PhD THESIS DECLARATION

I, the undersigned

FAMILY NAME	Schmidt
NAME	Peter Robert
Student ID no.	1824366

Thesis title:

Essays in Empirical Industrial Organization

PhD in	Economics
Cycle	30th
Student's Advisor	Prof. Jérôme Adda
Calendar year of thesis defence	2019

DECLARE

under my responsibility:

1) that, according to Italian Republic Presidential Decree no. 445, 28th December 2000, mendacious declarations, falsifying records and the use of false records are punishable under the Italian penal code and related special laws. Should any of the above prove true, all benefits included in this declaration and those of the temporary "embargo" are automatically forfeited from the beginning;

2) that the University has the obligation, according to art. 6, par. 11, Ministerial Decree no. 224, 30th April 1999, to keep a copy of the thesis on deposit at the "Biblioteche Nazionali Centrali" (Italian National Libraries) in Rome and Florence, where consultation will be permitted, unless there is a temporary "embargo" protecting the rights of external bodies and the industrial/commercial exploitation of the thesis;

3) that the Bocconi Library will file the thesis in its "Archivio Istituzionale ad Accesso Aperto" (Institutional Registry) which permits online consultation of the complete text (except in cases of temporary "embargo");

4) that, in order to file the thesis at the Bocconi Library, the University requires that the thesis be submitted online by the student in unalterable format to Società NORMADEC (acting on behalf of the University), and that NORMADEC will indicate in each footnote the following information:

- PhD thesis Essays in Empirical industrial Organization;
- by; Peter Robert Schmidt
- defended at Università Commerciale "Luigi Bocconi" – Milano in the year 2019;

- the thesis is protected by the regulations governing copyright (Italian law no. 633, 22nd April 1941 and subsequent modifications). The exception is the right of Università Commerciale "Luigi Bocconi" to reproduce the same, quoting the source, for research and teaching purposes;

5) that the copy of the thesis submitted online to Normadec is identical to the copies handed in/sent to the members of the Thesis Board and to any other paper or digital copy deposited at the University offices, and, as a consequence, the University is absolved from any responsibility regarding errors, inaccuracy or omissions in the contents of the thesis;

6) that the contents and organization of the thesis is an original work carried out by the undersigned and does not in any way compromise the rights of third parties (Italian law, no. 633, 22nd April 1941 and subsequent integrations and modifications), including those regarding security of personal details; therefore the University is in any case absolved from any responsibility whatsoever, civil, administrative or penal, and shall be exempt from any requests or claims from third parties;

7) that the thesis is not subject to "embargo", i.e. that it is not the result of work included in the regulations governing industrial property; it was not written as part of a project financed by public or private bodies with restrictions on the diffusion of the results; is not subject to patent or protection registrations.

Date 10.12.2018

Abstract

This thesis consists of three chapters on the Sharing Economy, its impact on users and existing industries, and the car market.

Sharing Economy online platforms enable peer-to-peer goods rental and service provision at large scale, offering potentially large efficiency gains. Using new data from the largest peer-to-peer car sharing platform in the U.S., in the first chapter I show that gender-specific differences in price setting are driven by gender homophily. Due to female renters' concerns for physical safety during the exchange, they prefer renting from female owners and owners with experience on the platform. Experienced female car owners hence face increased demand and charge a price premium of 13%, compared to men with similar cars in similar locations. Inexperienced female car owners charge 4% less than men and acquire a track record on the platform more quickly. Counterfactuals from a dynamic model of price setting show that requiring experienced female owners to lower their prices, while compensating them for lost income, improves outcomes for renters and female car owners, and is profitable for the platform.

In the second chapter, I examine the effect of a new type of car sharing on new car sales. Using an original administrative dataset on car registrations and a proprietary dataset provided by one of the Free-Floating Car Sharing companies, the staggered rollout of Free-Floating Car Sharing in German cities is exploited with a difference-in-difference methodology to estimate its impact on new car sales. One car sharing vehicle reduces annual new car sales by between three vehicles. This effect is driven by a reduction in sales of small, compact and medium-sized models and it is stronger when more women and older users use car sharing.

In the third chapter, I exploit the exogenous variation provided by the Volkswagen emissions scandal (Dieselgate) to identify how consumers value environmental friendliness of automobiles. The emissions scandal allows me to disentangle preferences for fuel efficiency from preferences for environmental friendliness. This paper documents the effects of the Dieselgate shock on the new and used car market and proposes a structural model to evaluate counterfactuals.

Acknowledgements

I thank my advisor, Jérôme Adda, as well as Greg Crawford, Fabiano Schivardi and Carlo Favero for their invaluable academic support. I am very grateful to my parents, who enabled me to pursue my studies in the best possible way. Finally, a special thank you to Zornitsa, who not only helped me complete the PhD, but made me do it with a smile on my face.

Chapter 1

Price Setting on Peer-to-Peer Sharing Platforms

1.1 Introduction

Sharing Economy online platforms enable peer-to-peer rental of durable goods and provision of services at large scale. The platforms take advantage of increasing smartphone penetration, improvements in online platform design and a growing acceptance of e-commerce in the population to match private durable goods owners and service providers with potential renters and customers[1]. Among the best known examples of Sharing Economy platforms are *Airbnb* (apartments), *Uber* (rides), *Turo* (cars) and *TaskRabbit* (freelance labor). The platforms have grown quickly and are now major players in their respective industries. For example, *Airbnb* has more room listings on its website than the five largest hotel chains combined[2]. *Uber* had revenues of $6.5 billion in 2016[3]. Its success has spurred substantial interest in the literature about the Sharing Economy as a new way of economic organization[4].

As privately owned goods are often underutilized and many individuals would like to work

[1]Peer-to-peer platforms for services are sometimes collectively referred to as the 'Gig Economy'.

[2]https://thespaces.com/airbnb-now-bigger-worlds-top-five-hotel-brands-put-together/

[3]http://fortune.com/2017/04/14/uber-revenue-2016/

[4]Sundararajan (2016) proclaims the end of economic allocation as we know it and the 'rise of crowd based capitalism'.

more hours, the Sharing Economy offers potentially enormous efficiency gains[5]. An emerging economic literature shows that the extent to which platform users realize these monetary gains depends on their personal characteristics, such as race and gender. An important question is whether these differences in price setting (on durable goods sharing platforms) and wages (on service platforms) are due to freely made decisions or due to discrimination, as the latter would warrant regulatory intervention.

In this paper, I use a new panel dataset of around 300 000 car listings by 40 000 unique owners to provide new empirical evidence towards the determinants of gender-specific differences in price setting on Sharing Economy platforms. I collected the dataset by scraping listings of privately owned vehicles available for rent on *Turo*, the largest peer-to-peer car sharing platform in the U.S., on a daily basis for over six months. The data consists of information about the vehicle, its location, the owner and previous renters. This dataset is uniquely suited to analyze the role of gender in the Sharing Economy, as it allows the econometrician to include virtually all the information available to platform participants in the analysis, alleviating concerns about biasing unobserved variables[6]. Furthermore, since car owners are observed at different points in time, their evolving behavior as they gain experience on the platform can be tracked.

I find that gender-specific differences in price setting on *Turo* are driven by preferences of female car renters. Female renters prefer to rent from female car owners and from owners with a long track record on the platform[7]. Consequently, experienced female car owners face increased demand and raise their prices by up to 13%, controlling for car, owner and location characteristics. Anticipating the ability to charge a premium, female car owners new to the platform lower their prices by around 6% to raise demand for their cars and accumulate trips more quickly. These differences in price setting across genders and experience levels are driven by changing behavior over time, not by time-invariant unobserved differences between own-

[5]An often cited statistic states that privately owned cars are idle 95% of the time (Shoup (2017)).

[6]Cars are mostly characterized by their model and age, whereas apartments, for example, can differ significantly even when controlling for their size and location. A potential renter (for example on *Airbnb*) is able to see these differences in the pictures the apartment owner uploaded, while there is no straightforward way for the econometrician to incorporate this information in the analysis, creating potentially severe omitted variable bias. Similarly, it is difficult to control for service quality on 'Gig Economy' platforms such as *Uber*, while this aspect is less important on a durable goods sharing platform such as *Turo*.

[7]As measured by the number of previously completed transactions (from now on referred to as trips).

ers with few and many trips. The price differences are more pronounced in areas with high rates of physical violence, indicating that the gender homophily among female car renters is driven by concerns for their physical safety[8]. Attitudes towards the role of a woman in society or the gender wage gap, which is used as a proxy for discrimination, are not able to explain local differences in the gender-specific pricing pattern.

While the price premium charged by experienced female car owners is not a case of market failure, but reflects demand side preferences, it in effect forces cautious female renters, who will only rent from experienced female car owners, to pay a surcharge. Sharing Economy platforms do have trust-building features in place, but the results in this paper show them to be insufficient to alleviate safety concerns by female renters[9]. The extent and scrutiny of the safety features in place depend on the platform-specific tradeoff between efficiency and user safety. More stringent safety requirements make the platform safer to use, but inhibit its convenience[10].

The second contribution of this paper is to evaluate how a change in the pricing system can improve user outcomes on Sharing Economy platforms in the presence of gender homophily on the demand side. To this end, I develop a dynamic single-agent model of price setting on peer-to-peer sharing platforms. The model is applied to the car sharing context and calibrated using data from *Turo*. The calibrated model serves two purposes. First, I show that the gender-specific pricing pattern documented on *Turo* can be explained by dynamic behavior of the car owners. Second, using counterfactual calculations, I evaluate a change to the platform's pricing system, under which experienced female car owners are required to keep their prices low, but are compensated for lost income. This intervention is shown to improve outcomes for car renters and female car owners, as well as to be profitable for the platform itself.

[8]Women may be more at risk when using Sharing Economy platforms due to the potential for sexual harrassment or assault. As of 2018, the two largest platforms *Airbnb* (https://www.theguardian.com/technology/2017/jul/27/airbnb-guest-sexual-assault-allegation) and *Uber* (http://money.cnn.com/2018/04/30/technology/uber-driver-sexual-assault/index.html) face lawsuits over alleged sexual assaults on female platform users.

[9]The main feature is a review and rating system, with which users may publicly evaluate the experiences they have had during a transaction with another user. Furthermore, some platforms allow users to link their platform profile to their social media accounts or require in-person verification or even criminal background checks.

[10]How to strike the right balance between safety and efficiency on online platforms is an open question and active area of research (Athey & Luca (2018))

These findings have more general implications for Sharing Economy platforms. The empirical evidence provided by this paper shows that gender homophily of female renters is an important characteristic of the demand side. Taking this as given, a change to the platform design, in this case to the pricing mechanism, can improve user outcomes. By making the listings of users from underrepresented groups more attractive, platforms can increase their usefulness to demand side members of the same group, and even increase participation and traffic on the platform. Proponents of the Sharing Economy highlight that it is more inclusive and equitable than existing allocation systems, because it enables the less well-off to access high quality goods at cheaper prices and allows those unable to work during regular business hours more flexibility. For the Sharing Economy to become a truly more equitable and inclusive economic alternative, all participants of the economy have to be able to use it. In particular, groups of the population that were disadvantaged or discriminated against in the regular economy cannot suffer the same fate in the Sharing Economy, if its platforms are to live up to their expectations.

This paper is structured as follows: The rest of this section introduces related literature and connects this paper to it. Section 1.2 introduces the *Turo* online platform and the different data sources used in the paper, Section 1.3 contains the reduced form empirical analysis and results, Section 1.4 develops and applies the model and evaluates the counterfactuals. Section 1.5 concludes.

Related Literature

Einav et al. (2016) provide an introduction to Sharing Economy business models, their economic implications and open questions on the topic. A number of papers provide empirical evidence about consumer behavior on peer-to-peer sharing platforms. Fraiberger & Sundararajan (2017) use company data from the hourly peer-to-peer car sharing platform *getaround*. They show theoretically that giving car owners the possibility to rent out their vehicles when they would otherwise sit idle, provides an additional incentive to purchase a car. This prediction is tested and validated with data from New York City. Their paper is concerned with the decision making at the extensive margin, i.e. whether consumers buy a car and participate in car sharing, whereas this paper looks at the intensive margin, i.e. once a consumer has a car and is participating in car sharing, what price should he or she set. Benjafaar et al. (2018) provide a theoretical framework for the effects of peer-to-peer sharing of durable goods on ownership, usage and welfare. A key finding is that consumers always benefit from the presence of sharing platforms, but the extent to the welfare gains depends heavily on usage patterns and demographics. Weber (2016) develops a theoretical model that jointly explains purchase decisions and price setting for durable goods in the presence of a peer-to-peer rental market. A large part of the existing literature on the Sharing Economy analyszes its impact on existing industries. Manco & Percoco (2018) show that *Uber* is not competing directly with taxi operators and its presence may even generate positive spillover effects. Similarly, Cramer & Krueger (2016) analyze the effect of *Uber* availability on taxi usage in New York City, Farronato & Fradkin (2017) examine the impact of *Airbnb* on hotel room occupancy and Seamans & Zhu (2014) quantify the effect of *Craigslist* on local newspaper ads. Gong et al. (2017) investigate the effect of the presence of *Uber* on durable good purchase decisions. Horton & Zeckhauser provide a framework for the joint determination of durable good ownership and rental in the Sharing Economy in general.

Close to this work is a series of papers looking at racial discrimination on *Airbnb*, the largest apartment sharing website. Edelman & Luca (2014), as well as Laouénan & Rathelot (2017) document that hosts from racial minorities charge less than white hosts for similar apartments. Both papers control for a large number of observables and conclude that minority hosts face discrimination and are forced to lower their prices in order to increase demand. Edelman et al. (2017) examine the demand side of the *Airbnb* marketplace. The authors conduct an experiment, during which they send out booking requests from artificially created accounts of users with different ethnicities. They find that booking requests from black users are 12% less likely to be successful, controlling for observables. Another peer-to-peer platform that attracted the attention of researchers is the long-distance ride sharing platform *Blablacar.* Farajallah et al. (2016) examine the dynamics of price setting behavior of drivers on the platform and, similarly as Lambin et al. (2017), find that drivers with Arabic-sounding names face discrimination by potential riders.

Although the Sharing (or Gig) Economy has often been heralded as a great equalizer for earnings of men and women, because of flexible hours and transparent remuneration schemes, there is little research on whether this prediction has actually come true. Cook et al. (2018) document and then unpack the substantial difference in hourly earnings between male and female *Uber* drivers. They use trip-level data to show that the gap is due to differential behavior across genders and not discrimination. Barzilay & Ben-David (2017) document a gender wage gap on a freelance labor online platform, even when they control for task and service provider observable characteristics. While this paper is the first to show that gender-specific pricing differences on peer-to-peer platforms are driven by gender homophily on the demand side, Zeltzer (2017) documents similar gender homophily in physician referral networks.

Closest to this paper is Bohren et al. (2018). The authors document belief-based discrimination towards women in an online forum for math problems. Their data show a reversal in payoffs which is very similar to what this paper finds; inexperienced female users on the forum receive fewer upvotes (a non-monetary currency on a website), than male users for content of the same quality. However, experienced female users receive more upvotes than their male peers for content of the same quality. Their setting however is purely internet-based, so that physical safety concerns cannot play a role. Even though the data used in this paper exhibits a similar reversal in payoffs for female users as they gain experience, I show that discrimination is not the cause of the reversal in this paper.

Finally, a small literature document price setting determinants on peer-to-peer sharing platforms. Consistent with the findings of this paper, particular importance is attributed to reputational signals. Teubner et al. (2017) quantify the price effect of favorable ratings on *Airbnb*. Gutt & Herrmann (2015) provide causal evidence that *Airbnb* hosts increase their prices after their previous ratings become publicly visible. As mentioned above, an important advantage of using data from a car sharing platform, such as *Turo*, over data on apartments (e.g. from *Airbnb*), is that by including the car observables in the analysis, the researcher is much more likely to control for all potential confounders.

1.2 The Data

1.2.1 Car rental listings on *Turo*

Turo is the largest daily peer-to-peer car sharing platform in the U.S. On the *Turo* website, private individuals may list their vehicles to be rented by other private individuals for a minimum duration of one day. According to *Turo*, the platform is active in more than 4 500 cities across the U.S., its users offer more than 800 different models and have average earnings of $720 per month. The platform claims that its prices are are on average 35% lower than traditional rental car agency prices.

The dataset used in the empirical part of this paper was obtained by scraping the *Turo* website daily from February to August 2018. One cross-sectional observation in the dataset corresponds to a listing of a car offered for rent on the *Turo* website[11]. The observed variables are model and model year of the listed car, daily price, included mileage, number of previous trips, average rating for those previous trips, the location of the car, the name of the owner, his description of the car, whether the 'Instant Booking' feature is enabled and the names and reviews of previous renters.

Table 1.1 shows summary statistics for the cross-section of listings. The majority of owners (83.7%) and renters (71%) on the platform are male. On average, they charge higher prices, have fewer previous trips, older cars, better ratings and fewer female renters than female car owners on the website. Finally, female owners are slightly more likely to have activated the 'Instant Booking' feature. Car owners who activate this feature forfeit the opportunity to screen potential applicants to rent their car, but must accept any renter.

Turo suggests a daily rental price to car owners, which is based on seasonality, geography, car characteristics and local competition, but crucially does not depend on the gender of the

[11]For details about the website and how the datast was constructed, please see the Appendix

owner[12]. Discrepancies in pricing across gender must therefore be due to differential deviations from the suggested price by men and women.

Listings by female car owners have lower prices in part because they tend to have smaller, cheaper cars than men. Table 1.2 shows the fraction of different car segments in the sample, within each segment the fraction of female car owners and the median price for listings with at least one trip in each segment. The table shows that female car owners are under-represented in the categories ”Large Sedan” and ”Electric”, which comprise the most expensive models[13]. There are relatively more female owners in the ”Small SUV” and ”Hybrid” segments, but models in these segments command lower prices. This illustrates that controlling for car characteristics in the empirical analysis is strictly necessary.

The distribution of the number of previous trips (and hence experience on the platform) is highly skewed. Figure 1.4 shows the histogram for the number of previous trips. The majority of listings has no or very few previous trips. The histogram for daily rental price is similarly skewed (Figure 1.5), although there is a noticeable number of very expensive vehicles (500$ or more per day).

Figure 1.6 offers a first look at the reversal in pricing behavior between male and female car owners as they gain experience on the platform. For low numbers of previous trips, male owners charge significantly higher prices, while among car owners with a large number of previous trips, this is reversed. Figure 1.6 does not control for car, location or other owner characteristics besides experience. Nonetheless, the pattern it displays will be confirmed by the more rigorous analysis in section 1.3.

1.2.2 Local Crime Data

To show that gender homophily on the demand side of *Turo* is driven by concerns for physical safety, I use local crime data. This data is available from the UCR datatool (U.S. Department of Justice (2014)). The number of rape cases in 2014 (the latest available year) per 100 000

[12] This was confirmed by a *Turo* representative. See also the description on the company website: https://blog.turo.com/news/using-data-to-drive-your-daily-price (Accessed 05.11.2018)

[13] Not only in terms of daily rental prices on *Turo*, but also in terms of list purchase prices. The category ”Electric”, for example, consists almost exclusively of expensive *Tesla* cars.

inhabitants of a city is used as a proxy for how unsafe a particular area is. This statistic is available for 83 U.S. cities and it is matched to *Turo* listings using the combination of city and state name. The crime data is more likely to be available for urban areas than rural ones. Summary statistics for the subsample where local crime data is available are in Table 1.3. The variables mostly exhibit the same differences across men and women as in the full sample.

1.2.3 Attitudes towards Women

Another channel potentially driving differential price setting behavior across genders are differential attitudes towards women across different geographical locations. To investigate this channel, data from the General Social Survey Wave 2016 (Smith et al. 2016) are used. This survey identifies the attitudes of respondents towards the role of women in society. The survey classifies respondents' places of residence into nine geographical regions in the U.S., so that differences across localities can be identified. It is therefore possible to match all *Turo* listings to the average survey response in the corresponding geographical area. The survey question used here asks, to which degree respondents agree with the following statement: "It is much better for everyone involved if the man is the achiever outside the home and the woman takes care of the home and family." The possible answers are 1 (Strongly Agree), 2 (Agree), 3 (Disagree) and 4 (Strongly Disagree).

1.2.4 Gender Wage Gap

Data on the gender wage gap by U.S. state from 2016 is provided by the National Women's Law Center (NWLC (2018)). The gender wage gap is used as a proxy for discrimination against women. It consists of the fraction of average female to male wages in percent. All *Turo* listings can be matched to the gender wage gap data using the state for which the listing was posted.

1.3 Reduced Form Strategy and Results

1.3.1 Renter Preferences

Figure 1.7 shows the proportion of previous renters who are female, for male and female car owners with different numbers of previous trips on *Turo*. Two observations on the demand side of *Turo* can be made: (i) Female renters prefer to rent from female owners and (ii), female renters prefer to rent from experienced owners of either gender. Taken together, these observations imply that experienced female owners on *Turo* face increased demand, as there are women among the population of potential renters who will only rent from experienced female car owners.

This homophily is not observed on the supply side, as female owners do not prefer to rent to female rather than male renters. The summary statistics (Table 1.1) show that female than male owners are equally likely to activate the ”Instant Booking” feature. Activating this feature removes the option for car owners to screen applicants before accepting or rejecting their rental request. Female owners are no more likely to use this feature than males, implying that they are on average as indifferent towards who rents their cars as men are.

1.3.2 Owner Behavior

Cross Sectional Analysis

The data section illustrated that observable characteristics of car listings on *Turo* differ between male and female car owners. Equation (1.1) shows how the observables are controlled for in an OLS regression setup, so that differing prices can be attributed to different behavior of male and female car owners.

$$log(p_{it}) = \lambda_c + \lambda_{sm} + \lambda_t + \beta_1 \cdot \mathbb{1}\{\text{i is female}\} + \beta_2 \cdot \mathbb{1}\{\text{i is female}\} \cdot \text{trips}_{it} + \gamma \cdot X_{it} + \epsilon_{it} \quad (1.1)$$

The subscript i stands for owner, c for location, s for car model, m for model year and t for the time of the listing. The dependent variable is the (log of) the daily price for a listing.

The λ terms denote fixed effects; λ_c and λ_t denote location and time fixed effects, respectively. λ_{sm} are car model-year interaction fixed effects[14]. The coefficients of interest are β_1 and β_2. β_1 multiplies a dummy variable equal to one for listings by a female car owner, while β_2 is the coefficient of the interaction term between the dummy for female and the number of previous trips a listing has. Together, these two coefficients show the differences in price setting between male and female car owners, and how the difference changes, as the owners accumulate more trips. Finally, X is a vector of observed control variables that may vary over time, which are the number of previous trips, the average rating and the allowed driving distance. ϵ denotes the error term.

The estimation results for equation (1.1) are in Table 1.4. Female owners charge lower prices than men when they have few previous trips and higher prices when they have many. The magnitude of the differences in pricing shrinks after including car model-year fixed effects, as men and women offer different cars on *Turo*. The number of trips for which there is no difference in prices is around 22.

A different way to show the results is to split the sample by the number of previous trips the listings have, as in Table (1.5). Using the entire sample, women charge 2.6% less for their vehicles, controlling for location, time, listing and vehicle characteristics. This result conceals large heterogeneity in the price difference between male and female car owners across different levels of experience on the platform. Columns (2) and (3) show that listings by female owners with few previous trips are less expensive than comparable listings by male owners; the results in columns (6) and (7) however demonstrate how this pattern is reversed for car owners with many previous trips. Female car owners with previous trips above the 75th percentile charge 6% more than their male counterparts (column (6); the premium increases to 13% for the most experienced female owners, who are above the 90th percentile in terms of previous trips (column(7)).

The differences in earnings due to the differing prices are substantive. Using the average daily rental price for car owners above the 90th percentile in terms of previous trips ($55), the es-

[14]For example, a *Fiat 500* built in 2014 receives a different dummy variable than a *Fiat 500* built in 2015

timated premium equals \$7.15 per day. Taking into account the high usage rate prevailing among these frequently rented cars, the differences in earnings are economically significant.

An explanation for why inexperienced female car owners charge lower prices than comparable men is that they anticipate the ability to charge a price premium once they have accumulated a track record on the platform. Therefore they lower prices in order to increase demand and accumulate trips quickly, so that they will perceived as an experienced platform user and be in a position to charge a price premium[15].

For the 50% of listings in the middle of the experience distribution, there is no statistically significant difference in price setting between male and female car owners (columns (4) and (5) of table 1.5).

Panel Analysis

The previous subsection showed that experienced female car owners on *Turo* charge higher prices, while inexperienced females charge lower prices than men with similar observable characteristics. This result could potentially be driven by selection; those female owners who managed to accumulate many trips may offer cars and service of particularly high quality, enabling them to charge a premium over experienced men. Female owners with few trips on the other hand have predominantly inferior offers and are hence not able to accumulate many trips and have to lower their prices. According to this explanation, the observed reversal in pricing of female owners compared to male owners is due to the fact that females with many previous trips differ from those with few trips in terms of the true quality of their listings, which is not observable by the econometrician.

The data does not support a selection based explanation of the pricing pattern. All car owners have to start with zero previous trips on the platform. If only the female owners with high quality listings manage to accumulate trips, the proportion of women among all car owners should decrease in the number of previous trips. Figure 1.9 shows that this is not the case; the number of female car owners is constant or throughout the distribution of previous trips on

[15] Figure 1.8 confirms that female car owners with few trips accumulate trips at a faster rate than their male counterparts.

the platform and even increases slightly among the very experienced owners.

Since the same car owners are observed multiple times and at different experience levels, it is possible to rule out a selection based explanation using panel regressions. For this purpose, the dataset is collapsed so that a unique observation is given by an owner-trip combination. Summary statistics for the panel dataset are in Table 1.6. Using the panel allows for the inclusion of owner fixed effects in a regression of the following form:

$$log(p_{it}) = \lambda_i + \beta \cdot \mathbb{1}\{\text{i is female}\} \cdot \{\text{trips}\} + \gamma \cdot X_{it} + \epsilon_{it} \tag{1.2}$$

where λ_i denotes owner fixed effects and the rest of the notation and the control variables are as above. Including owner fixed effects in the regression is crucial, as it controls for potentially unobserved differences in quality across owners, which were discussed above. The variation in the data used to identify the coefficient of interest, β, comes from within-owner changes in price setting over time, not from differences across owners, which are controlled for by the owner fixed effects. The estimation results are in Table 1.7. The positive and significant coefficient on the interaction term indicates that female owners increase their prices as they accumulate trips, while controlling for all time-invariant personal characteristics. The coefficients in column (2), which includes the squared number of previous trips, suggest that a given female car owner will increase her price by around 7% after accumulating 20 trips. This result confirms that the differential price setting behavior across genders is driven by changing behavior as owners accumulate trips, and not by a selection process.

Another way to see illustrate how changing behavior as owners gain experience causes the pricing pattern is given by Figure 1.10. It shows how car owners changed their prices after gaining another trip on *Turo*, depending on how many trips they had completed before. For car owners that had previous between one and twenty previous trips, there is no systematic difference in the price reaction to completing another trip. Both male and female car owners increase their price upon completion of a trip. However, completing trips number 21 and higher leads to differential price adjustments of men and women; while male car owners increase their prices less and even converge to keeping them the same after another trip, female car owners keep increasing their prices by between \$1 and ¢40 after each trip until trip number 80. Further trips don't seem to allow car owners to increase their prices any further.

1.3.3 The Mechanism

Crime and Safety

Section 1.3 so far has documented that experienced female car owners are charging higher prices than experienced male car owners, controlling for car, location and owner characteristics, while the opposite is true among inexperienced car owners. This pattern is due to the presence of female renters, who are willing to pay a higher price to rent from female car owners with long track records on *Turo*. This gender homophily on the part of female renters may be driven by concerns for safety during the transaction, as female renters assign a higher probability of being safe renting from a female car owner, than from a male car owner. Being perceived as a safe transaction partner should then be even more valuable in geographical areas with high levels of crime. Put differently, if it is safety concerns that drive the increased demand for experienced female owner's cars, their price premium should be higher and their dynamic behavior more pronounced in unsafe areas.

To test these predictions, I use data from the Uniform Crime Reporting Database. This data measures violent crime and rape rates in 83 American cities [16]. To examine how local safety conditions interact with gender-specific pricing behavior, I introduce a triple-interaction term between the dummy for female, the number of trips and the local crime rate. The regression setup is given by Equation (1.3).

$$\begin{aligned} lg(p_{it}) = \lambda_c + \lambda_{sm} + \lambda_t + \beta_1 \cdot \mathbb{1}\{\text{i is female}\} + \beta_2 \cdot \mathbb{1}\{\text{i is female}\} \cdot \text{trips}_{it} + \\ \beta_3 \cdot \mathbb{1}\{\text{i is female}\} \cdot \text{trips}_{it} \cdot \text{crime\_rate}_c + \gamma \cdot X_{it} + \epsilon_{it} \end{aligned} \tag{1.3}$$

where notation and control variables are as above. The results are in Table 1.8. Column (1)

[16] See Section 1.2.2 for a descrption of the data. The crime rate used here is number of rape cases per 100 000 inhabitants.

reproduces the familiar pricing pattern for the subset of the *Turo* listings that can be matched to the UCR local crime data, with lower predicted prices for female owners with no previous trips, but a positive coefficient on the interaction of the female dummy with the number of trips. Statistical significance is reduced due to the much smaller sample size. Including an interaction of the dummy for female with the crime statistic and a triple interaction term between the dummy for female, the crime statistic and the number of trips in column (3) indicates that local safety does have an impact on price setting behavior. When including the quadratic number of trips and its interaction with the female dummy and the crime rate in column (4), significance levels increase, indicating important non-linear relationships. Since the interaction terms are easier to interpret grapically, Figure 1.11 shows the predicted price difference between female and male car owners as a function of previous trips and local crime rates in percent, according to the parameter estimates of column (4) in Table 1.8.

For low levels of crime, Figure 1.11 shows that the relationship between differences in price and the number of previous trips is as before: inexperienced female car owners charge lower prices than their male counterparts, while experienced females charge a substantial premium over comparable men. The number of trips where prices are the same across genders, indicated by the white part of the surface in Figure 1.11 is around 20. As the local crime rate increases, the number of trips for which prices are the same for men and women increases as well, indicating that women keep their prices low for longer in unsafe areas. The motivation for doing so lies in the ultimately higher payoff for experienced female car owners in unsafe areas; Figure 1.11 shows that among car owners with high levels of experience, the price premium charged by women is larger, the more unsafe the area is.

Other Mechanisms

Attitudes towards a woman's role in society, as well as outright gender discrimination are alternative mechanisms potentially driving the documented homophily among female car renters and owners on *Turo*. Using regionally varying proxies for both, attitude and discrimination, this subsection documents that neither influences price setting on *Turo* in a significant way. For the attitude towards women's roles in society, the proxy is given by the response to the survey statement: "It is much better for everyone involved if the man is the achiever outside the home and the woman takes care of the home and family.". This question was asked in the General Social Survey 2016 (Smith et al. (2016)). To proxy for gender discrimination, I use data on the gender wage gap by U.S. state in 2016 provided by the National Women's Law Center (NWLC (2018))[17].

Tables 1.9 and 1.10 show the estimation results of regressions of equation (1.3), where the crime rate variable is replaced by the respective proxy variables for attitude and discrimination. Column (1) of both tables reproduces the main result: inexperienced female car owners charge lower prices than men while experienced female car owners charge a premium, controlling for the same observables as before. Columns (2) to (4) of both tables include interaction terms with the proxy variable for attitude and the gender wage gap. Neither variable interacts significantly with the dummy for female, or its interaction with the number of trips. This is taken as evidence that neither attitudes towards women's roles in society nor gender discrimination (as proxied for by the gender wage gap) have any explanatory power towards the differential pricing behavior of men and women on *Turo*.

1.4 The Model

This section develops a single-agent dynamic model of price setting on peer-to-peer sharing platforms for durable goods. The payoff for the goods owner depends on three factors; first, a personal characteristic, which can not be altered by the owner (e.g. gender or race). Second, it depends on their reputation on the platform (e.g. how many previous transactions they have or a personal rating), which the owner can influence over time. Finally, the payoff depends on the rental price set in each period. The model is then applied to the peer-to-peer car sharing setting and calibrated using data from *Turo*.

[17] See the data section for more details on the proxies used here.

1.4.1 Model Setup

Durable Goods Owners

I consider a dynamic single-agent problem in which a durable goods owner sets a rental price in each period. The owner possesses one unit of the good and can rent it out at most once per period for a maximum duration of one period. In each period, if the good is rented out, the owner receives flow utility $u(p)$ and zero otherwise. The probability of a renter arriving in a period depends on the owner's personal characteristic i, reputation s and chosen price p and is given by $\phi_i(p, s)$.

Time and Timing

Time is discrete and covers periods $t = 1, 2, 3, ..., \infty$. The owner is long-lived and active on the platform in each period. She chooses a rental price p for her good at the start of each period, after which a renter arrives (with probability $\phi_i(p, s)$) or not (with probability 1 - $\phi_i(p, s)$).

Payoffs and Reputation Evolution

In case a renter arrives, the owner enjoys flow utility u(p) and her reputation s is updated according to a function $h(\cdot)$; in case no renter arrives, the utility flow is zero and reputation stays the same. The evolution of the state variable can then be written as:

$$s' = \begin{cases} h(p, s, i) & \text{If a renter arrives} \\ s & \text{If no renter arrives} \end{cases} \tag{1.4}$$

where s' denotes the updated reputation. Depending on the application of the model and which measure for reputation is used, the update function $h(\cdot)$ potentially takes the rental price, previous reputation and the unalterable owner characteristic as inputs.

Tradeoffs and the Maximization Problem of the Owner

Each period, the owner chooses a price p in order to maximize the sum of expected flow utility this period and the sum of discounted expected future utility flows. The value function for an owner with characteristic i can be written as:

$$V_i(s) = \max_p \phi_i(p,s) \cdot u(p) + \beta \cdot E\left[V_i(s')|p,s\right] \quad (1.5)$$

where β denotes the discount factor.

In choosing a price for the current period, the owner balances two tradeoffs[18]. First, setting a higher price increases the utility flow in case a renter arrives, but decreases the probability of a renter arriving. And second, in setting a higher price and decreasing the probability of a renter arriving, the owner lowers the probability of her rating being updated, which alters the expected utility flow next period.

1.4.2 Application of the Model to the *Turo* Car Sharing Platform

To apply the model to the *Turo* car sharing platform and the question why female owners have different pricing stregies than men, the unalterable owner characteristic is their gender, which is either male or female: $i \in \{M, F\}$. The owner's control variable is the daily rental price for their vehicle. Finally, the measure for reputation is the number of previous trips owners have rented their cars out for on the platform. The evolution of this state variable s can be partly influenced by the owners through their prices. This is because car owners can influence the probability of a renter arriving, which is given by $\phi_i(p,s)$, by adjusting their prices. The state

[18]The following illustration of the owner's problem assumes (i) that the flow utility of the owner increases in the amount of money she is paid; (ii) that the probability of a renter arriving depends negatively on the asking price. Both assumptions will be true in most applications.

variable is updated whenever a renter arrives. More formally, it follows transition rule (1.6).

$$s' = \begin{cases} s+1 & \text{with probability } \phi_i(p,s) \\ s & \text{with probability 1- } \phi_i(p,s) \end{cases} \tag{1.6}$$

Substituting the expectation term in the general formulation of the value function 1.5 with the transition rule yields the following Bellman equation:

$$V_i(s) = \max_p \phi_i(p,s) \cdot u(p) + \beta \cdot [\phi_i(p,s) \cdot V_i(s+1) + (1-\phi_i(p,s)) \cdot V_i(s)] \tag{1.7}$$

Once the flow utility function $u(p)$, the discount factor β and the probability of a renter arriving, $\phi_i(p,s)$, are specified, the policy function, and hence the optimal car owner behavior, can be recovered by Value Function Iteration on equation (1.7).

The Probability of A Renter Arriving

The key variable in the model is the probability of a renter arriving in each period, given the owner's gender, number of previous trips and price, and is denoted by $\phi_i(p,s)$. This probability can be approximated using the panel structure of the *Turo* data. For each listing, a dummy variable is created which is equal to one, if the same listing has at least one more trip two weeks later, and zero if it still has the same number of trips[19]. Next, the listings for which the dummy is available are divided into 100 cells according to their price and the number of previous trips[20]. Finally, within each of the 100 cells, the averages of the dummy variable among car owners of both genders is taken. These 200 averages (100 price-trip cells for two genders) are divided by 14 (because the dummy variable was created by checking for additional trips after two weeks) to approximate the daily probability of a renter arriving as a function of price, number of trips and owner gender.

[19] This variable can be constructed for 31 444 listings, 5 680 of which were posted by women.

[20] The ten price bins are \$20-\$29, \$30-\$39, ... , \$110-\$119. The ten trip bins are 0-9, 10-19, ..., 90-99. 94% of observations, for which the dummy variable can be constructed, can be assigned to one of the cells.

Figures (1.13) and (1.14) show the approximations to $\phi_M(p, s)$ and $\phi_F(p, s)$, respectively. For both, male and female owners, the probability of a renter arriving decreases in price for all levels of previous trips. Figure (1.15) shows the difference in probability between female and male car owners. The probability of a renter arriving is higher for female car owners for almost all price-trip combinations. There is little difference between the genders for expensive cars ($80) or more at any level of previous trips. Among the cheaper listings, female car owners face a higher probability of a renter arriving than male owners, especially if they have many previous trips. This is consistent with the reduced form findings and the presence of some potential renters, who will only rent from experienced female owners.

Flow Utility and Discount Factor

The discount factor is set to $\beta = 0.99$ and the flow utility function is specified as logarithmic:

$$u(p) = log(p) \tag{1.8}$$

Value Function Iteration, Policy Function and Model Fit

Having specified all components of the model, the optimal car owner behavior can be recovered through Value Function Iteration on equation(1.7)[21]. Figure (1.16) shows the policy functions for male and female car owners. The absolute $ values of the functions cannot be easily interpreted, because the calibration of the model did not take the differences between which car models are offered by the owners into account. However, the differences in optimal dynamic behavior between male and female car owners can explain the differences in pricing decisions between them, which were documented in section 1.3 of this paper. For low amounts of previous trips, the policy function implies that dynamically optimizing female car owners charge less than male owners do. At around twenty trips, there is no difference between the optimal prices across genders. For larger amounts of previous trips, female car owners charge more than male owners, with the difference increasing in the number of previous trips.

[21] The tolerance for the convergence is set to 0.0001 and the 200 probabilities are linearly interpolated from the middle values of the price-trip combination bins to all other possible price-trip combinations.

In order to evaluate the model fit, Figure (1.17) shows the difference in optimal prices in percent between female and male owners as estimated by the model and the difference predicted by the regression of Table 1.4. For low levels of previous trips, the model underestimates the prices set on the platform significantly. However, the model and data predictions as to the price, at which the optimal price of female and male owners is the same, are very close, with the model predicting that number of trips to be 20 and the regression estimates predicting it to be 22. For higher numbers of previous trips, the model fits the data better, in particular for experience levels between 20 and 60 previous trips. According to the model, dynamically optimizing female car owners should set higher prices than they do; however, the calibration of the model did not take into account which kinds of cars the owners are offering. Since female car owners tend to offer smaller and cheaper models (see Table 1.2) than males, the direction in which the model deviates from the data is intuitive.

1.4.3 Counterfactuals

Experienced female car owners face increased demand for their vehicles, allowing them to raise their rental prices (Table 1.5). This increased demand comes from female renters (Figure 1.7), who are willing to pay a premium in order to be able to rent from a woman. Section 1.3.3 showed that this premium is the result of a willingness to pay for safety on the side of the female car renters. This section discusses potential changes to the pricing scheme of the platform that lower the prices charged by experienced female owners. This would enable renters to take advantage of listings they perceive as safe without having to pay a surcharge and draw in extra demand from renters who chose not to participate on the platform before the change. Besides potential profit motives, the platform may want to engage in activities promoting the perceived safety of its users, in order to be proactive in making their platforms safe for all potential participants.

First-Best Intervention

The intervention proposed here requires female car owners to charge prices not higher than the optimal price for male car owners with the same number of previous trips, as estimated by the calibrated model[22]. For each trip, the female owners would be compensated by *Turo* for the difference between their optimal price and the price they were forced to charge[23]. Following the intervention, car renters are better off, as they have access to the same listings by experienced female car owners on *Turo* at lower prices. Female car owners are not worse off, as the total amount of money they receive (the price plus the compensation) is the same as before the intervention and demand for their vehicles has increased, due to the lower prices faced by consumers[24].

For the intervention to be profitable for *Turo*, the increase in revenue due to higher demand for the female cars at the lower prices has to be larger than the cost of reimbursing all experienced female owners, whose cars were rented and would have charged higher prices. The daily expected cost of the intervention is given by equation (1.9):

$$\text{Expected Daily Cost} = \sum_{s:\ \Delta p_s^* > 0} \Delta p_s^* \cdot \bar{p}_{F,s} \cdot \phi_F(p_{M,s}^*, s) \cdot \text{Listings}_{F,s} \tag{1.9}$$

Δp_s^* denotes the percentage difference between the optimal prices by female and male car owners with s previous trips as estimated by the dynamic model (visualized by the solid line in Figure 1.17). This percentage difference is multiplied by the average price charged by female car owners with s previous trips, $\bar{p}_{F,s}$. To get the expected amount to be paid as compensation for a listing with s trips, it is neccessary to multiply by $\phi_F(p_{M,s}^*, s)$, the probability of a renter arriving for a female owner with s trips, who is forced to charge the optimal price for a male owner with s trips. The number is then multiplied by the number of listings by females with s previous trips, $\text{Listings}_{F,s}$ to obtain the expected daily cost for a particular s. Finally, in

[22]For each level of experience, this difference in percentage terms is visualized by the solid line in Figure 1.17.
[23]Charging even lower prices would still be allowed.
[24]Male car owners may potentially be worse off, if they suffer from increased competition from now cheaper offerings by female car owners. The model evaluated here ignores this competition effect. However, the competition effect is likely to be small, because female car owners account for only a small fraction of all listings (see Figure 1.9).

order to calculate the total daily expected cost of the intervention, the multiplication term is summed over all levels of s, for which female car owners charge higher prices than male car owners, given by all s for which $\Delta p_s^* > 0$ is true.

This expected cost has to be compared to the expected change in revenue resulting from the changed prices and hence demand for female car owner's listings. *Turo* earns money by charging a fee of 25% of the total price of each transaction. The change in revenue is given by equation (1.10)[25].

$$\text{Expected Revenue Change} = \sum_{s:\ \Delta p_s^* > 0} 0.25 \cdot (p_{M,s}^* \cdot \phi_F(p_{M,s}^*, s) - p_{F,s}^* \cdot \phi_F(p_{F,s}^*, s)) \cdot \text{Listings}_{F,s} \tag{1.10}$$

The expected revenue change compares the expected daily earnings after the intervention, when female owners charge the optimal price for male owners, $p_{M,s}^*$, and consequently face the probability of a renter arriving given by $\phi_F(p_{M,s}^*, s)$, to expected daily earnings without the intervention, when female owners charge $p_{F,s}^*$ and face rental probability $\phi_F(p_{F,s}^*, s)$. This expected difference for each level of previous trips s is then multiplied by the respective number of listings by female owners and finally summed across all levels of previous trips to obtain the total expected daily revenue change facing *Turo*.

Since all components of equations (1.9) and (1.10) are either observed in the data or estimates of the model, the cost of and revenue change following the intervention can be calculated in dollar terms. The daily total expected cost of the intervention is \$4 332, while daily total expected revenue increases by \$6 157. These estimates imply that the intervention would increase annual profits by \$669 410[26].

[25]As in the cost calculation, the expected change in revenue is computed assuming that there is no effect on the demand for listings unaffected by the intervention, i.e. listings by male owners and those with low levels of previous trips.

[26]This compares to estimated annual revenues of \$10 million according to https://www.crunchbase.com/organization/turo

Discussion and Alternative Intervention Designs

The counterfactual calculations show that implementing a change to the pricing scheme, following which female owners would be required to charge no more than the optimal prices of male car owners, results in a profit increase for *Turo*[27]. While not being outright illegal, as car owners who list their vehicles on *Turo* are not classified as employees, the platform may nonetheless find it difficult to implement the first-best intervention as described above, since male car owners may feel treated unfairly. Absent the possibility of implementing the first-best intervention, other schemes incentivizing experienced female car owners to charge lower prices are worth investigating.

Online platforms often award achievement badges to users who fulfill certain requirements[28]. Gender-specific achievement badges would be easier to implement than gender-specific compensation schemes. Alternatively, awarding a 'Fair Price Setter' badge to all experienced users, who abstained from charging premiums, would have the biggest impact on experienced female users, as they are charging the biggest premiums, and incentivize them to keep their prices low. Finally, while gender-specific cash compensations are difficult to implement, subsidizing insurance premiums, which have to be paid for each trip on *Turo*, or fuel purchases may be easier to realize.

1.5 Conclusion

Male and female car owners set different prices on a peer-to-peer car sharing platform. Experienced female owners set higher prices than men, while inexperienced female owners set lower prices. The differences are driven by excess demand for experienced female owner's listings. The excess demand comes from female renters who are concerned about their safety during the transaction and willing to pay extra in order to be able to rent from a woman as

[27] As mentioned above, this number overstates the true benefit, if the now cheaper listings by female owners compete with listings by male owners. On the other hand, Figure 1.17 shows that the model overestimates the price difference between experienced and inexperienced female car owners, implying that the true compensations that would have to be paid, and hence the cost of the intervention, would be lower.

[28] For example, *Turo* awards a badge to owners who had a certified photographer take the pictures of their vehicle they use on the platform.

opposed to a man.

A dynamic model of price setting on peer-to-peer sharing platforms is able to replicate the pricing pattern observed in the data. Counterfactual calculations show that requiring experienced female owners to set low prices, but reimbursing them for the lost income, benefits renters, female owners and increases platform profits. This is due to the fact that lowering prices for experienced female owner's cars generates additional demand and traffic on the platform.

This paper shows how gender homophily on the demand side can drive price setting behavior on peer-to-peer sharing platforms. Taking this homophily as a given feature of the demand side and adapting the platform's pricing scheme, so that female users no longer have to pay a surcharge for feeling safe, results in an improvement in outcomes for female renters and users, as well as for the platform itself. This insight can be applied to other peer-to-peer platforms, where some user groups are underrepresented and potentially homophilous. Subsidizing users belonging to these groups or otherwise encouraging their participation on the platform potentially improves outcomes of already active users and makes the platform attractive to a wider audience.

1.6 Appendix

Details about the *Turo* website and the data collection process

On the *Turo* website, a user inserts the location and time for which she wishes to rent a car (see Figure 1.1). Next, the results of the search appear, featuring a picture, the price, rating and number of previous trips for every suitable listing (see Figure **??**). To narrow down the selection, the user may filter the results using the functions on the left border of the screen (for example, to only display convertibles). After clicking on a listing, more information about that listing becomes visible (see Figure 1.3). The renter may then book a vehicle through the platform, which puts renter and owner in touch with each other. All communication between owners and renters, payment and insurance are handled and provided by the online platform[29][30].

100 mid-sized U.S. cities (with populations between 100 000 and 300 000) were chosen at random and all *Turo* listings in these cities were scraped daily. This was done in order to ensure that enough unique drivers were observed multiple times over the months. Choosing smaller cities would have resulted in fewer offered listings, choosing larger cities would have brought the risk of missing a driver, because *Turo* only displays 200 results for each search. Every day, after the 100 cities had been scraped, a random U.S. zip code was drawn (with replacement) from the universe of U.S. zip codes and all listings in that zip code area were scraped.

Only listings for which the first name of the owner unambiguously identifies the gender remain in the dataset. A first name is determined to be female (male), if 95% of respondents to the 1990 U.S. census with that name were female (male)[31].

[29] *Turo* claims to go great lengths to avoid side payments outside its digital infrastructure

[30] Both car and driver are insured independently from any pre-exisitng insurance policies for the duration of any trip initiated through *Turo*. The renters pay for this insurance.

[31] Listings by owners named ”Andrea”, for example, were dropped

Limitations of the Data Collection Process

The data collection process described in this section has several potential problems. First, the most attractive listings may be booked very quickly, which makes them less likely to appear in the dataset. Second, long-term rentals over several weeks or even months imply a high utilization rate and high earnings for the car owner, but only count for one trip, which is my measure of experience on the platform. Finally, after having concluded a trip, renters are not obliged to leave a review for the listing. This may bias my measure of the ratio of female renters, which is given by the fraction of reviewers with female names. This bias may prove problematic, if it is systematic, for example because female renters are more likely to leave a review to a female owner, which would lead to an overestimation of the degree of homophily among renters. However, Figure 1.12 shows that there is no systematic difference across genders in the likelihood of a review being left after a trip.

1.7 Figures

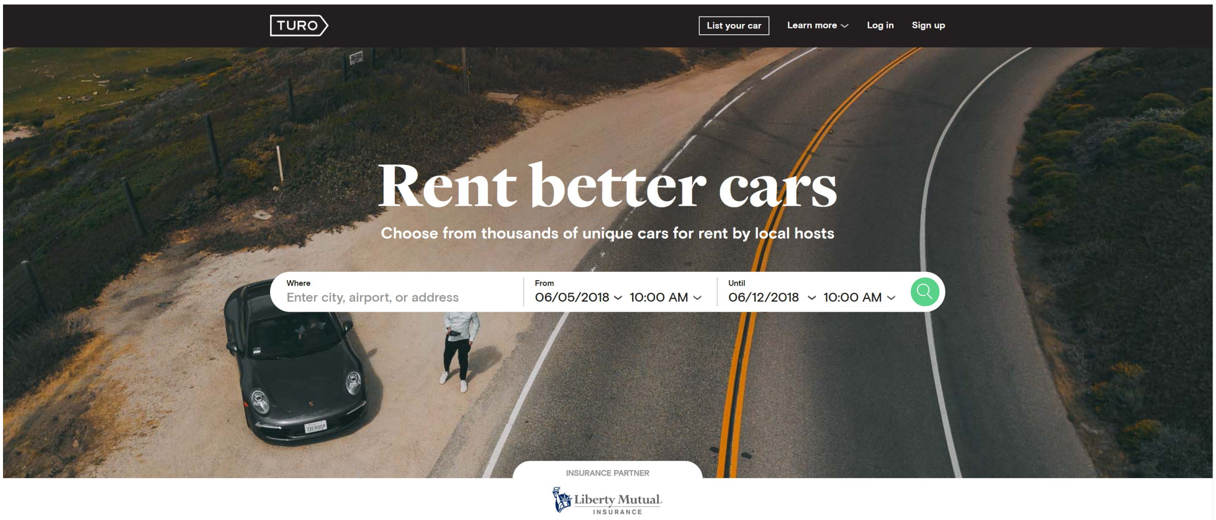

Figure 1.1: The *Turo* home page.

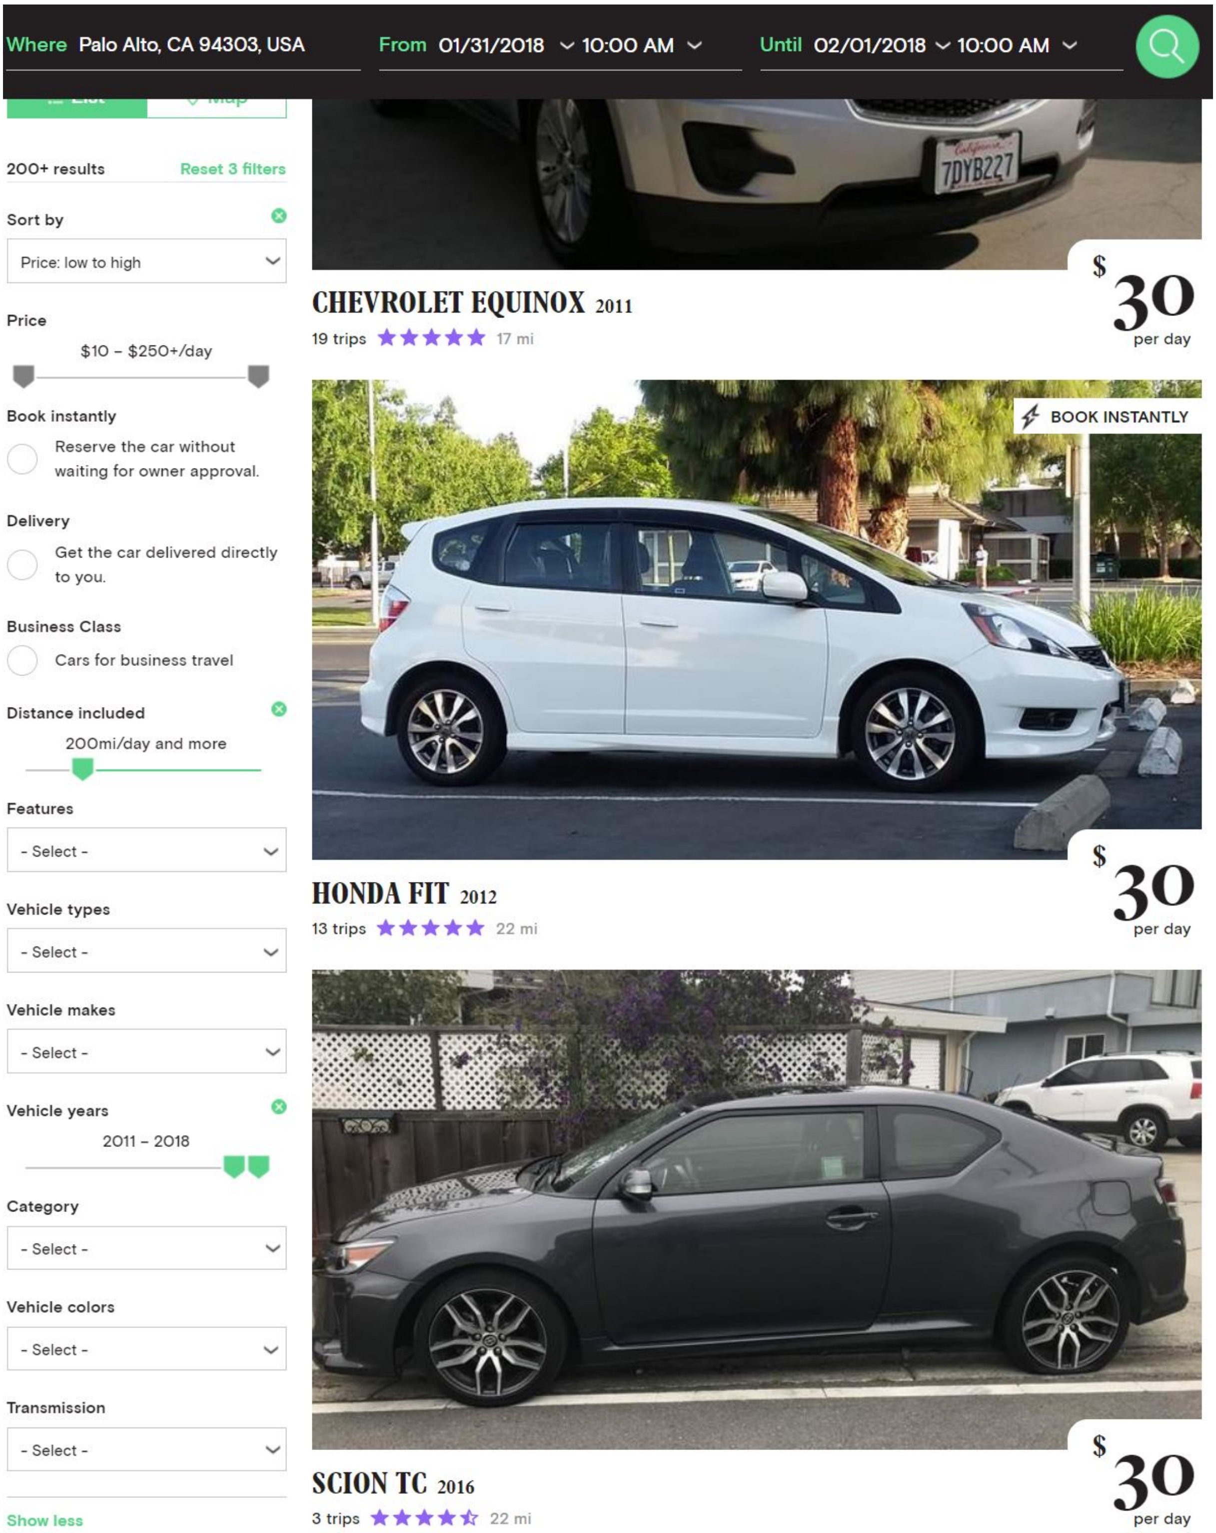

Figure 1.2: The *Turo* search results page.

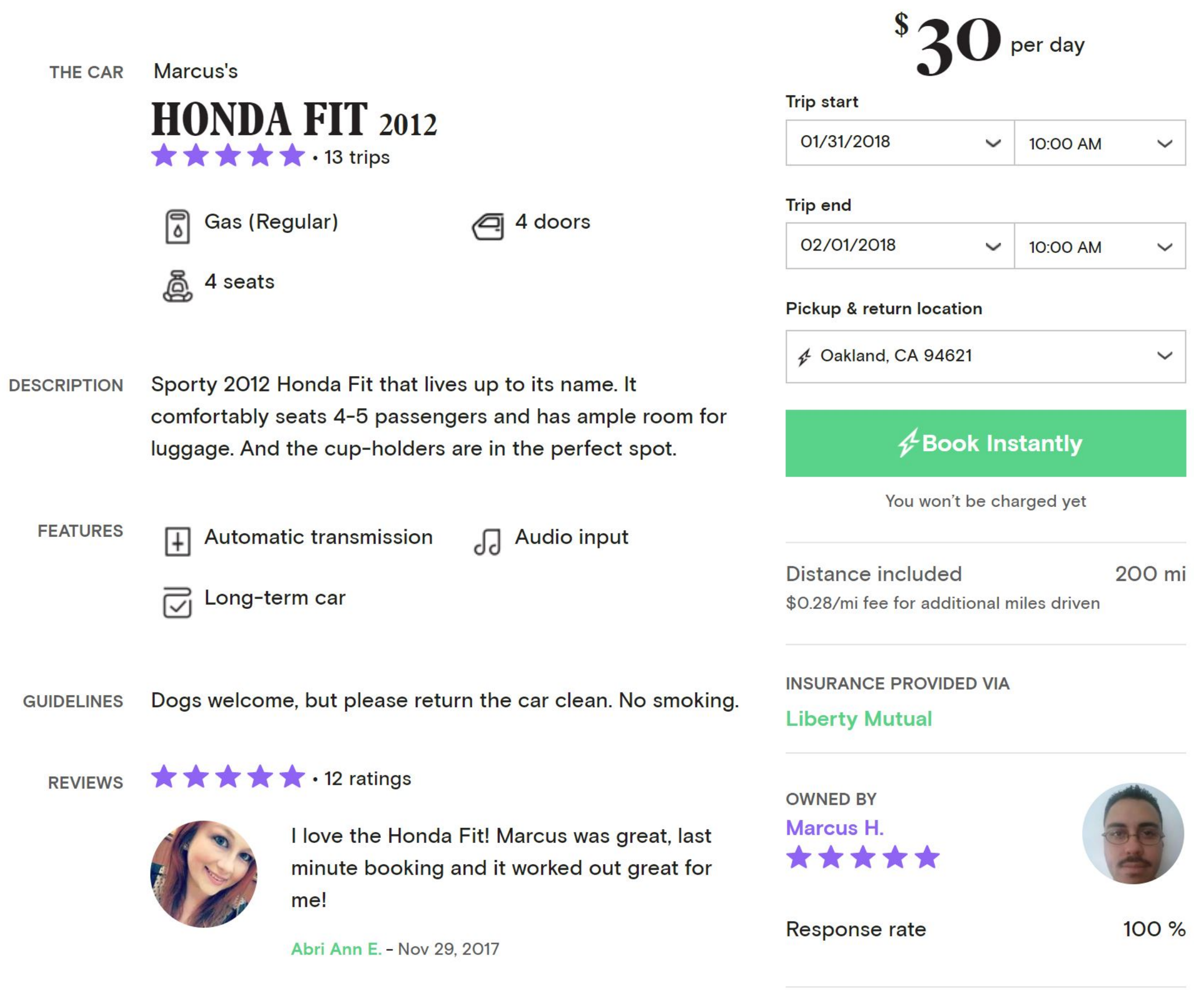

Figure 1.3: Example for a listing page. One cross-sectional observation in the dataset corresponds to the data visible in such a listing page. The observed variables include model and model year of the listed car, daily price, included mileage, number of previous trips, average rating for those previous trips, whether instant booking is enabled, the name of the owner, his/her description of the car and the names and reviews of previous renters.

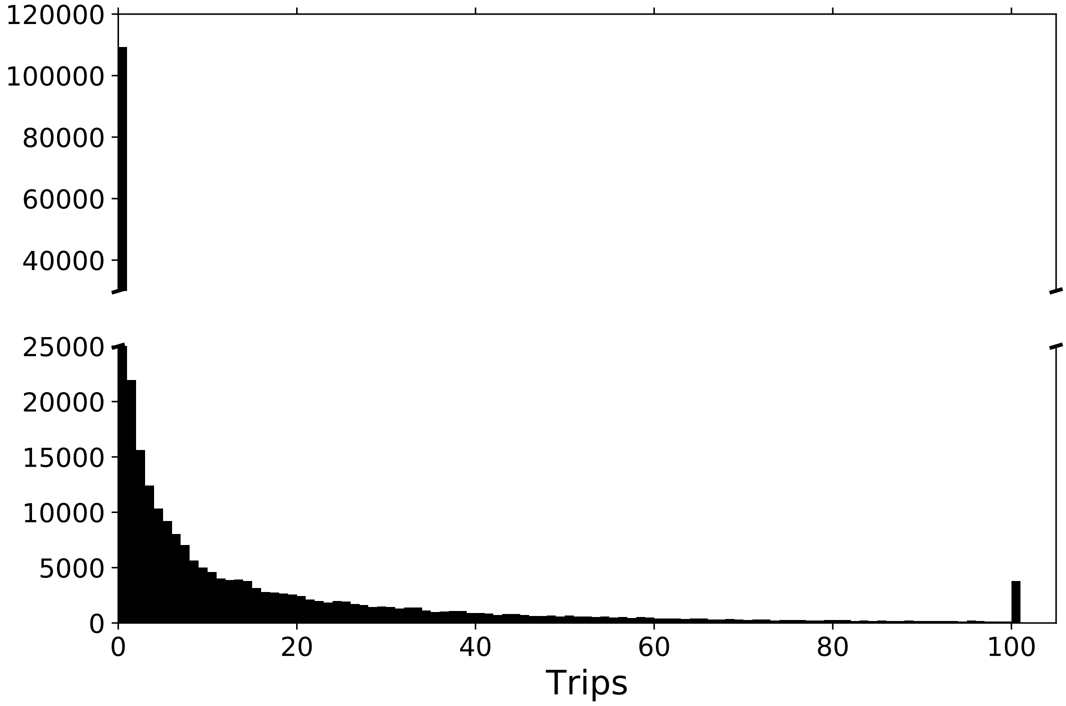

Figure 1.4: Frequency of listings in the sample by number of previous trips.

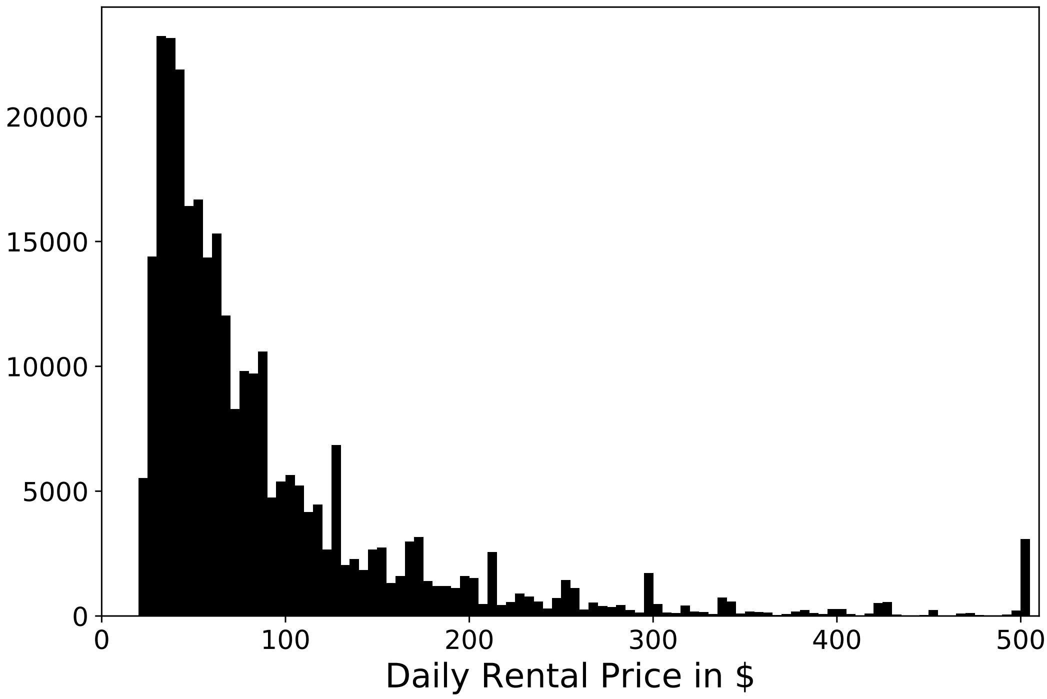

Figure 1.5: Frequency of prices in the sample by number of previous trips.

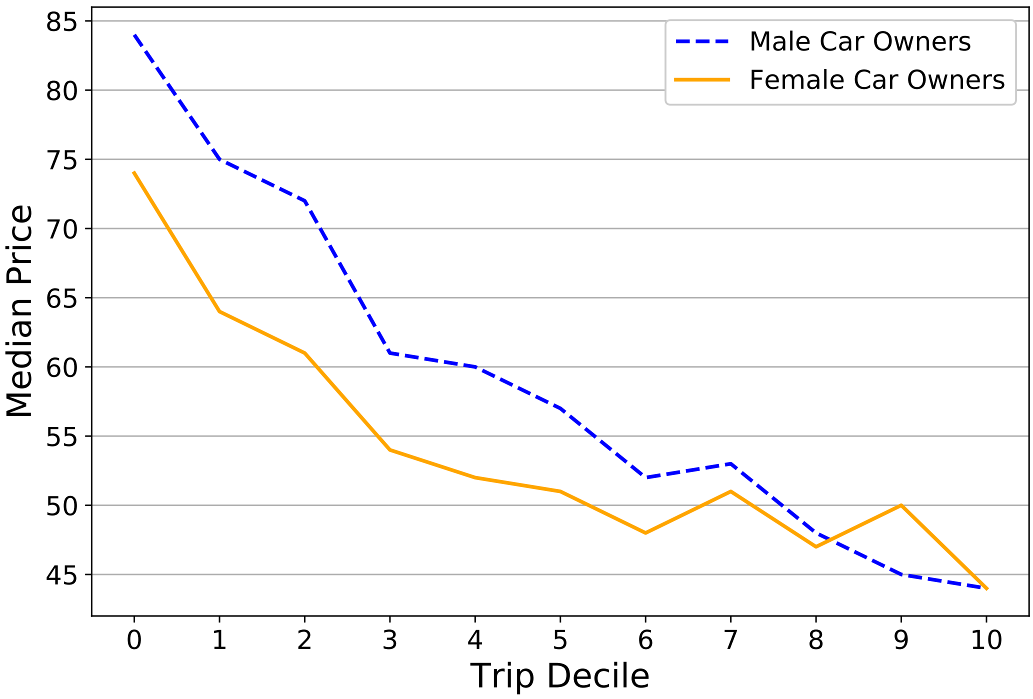

Figure 1.6: Median price by number of previous trips and owner gender.

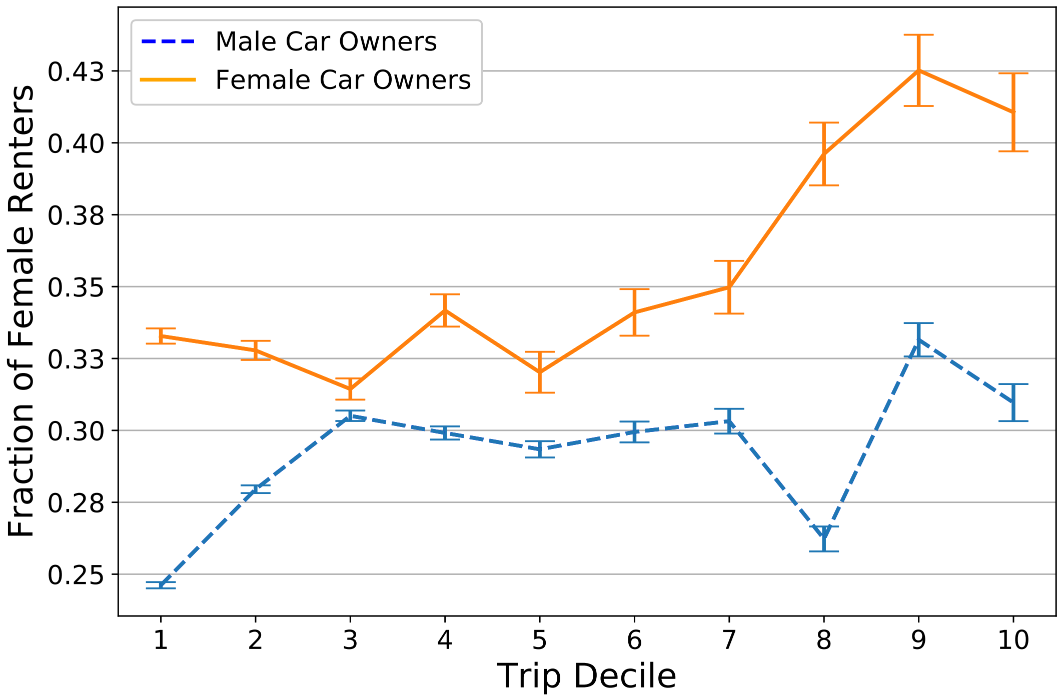

Figure 1.7: Fraction of female renters on previous trips by male and female owners and by number of previous trips.

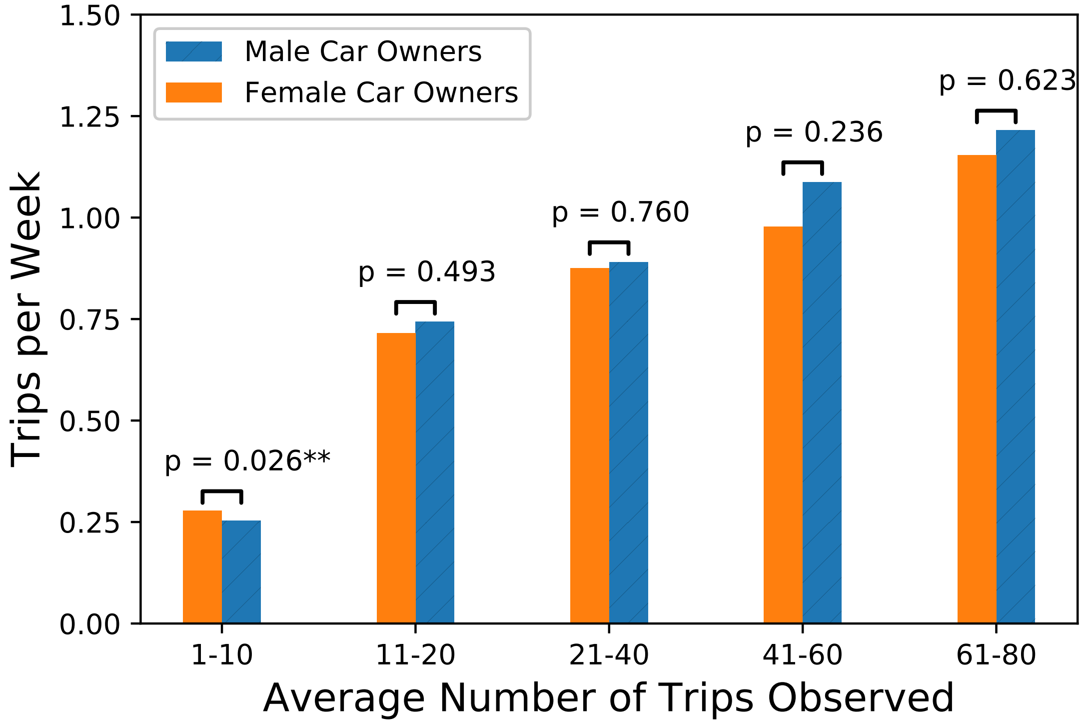

Figure 1.8: Average number of trips per week by owner gender. The numbers above the bar indicate the p-values of a two-sided t-test comparing the number of trips per week across owner genders.

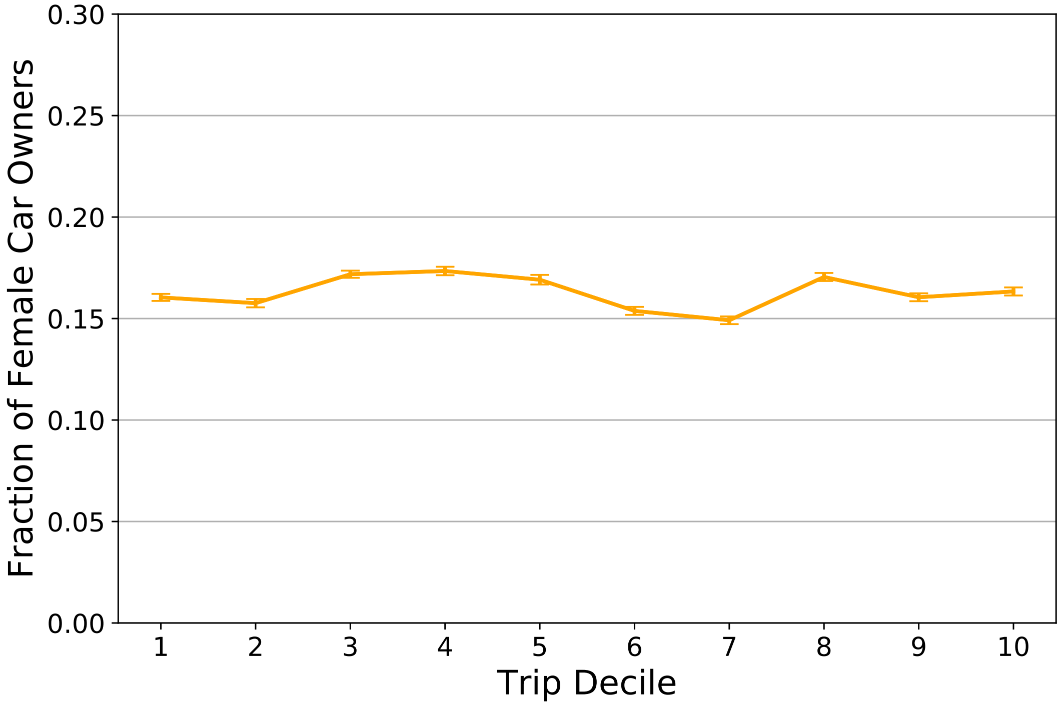

Figure 1.9: Fraction of female car owners by experience decile and owner gender.

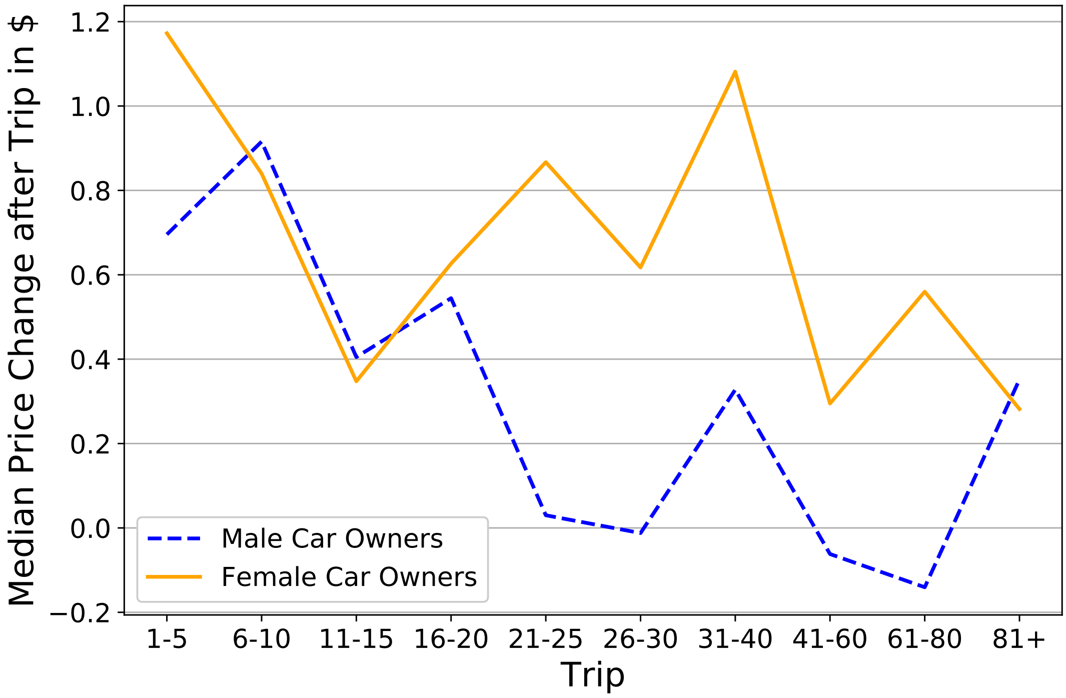

Figure 1.10: Average price change after trip by owner gender.

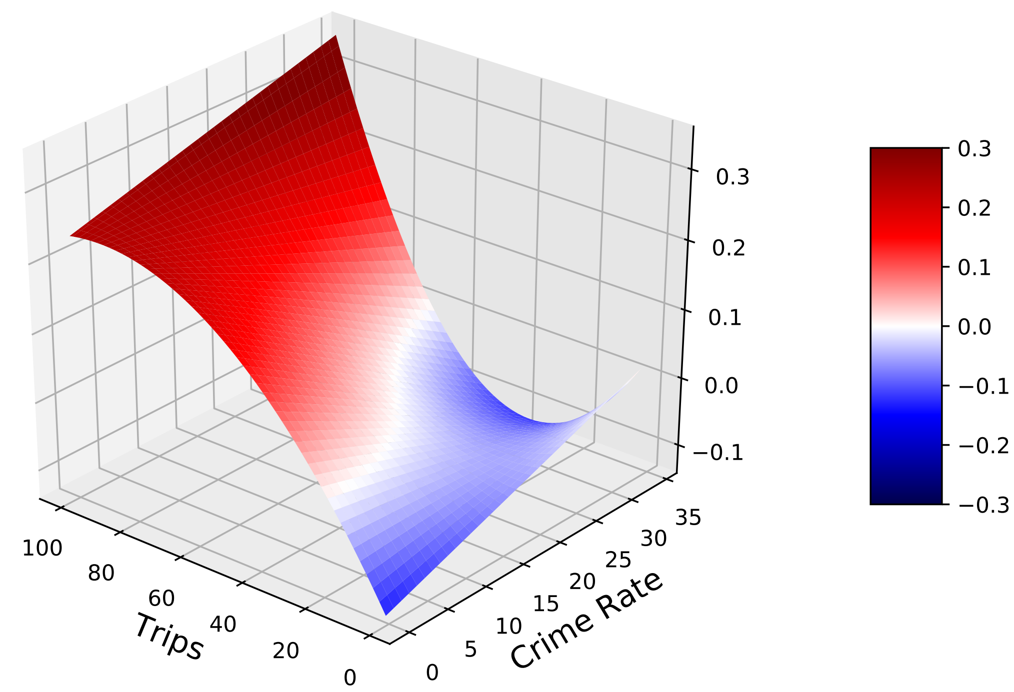

Figure 1.11: Predicted price difference in % between female and male car owners as a function of previous trips and local crime rate as predicted by the parameter estimates in column (4) of Table 1.8.

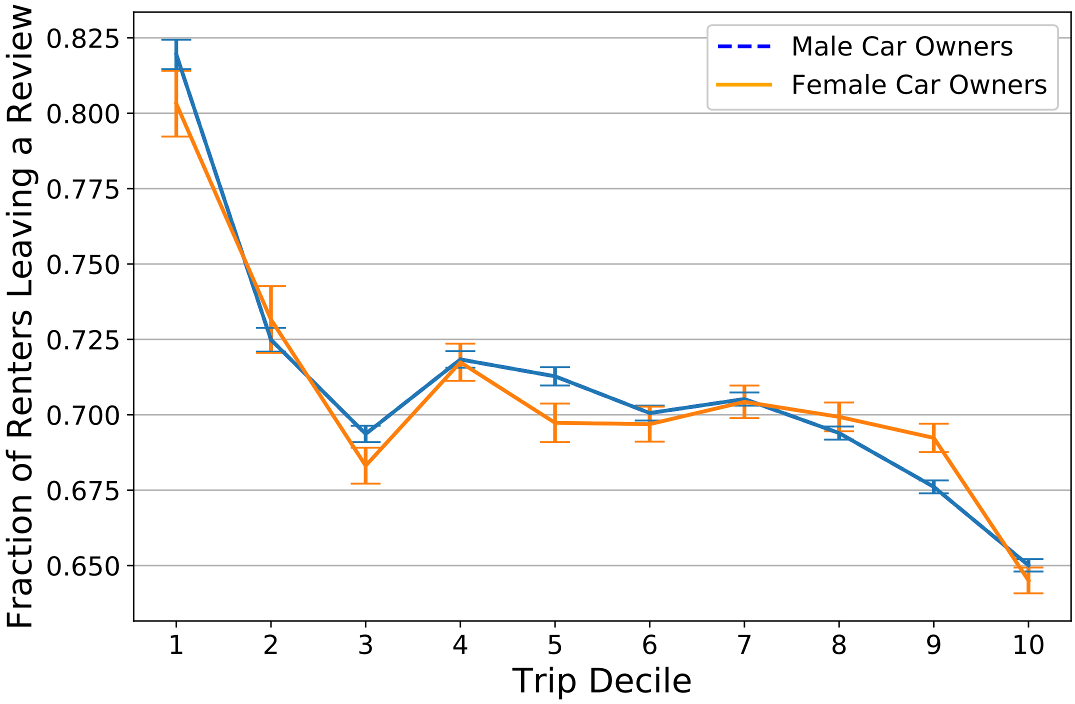

Figure 1.12: Number of reviews divided by number of trips by car owner gender.

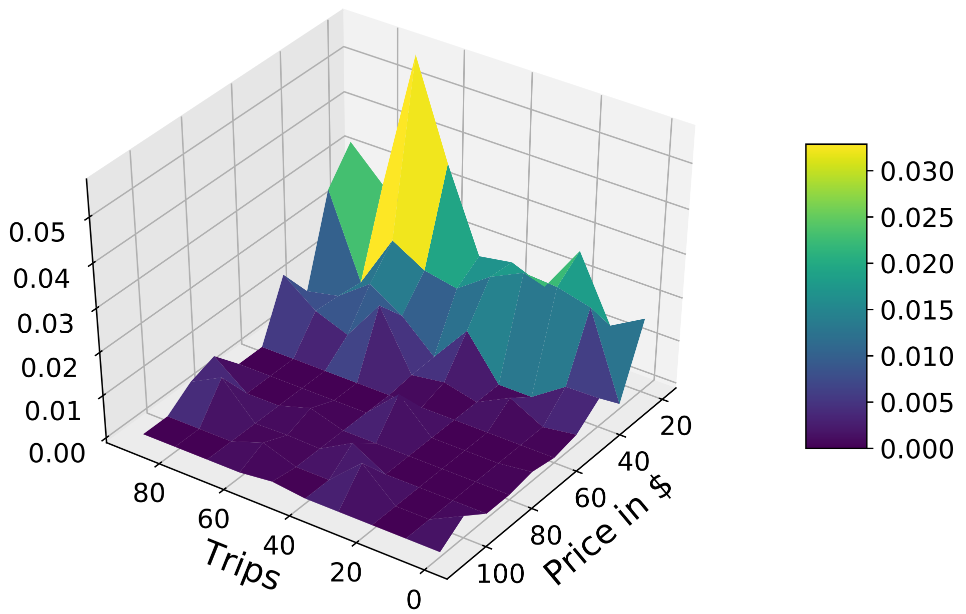

Figure 1.13: Daily probability of a renter arriving for male car owners for 100 trip-price combination cells.

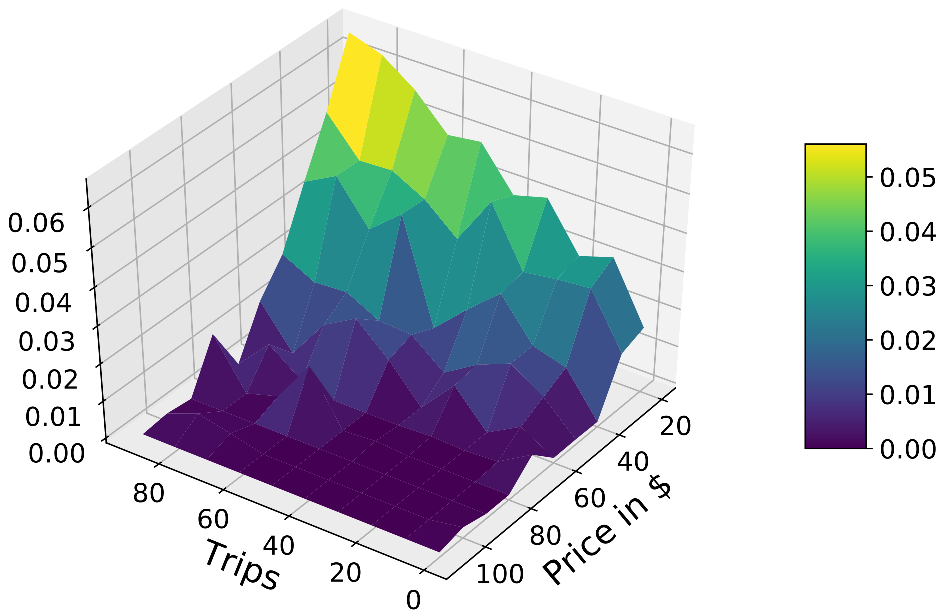

Figure 1.14: Daily probability of a renter arriving for female car owners for 100 trip-price combination cells.

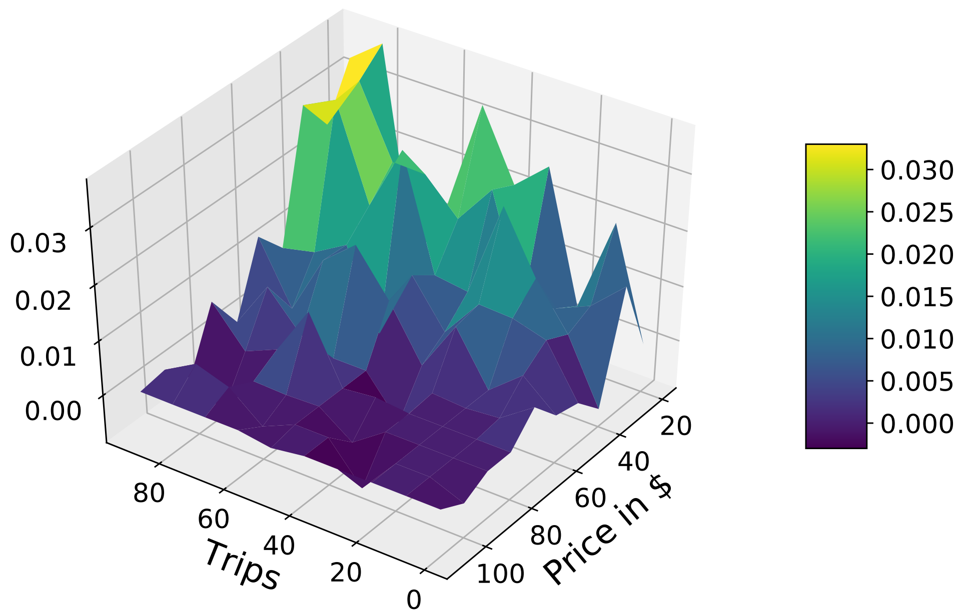

Figure 1.15: Difference in daily probabilities of a renter arriving between female and male car owners for 100 trip-price combination cells.

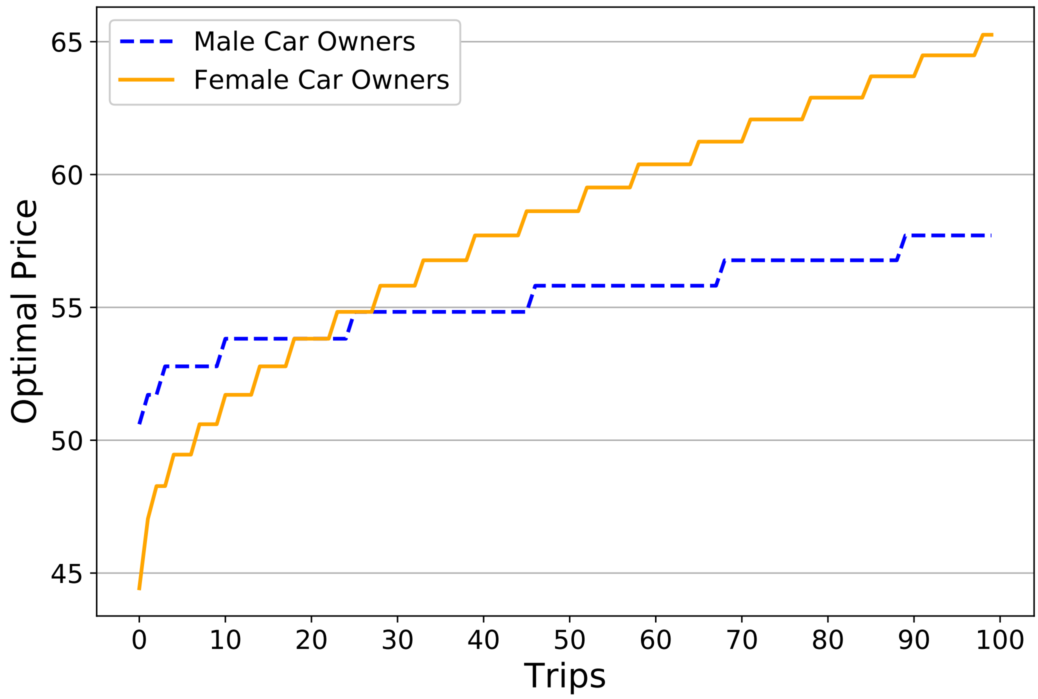

Figure 1.16: Policy function recovered through Value Function Iteration for the model in section 1.4.2.

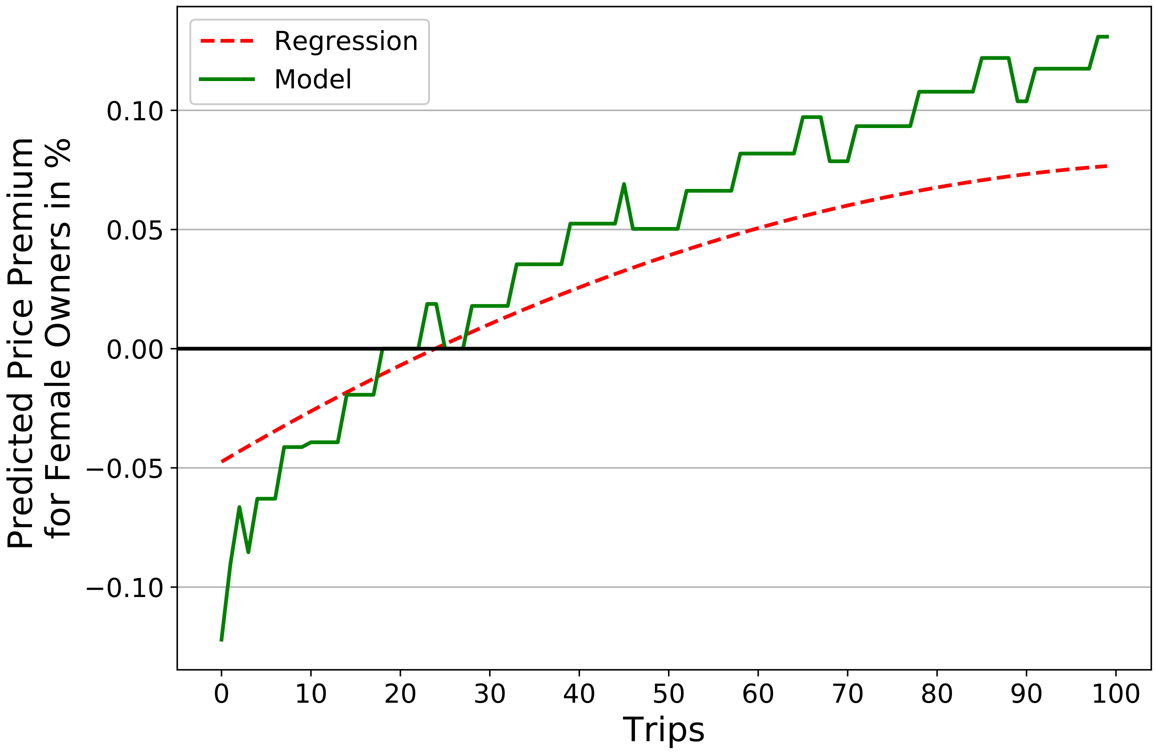

Figure 1.17: Differences in optimal prices in percent between male and female car owners as estimated by the poliy function of the model of section 1.4.2 and the regression results of Table 1.4.

1.8 Tables

	Male (N = 245 764)		Female (N = 47 898)		
Variable	Mean (SD)	Median	Mean (SD)	Median	Δ Mean
Price per Day (in $)	94.2 (94.1)	64	84.3 (84.3)	59	9.9***
Previous Trips	11.8 (22.5)	2	12.5 (23.9)	3	-0.7***
Distance Included (in m)	125.7 (82.3)	200	124.1 (85.9)	200	1.6***
Model Year	2013.1 (4.4)	2014	2013.4 (4.3)	2014	-0.3***
Rating	4.78 (0.63)	5	4.77 (0.62)	5	-0.01***
Fraction Female Renters	0.28 (0.32)	0.2	0.32 (0.34)	0.25	-0.03***
Instant Booking Activated	0.67 (0.47)	1	0.67 (0.47)	1	0.00
Attitude towards Women	2.97 (0.11)	2.96	2.98 (0.10)	2.99	-0.01***
Gender Wage Gap	0.82 (0.05)	0.82	0.82 (0.04)	0.82	-0.00***

Table 1.1: Summary statistics for the entire cross-sectional dataset. The variable "Fraction Female Renters" was constructed by calculating the fraction of reviewers of previous trips with female names. The "Attitude towards Women" variable refers to the reaction to the statement: "It is much better for everyone involved if the man is the achiever outside the home and the woman takes care of the home and family.". The possible reactions are 1 (Strongly Agree), 2 (Agree), 3 (Disagree) and 4 (Strongly Disagree). The gender wage gap variable measures the fraction of average female to average male wages in percent by U.S. state in 2016. The last column shows the difference in means across genders and whether the difference is significant in a two-sided t-test. The stars indicate the following significance values: * $p< 0.1$, ** $p<0.05$, *** $p<0.01$

Segment	Fraction in Sample	Fraction Female in Segment	Median Price in (\$/day) $\mid$ Trips > 0
Electric	4.4%	10.5%	140
Large Sedan	8.1%	14.1%	98
Other	7.9 %	11.7%	98
Medium SUV	14.0%	16.6%	85
Pickup Truck	5.2%	8.5%	58
Minivan	4.5%	16.5%	56
Small SUV	10.0%	23.9%	55
Medium Sedan	19.3%	16.6%	49
Hybrid	9.3 %	19.6%	40
Small Sedan	17.3%	16.8%	35

Table 1.2: Car segments and median prices. The first column shows the frequency with which models in each segment appear in the sample. The second column shows what fraction of owners in each segment is female. The third column shows the median asking price in each segment, given that the vehicle has been rented at least once before.

	Male (N = 51 726)		Female (N = 10 918)		
Variable	Mean (SD)	Median	Mean (SD)	Median	Δ Mean
Price per Day (in $)	91.5 (81.3)	68	80.5 (74.4)	59	11.0***
Previous Trips	11.7 (21.3)	3	11.9 (17.5)	3	-0.2*
Distance Included (in m)	126.2 (8.3)	200	125.1 (82.3)	200	1.1**
Model Year	2013.9 (4.1)	2015	2014.0 (4.2)	2015	-0.4***
Rating	4.83 (0.53)	5	4.71 (0.81)	5	-0.12***
Fraction Female Renters	0.27 (0.32)	0.2	0.33 (0.35)	0.25	-0.06***
Instant Booking Activated	0.71 (0.45)	1	0.76 (0.43)	1	-0.05***

Panel B: *UCR* Variables			
Variable (per 100k Inhabitants)	Mean	Median	SD
Violent Crime Cases	622.2	571.9	346.0
Rape Cases	27.5	27.1	13.7

Table 1.3: Summary statistics for the cross-sectional dataset where local crime data is available. The last column shows the difference in means across genders and whether the difference is significant in a two-sided t-test. The stars indicate the following significance values: * p< 0.1, ** p<0.05, *** p<0.01

	(1) lg(Price)	(2) lg(Price)	(3) lg(Price)	(4) lg(Price)	(5) lg(Price)
Female	-0.115***	-0.039***	-0.042***	-0.041***	-0.047***
	(0.021)	(0.011)	(0.009)	(0.009)	(0.010)
Miles Allowed	-0.000***	-0.000	-0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Previous Trips	-7.326***	-2.276***	-1.939***	-1.866***	-3.959***
$x10^{-3}$	(0.323)	(0.198)	(0.138)	(0.135)	(0.270)
Female * Previous Trips	1.785**	1.088**	1.297***	1.266***	2.219***
$x10^{-3}$	(0.695)	(0.435)	(0.290)	(0.292)	(0.506)
Previous Trips^2					1.843***
$x10^{-5}$					(0.207)
Female * Previous Trips^2					-0.976***
$x10^{-5}$					(0.303)
Model-Year FE	No	Yes	Yes	Yes	Yes
Zip Code FE	No	No	Yes	Yes	Yes
Calendar Day FE	No	No	No	Yes	Yes
Sample	Full	Full	Full	Full	Full
N	293 662	293 662	293 662	293 662	293 662
R^2	0.064	0.735	0.857	0.858	0.859

Table 1.4: Standard errors (in parentheses) are clustered at the zip code level. The stars indicate the following significance values: * p< 0.1, ** p<0.05, *** p<0.01

	(1) lg(Price)	(2) lg(Price)	(3) lg(Price)	(4) lg(Price)	(5) lg(Price)	(6) lg(Price)	(7) lg(Price)
Female	-0.026***	-0.068***	-0.043*	-0.002	0.027	0.064**	0.126***
	(0.008)	(0.015)	(0.025)	(0.023)	(0.024)	(0.027)	(0.034)
Miles Allowed	-0.000	-0.000***	0.000	-0.000	-0.000	0.000	0.000**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Previous Trips	-0.002***		-0.003	0.002	-0.001	-0.001***	-0.000***
	(0.000)		(0.005)	(0.002)	(0.001)	(0.000)	(0.000)
Rating			0.007	-0.023**	0.022*	0.055***	0.048***
			(0.007)	(0.012)	(0.011)	(0.012)	(0.013)
Model-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Zip Code FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Sample	Full	0 Trips	1-3 Trips	4-8 Trips	9-23 Trips	$>$ 24 Trips	$>$ 48 Trips
N	293 662	105 907	41 112	36 522	34 875	35 142	14 256
R^2	0.858	0.903	0.961	0.961	0.962	0.960	0.968

Table 1.5: Standard errors (in parentheses) are clustered at the zip code level. The specification of the regression model is as in equation (1.1). The subsamples are created using the percentiles of the number of previous trips conditional on having at least one previous trip. The 25th, 50th, 75th and 90th percentile of trips conditional on having at least one trip are 3, 8, 23 and 47. The stars indicate the following significance values: * $p< 0.1$, ** $p<0.05$, *** $p<0.01$

	Male (N = 34 416)		Female (N = 6 555)		
Variable	Mean (SD)	Median	Mean (SD)	Median	Δ Mean
Price per Day (in $)	66.4 (55.2)	47	61.6 (49.5)	46	4.8***
Previous Trips	24.3 (31.3)	13	24.8 (32.4)	13	-0.5
Distance Included (in m)	123.2 (78.3)	200	122.1 (82.1)	200	1.1***
Model Year	2013.2 (3.9)	2014	2013.3 (4.22)	2014	-0.1***
Rating	4.6 (1.1)	5	4.6 (1.1)	5	-0.0**
Fraction Female Renters	0.31 (0.31)	0.25	0.34 (0.32)	0.25	-0.03***
Instant Booking Activated	0.68 (1)	1	0.63 (0.48)	1	-0.05***
Times Observed	4.3 (1.9)	4	4.8 (2.1)	4	-0.5***

Table 1.6: Summary statistics for the panel dataset. One observation corresponds to an owner-trip combination. All owners that were only observed for one level of previous trips were dropped. Many owners are observed multiple times for the same number of previous trips. For those cases, all variable values are from the latest calendar day during which the owner was observed with that number of trips. The variable "Fraction Female Renters" was constructed by calculating the fraction of reviewers of previous trips with female names. The last column shows the difference in means across genders and whether the difference is significant in a two-sided t-test. The stars indicate the following significance values: * $p < 0.1$, ** $p<0.05$, *** $p<0.01$

	(1) lg(Price)	(2) lg(Price)
Miles Allowed	-0.000	-0.000
	(0.000)	(0.000)
Previous Trips	0.000	0.000
	(0.000)	(0.000)
Rating	0.005***	0.004***
	(0.001)	(0.001)
Female * Trips	2.27**	3.59**
x10^{-3}	(1.00)	(1.41)
Trips^2		-1.45
x10^{-6}		(3.07)
Female * Trips^2		-1.55**
x10^{-6}		(7.01)
Owner FE	Yes	Yes
Sample	Full	Full
N	40 971	40 971
R^2	0.978	0.978

Table 1.7: Standard errors (in parentheses) are clustered at the zip code level. The specification of the regression model is as in equation (1.2). The stars indicate the following significance values: * $p< 0.1$, ** $p<0.05$, *** $p<0.01$

	(1) lg(Price)	(2) lg(Price)	(3) lg(Price)	(4) lg(Price)
Female	-0.0261	-0.0738	-0.124**	-0.141**
	(0.021)	(0.049)	(0.054)	(0.057)
Miles Allowed	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)
Previous Trips	-0.002***	-0.002***	-0.002***	-0.002***
	(0.000)	(0.000)	(0.000)	(0.000)
Rating	-0.004**	-0.004**	-0.004**	-0.003**
	(0.002)	(0.002)	(0.002)	(0.002)
Female*Trips	0.506	0.75	4.68***	8.50***
$x10^{-3}$	(0.817)	(0.87)	(1.21)	(2.68)
Female*Crime Stat		1.58	3.71*	4.87**
$x10^{-3}$		(1.93)	(2.18)	(2.32)
Female*Crime Stat*Trips			-1.82***	-4.86***
$x10^{-4}$			(0.454)	(1.15)
Female*Trips^2				-4.71**
$x10^{-5}$				(1.89)
Female*Crime Stat*Trips^2				4.70***
$x10^{-6}$				(1.31)
Model-Year FE	Yes	Yes	Yes	Yes
Zip Code FE	Yes	Yes	Yes	Yes
Calendar Day FE	Yes	Yes	Yes	Yes
Sample	UCR available	UCR available	UCR available	UCR available
N	51 855	51 855	51 855	51 855
R^2	0.901	0.901	0.902	0.902

Table 1.8: Standard errors (in parentheses) are clustered at the zip code level. The specification of the regression model is as in equation (1.3). Model-year fixed effects refer to the car model and the year in which the car was built. The crime rate is rape cases per 100 000 inhabitants. The stars indicate the following significance values: * $p< 0.1$, ** $p<0.05$, *** $p<0.01$

	(1) lg(Price)	(2) lg(Price)	(3) lg(Price)	(4) lg(Price)
Female	-0.042***	-0.517	-0.451	-0.503
	(0.012)	(0.307)	(0.327)	(0.351)
Miles Allowed	-0.000	-0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)
Previous Trips	-0.002***	-0.002***	-0.002***	-0.004***
	(0.000)	(0.000)	(0.000)	(0.000)
Rating	-0.011***	-0.011***	-0.011***	-0.009***
	(0.001)	(0.001)	(0.001)	(0.001)
Female*Trips	1.32***	1.30***	-4.89	-5.04
x10^{-3}	(0.37)	(0.36)	(9.99)	(23.4)
Female*Attitude		0.159	0.137	0.153
		(0.103)	(0.110)	(0.118)
Female*Attitude*Trips			2.07	-0.92
x10^{-4}			(3.33)	(7.8)
Female*Trips^2				-0.993
x10^{-4}				(1.57)
Female*Attitude*Trips^2				3.00
x10^{-5}				(5.28)
Model-Year FE	Yes	Yes	Yes	Yes
Zip Code FE	Yes	Yes	Yes	Yes
Calendar Day FE	Yes	Yes	Yes	Yes
Sample	Full	Full	Full	Full
N	293 662	293 662	293 662	293 662
R^2	0.874	0.874	0.874	0.874

Table 1.9: Standard errors (in parentheses) are clustered at the zip code level. Model-year fixed effects refer to the car model and the year in which the car was built. The attitude variable refers to the reply to the statement: ”It is much better for everyone involved if the man is the achiever outside the home and the woman takes care of the home and family.”. The possible answers are 1 (Strongly Agree), 2 (Agree), 3 (Disagree) and 4 (Strongly Disagree). The stars indicate the following significance values: * $p< 0.1$, ** $p<0.05$, *** $p<0.01$

	(1) lg(Price)	(2) lg(Price)	(3) lg(Price)	(4) lg(Price)
Female	-0.042***	0.196	0.289	0.374
	(0.012)	(0.216)	(0.241)	(0.259)
Miles Allowed	-0.000	-0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)
Previous Trips	-0.002***	-0.002***	-0.002***	-0.002***
	(0.000)	(0.000)	(0.000)	(0.000)
Rating	-0.011***	-0.011***	-0.011***	-0.011***
	(0.001)	(0.001)	(0.001)	(0.001)
Female*Trips	1.32***	1.34***	-6.37	-8.25
$x10^{-3}$	(0.37)	(0.36)	(9.99)	(17.2)
Female*Wage Gap		-0.286	-0.398	-0.495
		(0.257)	(0.287)	(0.308)
Female*Wage Gap*Trips			0.009	0.028
			(0.008)	(0.02)
Female*Trips^2				1.69
$x10^{-4}$				(1.01)
Female*Wage Gap*Trips^2				-1.93
$x10^{-4}$				(1.21)
Model-Year FE	Yes	Yes	Yes	Yes
Zip Code FE	Yes	Yes	Yes	Yes
Calendar Day FE	Yes	Yes	Yes	Yes
Sample	Full	Full	Full	Full
N	293 662	293 662	293 662	293 662
R^2	0.874	0.874	0.874	0.874

Table 1.10: Standard errors (in parentheses) are clustered at the zip code level. Model-year fixed effects refer to the car model and the year in which the car was built. The gender wage gap variable measures the fraction of average female to average male wages in percent by U.S. state in 2016. The stars indicate the following significance values: * $p< 0.1$, ** $p<0.05$, *** $p<0.01$

Chapter 2

The Effect of Car Sharing on Car Sales

2.1 Introduction

Populations of large cities around the world are forecast to increase by 2.5 billion people by the year 2050[1]. Should urbanites have to rely on privately owned cars for transportation, traffic congestion, scarcity of parking as well as air pollution will cause massive welfare losses. It is therefore crucial to identify and implement at large scale alternatives to private vehicle ownership and use in large cities.

Car sharing is a promising way to enable a more efficient use of vehicles. It is part of a broader mega trend called the ”Sharing Economy”. Sharing Economy business models[2] enable access to durable goods without a transfer of ownership, thereby potentially increasing allocative efficiency. While both, Business-to-Consumer (B2C) and Peer-to-Peer (P2P) car sharing schemes have been operating for decades, advances in smartphone and internet platform technology have enabled technologically superior car sharing services to emerge in recent years. The most impactful new B2C car sharing service is called Free-Floating Car Sharing. It allows users to make one-way trips within a city and does not require them to drop the car sharing vehicle off where they picked it up, as was the case with the longer-established Station-Based Car Sharing. The new service has been picked up enthusiastically in German cities since the

[1]As estimated by the *United Nations World Urbanization Prospect*: https://esa.un.org/unpd/wup/publications/files/wup2014-highlights.pdf

[2]Well-known examples are *Airbnb* for housing, *Uber* for ride-sharing and *Turo* for Peer-to-Peer Car Sharing

early 2010s (see Figure 2.1).

This project examines whether the increase in Free-Floating Car Sharing usage has had an effect on private vehicle ownership. More specifically, it asks whether inhabitants of cities, where Free-Floating Car Sharing is offered, buy fewer new cars, because they substitute private car ownership with using Free-Floating Car Sharing. To answer this question, I collect data on launch dates and fleet sizes of Free-Floating Car Sharing providers in Germany and an original, detailed administrative panel dataset on new car registrations[3], to examine the reaction of new car sales to the introduction of Free-Floating Car Sharing. Furthermore, in order to test for heterogeneous effects across different groups of the population, I use a proprietary dataset provided by one of the Free-Floating Car Sharing companies, *DriveNow*, which contains the total number of users, as well as their age and gender in each city where *DriveNow* is active over time.

The two dominant providers of the service, *car2go* and *DriveNow*, staggered their launches and have varying fleet sizes across German cities (see Table 2.1 and Figure 2.2), creating variation in the availability and usefulness of Free-Floating Car Sharing, which allows the identification of its effect on new car sales using a difference-in-difference methodology. Exogeneity of the treatment is given, because launch dates and the timing of increases in fleet sizes were not freely chosen by the car sharing companies, but subject to regulatory approval. The car registration data is split at the year-month-city-model level, allowing for the identification of differential effects across car segments and models.

The results of this paper suggest that Free-Floating Car Sharing is indeed a substitute to private car ownership. An additional car sharing vehicle in a city leads to a decline in annual new car sales of between 2 and 4.5 cars, depending on the exact methodological specification. This decline is due to decreased sales of small, compact and medium-sized models; larger and/or more expensive models, such as SUVs, luxury cars or vans are not affected, as they serve purposes beyond pure intra-city transportation. The negative effect of Free-Floating Car Sharing on new car sales is stronger when its users are more likely to be female and older on

[3]Car registrations are used to proxy for car sales

average. This implies that in order for the effect Free-Floating Car Sharing on new car sales to be maximized, it has to appeal to groups of the population beyond its typical early adopter, who is young and male.

Both Free-Floating Car Sharing providers are affiliated with large automobile manufacturers. *car2go* is owned by *Daimler*, while *DriveNow* was founded as a joint venture between the *BMW Group* and the car rental company *Sixt*. Even though introducing new car sharing services might result in cannibalization of sales, the parent companies have several motives[4] to do so. One motive that is testable within the scope of this project is that Free-Floating Car Sharing may serve as advertisement. The vehicles used by *car2go* belong to the *smart* brand, which is also owned by *Daimler*. *DriveNow* offers its customers models of the *MINI* and *BMW* brands, both of which are owned by the *BMW Group*. Providing their vehicles to car sharing customers increases brand and model visibility. Car sharing customers obtain hands-on experience and learn about the quality of the models and brands of the car sharing vehicles. People that do not use car sharing themselves still see large numbers of these vehicles of the newest generation in and around the city center. The second question addressed by this project is therefore whether the parent companies of the Free-Floating Car Sharing providers enjoy advertising effects in the form of increased sales for the entire brand or the specific models offered through car sharing. The granularity of the dataset allows the application of a triple-difference methodology to provide robust empirical evidence of an advertising effect. Documentation and quantification of this effect are of significant interest to policy makers, who struggle with the correct regulation of new sharing economy business models, such as Free-Floating Car Sharing.

Due to the simultaneity of the substitution and the advertising effect, it is difficult to isolate the advertising effect on the absolute number of car sales. Instead, by looking at the market share of the different car models within their segment, it is possible to disentangle the two effects. The market share of one of the two models offered by *DriveNow*, the *BMW 1 Series*, indeed increases significantly in cities and at times when *DriveNow* offers its Free-Floating

[4]These include pre-emptive entry into the growing car sharing market, data collection, reducing average emissions by selling many small vehicles to their car sharing subsidiaries, improving the company image by participating in the Sharing Economy and testing new models on a young audience.

Car Sharing service[5]. Within the compact car model segment, the market share of the *BMW 1 Series* increases by 0.7% due to the presence of *DriveNow* from a baseline market share of 7.9%.

Relation to the Literature

This project contributes to the growing literature on the effects of novel Sharing Economy business models on established industries. Einav et al. (2016) provide an overview of Sharing Economy business models. Farronato & Fradkin (2017) look at the impact of *Airbnb* on hotel room occupancy. Cramer & Krueger (2016) analyze the effect of *Uber* availability on taxi usage in New York City. Seamans & Zhu (2014) quantify the effect of *Craigslist* on local newspaper ads. The paper closest to this project is Fraiberger & Sundararajan (2016), who examine the local impact of a P2P car sharing scheme on car ownership. Similarly, Gong et al. (2017) investigate the effect of the presence of *Uber* on durable good purchase decisions. Horton & Zeckhauser provide a framework for the joint determination of durable good ownership and rental in the Sharing Economy in general. In contrast, this project examines the impact of a B2C car sharing service on new car sales at the national level.

A second strand of the literature not only considers the impact of a new mobility option on an existing one, but tries to explain transportation and durable good purchase decisions in a unified framework. Gillingham et al. (2015) use Danish data to explain households driving and car purchase decisions. Brownstone & Fang (2014) use a structural model to relate residential density to car purchases of different types. Finally, Jong et al. (2004) provide an overview of early car ownership models.

Another related strand of the literature is concerned with car sharing itself. Münzel et al. (2017) provide a comprehensive description of the German car sharing market. Abhishek et al. (2016) provide a theoretical model for why a durable goods manufacturer (such as a car company) would engage in building a P2P rental market. Bellos et al. (2017) provide a theoretical rationale for why large automobile firms benefit from engaging in car sharing operations. Using BMW and Daimler as examples, they conclude that providing car sharing

[5]Due to data limitations, it is not possible to determine whether the other Free-Floating Car Sharing company, *car2go* enjoyed advertising effects as well.

services improves the fuel efficiency of the vehicles provided for car sharing. This in turn allows the car manufacturers to charge higher prices at the high end of the market, ultimately increasing profits.

Several papers consider the environmental impact of car sharing. Firnkorn & Müller discuss the environmental effects of the presence of *car2go* in Ulm, Germany. Using survey data, they quantify average emissions per car sharing user and conclude that Free-Floating Car Sharing may indeed have an impact on private car ownership. Martin & Shaheen (2016) also use *car2go* customer survey data, but in North America, to estimate that one *car2go* vehicle removes between 7 and 11 privately owned vehicles from the road. They also find that the dominant channel of reduction is not an increase in scrappage or sales by car sharing users, but delay or suppression of vehicle purchases. This indicates that examining the effect of Free-Floating Car Sharing on new car sales is indeed the correct approach to identify the effect of interest. According to the German Car Sharing association[6], one car sharing vehicle removes twenty privately owned vehicles from the streets of large German cities and car sharing users reduce private car ownership by 62%. Bert et al. (2016) forecast private vehicle sales in Germany to decline by 1.3% by the year 2021 due to the presence of car sharing. While these studies survey active car sharing users and extrapolate to the general population to estimate the impact of Free-Floating Car Sharing on private vehicle ownership, this project is the first to provide large-scale empirical revealed preference evidence.

Finally, this paper relates to other projects examining determinants of car sales. Adda & Cooper (2000) and Schiraldi (2011) estimate dynamic structural models to quantify the impact of scrapping subsidies on the car market. Grigolon et al. (2016) examine the effect of scrapping subsidies in Europe during the financial crisis. Grigolon et al. (2014) look at the impact of fuel tax changes on car purchases and the incidence of the tax. Finally, Nurski & Verboven (2016) consider the impact of exclusive dealership arrangements on car sales.

In light of the findings of the related literature, the main contribution of this paper lies in the revealed-preference evidence it provides towards the question of whether private car

[6]The report is available (in German) at https://carsharing.de/presse/pressemitteilungen/carsharing-jahresbilanz-2016-mehr-17-millionen-carsharing-nutzer (last accessed 07/03/2018)

ownership may be substituted by other forms of transportation, such as Free-Floating Car Sharing.

2.2 Identification Strategy

This section introduces the empirical strategies for identifying the effects of Free-Floating Car Sharing on new car sales. A difference-in-difference setup as described in section 2.2.1 identifies the substitution effect caused by consumers buying fewer new cars when Free-Floating Car Sharing is available. Along with the difference-in-difference setup, the triple-difference methodology in section 2.2.2 is used to investigate whether Free-Floating Car Sharing serves as advertising for particular car models or brands. For all regressions, the unit of observation is at the city-month-model level. Threats to the credible identification of the effects and how the research design of this paper addresses them are recognized in section 2.4.3.

2.2.1 Difference-in-Difference

$$y_{cts} = \alpha + \lambda_c + \lambda_t + \lambda_s + \beta \cdot X_{ct} + \delta \cdot D_{ct} + \epsilon_{cts} \tag{2.1}$$

c stands for city, t for a year-month combination and s for a car model. The outcome variable y is a measure of sales. α is a constant, the λ terms denote fixed effects and the vector X collects control variables. D is the treatment and δ is the coefficient of interest. ϵ denotes the error term.

In order to identify the effect of Free-Floating Car Sharing (the treatment) on car sales (the outcome), the regression setup of (2.1) relies on two differences. The first is geographical, as some cities were treated and others were not. The second is temporal, since all cities are observed before and after the launch of Free-Floating Car Sharing. The intuition for why this setup identifies the effect of interest is standard. The coefficient δ quantifies by how much the differences in new car sales before and after the introduction of Free-Floating Car Sharing differ between treated and control cities. By including a full set of fixed effects and additional controls that vary at the city-time level and are likely to influence new car sales, a causal interpretation of the coefficient becomes possible.

The Free-Floating Car Sharing providers staggered the launches of their services in German cities over time. This implies substantial variation in the start date of the treatment among the treated cities, improving identification with respect to the standard case, in which all units begin receiving treatment simultaneously. Furthermore, the car sharing fleet sizes differ across cities and over time (in absolute terms as well as per-capita), implying different treatment intensities, which helps quantify the coefficient of interest. Finally, in some of the treated cities only one of the two providers is active, creating further variation in treatment intensity.

2.2.2 Triple Difference

$$y_{cts} = \alpha + \lambda_{ct} + \lambda_{st} + \lambda_{cs} + \delta \cdot D_{cts} + \epsilon_{cts} \tag{2.2}$$

The notation is as in the previous subsection. Treatment now varies at the city-time-model level. To examine the effect of Free-Floating Car Sharing on sales of a particular car model, the treatment "switches on" only for observations of that car model at a time and in a city, when and where Free-Floating Car Sharing is offered. This allows the treatment to have differential effects across car models, enabling the identification of potential advertising effects for the models or brands offered through Free-Floating Car Sharing.

The λ terms in (2.2) denote interaction fixed effects; for example, λ_{ct} denote city-specific time fixed effects, so that each time period receives a different dummy variable in each city. Including a full set of interaction fixed effects allows to non-parametrically control for many potentially confounding factors.

2.3 The Data

This section describes the data used in the analysis. A more detailed description of how exactly the estimation sample is created can be found in the Appendix.

2.3.1 Car Sharing Data

Operation Launch Dates and Fleet Sizes

Information about launch dates (Table 2.1) and fleet sizes (Figure 2.2) of the two Free-Floating Car Sharing providers *car2go* and *DriveNow*, as well as which car models they use, is retrieved from public press releases and ”fact sheets”, available on the companies’ websites. Both companies offer their customers vehicles of brands which are owned by their their respective parent companies. *Daimler* owns the *smart* brand and *car2go*, which offered its customers the *smart fortwo* exclusively[7]. *DriveNow* (which is owned by the *BMW Group*) customers can choose between the *MINI Cooper* and the *BMW 1 Series*, the brands of both of which are owned by the *BMW Group* [8].

Car Sharing User Demographics

The second strand of data about Free-Floating Car Sharing in Germany comes from a proprietary dataset obtained from one of the car sharing providers, *DriveNow*. It consists of the total number of users and their age and gender by city and over time. Raw data or summary statistics may not be presented, as the dataset is proprietary. However, it is possible to state some general observations. First, the overall number of customers increases steadily over time[9]. Second, the majority of users is male. Finally, the average user is in his or her thirties.

Information about Free-Floating Car Sharing users is only available from one of the two providers in Germany, *DriveNow*. It is not necessarily the case that the demographics and total number of users of the competitor, *car2go*, are identical or even similar to those of *DriveNow*. However, in the analysis using this data (Section 2.4.1), I assume that the average age and the proportion of female Free-Floating Car Sharing users are the same across providers.

[7]Towards the end of the sample period (Summer 2016), *car2go* started offering some more up-scale *Mercedes* models in some cities. I ignore these in the estimation below, as they were only available for a very brief part of the sample period.

[8]Similarly as *car2go*, *DriveNow* started offering other models towards the end of the sample period, which I ignore here

[9]The dataset does not differentiate between active and inactive users however, so that it is possible that this trend is driven by users who sign up to Free-Floating car Sharing, but don’t actually use it.

2.3.2 Car Registration Data

New car registrations are observed monthly from January 2008 until December 2016 and are split at the model level and by registration district. Summary statistics can be found in Tables 2.3 and 2.4. As is common in the literature, new car registrations are used as a proxy for sales. The administrative process of a new car registration is necessary when a car is registered after its initial sale[10]. While a resident of a city is free to buy a new car anywhere, it has to be registered in the registration district where the resident lives. The dataset does not differentiate between private and commercial car registrations, so that for example the registrations by the Free-Floating Car Sharing providers are visible in the data (for an example see Figure 2.4). In the analysis, I drop the models used by the car sharing companies. As described in more detail in the appendix, for the cities in the estimation sample, city boundaries and registration district borders always coincide[11].

2.3.3 City Demographics

Demographic information by city and year is publicly available from the German Federal Statistical Office. I use the following variables as controls in some regressions: population size, percentage of the population that is male, average income, average disposable income and price per m^2 of land. Table 2.2 shows summary statistics for these control variables.

2.4 Results

2.4.1 Substitution Effect

Table 2.5 shows the difference-in-difference estimation results for the effect of Free-Floating Car Sharing on new car sales for different definitions of the treatment. In column (1), the treatment is a dummy variable equal to one, if one or both of the Free-Floating Car Sharing providers are active. The coefficient implies a reduction of monthly new car sales at the model

[10] The dataset contains no information about used car sales.
[11] More rural registration districts sometimes contain several smaller cities.

level of roughly one in the presence of Free-Floating Car Sharing. For the average number of car sharing vehicles and the average number of models available for sale, this implies that the average marginal annual effect of one Free-Floating Car Sharing vehicle on new car sales is a reduction of three fewer vehicles sold. This simple definition of treatment does not use differences in treatment intensity over time or across cities and the coefficient is statistically significant only at the 10% level.

In column (2), treatment is defined as the number of available Free-Floating Car Sharing vehicles. The treatment varies within cities over time, as well as across cities (see Figure 2.2). Using this additional information results in statistical significance of the treatment effect at the 1% level. The estimated average marginal effect of a car sharing vehicle on annual sales increases in absolute value to four fewer vehicles sold.

The availability and usefulness of Free-Floating Car Sharing does not only depend on the absolute number of available vehicles, but also on how large the area is, in which they can be driven, as well as on how many people share the vehicles. Therefore, in columns (3) and (4), treatment is defined as the number of available Free-Floating Car Sharing normalized by the size of the operating area of the car sharing providers and by population size, respectively. For these two definitions, the estimated effect is also highly statistically significant and the implied average marginal effect of a car sharing vehicle on annual new cars sales increases (in absolute terms) to -4.7 and -4.1 vehicles, respectively. Few of the controls are statistically significant. The sample cities did not experience large demographic changes during the sample period, so that there is little variation in the controls. Therefore, demographic influences are captured by the city-fixed effects.

The results are robust to several robustness checks (see Table 2.6), such as using the logarithm of new car sales as an outcome variable. Further, following Besley & Burgess (2004), I include city-specific linear time trends. City-specific linear trends pick up linear developments in new car sales which differ across cities. Including them in the regression allows treatment and control states to follow differential trends in a limited way (Angrist & Pischke (2009)). The results are robust to their incusion.

A concern is that treatment and control cities are not directly comparable, but differ in terms of new car sales in an important, but unobserved way. In order to address this concern, the regression is re-run after dropping the control cities from the sample. This is possible because of the variation in the launch dates of Free-Floating Car Sharing in the treated cities. The analysis using only treated cities relies on this variation in launch dates to identify the effect. The results remain significant and quantitatively similar, even when city-specific trends are introduced in the regresion using only treated cities. Depending on the exact specification, the estimated annual average effect of one car sharing vehicle on new car sales is between 1.7 and 4.5 fewer vehicles sold.

Effects by Car Segment

While all car models serve transportation needs, the purposes of different models still differ. A smaller, more agile car may be preferable for inner city trips, as maneuvering in tight spaces and parking are easier. For longer trips outside the city center, larger limousines or SUVs may be preferred. The effect of the introduction of Free-Floating Car Sharing on new car sales may therefore vary across the different car model segments.

Table 2.8 examines the possibility of differential effects of Free-Floating Car Sharing on new car sales of different segments. Each column presents the difference-in-difference results using the observations from one car segment. The introduction of Free-Floating Car Sharing only has statistically significant impacts on sales of car models in the *small*, *compact* and *mid-sized* segments, with the biggest part of the economic impact coming from reductions in sales of *small* and *mid-sized* Models. There is no statistically significant effect on *SUVs*, *luxury* models and *vans*. Tables 2.3 and 2.4 show the relative importance of the segments: the *small*, *compact* and *mid-sized* segments are the most popular segments both in control and treated cities.

Heterogeneous Effects by Car Sharing User Demographics

The previous subsection showed which types of cars are being bought less due to the availability of Free-Floating Car Sharing. A similar question is what kind of consumers forego the purchase of a new vehicle and use Free-Floating Car Sharing instead. In order to provide evidence on this empirical question, this subsection uses a proprietary dataset from one of the Free-Floating Car Sharing providers, *DriveNow*, which contains information on age and gender of the companies' customers by city and over time.

In order to allow for differential effects of Free-Floating Car Sharing depending on the characteristics of the customers, the baseline difference-in-difference setup introduced in equation (2.1) is augmented by interaction terms of the treatment with demographic variables. More formally, the regressions allowing for heterogeneous effects will have the following structure:

$$y_{cts} = \alpha + \lambda_c + \lambda_t + \lambda_s + \beta \cdot X_{ct} + \delta \cdot D_{ct} + \zeta \cdot \{\mathit{Frac\_Female}\}_{ct} + \gamma \cdot D_{ct} \cdot \{\mathit{Frac\_Female}\}_{ct} + \epsilon_{cts} \quad (2.3)$$

and

$$y_{cts} = \alpha + \lambda_c + \lambda_t + \lambda_s + \beta \cdot X_{ct} + \delta \cdot D_{ct} + \zeta \cdot \{\mathit{Mean\_Age}\}_{ct} + \gamma \cdot D_{ct} \cdot \{\mathit{Mean\_Age}\}_{ct} + \epsilon_{cts} \quad (2.4)$$

Frac_Female and *Mean_Age* refer to the fraction of female users and the average age of all users of Free-Floating Car Sharing in a city at a time, respectively. In these regressions, the baseline effects of treatment and demographic variable are given by δ and ζ, while the interaction effects are given by γ. These can be interpreted as how the effect of Free-Floating Car Sharing on new car sales changes for a given treatment intensity D_{ct} and varying proportions of female car sharing users and varying average age of car sharing users.

The estimation results of equations (2.3) and (2.4) are in Table 2.9. Columns (1) and (2) of both panels show that the inclusion of only the demographic variables does not change the baseline results. Column (4) of Panel A however shows that for a given number of available Free-Floating Car Sharing vehicles, the effect on new car sales becomes more negative with an increase in the fraction of female car sharing users. Similarly, column (4) of Panel B shows that the negative effect of Free-Floating Car Sharing on new car sales becomes larger with an older average age of the car sharing users. Both of these effects are strongly statistically significant.

At this point it is important to note that differences in the demographics of the car sharing populations may be endogenous, so that the results in Table 2.9 cannot be interpreted causally. Nonetheless, the correlations suggest that the effect of Free-Floating Car Sharing on new car sales is stronger, when the car sharing service is used by more women and older drivers. The implication is that in order to maximize the effect of Free-Floating Car Sharing on car sales, its use should be promoted among those parts of the population, which are under-represented among car sharing users: female and older drivers.

2.4.2 Advertising Effect

This subsection examines whether there is an advertising effect in the form of increased sales for the brands and models used by the Free-Floating Car Sharing providers.

Model-Level Advertising

Using a difference-in-difference regression setup, the impact of the presence of *DriveNow* on the sales of the models offered by *DriveNow* can be measured[12]. In each column of Table 2.10, the sample is restricted to observations of sales of the indicated car model. The treatment variable is a dummy indicating the presence of *DriveNow*[13] Columns (1) and (2) show no significant impact of the presence of *DriveNow* on the total number of sales of either of the two models offered by *DriveNow*. However, this does not allow the conclusion that there is no positive advertising effect, since the negative substitution effect documented in subsection 2.4.1 is acting simultaneously.

It is more instructive to examine the relative market shares of the *DriveNow* models within their segments. While overall sales of a model may fall due to the previously documented

[12] A potential advertising effect of *car2go* can not be examined, because it is not possible to disentangle private registrations from those made by *car2go* itself. Please see the sample details in the appendix for more information.

[13] It is not possible to construct a measure of treatment intensity at the model level. This is because while the overall number of cars provided by *DriveNow* in a city is known over time, the composition of the fleets (i.e. how many of the cars are *MINIs* and how many are *BMW 1 Series*) is unknown.

substitution effect caused by Free-Floating Car Sharing, sales of the *DriveNow* models may be flat or falling less strongly, which would increase their market share within the segment. Indeed, using the model market share within its respective segment as an outcome variable in columns (3) and (4) reveals that *BMW 1 Series* vehicles became relatively more popular when *DriveNow* was operating. The market share within the *Compact* car segment increased by around 0.7%, which is substantial as the average within segment market share of the *BMW 1 Series* was 7.9% before treatment commenced. For the *MINI*, there is no substantial change in market share.

In order to exclude the possibility that this result is driven by an alternative trend, such as an increased preference for premium compact cars in treated cities which coincides with the launch of *DriveNow*, I check for effects on the two closest competitors of the *BMW 1 Series*. Columns (5) and (6) of Table 2.10 show the effect of *DriveNow* on the segment market shares of the *Mercedes A-Class* and the *Audi A3*. Column (5) shows that sales of the compact *Mercedes* model *A-Class* suffered disproportionately from the introduction of *DriveNow*, while the market share of the *Audi A3* did not change significantly.

Since the dataset is split at the year-month-city-model level, an alternative way to check for the presence of an advertising effect of Free-Floating Car Sharing is to introduce a third difference. Above, identification of the effects relied on exploiting differences across time (before vs. after the introduction of Free-Floating Car Sharing) and location (treated vs. untreated citites). Table 2.11 presents the results of a triple-difference regression, where the third difference is across car models (offered through Free-Floating Car Sharing vs. not offered).

Even with the large number of fixed effects used in the regression, the triple-difference approach confirms the findings described above. No significant effect on the absolute number of sales is found for either *DriveNow* model. However, the market share of the *BMW 1 Series* within the *Compact* segment is estimated to increase by around 0.7% due to the presence of *DriveNow*. Among the two closest substitutes to the *BMW 1 Series*, the *Mercedes A-Class* suffered a significant loss in market share, while the share of *Audi A3*s was not affected.

Brand-Level Advertising

While Free-Floating Car Sharing enables customers to obtain hands-on experience with the particular models the providers choose to offer, it also increases the visibility of the brand of the models. Furthermore, car sharing users form opinions about the quality of the model they're driving and potentially extrapolate towards other models of the same brand. Such a spillover advertising effect would manifest itself in the form of increased sales of *BMW* models not offered through Free-Floating Car Sharing in the presence of *DriveNow*. Table 2.12 presents results of a difference-in-difference regression, which compares sales trends of *BMW* models in different segments around the introduction of *DriveNow* in treated and control cities.

Absolute sales in all three segments fall in the presence of *DriveNow*, but not significantly so for the *luxury* segment. For this segment, the market share of the *BMW* models increases by almost 1.3%, when and where *DriveNow* is offered. This result is statistically significant at the 5% level. *BMW* models in the *mid-sized* and *SUV* segments do not gain or lose market share significantly.

It is surprising to see a positive and significant effect of the presence of *DriveNow* on sales of luxury *BMW* models, as one would expect sales dynamics of these very expensive vehicles to be driven by other factors. To check whether this result is a statistical artifact or an actual effect, Table 2.13 shows the results of a triple-difference regression of sales and within-segment market shares on the presence of *DriveNow*. The treatment variable is a dummy equal to one for observations of sales of *BMW* models when and where *DriveNow* is active. This regression setup compares the effect of *DriveNow* on sales of *BMW* models to its effect on sales of other brands. The results are similar to the results of the difference-in-difference regression, but the effect of *DriveNow* on the market shares of *BMW* luxury models is no longer statistically significant. Overall, there is no robust evidence for the presence of a spillover advertising effect of *DriveNow* on *BMW* models not offered by the *DriveNow* car sharing service.

2.4.3 Addressing Threats to Identification

Visual Inspection of Trends

The staggered launch dates of Free-Floating Car Sharing across treated cities, the fact that there are two providers that start offering their services at different times and the changes in treatment intensity over time create substantial variation in the data and imply that the current setting lends itself very well to an analysis in a difference-in-difference regression framework. However, for the same reasons, visual inspection of trends, in particular a comparison of treated and control cities around the beginning of treatment is not as straightforward as in the textbook case. Nonetheless, Figure 2.5 shows the average (across car models and cities) residuals of a regression of new car sales on year-month dummies (as in equation (2.5)), as well as the dates at which Free-Floating Car Sharing was first offered in each treated city (indicated by vertical lines).

$$y_{cts} = \alpha + \lambda_t + \epsilon_{cts} \tag{2.5}$$

While sales are flat in the control cities throughout the sample period, there is a decline in sales in the years following the introduction of Free-Floating Car Sharing. In Figure 2.6, the launch months of Free-Floating Car Sharing are normalized to zero, and the average residuals for the treated cities are plotted around the launch date. There is a visible decline in new car sales, although it is not immediate, but occurs around one year after Free-Floating Car Sharing was first offered.

Threats to Identifying the Substitution Effect

The assignment of treatment with Free-Floating Car Sharing is the result of decisions made by the car sharing providers. If these companies freely decide which cities to enter, when and in which order to enter them, and how many cars to provide in each city at each point in time, the assumption of exogenous treatment assignment is potentially violated. This would be the case, if the car sharing providers choose the treated cities on the basis that their populations are more likely to substitute purchasing a new car with using Free-Floating Car Sharing. However, there is substantial evidence that the car sharing companies are not freely or strategically choosing along any of the three dimensions (which cities to enter, when to enter and how many vehicles to provide). First, both providers are only active in the largest German cities, indicating that population size, not strategic behavior drives the choice of which cities to operate in. To further alleviate concerns that treated cities are somehow different from the control group, in a robustness check (Table 2.6) all control group cities are dropped and the results are reproduced using only the variation in timing of entry. Second, the cities where Free-Floating Car Sharing was first launched (see Table 2.1) are the geographically closest large cities to the headquarters of the providers' parent companies (Ulm for *car2go*, Munich for *DriveNow*). This, along with substantial anecdotal evidence[14], suggests that logistic factors and political connectedness drive the decisions over the timing of entry into the cities. Finally, there is evidence that the Free-Floating Car Sharing providers are not free to provide as many vehicles as they desire, but only as many as local governments allow[15].

Differential trends in new car sales between treated and control cities would invalidate a difference-in-difference approach. To ensure that pre-treatment trends are parallel, in addition to the visual inspection described above, I reproduce the results using only observations from treated cities (this is possible due to the variation in entry timing) and include city-specific time trends (see Table 2.6) and include leads and lags of the treatment (see Table 2.7). The results remain qualitatively the same and quantitatively very similar to the baseline regression results.

Including city-specific time trends also helps alleviate concerns about another potentially confounding factor of the results: improvements in other forms of transportation[16] in treated cities, which coincide with the launch of Free-Floating Car Sharing. The results in Table 2.6

[14]For example, plans of *DriveNow* to enter Frankfurt were prevented by local government, as documented (in German) here: https://www.focus.de/regional/muenchen/auto-bmw-sixt-carsharing-expandiert-im-ausland_id_4509296.html (last accessed 26/02/2018)

[15]This newspaper article documents (in German) a 2015 decision by the Munich city council to increase the number of allowed car sharing vehicles https://www.carsharing-news.de/stadtrat-muenchen-carsharing/ (last accessed 26/02/2018)

[16]Other forms of transportation that are potential substitutes to privately owning a car are public transportation, station-based car sharing and bike sharing. Station-Based Car Sharing is present in treated and in control cities. Differential trends across the two groups are picked up by the city-specific time trends. Ride-sharing services, such as *Uber* or *Lyft* never operated at scale in Germany during the sample period.

confirm that the documented substitution effect is in fact due to Free-Floating Car Sharing.

Threats to Identifying the Advertising Effect

Identifying an advertising effect in the form of increased sales of the models or brands used by the Free-Floating Car Sharing providers is complicated by the simultaneous presence of the substitution effect. An effect on the absolute number of sales may thus be hard to quantify, as the two effects potentially offset each other. In order to disentangle them, I use the market shares of the car models within a city and segment as an outcome variable. This allows the identification of an advertising effect even in an environment, in which aggregate new car sales are falling due to the substitution effect.

Since treated cities are larger and denser than control cities, buyers in treated cities may have a pre-existing preference for buying the smaller models offered through Free-Floating Car Sharing. The city-model fixed effects in the triple-difference regressions control for such persistent local taste shocks non-parametrically. Similarly, some models may become more popular everywhere at certain points in time, for example because a new generation of a car model becomes available, or because of national advertising campaigns[17]. Such confounding factors are picked up by the model-time fixed effects.

Given that the parent companies (*Daimler* for *car2go*, *BMW* for *DriveNow*) are able to substantially influence the behavior of their car sharing providers, they may engage in local advertising campaigns or local price discounts accompanying the launch of the car sharing services. Local increases in new car sales of the models used by car sharing providers after their launches could then be driven not by car sharing customers learning about the models, but by the general population being more exposed to advertising and taking advantage of discounts. However, most car advertising likely takes place in national print media, television or online and none of these channels are local[18]. Neither are local discounts likely to play a substantial role, as customers from other cities cannot be excluded from taking advantage of

[17] Releases of new generations of a model as well as advertising campaigns are controlled by the parent companies of the Free-Floating Car Sharing providers, so it is conceivable that these decisions are taken jointly with decisions over car sharing activities.

[18] With the exception of local events and billboards, data on which is difficult to obtain.

price discounts in treated cities.

2.4.4 Discussion of Results

Putting an additional Free-Floating Car Sharing vehicle into a city is estimated to lower annual new car sales by between 1.7 and 4.7, depending on the exact specification. A back-of-the-envelope calculation[19] suggests that 1.5% of new car sales were avoided due to Free-Floating Car Sharing in the treated cities since the first Free-Floating Car Sharing launch. While this number suggests that the novel service has not changed the new car market fundamentally, the estimated impact is already bigger than Bert et. al (2016) predicted it to be in 2021.

Section 2.4.1 shows that the reduction in new car sales following the introduction of Free-Floating Car Sharing is driven by reductions in sales of small, compact and mid-sized vehicles. This result is intuitive; Free-Floating Car Sharing is a close substitute to the function of a small city car, as it enables the driver to make short intra-city trips with a vehicle that is easy to maneuver and park in small spaces. Free-Floating Car Sharing does not have an effect on SUVs, luxury cars or vans. These larger vehicles, for example SUVs, are not as useful in dense cities and more often driven outside the city center or for inter-city trips, both of which is not possible or convenient with Free-Floating Car Sharing. Luxury vehicles serve functions beyond pure transportation needs, for example as status symbols. Finally, vans are convenient for transporting many people or large items, both of which is impossible with the small vehicles provided through Free-Floating Car Sharing.

The results of section 2.4.1 suggest that the effect of Free-Floating Car Sharing on new car sales increases if the car sharing users are older and more likely to be female. It is not possible to interpret these results causally, as the demographics of the car sharing users are not exogenously imposed. Nonetheless, the correlations suggest that when Free-Floating Car Sharing manages to appeal to a broader population of customers, its effects on new car sales

[19]For each city and for each month, during which Free-Floating Car Sharing is offered, I multiply the estimated average marginal effect with the number of available car sharing vehicles in the treated cities, and sum across cities and months to obtain an estimate for how many total car purchases were avoided during the sample period. To approximate the number of new car sales that would have taken place in the absence of Free-Floating Car Sharing, for each city I multiply the number of periods in which the service was offered with the average number of new car sales before Free-Floating Car Sharing was launched.

are maximized.

The difference-in-difference and the triple difference methodologies delivered qualitatively identical and quantitatively very similar estimates of the model-level advertising effect, allowing to confidently conclude that the presence of *DriveNow* indeed has a positive advertising effect for the *BMW 1 Series*, reflected in an increased number of sales relative to competitor's models. Quantifying the value of the advertising effect for the *BMW Group* is not straightforward, since *DriveNow* does not increase sales of the *BMW 1 Series* in absolute terms, but only increases the model's market share. It is however possible to conduct another back-of-the-envelope calculation to compare the added value the *BMW Group* incurred by offering the *BMW 1 Series* to a counterfactual scenario, in which another car company provides the Free-Floating Car Sharing vehicles and hence where there is no positive advertising effect for the *BMW 1 Series* model. The calculation[20] implies an annual profit increase of around 1.01M€ for the *BMW Group* due to the increase in market share caused by the presence of *DriveNow*. To put this number into perspective, the annual report of the *BMW Group*[21] indicates a financial loss of *DriveNow* for 2015 of 6M€ and for 2016 of 15M€. The documented advertising effect does not make up for the entire financial loss to the *BMW Group*, but for a sizable portion of it. In this context it is interesting to note that in late 2016 and 2017, both *Mercedes* and *Audi*, the two closest competitors of *BMW*, started making their own models available through Free-Floating Car Sharing as well (*Mercedes* through *car2go* and *Audi* through a new service called *Drive by*). Unfortunately, these developments lie outside the sample period.

2.5 Conclusion

Using an original administrative dataset and exploiting the rollout pattern of Free-Floating Car Sharing in Germany with a difference-in-difference regression setup, this paper provides

[20] Table 2.11 shows that the presence of *DriveNow* increases the market share of the *BMW 1 Series* in the *compact* segment by 0.74%. Using average sales in the segment in treated cities when *DriveNow* is active, this is equivalent to an increase in annual sales of 375. Assuming a profit margin of 10% (as estimated by Cosar et al. (2018)) and a list price of 27 000€ for a new *BMW 1 Series* (this number was suggested by the *BMW* website as a conservative estimate) results in an annual profit increase of 1.01M€.

[21] https://www.bmwgroup.com/content/dam/bmw-group-websites/bmwgroup_com/ir/downloads/en/2016/BMW_GB16_en_Finanzbericht.pdf (last accessed 07/03/2018)

revealed-preference evidence for a negative effect of the introduction of Free-Floating Car Sharing on new car sales. One Free-Floating Car Sharing vehicle reduces annual new car sales by between 2 and 4.5 vehicles. This effect is driven by a reduction in sales of small, compact and mid-sized car models and is stronger when the car sharing users are older and more likely to be female.

The two Free-Floating Car Sharing providers are owned by large automobile manufacturers, and the car sharing users are offered vehicles produced by the parent companies. Using a difference-in-difference and a triple difference methodology, I find robust evidence that the market share of one of the car models increases, when the model is available for car sharing. This advertising effect of Free-Floating Car Sharing results in an economically sizable profit increase for the car manufacturer.

Policy makers in large cities are interested in reducing private car ownership, because they see it as a means towards goals such as alleviating scarcity of parking and traffic congestion, and to improve air quality. This paper shows that Free-Floating Car Sharing does indeed cause new car sales to fall, but whether it contributes to reaching the ultimate goals is not clear. Its users may well substitute walking short distances with using a car sharing vehicle, so that congestion and air quality do not necessarily improve. This is because although there are now fewer cars in a city, the car sharing vehicles are used more intensively. In certain cirumstances, Free-Floating Car Sharing may even make parking more scarce, for example because people that used to take public transportation to the city center on a weekend are now looking for a place to park their car sharing vehicle. More detailed data on trips and driving behavior of car sharing users is needed to address these potentially ambiguous effects of Free-Floating Car Sharing on aspects of city life, which are of significant importance to policy makers.

2.6 Appendix

A1: Estimation Sample

The utility of free-floating car sharing to a consumer hinges on whether there is an available vehicle within an acceptable walking distance when one is needed. The utility also increases with the number of desirable destinations within the operating area of the provider. Together, these factor imply that free-floating car sharing requires dense cities with a large enough population[22] According to a *DriveNow* executive[23], Free-Floating Car Sharing functions well only in cities with more than 500 000 inhabitants[24]. Therefore, the estimation sample only contains the 15 German cities with more than 500 000 inhabitants. For these cities, the registration districts coincide with city boundaries. The cities in the sample where free floating car sharing was never offered during the sample period are Bremen, Dortmund, Dresden, Essen, Hanover, Leipzig, Nuremberg and Recklinghausen. The cities in the sample were free-floating car sharing was launched at some point during the sample period are Berlin, Düsseldorf, Frankfurt, Hamburg, Cologne, and Stuttgart [25].

The car registration data does not differentiate between private and corporate car registrations. Consequentially, registrations of the car models used by *car2go* and *DriveNow* spike around the launch dates of the free-floating car sharing services (see Figure 2.4) and are elevated afterwards (as *car2go* and *DriveNow* continuously register replacement vehicles for their fleets. In order to ensure these company registrations don't distort the estimation of the substitution effect between buying a new car and using car sharing, I drop all models used by *car2go* and *DriveNow* from the sample. Since many large car rental companies, as well as *DriveNow*, register all of their vehicles in Munich, I drop Munich from the sample, to ensure that these corporate registrations don't mechanically drive the results. *car2go* registers all its vehicles in

[22]This is a potential explanation for why free floating car sharing has been more widely adopted in Europe that in the United States. Cities in the U.S. tend to be sprawling and often lack a well-defined city center.

[23]The interview can be found (in German) at http://www.carsharing-experten.de/drivenow-carsharing/drivenow-geschaeftsfueher-dr-andreas-schaaf-im-interview.html (last accessed 26/02/2018)

[24]Bert et al. (2016) also consider this number as the correct threshold.

[25]*car2go* tested free-floating car sharing among *Mercedes* employees and later offered it to the general population in Ulm from October 2010 until December 2014. As Ulm has fewer than 500 000 inhabitants, I drop it from the sample.

the cities where they are used (as can be seen in Figure 2.4), so it is not possible to estimate an advertising effect for the *smart* model.

The car models are assigned to segments following the methodology of the German Department of Motor Vehicles (*Kraftfahrtbundesamt Flensburg*).

Finally, in order to avoid the results being driven by outliers in car registrations, which are likely to be the result of corporate fleet registrations, for each car model, I replace all registration numbers above the 99th percentile with the 99th percentile.

A2: Additional Robustness Checks

In order to address some of the concerns raised in Subsection 2.4.3, Table 2.6 contains further difference-in-difference estimation results. In all columns of Table 2.6, treatment is defined as the number of available car sharing vehicles (as in column (2) of Table 2.5). The outcome variable is either monthly sales or the natural logarithm of monthly sales at the year-month-city-model level.

To ensure that the results are not driven by outliers, the outcome variable of the specification in column (1) is taken to be the logarithm of monthly new car sales. While the estimated average marginal effect of a free-floating car sharing vehicle on annual new car sales shrinks in absolute value, the effect remains statistically significant at the 1% level.

In order to exclude the possibility that results are driven by trends other than the introduction of free-floating car sharing, city-specific linear time trends are included in the specifications in columns (2), (3), (6) and (7). These trends vary at the same level as the treatment variable and would capture city-specific changes over time, that could threaten identification of the treatment effect of free-floating car sharing. Examples for such changes are improvements in public transport, bike sharing or station-based car sharing that coincide with the introduction of free-floating car sharing. The estimated treatment effect of free-floating car sharing is not driven by other city-specific changes however, as its estimated treatment effect remains statistically and economically significant in the specifications including city-specific linear trends.

Another potential threat to identification of the treatment effect is endogeneity in the treatment assignment. The free-floating car sharing providers might provide more vehicles in cities where they expect car ownership to decline faster, leading to non-random treatment intensity assignment. This concern is alleviated by the result of column (1) in Table (2.5), which shows a significant treatment effect even when the potentially endogenous decision over how many cars to provide in a city is excluded from the analysis. Not only the decision over how many cars to provide, but also the decision in which cities to operate potentially threatens the non-random treatment assignment assumption of the difference-in-difference identification strategy. In order to address this concern, the specifications in columns (4) to (7) of Table (2.6) drop from the sample those cities that were never treated with free-floating car sharing. Identification in these specifications comes only from the staggered entry dates of the free-floating car sharing providers (see Table (2.1)). The estimated treatment effect remains statistically significant and economically similar to specifications using the full sample, alleviating concerns about non-random treatment assignment[26].

In order to ensure that results are not driven by pre-existing trends, Table (2.7) contains results of regressions including leads and lags of the treatment. For all specifications, the effect of the treatment remains statistically significant at the 1% level, while the economic significance stays roughly the same compared to specifications without leads and lags of the treatment.

A3: Additional Information about Free-Floating Car Sharing

Free-Floating Car Sharing vehicles are not located in specific places, but are scattered across the area of a city. Any vacant car may be rented by a customer. The car key is inside the vehicle, which is opened and locked using a smartphone app[27]. Customers may drive to any destination within the Free-Floating Car Sharing operating area and park the vehicle in any available parking spot (public, residential, or paid). Customers usually pay by the minute, plus an initial, one-time registration fee. There is no substantial variation in the price paid

[26] See also the discussion in Subsection 2.4.3

[27] Figure 2.3 shows a screenshot of the *car2go* app. Available vehicles are indicated by the blue shapes.

by consumers for Free-Floating Car Sharing over time or across cities. The cities charge the Free-Floating Car Sharing providers for these parking privileges[28].

The main advantage of Free-Floating Car Sharing over other forms of car sharing lies in its convenience. A customer may use a Free-Floating Car Sharing vehicle to get from point A to point B and later back to point A. However, since she will do so in two separate trips, she will not be charged a rental fee for the duration of her stay at point B. Furthermore, she does not need to worry about the cost of parking during her stay at point B, since she may park the Free-Floating Car Sharing vehicle in any parking spot at no cost to her. Especially in dense urban areas, the combination of these two characteristics makes Free-Floating Car Sharing a better alternative to privately owning a car than existing car sharing schemes. Of course, the convenience of Free-Floating Car Sharing hinges on the probability of being able to find an available vehicle within an acceptable distance.

[28]The city of Munich, for example, charges *DriveNow* 900€ annually per car for parking, as reported (in German) here: https://www.carsharing-news.de/stadtrat-muenchen-carsharing/ (last accessed (03.07.2018)

2.7 Figures

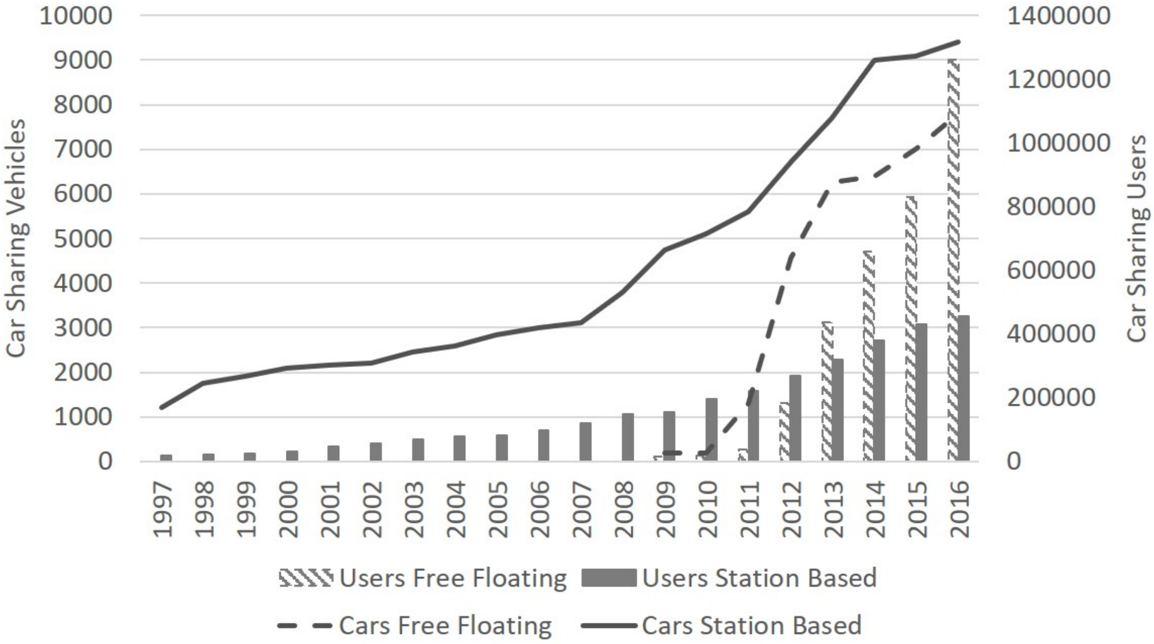

Figure 2.1: Available Car Sharing Vehicles (left axis) and Car Sharing Users (right axis) for Station-Based and Free-Floating Car Sharing in Germany over Time.

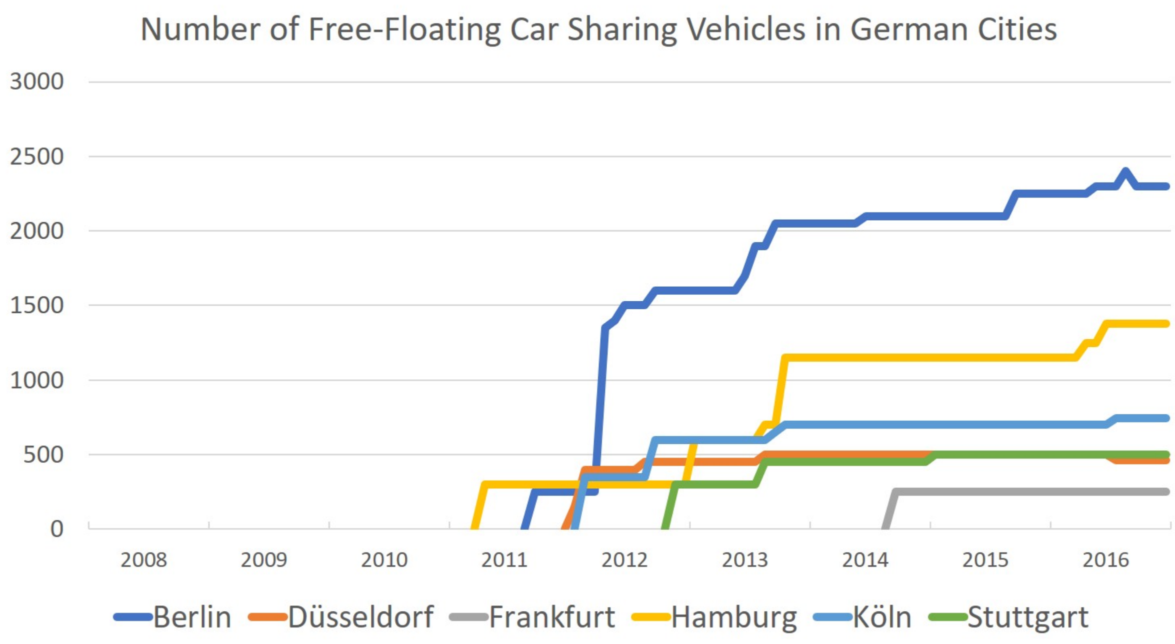

Figure 2.2: Combined number of available free-floating car sharing vehicles in the treated cities over time.

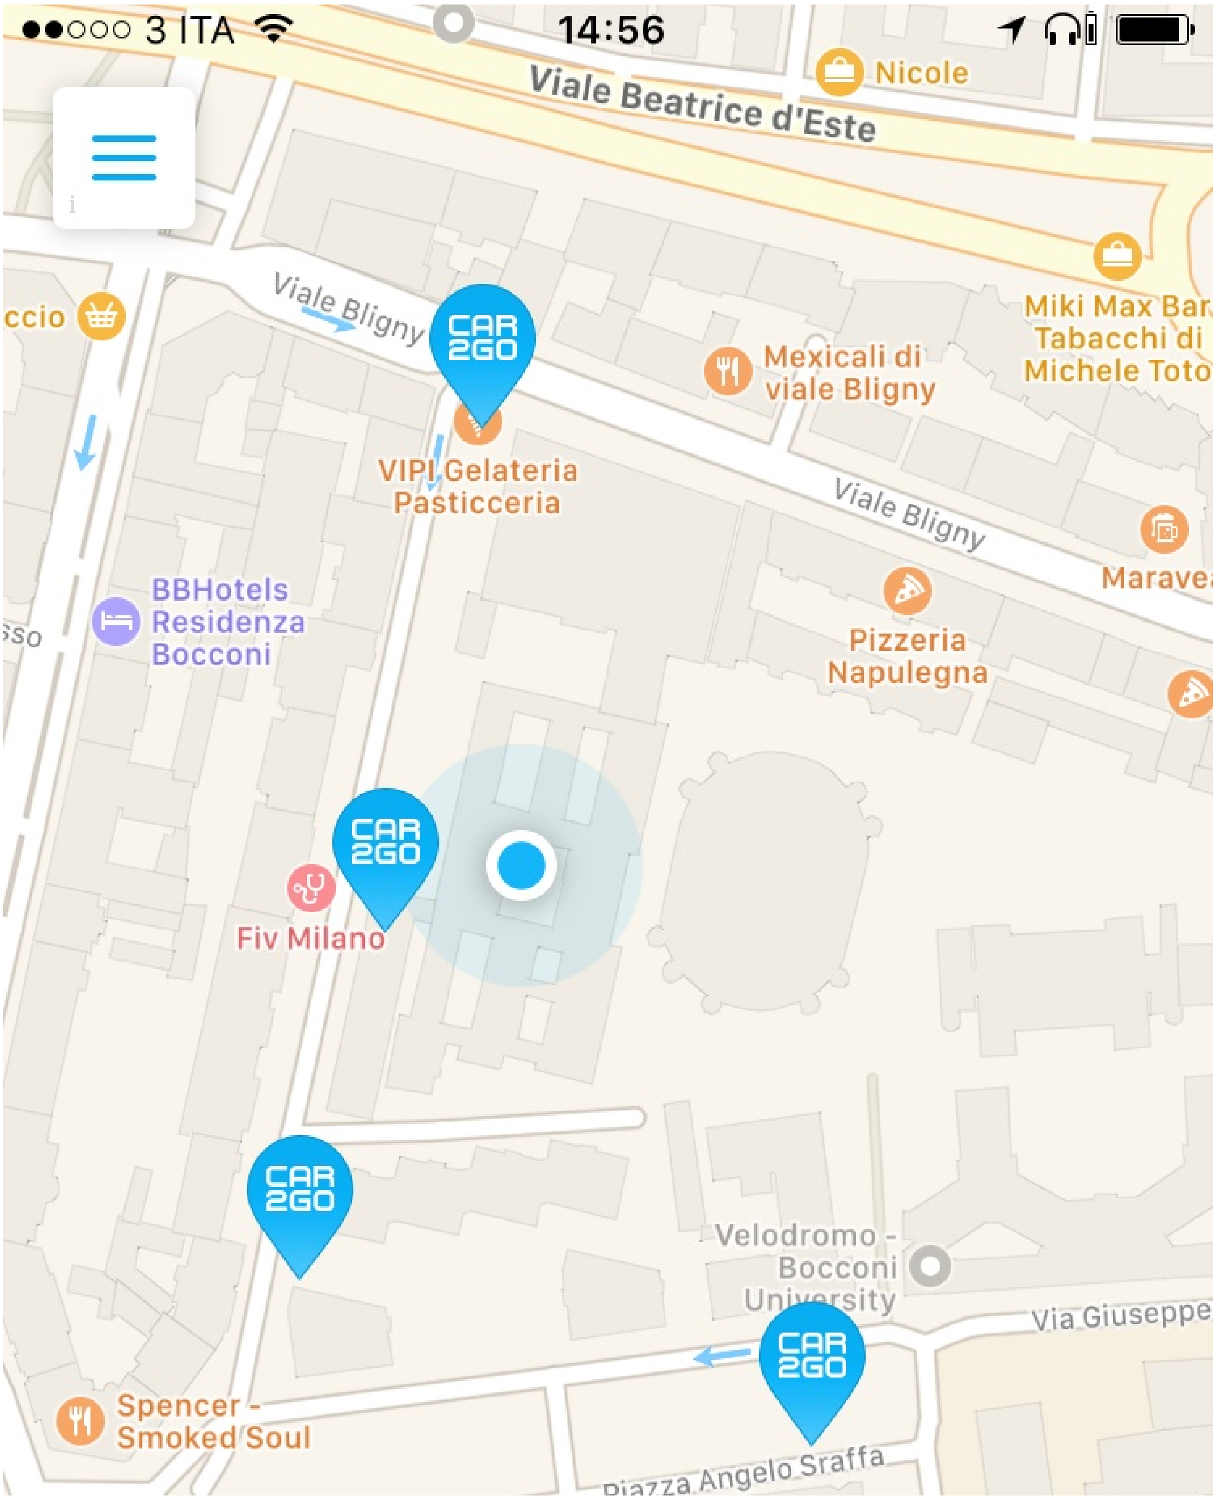

Figure 2.3: Screenshot of the *car2go* App.

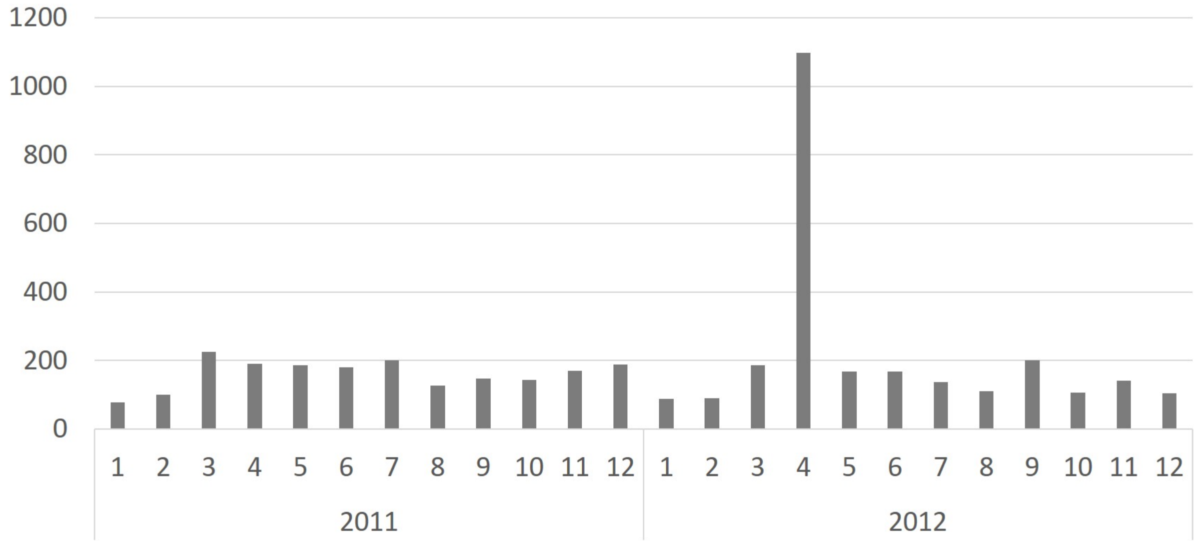

Figure 2.4: *smart fortwo* sales in Berlin 2011 and 2012. Registrations spike in April 2012, when *car2go* launched its Berlin operation with 1000 *smarts*.

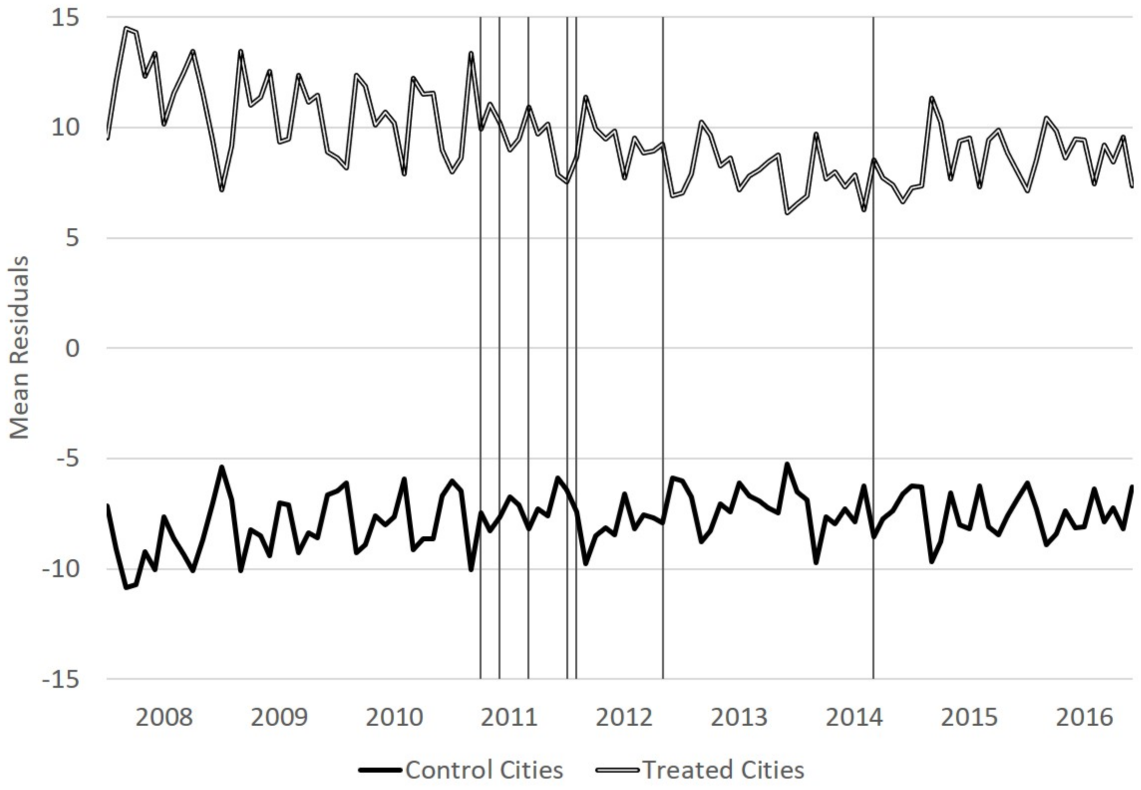

Figure 2.5: Mean residuals of a regression of new car sales on year-month dummies for treated and control cities. The vertical lines indicate the launch dates of the first Free-Floating Car Sharing provider in each treated city.

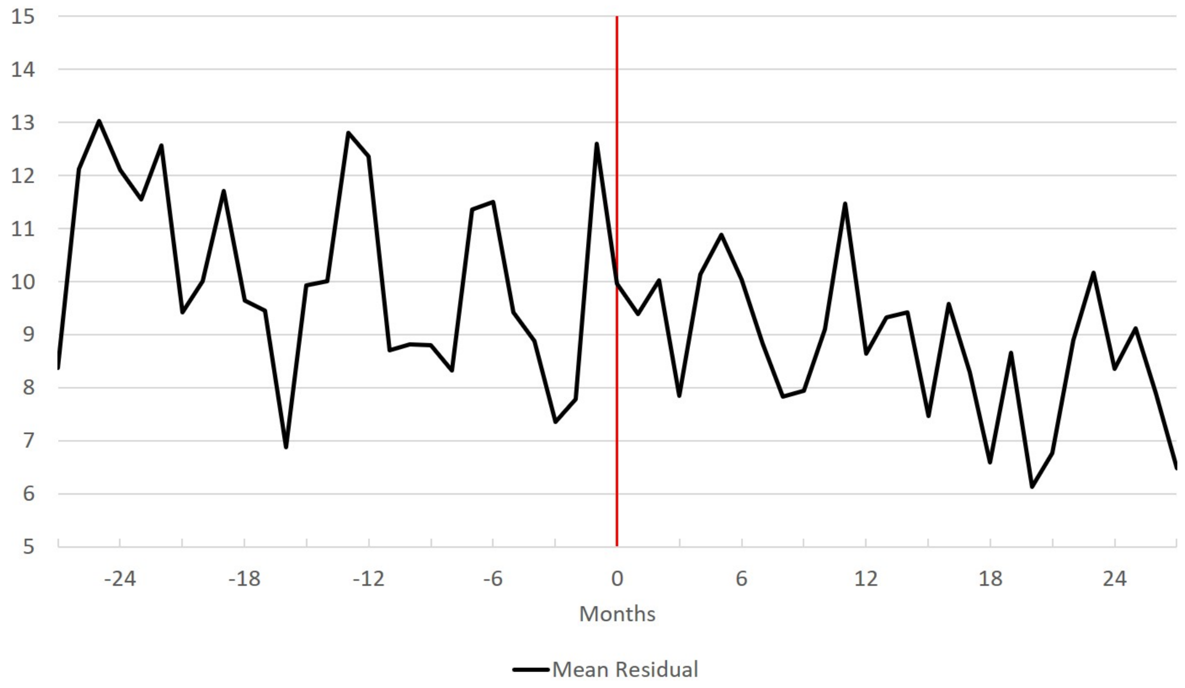

Figure 2.6: The month in which Free-Floating Car Sharing was first offered in a treated city is normalized to zero and indicated by the vertical red line. The graph shows the mean residuals of a regression of new car sales on year-month dummies around the launch date of the first Free-Floating Car Sharing provider in the treated cities.

2.8 Tables

Table 2.1: Launch dates, fleet sizes and sizes of operating areas (as of December 2016) of Free-Floating Car Sharing providers in German cities

	car2go	**DriveNow**
Berlin	04/2012 (1100 cars, 160km^2)	09/2011 (1200 cars, 161km^2)
Düsseldorf	02/2012 (250 cars, 80km^2)	01/2012 (210 cars, 55km^2)
Frankfurt	09/2014 (250 cars, 61km^2)	
Hamburg	04/2011 (780 cars, 102km^2)	10/2013 (580 cars, 90km^2)
Cologne	02/2012 (350 cars, 130km^2)	09/2012 (400 cars, 90km^2)
Munich	06/2013 (480 cars, 125km^2)	06/2011 (700 cars, 85km^2)
Stuttgart	11/2012 (490 cars, 153km^2)	
Ulm	10/2008 - 12/2014 (200 cars, 70km^2)	

car2go ran a test scheme for its free-floating car scharing service in Ulm before opening it up to the public. Service was discontinued at the end of 2014. According to the management (as cited in this German newspaper article:
https://www.swp.de/suedwesten/staedte/ulm/car2go-macht-in-ulm-dicht-pilotstadt-war-zu-klein-und-zu-teuer-21629207.html), the reasons for the discontinuation were lack of potential and the small population of Ulm (around 300 000 inhabitants).

Table 2.2: Summary Statistics 1 - Demographics and Car Sharing Vehicles

Panel A: Treated Cities (119 822 obs.)							
	Mean	Std. Dev.	Min	p25	p50	p75	Max
Population (in 1000s)	1 340	1 014	581	605	733	1 763	3 520
Percentage Male	48.8	0.005	47.6	48.4	48.8	49.0	49.8
Income (in 1000€)	27.6	3.6	19.2	26.3	28.5	30.0	32.3
Disposable Income (in 1000€)	20.6	1.9	16.3	19.1	21.0	22.1	23.7
Price per m^2 of Land (in €)	503	199	171	323	486	632	984
Number of Car Sharing Vehicles when Free-Floating Car Sharing is Offered							
car2go	527	325	250	250	350	700	1200
DriveNow	464	276	150	250	350	580	1200
Panel B: Control Cities (148 166 obs.)							
	Mean	Std. Dev.	Min	p25	p50	p75	Max
Population (in 1000s)	635	199	501	525	571	614	1 144
Percentage Male	48.6	0.004	47.8	48.4	48.5	48.9	49.7
Income (in 1000€)	21.3	2.63	16.6	19.2	21.5	23.4	27.1
Disposable Income (in 1000€)	20.4	1.87	16.0	18.7	20.4	21.8	23.4
Price per m^2 of Land (in €)	173	86	59	111	153	221	414.8

Summary statistics for control and treatment variables. The unit of observation is a year-month-city-model combination.

Table 2.3: Summary Statistics 2 - Monthly New Car Sales by Segment

Segment (Nr of Models)	% of Total Sales	Mean	Std. Dev.	Min	p25	p50	p75	Max
Panel A: Treated Cities (119 822 obs.)								
All Segments (187)	100	20.0	53.0	0	0	5	17	1 574
Small (45)	23.1	19.2	41.1	0	0	6	20	567
Compact (29)	25.0	32.7	93.1	0	1	6	27	1574
Mid-Sized (24)	27.0	41.9	76.6	0	1	10	49	691
SUV (49)	12.9	10.0	18.0	0	1	4	11	261
Luxury (14)	3.1	8.8	15.3	0	2	4	9	181
Van (26)	8.9	11.9	25.5	0	0	3	12	327
Monthly New Car Sales of Models Used by *car2go* and *DriveNow* (Before Treatment)								
smart fortwo (*car2go*)	2.2	82.2	65.7	6	26	55.5	138	288
MINI (*DriveNow*)	1.5	56.1	28.5	9	37	47	71	143
BMW 1 Series (*DriveNow*)	2.2	81.7	50.5	13	45	68	103	243
Panel B: Control Cities (148 166 obs.)								
Segment (Nr of Models)	% of Total Sales	Mean	Std. Dev.	Min	p25	p50	p75	Max
All Segments (187)	100	7.4	15.5	0	0	2	8	536
Small (45)	28.8	8.8	15.3	0	0	3	11	237
Compact (29)	26.6	12.6	26.3	0	0	4	13	536
Mid-Sized (24)	19.5	11.2	17.9	0	1	4	14	281
SUV (49)	14.4	4.0	7.1	0	0	2	5	152
Luxury (14)	2.1	2.1	2.9	0	0	1	3	33
Van (26)	8.6	4.6	8.0	0	0	2	6	130
Monthly New Car Sales of Models Used by *car2go* and *DriveNow* (Before Treatment)								
smart fortwo (*car2go*)	1.1	15.2	12.5	0	7	12	19	114
MINI (*DriveNow*)	1.6	21.5	12.0	0	12	19	28	88
BMW 1 Series (*DriveNow*)	2.3	31.9	20.5	2	17	26	42	161

Summary statistics for new car sales. The unit of observation is a year-month-city-model combination. The number of models refers to the average number of available models over time. The category SUV was only introduced in 2011. All models used by the car sharing companies are dropped from the segments.

Table 2.4: Aggregate Monthly New Car Sales by Segment per 10 000 Inhabitants

Panel A: Treated Cities							
Segment (Nr of Models)	Mean	Std. Dev.	Min	p25	p50	p75	Max
All Segments (187)	33.5	10.2	9.7	30.3	35.6	40.3	54.7
Small (45)	7.1	2.5	2.4	5.3	7.0	8.7	15.4
Compact (29)	7.8	3.2	2.3	5.6	7.5	9.5	18.8
Mid-Sized (24)	10.0	4.6	1.6	6.5	10.3	13.5	24.5
SUV (49)	4.4	1.9	0.8	3.0	4.3	5.8	10.2
Luxury (14)	1.3	1.4	0.2	0.5	0.8	1.2	6.9
Van (26)	2.7	1.2	0.5	1.7	2.7	3.5	7.2
Panel B: Control Cities							
Segment (Nr of Models)	Mean	Std. Dev.	Min	p25	p50	p75	Max
All Segments (187)	21.8	4.4	13.8	18.6	21.2	24.3	42.1
Small (45)	6.1	2.1	2.5	4.7	5.8	7.0	17.0
Compact (29)	5.8	1.6	2.9	4.8	5.5	6.5	14.5
Mid-Sized (24)	4.4	1.8	1.3	3.1	4.1	5.2	10.5
SUV (49)	3.1	1.2	0.8	2.0	3.0	4.0	8.2
Luxury (14)	0.5	0.2	0.0	0.3	0.4	0.5	1.5
Van (26)	1.9	0.7	0.6	1.4	1.9	2.3	5.2

New car sales aggregated over all models in a segment. The unit of observation is a year-month-city-segment combination. All models used by the car sharing companies are dropped.

Table 2.5: Substitution Effect Estimation Results for Different Measures of Treatment with Free-Floating Car Sharing

	(1) Sales	(2) Sales	(3) Sales	(4) Sales
Car Sharing Dummy	-1.096*			
	(0.573)			
Number of CS Cars $\times 10^{-3}$		-1.771***		
		(0.278)		
Number of CS Cars/			-0.316***	
CS Area in km			(0.068)	
Number of CS Cars/				-2524.6**
Population				(1077.7)
Effect of one CS Car on Annual Sales (in cars)	-3.0	-4.0	-4.7	-4.1
Population $\times 10^{-6}$	1.490	5.910	1.480	3.610
	(3.920)	(4.280)	(4.720)	(4.090)
Percentage Male	482.6***	225.1*	337.8***	464.1***
	(118.9)	(108.0)	(99.4)	(106.2)
Income $\times 10^{-4}$	1.939	-1.566	-0.349	2.952
	(4.029)	(3.096)	(2.853)	(3.774)
Disposable Income $\times 10^{-4}$	2.105	2.206	1.994	3.396
	(6.455)	(4.899)	(5.716)	(6.352)
Price per m^2	-2.083	-0.528	-1.081	-1.388
of Land $\times 10^{-3}$	(1.987)	(1.547)	(1.842)	(2.105)
Model FE	Yes	Yes	Yes	Yes
City FE	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes
N	267 988	267 988	267 988	267 988
R^2	0.468	0.468	0.468	0.468

Difference-in-Difference regression results. Standard errors in parentheses. Standard errors are clustered at the city level. All coefficients are weighted by population. The average number of available models in a month is 187. The average number of Free-Floating Car Sharing Vehicles in a treated city is 812. The effect of one Free-Floating Car Sharing vehicle on annual sales is calculated at these averages. * $p< 0.1$, ** $p<0.05$, *** $p<0.01$

Table 2.6: Substitution Effect Robustness Checks for Number of Free-Floating Car Sharing Vehicles as Treatment Intensity Measure.

Panel A

	(1) Sales	(2) Sales	(3) Sales
Number of CS Cars $\times 10^{-3}$	-1.130***	-1.431***	-1.156*
	(0.273)	(0.514)	(0.635)
Effect of one CS Car on Annual Sales (in cars)	-2.5	-3.1	-3.5
Controls	Yes	Yes	Yes
Model FE	Yes	Yes	Yes
City FE	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes
City-Specific Trend	Yes	No	Yes
Sample	Full	Only Treated Cities	Only Treated Cities
N	267 988	119 822	119 822
R^2	0.468	0.547	0.547

Panel B

	(1) lg(Sales)	(2) lg(Sales)	(3) lg(Sales)	(4) lg(Sales)
Number of CS Cars $\times 10^{-5}$	-3.77***	-9.54***	5.47***	-9.921***
	(-0.518)	(1.431)	(0.824)	(1.610)
Effect of one CS Car on Annual Sales (in cars)	-1.7	-4.4	-2.5	-4.5
Controls	Yes	Yes	Yes	Yes
Model FE	Yes	Yes	Yes	Yes
City FE	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes
City-Specific Trend	No	Yes	No	Yes
Sample	Full	Full	Only Treated Cities	Only Treated Cities
N	267 988	267 988	119 822	119 822
R^2	0.616	0.616	0.616	0.616

Difference-in-Difference regression results. Standard errors in parentheses. Standard errors are clustered at the city level. All coefficients are weighted by population. Control variables are: Population, percentage male, income, disposable income, price of land. Average number of available models per month: 187. Average number of Free-Floating Car Sharing Vehicles in treated city: 812. Average number of monthly sales per model: 20.2. Average sum of the sizes of the operating areas of the Free-Floating Car Sharing providers in treated cities: 145 km^2. Average population in treated cities: 1 346 387. The effect of one Free-Floating Car Sharing vehicle on annual sales is calculated at these averages. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 2.7: Substitution Effect Robustness Checks: Leads and Lags

	(1) Sales	(2) Sales	(3) Sales
Lead36 $\times 10^{-3}$		1.094	1.267
		(0.873)	(0.888)
Lead24 $\times 10^{-3}$		-0.814	-0.702
		(0.958)	(0.667)
Lead12 $\times 10^{-3}$		-0.787	-0.564
		(0.742)	(1.212)
Number of CS Cars $\times 10^{-3}$	-1.771***	-2.076***	-2.150***
	(0.278)	(0.423)	(0.661)
Lag12 $\times 10^{-3}$		1.274*	1.693***
		(0.596)	(0.218)
Lag24 $\times 10^{-3}$		0.019	0.568*
		(0.321)	(0.320)
Lag36 $\times 10^{-3}$		0.266	0.685
		(0.485)	(0.987)
Controls	Yes	Yes	Yes
Model FE	Yes	Yes	Yes
City FE	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes
City-Specific Trend	No	No	Yes
N	267 988	267 988	267 988
R^2	0.468	0.468	0.468

Difference-in-Difference regression results. Standard errors in parentheses. Standard errors are clustered at the city level. All coefficients are weighted by population. Control variables are: Population, percentage male, income, disposable income, price of land. Variables beginning with *Lead* (*Lag*) are constructed by leading (lagging) the treatment variable by the specified number of months. The treatment here is the number of available car sharing vehicles. * p< 0.1, ** p<0.05, *** p<0.01

Table 2.8: Substitution Effect Estimation Results by Car Segment

	(1) Sales Small	(2) Sales Compact	(3) Sales Mid-Sized	(4) Sales SUV	(5) Sales Luxury	(6) Sales Van
Number of CS Cars $\times 10^{-3}$	-2.991***	-3.234***	-6.092***	0.242	-0.271	-0.683
	(0.315)	(0.311)	(0.127)	(0.680)	(0.286)	(0.586)
Effect of one CS Car on Annual Sales (in cars)	-1.4	-1.0	-1.6	-	-	-
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Model FE	Yes	Yes	Yes	Yes	Yes	Yes
City FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Sample	Small Models	Compact Models	Mid-Sized Models	SUVs	Luxury Models	Vans
N	63 158	41 635	35 034	40 846	18 974	38 202
R^2	0.444	0.472	0.528	0.451	0.352	0.386

Difference-in-Difference regression results. Standard errors in parentheses. Standard errors are clustered at the city level. All coefficients are weighted by population. Control variables are: Population, percentage male, income, disposable income, price of land.The effect of one Free-Floating Car Sharing vehicle on annual sales in a segment is calculated for the average number of available models in that segment. * $p< 0.1$, ** $p<0.05$, *** $p<0.01$

Table 2.9: Substitution Effect Estimation Results with Heterogeneous Effects by Demographics

Panel A: Fraction Female Car Sharing Users				
	(1)	(2)	(3)	(4)
	Sales	lg(Sales)	Sales	lg(Sales)
Number of CS Cars $\times 10^{-3}$	-1.66**	-0.068***	-1.198*	-0.058**
	(0.761)	(0.021)	(0.679)	(0.021)
Fraction Female CS Users	-0.257	0.191	0.299	0.343**
	(4.242)	(0.149)	(5.152)	(0.124)
Number of CS Cars*			0.002	-0.485***
Fraction Female CS Users			(0.004)	(0.109)
Controls	Yes	Yes	Yes	Yes
Model FE	Yes	Yes	Yes	Yes
City FE	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes
N	267 988	267 988	267 988	267 988
R^2	0.468	0.616	0.468	0.616
Panel B: Average Age of Car Sharing Users				
	(1)	(2)	(3)	(4)
	Sales	lg(Sales)	Sales	lg(Sales)
Number of CS Cars $\times 10^{-3}$	-1.594**	-0.067***	0.812	-0.050*
	(0.597)	(0.013)	(0.459)	(0.024)
Mean Age CS Users $\times 10^{-3}$	5.775	1.551*	0.459	2.489**
	(2.658)	(0.868)	(3.159)	(0.838)
Number of CS Cars*			-2.410	-0.361***
Mean Age CS Users $\times 10^{-5}$			(2.631)	(0.092)
Controls	Yes	Yes	Yes	Yes
Model FE	Yes	Yes	Yes	Yes
City FE	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes
N	267 988	267 988	267 988	267 988
R^2	0.468	0.468	0.616	0.468

Difference-in-Difference regression results. Standard errors in parentheses. Standard errors are clustered at the city level. All coefficients are weighted by population. * $p< 0.1$, ** $p<0.05$, *** $p<0.01$

Table 2.10: Difference-in-Difference Model Advertising Effect Estimation Results

	(1)	(2)	(3)	(4)	(5)	(6)
	Sales	Sales	Market Share ×100	Market Share ×100	Market Share ×100	Market Share ×100
DriveNow	-7.026	-0.315	0.675***	-0.526	-1.950**	0.020
	(4.375)	(1.598)	(0.220)	(0.587)	(0.977)	(0.285)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
City FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Sample	*BMW 1 Series*	*MINI*	*BMW 1 Series*	*MINI*	*Mercedes A-Class*	*Audi A3*
N	1 426	1 426	1 426	1 426	1 426	1 426
R^2	0.773	0.882	0.626	0.574	0.700	0.516

Difference-in-Difference regression results. Standard errors in parentheses. Standard errors are clustered at the city level. All coefficients are weighted by population. Control variables are: Population, percentage male, income, disposable income, price of land.The treatment variable is a dummy variable equal to one when *DriveNow* is active in a city. The estimation samples consist of all observations of sales of a model (roughly 14 cities times 108 time periods). The outcome variable in columns (3)-(6) is the market share of the model within its segment (*Small* for the *MINI*, *Compact* for the three other models) in a city-month combination. * $p< 0.1$, ** $p<0.05$, *** $p<0.01$

Table 2.11: Triple Difference Model Advertising Effect Estimation Results

	(1)	(2)	(3)	(4)	(5)
	Sales	Market Share ×100	Sales	Market Share ×100	Market Share ×100
MINI	3.997	-0.262			
× *DriveNow*	(2.697)	(0.652)			
BMW 1 Series			-6.669	0.740***	0.665***
× *DriveNow*			(4.975)	(0.244)	(0.253)
Mercedes A-Class					-2.43**
× *DriveNow*					(0.774)
Audi A3					0.222
× *DriveNow*					(0.352)
Time-Model FE	Yes	Yes	Yes	Yes	Yes
Model-City FE	Yes	Yes	Yes	Yes	Yes
City-Time FE	Yes	Yes	Yes	Yes	Yes
Sample	All Small	All Small	All Compact	All Compact	All Compact
N	64 584	64 584	45 094	45 094	45 094
R^2	0.810	0.751	0.891	0.892	0.892

Triple Difference regression results. Standard errors in parentheses. Standard errors are clustered at the city level. All coefficients are weighted by population. The outcome variable in columns (1) and (3) is total sales, in all other columns it is the market share of a model within its segment. The treatment variables are dummies equal to one for observations of sales of the particular model and in cities and at times when *DriveNow* was active. * $p< 0.1$, ** $p<0.05$, *** $p<0.01$

Table 2.12: Difference-in-Difference Brand Advertising Effect Estimation Results

	(1)	(2)	(3)	(4)	(5)	(6)
	Sales	Market Share $\times 100$	Sales	Market Share $\times 100$	Sales	Market Share $\times 100$
DriveNow	-10.55**	0.117	-0.704*	-0.125	-0.860	1.29**
	(3.707)	(0.179)	(0.391)	(0.235)	(0.558)	(0.572)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
City FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Sample	BMW Mid-Sized	BMW Mid-Sized	BMW SUV	BMW SUV	BMW Luxury	BMW Luxury
N	4 890	4 890	5 829	5 829	2 852	2 852
R^2	0.299	0.117	0.355	0.346	0.426	0.295

Difference-in-Difference regression results. Standard errors in parentheses. Standard errors are clustered at the city level. All coefficients are weighted by population. Control variables are: Population, percentage male, income, disposable income, price of land. The treatment variable is a dummy variable equal to one when *DriveNow* is active in a city. The estimation samples consist of all observations of BMW sales of a segment. The outcome variable is either total sales or the market share of the models within the segment in a city-month combination. * $p< 0.1$, ** $p<0.05$, *** $p<0.01$

Table 2.13: Triple Difference Brand Advertising Effect Estimation Results

	(1)	(2)	(3)	(4)	(5)	(6)
	Sales	Market Share $\times 100$	Sales	Market Share $\times 100$	Sales	Market Share $\times 100$
BMW $\times$ *DriveNow*	-4.885	-0.081	-2.001**	-0.018	-0.199	1.57
	(3.504)	(0.268)	(0.900)	(0.283)	(0.451)	(0.906)
Time-Model FE	Yes	Yes	Yes	Yes	Yes	Yes
Model-City FE	Yes	Yes	Yes	Yes	Yes	Yes
City-Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Sample	All Mid-Sized	All Mid-Sized	All SUV	All SUV	All Luxury	All Luxury
N	37 073	37 073	71 935	71 935	17 542	17 542
R^2	0.885	0.895	0.790	0.783	0.769	0.653

Triple Difference regression results. Standard errors in parentheses. Standard errors are clustered at the city level. All coefficients are weighted by population. The treatment is a dummy equal to one for BMW models in cities and at times where and when *DriveNow* is active. The outcome variable is either total sales or the market share of the models within the segment in a city-month combination. * $p< 0.1$, ** $p<0.05$, *** $p<0.01$

Chapter 3

Environmental Consciousness in the Car Market

3.1 Introduction

The use of automobiles powered by fossil fuels accounts for a large amount of emitted greenhouse gases and is therefore an important contributor to global warming. Since fuel-efficiency, which is collinear with pollutant emissions, varies widely across car models and fuel types, consumer choices along these two dimensions are an important determinant of total emissions. Understanding how consumers value environmental friendliness of cars is therefore a prerequisite for informed policy decision making.

Identifying environmental consciousness in the car market is challenging, because of the one-to-one correlation of pollutant emissions and fuel efficiency (when both are measured for a given distance driven). When a consumer is observed to purchase a low-emission (which is the same as saying fuel-efficient) vehicle, it is not possible to say whether she chose this model because of her preference for a clean environment or her unwillingness to pay large amounts of money for gas.

In this paper I exploit an exogenous shifter of beliefs over environmental friendliness of a subset of car models, the Volkswagen (VW) emissions scandal (popularly known as "Dieselgate"). In September 2015 it became widely known that due to the installation of illegal "defeat devices", the emissions of some Volkswagen models with diesel engines were by orders of magnitude higher than the legally allowed limits and what they were thought to be. However, crucially, consumers' perceptions about the fuel efficiency of the affected car models did not change. This variation allows the disentanglement of preferences for fuel efficiency and environmental friendliness.

Policy Implications

For a policy maker interested in inducing consumers to switch to environmentally friendly car models, it is important to know how much the consumers inherently value this environmental friendliness. If they already have a high valuation of a clean environment, further subsidization of clean cars is likely going to be ineffective. However, if they are indifferent towards the environment, but somewhat price sensitive, subsidization of clean cars (or taxation of dirty ones) could be an effective policy instrument.

Expectations over Future Regulation and Resale Values

Choosing and buying a car is an important and infrequent decision taken by households. One of the most important determinants of this decision are considerations over the future resale value of the vehicle. The Dieselgate scandal and the massive regulatory uncertainty that came with it impacted consumers' expectations over future resale values. Any model of the car market has to take this into account by including the dynamic aspect of the car purchase decision.

3.2 The Data

Quantities

The empirical analysis in this paper uses data from Germany. The quantity data[1] described in this subsection is provided by the the German Department of Motor Vehicles ("Kraftfahrtbundesamt Flensburg"). As is common in the literature, first-time registrations of vehicles are used as proxies for new car sales[2] and changes of ownership for used car sales. Both are observed monthly from January 2010 until December 2017 and are split at the registration district[3] and at the model-fuel level. The observed characteristics of the car are fuel type (petrol, diesel or other), engine displacement, fuel consumption and CO_2 emissions. For used cars I also observe the year during which they were first registered. For private buyers I observe age and gender; for commercial buyers I see their occupation or industry.

I also observe the stock of owned cars at the city-model-type level, but only at annual frequency from 2010 until 2017. This will allow me to back out how many and which cars were scrapped and hence give a full account of car ownership in Germany during the sample period.

Prices

Monthly[4] new (list) and used car price data from 2010 until 2017 is provided by the market research company *Deutsche Automobil Treuhand*, who survey more than 70 000 German car dealerships. Crucially, used car prices are not only functions of observable characteristics, but average market transaction prices. They are available for vehicles built in the preceding 7 years[5]. All prices are observed at the model-engine level (resulting in a cross-section of

[1]Note that the data described here is available for purchase, but not available for the purposes of this draft (as of June 2018). The quantity data used in this draft is more coarse and described in section 3.3.2

[2]In the paper, sales and registrations are synonymous

[3]There are around 400 registration districts in Germany. While the boundaries of larger cities coincide with their registration districts, rural registration districts sometimes contain several smaller cities. Registration districts are referred to as cities in the paper

[4]The data is available for purchase on a monthly level. As of June 2018, I have 12 cross-sections from the dates indicated on the horizontal axis of Figure 3.3

[5]For an example of the structure of the data, please see Figure 3.1

between 2 000 and 6 000 prices). I collapse the data at the model-fuel level using the median price across all engine specifications, resulting in a cross-section of around 500 model-fuel combinations. Not all models are available at all times and some models are only offered with petrol, not with diesel engines.

Other Data

Demographics as well as election results at the city level are publicly available from the Federal Statistical Office. I also use a 2014 survey about attitudes towards the environment[6] in Germany.

3.3 Reduced Form Evidence

3.3.1 Prices

This section examines the effect of Dieselgate on new and used car prices. As a first step, the reactions of prices of cars with and without diesel engines to Dieselgate are compared in a difference in difference setting. I then introduce a third difference across brands in order to allow for differential effects on models produced by *Volkswagen* compared to models produced by other car manufacturers.

Difference-in-Difference Specification

$$price_{jft} = \alpha + \lambda_j + \lambda_t + \beta * \mathbb{1}\{Diesel\}_f + \delta * \mathbb{1}\{Diesel\}_f * \mathbb{1}\{\text{After 09/2015}\}_t + \epsilon_{jft} \quad (3.1)$$

The unit of observation is a model-fuel-year combination, where model is denoted by j, fuel type by f and time by t. α denotes the constant, λ_j and λ_t are model and time fixed effects, δ is the coefficient of interest and ϵ is the error term.

[6] Bundesministerium für Umwelt, Naturschutz, Bau und Reaktorsicherheit (BMUB), Berlin (2015): Umweltbewusstsein in Deutschland 2014

Triple-Difference Specification

$$price_{jft} = \alpha + \lambda_{jt} + \lambda_{tf} + \lambda_{jf} + \delta * \mathbb{1}\{VW\}_j * \mathbb{1}\{Diesel\}_f * \mathbb{1}\{\text{After 09/2015}\}_t + \epsilon_{jft} \quad (3.2)$$

The λ terms now denote interaction fixed effects. Including model-time fixed effects λ_{jt}, for example, implies that for each model, a separate time dummy variable is included for each time period.

Results and Comments

The difference in difference regression results are in Table 3.1 and the triple-difference results are in Table 3.2. The first row of table 3.1 shows that list and used prices are higher for diesel cars (due to higher taxes and the fact that diesel engines are more expensive to produce). Examining jointly the second row and the information about average new and used car prices, it is evident that the price reaction to Dieselgate in percentage terms is far larger for used cars compared to new cars (around 3% for new cars and 10% for used cars.

This effect can also be shown graphically; Figures 3.3 and 3.4 show price increases of non-diesel new and used cars, which are likely due to increased demand from consumers substituting away from diesel cars, because of concerns about future diesel regulation or about a drop in resale values of diesel cars.

The *Volkswagen* brand was by far the most affected brand during the *Dieselgate* scandal. It is therefore conceivable, that consumers not only substituted away from diesel cars in general, as the difference-in-difference results suggest, but also that *Volkswagen* models were even more affected. In order to allow for a differential effect, I introduce a third difference across car brands.

However, figures 3.5 and 3.6, as well as the triple difference results in Table 3.2 show that prices of VW models were not affected differently by the Dieselgate scandal with respect to prices of cars of other brands.

In terms of timing of the effect of the scandal on car prices, figures 3.7 and 3.8 show that there was no statistically significant price effect until late 2017. German consumers did not react to the revelations in the U.S. immediately, but it took some time for the implications of Dieselgate for the German car market to emerge more clearly.

3.3.2 Quantities

The quantity data described in Section 1.2 is available for purchase. In this draft, only a more coarse version of the data is available. It consists of monthly (from 2010 to 2016) new car registrations (no data on used cars is available), split by brand, model and city (but not by fuel type). Therefore, the analysis here will focus on how the Dieselgate scandal impacted VW new car sales as opposed to sales of other brands. In other words, the analysis in this subsection checks whether VW sales suffered disproportionately in the aftermath of the Dieselgate scandal. This is done in a difference-in-difference setting, the differences being over time and across brands and the outcome variable being new car sales.

Different parts of the population may have reacted to the Dieselgate scandal to different extents. More specifically, it is conceivable that more environmentally conscious consumers, or consumers to which the environmental footprint of cars is more salient, had a different reaction to Dieselgate compared to consumers indifferent towards the environment. To test whether this is the case, I introduce a third difference across how ”environmentally conscious” consumers in different cities are. The true value is unobserved, so I use different proxies (described below) for it. To keep the number of fixed effects tractable, and since I cannot differentiate between fuel types and am interested in effects on the VW brand, I collapse the data at the brand level.

Proxies for Environmental Consciousness of Consumers in Different Cities

The first proxy is the voting share of the Green Party from the federal elections in 2013 and 2017 (available at the city level). The Green Party has traditionally been associated with environmental protection, so that its voting share can be used as a proxy for how important this topic is to voters in a constituency.

A second proxy is given by responses to a 2014 survey[7] about protection of the environment. Survey responses are split at the state level (Germany has 16 states), so I aggregate the data at that level. One proxy from the survey is the proportion of respondents who listed the environment as one of the top two problems Germany is facing. The other proxy is the average rating respondents assigned to the quality of the environment in their hometown, ranging from 1 (very good) to 4 (very bad).

Finally, instead of "taste" for environmental protection, it is also possible that citizens of some cities reacted differently to the Dieselgate scandal because the issue of environmental harm caused by privately owned (diesel) vehicles is more salient to them. Therefore, a third group of proxies is given by the intensity of reporting on the Dieselgate scandal in local newspapers (work in progress, not included in this draft). As opposed to many other countries, where a small number of national newspapers dominate the market, local newspapers are widely read in Germany. A measure for how salient the Dieselgate scandal is in different localities could be given by the share of scandal-related articles in local newspapers. More sophisticated machine-learning techniques would allow the construction of a measure of not only salience, but also sentiment towards the Dieselgate scandal in the different cities. The strategy of using newspaper articles as proxies for salience of current events has been used recently by Fouka & Voth (2018).

Difference-in-Difference Specification

$$sales_{bct} = \alpha + \lambda_c + \lambda_t + \beta * \mathbb{1}\{VW\}_b + \delta * \mathbb{1}\{VW\}_b * \mathbb{1}\{\text{After 09/2015}\}_t + \epsilon_{bct} \tag{3.3}$$

b stands for brand and c for city/state. The rest of the notation is as above.

Triple-Difference Specification

$$sales_{bct} = \alpha + \lambda_{bt} + \lambda_{tc} + \lambda_{bc} + \delta * \text{Green}_c * \mathbb{1}\{VW\}_b * \mathbb{1}\{\text{After 09/2015}\}_t + \epsilon_{bct} \tag{3.4}$$

The λ terms now denote interaction fixed effects. Including model-time fixed effects λ_{jt},

[7] Bundesministerium für Umwelt, Naturschutz, Bau und Reaktorsicherheit (BMUB), Berlin (2015): Umweltbewusstsein in Deutschland 2014.

for example, implies that for each model, a separate time dummy variable is included for each time period.

Results and Comments

Figure 3.9 shows that the negative quantity reaction took place towards the end of the observed data, in late 2016. Similar as with prices, it seems that the revelations from the U.S. did not have an immediate impact on sales. moving to the difference-in-difference results, the first row of Table 3.3 shows that VW is a very popular brand in Germany. Introducing the interaction variable in columns (3) and (4) of Table 3.3 shows that there was at most a very small negative differential reaction of VW sales compared to the sales of other brands. Only when the log of sales is used as an outcome variable is there a statistically significant effect at the 10% level. This is consistent with anecdotal evidence that VW has not suffered big declines in sales after Dieselgate. In fact, 2017 was a record year for VW in terms of sales and revenue[8].
Now a third difference across environmental consciousness is introduced, in order to allow for differential effects for more "green" parts of the population. The first two columns of Table 3.4 show that VW sales after Dieselgate were *higher* in cities, where the Green Party was more successful in the two most recent federal elections. The coefficients are highly statistically significant. Confounding demographic factors correlated with preferences for both the Green Party and VW vehicles are picked up by brand-city fixed effects. Demographic factors changing over time are picked up by city-time fixed effects. I take this result as evidence that the VW brand did not suffer disproportionately in "green" cities, but actually did better compared to other cities. However, it is not possible to conclude that there is no environmental consciousness among consumers, because they may have switched from diesel to petrol cars, which is not observable with the data available in this draft.
Columns (3) and (4) use survey responses to questions asking about the importance of environmental protection and the state of the environment as proxies for environmental consciousness in the population. These are only available by state, so that the data has to be collapsed at the

[8] https://global.handelsblatt.com/companies/despite-dieselgate-another-record-year-vw-891450

state level. The use of neither proxy results in statistically significant results, which could be due to the fact that not much identifying variation is left when doing the analysis at the state level. In conclusion, there is no evidence that German car buyers "punished" VW after Dieselgate by substituting towards other car makers.

3.4 A Structural Model of the Car Market

3.4.1 Demand for Cars

Let i stand for an individual/household, j for a car model, f for the fuel type, J for the set of available car models, t for time, x for car characteristics, p for price, a for age of car, α for the price coefficient, γ for the discount factor, C for operating costs and t for a market, defined as a city-year-month combination.

Individuals derive utility from car ownership from car characteristics interacted with their demographics and the rental price of the vehicle in time t, which captures depreciation:

$$u_{ijft} = x_{jft}\beta_i^x - \alpha_i(p_{jft} - \gamma\mathbb{E}[p_{jft+1}] + C_{jft}) + \xi_{jft} + \epsilon_{ijft} \tag{3.5}$$

ϵ denotes an i.i.d. error term distributed according to a Type 1 Extreme Value distribution. Modeling vehicles as being rented each period is equivalent to assuming a perfect used-car market without transaction costs. This assumption has been used in the literature, for example by Bento et al. (2009).

The formulation above is close to a standard random coefficients model. A key assumption is that each consumer chooses in each period from all available alternatives (including the outside option) the model j which maximizes her utility. Let $\beta_i = (\beta_i^x, \alpha_i)$ be the random utility coefficients, which come from a distribution $F_\beta(\beta;\theta)$, where θ is a vector of distributional parameters, i.e. means and a variance-covariance matrix. This distribution is independent from the idiosyncratic logit error ϵ_{ijt}. Using the aggregation property of the logit error term, the predicted market share for the model-fuel combination j is the probability that j maximizes the consumers utility and can be expressed as:

$$s_{jt}(\xi;\theta) = \int_{\beta} \frac{\exp(x_{jft}\beta_i^x - \alpha_i(p_{jft} - \gamma\mathbb{E}[p_{jft+1}] + C_{jt}) + \xi_{jt})}{1 + \sum_{k\in J}\exp(x_{kt}\beta_i^x - \alpha_i(p_{kt} - \gamma\mathbb{E}[p_{kt+1}] + C_{kt}) + \xi_{kt})} dF(\theta) \tag{3.6}$$

3.4.2 Estimation

Let observed sales be denoted by q_{jt}. Estimation proceeds by setting observed sales equal to the market shares predicted by the model, multiplied by the number of potential customers in market t (denoted by I_t):

$$q_{jft} = s_{jft}(\xi;\theta) * I_t \tag{3.7}$$

ξ_{jft} is a product characteristic unobserved by the econometrician, but known to the consumer when she decides which car to buy. Similarly as in Grigolon et. al (2017), it is modelled as:

$$\xi_{jft} = \xi_j + \xi_t + \widetilde{\xi}_{jft} \tag{3.8}$$

ξ_j are model fixed effects and ξ_t are city fixed effects interacted with a time trend. Following the contraction mapping procedure by Berry (1994) and Berry et al. (1995), the market share system (3.7) is inverted to obtain an expression for $\widetilde{\xi}_{jft}$ as a function of observed market shares and parameters. $\widetilde{\xi}_{jft}$ is a proxy for an unobserved quality component of a car-fuel combination, which is likely correlated with its price. The identifying assumption is that $\widetilde{\xi}_{jft}$ is mean independent of the observed car characteristics x_t, so that:

$$\mathbb{E}\Big(h_{jft}(x_{jft})\widetilde{\xi}_{jft}\Big) = 0 \tag{3.9}$$

where $(h_{jft}(x_t)$ is an instrument function formed from variables observed at time t. Estimation then proceeds by finding the parameters that solve:

$$\min_{\theta} \widetilde{\xi}(\theta)' h(x)' \Omega h(x) \widetilde{\xi}(\theta) \tag{3.10}$$

where the variables are now vectors stacking all markets and car models and Ω is a weighting matrix.

3.5 Conclusion

This paper uses the exogenous changes in consumers perceptions of the environmental friendliness of some vehicles brought about by the Dieselgate scandal to identify environmental consciousness in the car market. Reduced form evidence shows that after it became known that VW diesel cars emit more pollutants than officially stated, prices of all non-diesel cars increased, while there was no large reaction in prices of diesel cars. Prices of VW models were not differentially affected compared to other car brands. Car buyers did not substitute away from VW in the aftermath of the scandal and there was no differential reaction across buyers with different valuations for the environment (as measured by different proxies).

Finer data, in particular registration data which is split geographically and at the model-fuel level is necessary to estimate a structural model of the car market. This in turn allows to answer the research question of the paper (how much consumers value the environmental friendliness of their cars) and to compute counterfactuals in order to evaluate many of the policies which are currently being discussed.

3.6 Tables

	(1) List Price	(2) List Price	(3) Used Price	(4) Used Price
$\mathbb{1}\{\text{Diesel}\}$	785.5***	1046.1***	458.2***	828.3***
	(107.7)	(124.4)	(80.6)	(93.4)
$\mathbb{1}\{\text{Diesel}\}\times$ $\mathbb{1}\{\text{After 09/2015}\}$		-979.5***		-1335.4***
		(234.7)		(172.4)
Model FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
Mean(Y)	29 560	29 560	14 048	14 048
SD(Y)	17 117	17 117	9 925	9 925
N	6 428	6 428	6 542	6 542
R^2	94.8%	94.9%	91.1%	91.2%

Table 3.1: Difference-in-Difference regression results. Standard errors in parentheses. Used Price refers to the median price across used cars of age 1 to 7.

	(1) List Price	(2) Used Price
$\mathbb{1}\{\text{Diesel}\} \times \mathbb{1}\{\text{VW}\}\times$	-1052.5	344.2
$\times\mathbb{1}\{\text{After 09/2015}\}$	(896.3)	(629.2)
Model-Diesel FE	Yes	Yes
Model-Time FE	Yes	Yes
Diesel-Time FE	Yes	Yes
Mean(Y)	29 560	14 048
SD(Y)	17 117	9 925
N	5 471	5 565
R^2	97.8%	96.1%

Table 3.2: Triple Difference regression results. Standard errors in parentheses. Used Price refers to the median price across used cars of age 1 to 7.

	(1)	(2)	(3)	(4)
	Sales	lg(Sales)	Sales	lg(Sales)
$\mathbb{1}\{\text{VW}\}$	108.9***	2.665***	109.2***	2.671***
	(0.648)	(0.007)	(0.714)	(0.008)
$\mathbb{1}\{\text{VW}\}\times$			-1.727	-0.034*
$\mathbb{1}\{\text{After 09/2015}\}$			(1.696)	(0.019)
Controls	Yes	Yes	Yes	Yes
City FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
Mean (Y)	18.8	1.59	18.8	1.59
SD (Y)	120.9	1.49	120.9	1.49
N	1 007 004	1 007 004	1 007 004	1 007 004
R^2	10.6%	25.9%	10.6%	25.9%

Table 3.3: Difference-in-Difference regression results. Standard errors in parentheses. The controls are population, percentage male, average age, income, disposable income, price of land in m^2

	(1) Sales	(2) Sales	(3) Sales	(4) Sales
Green × $\mathbb{1}\{\text{VW}\}\times$ $\mathbb{1}\{\text{After 09/2015}\}$	1.867*** (0.200)	1.511*** (0.165)	530.1 (407.0)	-21.8 (143.7)
Brand-City/State FE	Yes	Yes	Yes	Yes
City/State-Time FE	Yes	Yes	Yes	Yes
Time-Brand FE	Yes	Yes	Yes	Yes
Proxy for Green	Green Party 2013	Green Party 2017	top2	loc_qual
Mean(Y)	18.8	18.8	457.8	457.8
SD(Y)	120.9	120.9	1 157.5	1 157.5
N	980 220	980 220	41 664	41 664
R^2	90.5%	90.5%	96.9%	96.9%

Table 3.4: Triple Difference regression results. Standard errors in parentheses. The proxy for how green the population of a city/state is is given by the voting share of the Green Party in the federal elections 2013 and 2017 in the first two columns. This data is available at the city level. In columns (3) and (4), survey responses, which are available only at the state level, are used as the proxy. In column (3) the proxy is equal to the share of survey respondents who listed protecting the environment as one of their top two problems Germany is facing. The proxy in column (4) is the average rating given by respondents to the quality of the environment in their city. The scale ranges from 1 (very good) to 4 (very bad).

3.7 Figures

Gebrauchtwagenpreise 2016

Fahrzeug/Modell	kW (PS)	km-Kl.	Neu-preis	Durchschnittspreise in Euro für Baujahr 2014	2013	2012	2011	2010	2009	2008
BMW										
123 d DPF 5-türig Euro 5 (E81/87)	150 (203)	IV	**32650**				14250	13200	11800	
114 i 5-türig (F20/21)	75 (101)	III	**22650**		12700	11650				
116 i 5-türig (F20/21)	100 (135)	III	**23850**		13700	12500	11450			
116 i 5-türig (F20/21) (Euro 6)	100 (135)	III	**24000**	15150	13850					
118 i 5-türig (F20/21)	125 (169)	III	**26750**		15300	14100	12900			
118 i 5-türig (F20/21) (Euro 6)	125 (169)	III	**27000**	17050	15450					
125 i 5-türig (F20/21)	160 (217)	III	**31000**		18250	16850				
125 i 5-türig (F20/21) (Euro 6)	160 (217)	III	**31000**	20150	18250					
114d 5-türig (F20/21)	70 (95)	IV	**24900**	15400	13850	12750				
116d 5-türig (F20/21)	85 (115)	IV	**26100**	15750	14400	13200	12050			
118d 5-türig (F20/21)	105 (142)	IV	**27400**	17250	15700	14450	13200			
120d 5-türig (F20/21)	135 (183)	IV	**29400**	18600	17050	15750	14500			
125d 5-türig (F20/21)	160 (217)	IV	**33500**	22750	20950	19400				
125 i Coupe (E82)	160 (218)	III	**32550**				16500	14900	13450	11900
135 i Coupe (E82)	225 (306)	III	**40100**					17150	15400	13950

Figure 3.1: Used Car Prices as of October 2016 at the model-engine level. They are available for vehicles of age 1 to 7 years.

ADAC Autokosten 2016 - alle Hersteller			Kosten pro Monat in Euro und Cent/km					
Marke / Modell:	Leistung kW	Listen-preis	Fix-kosten	Werkst.-kosten	Betriebs-kosten	Wert-verlust	Gesamtkosten Monat	Cent
BMW 115 EUR/h								
225d Cabrio Steptronic	165	42950	141	59	86	499	**785**	**62,8**
318i Advantage	100	31450	112	53	109	465	**739**	**59,1**
320i Advantage	135	35800	114	67	116	506	**803**	**64,2**
330i Advantage	185	40050	120	85	126	571	**902**	**72,2**
330e Advantage Steptronic	185	43500	109	92	91	596	**888**	**71,0**
340i Advantage	240	46250	145	97	148	642	**1032**	**82,6**
316d Advantage	85	33850	125	53	77	480	**735**	**58,8**
318d Advantage	110	34750	125	60	78	482	**745**	**59,6**
320d EfficientDynamics Edition Advantage	120	38300	137	61	77	511	**786**	**62,9**
320d Advantage	140	38300	138	62	78	511	**789**	**63,1**
325d Advantage	165	41400	163	81	86	553	**883**	**70,6**
330d Advantage Steptronic	190	47350	172	86	90	631	**979**	**78,3**

Figure 3.2: Costs of owning a new vehicle as of October 2016, observed at the model-engine level. Costs are broken down into fixed cost, maintenance, operating cost and depreciation. The last column shows the total cost in Euro per month as well as Cent per km (assuming an annual mileage of 15 000km).

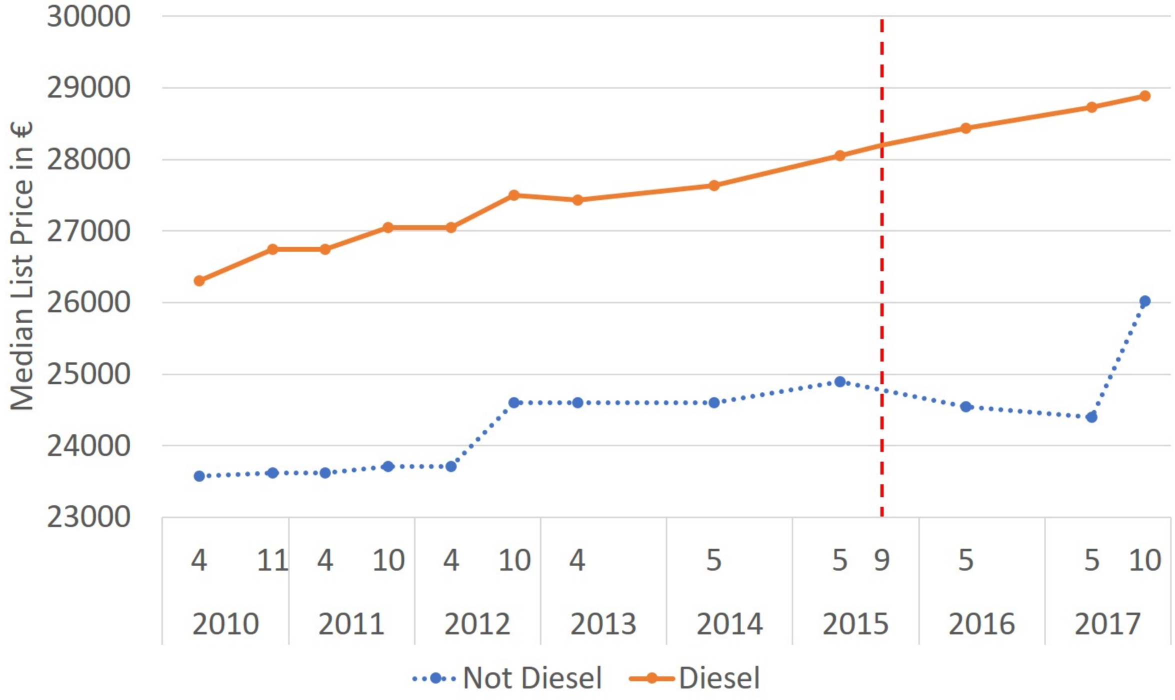

Figure 3.3: Median (across car models) list prices over time. The dashed red line indicates September 2015, when the Dieselgate Scandal began.

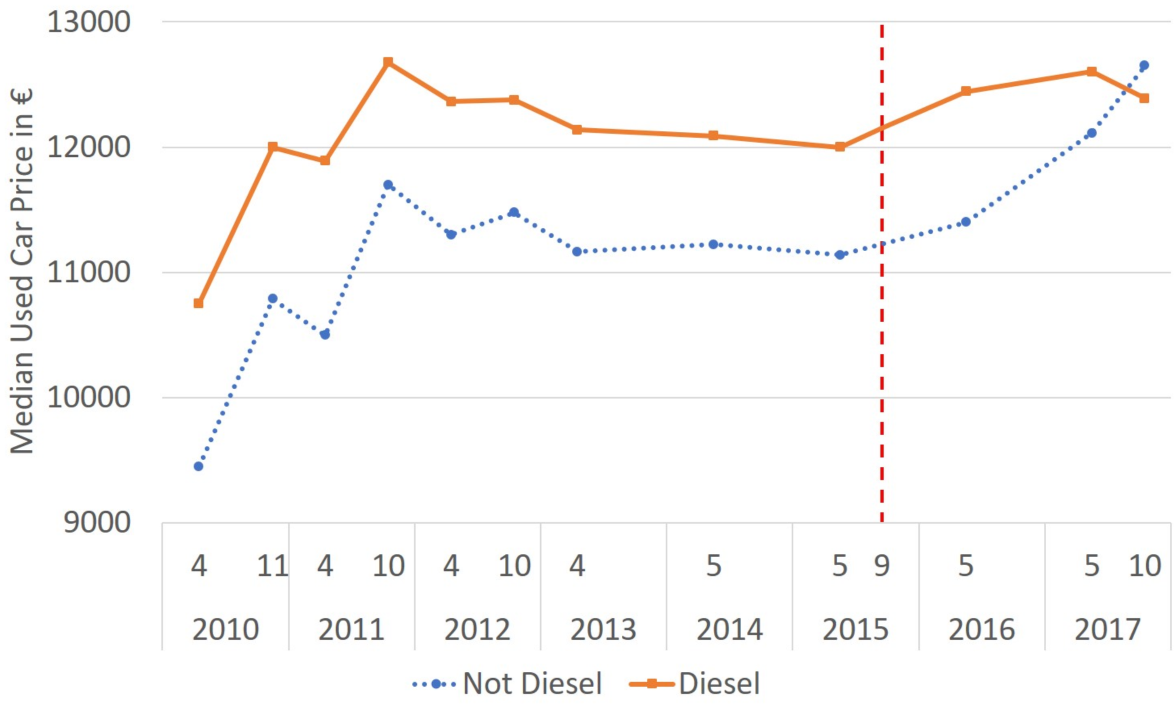

Figure 3.4: Median (across car models and model years) used car prices over time. The dashed red line indicates September 2015, when the Dieselgate Scandal began.

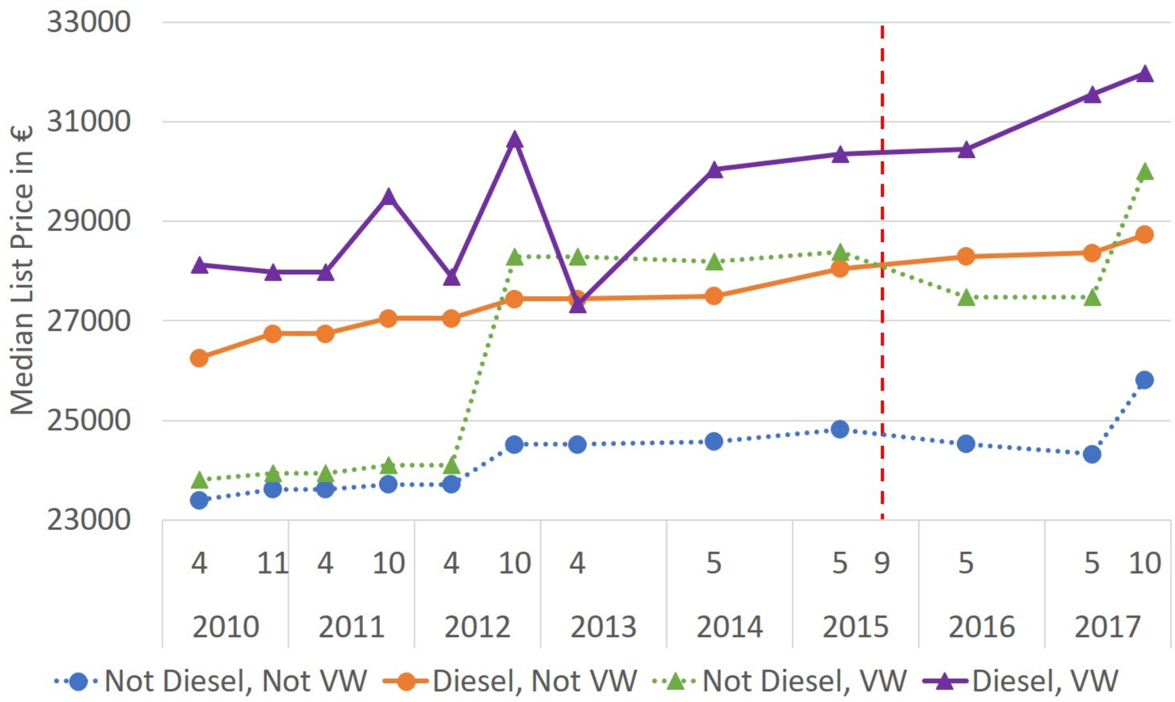

Figure 3.5: Median (across car models) list prices over time. The dashed red line indicates September 2015, when the Dieselgate Scandal began.

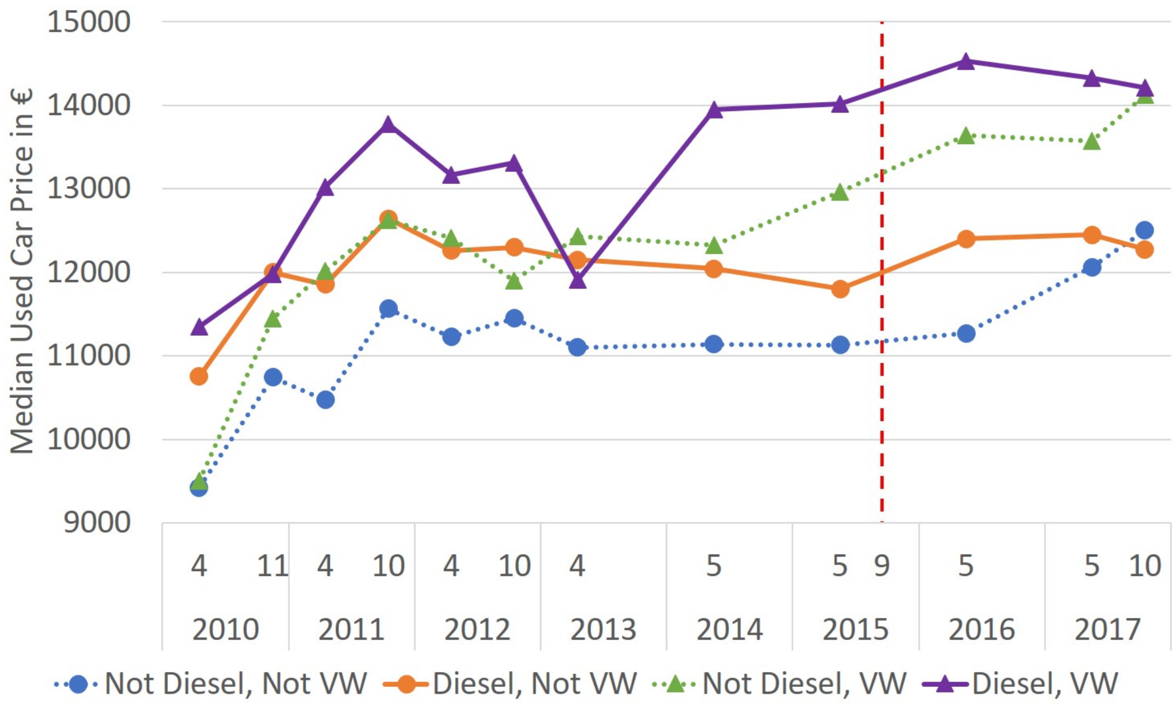

Figure 3.6: Median (across car models and model years) used car prices over time. The dashed red line indicates September 2015, when the Dieselgate Scandal began.

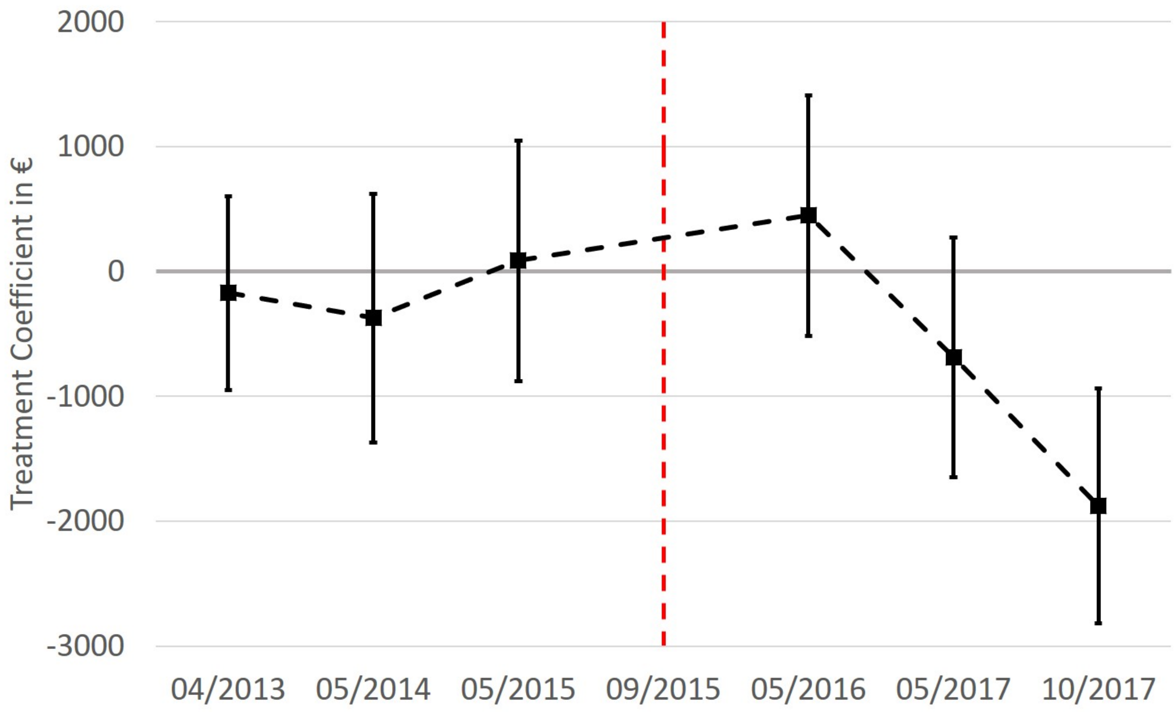

Figure 3.7: Treatment coefficients for the difference-in-difference regression for new car list prices including leads and lags of the treatment.

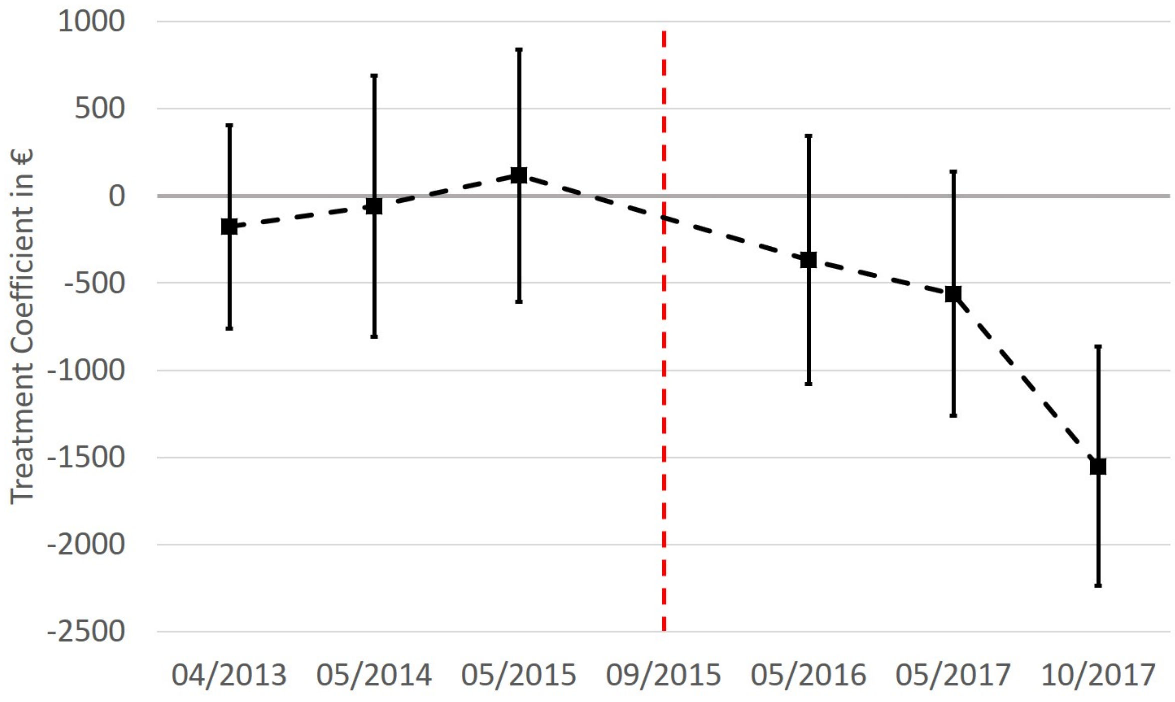

Figure 3.8: Treatment coefficients for the difference-in-difference regression for used car prices including leads and lags of the treatment.

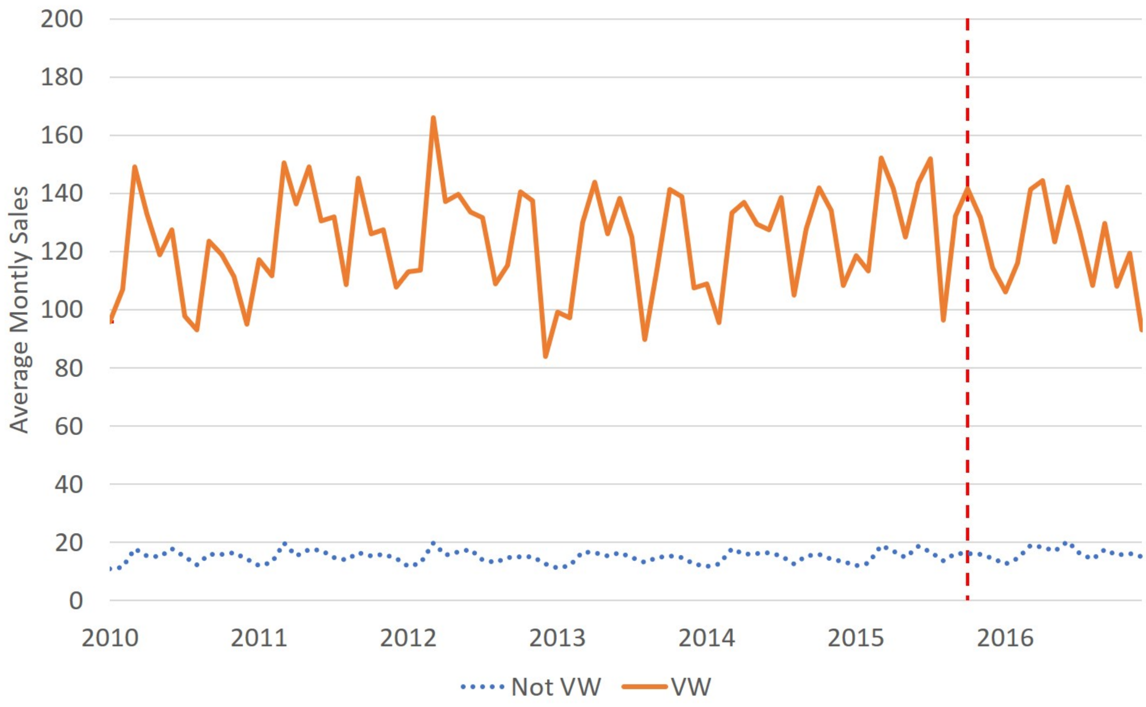

Figure 3.9: Average monthly car sales by brand in German cities.

Bibliography

Abhishek, V., Guajardo, J. A. & Zhang, Z. (2016), 'Business models in the sharing economy: Manufacturing durable goods in the presence of peer-to-peer rental markets'.

Adda, J. & Cooper, R. (2000), 'Balladurette and juppette: A discrete analysis of scrapping subsidies', *Journal of political Economy* **108**(4), 778–806.

Angrist, J. & Pischke, J.-S. (2009), 'Mostly harmless econometrics: an empiricists guide'.

Bellos, I., Ferguson, M. & Toktay, L. B. (2017), 'The car sharing economy: Interaction of business model choice and product line design', *Manufacturing & Service Operations Management* **19**(2), 185–201.

Benjaafar, S., Kong, G., Li, X. & Courcoubetis, C. (2018), 'Peer-to-peer product sharing: Implications for ownership, usage, and social welfare in the sharing economy', *Management Science* .

Bento, A. M., Goulder, L. H., Jacobsen, M. R. & Von Haefen, R. H. (2009), 'Distributional and efficiency impacts of increased us gasoline taxes', *American Economic Review* **99**(3), 667–99.

Berry, S., Levinsohn, J. & Pakes, A. (1995), 'Automobile prices in market equilibrium', *Econometrica: Journal of the Econometric Society* pp. 841–890.

Berry, S. T. (1994), 'Estimating discrete-choice models of product differentiation', *The RAND Journal of Economics* pp. 242–262.

Bert, J., Collie, B., Gerrits, M. & Xu, G. (2016), 'What's ahead for car sharing?', *Boston Consulting Group* .

Besley, T. & Burgess, R. (2004), 'Can labor regulation hinder economic performance? evidence from india', *The Quarterly journal of economics* **119**(1), 91–134.

BMW Group (2015), 'Annual report 2015', pp. 126–128.

BMW Group (2016), 'Annual report 2016', pp. 125–127.

Bohren, J. A., Imas, A. & Rosenberg, M. (2018), 'The dynamics of discrimination: Theory and evidence', *Working Paper* .

Brownstone, D. & Fang, H. (2014), 'A vehicle ownership and utilization choice model with endogenous residential density', *Journal of Transport and Land Use* **7**(2), 135–151.

Cohen, M. & Sundararajan, A. (2015), 'Self-regulation and innovation in the peer-to-peer sharing economy', *U. Chi. L. Rev. Dialogue* **82**, 116.

Cook, C., Diamond, R., Hall, J., List, J. A., Oyer, P. et al. (2018), The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers, Technical report.

Correia, S. (2016), Linear models with high-dimensional fixed effects: An efficient and feasible estimator, Technical report.

Coşar, A. K., Grieco, P. L., Li, S. & Tintelnot, F. (2018), 'What drives home market advantage?', *Journal of International Economics* **110**, 135–150.

Cramer, J. & Krueger, A. B. (2016), 'Disruptive change in the taxi business: The case of uber', *American Economic Review* **106**(5), 177–82.

Edelman, B. G. & Luca, M. (2014), 'Digital discrimination: The case of airbnb. com', *Working Paper* .

Edelman, B., Luca, M. & Svirsky, D. (2017), 'Racial discrimination in the sharing economy: Evidence from a field experiment', *American Economic Journal: Applied Economics* **9**(2), 1–22.

Einav, L., Farronato, C. & Levin, J. (2016), 'Peer-to-peer markets', *Annual Review of Economics* **8**, 615–635.

Farajallah, M., Hammond, R. G. & Pénard, T. (2016), 'What drives pricing behavior in peer-to-peer markets? evidence from the carsharing platform blablacar', *Working Paper* .

Farronato, C. & Fradkin, A. (2017), The welfare effects of peer entry in the accommodations market: The case of airbnb, Technical report, Working paper.

Firnkorn, J. & Müller, M. (2011), 'What will be the environmental effects of new free-floating car-sharing systems? the case of car2go in ulm', *Ecological Economics* **70**(8), 1519–1528.

Fouka, V. & Voth, H.-J. (2018), 'Reprisals remembered: German-greek conflict and car sales during the euro crisis', *Working Paper* .

Fraiberger, S. P. & Sundararajan, A. (2017), 'Peer-to-peer rental markets in the sharing economy', *Working Paper* .

Gillingham, K., Iskhakov, F., Munk-Nielsen, A., Rust, J. & Schjerning, B. (2015), 'A dynamic model of vehicle ownership, type choice, and usage', *Unpublished Working Paper.¡ http://bschjerning. com/papers/iruc. pdf* .

Gong, J., Greenwood, B. N. & Song, Y. (2017), 'Uber might buy me a mercedes benz: An empirical investigation of the sharing economy and durable goods purchase'.

Grigolon, L., Leheyda, N. & Verboven, F. (2016), 'Scrapping subsidies during the financial crisis—evidence from europe', *International Journal of Industrial Organization* **44**, 41–59.

Grigolon, L., Reynaert, M. & Verboven, F. (2014), 'Consumer valuation of fuel costs and the effectiveness of tax policy: Evidence from the european car market'.

Gutt, D. & Herrmann, P. (2015), Sharing means caring? hosts' price reaction to rating visibility., *in* 'ECIS'.

Horton, J. J. & Zeckhauser, R. J. (2016), Owning, using and renting: Some simple economics of the" sharing economy", Technical report, National Bureau of Economic Research.

Jong, G. D., Fox, J., Daly, A., Pieters, M. & Smit, R. (2004), 'Comparison of car ownership models', *Transport Reviews* **24**(4), 379–408.

Lambin, X. & Palikot, E. (2017), 'Does reputation hinder entry? study of statistical discrimination on a platform', *Working Paper* .

Laouénan, M. & Rathelot, R. (2017), 'Ethnic discrimination on an online marketplace of vacation rental', *Working Paper* .

Manco, F. & Percoco, M. (2018), 'Spaces of competition in the sharing economy: taxi markets and uber entry in nyc'.

Martin, E. & Shaheen, S. (2016), 'Impacts of car2go on vehicle ownership, modal shift, vehicle miles traveled, and greenhouse gas emissions: an analysis of five north american cities', *Transportation Sustainability Research Center, UC Berkeley* .

Münzel, K., Boon, W., Frenken, K. & Vaskelainen, T. (2017), 'Carsharing business models in germany: characteristics, success and future prospects', *Information Systems and e-Business Management* pp. 1–21.

Nurski, L. & Verboven, F. (2016), 'Exclusive dealing as a barrier to entry? evidence from automobiles', *The Review of Economic Studies* **83**(3), 1156–1188.

NWLC (2018), 'The wage gap by state for women overall 2016', *National Women's Law Center* .

Renan Barzilay, A. & Ben-David, A. (2017), 'Platform inequality: Gender in the gig-economy', *Working Paper* .

Schiraldi, P. (2011), 'Automobile replacement: a dynamic structural approach', *The RAND journal of economics* **42**(2), 266–291.

Seamans, R. & Zhu, F. (2013), 'Responses to entry in multi-sided markets: The impact of craigslist on local newspapers', *Management Science* **60**(2), 476–493.

Shoup, D. (2017), *The High Cost of Free Parking: Updated Edition*, Routledge.

Smith, T. W., Davern, M., Freese, J. & Hout, M. (2016), 'General social surveys, 1972-2016', *NORC at the University of Chicago* .

Sundararajan, A. (2016), *The sharing economy: The end of employment and the rise of crowd-based capitalism*, Mit Press.

Teubner, T., Hawlitschek, F. & Dann, D. (2017), 'Price determinants on airbnb: How reputation pays off in the sharing economy.', *Journal of Self-Governance & Management Economics* **5**(4).

U.S. Department of Justice, F. B. I. (2014), 'Uniform crime reporting statistics'.

Weber, T. A. (2016), 'Product pricing in a peer-to-peer economy', *Journal of Management Information Systems* **33**(2), 573–596.

Zeltzer, D. (2017), 'Gender homophily in referral networks: Consequences for the medicare physician earnings gap'.