Affiliated = other-brand-affiliated dealers: Toyota-Daihatsu, Honda, Subaru, Suzuki, Mazda)																
Listing price (1K JPY)	1,138	950	1,152	1,326	1,143	955	1,159	1,331	1,073	884	1,072	1,273	1,075	874	1,080	1,281
Mileage (1K km)	6.3	2.0	3.9	7.0	8.4	2.6	4.4	9.5	11.7	3.0	7.0	16.0	16.5	4.4	11.8	24.0
Vehicle age (years)	1.05	1.00	1.00	1.00	1.11	0.00	1.00	1.00	1.55	1.00	1.00	2.00	1.84	1.00	2.00	3.00
Listing duration (months)	2.3	1.0	2.0	3.0	2.6	1.0	2.0	3.0	3.2	1.0	2.0	4.0	3.3	1.0	2.0	4.0
Repaired (share)	0.005				0.008				0.128				0.134			
Free warranty (share)	0.991				0.987				0.838				0.843			
Paid warranty (share)	0.004				0.001				0.114				0.094			
Observations	51,343				50,619				19,702				23,288			
Dealers	2,564				2,782				3,345				3,674			

The unit is the first observed listing of each car. Listing prices are in thousand JPY, CPI-deflated to 2016 prices; mileage is in thousand km; vehicle age is the listing year minus the registration year; listing duration is the number of monthly snapshots from the car's first to its last listing month, inclusive. Share rows are shares of listings (the repair share among listings that record the repair history). Dealers is the number of dealers with a first listing in the cell. Affiliated dealers are the franchised dealers of the car's own brand: Nissan's dealers for the affected cars in Panel A and, for the unaffected cars in Panel B, the franchised dealers of Toyota and Daihatsu (counted as one brand), Honda, Subaru, Suzuki, and Mazda that list a car of their own brand. Sample: first observed listings of Nissan-affected and unaffected kei cars of vintage 2013–2016, first listed August 2015 to December 2016, with 500 km to under 100,000 km on the odometer and listing prices trimmed at the top and bottom 1 percent within vehicle group and listing month, excluding Mitsubishi dealers and one nationwide chain that never records where a car went. The sample keeps affected cars at Nissan-affiliated and independent dealers and unaffected cars at independent dealers and, when the car carries the dealer's own brand (Toyota and Daihatsu counted as one), at other-brand-affiliated dealers; this drops 10,657 first listings (6.1 percent): 1,741 unaffected cars at Nissan-affiliated dealers, 725 affected cars and 8,191 unaffected cars of another brand at other-brand-affiliated dealers. Pre covers first listings from August 2015 to April 2016 and post those from May to December 2016.

Table 3 reports how the new listings of the reporting sample leave the dealer's inventory within 3 months in Panel A and within 6 months in Panel B: the percentage sold to a household and the percentage that leaves for any other destination, that is, a sale at auction, a sale to another dealer, a transfer to another store of the same company, or an export, which accounts for about 0.5 percent of these other exits. The two dealer types differ in both channels. Within 3 months, independent dealers sell a larger share of their new listings to households than affiliated dealers, 63.6 against 49.6 percent of the affected cars first listed before the disclosure and 57.6 against 50.2 percent of the unaffected cars, whereas affiliated dealers send far more of their cars to other destinations, 37.9 against 9.7 percent of affected cars and 30.6 against 14.1 percent of unaffected cars. Extending the window to 6 months raises the household share of affected cars little at affiliated dealers, from 49.6 to 54.7 percent, but considerably at independent dealers, from 63.6 to 77.5 percent; the difference between dealer types is smaller for unaffected cars. A car that neither sells nor leaves through another channel within 3 months stays on an independent dealer's lot and is usually sold to a household later. Because these raw comparisons confound composition and market-wide changes, Section 4 turns to regression-adjusted comparisons against the unaffected control group.

Table 3: Exits of New Listings within 3 and 6 Months

	Affiliated						Independent					
	Pre			Post			Pre			Post		
Vehicle group	N	Household	Other exit	N	Household	Other exit	N	Household	Other exit	N	Household	Other exit
Panel A: Exit within 3 months												
Nissan-affected kei	4,103	49.6	37.9	4,086	52.0	36.7	1,551	63.6	9.7	2,430	61.9	12.3
Unaffected kei	17,810	50.2	30.6	23,702	51.0	25.9	14,967	57.6	14.1	23,760	55.5	16.7
Panel B: Exit within 6 months												
Nissan-affected kei	4,103	54.7	42.9	4,085	56.2	41.9	1,551	77.5	13.2	2,428	75.7	14.8
Unaffected kei	17,807	58.7	37.7	23,689	61.0	34.8	14,966	71.5	18.0	23,748	68.4	21.3

N is the number of new listings (the first listing month of each car) in the cell. Household and other exit are the percentages of these listings that leave the dealer within the listing month and the next 2 (Panel A) or 5 (Panel B) months, with a recorded sale to a household and with any other destination, such as an auction, another dealer, or another store of the same company; the remaining listings are still listed at the end of these months. Sample: new listings of Nissan-affected and unaffected kei cars of vintage 2013–2016 with 500 km to under 100,000 km on the odometer at dealers that report an exit destination in every month from August 2015 to December 2016 and in each month of the car's 3 (Panel A) or 6 (Panel B) months, excluding Mitsubishi dealers and one nationwide chain that never records where a car went. Affiliated dealers are Nissan-affiliated dealers in the Nissan-affected kei row and other-brand-affiliated dealers (the franchised dealers of Toyota and Daihatsu, counted as one brand, Honda, Subaru, Suzuki, and Mazda that list a car of their own brand) in the unaffected kei row. In both panels, pre covers cars first listed August 2015 to February 2016 and post those first listed May to December 2016; cars first listed in March or April 2016 are excluded. In Panel B, the 6 months of cars first listed December 2015 to February 2016 extend past the April 20 announcement.

4 Descriptive Analysis

This section documents how the disclosure changed the prices of affected vehicles and the rates at which they sold, in the wholesale and retail markets. Throughout, I compare the affected Nissan models with unaffected kei cars of the same vintages before and after the disclosure on April 20, 2016, in two steps. Event studies trace the difference between the two groups month by month, with the last period before the disclosure as the reference period. A second set of specifications

then aggregates the months into a pre period and a post period, which is split into a short run, the first 3 months after the disclosure, May to July 2016, and a medium run, August to December 2016; the split separates the first months, in which dealers reprice and restock, from the months that follow. At retail, the short run consists of the cars first listed in June and July 2016, because the May 2016 listings include cars priced before the disclosure, and the consumer-sale analysis adds the April 2016 entrants, which met the disclosure at prices set before it, as a very short run. Section 4.1 applies these specifications to the wholesale market, where the outcomes are the transaction price and whether a lot goes unsold. Sections 4.2 and 4.3 turn to the retail market, where the outcomes are the price at which a vehicle is first listed and the probability that a new listing sells to a household within 3 months. For the retail outcomes, I also estimate the changes separately for affected vehicles at Nissan-affiliated and at independent dealers, comparing each with its own pre-disclosure level, while the unaffected vehicles at both types of dealers form the comparison group. Because reputational spillovers could touch other kei cars as well, all estimates are changes for affected vehicles relative to unaffected kei cars rather than absolute effects.

4.1 Wholesale Auction Prices and Unsold Rates

On the auction calendar of Section 3.3, the period labelled April 2016 is the last before the disclosure and serves as the reference period, so the post period begins in May 2016, as at retail. The event-study specification is

$$Y_i = \sum_{\tau \neq \text{Apr. 2016}} \beta_\tau \, \text{Affected}_i \times \mathbf{1}[t_i = \tau] + X_i'\gamma + \text{FEs} + \varepsilon_i, \tag{1}$$

where Affected_i is a dummy variable that equals 1 if vehicle i is one of the affected Nissan models and 0 if it is an unaffected kei car, t_i is the period in which vehicle i is observed, the auction period at wholesale and the month of first listing at retail, and $\mathbf{1}[t_i = \tau]$ equals 1 if that period is τ. Each β_τ is thus the difference in the outcome between affected and unaffected vehicles in period τ relative to April 2016, the reference period. At wholesale, Y_i is the transaction price[18] or an unsold dummy that equals 1 if the lot does not sell and 0 otherwise, whose coefficients are reported in percentage points. The controls X_i are dummies for the mileage bins of Section 3.3, vehicle-age bins (0, 1, and 2 or more years), indicators for high, middle, and missing condition scores (Appendix B), and a repair-history indicator, and the term FEs collects auction-period, model-code by grade (trim) by vintage, and auction-venue fixed effects.[19] Standard errors are clustered at the model-code by vintage level. The aggregated specification is

$$Y_i = \beta_{\text{SR}} \, \text{Affected}_i \times \text{PostSR}_{t_i} + \beta_{\text{MR}} \, \text{Affected}_i \times \text{PostMR}_{t_i} + X_i'\gamma + \text{FEs} + \varepsilon_i, \tag{2}$$

where PostSR_{t_i} equals 1 if period t_i falls in the short run, May to July 2016, PostMR_{t_i} equals 1 if it falls in the medium run, August to December 2016, and each equals 0 otherwise. The period

[18] Table 12 in Appendix A reports the short-run and medium-run estimates for the log transaction price.

[19] In this and the following regressions, observations that form a fixed-effect group on their own are dropped, so the numbers of observations in the tables are up to 2.5 percent smaller than the sample sizes given in the text.

fixed effects absorb the levels of the two post indicators, and the vehicle fixed effects absorb the level of Affected_i.

Figure 1 shows that, relative to unaffected kei cars, transaction prices of affected vehicles fall by 58.8 thousand JPY in the auction period that opens on the announcement, recover to declines of 20.0 to 25.2 thousand JPY between August and November 2016, and fall by 31.6 thousand JPY in December, remaining below the reference level throughout. The share of affected lots that go unsold jumps by 12.3 percentage points in the first post-disclosure period and then returns to its pre-disclosure level.

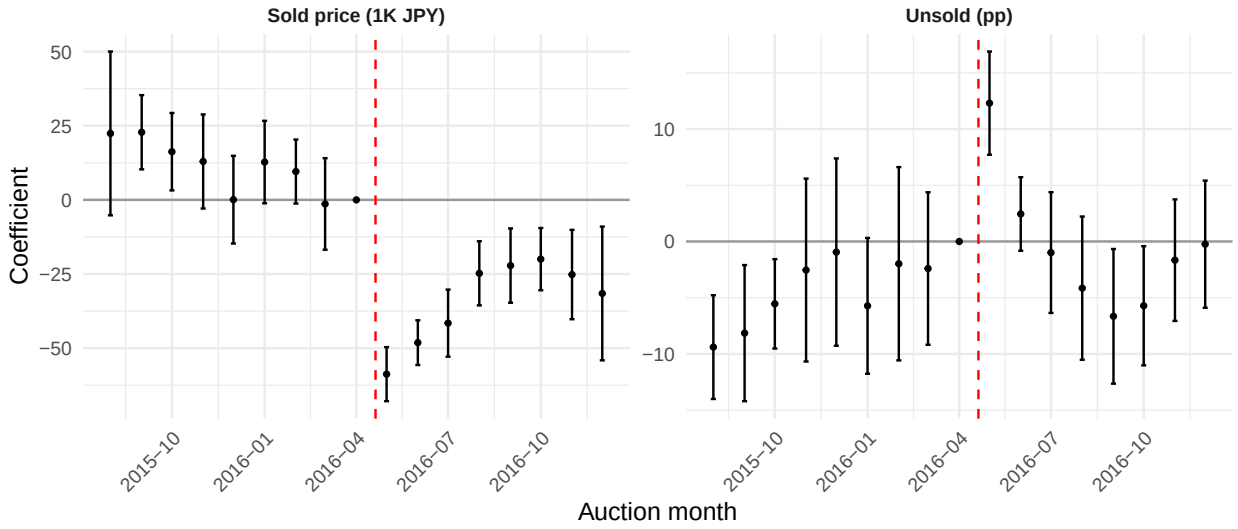


Figure 1: Wholesale Auction Event Studies.

The left panel plots estimates from equation (1) for the transaction price in thousand JPY at 2016 prices, the right panel for the unsold dummy in percentage points. Auction periods run from the 20th of the previous calendar month to the 19th and are labelled by the month in which they end. The dashed vertical line marks April 20, 2016, the disclosure date; the reference period is April 2016. Error bars show 95% confidence intervals based on standard errors clustered at the model-code by vintage level.

Table 4 summarizes the average changes. Relative to unaffected kei cars, wholesale prices of affected vehicles fall by 57.26 thousand JPY in the short run and by 33.56 thousand JPY in the medium run, with standard errors of 7.09 and 8.35, so most of the initial decline persists through the end of 2016. The unsold rate rises by 7.76 percentage points in the short run but changes by a statistically insignificant −0.14 percentage points in the medium run, with standard errors of 2.24 and 2.55. Once the first weeks had passed, affected vehicles sold at auction at their usual rate, but at lower prices. Appendix Table 16 splits these estimates between Nissan's own auction network, where the affected brand's franchised dealers sell their trade-ins, and all other venues. The price declines at the two types of venue differ by a statistically insignificant 8.24 thousand JPY in the short run and −0.23 thousand JPY in the medium run, with standard errors of 7.81 and 9.25. The unsold rate rises more at Nissan's venues, by 6.91 and 6.35 percentage points with standard errors of 5.84 and 5.90, but this difference is not statistically significant in either horizon.

Table 4: Wholesale Auction Changes

	Sold price (1K JPY)	Unsold (pp)
Nissan Affected × Post SR, May–Jul 2016	-57.26***	7.76***
	(7.09)	(2.24)
Nissan Affected × Post MR, Aug–Dec 2016	-33.56***	-0.14
	(8.35)	(2.55)
N	148,714	222,999
Adj. R2	0.90	0.22
FE: Auction month	X	X
FE: Model × grade × vintage	X	X
FE: Venue	X	X

The sold price, in thousand JPY at 2016 prices, is estimated on sold listings with a positive price; the unsold indicator, 100 if the car is not sold and 0 otherwise, on all listings, with the same right-hand side. $*$ p $<$ 0.1, $**$ p $<$ 0.05, $***$ p $<$ 0.01. Standard errors clustered at the model-code by first-registration-year level are in parentheses. Controls: mileage bins, vehicle-age bins (0/1/2+ years), missing, high, and middle condition-score indicators, repair history. Fixed effects: auction month, model code by grade by registration year, auction venue. Affected kei cars are the Nissan-badged models subject to the fuel-economy misconduct disclosure (Dayz and Dayz Roox) of vintage 2013–2016; the comparison group is unaffected kei cars of the main auction sample (500 km to under 100,000 km on the odometer). The sample covers the auction months August 2015 to December 2016; an auction month runs from the 20th of the previous calendar month to the 19th and is labelled by the month in which it ends. Pre covers August 2015 to April 2016 and post the auction months from May 2016 on; the May 2016 auction month opens on April 20, 2016. The short run is May to July 2016 and the medium run August to December 2016, both measured against the same pre period.

4.2 Retail Listing Prices

For each vehicle in the price sample of Section 3.3, the outcome is the listing price in the first monthly snapshot that contains it, the price that the dealer sets when the vehicle enters the platform.[20] Snapshots are taken on the first day of each month, so April 2016 is the last snapshot before the disclosure and the reference month, and the post period begins with the listings first observed in May 2016. The May 2016 cohort, however, mixes cars priced before and after the disclosure. The May 1 snapshot admits every car first published after the April 1 snapshot, including cars published before April 20 at pre-disclosure prices, and the publication date that the platform records cannot separate the two groups, because it is imputed as the snapshot date for a large share of affiliated dealers' listings. The event studies keep every listing month, whereas the aggregated specifications measure the short run with the listings first observed in June and July 2016 and

[20] In a given month, 16.5 percent of the price-sample listings that were also listed in the previous snapshot show a different asking price from that snapshot. Among the listings observed in at least 2 consecutive snapshots by December 2016, 29.1 percent change their price at least once.

exclude the 9,998 listings first observed in May 2016, so that they use 153,554 of the 163,552 first listings. The pooled specifications are equations (1) and (2) with the listing price as Y_i,[21] controls for mileage bins, binned vehicle age, repair history, and free and paid warranty coverage, and listing-month, model-code by grade by vintage, and dealer fixed effects, with standard errors clustered by dealer. With dealer fixed effects, the coefficients measure repricing within dealers rather than shifts of listings across dealers.

To compare dealer types, the aggregated specification becomes

$$\begin{aligned} Y_i = {} & \text{Affected}_i \times \text{Affiliated}_i \times \left(\beta_{N,\text{VSR}} \, \text{PostVSR}_{t_i} + \beta_{N,\text{SR}} \, \text{PostSR}_{t_i} + \beta_{N,\text{MR}} \, \text{PostMR}_{t_i}\right) \\ & + \text{Affected}_i \times \text{Independent}_i \times \left(\beta_{I,\text{VSR}} \, \text{PostVSR}_{t_i} + \beta_{I,\text{SR}} \, \text{PostSR}_{t_i} + \beta_{I,\text{MR}} \, \text{PostMR}_{t_i}\right) \\ & + X_i'\gamma + \text{FEs} + \varepsilon_i, \end{aligned} \tag{3}$$

where Affiliated_i equals 1 if vehicle i is listed by an affiliated dealer, which for an affected car is a Nissan-affiliated dealer, Independent_i equals 1 if it is listed by an independent dealer, and Y_i is the listing price here and the consumer-sale dummy in Section 4.3. The very-short-run indicator PostVSR_{t_i} equals 1 for the cars first listed in April 2016 and enters only the sale regressions, because the April 2016 listing prices were set before the disclosure and therefore belong to the pre period of the price regressions. The dealer and vehicle fixed effects absorb the levels of the two dealer-type indicators and of their interactions with Affected_i, because the affected cars at affiliated dealers are exactly the cars at Nissan-affiliated dealers, and the listing-month fixed effects absorb the levels of the post indicators. Because the listing-month fixed effects are common to both dealer types, each dealer-type coefficient is measured against the market-wide change for unaffected kei cars. The event-study version by dealer type replaces the post indicators with the month indicators $\mathbf{1}[t_i = \tau]$.

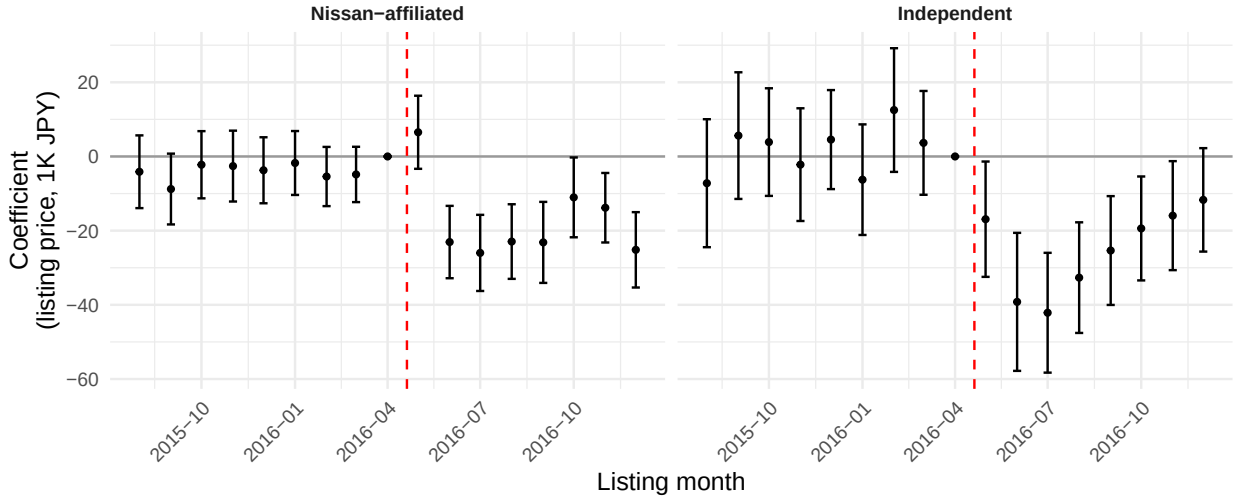


Figure 2: Retail Listing Price Event Study by Dealer Type.

Each panel plots estimates from the event-study version of equation (3), in thousand JPY at 2016 prices. The dashed vertical line marks April 20, 2016, the disclosure date; the reference month is April 2016. Error bars show 95% confidence intervals based on standard errors clustered at the dealer level.

[21] Table 11 in Appendix A reports the short-run and medium-run estimates for the log listing price.

Figure 2 plots the event-study estimates of the listing price by dealer type.[22] At affiliated dealers, the coefficient is slightly positive in May 2016, consistent with a May cohort that mixes cars published before and after the disclosure, and lies between 11.0 and 26.0 thousand JPY below zero from June 2016 on. At independent dealers, by contrast, prices already fall by 16.9 thousand JPY in May 2016, reach their largest decline of 42.1 thousand JPY in July, and recover to a statistically insignificant decline of 11.7 thousand JPY in December, so the difference between the two dealer types is largest right after the disclosure.

Table 5 reports the short-run and medium-run estimates. Listing prices of affected vehicles fall by 25.38 thousand JPY in the short run and by 17.29 thousand JPY in the medium run. The decline is larger at independent dealers in both horizons, 42.87 against 20.88 thousand JPY at Nissan-affiliated dealers in the short run and 22.02 against 15.67 thousand JPY in the medium run, but it also eases far more there, so the gap between the two dealer types narrows from about 22 to about 6 thousand JPY.

4.3 Consumer Sale Rate

Listing prices are set by dealers, so I also ask whether affected vehicles sold to households at a different rate. The outcome equals 1 if a new listing sells to a household within 3 months and 0 otherwise, whether the car leaves for another destination or is still listed, with coefficients in percentage points.[23] A listing counts as sold if it is gone by the third monthly snapshot after its first one and its recorded destination is a sale to a consumer. Listings are grouped by their first month into three post periods. The very short run is the April 2016 entrants, which were priced at the April 1 snapshot, before the disclosure, and faced the disclosure for 72 of the 91 days of their 3-month window, from April 1 to July 1. The short run covers the June and July 2016 entrants and the medium run the August to December 2016 entrants. The March 2016 entrants are excluded because most of their window precedes the disclosure, and the May 2016 entrants for the reason given in Section 4.2. The event-study version excludes only the March entrants, so it shows the April 2016 entrants and every later cohort, including May, and uses February 2016 as its reference month; the window of the February entrants ends with the May 2016 snapshot, but only its last 11 of 90 days fall after April 20.

The estimation sample consists of the new listings of the reporting sample of Section 3.3 whose dealer records an exit destination in each month of the 3-month window, 92,480 of the 146,874 new listings in the remaining entry months that pass the dealer-type rule, which drops 9,503 new listings, 6.1 percent. Because the share of dealers that record destinations differs across dealer types and, among other-brand-affiliated dealers, across brands, I reweight this sample to these 146,874 new

[22] The pooled event study, plotted in Appendix Figure 6, shows that listing prices of affected vehicles are unchanged in May 2016, the first snapshot after the disclosure, fall by 23.9 to 30.1 thousand JPY from June to September 2016, and remain 13.9 to 21.0 thousand JPY lower from October to December.

[23] Appendix A reports an alternative measure, the monthly hazard of a sale to a household, which follows every listed vehicle, including those listed before the disclosure. On this measure, the hazard of affected vehicles at independent dealers changes by −0.013 in a short run that starts in April 2016 and by −0.026 in the medium run, the latter significant at the 5 percent level, and at Nissan-affiliated dealers by 0.000 and 0.005.

Table 5: Retail Listing Price and 3-Month Consumer Sale Changes

	Listing price (1K JPY)		3-month consumer sale (pp)	
	Pooled	By dealer type	Pooled	By dealer type
Panel A: Very short run, Apr 2016 entrants (consumer sale only)				
Nissan Affected × Post VSR			-3.70	
			(2.33)	
Nissan Affected × Nissan Affiliated × Post VSR				-3.70
				(2.72)
Nissan Affected × Independent × Post VSR				-3.76
				(3.13)
Panel B: Short run, Jun–Jul 2016				
Nissan Affected × Post SR	-25.38***		5.96***	
	(3.03)		(2.01)	
Nissan Affected × Nissan Affiliated × Post SR		-20.88***		5.54**
		(3.50)		(2.37)
Nissan Affected × Independent × Post SR		-42.87***		7.67***
		(5.44)		(2.60)
Panel C: Medium run, Aug–Dec 2016				
Nissan Affected × Post MR	-17.29***		-0.45	
	(2.21)		(1.45)	
Nissan Affected × Nissan Affiliated × Post MR		-15.67***		0.38
		(2.68)		(1.76)
Nissan Affected × Independent × Post MR		-22.02***		-2.32
		(3.46)		(1.92)
N	150,975	150,975	90,160	90,160
R2 (within)	0.13	0.13	0.00	0.00
FE: Listing month	X	X	X	X
FE: Model × grade × vintage	X	X	X	X
FE: Dealer	X	X	X	X

$*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$. Standard errors clustered by dealer are in parentheses; for the consumer sale they hold the weights fixed. The listing price is in thousand JPY at 2016 prices. The consumer sale is a dummy for a sale to a household within the car's first listing month or the next 2 months, estimated as a weighted linear probability model with coefficients in percentage points. Samples (Section 3.3): for the price, the 153,554 first listings of the price sample first listed August 2015 to April 2016 or June to December 2016; for the sale, the new listings of the reporting sample first listed August 2015 to February 2016, in April 2016, or June to December 2016 whose dealer reports destinations in each month of the car's 3 months, weighted within period, dealer type, and vehicle group to the badge mix of all new listings (Section 4.3). Both samples exclude the listings outside the two dealer-type cells (10,657 first listings, 6.1 percent, of the price sample; 9,503 new listings, 6.1 percent, of the sale sample). Each post period is measured against the pre period, August 2015 to April 2016 for the price and, by first listing month, August 2015 to February 2016 for the sale. The very short run (Panel A) is the April 2016 entrants. They were priced before the April 20 disclosure, so the price has no very-short-run term, and they faced the disclosure for 72 of the 91 days of their 3-month window, April 1 to July 1; under a constant sale hazard within the window, dividing a Panel A coefficient by 72/91 gives, to first order, the change for a cohort exposed to the disclosure for its whole window. The table reports the regression coefficients, not these scaled values. The short run (Panel B) is the cars first listed in June and July 2016 and the medium run (Panel C) those first listed August to December 2016. The cars first listed in May 2016 are excluded from both outcomes: the May 1 snapshot admits cars first published after the April 1 snapshot, so this cohort mixes cars published before and after the disclosure, and the publication date is imputed as the snapshot date for a large share of affiliated dealers' listings. The March 2016 entrants, whose window lies mostly before the disclosure, are excluded from the sale. Controls: repair history, mileage bins, vehicle-age bins (0/1/2+ years), and, for the price, free and paid warranty indicators. Fixed effects: (first) listing month, model code by grade by vintage, dealer. The comparison group is the unaffected kei cars; affiliated dealers are the franchised dealers of the car's own brand, so the affected cars at affiliated dealers are those at Nissan-affiliated dealers.

listings: within each period, dealer type, and vehicle group, each listing receives the ratio of all new listings of its car's brand to those in the estimation sample. As the cars at other-brand-affiliated dealers carry the dealer's own brand, these ratios correct for the reporting rates of the individual brand networks, assuming that, within a brand, cars sell at the same rate at dealers that record destinations and at those that do not.

The specifications are equations (2) and (3) and their event-study versions, with the sale dummy as the outcome, these weights, and the price controls except the warranty indicators, and with standard errors clustered by dealer; in the pooled aggregated specification, a term $\beta_{\mathrm{VSR}}\,\mathrm{Affected}_i \times \mathrm{PostVSR}_{t_i}$ captures the very short run. The listing price is excluded from the controls, since the price and the sale probability are jointly determined, so the estimates measure the change in sales inclusive of the price changes documented above. Appendix Figures 7 and 8 plot the event-study estimates by entry month; the pooled series is -3.70 percentage points for the April 2016 entrants, peaks at 6.78 percentage points for the June entrants, and lies between -2.24 and 1.39 percentage points from the August entrants on.

In the aggregated estimates of Table 5, the 3-month consumer-sale probability of affected vehicles changes in the very short run by -3.70 percentage points at Nissan-affiliated dealers and by -3.76 percentage points at independent dealers, with standard errors of 2.72 and 3.13, and by -3.70 percentage points pooled, none of them statistically significant. Because the April 2016 entrants were exposed to the disclosure for only part of their window, these coefficients understate in absolute value the change for a cohort exposed to the disclosure throughout its window; divided by 72/91, they correspond to -4.67 and -4.75 percentage points at the two dealer types and -4.68 percentage points pooled, which remain statistically insignificant because their standard errors scale by the same factor.[24] In the short run, by contrast, the sale probability rises by 5.54 percentage points at Nissan-affiliated dealers and by 7.67 percentage points at independent dealers, significant at the 5 and 1 percent levels, and by 5.96 percentage points pooled, significant at the 1 percent level. The medium-run changes of 0.38 and -2.32 percentage points at the two dealer types, with standard errors of 1.76 and 1.92, and of -0.45 percentage points pooled are not statistically significant. Together with the price estimates in the same table, these results suggest that affected vehicles sold to households less often at the prices set before the disclosure, although this decline is imprecisely estimated. Once dealers had repriced, affected vehicles were listed at lower prices and sold more often at both dealer types in the short run, whereas in the medium run they remained 15.67 and 22.02 thousand JPY cheaper at affiliated and independent dealers and their sale probability showed no statistically significant change.

[24] The window of an April 2016 entrant runs from April 1 to July 1, 91 days, of which the first 19, a share $s = 19/91$, precede the disclosure. If the hazard of a sale to a household is constant before the disclosure and constant after it, the probability that the car is not sold by the end of the window is $(1 - q^{\mathrm{pre}})^s (1 - q^{\mathrm{shock}})^{1-s}$, where q^{pre} and q^{shock} are the 3-month sale probabilities of a cohort that does not face the disclosure and of one that faces it throughout its window. To first order in $q^{\mathrm{shock}} - q^{\mathrm{pre}}$, the change in the observed sale probability is therefore $(1 - s)$ times the change for a fully exposed cohort, so dividing a very-short-run coefficient and its standard error by $1 - s = 72/91$ gives that change. Section 6.3 applies the same blend to the model's period-0 cohort.

5 Model

Section 4 examined the wholesale and retail markets separately, but dealers tie the two markets together. Dealers buy cars at wholesale, and they return there the cars that households do not buy. The wholesale price is therefore both the cost of a car and the value of an unsold one. In turn, the retail demand that dealers expect determines what they are willing to pay at wholesale. This section builds a model of this link. Section 6 takes it to the data, and Section 8 uses it to ask how the organization of the wholesale market shapes the pass-through of the scandal to wholesale and retail outcomes. Section 5.1 describes the environment, Sections 5.2–5.4 solve it, and Section 5.5 compares the stationary equilibria before and after a demand shock.

5.1 Environment

The model is an overlapping-generations model of used car dealers. A dealer lives for at most 2 periods: it is *young* in the period in which it enters and *old* in the next period, after which it exits. In every period, a continuum of young dealers of mass D_d of each dealer type $d \in \{N, I, B\}$ enters: Nissan-affiliated dealers (N), independent dealers (I), and other-brand-affiliated dealers (B). Individual dealers are indexed by i. Cars come in two brands, $b \in \{A, U\}$: brand A stands for the affected Nissan kei models and brand U for the unaffected kei models of the six other manufacturers. Cars of a brand are identical, and all cars are used cars. Each dealer has a capacity of at most one car, so a physical dealership that lists many cars corresponds to as many dealers in the model. A dealer of dealer type d can hold a car of brand $b \in \mathcal{J}_d$ or no car, where $\mathcal{J}_N = \{A\}$, $\mathcal{J}_I = \{A, U\}$, and $\mathcal{J}_B = \{U\}$: affiliated dealers hold only their own brand, and independent dealers hold either brand. The model thus has four retail cells, (N, A), (I, A), (I, U), and (B, U).

On entry, dealer i of dealer type d draws its market demand $\xi_i \sim G_d$ for the brands that it can hold: ξ_{iA} for a Nissan-affiliated entrant, ξ_{iU} for an other-brand-affiliated entrant, and a possibly correlated pair (ξ_{iA}, ξ_{iU}) for an independent entrant. The demand ξ_{ib} is the valuation that the dealer's household attaches to a car of brand b (Section 5.2).

In every period, a mass $E_b > 0$ of cars of brand b enters the model; I call E_b the inflow of brand b. The inflow arrives as the endowments of entering dealers, such as trade-ins, so some young dealers enter with a car and the rest without one. Two kinds of shares describe how the inflow is allocated. The endowment share ω_b splits the inflow of brand b between affiliated and independent entrants: a share ω_A of the brand-A inflow goes to Nissan-affiliated entrants and the remaining share $1 - \omega_A$ to independent entrants, and a share ω_U of the brand-U inflow goes to other-brand-affiliated entrants and the remaining share $1 - \omega_U$ to independent entrants. The endowment rate γ_{db} is the share of the entrants of dealer type d that are endowed with a car of brand b. Within a dealer type, endowments are assigned at random, so an entrant of dealer type d enters with a car of brand b with probability γ_{db} and without a car with probability $1 - \sum_{b \in \mathcal{J}_d} \gamma_{db} > 0$. The two kinds of shares are linked because the endowed entrants of a dealer type hold exactly the part of

the inflow allocated to that dealer type:

$$\gamma_{NA} D_N = \omega_A E_A, \qquad \gamma_{IA} D_I = (1-\omega_A) E_A, \qquad \gamma_{BU} D_B = \omega_U E_U, \qquad \gamma_{IU} D_I = (1-\omega_U) E_U. \quad (4)$$

Endowments are assigned independently of demand, so apart from their dealer type and endowment, dealers differ only in their demand.

There is a single, competitive wholesale market, in which dealers trade cars of both brands: a car of brand b trades at the wholesale price w_b. I study the stationary state, in which w_A and w_U are constant over time. Section 6.3 describes the equilibrium when the wholesale prices vary over time, which the estimation uses for the months after the disclosure.

A dealer's life has two stages, which Figure 3 shows for a Nissan-affiliated dealer in panel (a) and an independent dealer in panel (b). In the wholesale stage, on entry, an affiliated entrant trades only its own brand: if endowed, it keeps its car or sells it and exits, and if not, it buys a car or exits. An independent entrant has one more option in each case: if endowed, it can also sell its car and buy a car of the other brand, and if not, it can buy a car of either brand. An entrant that holds a car then posts a price in the retail stage, as the monopolist in its own retail market, which has a single household. If the household buys the car, the dealer exits. If not, the dealer cannot offer the car at retail again: it carries the car into the next period and, now old, sells it at wholesale and exits.

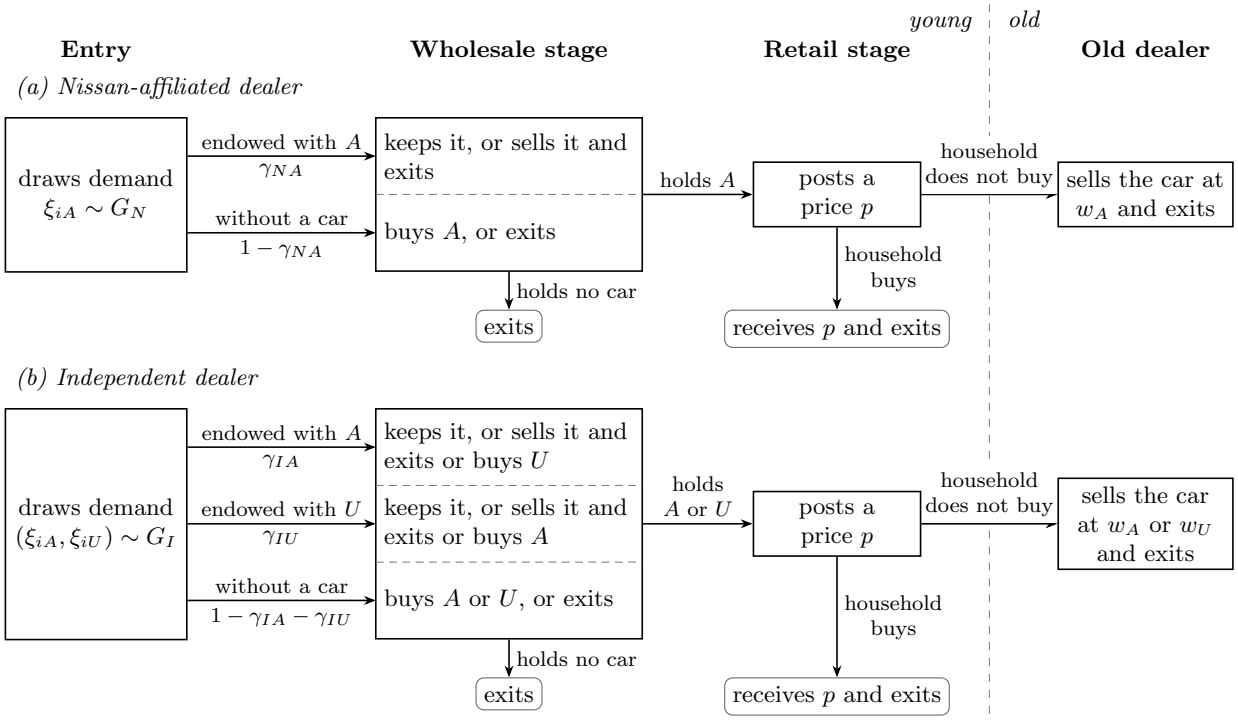


Figure 3: Timing of a Dealer's Life by Dealer Type.

Panel (a) shows a Nissan-affiliated dealer and panel (b) an independent dealer; an other-brand-affiliated dealer is the same as in panel (a) with brand U in place of brand A. The labels on the arrows from entry give an entrant's endowment and its probability. A dealer is young to the left of the dashed line and old to the right of it.

5.2 Retail Stage

I solve a dealer's problem backward and begin with the retail stage. Consider a young dealer i that holds a car of brand b and faces demand ξ_{ib}. Its household obtains utility

$$u_i = \xi_{ib} - p + \epsilon_i \tag{5}$$

from buying the car at price p and utility 0 from not buying it. The taste shock ϵ_i is logistic with scale σ_ϵ and is drawn after the dealer posts its price, so the dealer does not observe it. The household buys if $u_i \geq 0$, so the car sells at price p with probability $\exp((\xi_{ib} - p)/\sigma_\epsilon)/[1 + \exp((\xi_{ib} - p)/\sigma_\epsilon)]$. The sale probability decreases in the price, and the coefficient on the price is $1/\sigma_\epsilon$, the inverse of the household demand scale. If the car does not sell, the dealer sells it at wholesale in the next period, which it discounts by the factor $\delta \in (0, 1)$, so it values an unsold car at its liquidation value δw_b, the discounted wholesale price. The dealer chooses the price that maximizes its expected profit, and the value of holding the car is

$$V(\xi_{ib}, w_b) = \max_p \left\{ \underbrace{\frac{\exp\big((\xi_{ib} - p)/\sigma_\epsilon\big)}{1 + \exp\big((\xi_{ib} - p)/\sigma_\epsilon\big)}}_{\Pr(\text{sale})} p + \underbrace{\frac{1}{1 + \exp\big((\xi_{ib} - p)/\sigma_\epsilon\big)}}_{\Pr(\text{no sale})} \delta w_b \right\}. \tag{6}$$

Lemma 1 (Retail pricing)**.** *Problem* (6) *has a unique solution* p^**, which satisfies*

$$p^* = \delta w_b + \sigma_\epsilon \left[1 + \exp\left(\frac{\xi_{ib} - p^*}{\sigma_\epsilon} \right) \right]. \tag{7}$$

Denote this solution by $p(\xi_{ib}, w_b)$ *and the probability that the car sells at this price by*

$$q(\xi_{ib}, w_b) = \frac{\exp\big((\xi_{ib} - p(\xi_{ib}, w_b))/\sigma_\epsilon\big)}{1 + \exp\big((\xi_{ib} - p(\xi_{ib}, w_b))/\sigma_\epsilon\big)}. \tag{8}$$

The value of holding the car is $V(\xi_{ib}, w_b) = p(\xi_{ib}, w_b) - \sigma_\epsilon$*. The value increases in demand and in the wholesale price, with*

$$\frac{\partial V}{\partial \xi_{ib}} = q(\xi_{ib}, w_b), \qquad \frac{\partial V}{\partial w_b} = \delta\big[1 - q(\xi_{ib}, w_b)\big], \tag{9}$$

and it ranges over $(\delta w_b, \infty)$*, so for every* $v > \delta w_b$ *there is a unique demand* $V^{-1}(v; w_b)$ *at which the value of holding the car equals* v*.*

Proof. See Appendix F.1.

By (7), the dealer prices above the liquidation value, and by more when the car is more likely to sell. The price $p(\xi_{ib}, w_b)$ and the sale probability $q(\xi_{ib}, w_b)$ both increase in demand, whereas a higher wholesale price raises the price and lowers the sale probability.[25]

5.3 Wholesale Stage: Inventory Choice

I now return to the wholesale stage. By Lemma 1, a young dealer that holds a car of brand b values it at $V(\xi_{ib}, w_b)$, whether it kept or bought the car. Its options depend on its dealer type and on

[25] See Appendix F.1 for the proof.

whether it entered with a car, so I derive the inventory choice of affiliated and independent dealers in turn.

5.3.1 Affiliated Dealers

Without loss of generality, consider a Nissan-affiliated entrant, which can hold only brand A. Its options at the wholesale stage and their payoffs depend on whether it is endowed:

$$\begin{aligned} &\text{endowed with a car:} && \max\Big\{\underbrace{V(\xi_{iA}, w_A)}_{\text{keep}},\ \underbrace{w_A}_{\text{sell and exit}}\Big\}, \\ &\text{without a car:} && \max\Big\{\underbrace{V(\xi_{iA}, w_A) - w_A}_{\text{buy}},\ \underbrace{0}_{\text{exit}}\Big\}. \end{aligned}$$

An endowed entrant keeps its car if and only if $V(\xi_{iA}, w_A) \geq w_A$, and an entrant without a car buys one if and only if $V(\xi_{iA}, w_A) - w_A \geq 0$. The two conditions coincide: given its demand, whether an entrant holds a car does not depend on its endowment, because an endowed entrant can sell its car at the same wholesale price at which it would buy one. The endowment determines only whether the dealer trades at wholesale.

Given the wholesale price $w_A > 0$, the value of holding a car increases in demand (Lemma 1), so an entrant holds a car if and only if its demand is at least the cutoff at which the value of holding the car equals the wholesale price,

$$\xi_{iA} \geq \bar{\xi}_A(w_A), \qquad \text{where} \qquad \bar{\xi}_A(w_A) = V^{-1}(w_A; w_A), \tag{10}$$

which exists because w_A lies in the range $(\delta w_A, \infty)$ of the value.[26] Define the inventory policy $j_N(\xi_i; \boldsymbol{w}) \in \{A, \emptyset\}$ as the car that a Nissan-affiliated entrant with demand ξ_i holds, where $\emptyset$ means no car:

$$j_N(\xi_i; \boldsymbol{w}) = \begin{cases} A & \text{if } \xi_{iA} \geq \bar{\xi}_A(w_A) \quad \text{(holds a car)}, \\ \emptyset & \text{if } \xi_{iA} < \bar{\xi}_A(w_A) \quad \text{(exits without a car)}. \end{cases} \tag{11}$$

Differentiating $V(\bar{\xi}_A(w_A), w_A) = w_A$ with (9) gives $\mathrm{d}\bar{\xi}_A / \mathrm{d}w_A = [1 - \delta(1-q)]/q > 0$, where $q = q(\bar{\xi}_A(w_A), w_A)$ is the sale probability at the cutoff. A higher wholesale price therefore raises the cutoff, so fewer dealers hold the brand: the price that an entrant pays for a car, or forgoes by keeping one, rises one for one, whereas the value of holding the car rises only at the rate $\delta(1-q)$. Let

$$h_{NA}(\boldsymbol{w}) = 1 - G_N\big(\bar{\xi}_A(w_A)\big)$$

denote the holding share of Nissan-affiliated dealers, the share of their entrants whose demand is above the cutoff. Because endowments are assigned independently of demand, the holding share is the same among entrants with and without an endowment. The mass of entrants that keep their

[26] Appendix F.2 gives the cutoff in closed form.

endowment, the mass that sell it and exit, and the mass that buy a car at wholesale are therefore

$$\underbrace{\gamma_{NA} D_N \, h_{NA}(\boldsymbol{w})}_{\text{keep the endowment}}, \qquad \underbrace{\gamma_{NA} D_N \big[1 - h_{NA}(\boldsymbol{w})\big]}_{\text{sell it and exit}}, \qquad \underbrace{(1-\gamma_{NA}) D_N \, h_{NA}(\boldsymbol{w})}_{\text{buy a car}}, \tag{12}$$

and the remaining $(1-\gamma_{NA})D_N[1-h_{NA}(\boldsymbol{w})]$ entrants exit without a car. The same applies to other-brand-affiliated entrants and brand U, with the cutoff $\bar{\xi}_U(w_U) = V^{-1}(w_U; w_U)$, the policy $j_B(\xi_i; \boldsymbol{w}) \in \{U, \emptyset\}$, and the holding share $h_{BU}(\boldsymbol{w}) = 1 - G_B\big(\bar{\xi}_U(w_U)\big)$.

5.3.2 Independent Dealers

An independent entrant's options at the wholesale stage and their payoffs are, with b' denoting the brand other than b,

$$\text{endowed with a car of brand } b\text{:} \quad \max\Big\{ \underbrace{V(\xi_{ib}, w_b)}_{\text{keep}}, \ \underbrace{w_b}_{\text{sell and exit}}, \ \underbrace{w_b - w_{b'} + V(\xi_{ib'}, w_{b'})}_{\text{sell and buy } b'} \Big\},$$

$$\text{without a car:} \qquad \max\Big\{ \underbrace{V(\xi_{iA}, w_A) - w_A}_{\text{buy } A}, \ \underbrace{V(\xi_{iU}, w_U) - w_U}_{\text{buy } U}, \ \underbrace{0}_{\text{exit}} \Big\}.$$

Each payoff of an entrant endowed with a car of brand b is w_b plus the payoff of an entrant without a car that ends up holding the same car, a car of brand b, a car of brand b', or none, so entrants with the same demand ξ_i make the same choice whether or not they are endowed. Given the wholesale prices, an independent entrant holds the brand with the larger net value, $V(\xi_{ib}, w_b) - w_b$, if that value is nonnegative and no car otherwise:

$$j_I(\xi_i; \boldsymbol{w}) = \begin{cases} A & \text{if } V(\xi_{iA}, w_A) - w_A \geq \max\big\{V(\xi_{iU}, w_U) - w_U, \, 0\big\}, \\ U & \text{if } V(\xi_{iU}, w_U) - w_U > \max\big\{V(\xi_{iA}, w_A) - w_A, \, 0\big\}, \\ \emptyset & \text{otherwise.} \end{cases} \tag{13}$$

This policy is a cutoff rule in each demand. An entrant whose demands are both below the cutoffs of (10), $\xi_{iA} < \bar{\xi}_A(w_A)$ and $\xi_{iU} < \bar{\xi}_U(w_U)$, holds no car. Among the other entrants, those whose demand for brand A is above a boundary that depends not only on w_A but also on ξ_{iU} and w_U hold brand A, and the rest hold brand U.[27] Let

$$h_{Ib}(\boldsymbol{w}) = \int \mathbf{1}\{j_I(\xi_i; \boldsymbol{w}) = b\} \, \mathrm{d}G_I(\xi_i), \qquad b \in \{A, U\},$$

denote the holding share of brand b, the share of independent entrants whose demands lie in the region in which they hold brand b; the share that holds no car is $1 - h_{IA}(\boldsymbol{w}) - h_{IU}(\boldsymbol{w}) = G_I(\bar{\xi}_A(w_A), \bar{\xi}_U(w_U))$, the joint distribution function at the two cutoffs. Because endowments are independent of demand, the mass of entrants endowed with a car of brand b that keep it, the mass that sell it and exit, the mass that sell it and buy a car of brand b', and the mass of entrants

[27] Appendix F.2 gives the boundary in closed form.

without a car that buy a car of brand b are

$$\underbrace{\gamma_{Ib} D_I \, h_{Ib}(\boldsymbol{w})}_{\text{keep the endowment}}, \qquad \underbrace{\gamma_{Ib} D_I \big[1 - h_{IA}(\boldsymbol{w}) - h_{IU}(\boldsymbol{w})\big]}_{\text{sell it and exit}}, \qquad \underbrace{\gamma_{Ib} D_I \, h_{Ib'}(\boldsymbol{w})}_{\text{sell it and buy brand } b'}, \\ \underbrace{(1 - \gamma_{IA} - \gamma_{IU}) D_I \, h_{Ib}(\boldsymbol{w})}_{\text{buy a car of brand } b}. \tag{14}$$

5.4 Stationary Equilibrium

I now aggregate the inventory choices over dealers to obtain wholesale demand and supply. Without loss of generality, consider brand A, which Nissan-affiliated and independent dealers hold; the derivation for brand U is the same with other-brand-affiliated dealers in place of Nissan-affiliated dealers. Each young dealer that holds a car of brand b sells it to its household with the probability $q(\xi_{ib}, w_b)$ of Lemma 1. Let

$$Q_{db}(\boldsymbol{w}) = \mathbb{E}\left[q(\xi_{ib}, w_b) \,\middle|\, j_d(\xi_i; \boldsymbol{w}) = b\right]$$

denote the mean sale probability of the young dealers of dealer type d that hold a car of brand b, with the expectation over the demand $\xi_i \sim G_d$. A mass $N_A(\boldsymbol{w}) = D_N h_{NA}(\boldsymbol{w}) + D_I h_{IA}(\boldsymbol{w})$ of young dealers end up holding a car of brand A at the wholesale stage, and their mean sale probability $Q_A(\boldsymbol{w})$ satisfies $N_A(\boldsymbol{w}) Q_A(\boldsymbol{w}) = D_N h_{NA}(\boldsymbol{w}) Q_{NA}(\boldsymbol{w}) + D_I h_{IA}(\boldsymbol{w}) Q_{IA}(\boldsymbol{w})$. Of these dealers, $N_A(\boldsymbol{w}) Q_A(\boldsymbol{w})$ sell the car to their household and exit, and the remaining $N_A(\boldsymbol{w})[1 - Q_A(\boldsymbol{w})]$ carry it into the next period and sell it at wholesale as old dealers.

Wholesale demand for brand A consists of the cars that young dealers buy. By (12) and (14), these are bought by the Nissan-affiliated entrants without a car whose demand is above the cutoff, by the independent entrants without a car for which brand A has the larger net value, and by the independent entrants endowed with a car of brand U that sell it and buy a car of brand A:

$$\begin{aligned} \mathcal{D}_A(\boldsymbol{w}) &= \underbrace{(1 - \gamma_{NA}) D_N \, h_{NA}(\boldsymbol{w})}_{\text{Nissan-affiliated, buy}} + \underbrace{(1 - \gamma_{IA} - \gamma_{IU}) D_I \, h_{IA}(\boldsymbol{w})}_{\text{independent, buy}} + \underbrace{\gamma_{IU} D_I \, h_{IA}(\boldsymbol{w})}_{\text{independent, sell } U \text{ and buy } A} \\ &= (1 - \gamma_{NA}) D_N \, h_{NA}(\boldsymbol{w}) + (1 - \gamma_{IA}) D_I \, h_{IA}(\boldsymbol{w}). \end{aligned} \tag{15}$$

Wholesale supply of brand A consists of the endowments that young dealers sell, by the Nissan-affiliated entrants endowed with a car whose demand is below the cutoff and by the independent entrants endowed with a car of brand A that sell it, whether they exit or buy a car of brand U, plus the cars of brand A that old dealers sell at wholesale. In the stationary state, the young dealers of the previous period faced the same wholesale prices, so the old dealers sell $N_A(\boldsymbol{w})[1 - Q_A(\boldsymbol{w})]$ cars:

$$\begin{aligned} \mathcal{S}_A(\boldsymbol{w}) = &\underbrace{\gamma_{NA} D_N \big[1 - h_{NA}(\boldsymbol{w})\big]}_{\text{Nissan-affiliated, sell}} + \underbrace{\gamma_{IA} D_I \big[1 - h_{IA}(\boldsymbol{w}) - h_{IU}(\boldsymbol{w})\big]}_{\text{independent, sell and exit}} \\ &+ \underbrace{\gamma_{IA} D_I \, h_{IU}(\boldsymbol{w})}_{\text{independent, sell } A \text{ and buy } U} + \underbrace{N_A(\boldsymbol{w}) \big[1 - Q_A(\boldsymbol{w})\big]}_{\text{old dealers}} \\ = &\, \gamma_{NA} D_N \big[1 - h_{NA}(\boldsymbol{w})\big] + \gamma_{IA} D_I \big[1 - h_{IA}(\boldsymbol{w})\big] + N_A(\boldsymbol{w}) \big[1 - Q_A(\boldsymbol{w})\big]. \end{aligned} \tag{16}$$

The wholesale market clears for brand A when $\mathcal{D}_A(\boldsymbol{w}) = \mathcal{S}_A(\boldsymbol{w})$. The endowment rates drop out of this condition: every young dealer that holds a car of brand A either bought it or kept its endowment, and every endowment is either kept or sold, so the cars bought net of the endowments sold equal the cars held net of all endowments, which add up to the inflow by (4):

$$\begin{aligned}\mathcal{D}_A(\boldsymbol{w}) - \mathcal{S}_A(\boldsymbol{w}) &= \underbrace{D_N h_{NA}(\boldsymbol{w}) + D_I h_{IA}(\boldsymbol{w})}_{N_A(\boldsymbol{w})} - \underbrace{(\gamma_{NA} D_N + \gamma_{IA} D_I)}_{E_A} - N_A(\boldsymbol{w})\big[1 - Q_A(\boldsymbol{w})\big] \\ &= N_A(\boldsymbol{w})\, Q_A(\boldsymbol{w}) - E_A.\end{aligned}$$

The market therefore clears if and only if household sales of brand A absorb exactly its inflow:

$$N_A(\boldsymbol{w})\, Q_A(\boldsymbol{w}) = E_A. \tag{17}$$

Because old dealers sell at wholesale every car that their household did not buy, the young dealers can hold every car present, the inflow and the cars that old dealers sell, only if households buy as many cars as arrive. The same derivation for brand U gives $N_U(\boldsymbol{w})Q_U(\boldsymbol{w}) = E_U$.

Definition 1 (Stationary equilibrium)**.** *A stationary equilibrium is a pair of wholesale prices* $\boldsymbol{w} = (w_A, w_U)$ *such that, when every dealer prices according to Lemma 1 and chooses its inventory according to* (11) *and* (13)*, the wholesale market clears for each brand:* $N_b(\boldsymbol{w})Q_b(\boldsymbol{w}) = E_b$ *for* $b \in \{A, U\}$.

Because the endowment rates drop out of the clearing conditions, neither they nor the endowment shares affect the equilibrium wholesale and retail prices, inventory choices, and sale probabilities; they determine only the volume of wholesale trade in (15) and (16), and they matter under the counterfactual scenarios of Section 8, which restrict how endowed dealers trade.

Proposition 1 (Existence and uniqueness)**.** *Under Assumptions 1 and 2 in Appendix F.3, a stationary equilibrium exists, and both wholesale prices are positive. Under Assumption 3, it is unique.*

Proof. See Appendix F.3.

Assumption 1 requires the demand distributions to have positive densities, and Assumption 2 requires dealers to be numerous enough that, at zero wholesale prices, household sales of each brand would exceed its inflow. Household sales of a brand fall in its own wholesale price and rise in the other brand's, through the independent dealers that switch brands, so some pair of prices clears the market for both brands. Assumption 3 requires the own effect of each wholesale price on household sales to exceed its cross effect.

5.5 A Demand Shock: Comparing Stationary Equilibria

The scandal lowers the valuation that households attach to affected cars. After the disclosure, the household of a dealer of dealer type $d \in \{N, I\}$ obtains utility

$$u_i = \xi_{iA} - \lambda_d - p + \epsilon_i$$

from buying an affected car at price p, in place of (5), whereas its utility from an unaffected car does not change. The demand loss λ_d lowers the demand for affected cars of every dealer of dealer type d by the same amount, and it can differ between Nissan-affiliated and independent dealers. The loss is permanent, so in the long run the market moves from the pre-scandal stationary equilibrium to a post-scandal stationary equilibrium, which solves the clearing conditions (17) with the lowered demands.[28]

To show how the equilibrium responds to the loss, I solve the model numerically. I assume demand is normal: a Nissan-affiliated entrant draws its demand for affected cars, $\xi_{iA} \sim \mathcal{N}(\mu_{NA}, \sigma_A^2)$, an other-brand-affiliated entrant its demand for unaffected cars, $\xi_{iU} \sim \mathcal{N}(\mu_{BU}, \sigma_U^2)$, and an independent entrant its two demands jointly,

$$\begin{pmatrix} \xi_{iA} \\ \xi_{iU} \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} \mu_{IA} \\ \mu_{IU} \end{pmatrix}, \begin{pmatrix} \sigma_A^2 & \rho\,\sigma_A\sigma_U \\ \rho\,\sigma_A\sigma_U & \sigma_U^2 \end{pmatrix} \right). \tag{18}$$

Table 6 sets the parameters to round values near the baseline estimates of Section 7.1. Table 7 compares the pre-scandal equilibrium with three post-scandal equilibria: a common loss, $\lambda_N = \lambda_I = 50$ thousand JPY, and a loss of 10 thousand JPY at Nissan-affiliated dealers together with a loss of 50 or 100 thousand JPY at independent dealers.

Table 6: Parameter Values of the Numerical Illustration

Parameter	Symbol	Value
Panel A: Dealer Masses and Inflows		
Dealer mass, Nissan-affiliated	D_N	23,000
Dealer mass, independent	D_I	115,000
Dealer mass, other-brand-affiliated	D_B	90,000
Inflow, affected	E_A	1,700
Inflow, unaffected	E_U	11,300
Panel B: Demand and Discounting		
Mean demand, affected cars at Nissan-affiliated dealers	μ_{NA}	470
Mean demand, affected cars at independent dealers	μ_{IA}	30
Mean demand, unaffected cars at independent dealers	μ_{IU}	470
Mean demand, unaffected cars at other-brand-affiliated dealers	μ_{BU}	650
Demand standard deviation, affected cars	σ_A	380
Demand standard deviation, unaffected cars	σ_U	300
Demand correlation	ρ	0.4
Household demand scale	σ_ϵ	70
Discount factor	δ	0.98

Values used in Section 5.5, rounded from the baseline Stage-1 estimates of Table 18 in Appendix D; the model has no holding cost. Dealer masses and inflows are numbers of entering dealers and of cars per 3-month model period; the equilibrium depends on them only through their ratios. Mean demands, demand standard deviations, and the household demand scale are in thousand JPY.

With a common loss, the shock passes through to prices rather than quantities. The wholesale price of affected cars falls by 49.1 thousand JPY, almost as much as the loss, and the mean

[28] See Section 6.3 for the path between the two equilibria.

Table 7: Stationary Equilibria Before and After a Demand Shock

		Post-scandal, (λ_N, λ_I)		
Outcome	Pre-scandal	(50, 50)	(10, 50)	(10, 100)
Wholesale price				
Affected	951.8	902.7	930.6	919.1
Unaffected	944.3	944.3	944.4	944.6
Listing share (%)				
Nissan-affiliated / affected	11.27	11.41	11.93	12.59
Independent / affected	0.72	0.74	0.59	0.44
Independent / unaffected	6.45	6.45	6.46	6.48
Other-brand-affiliated / unaffected	18.18	18.18	18.17	18.15
Mean retail price				
Nissan-affiliated / affected	1,110.8	1,060.6	1,090.5	1,080.0
Independent / affected	1,080.7	1,030.6	1,057.6	1,044.0
Independent / unaffected	1,070.3	1,070.3	1,070.4	1,070.5
Other-brand-affiliated / unaffected	1,088.8	1,088.8	1,088.9	1,089.0
Retail price standard deviation				
Nissan-affiliated / affected	103.6	103.1	104.6	105.8
Independent / affected	72.0	71.5	70.2	68.1
Independent / unaffected	67.1	67.1	67.1	67.0
Other-brand-affiliated / unaffected	84.7	84.7	84.7	84.7
Sale probability (%)				
Nissan-affiliated / affected	51.05	50.39	51.00	51.08
Independent / affected	45.23	44.46	44.56	43.94
Independent / unaffected	44.68	44.68	44.66	44.64
Other-brand-affiliated / unaffected	48.81	48.81	48.81	48.81

Stationary equilibria of the model at the parameter values of Table 6, before the scandal and after demand losses λ_N and λ_I, in thousand JPY, at Nissan-affiliated and independent dealers. A cell is a dealer type and a brand. The listing share is the share of a dealer type's entrants that list a car of the brand, the holding share h_{db}. The mean and the standard deviation of the retail price, in thousand JPY, are taken over the dealers that list a car, and the sale probability is their mean sale probability Q_{db}.

retail price of affected cars falls by 50.2 thousand JPY at Nissan-affiliated dealers and by 50.1 at independent dealers. Because the wholesale price falls by about as much as households' valuation, the value of holding an affected car net of its price barely changes, so the listing shares change by less than 0.2 percentage points and the sale probabilities by less than 1 percentage point.

When the losses differ, the single wholesale price pools them. With $\lambda_N = 10$ and $\lambda_I = 50$, the wholesale price of affected cars falls by 21.2 thousand JPY, so the retail price of affected cars falls by 20.3 thousand JPY at Nissan-affiliated dealers, twice the loss of their households, and by 23.0 at independent dealers, less than half of theirs. With $\lambda_I = 100$, the two declines are 30.7 and 36.7 thousand JPY. Listings move toward the dealer type with the smaller loss: as λ_I rises from 50 to 100, the listing share of affected cars rises from 11.93 to 12.59 percent at Nissan-affiliated dealers and falls from 0.59 to 0.44 percent at independent dealers, against 11.27 and 0.72 percent before

the scandal.

The shock to affected cars barely spills over to unaffected cars. Under the common loss, the unaffected outcomes stay at their pre-scandal values. With $\lambda_N = 10$ and $\lambda_I = 100$, the wholesale price of unaffected cars rises by 0.3 thousand JPY, and the listing share of unaffected cars at independent dealers rises by 0.03 percentage points while that of affected cars falls by 0.29 points: few of the independent entrants that stop holding an affected car value unaffected cars enough to hold one instead.

6 Estimation and Identification

A model period lasts 3 months, and the model's four cells are the dealer–brand cells of the analysis sample of Section 3.3: affected Nissan cars at Nissan-affiliated dealers, affected and unaffected cars at independent dealers, and unaffected cars of their own brand at other-brand-affiliated dealers. Prices, their dispersion, and entry counts come from all first listings of the analysis sample, and consumer sales within 3 months and the covariance between prices and sales come from the new listings of the reporting sample of Section 3.3, weighted as in Section 4.3. Because cars of a brand are homogeneous in the model, the price moments use prices standardized to a common reference car: a hedonic regression removes the part of each retail and wholesale price that the car's age, mileage, repair history, model–grade–vintage, and month explain, and the predicted price of the reference car is added back.[29] The post-disclosure changes that the estimation targets use observed prices.

In the first stage, I estimate two specifications: the baseline model of Section 5, in which dealers of a type differ only in their demand, and an extension with a holding cost of unsold cars, because their comparison shows which features of the pre-scandal data require this second dimension of heterogeneity; the second stage and the counterfactual analysis use the extension. Section 6.1 introduces the holding cost, Section 6.2 specifies the distributions and sets four parameters outside the estimation, Section 6.3 describes the transition path after the disclosure, Section 6.4 describes the two stages, and Section 6.5 discusses identification.

6.1 Holding Cost of an Unsold Car

The model of Section 5 lets dealers differ only in demand, so a dealer with higher demand posts a higher price and sells with a higher probability (Section 5.2). The estimated model adds a second dimension of heterogeneity, which lets dealers with the same demand post different prices: a dealer whose car does not sell at retail holds it until the wholesale market of the next period and incurs a holding cost c_i while it does. The cost is drawn from a distribution H_d specific to the dealer type after the inventory choice, so a dealer learns its cost only once it holds a car. It is paid only if the car does not sell, and it accrues in the period of the failed sale, so it is not discounted, whereas the wholesale sale takes place in the next period. The cost is independent of demand and of the

[29] See Appendix D.1 for more details on the price standardization.

endowment, and H_d has support on $\mathbb{R}$: a negative cost is a benefit of keeping the car for one more period.

With the cost, the dealer values an unsold car at its liquidation value $\delta w_b - c_i$, and the value of holding a car of brand b becomes

$$V(\xi_{ib}, c_i; w_b) = \max_{p \geq 0} \left\{ \underbrace{\frac{\exp\big((\xi_{ib} - p)/\sigma_\epsilon\big)}{1 + \exp\big((\xi_{ib} - p)/\sigma_\epsilon\big)}}_{\Pr(\text{sale})} p + \underbrace{\frac{1}{1 + \exp\big((\xi_{ib} - p)/\sigma_\epsilon\big)}}_{\Pr(\text{no sale})} (\delta w_b - c_i) \right\}, \tag{19}$$

whose solution p^* satisfies

$$p^* = \delta w_b - c_i + \sigma_\epsilon \left[1 + \exp\left(\frac{\xi_{ib} - p^*}{\sigma_\epsilon} \right) \right], \qquad V(\xi_{ib}, c_i; w_b) = p^* - \sigma_\epsilon, \tag{20}$$

unless the cost is so high that the liquidation value is far below zero, in which case the dealer posts $p^* = 0$.[30] As in Lemma 1, denote the solution by $p(\xi_{ib}, c_i; w_b)$ and the probability that the car sells at this price, defined as in (8), by $q(\xi_{ib}, c_i; w_b)$. The cost enters the pricing problem only through the liquidation value: a dealer with a higher cost values an unsold car less, so it posts a lower price and sells with a higher probability. Across dealers that differ in cost, prices and sale probabilities therefore move in opposite directions, whereas across dealers that differ in demand they move together.

At the wholesale stage, a dealer does not know the cost that it will draw, so it values holding a car of brand b at the expected value

$$\overline{V}_d(\xi_{ib}; w_b) = \int V(\xi_{ib}, c; w_b) \, \mathrm{d}H_d(c), \tag{21}$$

which, like V, increases in demand and in the wholesale price. The inventory choice of Section 5.3 applies with $\overline{V}_d$ in place of V. An affiliated entrant holds a car if and only if $\overline{V}_d(\xi_{ib}; w_b) \geq w_b$, that is, if its demand is at least the cutoff $\bar{\xi}_{db}(w_b) = \overline{V}_d^{-1}(w_b; w_b)$, which now depends on the dealer type through H_d and exists when the mean holding cost is positive.[31] An independent entrant holds the brand with the larger net value $\overline{V}_I(\xi_{ib}; w_b) - w_b$ if that value is nonnegative and no car otherwise, and the endowment does not affect the choice, for the same reason as before. Because the cost is drawn after the choice, the cost distribution among the dealers that hold a car is H_d for either brand, whereas demand is selected as before.

The holding shares and wholesale demand and supply of Section 5.4 keep their form under these policies. Only the mean sale probabilities change, because each holder now sells with a probability that depends on its cost: with $\bar{q}_d(\xi_{ib}; w_b) = \int q(\xi_{ib}, c; w_b) \, \mathrm{d}H_d(c)$ the sale probability averaged over the cost,

$$Q_{db}(\boldsymbol{w}) = \mathbb{E}\left[\bar{q}_d(\xi_{ib}; w_b) \,\middle|\, j_d(\xi_i; \boldsymbol{w}) = b \right]. \tag{22}$$

The clearing condition (17) is unchanged, so Definition 1 applies with the pricing problem (19) and the policies above, and Proposition 1 holds under Assumptions 1–3 applied to the household sales

[30] See Appendix F.4 for more details on the pricing problem with a holding cost.

[31] See Appendix F.4 for more details on the existence of the cutoff.

$N_b(\boldsymbol{w})Q_b(\boldsymbol{w})$ with these mean sale probabilities.[32]

6.2 Econometric Specification

6.2.1 Distributions of Demand and Holding Cost

Demand follows the normal distributions (18) of Section 5.5. The mean demand μ_{db} is specific to each cell, so it lets the level of demand for a brand differ across the dealer types that hold it. The standard deviation σ_b is specific to each brand and common across dealer types, and the correlation ρ is set outside the estimation. In the extension, the holding cost is normal, $c_i \sim \mathcal{N}\big(\mu_g^c, (\sigma_g^c)^2\big)$, where the cost group is $g = \text{aff}$ for the two affiliated types, N and B, and $g = I$ for independent dealers, so H_d is the same for the two affiliated types. It is not truncated, so some dealers draw a negative cost. In the baseline, the cost is zero for every dealer. The demand losses λ_N and λ_I of Section 5.5 are not restricted in sign, and the demand for unaffected cars does not change.

6.2.2 Externally Set Parameters

Four parameters are set outside the estimation and held fixed throughout: the dealer masses D_d of the three dealer types, the inflow E_b of each brand, the correlation ρ of an independent dealer's demands for the two brands, and the discount factor δ.[33]

Dealer mass D_d. Each model dealer holds one car, so I count every new listing at a dealership as one entering dealer, whatever the car: a dealer's outside option in the model corresponds to listing a car outside the model, such as a regular car or a kei car of an older vintage. With A_{it} the new listings of dealership i in month t and K_{it} its new listings of kei cars, the dealer mass per 3-month period over the 17 months of August 2015 to December 2016 is

$$\widehat{D}_d = \frac{3}{17} \sum_{t=\text{Aug. 2015}}^{\text{Dec. 2016}} \sum_{i:\, d(i)=d} \mathbf{1}\left\{ \frac{\sum_s K_{is}}{\sum_s A_{is}} \geq 0.05 \right\} A_{it}, \tag{23}$$

which counts the dealerships whose kei cars account for at least 5 percent of their new listings.[34]

Inflow E_b. By the clearing condition (17), stationary household sales equal the inflow, so I measure the inflow as consumer sales per model period before the scandal. With $\widehat{N}_{db}$ the new listings per period of the entrants of August 2015 to January 2016 in cell (d, b), whose 3-month outcome windows close before the disclosure month, and $\widehat{q}_{db}$ their 3-month consumer-sale probability,

$$\widehat{E}_b = \sum_d \widehat{N}_{db}\, \widehat{q}_{db}.$$

[32] See Appendix F.3 for more details.
[33] See Panel A of Table 18 in Appendix D for their values.
[34] See Appendix D.2 for more details on the listings that the counts cover.

The inflow stays at this value in every period, before and after the scandal, so the scandal acts only through the demand losses.[35]

Demand correlation ρ. Because each model dealer holds one car, no moment pins ρ. I set it to the correlation between the residual prices of an affected and an unaffected car listed by the same independent dealership. A residual price has a dealership component, common to the dealership's cars of a brand, and a car-specific component, which is uncorrelated across cars. The dealership components of the two brands are highly correlated, 0.91, but they account for only 50 and 37 percent of the residual variance of an affected and an unaffected car, so the correlation between the residual prices of two such cars is $0.91 \times (0.50 \times 0.37)^{1/2} = 0.39$, and I set $\rho = 0.4$.[36]

Discount factor δ. I set $\delta = 0.98$ per 3-month model period, about 8 percent per year. The discount factor enters the model only through the liquidation value of an unsold car, $\delta w_b - c_i$, where it is hard to separate from the mean holding cost, so I fix it.

6.3 Computing the Transition Path

The stationary equilibria of Section 5.5 describe the market before the scandal and its long-run state after it. The estimates of Section 4, however, measure the changes in the first 8 months after the disclosure, during which the market moves from one stationary equilibrium to the other, so I solve the model along this transition path. I compute the path under three assumptions. First, the disclosure arrives in period 0 after the wholesale stage: the young dealers of period 0 chose their inventory and posted their prices before the disclosure, and their households learn of the scandal while the cars are on the lot. This surprise cohort therefore sells fewer affected cars than it would at prices set after the disclosure and leaves more unsold cars for period 1 than in the pre-scandal equilibrium. Its data counterpart, the April 2016 entrants, faced the disclosure for 72 of the 91 days of their 3-month window, so the period-0 cohort faces the loss for the same share of its retail stage. Second, from period 1 on, entering dealers know the losses and the whole path of wholesale prices. Third, the market reaches the post-scandal stationary equilibrium after $T = 8$ periods, 2 years.

Two links tie each period to its neighbors. Unsold cars move forward: the cars that a cohort does not sell are supplied at wholesale in the next period, so the wholesale supply of a period depends on the prices of the period before. Values move backward: an entrant values a car by the price at which it would sell the car at wholesale in the next period if its household does not buy it, so the wholesale demand of a period depends on the prices of the period after. Neither forward simulation, which needs the future prices, nor backward induction, which needs the past unsold cars, can therefore solve the path one period at a time. I solve the 16 wholesale prices of periods 1 to 8 jointly from the 16 clearing conditions, one for each brand and period, which pin

[35] See Appendix D.2 for more details on the measurement of the inflow.

[36] See Appendix D.2 for the sample, the hedonic regression, and the correction for sampling noise.

the path between the unsold cars of the surprise cohort at its start and the post-scandal stationary equilibrium at its end. The path is locally unique, and although I cannot show that it is globally unique, the solution method reaches the same path from a wide range of starting paths.[37]

6.4 Estimation Procedure

The parameters of the model fall into two groups. The parameters of the pre-scandal stationary equilibrium, θ, are the four mean demands, the two demand standard deviations, and the household demand scale, $\theta = (\mu_{NA}, \mu_{IA}, \mu_{IU}, \mu_{BU}, \sigma_A, \sigma_U, \sigma_\epsilon)$, seven parameters in the baseline; the extension adds the means and standard deviations of the two holding-cost distributions, μ^c_{aff}, σ^c_{aff}, μ^c_I, and σ^c_I, for 11 parameters. The scandal is described by the two demand losses, $\lambda = (\lambda_N, \lambda_I)$. Stage 1 fits θ to 22 moments of the pre-scandal data. Stage 2 holds $\widehat{\theta}$ fixed and fits λ, on the transition path of Section 6.3, to twelve changes for affected cars relative to unaffected cars after the disclosure, estimated in Section 4. In both stages, the wholesale prices are solved again from the clearing conditions at every candidate parameter value, so every candidate is evaluated at an equilibrium. The endowment shares enter neither stage, because the equilibrium does not depend on them (Section 5.4). Table 8 lists the moments of Stage 1 and the targets of Stage 2.

Table 8: Moments of Stage 1 and Targets of Stage 2

Moments	Number	Data
Stage 1: moments of the pre-scandal equilibrium		
Log listing rates by dealer type	3	All first listings
Affected share of independent dealers' listings	1	All first listings
Initial price means	4	All first listings, standardized prices
Initial price standard deviations	4	All first listings, standardized prices
3-month consumer-sale probabilities	4	Dealers that report exit destinations, weighted
Initial price–consumer-sale covariances	4	Dealers that report exit destinations, weighted, standardized prices
Wholesale prices	2	Sold auction lots, standardized prices
Total	22	
Stage 2: changes for affected cars relative to unaffected cars after the disclosure		
Retail listing prices, $d \in \{N, I\}$, short and medium run	4	Equation (3), Table 5
3-month consumer sales, $d \in \{N, I\}$, very short, short, and medium run	6	Equation (3), Table 5
Wholesale price, short and medium run	2	Equation (2), Table 4
Total	12	

[37] See Appendix E for more details on the equilibrium conditions along the path, on how they are solved, and on the uniqueness of the path.

6.4.1 Stage 1: The Pre-Scandal Equilibrium

The retail moments of Stage 1 use the entrants of August 2015 to January 2016, and the two wholesale prices are the mean standardized prices of the lots sold in the auction months of August 2015 to January 2016.[38] Model moments average over the dealers that hold a car, so they condition on the inventory choice: the model's mean price in cell (d, b) is the mean of $p(\xi_{ib}, w_b)$ over the entrants of dealer type d that hold brand b, and the price standard deviation, the sale probability $Q_{db}(\boldsymbol{w})$, and the price–sale covariance are defined in the same way, the covariance with the sale probability $q(\xi_{ib}, w_b)$ in place of the consumer-sale indicator. With the holding cost, prices and sale probabilities are also averaged over the cost. The log listing rate of dealer type d is $\log \sum_{b \in \mathcal{J}_d} h_{db}(\boldsymbol{w})$, and the affected share of independent dealers' listings is $h_{IA}(\boldsymbol{w})/[h_{IA}(\boldsymbol{w}) + h_{IU}(\boldsymbol{w})]$. At any candidate θ, the pre-scandal wholesale prices $\boldsymbol{w}(\theta)$ solve the clearing conditions exactly, $N_b(\boldsymbol{w}(\theta); \theta)\, Q_b(\boldsymbol{w}(\theta); \theta) = \widehat{E}_b$ for both brands, and Stage 1 minimizes

$$J(\theta) = \frac{1}{22} \sum_{\ell=1}^{22} \left[\frac{m_\ell\big(\theta, \boldsymbol{w}(\theta)\big) - \widehat{m}_\ell}{s_\ell} \right]^2 ,$$

where $\widehat{m}_\ell$ is an empirical moment, m_ℓ its model counterpart, and s_ℓ its sampling standard error. The criterion weights the residuals by $W = \operatorname{diag}(1/s_\ell^2)$, which is diagonal although the moments are correlated, so its value measures the fit and is not a chi-squared statistic.[39]

6.4.2 Stage 2: Post-Scandal Demand Losses

The twelve targets of Stage 2 are the aggregated estimates of Section 4 listed in Table 8, each the change for affected cars relative to unaffected cars: the retail listing price at the two dealer types in the short and the medium run, the 3-month consumer-sale probability at the two dealer types in the very short, the short, and the medium run, and the wholesale price in the short and the medium run.[40] The very-short-run sale targets are the regression coefficients before the exposure adjustment of Section 4.3, because the model's period-0 cohort faces the loss for the same share of its retail stage. Each target is matched to the model at the periods that its data window covers: the very short run is the surprise cohort of period 0; the short run is period 1, whose wholesale counterpart, the auction months May to July 2016, covers exactly the first 3 months after the disclosure; and the medium run, August to December 2016, combines periods 2 and 3 with month weights of 3 and 2.

The model counterpart of each retail target is the same kind of regression coefficient, computed

[38] In the cells (N, A), (I, A), (I, U), and (B, U), the price moments rest on 5,345, 1,307, 12,505, and 34,466 first listings and the sale moments on 3,557, 1,280, 12,094, and 14,909 new listings of the dealers that report exit destinations.

[39] See Appendix D.4 for more details on the sampling standard errors and on the standard errors of the estimates.

[40] The wholesale targets control for vehicle age with a dummy for each year of age rather than the 0/1/2+ bins of Table 4, which moves the short-run coefficient from −57.26 to −57.28 thousand JPY and the medium-run coefficient from −33.56 to −33.57; the retail targets equal the estimates of Table 5. The sampling standard errors of all twelve targets come from the multiplier draws of Appendix D.4, so they differ from the clustered standard errors of the two tables.

from model outcomes. For a candidate λ, I stack the four retail cells of the pre-scandal equilibrium, $s = \text{pre}$, and the four cells of the horizon, $s = h$, and regress the cell means $\overline{y}_{db,s}$, either the price or the sale probability, on cell fixed effects, a common post-period indicator, and interactions of the affected brand with the two dealer types and the post period,

$$\overline{y}_{db,s} = \alpha_{db} + \tau \,\text{post}_s + \sum_{d' \in \{N,I\}} \Delta m_{d'} \,\mathbf{1}\{d = d',\, b = A\}\,\text{post}_s + e_{db,s}, \qquad s \in \{\text{pre}, h\}, \tag{24}$$

with each cell weighted by its listings. As in equation (3), the common post-period coefficient, which the two unaffected cells identify, makes each Δm_d a comparison with unaffected cars. The wholesale counterparts are the changes in the spread between the wholesale prices of affected and unaffected cars, in period 1 for the short run and averaged over periods 2 and 3 with the same month weights for the medium run. The four price coefficients, the six sale coefficients, and the two wholesale changes form the twelve model coefficients Δm_ℓ. Holding $\widehat{\theta}$ and the externally set parameters fixed, Stage 2 solves

$$\widehat{\lambda} = \arg\min_{\lambda} \;\frac{1}{12}\sum_{\ell=1}^{12}\left[\frac{\Delta m_\ell(\widehat{\theta}, \lambda) - \widehat{\beta}_\ell}{s_\ell}\right]^2, \tag{25}$$

where $\widehat{\beta}_\ell$ is a target and s_ℓ its sampling standard error, re-solving the post-scandal stationary equilibrium and the whole transition path at every candidate.[41]

6.5 Identification

In Stage 1, the selection and pricing structure of the model ties each block of moments to a group of parameters. A type-d entrant holds brand b only if its demand exceeds the cutoff $\bar{\xi}_{db}$, which equals $\bar{\xi}_b$ in the baseline, so the three listing rates and the affected share of independent dealers' listings pin the standardized cutoffs $(\bar{\xi}_{db} - \mu_{db})/\sigma_b$ of the four cells. Given the selection, the levels of prices and sale probabilities pin the mean demands and the household demand scale through the pricing condition (20), $p^* = \delta w_b - c_i + \sigma_\epsilon[1 + \exp((\xi_{ib} - p^*)/\sigma_\epsilon)]$, under which the markup over the liquidation value rises with the sale probability at a rate set by σ_ϵ. The dispersion of prices within a cell pins the demand standard deviations, and the price–sale covariance pins the holding cost. Without the holding cost, the covariance is nonnegative in every cell, because a dealer with higher demand posts a higher price and sells with a higher probability (Lemma 1). In the data, by contrast, the covariance is negative in all four cells (Table 9), so the baseline cannot fit it, whereas the holding cost, which moves prices and sale probabilities in opposite directions, lets the second specification do so.

In Stage 2, the two losses are identified from how the twelve coefficients respond to them. Because both dealer types trade affected cars in one wholesale market, a loss at either type lowers the wholesale price of affected cars and, through it, the retail prices at both types (Section 5.5), so the four retail price coefficients and the two wholesale coefficients pin down mainly a weighted sum of the two losses. The difference between the losses is pinned down by the gap between the retail

[41] See Appendix D.3 for more details on how the equilibria and the estimates are computed.

declines at the two dealer types and by the very-short-run sale changes, each of which responds only to the loss of its own dealer type, because the period-0 cohort faces the loss at prices posted before the disclosure. From period 1 on, dealers reprice, and the sale coefficients barely respond to either loss.

7 Estimation Results

This section reports the parameter estimates of the two stages and then the fit of each stage to its moments.

7.1 Parameter Estimates

The Stage-1 parameters describe entering dealers before the inventory choice selects which of them hold a car.[42] In the baseline, the mean demands are 468.3 thousand JPY for affected cars at Nissan-affiliated dealers, 27.5 and 471.0 thousand JPY for affected and unaffected cars at independent dealers, and 649.3 thousand JPY for unaffected cars at other-brand-affiliated dealers. The standard deviation of demand is 377.5 thousand JPY for affected cars and 302.8 thousand JPY for unaffected cars, and the household demand scale is 68.1 thousand JPY. With the holding cost, the mean demands are 648.5, 356.2, 747.1, and 813.3 thousand JPY in the same order, the standard deviations of demand are 286.9 and 200.8 thousand JPY, and the household demand scale is 138.1 thousand JPY. The holding cost has a mean of 72.8 thousand JPY at affiliated dealers and 160.4 thousand JPY at independent dealers, with standard deviations of 125.2 and 221.0 thousand JPY, so that 28 and 23 percent of the costs are negative.

The demand losses, estimated with the holding cost, are 11.1 thousand JPY at Nissan-affiliated dealers, with a standard error of 7.1, and 55.2 thousand JPY at independent dealers, with a standard error of 20.7, so the households of independent dealers lowered their valuation of an affected car by five times as much as the households of Nissan-affiliated dealers, by about 5 percent of the mean listing price against 1 percent. The loss at independent dealers is estimated much less precisely, because only the gap between the retail declines at the two dealer types and the very-short-run sale changes separate the two losses (Section 6.5).[43] The three starts reach the same point, and the criterion, the mean squared standardized residual over the twelve targets, is 5.46; Section 7.2 shows which targets drive it.

[42] See Table 18 in Appendix D for all parameter values, estimates, and standard errors.

[43] At the estimates, each of the four retail price coefficients and the two wholesale coefficients falls by 0.75 to 0.87 thousand JPY per thousand JPY of the loss at Nissan-affiliated dealers and by 0.22 to 0.29 thousand JPY per thousand JPY of the loss at independent dealers, in nearly the same proportions for all six, and a loss of 1 thousand JPY lowers the very-short-run sale probability of its own dealer type by 0.13 percentage points at Nissan-affiliated dealers and by 0.12 percentage points at independent dealers and leaves the other type's unchanged. At the Stage-1 estimates, the Jacobian of the 22 standardized moments with respect to the parameters has full column rank in both specifications, seven in the baseline and 11 with the holding cost, and at the Stage-2 estimates, the Jacobian of the twelve standardized targets with respect to the two losses has rank 2, with a condition number of 12.3. Full local rank does not establish global identification.

7.2 Model Fit

Table 9 compares the 22 Stage-1 moments with the data. The criterion is 110.15 in the baseline and 8.05 with the holding cost. The baseline cannot fit the price–sale covariances: its covariances are positive in all four cells, between 10.5 and 17.2 thousand JPY, whereas the observed covariances lie between −10.8 and −5.1 thousand JPY, so that the residuals range from 11.63 to 23.48. It also understates the price standard deviations and the consumer-sale probabilities at independent dealers. With the holding cost, the model's covariances are negative in all four cells, between −13.3 and −1.0 thousand JPY, and the largest residuals are in the mean prices: the model places the mean prices at affiliated and independent dealers closer together than the data, overstating them at independent dealers, 1,123.3 and 1,118.6 thousand JPY against 1,083.5 and 1,090.6 thousand JPY (6.45 and 6.00), and understating them at the two affiliated types, 1,134.3 and 1,121.1 thousand JPY against 1,146.3 and 1,134.8 thousand JPY (−2.61 and −4.46).

Table 10 compares the twelve targets with the model coefficients at the estimates, horizon by horizon. The model reproduces the medium-run targets and the very-short-run sale changes but not the short-run pattern. In the medium run, the retail price declines are 22.6 and 25.4 thousand JPY at Nissan-affiliated and independent dealers against 15.7 and 22.0 in the data, and the wholesale price of affected cars relative to unaffected cars falls by 23.8 thousand JPY against 33.6. In the very short run, the period-0 cohort meets the loss at prices posted before the disclosure, and its sale probability falls by 1.45 percentage points at Nissan-affiliated dealers and by 6.86 at independent dealers, against 3.70 and 3.76 in the data, both within one standard error. In the short run, by contrast, the model's price declines are almost the same as its medium-run declines, so it understates the decline at independent dealers, 42.9 thousand JPY in the data, by 17 thousand JPY (3.65); its sale probabilities are unchanged from period 1 on, whereas in the data they rise by 5.5 and 7.7 percentage points; and its wholesale decline, 24.6 thousand JPY, is less than half of the 57.3 in the data (4.74), the largest residual of the twelve.

The model has almost no short run because its only state variable is the stock of unsold cars, and dealers' demand for inventory is elastic. The surprise raises the affected cars that old dealers sell at wholesale at the start of period 1 by 5.0 percent above their pre-scandal level, and young dealers absorb them at a wholesale price only 1.5 thousand JPY below its post-scandal stationary level, because many entrants have demands just below the cutoff.[44] A dealer that buys in period 1 prices against the wholesale price of period 2, which is already close to its post-scandal level, so the low period-1 price changes who holds affected cars but not the prices they post. With a discount factor near one, the marginal dealer's willingness to pay at wholesale moves one for one with the retail value of the car, so a permanent loss in households' valuation cannot widen the gap between the retail and the wholesale decline at any horizon. The short-run pattern thus points to a transitory reduction in dealers' willingness to pay for affected cars beyond the demand loss, such as a temporary premium for the risk of holding them or a wholesale market that did not clear in the weeks after the disclosure, which the model leaves out. Because the losses are permanent,

[44] See Appendix E.3 for the transition path at the estimates.

Table 9: Stage-1 Fit: Data and Model Moments

			Baseline		With holding cost	
Moment	Data	Sampling SE	Model	Residual	Model	Residual
Listing rates and shares						
Log listing rate: Nissan-affiliated	-2.170	0.030	-2.188	-0.58	-2.142	0.97
Log listing rate: Independent	-2.812	0.034	-2.652	4.74	-2.847	-1.01
Log listing rate: Other-brand-affiliated	-1.656	0.025	-1.733	-3.08	-1.688	-1.30
Affected share of independent dealers' listings (%)	9.5	0.4	9.9	1.04	9.6	0.27
Mean price (thousand JPY)						
Nissan-affiliated / affected	1,146.3	4.6	1,103.7	-9.26	1,134.3	-2.61
Independent / affected	1,083.5	6.2	1,073.2	-1.67	1,123.3	6.45
Independent / unaffected	1,090.6	4.7	1,075.3	-3.30	1,118.6	6.00
Other-brand-affiliated / unaffected	1,134.8	3.1	1,093.7	-13.45	1,121.1	-4.46
Price SD (thousand JPY)						
Nissan-affiliated / affected	96.0	1.8	103.9	4.31	94.1	-1.08
Independent / affected	117.7	3.8	72.1	-11.95	112.2	-1.44
Independent / unaffected	115.3	2.3	68.9	-20.52	109.4	-2.63
Other-brand-affiliated / unaffected	82.2	0.8	86.4	5.29	83.0	1.11
Consumer-sale probability (percent)						
Nissan-affiliated / affected	49.1	1.7	51.7	1.56	49.4	0.18
Independent / affected	63.3	1.8	45.8	-9.60	57.5	-3.20
Independent / unaffected	57.0	1.0	45.8	-11.18	56.6	-0.40
Other-brand-affiliated / unaffected	44.8	1.3	49.8	3.98	47.2	1.88
Price–sale covariance (thousand JPY)						
Nissan-affiliated / affected	-9.8	1.8	17.2	15.23	-1.0	4.98
Independent / affected	-10.8	1.9	11.0	11.63	-12.0	-0.65
Independent / unaffected	-9.8	0.9	10.5	23.48	-13.3	-4.03
Other-brand-affiliated / unaffected	-5.1	1.3	13.9	14.73	-3.9	0.96
Wholesale price (thousand JPY)						
Affected	931.5	2.4	946.1	5.98	928.7	-1.16
Unaffected	932.2	1.8	949.3	9.62	932.7	0.30

Data and model moments at the Stage-1 estimates of Panels B and C of Table 18, for the baseline specification and the specification with a holding cost. A cell is a dealer type's new listings of one brand. The log listing rate is the log of a dealer type's new listings per model period per unit of its dealer mass. Mean prices, price SDs, and price–sale covariances use initial listing prices, and wholesale prices use auction prices, both standardized to the reference car of Appendix D.1, in thousand JPY at 2016 prices. Consumer-sale probabilities, the share sold to a household within 3 months, and price–sale covariances use the new listings of the dealers that report exit destinations, weighted to the badge mix of all first listings. The sampling SE is the standard deviation of the data moment over 100 paired multiplier draws. The residual is the model value minus the data value, divided by the sampling SE. The criterion, the mean squared residual, is 110.15 in the baseline and 8.05 with the holding cost.

Table 10: Stage-2 Targets and Model Coefficients by Horizon

Target	Data	Sampling SE	Model	Standardized residual
Retail listing price (thousand JPY)				
Nissan-affiliated, short run	-20.9	3.4	-23.4	-0.74
Nissan-affiliated, medium run	-15.7	2.6	-22.6	-2.68
Independent, short run	-42.9	4.6	-26.2	3.65
Independent, medium run	-22.0	2.8	-25.4	-1.20
Consumer sale within 3 months (percentage points)				
Nissan-affiliated, very short run	-3.70	2.62	-1.45	0.86
Nissan-affiliated, short run	5.54	2.28	-0.05	-2.44
Nissan-affiliated, medium run	0.38	1.70	0.02	-0.21
Independent, very short run	-3.76	3.27	-6.86	-0.95
Independent, short run	7.67	2.47	-0.30	-3.22
Independent, medium run	-2.32	1.96	-0.25	1.05
Wholesale price, affected minus unaffected (thousand JPY)				
Short run	-57.3	6.9	-24.6	4.74
Medium run	-33.6	8.0	-23.8	1.22

Data: coefficients of the aggregated specification by dealer type for retail listing prices and consumer sales (Table 5) and of the pooled aggregated specification for wholesale prices (Table 4). The very short run, used only for consumer sales, is the cohort of cars first listed in April 2016, which were priced before the disclosure and exposed to it for 72 of the 91 days of their 3-month sale window; the short run is the auction months of May to July 2016 and the cars first listed in June and July 2016; the medium run is August to December 2016. The sampling SE is the standard deviation of the coefficient over 100 paired multiplier draws, clustered by dealership for retail outcomes and by model code and registration year at auction; it standardizes the residual in the Stage-2 criterion and differs from the analytic clustered standard error of those tables. Model: the same regression coefficients computed on the model's transition path at the estimates of Panel D of Table 18 (Section 6.4.2), with cell means weighted by listings and with the model's period-0 cohort, the counterpart of the April 2016 cohort, exposed to the demand loss for the same share of its retail stage. The standardized residual is the model value minus the data value, divided by the sampling SE. Prices are in thousand JPY at 2016 prices; consumer sales (a sale to a household within 3 months) are in percentage points.

the counterfactual analysis of Section 8 compares stationary equilibria, for which the medium-run targets are the closest data counterpart.

8 Counterfactual Analysis

This section asks how frictions in the wholesale market shape the pass-through of the scandal's demand shock to retail prices, sale probabilities, and wholesale prices. In the baseline, the wholesale market is integrated: every dealer can sell its endowment there at once and buy any car that its type can hold, at the same two prices. As a result, the endowment affects neither the inventory choice nor the retail price (Section 5.4), and the single wholesale price of affected cars pools the losses of the two dealer types (Section 6.5).

I compare this baseline with two counterfactual scenarios. In the first scenario, segmented wholesale markets, the wholesale market of each brand is divided by dealer type, so that dealers trade only with dealers of their own type. In the second scenario, mandatory retail of endowments, an entrant cannot sell its endowment at wholesale on entry. It must offer the car at retail first, as if a dealer that takes a car in trade had to sell it itself, and only a car that does not sell goes to the wholesale market in the next period, as every unsold car does.

The first scenario has a counterpart in the organization of wholesale auto auctions in Japan. Manufacturers operate their own auction networks, such as Nissan's NAA, encourage their affiliated dealers to sell and buy there, and restrict the participation of independent dealers. In practice, however, these networks are not closed: some independent dealers take part in manufacturer-operated venues, affiliated dealers also sell at the other venues, and the scandal's price response does not differ between Nissan's venues and the others (Appendix Table 16).[45] The first scenario thus corresponds to manufacturers that bar their affiliated dealers from other auctions and exclude independent dealers from their own venues entirely.

Unlike the baseline, both scenarios make the division of the inflow across dealer types matter. Because endowments are not observed, I set the two endowment shares, which give the part of each brand's inflow endowed to its affiliated network, to a common value, $\omega_A = \omega_U = \omega \in \{0.5, 0.7, 0.9\}$, which stays fixed before and after the scandal. Section 8.1 sets out how I compare each scenario with the baseline, and Sections 8.2 and 8.3 formalize the two scenarios and report their results.

8.1 Framework

Each scenario uses the Stage-1 estimates of the specification with the holding cost and its demand losses, together with the externally set parameters, without re-estimating them.[46] The comparison concerns the long run. For each scenario and endowment share, I solve the pre-scandal and the post-scandal stationary equilibria under the scenario's rules and compute the change from the one to the other for affected cars relative to unaffected cars. For retail prices and sale probabilities, the changes are the coefficients of regression (24) with the four post-scandal cells in place of the horizon cells, each cell weighted by its listings in its own equilibrium. For wholesale prices, the change is that of the spread $w_A - w_U$ of the single wholesale market of each brand or, under segmented wholesale markets, where a superscript marks the dealer type whose market sets the price, the changes in the spread across the markets of the two affiliated networks, $w_A^N - w_U^B$, which compares prices in two different markets, and in the spread within the independent market, $w_A^I - w_U^I$. The baseline values of the same changes are the estimated model's long-run prediction, the difference

[45] Matching retail listings to later auction sales by vehicle identifier shows that, of the cars that Nissan-affiliated dealers first listed between April 2015 and March 2017 and sold at auction after the listing ended, 41.0 percent of the affected cars and 45.0 percent of the unaffected cars were sold at NAA and the rest at other venues. On the buying side, of the affected cars that independent dealers bought at auction and then listed, 10.5 percent came from NAA. The match covers only the auction houses that record the identifier. In Appendix Table 16, the difference between the price responses at Nissan's venues and at the other venues is −0.23 thousand JPY in the medium run, with a standard error of 9.25, and 8.24 thousand JPY in the short run, with a standard error of 7.81.

[46] See Table 18 in Appendix D for their values.

between the pre-scandal equilibrium and the post-scandal stationary equilibrium that ends the transition path of Section 6.3. Figures 4 and 5 place the baseline and the scenario at the three endowment shares side by side for each outcome. Because the scenarios also change the pre-scandal equilibrium, I compare only these changes and not the levels of prices and sales, and because the data's last window, the medium run, is not a stationary state, the figures have no data counterpart.

8.2 Wholesale Markets Segmented by Dealer Type

Under segmented wholesale, there are three wholesale markets: Nissan-affiliated dealers trade affected cars among themselves at the price w_A^N, other-brand-affiliated dealers trade unaffected cars among themselves at w_U^B, and independent dealers trade both brands among themselves at w_A^I and w_U^I. Each entrant chooses as in Section 5.3 at the prices of its own market, and an unsold car returns to the market of its dealer's type. The inflow that enters each market as endowments must therefore be absorbed by the household sales of the same dealer type. By the clearing argument of Section 5.4 applied to each market, the four prices clear the markets when

$$D_N h_{NA} Q_{NA} = \omega_A E_A, \qquad D_B h_{BU} Q_{BU} = \omega_U E_U,$$
$$D_I h_{IA} Q_{IA} = (1-\omega_A) E_A, \qquad D_I h_{IU} Q_{IU} = (1-\omega_U) E_U.$$

Each affiliated market clears through a single equation in its own price, and the independent market is the model of Section 5.4 with $D_N = D_B = 0$ and the inflows $(1-\omega_A)E_A$ and $(1-\omega_U)E_U$, so Proposition 1 applies market by market. At the endowment shares implied by the baseline, $\omega_A = 0.79$ and $\omega_U = 0.70$, the shares of each brand's household sales that its affiliated dealers make in the baseline pre-scandal equilibrium, the pre-scandal equilibrium under segmentation coincides with the baseline, so the comparison at these shares reflects only how segmentation changes the pass-through of the shock; its results are within 0.1 thousand JPY of those at the three common shares.

Figure 4 compares the long-run changes under segmentation with the baseline at the three endowment shares. The share barely matters here: across the three shares, every price change varies by less than 0.1 thousand JPY, because each market clears its own inflow and the losses do not depend on the share, so the following numbers are those at $\omega = 0.7$. Because each market's price reflects only the loss of its own dealer type, the wholesale price of affected cars relative to unaffected cars falls by 10.8 thousand JPY across the two affiliated networks, less than half of the baseline decline of 23.1, and by 54.4 thousand JPY within the independent market, close to the loss of 55.2 at independent dealers. The retail price of affected cars accordingly falls by 10.9 thousand JPY at Nissan-affiliated dealers against 21.9 in the baseline and by 54.5 thousand JPY at independent dealers against 24.7, so the gap between the price declines at the two dealer types widens from 2.8 thousand JPY in the baseline to 43.6, close to the gap of 44.1 between the two losses, whereas the sale probabilities change by less than 0.2 percentage points at both dealer types, as in the baseline.

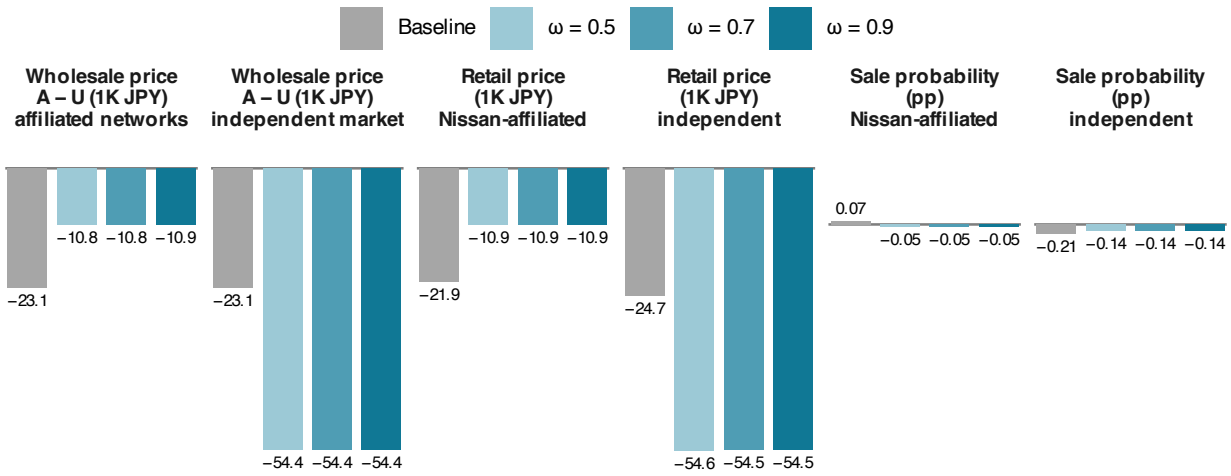


Figure 4: Long-Run Changes under Wholesale Markets Segmented by Dealer Type.
Each panel shows the change from the pre-scandal to the post-scandal stationary equilibrium for affected cars relative to unaffected cars, in the baseline and under segmented wholesale markets at the endowment shares $\omega = 0.5$, 0.7, and 0.9: the wholesale price spread between affected and unaffected cars in thousand JPY at 2016 prices, which under segmentation is the spread across the two affiliated networks, $w_A^N - w_U^B$, in the first panel and the spread within the independent market, $w_A^I - w_U^I$, in the second, both compared with the baseline spread $w_A - w_U$; the retail listing price at Nissan-affiliated and at independent dealers in thousand JPY; and the 3-month consumer-sale probability at the two dealer types in percentage points. The retail changes are the coefficients of regression (24) on the pre- and post-scandal cells. The baseline does not depend on the endowment shares.

8.3 Mandatory Retail of Endowments

Under mandatory retail, only the entrants without a car buy at wholesale on entry. Their masses are

$$D_{N\emptyset} = D_N - \omega_A E_A, \qquad D_{I\emptyset} = D_I - (1-\omega_A)E_A - (1-\omega_U)E_U, \qquad D_{B\emptyset} = D_B - \omega_U E_U,$$

which equal $(1-\gamma_{NA})D_N$, $(1-\gamma_{IA}-\gamma_{IU})D_I$, and $(1-\gamma_{BU})D_B$ by (4). These entrants choose by the policies of Section 5.3 at the wholesale prices of this scenario, whereas an endowed entrant lists its car whatever its demand. Listings and their mean sale probability $Q_{db}^{(2)}$ in cell (d,b) are therefore given by

$$N_{db}^{(2)} = \gamma_{db} D_d + D_{d\emptyset} h_{db}(\boldsymbol{w}),$$
$$N_{db}^{(2)} Q_{db}^{(2)} = \gamma_{db} D_d \, \mathbb{E}\left[\bar{q}_d(\xi_{ib}; w_b)\right] + D_{d\emptyset} h_{db}(\boldsymbol{w}) \, Q_{db}(\boldsymbol{w}),$$

where the superscript marks the scenario. The first expectation is over the whole demand distribution, because endowments are assigned independently of demand. Wholesale demand for brand b consists of the cars that entrants without a car buy, $\sum_d D_{d\emptyset} h_{db}(\boldsymbol{w})$, and wholesale supply consists only of the listings of the previous period that did not sell, $\sum_d N_{db}^{(2)}[1 - Q_{db}^{(2)}]$. Because the endowed

listings of brand b add up to its inflow by (4), equating demand and supply gives

$$\sum_d D_{d\emptyset} h_{db}(\boldsymbol{w}) = \sum_d N_{db}^{(2)} \left[1 - Q_{db}^{(2)}\right] \quad \Longleftrightarrow \quad \sum_d N_{db}^{(2)} Q_{db}^{(2)} = E_b,$$

the clearing condition of Section 5.4 with the endowed entrants' sales included. These household sales keep the properties of Lemma 2, because the endowed entrants' sales fall in w_b, vanish as w_b grows, and do not depend on the other brand's price, so Proposition 1 applies with its assumptions stated for them. The holding cost is drawn after the wholesale stage, as in the baseline.

Figure 5 compares the long-run changes under mandatory retail with the baseline at the three endowment shares. The wholesale spread falls by 22.5 to 22.9 thousand JPY against 23.1 in the baseline, so the wholesale price that both dealer types face changes almost as in the baseline, and the differences arise at retail from the composition of listings. At independent dealers, the retail price of affected cars falls by 38.7 to 42.8 thousand JPY against 24.7 in the baseline, and their sale probability falls by 3.7 to 4.8 percentage points against 0.2, whereas at Nissan-affiliated dealers the price falls by 18.9 to 19.8 thousand JPY against 21.9 and the sale probability rises by 0.7 to 0.8 percentage points against 0.1. An endowed entrant lists its car whatever its demand, so endowed cars carry lower prices and sell less often than the cars that entrants choose to buy, and the scandal changes the share of endowed cars among listings. At independent dealers, fewer entrants choose to hold an affected car after the scandal, so the share of endowed cars among their affected listings rises, from 48.3 to 56.8 percent at $\omega = 0.7$, and the mean price and sale probability of these listings fall by more than in the baseline. At Nissan-affiliated dealers, where the lower wholesale price leads more entrants to hold an affected car, the share falls, from 33.4 to 31.8 percent at $\omega = 0.7$, and the mean price falls by less. The endowment share matters in this scenario because it sets how many endowed cars enter the listings of each dealer type; at the shares implied by the baseline, the changes are within 0.3 thousand JPY and 0.1 percentage points of those at $\omega = 0.7$.

Read together, the two scenarios show what the integrated wholesale market does when the shock hits. It pools the losses of the two dealer types into one wholesale price, so that the retail decline at Nissan-affiliated dealers comes largely through a wholesale price that reflects the larger loss at independent dealers, and the decline at independent dealers stays far below their own households' loss. It also lets dealers pass unwanted cars on rather than discount them, so that the cars that households rarely buy do not weigh on retail prices and sale probabilities, which remain pinned by the inflow of trade-ins.

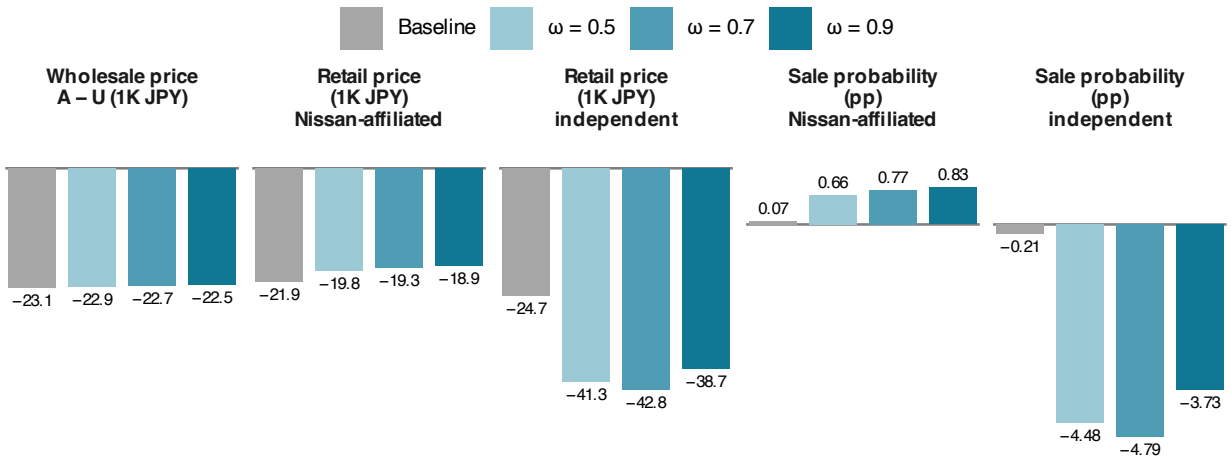


Figure 5: Long-Run Changes under Mandatory Retail of Endowments.
Each panel shows the change from the pre-scandal to the post-scandal stationary equilibrium for affected cars relative to unaffected cars, in the baseline and under mandatory retail of endowments at the endowment shares $\omega = 0.5$, 0.7, and 0.9: the wholesale price spread between affected and unaffected cars in thousand JPY at 2016 prices, the retail listing price at Nissan-affiliated and at independent dealers in thousand JPY, and the 3-month consumer-sale probability at the two dealer types in percentage points. The retail changes are the coefficients of regression (24) on the pre- and post-scandal cells. The baseline does not depend on the endowment shares.

9 Conclusion

This paper studies how used car dealers adjust prices and inventories when a product-harm crisis hits and what role the interdealer wholesale market plays in these responses, using the 2016 Mitsubishi fuel economy scandal as the shock. Wholesale auction records and retail listing data make both layers of the dealer market observable: the auctions where dealers trade with one another and the platform where they face households. Comparing the affected Nissan models with unaffected kei cars before and after the disclosure, I find that prices fell in both markets, at retail by less at Nissan-affiliated dealers than at independent dealers, whereas quantities moved only briefly: the unsold rate at auction rose for a few weeks, the affected cars priced before the disclosure sold less often, and once dealers had repriced, affected cars sold more often for a few months before sales returned to their previous level. The wholesale market thus absorbed the shock through prices rather than through quantities.

To account for the link between the two markets, I build an overlapping-generations model in which dealers trade cars with one another at wholesale, sell them to households as monopolists, and resell unsold cars at wholesale, and in which they differ in demand and in the cost of holding a car that does not sell. The wholesale price is both the cost of a car and the value of an unsold one, so the two markets are determined jointly, and in a stationary equilibrium household sales are pinned by the inflow of trade-ins, so a demand shock is absorbed by prices rather than by sales. Estimated on pre-scandal data and on the changes over the 8 months after the disclosure, matched along

the model's transition path, the model attributes the scandal to a loss in households' valuation of affected cars that is concentrated at independent dealers and reaches Nissan-affiliated dealers through the wholesale price that both dealer types share. It captures the medium-run changes but misses the short run, when the wholesale price fell by far more than it would in the medium run and sales briefly rose; a temporary fall in what dealers were willing to pay for affected cars, over and above the loss in households' valuation, would account for this pattern and is a natural extension of the model. Two counterfactual scenarios, one that segments the wholesale market by dealer type and one that requires dealers to offer their trade-ins at retail, show how the organization of the wholesale market shapes the pass-through: segmentation confines each dealer type's loss to its own market and pulls the retail price responses of the two dealer types apart, whereas mandatory retail of trade-ins changes the composition of listings and, through it, the price and sale responses at independent dealers.

References

Aizawa, Ayako. 2020. "Why fuel economy fraud happened in the Japanese automotive industry?" *Annals of Business Administrative Science*, 19(3): 111–125.

Astvansh, Vivek, Kersi Antia, and Gerard Tellis. 2025. "Product recall: a synthesis of multidisciplinary findings, and research directions." *Marketing Letters*, 36(1): 65–77.

Ater, Itai, and Nir S. Yoseph. 2022. "The Impact of Environmental Fraud on the Used Car Market: Evidence from Dieselgate*." *The Journal of Industrial Economics*, 70(2): 463–491.

Bachmann, Rüdiger, Gabriel Ehrlich, Ying Fan, Dimitrije Ruzic, and Benjamin Leard. 2023. "Firms and Collective Reputation: A Study of the Volkswagen Emissions Scandal." *Journal of the European Economic Association*, 21(2): 484–525.

Blanchard, Olivier Jean, and Charles M. Kahn. 1980. "The Solution of Linear Difference Models under Rational Expectations." *Econometrica*, 48(5): 1305–1311.

Che, Xiaogang, Hajime Katayama, and Peter Lee. 2023. "Product-Harm Crises and Spillover Effects: A Case Study of the Volkswagen Diesel Emissions Scandal in eBay Used Car Auction Markets." *Journal of Marketing Research*, 60(2): 409–424.

Cleeren, Kathleen, Marnik G. Dekimpe, and Harald J. Van Heerde. 2017. "Marketing research on product-harm crises: a review, managerial implications, and an agenda for future research." *Journal of the Academy of Marketing Science*, 45(5): 593–615.

Colliard, Jean-Edouard, Thierry Foucault, and Peter Hoffmann. 2021. "Inventory Management, Dealers' Connections, and Prices in Over-the-Counter Markets." *The Journal of Finance*, 76(5): 2199–2247.

Duffie, Darrell, Nicolae Gârleanu, and Lasse Heje Pedersen. 2005. "Over-the-Counter Markets." *Econometrica*, 73(6): 1815–1847.

Dunne, Peter G., Harald Hau, and Michael J. Moore. 2015. "Dealer Intermediation Between Markets." *Journal of the European Economic Association*, 13(5): 770–804.

Gavazza, Alessandro. 2011. "The Role of Trading Frictions in Real Asset Markets." *American Economic Review*, 101(4): 1106–1143.

Gavazza, Alessandro. 2016. "An Empirical Equilibrium Model of a Decentralized Asset Market." *Econometrica*, 84(5): 1755–1798.

Gavazza, Alessandro, Alessandro Lizzeri, and Nikita Roketskiy. 2014. "A Quantitative Analysis of the Used-Car Market." *American Economic Review*, 104(11): 3668–3700.

Gavazza, Alessandro, and Alessandro Lizzeri. 2021. "Frictions in product markets." In *Handbook of Industrial Organization*. Vol. 4, 433–484. Elsevier.

Genesove, David. 1993. "Adverse Selection in the Wholesale Used Car Market." *Journal of Political Economy*, 101(4): 644–665.

Genesove, David. 1995. "Search at Wholesale Auto Auctions*." *The Quarterly Journal of Economics*, 110(1): 23–49.

Huang, Guofang, Hong Luo, and Jing Xia. 2019. "Invest in Information or Wing It? A Model of Dynamic Pricing with Seller Learning." *Management Science*, 65(12): 5556–5583.

Hugonnier, Julien, Benjamin Lester, and Pierre-Olivier Weill. 2020. "Frictional Intermediation in Over-the-Counter Markets." *The Review of Economic Studies*, 87(3): 1432–1469.

Japan Automobile Manufacturers Association. 2016. "Survey on the Use of Kei Cars, Fiscal Year 2015." In Japanese.

Japan Automobile Manufacturers Association. 2018. "The Motor Industry of Japan 2018." Available at `https://www.jama.org/wp-content/uploads/2018/07/mij2018.pdf`.

Japan Used Car Dealers Association. 2021. "Used Car Sales toward Carbon Neutrality." In Japanese. Available at `https://www.mlit.go.jp/jidosha/content/001400613.pdf`.

Larsen, Bradley, and Anthony Lee Zhang. 2026. "Quantifying Bargaining Power Under Incomplete Information: A Supply-Side Analysis of the Used-Car Industry." *Working paper*.

Larsen, Bradley J. 2021. "The Efficiency of Real-World Bargaining: Evidence from Wholesale Used-Auto Auctions." *The Review of Economic Studies*, 88(2): 851–882.

Larsen, Bradley J, Carol Hengheng Lu, and Anthony Lee Zhang. 2026. "Intermediaries in Bargaining: Evidence from Business-to-Business Used-Car Inventory Negotiations." *Working paper*.

Li, Fei, Charles Murry, Can Tian, and Yiyi Zhou. 2024. "Inventory Management in Decentralized Markets." *International Economic Review*, 65(1): 431–470.

Li, Fei, Charles Murry, Can Tian, and Yiyi Zhou. 2025. "Intermediaries in Decentralized Markets: Evidence from Used-Car Transactions."

Liu, Yan, and Venkatesh Shankar. 2015. "The Dynamic Impact of Product-Harm Crises on Brand Preference and Advertising Effectiveness: An Empirical Analysis of the Automobile Industry." *Management Science*, 61(10): 2514–2535.

Ministry of Land, Infrastructure, Transport and Tourism. 2015. "Future Vision for the Utilization of Automobile-Related Information." In Japanese. Available at `https://www.mlit.go.jp/common/001066883.pdf`.

Mitsubishi Motors Corporation. 2016. "Investigation Report on Fuel Economy Misconduct." In Japanese. Available at `https://web.archive.org/web/20161019111746/http://www.mitsubishi-motors.com/content/dam/com/ir_jp/pdf/irnews/2016/20160802-02.pdf`.

Murry, Charles, and Yiyi Zhou. 2020. "Consumer Search and Automobile Dealer Colocation." *Management Science*, 66(5): 1909–1934.

Overby, Eric, and Chris Forman. 2015. "The Effect of Electronic Commerce on Geographic Purchasing Patterns and Price Dispersion." *Management Science*, 61(2): 431–453.

Recruit Carsensor Co., Ltd. 2018. "Carsensor Used Car Purchasing Behavior Survey 2017." In Japanese. Available at `https://www.recruit.co.jp/newsroom/recruit-mp/news/180516_01.pdf`.

Strittmatter, Anthony, and Michael Lechner. 2020. "Sorting in the used-car market after the Volkswagen emission scandal." *Journal of Environmental Economics and Management*, 101: 102305.

Subramanian, Hemang, and Eric Overby. 2017. "Electronic Commerce, Spatial Arbitrage, and Market Efficiency." *Information Systems Research*, 28(1): 97–116.

Tadelis, Steven, and Florian Zettelmeyer. 2015. "Information Disclosure as a Matching Mechanism: Theory and Evidence from a Field Experiment." *American Economic Review*, 105(2): 886–905.

Tanaka, Shinsuke. 2020. "When tax incentives drive illicit behavior: The manipulation of fuel economy in the automobile industry." *Journal of Environmental Economics and Management*, 104: 102367.

USS Co., Ltd. 2021. "Integrated Report 2021." In Japanese. Available at `https://www.ussnet.co.jp/ir/library/annual/pdf/2021/2021-00.pdf`.

USS Co., Ltd. 2022. "USS Auto Auction Rules." In Japanese. Available at `https://www.uss-engine.com/info_items/rules/auction_20220301.pdf`.

van den Bijgaart, Inge, and Davide Cerruti. 2024. "The Effect of Information on Market Activity: Evidence from Vehicle Recalls." *Review of Economics and Statistics*, 106(1): 230–245.

Van Heerde, Harald, Kristiaan Helsen, and Marnik G. Dekimpe. 2007. "The Impact of a Product-Harm Crisis on Marketing Effectiveness." *Marketing Science*, 26(2): 230–245.

A Additional Results

A.1 Log Prices

Tables 11 and 12 re-estimate the aggregated price specifications of Tables 5 and 4 with the log price as the outcome, on the same samples and with the same controls and fixed effects. Retail listing prices of affected vehicles fall by 0.023 log points in the short run and by 0.016 log points in the medium run, by 0.018 and 0.015 log points at Nissan's affiliated dealers and by 0.040 and 0.020 log points at independent dealers, whereas wholesale prices fall by 0.068 log points in the short run and by 0.038 log points in the medium run. The log specifications thus reproduce the ranking of the level estimates in both horizons: the wholesale decline exceeds the retail decline, and the retail decline is smaller at affiliated than at independent dealers.

Table 11: Retail Listing Price Relative Changes, Log Price

	Pooled	By dealer type
Panel A: Short run, Jun–Jul 2016		
Nissan Affected × Post SR	-0.023***	
	(0.003)	
Nissan Affected × Nissan Affiliated × Post SR		-0.018***
		(0.003)
Nissan Affected × Independent × Post SR		-0.040***
		(0.005)
Panel B: Medium run, Aug–Dec 2016		
Nissan Affected × Post MR	-0.016***	
	(0.002)	
Nissan Affected × Nissan Affiliated × Post MR		-0.015***
		(0.002)
Nissan Affected × Independent × Post MR		-0.020***
		(0.003)
N	150,975	150,975
R2 (within)	0.14	0.14
FE: Listing month	X	X
FE: Model × grade × vintage	X	X
FE: Dealer	X	X

The outcome is the log listing price; sample, periods and right-hand side are those of the listing-price columns of Table 5. $*$ p $<$ 0.1, $**$ p $<$ 0.05, $***$ p $<$ 0.01. Standard errors clustered by dealer are in parentheses. Sample: first observed listings of Nissan-affected and unaffected kei cars of vintage 2013–2016, first listed August 2015 to December 2016, with 500 km to under 100,000 km on the odometer and listing prices trimmed at the top and bottom 1 percent within vehicle group and listing month, excluding Mitsubishi dealers and one nationwide chain that never records where a car went. The sample keeps affected cars at Nissan-affiliated and independent dealers and unaffected cars at independent dealers and, when the car carries the dealer's own brand (Toyota and Daihatsu counted as one), at other-brand-affiliated dealers; this drops 10,657 first listings (6.1 percent): 1,741 unaffected cars at Nissan-affiliated dealers, 725 affected cars and 8,191 unaffected cars of another brand at other-brand-affiliated dealers. Pre covers August 2015 to April 2016 (April 2016 is the event-study reference month). The short run is the listings first observed in June and July 2016 and the medium run those first observed August to December 2016, both measured against the same pre period; the listings first observed in May 2016 are excluded (Table 5). Affiliated dealers are the franchised dealers of the car's own brand, so the affected cars at affiliated dealers are those at Nissan-affiliated dealers. Controls: repair history, mileage bins, free and paid warranty indicators, binned vehicle age (0/1/2+ years). Fixed effects: listing month, model code by grade by vintage, dealer. The comparison group is the unaffected kei cars. In the by-dealer-type column, the dealer and vehicle fixed effects give the affected cars at each dealer type a level of their own, because every car listed by a Nissan-affiliated dealer is affected.

Table 12: Wholesale Auction Relative Price Changes, Log Sold Price

	Log sold price
Nissan Affected × Post SR, May–Jul 2016	-0.068***
	(0.008)
Nissan Affected × Post MR, Aug–Dec 2016	-0.038***
	(0.007)
N	148,714
Adj. R2	0.89
FE: Auction month	X
FE: Model × grade × vintage	X
FE: Venue	X

The outcome is the log sold price; the sample (sold listings with a positive price) and the right-hand side are those of the sold-price column of Table 4. $*$ p $<$ 0.1, $**$ p $<$ 0.05, $***$ p $<$ 0.01. Standard errors clustered at the model-code by first-registration-year level are in parentheses. Controls: mileage bins, vehicle-age bins (0/1/2+ years), missing, high, and middle condition-score indicators, repair history. Fixed effects: auction month, model code by grade by registration year, auction venue. Affected kei cars are the Nissan-badged models subject to the fuel-economy misconduct disclosure (Dayz and Dayz Roox) of vintage 2013–2016; the comparison group is unaffected kei cars of the main auction sample (500 km to under 100,000 km on the odometer). The sample covers the auction months August 2015 to December 2016; an auction month runs from the 20th of the previous calendar month to the 19th and is labelled by the month in which it ends. Pre covers August 2015 to April 2016 and post the auction months from May 2016 on; the May 2016 auction month opens on April 20, 2016. The short run is May to July 2016 and the medium run August to December 2016, both measured against the same pre period.

A.2 Pooled Retail Listing-Price Event Study

Figure 6 plots the pooled event study of the listing price from equation (1), with the sample and controls of Section 4.2. Listing prices of affected vehicles are unchanged in May 2016, the first snapshot after the disclosure, likely because the May cohort mixes cars published before and after April 20, then fall by 23.9 to 30.1 thousand JPY from June to September 2016 and remain 13.9 to 21.0 thousand JPY lower from October to December.

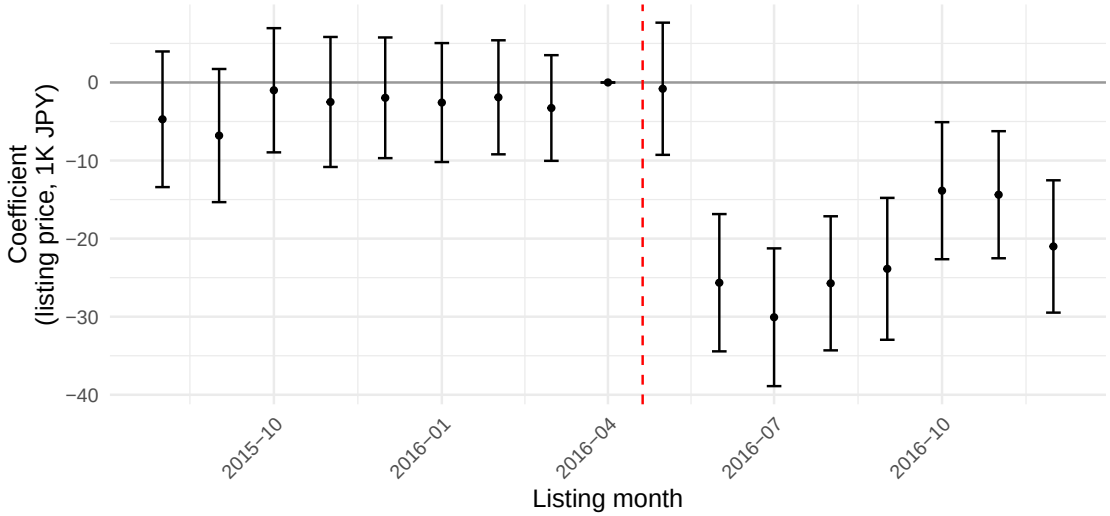


Figure 6: Retail Listing Price Event Study.
The figure plots estimates from equation (1), in thousand JPY at 2016 prices. The dashed vertical line marks April 20, 2016, the disclosure date; the reference month is April 2016. Error bars show 95% confidence intervals based on standard errors clustered at the dealer level.

A.3 Consumer-Sale Event Studies

Figures 7 and 8 plot the event-study estimates of the 3-month consumer-sale probability by entry month, pooled and by dealer type, from the event-study versions of the specifications of Section 4.3. Unlike the aggregated specifications, they keep the May 2016 entrants, and they plot the April 2016 entrants, which were priced before the disclosure, at April 1, just before the disclosure line. The pooled series is −3.70 percentage points for the April entrants, with a standard error of 2.98, rises to 6.78 percentage points for the June 2016 entrants, and lies between −2.24 and 1.39 percentage points from the August entrants on. By dealer type, the April coefficients are −3.01 percentage points at affiliated dealers and −6.42 percentage points at independent dealers, with standard errors of 3.46 and 4.40, neither statistically significant. At affiliated dealers, the coefficient for the June entrants, 7.37 percentage points, is the largest and the only statistically significant one, at the 5 percent level, and at independent dealers the coefficients are positive for the May to July 2016 entrants and between −6.96 and −5.71 percentage points from the September entrants on, none of them statistically significant.

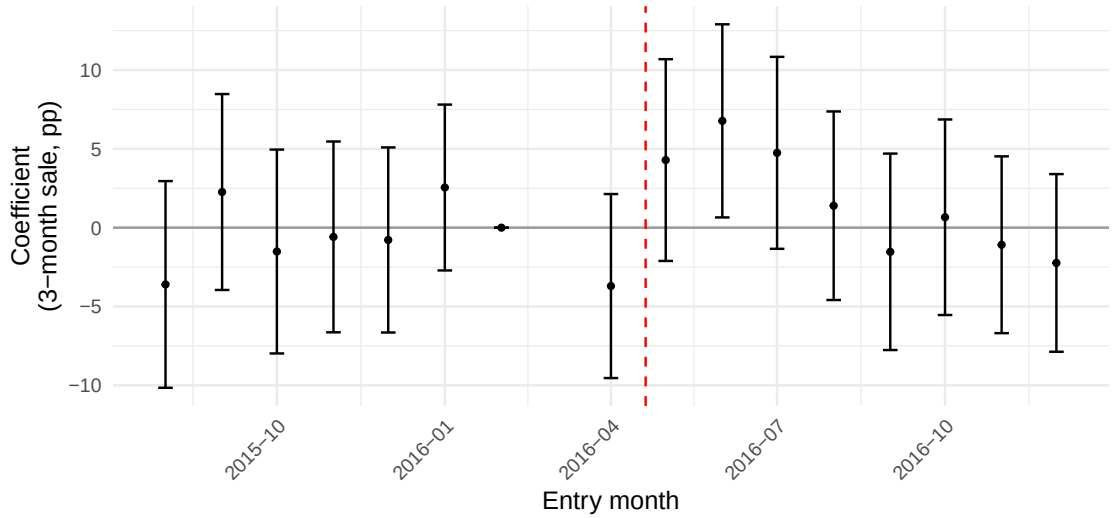


Figure 7: Consumer-Sale Event Study.

The figure plots estimates of the probability, in percentage points, that a new listing sells to a household within 3 months, by entry month, from the event-study version of equation (2). The dashed vertical line marks April 20, 2016, the disclosure date; the reference month is February 2016, and the March 2016 entrants are excluded. The April 2016 entrants, priced before the disclosure, are plotted at April 1. Error bars show 95% confidence intervals based on standard errors clustered at the dealer level.

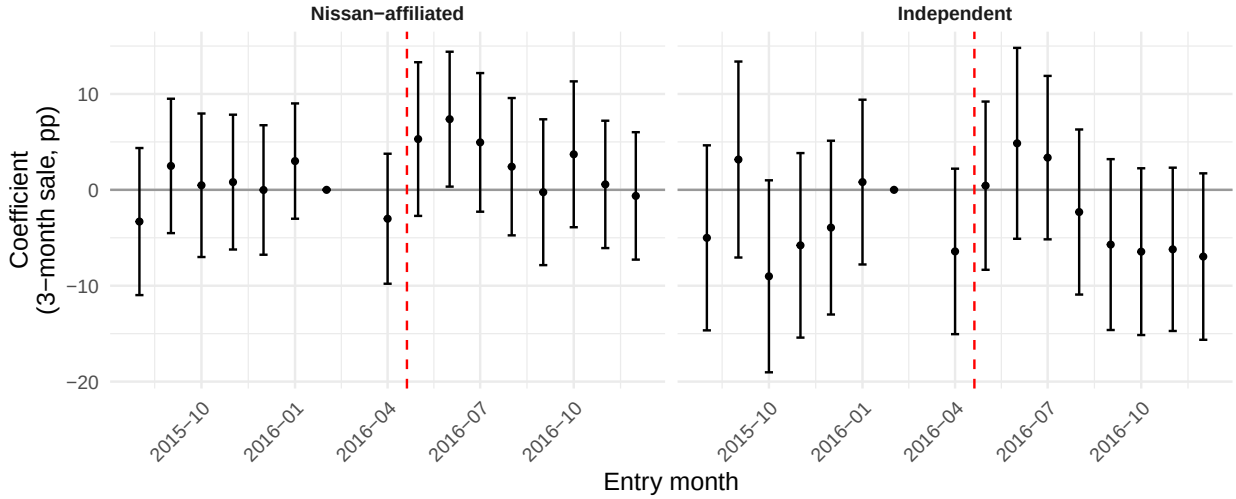


Figure 8: Consumer-Sale Event Study by Dealer Type.

Each panel plots estimates by entry month from the event-study version of equation (3), in percentage points. The dashed vertical line marks April 20, 2016, the disclosure date; the reference month is February 2016, and the March 2016 entrants are excluded. The April 2016 entrants, priced before the disclosure, are plotted at April 1. Error bars show 95% confidence intervals based on standard errors clustered at the dealer level.

A.4 Monthly Consumer-Sale Hazard

An alternative measure of consumer sales is the monthly hazard of a sale to a household in the listing-month panel of the reporting sample, which follows every listed vehicle from its first to its last listing month and includes vehicles listed on both sides of the disclosure. Because the outcome of the April 2016 listing month is realized between the April and May snapshots, April 2016 is the first post month of this design, and its short run covers April to July 2016. On the August 2015 to December 2016 window, after the dealer-type rule of Section 3.3 and after dropping the 48 vehicles whose dealer type changes within the window, the panel contains 271,465 listing-months of 112,689 vehicles at 7,610 dealers. Table 13 reports the mean monthly consumer-sale hazard by dealer type and vehicle group, classifying inventory by the cohort of its first listing, censoring the exits of listings that span the disclosure, and dropping April 2016; Table 14 counts every terminal destination instead. Nissan's affiliated network turns its inventory over much faster than independents once non-consumer exits are counted, reflecting its heavier use of the interdealer channel documented in Appendix C.

Table 15 reports the aggregated estimates for the hazard, with binned vehicle-age dummies, mileage bins, months-since-first-listing bins, and a repair-history indicator as controls and the fixed effects of the main specifications. The pooled hazard of affected vehicles changes by -0.004 in the short run, April to July 2016, and by -0.008 in the medium run, neither statistically significant. At Nissan's affiliated dealers, the estimates are 0.000 and 0.005, whereas at independent dealers the hazard changes by -0.013 in the short run and by -0.026 in the medium run, the latter significant at the 5 percent level. The hazard design therefore points to a lasting decline in consumer sales at independent dealers that the 3-month sale probability of new listings, which misses the post-disclosure sales of the vehicles listed before April 2016 and every sale after a listing's third month, does not show at conventional significance levels.

Table 13: Mean Monthly Consumer Sale Hazard

Vehicle group	Dealer type	Pre	Post
Nissan-affected kei	Affiliated	0.267	0.280
Nissan-affected kei	Independent	0.302	0.286
Unaffected kei	Affiliated	0.255	0.246
Unaffected kei	Independent	0.258	0.253

Share of listing-months in which the car is sold to a household. Sample: listing-months of Nissan-affected and unaffected kei cars of vintage 2013–2016 with 500 km to under 100,000 km on the odometer, August 2015 to December 2016, excluding Mitsubishi dealers and keeping only the dealers that record an exit destination in every month of that window; a car whose destination is not recorded is treated as an internal transfer. As for the listing price, the sample keeps affected cars at Nissan-affiliated and independent dealers, unaffected cars at independent dealers, and unaffected cars of the dealer's own brand at other-brand-affiliated dealers, all read at the car's first listing, and only the listings that stay at one dealer type within the window; this drops 9,907 listings (8.1 percent) with 24,230 listing-months (8.2 percent): 1,437 unaffected cars at Nissan-affiliated dealers, 669 affected cars and 7,753 unaffected cars of another brand at other-brand-affiliated dealers, and 48 listings that change dealer type. April 2016 is dropped. Pre runs from August 2015 through March 2016, with the exits of listings that span the announcement set to zero; post runs from May to December 2016 and covers only listings first observed from May 2016. Affiliated dealers are the franchised dealers of the car's own brand: Nissan's dealers for affected cars and the other manufacturers' dealers for unaffected cars.

Table 14: Mean Monthly Exit Hazard, All Terminal Destinations

Vehicle group	Dealer type	Pre	Post
Nissan-affected kei	Affiliated	0.469	0.483
Nissan-affected kei	Independent	0.355	0.348
Unaffected kei	Affiliated	0.417	0.381
Unaffected kei	Independent	0.324	0.333

Share of listing-months in which the car leaves the lot for any destination: sale to a household, auction, another dealer, export, scrap, or an internal transfer. Sample, dealer types and period cuts as in the consumer sale-hazard table.

Table 15: Consumer Sale Hazard Relative Changes

	Pooled	By dealer type
Panel A: Short run, Apr–Jul 2016		
Nissan Affected × Post SR	-0.004	
	(0.009)	
Nissan Affected × Nissan Affiliated × Post SR		0.000
		(0.013)
Nissan Affected × Independent × Post SR		-0.013
		(0.011)
Panel B: Medium run, Aug–Dec 2016		
Nissan Affected × Post MR	-0.008	
	(0.008)	
Nissan Affected × Nissan Affiliated × Post MR		0.005
		(0.010)
Nissan Affected × Independent × Post MR		-0.026**
		(0.011)
N	270,392	270,392
R2 (within)	0.00	0.00
FE: Listing month	X	X
FE: Model × grade × vintage	X	X
FE: Dealer	X	X

$*$ p $<$ 0.1, $**$ p $<$ 0.05, $***$ p $<$ 0.01. Standard errors clustered by dealer are in parentheses. The outcome is 1 in the month in which a listed car is sold to a household and 0 otherwise, including the month in which it leaves for any other destination. Sample: listing-months of Nissan-affected and unaffected kei cars of vintage 2013–2016 with 500 km to under 100,000 km on the odometer, August 2015 to December 2016, excluding Mitsubishi dealers and keeping only the dealers that record an exit destination in every month of that window; a car whose destination is not recorded is treated as an internal transfer. As for the listing price, the sample keeps affected cars at Nissan-affiliated and independent dealers, unaffected cars at independent dealers, and unaffected cars of the dealer's own brand at other-brand-affiliated dealers, all read at the car's first listing, and only the listings that stay at one dealer type within the window; this drops 9,907 listings (8.1 percent) with 24,230 listing-months (8.2 percent): 1,437 unaffected cars at Nissan-affiliated dealers, 669 affected cars and 7,753 unaffected cars of another brand at other-brand-affiliated dealers, and 48 listings that change dealer type. Post begins in April 2016, whose exits straddle the April 20 announcement, so pre covers August 2015 to March 2016. The short run is April to July 2016 and the medium run August to December 2016, both measured against the same pre period. Affiliated dealers are the franchised dealers of the car's own brand, so the affected cars at affiliated dealers are those at Nissan-affiliated dealers. Controls: binned vehicle age (0/1/2+ years), mileage bins, months-since-first-listing bins, repair history. Fixed effects: listing month, model code by grade by vintage, dealer. The comparison group is the unaffected kei cars. In the by-dealer-type column, the dealer and vehicle fixed effects give the affected cars at each dealer type a level of their own, because every car listed by a Nissan-affiliated dealer is affected.

A.5 Auction Lots at Nissan's Venues

Seller identities are not observed at auction, but Nissan's NAA is operated by Nissan and used mainly by its franchised dealers to sell their trade-ins.[47]

Table 16 asks whether the scandal's effect differs between Nissan's venues and the others. It re-estimates the aggregated specification of Section 4.1 with the affected coefficients split by venue type, together with a level for affected lots at Nissan's venues and interactions of the Nissan-venue indicator with each horizon, so that each coefficient compares affected lots with unaffected lots sold at the same type of venue. Relative to unaffected lots at the same type of venue, prices of affected lots fall by 50.72 thousand JPY at Nissan's venues and by 58.96 thousand JPY at other venues in the short run, and by 35.81 and 35.57 thousand JPY in the medium run; the differences of 8.24 and −0.23 thousand JPY, with standard errors of 7.81 and 9.25, are not statistically significant. The unsold rate of affected lots rises by 14.49 percentage points at Nissan's venues against 7.57 percentage points elsewhere in the short run and by 6.13 against −0.21 percentage points in the medium run, and the differences of 6.91 and 6.35 percentage points, with standard errors of 5.84 and 5.90, are not statistically significant either. The scandal's effect on affected lots thus does not differ between the venues where Nissan's dealers trade and the rest of the market, whereas the unsold rate of unaffected lots at Nissan's venues falls after the disclosure relative to unaffected lots elsewhere, by 8.17 percentage points in the short run, significant at the 5 percent level, and by a statistically insignificant 8.49 percentage points in the medium run.

[47] Based on conversations with industry participants.

Table 16: Wholesale Auction Changes by Venue Type

	Sold price (1K JPY)	Unsold (pp)
Nissan Affected × NAA × Post SR, May–Jul 2016	-50.72***	14.49***
	(10.11)	(5.42)
Nissan Affected × NAA × Post MR, Aug–Dec 2016	-35.81**	6.13
	(13.87)	(5.96)
Nissan Affected × Other venues × Post SR, May–Jul 2016	-58.96***	7.57***
	(6.83)	(2.35)
Nissan Affected × Other venues × Post MR, Aug–Dec 2016	-35.57***	-0.21
	(7.65)	(2.75)
NAA − Other, May–Jul 2016	8.24	6.91
	(7.81)	(5.84)
NAA − Other, Aug–Dec 2016	-0.23	6.35
	(9.25)	(5.90)
NAA × Post SR, May–Jul 2016	3.44	-8.17**
	(7.06)	(3.67)
NAA × Post MR, Aug–Dec 2016	14.81***	-8.49
	(5.60)	(5.37)
N	148,714	222,999
Adj. R2	0.90	0.23
FE: Auction month	X	X
FE: Model × grade × vintage	X	X
FE: Venue	X	X

NAA is Nissan's auction network; the other venues are all remaining venues, including Toyota's TAA and Honda AA. NAA × Post SR and NAA × Post MR are the short-run and medium-run changes for unaffected lots at NAA relative to those at the other venues, so each affected coefficient compares affected and unaffected lots at the same type of venue. The regressions also include the affected × NAA level, which no fixed effect absorbs; the venue fixed effects absorb the NAA level. NAA − Other is the difference between the two affected rows of the same window, with its standard error from the clustered variance matrix. Outcomes, samples, controls, fixed effects, and windows are those of Table 4. Standard errors clustered by model code and first-registration year are in parentheses; $*$ p < 0.1, $**$ p < 0.05, $***$ p < 0.01.

B Auction Condition Scores

Before each auction, inspectors at the venue examine every listed car and assign it an overall condition score. On the scale used by the USS group, for example, S marks a car within about 1 year of registration with fewer than 10,000 kilometers and almost no damage. The scores 6, 5, 4.5, and 4 apply to cars with fewer than 30,000, 50,000, 100,000, and 150,000 kilometers, respectively, and with progressively more visible damage; 3.5 and 3 mark cars that need panel work or more extensive repair, 2 a car in poor condition, and 1 a car that has been flooded or burned.[48] Cars with a repair history, whose frame has been repaired or replaced, are graded R instead, or RA when the repair is minor. The inspection covers the exterior, such as scratches, dents, and rust, the interior, such as tears, burns, and odors, and the mechanical condition of the engine and suspension. Because people assign the scores, the same car can receive different scores at different venues.

The auction data record the overall score as S or a number. Every car with a repair history carries a score of 0, as do other cars without a numeric grade: 38.7 percent of the listings with a score of 0 have a repair history, and no listing with any other score has one. I group the scores into four categories: high, for S and scores of 5 or above, which covers 43.9 percent of the 223,120 listings of the auction sample; middle, for 4 and 4.5, at 38.4 percent; low, for 3.5 and below, at 4.8 percent; and missing, for a score of 0, at 12.9 percent. The regressions of Section 4.1 include indicators for the high, middle, and missing categories, with the low category as the omitted group, together with the repair-history indicator.

C Retail Destinations

This appendix examines where vehicles go when they leave a dealer's inventory. The destination margin complements the consumer-sale analysis of Section 4.3: a dealer can respond to the scandal not only by repricing but by diverting affected vehicles away from consumers, toward auctions, other dealers, or stores within the same company, which is the inventory-return margin at the center of the model in Section 5.

C.1 Descriptive Destination Shares

The data record a terminal destination for each inventory exit: sale to a consumer, sale at auction, sale to another dealer without going through auction, transfer to another store within the same company, export, or scrap. I group auction sales, dealer sales, exports, and scrap into an "other sale" category, while internal transfers remain separate because they keep the vehicle inside the same dealer organization. Table 17 reports terminal destination shares for the listings of the reporting sample that are listed between August 2015 and December 2016, the listing-month panel of Appendix A, by vehicle group, dealer type, and period, with the affiliated dealers of Section 3.3

[48] Goo-net Magazine, web article, updated September 28, 2023, `https://www.goo-net.com/magazine/contents/check-point/25749/`.

(Nissan's dealers for affected cars and the dealers of the car's own brand for unaffected cars), classifying listings as pre-scandal if first listed by April 2016 and post-scandal otherwise; because the reporting sample keeps only dealers that record destinations in every month from August 2015 through December 2016, fewer than 3 percent of its listings lack a recorded destination. Consumer sale is the main exit route for affected vehicles in every cell, but its share differs sharply across dealer types: roughly eight in ten affected-vehicle exits at independent dealers are consumer sales, against fewer than six in ten at Nissan's affiliated dealers, which route a third of their affected vehicles to other sales and about a tenth to internal transfers. Unaffected kei cars show the same contrast with a smaller gap: before the disclosure, consumer sales account for 0.652 of their exits at the affiliated dealers of their own brand against 0.795 at independent dealers, other sales for 0.219 against 0.165, and internal transfers for 0.129 against 0.040. These shares change little between the pre and post periods. The interdealer channel is therefore quantitatively important, especially at affiliated dealers, which supports modeling the return of unsold inventory to the wholesale market as a first-order margin.

Table 17: Terminal Destination Shares

	Affiliated		Independent	
Destination	Pre	Post	Pre	Post
Panel A: Nissan-affected kei				
Consumer	0.559	0.571	0.845	0.841
Other sale	0.342	0.336	0.127	0.112
Internal transfer	0.098	0.093	0.028	0.047
Listings with a recorded destination	6,158	4,127	2,357	2,424
Listings without a recorded destination	39	11	12	30
Panel B: Unaffected kei				
Consumer	0.652	0.654	0.795	0.767
Other sale	0.219	0.210	0.165	0.179
Internal transfer	0.129	0.136	0.040	0.054
Listings with a recorded destination	24,398	23,205	23,472	23,739
Listings without a recorded destination	1,602	734	141	240

The unit is a listing of the sale-hazard sample, assigned to the dealer type and the period of its first listing month in the window. As in the summary statistics, pre covers listings first listed from August 2015 through April 2016 and post those first listed from May to December 2016. Shares are among listings with a recorded destination and sum to one within each column of a panel. Other sale covers auction, sale to another dealer, export, and scrap; internal transfer is a move to another store of the same company. The sale-hazard analysis treats a listing without a recorded destination as an internal transfer. Sample: listings of Nissan-affected and unaffected kei cars of vintage 2013–2016 in the listing panel of August 2015 to December 2016, with 500 km to under 100,000 km on the odometer, excluding Mitsubishi dealers and one nationwide chain that never records where a car went, at dealers that record an exit destination in every month from August 2015 to December 2016. As for the listing price, the sample keeps affected cars at Nissan-affiliated and independent dealers, unaffected cars at independent dealers, and unaffected cars of the dealer's own brand at other-brand-affiliated dealers, all read at the car's first listing, and only the listings that stay at one dealer type within the window; this drops 9,907 listings (8.1 percent) with 24,230 listing-months (8.2 percent): 1,437 unaffected cars at Nissan-affiliated dealers, 669 affected cars and 7,753 unaffected cars of another brand at other-brand-affiliated dealers, and 48 listings that change dealer type. Affiliated dealers are the franchised dealers of the car's own brand: Nissan's dealers in Panel A and the other manufacturers' dealers in Panel B.

D Estimation Details

This appendix gives the details of the estimation of Section 6: the price standardization, the measurement of the externally set parameters, the computation, the inference, and the table of parameter values and estimates.

D.1 Price Standardization

In the model, cars of a brand are homogeneous, so each dealer has one price. In the data, the price of a car also depends on its age, mileage, repair history, model, grade, vintage, and listing month. The moments on prices therefore use prices standardized to a reference car: I remove the part of each price that these attributes explain and add back the value of the same attributes for the reference car. For the retail and the wholesale market separately, I estimate a hedonic regression of the real price on age, mileage, and repair-history controls, with model–grade–vintage and month fixed effects and, in retail, dealership fixed effects. The regression also includes affected-by-month terms, by dealer type in retail, which absorb the scandal's price path. The retail regression uses all first listings of the analysis sample, which cover the 17 listing months of August 2015 to December 2016, and the wholesale regression uses all sold lots of the auction sample, from the auction months of August 2015 to December 2016. Let $\hat{p}_i$ denote the price that the regression predicts for car i from its attributes, its model–grade–vintage, and its month alone, without the dealership fixed effect and the affected-by-month terms. The standardized price is

$$\tilde{p}_i = p_i - \hat{p}_i + \bar{p},$$

where $\bar{p}$ is the predicted price of the reference car.

The reference car averages the attributes of the retail sample over cars that also trade at auction: $\bar{p}$ is the average of $\hat{p}_i$ over the combinations of brand, model–grade–vintage, age, mileage, repair history, and month that occur both among all first listings and among sold auction lots, weighted by their frequency among the first listings of August 2015 to December 2016 and pooled over both brands, both periods, and all dealer types. Each market's regression gives its own $\hat{p}_i$ and $\bar{p}$, evaluated at the same reference car, so retail and wholesale standardized prices refer to the same car. The standardized retail price keeps the dealership fixed effect and the affected-by-month terms, so it still reflects differences across dealers and the scandal's price path. Standardized prices enter the moments on retail and wholesale prices, in the means, standard deviations, and price–sale covariances of the initial price and in the two wholesale prices, whereas the changes after the disclosure that the estimation also targets (Section 6.4.2) use observed prices.

D.2 Externally Set Parameters

Dealer mass D_d. In (23), A_{it} counts all new listings of dealership i in month t, regular cars and kei cars of every vintage, and K_{it} its new listings of kei cars of every vintage, both with mileage

from 50 kilometers to below 100,000 kilometers. The counts cover the 17 months of August 2015 to December 2016, and because a model period lasts 3 months, the dealer mass per period is 3/17 of the total. The sum excludes the nationwide chain that never records destinations, and the dealer mass stays fixed before and after the scandal.

Inflow E_b. For each cell, I count the new listings of the entrants of August 2015 to January 2016 and their consumer sales within 3 months, the latter measured among the dealers that report exit destinations and reweighted to all listings, and I divide both by the 6 months of the window and multiply by 3 to express them per model period. The months after the disclosure are not a stationary state, so their sales do not measure an inflow.

Demand correlation ρ. For the 414 independent dealerships that listed at least three affected and three unaffected cars in the pre-scandal listing months of April 2015 to March 2016, I compute the correlation across the two brands of the dealership's mean residual listing price, where the residual comes from a hedonic regression of the price on the car's attributes with fixed effects for each combination of dealer type, brand, region, and month. The split-half method corrects this correlation for the sampling noise of the dealership means: it splits each dealership's cars of a brand into two halves, estimates the variance of the dealership component as the covariance of the two half means, and divides the cross-brand covariance of the dealership means by the geometric mean of these variances. The corrected dealership-level correlation is 0.91. The model's ρ is the correlation of the two demands that one entrant draws for two different cars, so it also carries the car-specific part of the residual, which is uncorrelated across cars. Scaling the dealership-level correlation by the dealership share of the residual variance of a single listing, 0.50 for affected and 0.37 for unaffected cars, gives $0.91 \times \sqrt{0.50 \times 0.37} = 0.39$, with a bootstrap standard error of 0.02. Weighting all dealerships equally gives 0.35, and using only the first listing month of each car gives 0.38. The hedonic regression and sample of this calibration come from an earlier version of the analysis and differ in their window and controls from the price standardization of Appendix D.1.

D.3 Computation

Every evaluation of either criterion solves a stationary equilibrium, in the following nested steps, and every evaluation of the Stage-2 criterion also solves the transition path.

1. *Cutoffs.* Given θ and $\boldsymbol{w}$, compute the cutoffs. In the baseline, they have the closed forms of Appendix F.2. With the holding cost, the cutoff of an affiliated type solves $\overline{V}_d(\bar{\xi}_{db}; w_b) = w_b$ by a safeguarded Newton–bisection method, and the brand boundary of independent dealers solves the same equation with the net value of the other brand in place of zero.

2. *Integration.* Integrate over the standardized demand for the held brand by composite Gauss–Legendre quadrature, with panels at the cutoffs (22 nodes per panel on the estimation grid and 32 on a finer check grid). For independent dealers, the other brand's demand is integrated

exactly through the conditional normal distribution function. The expectation over the holding cost uses Gauss–Hermite quadrature (49 nodes; 65 on the check grid). The holding shares, the cell moments, and household sales come from the same nodes.

3. *Wholesale prices.* Solve $N_A(\boldsymbol{w})Q_A(\boldsymbol{w}) = E_A$ and $N_U(\boldsymbol{w})Q_U(\boldsymbol{w}) = E_U$ jointly by damped Newton iteration. Starting from the previous candidate's prices or the observed wholesale prices, repeat steps 1–2 at the current prices, compute the residuals $E_b - N_b(\boldsymbol{w})Q_b(\boldsymbol{w})$, and update (w_A, w_U) by a Newton step whose finite-difference Jacobian recomputes the cutoffs at the perturbed prices; each step is capped at $3\sigma_\epsilon$ and halved until the residuals fall. By the monotonicity of Lemma 2, a price below its clearing value makes household sales exceed the inflow, so the step raises it, and conversely. The iteration stops when the residuals, scaled by the inflow, fall below 10^{-12}, which takes a few iterations.

4. *Transition path.* Given λ, solve the post-scandal stationary equilibrium by step 3 with the lowered mean demands, then the $2T$ prices of the path from the stacked clearing conditions (28) by Newton iteration, as described in Appendix E.2. Steps 1–2 evaluate each period at its own pair of prices, $(\boldsymbol{w}_t, \boldsymbol{w}_{t+1})$; the Jacobian is computed by central finite differences and is block tridiagonal; each step is capped at $3\sigma_\epsilon$ and halved until the sum of squared scaled residuals falls, and the iteration stops when every scaled residual is below 10^{-12}. Within Stage 2, each candidate starts from the previous candidate's path, shifted by the change in the post-scandal prices.

5. *Parameters.* Stage 1 minimizes $J(\theta)$ by L-BFGS-B, with the money parameters in thousands and the standard deviations in logs, using the gradient of the standardized residuals, whose Jacobian includes the response of the wholesale prices. The search runs from four starting values of σ_ϵ between 100 and 400 thousand JPY on a coarse grid, and the best start is refitted on the estimation grid. Stage 2 minimizes (25) over (λ_N, λ_I) by BFGS with a finite-difference gradient from three starts, which reach the same minimum.

A fit is accepted if the projected gradient of its criterion is below 0.01 and solving the equilibrium on the check grid changes no moment by more than 0.1 sampling standard errors; a Stage-2 fit must in addition change no coefficient by more than 0.01 sampling standard errors when the path is extended to $T = 12$. Every reported fit passes. Neither Stage-1 fit has a parameter at a bound, the projected gradient of the criterion is 0.0011 in the baseline and 0.0006 with the holding cost, and refining the quadrature changes no moment by more than 10^{-10} sampling standard errors; the projected gradient of the Stage-2 fit is below 10^{-6}.

D.4 Inference

Sampling uncertainty comes from 100 paired cluster multiplier draws. Each draw assigns an independent exponential weight with mean one to every dealership and to every auction model–vintage cluster, shared across the linked samples, and recomputes the price standardization, the reporting

weights, the empirical moments, and the Stage-2 targets. The sampling standard error s_ℓ of each moment and target is its standard deviation across the draws. Standard errors of the parameters come from a local linear mapping of each draw into the estimates. With J the Jacobian of the 22 model moments with respect to θ at $\widehat{\theta}$, including the response of the wholesale prices, and $\Delta\widehat{m}_r$ the deviation of the empirical moments in draw r from their point values, the draw moves the Stage-1 estimates by

$$\Delta\theta_r = (J'WJ)^{-1}J'W\,\Delta\widehat{m}_r,$$

and the standard error of each parameter is the standard deviation of $\Delta\theta_r$ across the draws. The Stage-2 standard errors apply the same mapping to the deviations of the twelve targets, with the 12×2 Jacobian of the model coefficients Δm_ℓ with respect to λ in place of J. These standard errors condition on the dealer masses and the inflow, which each draw also recomputes but which the mapping holds at their point values, and on ρ and δ; the Stage-2 standard errors also condition on the Stage-1 estimates.

D.5 Parameter Estimates

Table 18 reports the externally set parameters (Panel A), the Stage-1 estimates of the baseline and of the specification with the holding cost (Panels B and C), and the Stage-2 estimates of the demand losses (Panel D). Section 7.1 discusses them.

Table 18: Parameter Values and Estimates

Parameter	Symbol	Value	SE
Panel A: Externally Set Parameters			
Dealer mass, Nissan-affiliated	D_N	23,413	
Dealer mass, independent	D_I	114,986	
Dealer mass, other-brand-affiliated	D_B	90,238	
Inflow, affected	E_A	1,726	
Inflow, unaffected	E_U	11,281	
Demand correlation	ρ	0.4	
Discount factor	δ	0.98	
Panel B: Stage-1 Estimates, Baseline			
Mean demand, affected cars at Nissan-affiliated dealers	μ_{NA}	468.3	8.5
Mean demand, affected cars at independent dealers	μ_{IA}	27.5	13.0
Mean demand, unaffected cars at independent dealers	μ_{IU}	471.0	5.5
Mean demand, unaffected cars at other-brand-affiliated dealers	μ_{BU}	649.3	3.6
Demand standard deviation, affected cars	σ_A	377.5	4.9
Demand standard deviation, unaffected cars	σ_U	302.8	2.5
Household demand scale	σ_ϵ	68.1	1.1
Panel C: Stage-1 Estimates, with Holding Cost			
Mean demand, affected cars at Nissan-affiliated dealers	μ_{NA}	648.5	19.0
Mean demand, affected cars at independent dealers	μ_{IA}	356.2	30.1
Mean demand, unaffected cars at independent dealers	μ_{IU}	747.1	24.7
Mean demand, unaffected cars at other-brand-affiliated dealers	μ_{BU}	813.3	16.7
Demand standard deviation, affected cars	σ_A	286.9	11.0
Demand standard deviation, unaffected cars	σ_U	200.8	11.5
Household demand scale	σ_ϵ	138.1	3.6
Holding-cost mean, affiliated dealers	μ^c_{aff}	72.8	6.3
Holding-cost mean, independent dealers	μ^c_I	160.4	9.6
Holding-cost standard deviation, affiliated dealers	σ^c_{aff}	125.2	4.7
Holding-cost standard deviation, independent dealers	σ^c_I	221.0	6.7
Panel D: Stage-2 Estimates, Demand Losses (with Holding Cost)			
Demand loss, Nissan-affiliated dealers	λ_N	11.1	7.1
Demand loss, independent dealers	λ_I	55.2	20.7

Panel A lists the parameters set outside the estimation (Section 6.2.2). Dealer masses and the inflow are numbers of entering dealers and of cars per 3-month model period; the inflow is measured from the entrants of August 2015 to January 2016 and held fixed in every period. Panels B and C report minimum-distance estimates of the pre-scandal stationary equilibrium from the 22 moments of Table 9, in thousand JPY at 2016 prices. In the baseline, dealers of a type differ only in demand. The specification with a holding cost adds a normal holding cost of an unsold car, drawn after the inventory choice from one distribution for the two affiliated dealer types and one for independent dealers; at the estimates, 28 and 23 percent of the holding costs of affiliated and independent dealers are negative. The criterion, the mean squared standardized residual, is 110.15 in Panel B and 8.05 in Panel C. Panel D reports the demand losses, in thousand JPY, fitted to the twelve targets of Table 10 on the model's transition path (Section 6.4.2), holding Panels A and C fixed, with the inflow at its pre-scandal value; the criterion is 5.46. All three fits pass the acceptance checks. Standard errors map 100 paired multiplier draws of the data moments into changes in the estimates through the Jacobian of the model moments at the estimates; they hold Panel A fixed, and those of Panel D also hold Panel C fixed. See Appendix D.4 for more details.

E Transition Path

This appendix states the equilibrium conditions along the transition path of Section 6.3, describes how they are solved, and shows the path at the estimates. Each stage of a dealer's life is as in Section 5, with next period's wholesale price in place of a constant one.

E.1 Equilibrium Path

Timing. Periods are indexed by t. In periods $t \leq -1$, the market is in the pre-scandal stationary equilibrium, with wholesale prices $\boldsymbol{w}^{\text{pre}}$. The disclosure arrives in period 0 after the wholesale stage: the young dealers of period 0 chose their inventory and posted their prices under pre-scandal expectations, and their households learn of the scandal while the cars are on the lot. From period 0 on, the household of a type-d dealer values an affected car at $\xi_{iA} - \lambda_d$, permanently. From period 1 on, entering dealers know the losses and the whole path of wholesale prices $\{\boldsymbol{w}_t\}_{t \geq 1}$, so they have perfect foresight, and the inflow of each brand stays at E_b in every period.

Retail stage. A young dealer that holds a car of brand b in period $t \geq 1$ sells the car at wholesale in period $t+1$ if its household does not buy it, so it values an unsold car at $\delta w_{b,t+1} - c_i$, and its price satisfies the condition (20) with next period's wholesale price in place of a constant one:

$$p^* = \delta w_{b,t+1} - c_i + \sigma_\epsilon \left[1 + \exp\left(\frac{\xi_{ib} - \lambda_d \mathbf{1}\{b = A\} - p^*}{\sigma_\epsilon} \right) \right], \tag{26}$$

in which the household's valuation, $\xi_{iA} - \lambda_d$ for an affected car and ξ_{iU} for an unaffected one, takes the place of demand. The value of holding the car is $V(\xi_{ib} - \lambda_d \mathbf{1}\{b = A\}, c_i; w_{b,t+1})$ of equation (19), and its expectation over the holding cost is $\overline{V}_d(\xi_{ib} - \lambda_d \mathbf{1}\{b = A\}; w_{b,t+1})$ of equation (21).

Wholesale stage. An entrant of period t buys a car at $w_{b,t}$, or forgoes this price by keeping its endowment, and expects $\overline{V}_d(\cdot\,; w_{b,t+1})$ from holding it. Its net value of holding brand b is therefore $\overline{V}_d(\xi_{ib} - \lambda_d \mathbf{1}\{b = A\}; w_{b,t+1}) - w_{b,t}$, in which the price paid for the car and the price expected for it if it does not sell differ. The inventory policies of Sections 5.3 and 6.1 apply to this net value: an affiliated entrant holds a car if and only if the net value is nonnegative, that is, if its demand net of the loss is at least the cutoff $\overline{V}_d^{-1}(w_{b,t}; w_{b,t+1})$, and an independent entrant holds the brand with the larger nonnegative net value. As in the stationary state, the endowment does not affect the choice. The mass $N_{b,t}$ of young dealers that hold a car of brand b, the mean price and the mean sale probability of each cell, and the mean sale probability $Q_{b,t}$ of the dealers that hold brand b follow as in Section 5.4, and all of them depend on the wholesale prices of periods t and $t+1$, $(\boldsymbol{w}_t, \boldsymbol{w}_{t+1})$.

Market clearing. Off the stationary state, the cars that the previous cohort did not sell are a state variable. Let $R_{b,t}$ denote the cars of brand b that the old dealers of period t bring to the

wholesale market, the cars that the previous cohort held and did not sell to households,

$$R_{b,t+1} = N_{b,t}\big[1 - Q_{b,t}\big]. \tag{27}$$

The wholesale market of brand b clears in period t when the young dealers hold every car present, the inflow and these returns, as in the stationary state:

$$N_{b,t}(\boldsymbol{w}_t, \boldsymbol{w}_{t+1}) = E_b + R_{b,t}. \tag{28}$$

Given the returns $R_{b,1}$ that the period-0 cohort leaves, an equilibrium path is a sequence of wholesale prices $\{\boldsymbol{w}_t\}_{t\geq 1}$ such that, when every dealer prices by (26) and chooses its inventory by the policies above, condition (28) holds for both brands in every period $t \geq 1$, with the returns evolving by (27). A stationary equilibrium is a constant solution of these conditions: with $\boldsymbol{w}_t = \boldsymbol{w}_{t+1} = \boldsymbol{w}$ and $R_{b,t+1} = R_{b,t}$, substituting (27) into (28) gives $N_b(\boldsymbol{w})Q_b(\boldsymbol{w}) = E_b$, the clearing condition of Definition 1. Along the path, by contrast, household sales need not equal the inflow in any period. A cohort that sells fewer cars than arrive passes the difference on to the next cohort as returns, which next period's prices must induce young dealers to absorb.

The surprise cohort. The cohort of period 0 holds the pre-scandal inventory and has posted the pre-scandal prices $p(\xi_{ib}, c_i; w_b^{\text{pre}})$. Its affected cars sell with the probability that the lowered valuation implies at the posted price p, $\exp((\xi_{iA} - \lambda_d - p)/\sigma_\epsilon)/[1 + \exp((\xi_{iA} - \lambda_d - p)/\sigma_\epsilon)]$, whereas its unaffected cars sell as before. The data counterpart of this cohort is the April 2016 entrants (Section 4.3), which were priced at the April 1 snapshot and faced the disclosure for 72 of the 91 days of their 3-month window. I therefore let the period-0 cohort face the loss for the share $1 - s$ of its retail stage, with $s = 19/91$. Under a constant sale hazard within the period, the probability that an affected car remains unsold at the end of the stage is $[1 - q^{\text{pre}}]^s[1 - q^{\text{shock}}]^{1-s}$, where q^{pre} and q^{shock} are the sale probabilities at the posted price before and after the loss.[49] The cars that this cohort does not sell, $R_{b,1} = N_b^{\text{pre}}(1 - Q_b^0)$ with Q_b^0 its mean sale probability, open the path. For affected cars, they exceed the pre-scandal returns, because fewer of them sell at the pre-scandal prices; at the estimates, the excess is 5.0 percent.

E.2 Computation

The path converges to the post-scandal stationary equilibrium $\boldsymbol{w}^{\text{post}}$, which solves $N_b(\boldsymbol{w})Q_b(\boldsymbol{w}) = E_b$ with the lowered mean demands. I truncate the path at $T = 8$ periods, 2 years, by assuming that the market is in this equilibrium from period $T + 1$ on, $\boldsymbol{w}_{T+1} = \boldsymbol{w}^{\text{post}}$, and solve the $2T$ wholesale prices of periods 1 to T jointly from the $2T$ clearing conditions (28). Each period is linked to the one before it, through the returns that the previous cohort leaves, and to the one after it, through

[49] Both sale probabilities are evaluated dealer by dealer, at each dealer's demand and holding cost, and averaged over the dealers that hold a car in the pre-scandal equilibrium. The same blend, applied to the data, converts the very-short-run estimates of Section 4.3 into the change for a fully exposed cohort.

the value of an unsold car, so the conditions cannot be solved period by period in either direction. I stack them and solve them by Newton's method. Because each clearing condition involves the prices of only three periods, its own, the one before, and the one after, the Jacobian of the stacked conditions is block tridiagonal, and I compute it by finite differences, starting from a flat path at $\boldsymbol{w}^{\text{post}}$. Extending the path to $T = 12$ changes no targeted coefficient by more than 10^{-4} sampling standard errors. Appendix D.3 gives the step sizes and tolerances.

The path is locally unique. Linearized around the post-scandal stationary equilibrium, the map from the state of one period, the unsold cars $(R_{A,t}, R_{U,t})$ and the wholesale prices $(w_{A,t}, w_{U,t})$, to the state of the next has two eigenvalues outside the unit circle, 1.85 and 1.82, as many as the forward-looking prices, and two inside, 0.56 and 0.55, as many as the predetermined stocks of unsold cars, so a unique path converges to this equilibrium from nearby states (Blanchard and Kahn, 1980). The Jacobian of the stacked clearing conditions is also nonsingular at the solution. I cannot show that the path is globally unique, but Newton's method reaches the same path, with prices that differ by less than 10^{-10} thousand JPY, from 27 starting paths: flat paths at the pre- and post-scandal prices and at 100 and 200 thousand JPY above and below the post-scandal prices, a linear path from the pre- to the post-scandal prices, and 20 paths that add independent normal deviations with standard deviations of 50 and 100 thousand JPY to the post-scandal prices in every period.

E.3 Path at the Estimates

Figure 9 plots the transition path at the estimates of Table 18: the wholesale price of affected cars and the affected cars that old dealers return to the wholesale market in each model period, together with the pre- and post-scandal stationary values. Section 7.2 discusses the path.

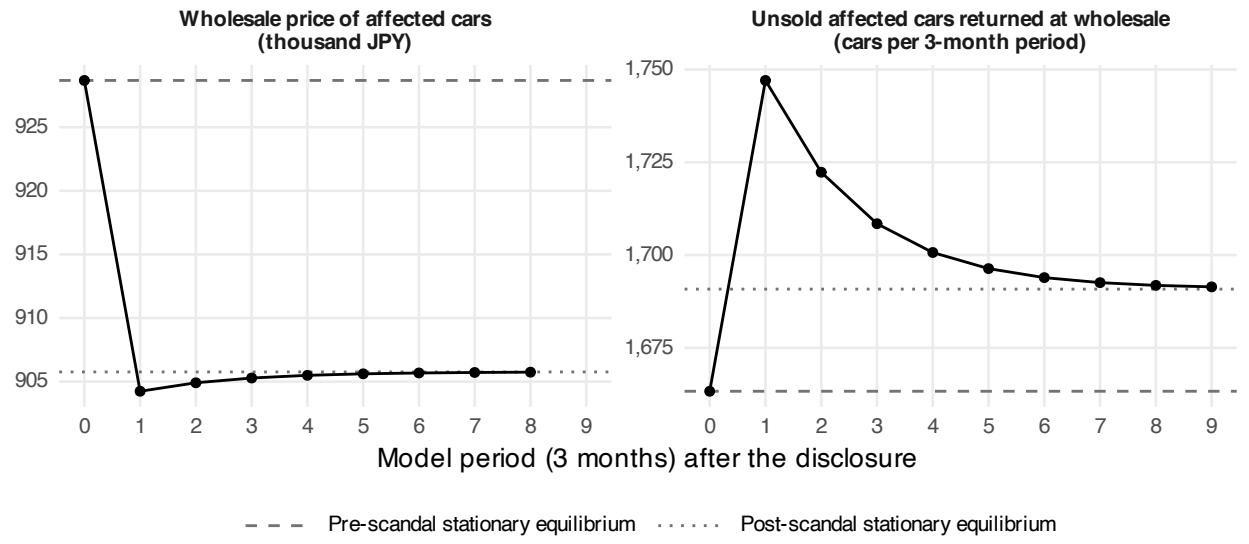


Figure 9: Transition Path at the Estimates.

The left panel plots the wholesale price of affected cars by model period, in thousand JPY at 2016 prices; period 0 is the surprise cohort, which bought at the pre-scandal price. The right panel plots the affected cars that old dealers return to the wholesale market at the start of each period, in cars per 3-month period, which the young dealers of the period absorb together with the inflow. Dashed lines mark the pre-scandal stationary values and dotted lines the post-scandal ones. One model period lasts 3 months, and period 1 begins on April 20, 2016, the disclosure date.

F Assumptions and Proofs

F.1 Proof of Lemma 1

Fix (ξ_{ib}, w_b), and write $L = \delta w_b \geq 0$ for the liquidation value and $q(p) = \exp((\xi_{ib} - p)/\sigma_\epsilon)/[1 + \exp((\xi_{ib} - p)/\sigma_\epsilon)]$ for the probability that the car sells at price p. The objective of problem (6) can be written as

$$\Pi(p) = L + q(p)\,(p - L), \tag{29}$$

with

$$q'(p) = -\frac{q(p)\,[1 - q(p)]}{\sigma_\epsilon}.$$

Differentiating gives

$$\Pi'(p) = q(p)\left[1 - \frac{[1 - q(p)]\,(p - L)}{\sigma_\epsilon}\right].$$

For $p \leq L$ the bracket is positive, so $\Pi'(p) > 0$. For $p > L$, the term $(p - L)[1 - q(p)]$ is strictly increasing in p, since

$$\frac{\partial}{\partial p}\left\{(p - L)\,[1 - q(p)]\right\} = [1 - q(p)] + (p - L)\,\frac{q(p)[1 - q(p)]}{\sigma_\epsilon} > 0,$$

and it rises from zero at $p = L$ to infinity as $p \to \infty$. The equation $(p - L)[1 - q(p)] = \sigma_\epsilon$ therefore has a unique solution p^*, and $\Pi'(p)$ changes sign exactly once, from positive to negative, at p^*, which is therefore the unique maximizer. Because $p^* > L \geq 0$, the maximizer is interior, and no boundary case arises.

Because $1 - q(p) = 1/[1 + \exp((\xi_{ib} - p)/\sigma_\epsilon)]$, the condition $(p^* - L)[1 - q(p^*)] = \sigma_\epsilon$ is (7), and the sale probability (8) is $q(\xi_{ib}, w_b) = q(p^*)$, written q in the rest of the proof. Substituting the condition into the objective gives

$$V = L + q\,(p^* - L) = L + \sigma_\epsilon \frac{q}{1 - q} = p^* - \sigma_\epsilon.$$

Because $\partial q(p)/\partial \xi_{ib} = q(p)[1 - q(p)]/\sigma_\epsilon$ at a given price and $p^* - L = \sigma_\epsilon/(1 - q)$, the envelope theorem applied to (29) yields $\partial V/\partial \xi_{ib} = q$ and $\partial V/\partial L = 1 - q$, so $\partial V/\partial w_b = \delta(1 - q)$, which gives (9). Because $p(\xi_{ib}, w_b) = V + \sigma_\epsilon$, the price has the same derivatives, $\partial p/\partial \xi_{ib} = q < 1$ and $\partial p/\partial w_b = \delta(1 - q) > 0$. The sale probability (8) increases in $\xi_{ib} - p(\xi_{ib}, w_b)$, which therefore rises in ξ_{ib} and falls in w_b, so the sale probability increases in ξ_{ib} and decreases in w_b.

For the range, the middle expression for V shows that $V > \delta w_b$. As $\xi_{ib} \to -\infty$, $q = q(p^*) < q(L) \to 0$ because $p^* > L$ and $q(p)$ decreases in p, so $V \to \delta w_b$. As $\xi_{ib} \to \infty$, $V \geq \Pi(\xi_{ib}/2) \to \infty$, because at the price $\xi_{ib}/2$ the margin over L grows without bound and the car sells with probability $q(\xi_{ib}/2) = \exp(\xi_{ib}/(2\sigma_\epsilon))/[1 + \exp(\xi_{ib}/(2\sigma_\epsilon))] \to 1$. Because $V(\cdot, w_b)$ is continuous and strictly increasing, with $\partial V/\partial \xi_{ib} = q > 0$, it maps $\mathbb{R}$ onto $(\delta w_b, \infty)$, and for every $v > \delta w_b$ the equation $V(\xi_{ib}, w_b) = v$ has a unique solution $V^{-1}(v; w_b)$. □

F.2 Closed Forms of the Cutoffs

With logistic taste shocks, the inverse of the value function has a closed form. Solving the first-order condition (7) for demand gives $\xi_{ib} = p^* + \sigma_\epsilon \log[(p^* - \delta w_b - \sigma_\epsilon)/\sigma_\epsilon]$. Substituting $p^* = V + \sigma_\epsilon$ gives the inverse in closed form,

$$V^{-1}(v; w_b) = v + \sigma_\epsilon \left[1 + \log \frac{v - \delta w_b}{\sigma_\epsilon}\right], \qquad v > \delta w_b. \tag{30}$$

Hence the cutoff (10) is $\bar{\xi}_b(w_b) = w_b + \sigma_\epsilon\big[1 + \log\big((1 - \delta)w_b/\sigma_\epsilon\big)\big]$. The dealer at the cutoff posts the price $w_b + \sigma_\epsilon$ and sells with probability $q = (1 - \delta)w_b/[(1 - \delta)w_b + \sigma_\epsilon]$, the slope of the cutoff is $\mathrm{d}\bar{\xi}_b/\,\mathrm{d}w_b = 1 + \sigma_\epsilon/w_b$, which equals $[1 - \delta(1 - q)]/q$ at the cutoff, and $\bar{\xi}_b(w_b) \to -\infty$ as $w_b \to 0$. The boundary between the regions of the two brands in (13) is the demand for brand A at which brand A is worth exactly as much as the better of brand U and no car: with $s = \max\{V(\xi_{iU}, w_U) - w_U, 0\}$, an independent entrant holds brand A if and only if $V(\xi_{iA}, w_A) - w_A \geq s$, that is, if and only if $\xi_{iA} \geq \tilde{\xi}_A(\xi_{iU}; \boldsymbol{w}) = V^{-1}(w_A + s; w_A)$, which by (30) is

$$\tilde{\xi}_A(\xi_{iU}; \boldsymbol{w}) = w_A + s + \sigma_\epsilon \left[1 + \log \frac{(1 - \delta)w_A + s}{\sigma_\epsilon}\right], \qquad s = \max\big\{V(\xi_{iU}, w_U) - w_U,\, 0\big\}. \tag{31}$$

For $\xi_{iU} \leq \bar{\xi}_U(w_U)$, $s = 0$ and the boundary equals $\bar{\xi}_A(w_A)$. Above $\bar{\xi}_U(w_U)$, s increases in ξ_{iU} and decreases in w_U by (9), and the boundary increases in s and in w_A, so it rises with the demand for brand U and with the wholesale price of brand A and falls with the wholesale price of brand U. Equivalently, the pairs $\big(V^{-1}(w_A + s; w_A), V^{-1}(w_U + s; w_U)\big)$, $s \geq 0$, at which both brands have the net value s, trace the boundary, starting at $\big(\bar{\xi}_A(w_A), \bar{\xi}_U(w_U)\big)$. At $\boldsymbol{w} = \mathbf{0}$, the two coordinates coincide, so the boundary is the line $\xi_{iA} = \xi_{iU}$. These closed forms rely on the absence of a holding cost: with the holding cost of the estimated model, the cutoffs are defined by the same indifference conditions but have no closed form.

F.3 Existence and Uniqueness

This subsection states the assumptions of Proposition 1 and proves it. Throughout, $\Phi_b(\boldsymbol{w}) = N_b(\boldsymbol{w})Q_b(\boldsymbol{w})$ denotes the household sales of brand b, so that the market clears for brand b when $\Phi_b(\boldsymbol{w}) = E_b$ by (17).

Assumption 1 (Demand distributions)**.** *The distributions G_N and G_B have continuous and positive densities on $\mathbb{R}$, and G_I has a continuous and positive density on $\mathbb{R}^2$.*

By Appendix F.2, $\bar{\xi}_b(w_b) \to -\infty$ as $w_b \to 0$, so at zero wholesale prices every entrant holds a car. At $\boldsymbol{w} = \mathbf{0}$, the boundary between the two brands is the line $\xi_{iA} = \xi_{iU}$, so an independent entrant holds brand A if and only if $\xi_{iA} \geq \xi_{iU}$. Let

$$\phi_{db} = h_{db}(\mathbf{0})\, Q_{db}(\mathbf{0})$$

denote the share of the entrants of dealer type d that, at zero wholesale prices, hold a car of brand b and sell it to their household, so that household sales at zero prices are $\Phi_b(\mathbf{0}) = \sum_{d:\, b \in \mathcal{J}_d} \phi_{db} D_d$.

Assumption 2 (Dealers)**.** *The dealer masses are large relative to the inflow:*

$$\phi_{NA} D_N + \phi_{IA} D_I > E_A, \qquad \phi_{IU} D_I + \phi_{BU} D_B > E_U.$$

That is, at zero wholesale prices, the entrants would sell more cars of each brand to households than its inflow, $\Phi_b(\mathbf{0}) > E_b$.

The condition is a lower bound on the dealer masses given the inflow; it holds, for example, if $D_N > E_A/\phi_{NA}$ and $D_B > E_U/\phi_{BU}$, so that the affiliated entrants of each brand alone could sell its inflow at zero prices.

Lemma 2 (Properties of household sales)**.** *Under Assumption 1, the household sales of the two brands, $\Phi(\boldsymbol{w}) = (\Phi_A(\boldsymbol{w}), \Phi_U(\boldsymbol{w}))$, have the following properties.*

(i) Φ is continuously differentiable.

(ii) Each Φ_b strictly decreases in its own wholesale price w_b, through three margins:

- *the intensive margin: every dealer that keeps holding brand b sells with a lower probability;*
- *the exit margin: entrants whose demand falls below the higher cutoff $\bar{\xi}_b(w_b)$ stop holding a car;*
- *the switching margin: some independent dealers switch to the other brand.*

(iii) Each Φ_b is nondecreasing in the other brand's wholesale price, because the only cross effect is the switching of independent dealers. The two brands are therefore gross substitutes.

(iv) Each Φ_b vanishes as its own wholesale price grows, whatever the other wholesale price. Formally, let $\Phi_A(w_A, \infty)$ denote the limit of Φ_A as $w_U \to \infty$. Then $\Phi_A(w_A, w_U) \leq \Phi_A(w_A, \infty)$, and $\Phi_A(w_A, \infty) \to 0$ as $w_A \to \infty$. The same holds for Φ_U, with $\Phi_U(\infty, w_U)$ defined likewise.

Proof of Lemma 2. (i) By Assumption 1, the densities of G_d are continuous, and by Lemma 1 the value V, which determines the cutoffs (10) and (31) and hence the policies j_d, and the sale probability $q(\xi_{ib}, w_b)$ are continuously differentiable. The holding shares and the household sales $\Phi_b(\boldsymbol{w}) = \sum_d D_d h_{db}(\boldsymbol{w}) Q_{db}(\boldsymbol{w})$, which integrate q over the region of demand in which each dealer type holds brand b, are therefore continuously differentiable in the prices.

(ii) Raising w_b lowers the sale probability $q(\xi_{ib}, w_b)$ of every holder of brand b (Appendix F.1). This intensive margin alone makes the decrease strict, since a positive mass of dealers holds brand b at any price. Raising w_b also lowers the net value $V(\xi_{ib}, w_b) - w_b$, whose slope in w_b is $-[1 - \delta + \delta q(\xi_{ib}, w_b)] \leq -(1 - \delta) < 0$ by (9), and this shrinks the set of dealers that hold brand b: the cutoff $\bar{\xi}_b(w_b)$ rises, so fewer dealers hold a car at all (the exit margin), and the boundary $\tilde{\xi}_A$ shifts so that fewer independent dealers prefer brand b to the other brand (the switching margin). Both changes lower Φ_b further.

(iii) Raising the other brand's price $w_{b'}$ lowers the net value of brand b' and leaves the net value of brand b and the sale probability $q(\xi_{ib}, w_b)$ unchanged. It does not affect the affiliated dealers that hold brand b, and it can only enlarge the set of independent dealers that hold brand b, so Φ_b does not fall.

(iv) The function Φ_A is bounded by $D_N + D_I$ and, by (iii), nondecreasing in w_U, so the limit $\Phi_A(w_A, \infty)$ exists and $\Phi_A(w_A, w_U) \leq \Phi_A(w_A, \infty)$. In this limit, no independent dealer holds brand U, because the net value of brand U tends to $-\infty$ for every demand, so $\Phi_A(w_A, \infty)$ integrates $q(\xi_{iA}, w_A)\, \mathbf{1}\{\xi_{iA} \geq \bar{\xi}_A(w_A)\}$ over both dealer types. As $w_A \to \infty$, the net value $V(\xi_{iA}, w_A) - w_A$ tends to $-\infty$ for every demand, because its slope in w_A is at most $-(1 - \delta)$, so this integrand converges to zero pointwise. Dominated convergence then gives $\Phi_A(w_A, \infty) \to 0$. The argument for Φ_U is symmetric. □

Because E_A and E_U are positive, Lemma 2(iv) implies that some price $\overline{w}$ satisfies $\Phi_A(\overline{w}, \infty) < E_A$ and $\Phi_U(\infty, \overline{w}) < E_U$. Above $\overline{w}$, household sales of either brand fall short of its inflow, whatever the other price.

Assumption 3 (Own effects dominate cross effects)**.** *At every* $\boldsymbol{w} \in [0, \overline{w}]^2$*,*

$$-\frac{\partial \Phi_A}{\partial w_A} > \frac{\partial \Phi_U}{\partial w_A}, \qquad -\frac{\partial \Phi_U}{\partial w_U} > \frac{\partial \Phi_A}{\partial w_U} :$$

a rise in one brand's wholesale price reduces household sales of that brand by more than it raises household sales of the other.

Because the cross effects come only from independent dealers that switch brands, whereas the own effects count the same switchers, weighted by their sale probability for the brand they leave, plus the intensive and exit margins, the condition can fail only if switchers sell the brand they move to with a much higher probability than the brand they leave.

Proof of Proposition 1. By (17), a stationary equilibrium is given by prices $\boldsymbol{w} \in [0, \infty)^2$ that solve the clearing conditions $\Phi_A(\boldsymbol{w}) = E_A$ and $\Phi_U(\boldsymbol{w}) = E_U$, together with the policies of Lemma 1, (11), and (13) at these prices. The inflow is positive, and Assumption 2 gives $E_A < \Phi_A(0, 0)$ and $E_U < \Phi_U(0, 0)$.

Every equilibrium lies in the rectangle $[0, \overline{w}]^2$. To see this, suppose that $w_A > \overline{w}$. Then $\Phi_A(\boldsymbol{w}) \leq \Phi_A(w_A, \infty) \leq \Phi_A(\overline{w}, \infty) < E_A$ by Lemma 2(ii) and (iv), so the market for brand A does not clear. The case $w_U > \overline{w}$ is symmetric.

For existence, fix $w_U \in [0, \overline{w}]$. By Lemma 2(iii), $\Phi_A(0, w_U) \geq \Phi_A(0, 0) > E_A$. By the choice of $\overline{w}$, $\Phi_A(\overline{w}, w_U) \leq \Phi_A(\overline{w}, \infty) < E_A$. Because $\Phi_A(\cdot, w_U)$ is continuous and strictly decreasing, the equation $\Phi_A(\hat{w}_A(w_U), w_U) = E_A$ has a unique solution $\hat{w}_A(w_U) \in (0, \overline{w})$, which is continuous in w_U. The map $w_U \mapsto \Phi_U(\hat{w}_A(w_U), w_U)$ is then continuous. At $w_U = 0$ it is at least $\Phi_U(0, 0) > E_U$, because $\hat{w}_A(0) \geq 0$ and the brands are gross substitutes, and at $w_U = \overline{w}$ it is at most $\Phi_U(\infty, \overline{w}) < E_U$. By the intermediate value theorem, $\Phi_U(\hat{w}_A(w_U^*), w_U^*) = E_U$ for some $w_U^* \in (0, \overline{w})$, and $(\hat{w}_A(w_U^*), w_U^*)$ is an equilibrium.

For uniqueness, suppose that $\boldsymbol{w} \neq \boldsymbol{w}'$ are both equilibria. Because Φ is continuously differentiable and the segment between the two points lies in the rectangle, $\Phi(\boldsymbol{w}') - \Phi(\boldsymbol{w}) = \bar{J}(\boldsymbol{w}' - \boldsymbol{w})$, where $\bar{J} = \int_0^1 J(\boldsymbol{w} + t(\boldsymbol{w}' - \boldsymbol{w}))\, \mathrm{d}t$ averages the Jacobian $J = \partial \Phi / \partial \boldsymbol{w}'$ along the segment. Assumption 3 states that J is strictly diagonally dominant by columns. The integration preserves this property, so $\bar{J}$ is nonsingular by the Levy–Desplanques theorem. Since $\Phi(\boldsymbol{w}') = \Phi(\boldsymbol{w}) = (E_A, E_U)'$, the equation $\bar{J}(\boldsymbol{w}' - \boldsymbol{w}) = 0$ implies $\boldsymbol{w}' = \boldsymbol{w}$, a contradiction. □

Extension with a holding cost. The proof uses only Assumptions 2 and 3 and the four properties of Lemma 2, so Proposition 1 holds for the household sales of Section 6.1, with the mean sale probabilities (22), once Lemma 2 is verified for them. By Appendix F.4, $\overline{V}_d$ is continuously differentiable with $\partial \overline{V}_d / \partial w_b = \delta(1 - \bar{q}_d)$, so the net value $\overline{V}_d(\xi_{ib}; w_b) - w_b$ has slope $-[1 - \delta(1 - \bar{q}_d)] \leq -(1 - \delta) < 0$ in w_b, and $\bar{q}_d$ is continuously differentiable with $\partial \bar{q}_d / \partial w_b < 0$, because the sale probability $q = q(\xi_{ib}, c; w_b)$ has $\partial q / \partial w_b = -\delta q(1 - q)^2 / \sigma_\epsilon < 0$ at interior solutions, which have positive probability under H_d, and $\partial q / \partial w_b = 0$ at boundary solutions. The arguments

for (i)–(iv) then apply with $\overline{V}_d$, $\bar{q}_d$, and the cutoffs defined from $\overline{V}_d$ in place of V, $q(\xi_{ib}, w_b)$, and the cutoffs of Section 5.3.

F.4 Retail Pricing with a Holding Cost

Fix (ξ_{ib}, c, w_b) and write $L = \delta w_b - c$ for the liquidation value, which can be negative, and $q(p)$ for the probability that the car sells at price p, as in Appendix F.1. The objective of problem (19) is $\Pi(p) = L + q(p)(p - L)$, as in Appendix F.1 with L in place of δw_b, so $\Pi'(p)$ changes sign exactly once, from positive to negative, at the unique solution $\hat{p} > L$ of $(p - L)[1 - q(p)] = \sigma_\epsilon$. On the feasible set $p \geq 0$, the solution is $p^* = \hat{p}$ if $\hat{p} \geq 0$, and $p^* = 0$ if $\hat{p} < 0$, which requires $L < 0$; the maximizer is unique in both cases. Write $q = q(p^*)$ for the sale probability $q(\xi_{ib}, c; w_b)$ at the solution. At an interior solution, $(p^* - L)(1 - q) = \sigma_\epsilon$, which is the first equation in (20) because $1-q(p) = 1/[1+\exp((\xi_{ib}-p)/\sigma_\epsilon)]$, and substituting it into the objective gives $V = L+\sigma_\epsilon q/(1-q) = p^* - \sigma_\epsilon$. The envelope theorem gives $\partial V/\partial \xi_{ib} = q$ and $\partial V/\partial L = 1 - q$, so $\partial V/\partial w_b = \delta(1 - q)$ and $\partial V/\partial c = -(1 - q)$; because $p^* = V + \sigma_\epsilon$, the price has the same derivatives, and in particular $\partial p^*/\partial c = -(1 - q) < 0$. Differentiating $q = \exp((\xi_{ib} - p^*)/\sigma_\epsilon)/[1 + \exp((\xi_{ib} - p^*)/\sigma_\epsilon)]$ then gives $\partial q/\partial c = q(1 - q)^2/\sigma_\epsilon > 0$, $\partial q/\partial \xi_{ib} = q(1 - q)^2/\sigma_\epsilon > 0$, and $\partial q/\partial w_b = -\delta q(1 - q)^2/\sigma_\epsilon < 0$. At a boundary solution, $V = [1 - q(0)]L$ with $q = q(0)$, so $\partial V/\partial w_b = \delta(1 - q)$ and $\partial V/\partial c = -(1 - q)$ as before, $\partial V/\partial \xi_{ib} = -L\, q(0)[1 - q(0)]/\sigma_\epsilon > 0$, and q does not depend on w_b or c; the two cases agree where the solution switches, because $(-L)[1 - q(0)] = \sigma_\epsilon$ there. Hence V is continuously differentiable and strictly increasing in ξ_{ib}, and by dominated convergence so is the expected value (21), with $\partial \overline{V}_d/\partial w_b = \delta(1-\bar{q}_d)$. As $\xi_{ib} \to -\infty$, $V \to L$ and $\overline{V}_d \to \delta w_b - \int c \, \mathrm{d}H_d(c)$, and as $\xi_{ib} \to \infty$, $\overline{V}_d \to \infty$, so the cutoff $\overline{V}_d^{-1}(w_b; w_b)$ of Section 6.1 exists whenever $\int c \, \mathrm{d}H_d(c) > -(1 - \delta)w_b$, in particular whenever the mean holding cost is positive.