

Gender price gaps and competition: Evidence from a correspondence study

Margarita Machelett*

Bank of Spain

November, 2019

I implement a large-scale field experiment in the US auto repair industry to study the existence and structure of gender-based price discrimination in service markets. Women receive price quotes that are 2 percent (over 9 dollars) higher than men. These differences disappear once women signal low search costs, suggesting statistical rather than taste-based discrimination. Price requests that appear to come from high-income households raise quotes for men but not women, also eliminating the gender gap. The price gap also falls with the number of nearby repair shops, suggesting that market competition alleviates gender-based price discrimination.

JEL codes: C93, D4, J16, J18, J71

Keywords: competition, discrimination, field experiment, gender

*I am very grateful for helpful comments and feedback from Emily Oster, John Friedman, Brian Knight, Anna Aizer, Vivek Bhattacharya, Pedro Dal Bó, Sebastián Fanelli, Bobby Pakzad-Hurson, Dan-Olof Rooth, Jesse Shapiro, Neil Thakral, Matt Turner, Lorenzo Adelco, Samuele Giambra, Florian Gunsilius, Ryan Kessler, Samsun Knight, Ted Lawless, Stefano Pollini, Catalina Ravizza, and seminar participants at the European Winter Meeting of the Econometric Society 2019, 2019 Advances with Field Experiments Conference, University of Milan, SOFI, University of Alicante, Universidad de las Islas Baleares, IE University, Universidad Autónoma de Madrid, Charles River and Associates, and the Bank of Spain. I also thank Luciana Casiutti, Josefina Higgs and Santiago Olivieri for help in processing the data. All errors are my own. *Email:* margarita_machelett@alumni.brown.edu

1 Introduction

In recent years, the fight for women’s rights in the US has expanded from a battle for access to political and economic spheres into a broader-based struggle for gender equity and equality. In the workplace, this recent change is most clearly manifest in the form of the #MeToo and Time’s Up movements, where advocates are pushing both for freedom from harassment and for fairness in terms of rights, benefits, obligations, and opportunities (ABC Of Women Worker’s Rights And Gender Equality, Christos, 2009). In part, this stems from a consensus that the failure to achieve equal treatment results from persistent discrimination based on gender. From an economic point of view, gender discrimination can manifest itself in many forms, affecting women both as workers and consumers.

Despite the existence of many studies on gender-based pay inequality, considerably less attention has been devoted to inequities in product and service markets. These can, however, have serious consequences for welfare and equity of women. For example, measures of consumer prices in the US typically reveal that products for women are more expensive than similar products for men. A recent study by the New York City Department of Consumer Affairs reports that women products cost on average 7% more than similar products for men (Bessendorf, 2015). A prior study by the State of California further claims that women paid an annual “gender tax” of approximately \$1,351 for the same goods and services as men or about a \$15 billion for all women in California¹.

Some argue that these price differences are evidence of sellers’ prejudice against women or sellers’ inferring that women have a higher willingness to pay (Ferrel, Kapelianis, Ferrel and Rowland, 2016). This seller’s prejudice corresponds to Becker (1957) taste-based

¹This study led California to become the first state to enact a bill to protect consumers from price discrimination for services (CA State Senate 1995, Gender Tax Repeal Act of 1995, AB 1100, Aug 31, 1995). Following suit, many states and counties also passed bills to prohibit businesses from charging different prices for products or services based solely on the customer’s gender.

discrimination model, where sellers have a distaste for providing services and products to women. To compensate for this distaste, women have to pay higher prices than men. In contrast, sellers use observable information about customers to infer their willingness to pay in a statistical discrimination framework (Arrow, 1973, and Phelps 1972). In this setting, sellers may consider that, absent of additional information, women have a higher average willingness to pay. Therefore, women can be charged higher prices than men. Both the existence of gender-based discrimination and the motives behind it, however, are in fact very difficult to prove. In many cases, the products being compared are not identical, and, moreover, it is frequently impossible to disprove that differences in pricing arise from differences observable to sellers but not to researchers. In the case of service markets, this difficulty is even greater as the lack of information about prices paid by each consumer creates yet another obstacle to inference.

In this paper, I circumvent these difficulties by conducting a large-scale field experiment on the US car repair market to test both the extent and patterns of gender-based price discrimination. Building on the correspondence testing methodology (Bertrand and Mullainathan 2004, Jowell and Prescott-Clarke, 1970 and Riach and Rich 1991), I sent more than 57,400 emails to shops requesting an estimate for the same standard car repair.

Importantly, I randomized two components of the email content observed to shops. I varied the perceived gender of the consumer by using distinctively male and female names. This variation effectively allows me to measure gender-price gaps. I also varied other customer attributes observed to shops, such as having low search costs, high implicit income or no car repair knowledge. That is, I added pieces of information to a “baseline” email revealing that the customer is searching for other estimates to signal low search costs and informed customers. I also revealed that the customer is not knowledgeable on automobile repairs to signal an a customer that may be easier to overcharge. Finally, I signaled that the customer is highly-educated as a proxy income. This signals a customer with higher

opportunity cost of time and willingness to pay. These customer type signals sought, in particular, to shift shop's priors on customers willingness to pay, and help distinguish between taste-based and statistical discrimination motives. In a taste-based discrimination framework, women would always pay a premium to compensate for sellers' distaste from serving them, whereas price gaps would disappear with additional information under a statistical discrimination framework.

Overall, the experiment resulted in a sample with 10,323 valid price estimates, corresponding to a response rate of close to one quote for every five emails sent. I find clear evidence of gender-based price discrimination in the "baseline" group. Indeed, women in this group receive price estimates of roughly 486 dollars - 9.4 dollars more than men. Reassuringly, this statistically significant difference is robust to the inclusion of additional controls, such as geographic shop characteristics and elements of the field experiment design.

Furthermore, this gender price gap is not driven by differential response rates or changes in the composition of the provided quotes. I find no evidence of gender gaps in non-monetary outcomes. For instance, using the information on the detailed quote components, I find that shops do not include additional parts (on top of labor, a radiator, and coolant) depending on customer's gender. There is also no evidence of differential treatment in terms of discounts, price matching, free inspections, and other such offers. This non-price discrimination dimension is also relevant in my setting, in particular, because car repair services are a typical example of credence goods (Dulleck and Kerschbamer, 2006). That is, consumers can observe utility they derive from the good ex-post, but they cannot judge whether the type or quality of the good they have received is ex-ante needed, giving rise to potentially inefficient over-or under-treatment as well as overcharging (Darby and Karni, 1973, and Dulleck, Kerschbamer and Sutter, 2011).

Importantly, the theoretical relationship between competition and price discrimination is ambiguous, making this an empirical question. In this paper, I rely on the detailed cross-

sectional nature of the data - where I gathered all car repair listings in the US - to construct measures of competition based on the number of competitors that are near each shop. I relate these measures to my experimental results and find a significant negative association between competition levels and gender price gaps. Car shops without nearby competitors give quotes to women roughly 15 dollars more than men, and this gap is on average 1.5 dollars lower for every additional competitor within a 1-kilometer radius. It is worth noting that Becker's taste-based model predicts that biased sellers would be competed away in a competitive market. As I show in a simple framework, this finding can also be rationalized in a statistical discrimination framework in which sellers are allowed to price discriminate based on customer characteristics but also with the ease in which customers can switch between shops. Consistent with this result, I find that shops with presumably more market power (franchises and dealerships) tend to increase gender price gaps relatively more than independent shops.

In contrast to the observed gender-price gap in the *baseline* group, customers in the other groups receive price estimates that do not vary significantly by perceived gender - although price shifts vary by group. Women get lower estimates once they signal low search costs (*search* group customers), while men estimates remain the same, effectively closing the gender price gap. By contrast, signaling high income does hurt men as they receive higher price estimates, while female estimates do not increase. These results suggest that while being a woman customer may lead shops to charge a higher price, gender becomes less relevant once shops learn additional information such as customer's willingness to search or their implicit income. Hence, discrimination appears more likely to stem from signal extraction problems (Aigner and Cain, 1977; Arrow, 1973; Phelps, 1972), rather than purely "taste-based" motives (Becker, 1957).

Following Antonovics and Knight (2009) and Anwar and Fang (2006), I test for the existence of in-group biases by tracking perceived employees gender. I do not find evi-

dence suggesting that employee gender influences the price estimate outcome, and women customers do not get better deals with women employees, providing additional support towards a statistical discrimination motive for price discrimination.

This paper contributes to the literature on discrimination in consumer markets. While most of this literature focuses on differential call-back rates- predominantly in the housing and rental markets (e.g., Ahmed and Hammarstedt 2008, Bosch et al. 2010 Edelman et al. 2017, Ewens, Tomlin and Wang 2014, and Hanson and Hawley 2011)- there is a reduced but growing literature that focuses on discrimination in consumer transactions (involving monetary outcomes).² This paper particularly relates to this latter literature, allowing the detection of discrimination beyond differences in reply rates to measuring the magnitude of discrimination's effects. In this case, measuring how much of the differences in provided estimates between men and women is due to discrimination. This paper contributes to a small but growing literature that studies discrimination in consumer transactions and provides new and robust evidence in favor of the hypothesis that price discrimination is both extant and economically meaningful. It extends the findings of academic studies that show race- and gender-based discrimination in auto and credit markets (e.g., Ayres and Siegelman 1995, Pope and Sydnor 2011 and Ravina 2007). In contrast to these prior studies, the observed price gap is lower, but still provides contrary evidence to non-experimental recent studies that have not found evidence of discrimination in online markets (Morton, Zettelmeyer and Silva-Risso, 2003).

Second, this paper contributes to the literature studying the mechanisms of discrimination. In line with the findings of Castillo, Petrie, Torero and Vesterlund (2013), Gneezy, List and Price (2012), and Busse, Israeli and Zettelmeyer (2017), this paper finds that price discrimination meaningfully changes in response to changes in the perceived customer profile. Castillo, Petrie, Torero and Vesterlund (2013) send six paired men and women testers

²See Bertrand and Dufflo (2017), Riach and Rich (2002), and Guryan and Charles (2013) for reviews of studies on discrimination.

to negotiate taxi fares using similar scripts in Lima, Peru. Differences in fares disappear once the testers reject the first taxi fare and proceed to the second taxi in line. Busse, Israeli and Zettelmeyer (2017) implement a field experiment in which four men and five women callers request price quotes for a radiator replacement from repair shops in heavily populated areas. They find that price gaps disappear once the average US repair price is referenced. Gneezy, List and Price (2012) send six disabled and six non-disabled testers to request price quotes at 36 auto body shops in Chicago and find that differentials also disappear once individuals mention they are getting other quotes. Given the scale of this experiment, this paper further advances upon the findings of these studies by interacting gender with other customer characteristics and combining the experimental results with novel geographic, shop, and market characteristics.

From a methodological point of view, the correspondence study approach also ensures more robust comparability across treatment groups than was available in prior studies, guaranteeing that any observed differences are caused solely by the gender and customer type trait manipulation.³ This methodology also allowed me to reach shops across the US, expanding on the previously more localized studies.

I also contribute to the empirical literature on price dispersion and competition, given the scale and rich geographical variation of this experiment. In contrast to a number of prior observational studies showing that price dispersion varies directly with the level of competition in a market (e.g., Borenstein and Rose 1994, Stavins 2001 and Shepard 1991)-in particular in the airline and retail gasoline market-, this paper joins the studies finding that increases in competition decrease price dispersion (Barron, Taylor and Umbeck 2004, Gerardi and Shapiro 2009). Importantly, my setting differs from the previous literature by relating competition to experimental results, and by allowing sellers to directly price discriminate between observable consumers characteristics, rather than having consumers

³Many of the weaknesses of audit studies have been discussed in Heckman (1998) and Siegelman and Heckman (1993).

self-select into different products.⁴

Finally, this paper relates to the literature on credence good markets, in particular to Schneider (2012), who conducts visits to car repair shops knowing in advance the set of car defects. He finds evidence of poor diagnosis, which does not vary when posing as a one-time versus a repeat-business customer.⁵ My analysis confirms these observations, in that recommended repairs do not vary with customer characteristics. In a broader perspective, I complement this literature by providing further evidence on how prices vary with other consumer characteristics.

The remainder of the paper proceeds as follows. The next section gives an overview of the experimental design, implementation and data collection. Section 3 describes the experiment data and discusses its main concerns. Section 4 presents the empirical results. Section 5 concludes.

2 Experimental design

Building on the correspondence testing method, I conducted a large-scale field experiment in the US car repair market to test if there is a gender-based price gap and how it varies with additional information about customer and market characteristics.⁶

This approach has several advantages with respect to prior studies using audit and phone call testing methods. It ensures more robust comparability across treatment groups by eliminating differences in individual characteristics observed to shops but not to researchers.

⁴Doleac and Stein (2013) and Nunley, Owens and Howard (2011) field experiments in online auctions examine how consumers discriminate based on sellers perceived race. Nunley et al. (2011) define competition by space display of shops products, Doleac and Stein (2013) by the bids consumers make. To the best of my knowledge, these are the only other studies relating experimental results to competition, but defining competition by demand measures.

⁵See Kerschbamer and Sutter (2017) for a survey of recent laboratory and field experiments on credence goods

⁶This project does not require IRB approval since it does not meet the definition of human subjects research as defined in Title 45 CFR 46.102(f), as verified with Brown Human Research Protection Program office.

Another advantage of this approach is that it allows for a larger sample size due to its low marginal cost. Conducting this experiment at such large scale provides more precise estimates. It allows the combination of experimental price results with larger variation in market characteristics, in particular with varying the degrees of market competition.

For this experiment, I constructed a comprehensive database of car repair businesses in the US by scraping the online Yellow Pages (www.yellowpages.com). For each shop, I retrieved information such as the email address, geographic location, and type of services provided. The data allow for the exploration of which geographic characteristics and, importantly, whether competition levels are associated with increased gender-based discrimination.

Between July and August 2018, I sent emails to more than 57,400 car repair shops using customized scripts. In these emails, I requested a quote for a radiator replacement for the same car. Using the same standard repair and car ensures that the quote provided is for the same service and therefore, always equally costly, regardless of customer gender. Importantly, I randomized the informational content of the emails while perfectly controlling all information about a customer observed to each shop. I randomized the gender of the customer, but I also varied additional information about the customers to signal implicit household income, knowledge about car repairs and low search costs. These signals are intended to shift shops priors on customers willingness to pay and test whether this information becomes more relevant than information on gender. These variations added to a standard (*baseline* type) customer represent a customer who is searching for other estimates (*search* type), a customer who is not knowledgeable on car repairs (*uninformed* type), or a customer who is highly educated to proxy for income (*income* type).

Overall, half of the shops had any reply to my emails (including all types of replies), and I obtained a sample with 10,323 valid price estimates, corresponding to a response rate of close to one quote for every five emails sent. The remainder of the section reviews

the experimental design details.


2.1 Auto repair shops

The car repair market provides an ideal setting to study price discrimination by gender. This is a typical credence good market, where there is a typical suspicion that shops can overcharge customers, and potentially treat different customers differently (e.g., *informed* or men customers could be treated differently). Thereby, implying a shop might have market power and potentially have the ability to segment customers. Furthermore, it is very difficult to resell a service once it is provided, leaving arbitrage opportunities minimized. These features make price discrimination feasible (Stole, 2007). With this in mind, I collected information on all US auto repair shops listed in the online Yellow Pages (YP) (www.yellowpages.com). The YP offers a comprehensive listing of businesses nationwide and are used by nearly 60 million customers in the US each month, with auto repair shops being one of the popular categories. As of June 2018, YP had over 230,000 listings in the auto repair category of which 40 percent have an available email address. Each listing also details information on the shop's name, location, website, and type of services provided.

I imposed several restrictions on this collected sample. First, since my inquiry is about a radiator replacement, I used the information on the type of services provided and shop names to keep shops specialized in mechanic car repairs. For instance, I dropped body paint shops, or repair truck and RV shops. I also kept shops with unique email addresses to avoid the possibility of contacting the same repair shop twice. My final sample counts approximately 58,300 auto repair businesses. Figure 1 shows the location of the shops in my sample. Shops are dispersed across the US, with more facilities located in more popular areas.⁷

⁷Alaska and Hawaii are included in the sample. Shops in Puerto Rico are excluded from the sample as they would expect to interact with their customers in Spanish.

Figure 1: Distribution of shops



Note: This figure shows the distribution of the 58,324 shops included the final sample. Source: Online Yellow Pages.

The information retrieved from the YP allows me to construct other variables of interest. First, I classified each repair shop as independent and non-independent (franchises and dealerships).⁸ This allows me to test if price patterns vary with shop types. Independent shops may have more discretion setting up prices as they do not have stringent pricing rules, and it is easier to reach the mechanic in charge of a repair directly.

Second, I built market competition measures using the YP geographic location of each shop and GPS software. With Google Maps geocoding API, I converted addresses into geographic coordinates (latitude and longitude), and then used QGIS to compute each shop's distance to its competitors to obtain the number of nearby competitors in alternative distance radiuses. This is meant to explore the correlation between competition and price

⁸I gathered listings of US franchises from online searches in websites such as entrepreneur.com and franchisegator.com. All franchises are available in the YP, and a rough comparison suggests that the YP have in fact more shops than those currently open according to online searches. These likely correspond to closed businesses that are still listed in the YP.

gaps. Finally, geographic shop locations are used to match each shop to their neighborhood.

2.2 Email design

As the second step of this experiment, I used this comprehensive database of car repair businesses to send emails inquiring about the cost of a standard car repair of one of US most popular cars (replacing a Honda Accord radiator). Importantly, I built a baseline email script and randomized its informational content along two crucial dimensions. I varied the perceived customer's gender by using names that are distinctively recognized as female or male. I also varied the signals about customer *types*, distinguishing between customers with no car knowledge, searching for other quotes, or with implicit high income.

Using these additional variations to the email content can help distinguish the motives for discrimination. Following Becker (1957) taste-based discrimination framework, if shops have a distaste towards women customers, they would compensate for this disutility by providing women with higher quotes for the same type of repair. When the fraction of shops discriminating against women is large enough, a gender-based price differential would exist reflecting this distaste parameter. Importantly, a gendered price gap would be observed across all types of customers. In contrast, imperfect information could lead to gendered price gaps in a statistical discrimination framework (Aigner and Cain 1977, Arrow, 1973, and Phelps 1972). Discrimination would be the result of a signal extraction problem, where -absent of additional information- shops use customer's gender to infer how profitable it could be to sell to them. But, once shops have other valuable customer-specific information (such as the customer *type*), they may rely less on gender and more on these additional signals to provide quotes. In this second framework, for instance, revealing to have low search costs could help customers get more competitive prices. If shops prior were that women have higher search costs in this industry, then this signal would be more helpful for women than men. Importantly, under statistical discrimination, gendered price gaps would

not be observed across all customer types.

2.2.1 Customer gender variation

The first variation I introduced refers to the perceived customer gender. I created eighteen email accounts using distinctively female and male names. The eighteen names are drawn from Bertrand and Mullainathan (2004) and Levitt and Dubner (2006), all of which are popular and are associated with a white-sounding name. I also paired these names with popular white-sounding last names from Bertrand and Mullainathan (2004).⁹ Choosing first *and* last names that are predominantly associated with white individuals allows me to use these uniquely as signals of customer gender, ruling out the possibility that changes in price estimates are associated with changes in race. By the same token, I chose to use the same Gmail email domain for all accounts.

Each customer name is observed by sellers in the email address, email account name, and email content signature.

2.2.2 Customer types variation

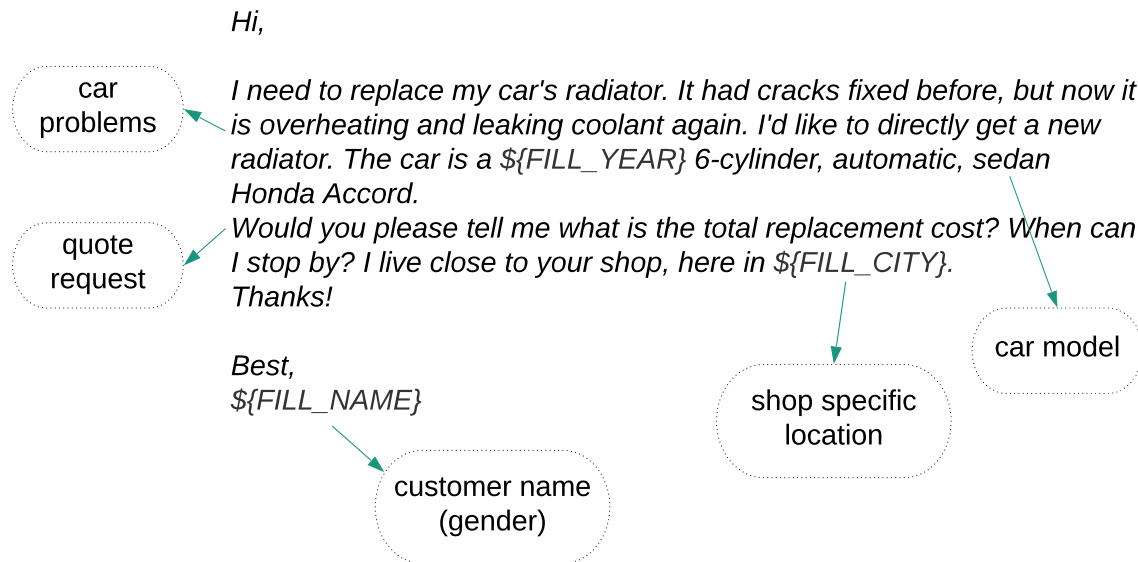
In addition to signaling customers gender, I added variations of customer attributes to a *baseline* email content (i.e., script). These variations to a *baseline* customer type are either a *search*, *income* or an *uninformed* type. They aim at shifting each seller's prior of

⁹ I use all the white-sounding first names in Bertrand and Mullainathan (2004) except for Meredith, Laurie, Brad, Brett, Jay and Todd as they have a smaller relative likelihood of being identified with the intended gender, according to online searches (e.g., gpeters.com). To extend the list of names, I chose six additional white-sounding names from Levitt and Dubner (2006). The selected names are consistent with Bertrand and Mullainathan (2004) study. That is, in addition to using popular female and male names in the US, each of these additional names is also among the 100 most popular baby boy and girls names as in Massachusetts between 1970 and 1986 (based on the Social Security Administration database on birth counts by year, state, name and gender). The complete list of names is Allison, Amy, Anne, Brendan, Carrie, Colin, Dustin, Emily, Geoffrey, Greg, Heather, Kristen, Matthew, Neil, Sarah, and Scott, and the last names are Baker, Kelly, McCarthy, Murray, O'Brien, Ryan, Sullivan and Walsh. The last names were randomly matched to the first names. Finally, all accounts use the same Gmail domain and a number or letter randomly added to the email address due to unavailability of first-name last-name only e-mail combinations.

consumers willingness to pay and are used to compare how gender price gaps vary, with the baseline type as the comparison group as specified in the Pre-Analysis-Plan (RCT ID: AEARCTR-0003279). While scripts vary the conveyed information on knowledge, income and quote search efforts, they are otherwise identical in structure and content.

Figure 2 presents a template script used with the *baseline* customer. This customer script mentions a car model, radiator problems, asks for a total cost estimate for its replacement, and also mentions living close to the shop. The car is one of the most popular US cars - a Honda Accord. The problems are leaking coolant and overheating. Both problems can be associated with a faulty radiator, according to the Federal Trade Commission and RepairPal website, among others. The reference to the residence near the shop is used to increase the apparent veracity of the email.

Figure 2: Baseline customer template script



Note: This figure shows the script template used for a baseline-type customer. Text between " $\{\}$ " is filled with specific randomized component.

For each of the three additional customer types, I modified the baseline email script

as follows. First, for *uninformed* type customers, I described the radiator problems using vague, non-technical terms - for example, I mentioned that the car “leaves green puddles” instead of “leaking coolant,” and that the “temperature thing” is raising instead of the “gauge.” With this language I aim to signal that the customer has no car knowledge, giving more discretion to the seller to increasing prices or offering unnecessary additional repairs. Second, for the *search* type customers, I added a sentence where I state to be searching for other price estimates from nearby shops. In a subset of these emails, I also mention having checked a website (RepairPal) with the reference area prices for this standard car repair, signaling further market knowledge. Thus, with this quote type customer, I aim to signal that customers have low-search costs, making sure shops know with certainty that customers are exerting some effort to obtain more competitive estimates. Last, for the *income* type customers, I add the title “Ph.D.” to the email signature. This high education signal intends to proxy for high income, signaling shops that this customer has a higher opportunity cost of time.

The script templates were pre-tested in a pilot carried out in June 2018, and all templates are available in Appendix 6.B.

2.2.3 Other email items variation

Finally, in addition to varying customer’s gender and type, I varied other email features used to contact each shop. These variations are the script version for each customer type, the email subject, and the car year. There are two scripts for each customer type that provide the same information but are reorganized and slightly differently worded. The email subjects all refer to a radiator, but in some of them I use the words “new,” “replacement” or “change.” While the car mentioned in the emails belongs to the same car generation, I varied the car years mentioning it is either a 2009 or a 2010 model.

The purpose of changing these additional email features is to avoid sending an identical

email to shops while not affecting shop price estimates. All price estimate differences should arise due to the randomly assigned gender and customer type used to contact each shop. To ensure this, I include these additional email features as controls in regression specifications as specified in the Pre Analysis Plan. As expected, the discrimination pattern results are robust to their inclusion. Moreover, the coefficients of these control variables are never statistically significant.

2.3 Randomization

The randomization assignment of email scripts to car repair shops was done unconditionally on any further characteristics (e.g., geographic or shop characteristics). As a first step, half of the shops were assigned to be contacted by a woman customer and the other half by a man, using one of the created email accounts. Next, I randomly assigned one of the four customer types to each shop within each gender. I used optimal shares obtained from power calculations. These power calculations were obtained using pilot observations and weighted more the baseline and quote types, resulting in a distribution of shops to be contacted by 18 percent baseline, 16 percent quote, 8 percent uninformed, and 8 percent income scripts of each gender. Lastly, I reshuffled the sample and assigned one of each remaining email items (car year, script version and the subject of the email). Appendix 6.A details this randomization procedure.

2.4 Follow-up of email replies

Besides following a protocol to send emails, I also followed pre-specified email guidelines to follow up the shop's responses. Importantly, I drafted scripted replies in the pilot in anticipation to shop's expected inquiries, and I allowed for one reply to each shop when they did not provide a price estimate in their first reply. These guidelines guarantee that

all treatment groups behave identically, and thus, observed price estimates only vary with information on customer’s gender and types. Appendix 6.C details these rules and includes the reply template scripts.

2.5 Implementation and data collection

The previous subsections described the experimental design. In this subsection, I describe the implementation and data collection used to guarantee comparability across groups.

I used automated tools to contact each shop via email. This wave of emails took place from July through August 2018, sending an average of 8,000 weekly emails between from Monday through Friday between 10:30 - 11:30 AM Eastern time. Overall, I sent approximately 57,400 emails. This resulted in a response rate of one half and a valid price estimate in close to one out of five emails sent, yielding 10,323 price estimates.

Once the interaction with the shops was completed, I proceeded to process the information from each reply and then match it to the shop data, which contains each treatment group and email details assigned to a shop as well as shop characteristics.

The information collected can be grouped into three main categories: reply type and employee first name, the salience of experimental design features, and price and service information. The first category records if, when and which type of answer is received. The type of answer includes a classification of each email content according to whether the email was invalid, the shop does not perform the service, the shop provides the service and gives a quote, and the shop provides the service but does not provide a quote. This category also tracks how much effort customers exert to receive a quote, measured by an indicator equal to one if a second email (reply) is sent to a shop before receiving a quote. Employee first name is recorded to assign them their likely gender, of which 80 percent of the employees have an assigned gender and, 17.5 percent correspond to woman first names. The gender is assigned based on the SSA National Data on the relative name frequencies by gender in

the population of U.S. births from 1940 through 2005. If a name is associated with women in more than 70 percent of the occurrences, then that name is classified as corresponding to a woman.

The salience of experimental design features measures if shops replies mention relevant content from the original email. This content intends to capture items from the experimental design, thus I constructed indicator variables for mentioning the customer's name, gender (e.g., Mr, Miss), and profession in the case of income type customers. I also created an indicator variable to record when the shop's reply makes a reference to RepairPal. I restricted this sample to replies from shops that perform the requested radiator replacements, as these integrate the sample of interest.

Finally, the pricing information collected includes price and service details, as well as other offers customers receive. The total price estimate is the primary outcome of interest and was processed using pre-specified rules. I defined total price estimates as follows: the total discounted price provided in a quote, the average price whenever price ranges are provided, the first price provided by a shop whenever more the shop provided more than one quote in separate emails. I also kept the estimate provided by a shop regardless of additional price-match offers. I exclude estimates defined as invalid if shops explicitly excluded either the price for the labor or the radiator, prices were extremely low or high (below 100 dollars or above 2,000 dollars), or estimates had a price range ratio over twice the price range and no further explanation about price components was given. These thresholds are arbitrary but conservative values. Overall, 119 price estimates are dropped from the sample as they are considered invalid. In addition, I recorded price component details whenever available (i.e., the estimate includes a thermostat replacement, or additional radiator hoses).

Regarding the service details, I recorded shop offers such as price-match of competitor estimates, mentions of a warranty on their service, discounts, or any additional offers such as a car loaner, shuttle service, and financing options. These service details are collected

to test whether shops change their behavior with each treatment group.

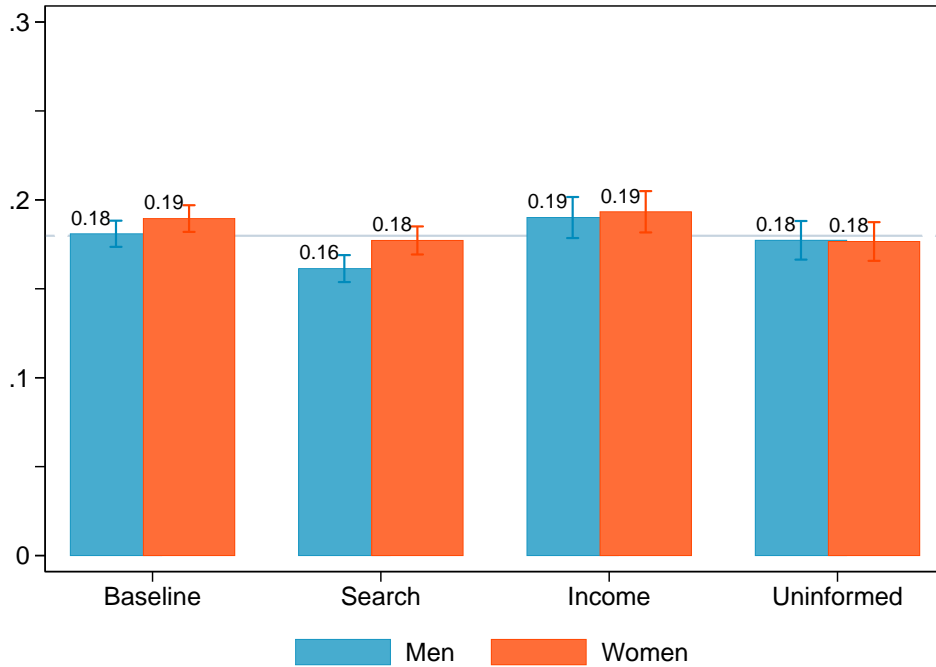
3 Final sample and design validation

The experiment resulted in a response rate of one half, and a valid price estimate in close to one out of five emails, yielding 10,323 valid price estimates. While the aggregated follow-up rates are similar across groups, there are some minor differences in the rate of provided quotes, as illustrated in Figure 3. These small differences may introduce selection problems, which are discussed in the next section. The differences are, that women in the *baseline* customer type groups receive 0.9 percentage point more replies than *baseline* men customers. This *baseline* group difference however, not significant when comparing the raw data or adjusted data. The second difference across genders is that in the *search* type groups women also receive more replies (1.5 percentage points). Within gender groups, women in the *uninformed* group receive an average of 1.2 percentage point fewer estimates than in the *baseline* group. One could speculate that this difference may be explained by shops being slightly less confident on the need to replace a radiator when an *uninformed* woman customer contacts them. Lastly, there is a 1.5 percentage point decrease in the reply rate for customers in the *search* group relative to the *baseline* group.

There are other three main issues of potential concern for the sample of shops that perform the requested repair and reply to my emails. First, are shops used to communicating by email and providing estimates? Second, is the experimental design valid - that is, do shops perceive as salient the variations by gender and customer type? And third, is the randomization balanced across treatments?

For the first concern, I document that shops performing Honda radiator replacements in the final sample use email communication on a frequent basis. Table 1 shows that more than 90 percent of shop replies are received within two days since the first email

Figure 3: Price estimate reply rates



Notes: This figure shows average shop replies with price estimates across gender and customer types: Baseline, search, income, and uninformed. The plot shows the 95% confidence intervals and the sample mean reply rate is marked by a grey line within the plot.

inquiry, suggesting that email communication must be common enough for them to reply promptly to the inquiries. Accordingly, less than one percent of shops mention they do not usually correspond by mail with customers. Furthermore, more than 80 percent of quotes are obtained without the need of sending a follow-up email and 13 percent of replies also provide an estimate in an attached document (oftentimes through their internet quote system). Altogether, this suggests that shops in the sample use and monitor their email frequently, are used to communicating with customers by email, and do not seem to find the email price requests suspicious.

With respect to the second concern about the experimental design saliency, I find evidence suggesting that employees pay attention to the email content. Table 1 shows that

half of the replies mention the customer’s name, indicating that this feature is salient to shops. Furthermore, the name to whom the reply is addressed to corresponds to the assigned name found in the email signature, and not to that in the email address - which is a combination of the name and an additional letter or number (e.g., shops mention Geoffrey instead of Mgeoffrey observed in the email murray.mgeoffrey@gmail.com). About three percent of shops make a reference to the customer’s gender, another three percent to the professional title (in the case of *income* customer types), and about two percent of shops explicitly mention RepairPal website after the initial email does so. In these latter cases, shops usually refer to where their estimates stand relative to RepairPal area price range. Finally, shops include car details mentioned in the initial email when they send an attached

Table 1: Summary statistics: Shops email use and salient design items

	Total (1)	Men (2)	Women (3)	<i>p</i> -value (4)	N all (5)
A. Reply detail					
On 1st day	0.77	0.77	0.76	0.49	10,323
On 2nd day	0.14	0.13	0.15	0.03	10,323
Not use email	0.002	0.002	0.002	0.41	10,323
Price on 1st reply	0.83	0.83	0.82	0.19	10,323
Price on attachment	0.13	0.12	0.14	0.07	10,323
B. Reply content					
Name	0.55	0.52	0.57	0.00	10,323
Gender	0.03	0.04	0.02	0.00	10,323
Profession	0.03	0.04	0.02	0.07	1,704
Repair Pal	0.02	0.03	0.01	0.00	1,436

Note: This table reports the shares of responses from shops that do Honda radiator replacements for all customers (column 1), men (column 2), and women (column 3). Column 4 reports the *p*-value from each statistic means comparison test by gender. Column 5 reports the total number of observations. Price on 1st reply row restricts observations to shops that provide valid quotes, Profession row restricts to observations where *income* treatment is assigned to each shop, and Repair Pal row restricts to observations where initial email uses the *search* customer type treatment with the “repair pal” search reference.

document with an estimate. This is not a feature intended to be salient, but suggests that employees read emails carefully.

Regarding the third and last concern, I provide evidence that the design is valid. First, the geographical distribution of shops across groups seems balanced both by gender and within gender-customer type groups (see Figures A.3). All groups are distributed similarly across the US, with more density in more populated markets. Reassuringly, there are no statistically significant differences by gender or gender-customer type with respect to the share of invalid emails, which further validates the randomization (See Figure A.4). Finally, 1.7 percent of emails originally assigned to a shop were not sent due to implementation issues such as having one email account disabled. This discrepancy between assigned and sent emails, however, did not modify the distribution of treatment groups.


4 Results

4.1 Is there a gender price-gap?

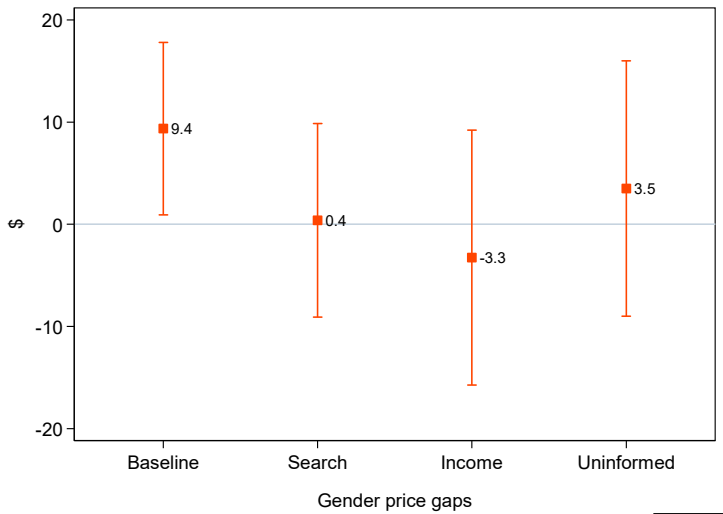
Figure 4 presents the main results. Panel (a) shows the mean price estimates by gender and customer type, while panel (b) plots the gender price-gap for each type. This figure shows that there is a gap for customers in the *baseline* group. Indeed, women in this group receive price estimates of roughly 486 dollars - about 9.4 dollars more than men. As later detailed, this gender price gap is not driven by differential response rates, changes in the composition of the provided quotes, and is robust to alternative specifications and alternative measures of total quotes provided to shops.

The implied gender gap in the baseline group is a small but a significant amount per repair. The 9.4 dollar gender differential in the *baseline* group represents a 2 percent gap. This represents clear evidence that gender-based price discrimination still persists, but also

Figure 4: Price estimates



A. Price Levels



B. Gender price gaps

Note: These figures show the price estimate differences across gender and customer type. Panel A plots raw average price estimates with their 95% confidence intervals separately for each group. The sample average price estimate is marked by a grey line within the plot. Panel B plots the average gender price gaps with their 95% confidence intervals for each customer type. Positive values imply a higher price estimate for women with respect to men in that group, and the case with no gender price gaps ($y\text{-axis} = 0$) is marked by a grey line within the plot.

that it is significantly lower than other estimates in the academic and policy literature. In a related setting, Busse, Israeli and Zettelmeyer (2017) find a gender price gap of almost six percentage dollar increase for women.¹⁰ This difference may be due to the changes in discrimination over time, as Busse, Israeli and Zettelmeyer (2017) experiment took place six years before, in 2012. Another important feature that may be driving the difference is that I used email communication while they called shops. Using email communication may increase transparency and decrease shops discretionary behavior as estimates are written, and help protect disadvantaged groups (Morton, Zettelmeyer and Silva-Risso, 2003). Thus, my results suggest that internet communication is particularly beneficial to individuals whose characteristics disadvantage them in personal interactions, although it does not completely remove price differences.¹¹

The results are robust to the inclusion of experimental design components as well as to state and commuting zone fixed effects. Appendix Section 9 replicates the analysis in a simple regression framework. Figure A.8 plots point estimates from a regression-adjusted specification of price estimates for each group relative to men in the baseline, and shows the main results are similar. Table 2 shows additional robustness checks.¹² Column 1 controls for state fixed effects, additional varying email items (car year, subject controls and script version used), and day of week and week fixed effects indicating when each shop is contacted, as pre-specified in the analysis plan. Column 2 adds a non-independent shop type control (i.e., franchises and dealerships) since their price levels are usually higher, and column 3 adds Commuting Zone (CZ) fixed effects and an additional indicator dummy

¹⁰The cleanest comparison between Busse, Israeli and Zettelmeyer (2017) and this paper is between my *baseline* group and their *uninformed* group, where agents request a quote without providing a reference price. This is the number I am reporting here.

¹¹One should be careful, however, of drawing too strong conclusions from this comparison, since there are important differences in the samples used in both studies. In particular, Busse, Israeli and Zettelmeyer (2017) focus on very populated areas, while I obtained estimates from 10,313 shops across the US, including less populated areas. Furthermore, since their study was implemented in 2012, one cannot rule out that the decrease in price discrimination is partially driven by changes over time, as suggested by Edelman et al. (2017) in a related setting.

¹²See Appendix A.6 for complete results

for the shops that were not matched to any commuting zone. These smaller geographic areas are included to control for economic diversity better. Overall, the regression-adjusted results are consistent with the raw price differences. Indeed, the gender-based premium in the baseline group remains significant across specifications and in fact increases to an average of 10 dollars.

Table 2: Robustness: Impact of gender and customer type on prices

	Dependent variable: Price estimates		
	(1)	(2)	(3)
Women	10.026** (4.236)	9.506** (4.115)	9.689** (4.176)
Search	-0.816 (4.555)	-0.175 (4.440)	0.389 (4.531)
Search × Women	-9.821 (6.396)	-11.317* (6.188)	-11.101* (6.291)
Income	12.772** (5.298)	12.248** (5.173)	13.999*** (5.351)
Income × Women	-13.792* (7.560)	-12.719* (7.349)	-13.856* (7.586)
Uninformed	-1.568 (5.450)	-3.169 (5.267)	-1.988 (5.367)
Uninformed × Women	-6.012 (7.601)	-5.000 (7.321)	-5.597 (7.500)
Controls	State FE, mail items, DOW, week FE	+ Shop type	+ CZ FE
Observations	10,323	10,323	10,323

Note: This table reports coefficients from a regression of total price estimates on an indicator for woman customer, and indicator for each customer type, and their interaction with woman customer, with baseline customer type as the omitted categories. Column 1 includes state fixed effects, other email varying items: car year, subject and script number controls, and day of week and week fixed effects - indicating when the email was sent to a shop. Column 2 includes a non-independent shop control (non-independent shops: dealership and franchises). Column 3 adds CZ fixed effects, with a dummy for observations not linked to a CZ. Robust standard errors in parenthesis, with significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Reassuringly, the results are also robust to alternative definitions of total quotes, such as excluding discounts in quotes (see Appendix Figure A.9). There is also no significant difference in the rate of replies received by women and men customers in the *baseline* group (see Figure 3), or a differential amount of components included in the quotes, on top of the radiator, antifreeze and labor (see Figure 5), as further detailed in next subsection.

4.2 How do prices vary with additional information?

By contrast to the observed gender gaps in the *baseline* group, gender gaps disappear once additional information is observed by shops. Customers in the *search*, *income*, and *uninformed* groups receive price estimates that do not vary significantly by perceived gender but price shifts have different patterns. This suggests that being a woman customer may lead shops to charge a higher price, but gender becomes less relevant once shops learn additional information, such as customer willingness to search or income. Indeed, Figure 4 shows that shops change their quotes based on perceived income and search costs. Signaling low search costs (i.e., *search* group) helps women get an estimate of 475.7 dollars, 10.6 dollars lower than their *baseline*. Men, by contrast, do not profit from this signal. Their quotes remain the same, suggesting that being a male and having low-search costs in this market are considered substitutes. By contrast, signaling high income seems to particularly hurt men, who receive estimates of 488.9 dollars, similar to women in the baseline group. Women quotes in this group remain the same. Finally, revealing ignorance on car repairs do not seem to have an impact on prices, although these estimates are noisier. Overall, these results could suggest that shops priors are that, absent of additional information, woman customers have a higher willingness to pay. When information about search costs and income is provided, shops update their priors resulting in no price differentials by gender.

Reassuringly, results are robust to the inclusion of controls, and none of the additional experimental design features are relevant to explain variation in prices. The price decrease

for women in the quote search group is marginally significant, while the price increase for men in the income group is statistically significant and varies between an average of 12.2 to 14 dollars relative to men in the baseline group (See Appendix A.6).

Are there differences in reply rates?

A potential concern may arise from the differences in the rate of provided quotes, although these differences are not statistically significant in all specifications, as discussed in Section 2.5. This may introduce a problem of sample selection bias. For instance, women in the search group are 1.2 percentage points less likely to receive a reply with an estimate than women in the baseline group. It is possible that any given shop does not price-discriminate between men and women, but that the most expensive shops only reply to women. To address this concern, I implement a correction for selection following Heckman’s two-step approach. Replies might vary with day of the week, and population nearby, whereas prices might not. For instance, shops might be more busy and find it harder to reply on Fridays, whereas quotes should not vary with this. For completeness, I do this for the comparison groups that have significant adjusted differences in the rate of quote replies, or whose quote estimates vary between groups. Table A.7 in the Appendix shows the results, which suggest that gender-price gap in the baseline group is not driven by selection bias.¹³

Are there gender differences in diagnostic recommendations and customer service?

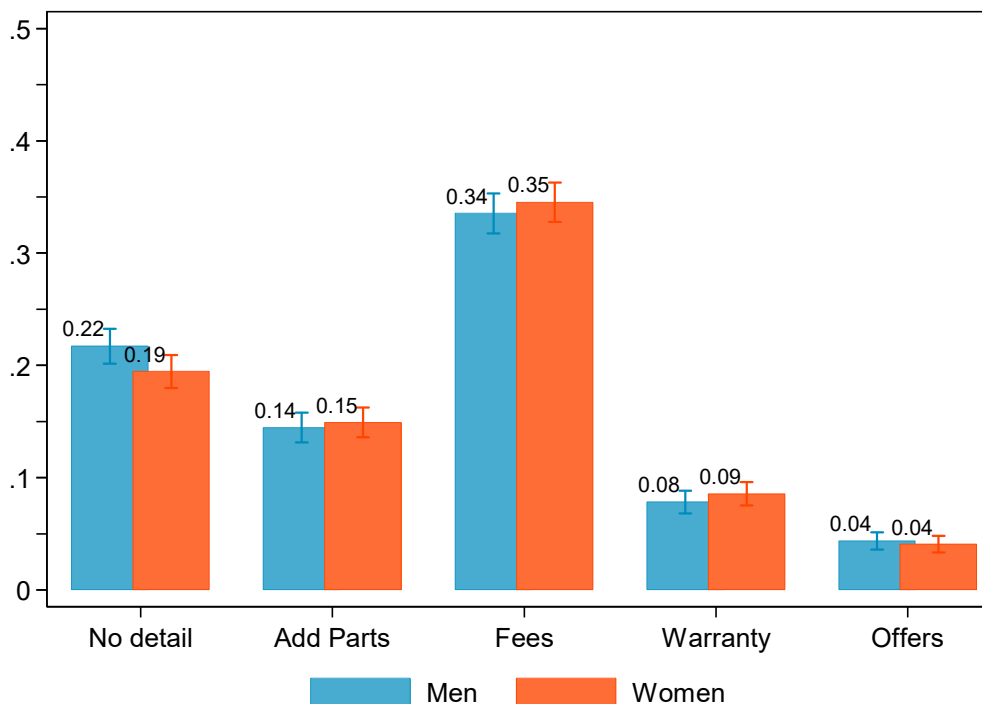
Shops can directly overcharge customers, but also vary prices by over-treatment of the repair. That is, by offering more than the necessary repairs. In my setting, customers specifically request a quote for a radiator but mechanics can also include the quote of additional parts, on top of labor, coolant and a new radiator in a quote. This brings

¹³I also implemented a non-parametric bounds estimation following Lee (2009). Unfortunately, unlike Heckman’s parametric correction, I can only use a limited set of covariates to tighten the estimates of the bounds. As a result, the difference between both bounds is large. I cannot reject the null of no effect in the bounds for any of the groups with significant differences in the rate of quote replies.

another channel through which shops can treat customers differently. Prices might be similar, but would shops include more items when providing quotes to women? Shops can treat customer differently by also changing complementary services they offer. Overall, I do not find differential recommendations by gender. Figure 5 shows that there is no difference in rates in which taxes and fees, additional replacement parts, warranties, and additional offers are provided. This suggests that shops do not provide additional components to women, nor do they seem to compete more aggressively to earn their business. There are also no differences in the amount of other recommended replacements by customer types. Table A.11 in the Appendix shows that shops are less likely to provide a detailed quote to *uninformed* men customers, while they are marginally more likely to offer these to *high-income* type customers and to women in the *baseline* group. This difference, however, does not translate into decreases (or increases) in additional recommended replacements relative to men customers in the *baseline* group. In the case of *uninformed* customers, it could suggest more skepticism in the customer self-diagnosis or suggest a smaller commitment from shops to keep effective prices close to estimated prices. The observed lack of differences in quoted parts by customer type resembles Schneider (2012) findings. In a different but related setting, he does not find evidence that mechanics suggest more or fewer repairs among one-time and repeat-business type customers.

Though customers do not receive additional parts differently, there are a few small but significant differences in the inclusion of taxes and fees, and additional services. *High-income* type customers receive more quotes that include shop fees and taxes than customers in the *baseline* group. This increase does not explain the observed differences in price estimates. The second difference is that men searching for nearby estimates are more likely to be quoted a service with warranty and additional offers, which include discounts, price matching, and financing options. This could suggest that when man customers signal the potential to substitute stores shops also try to compete more in the quality of customer

Figure 5: Additional parts and service offers



Notes: This figure shows share of estimates that: i) do not provide any detail on the components included (*No detail*) ii) include additional replacement parts on top of a radiator, coolant and labor (*Add parts*) iii) include taxes and fees (*Add parts*), iv) have a warranty (*Warranty*) and, v) include additional offers such as discounts, financing options and price-matching (*Offers*). The plot shows the 95% confidence intervals, obtained with robust standard errors.

care. Importantly, these non-monetary findings indicate the discrimination in prices is not driven by differential responses or additional repairs included in the quotes.

Are there quote differences by shop characteristics?

An additional exercise tries to distinguish between taste-based and statistical discrimination motives, as empirically tested in the literature (Fershtman and Gneezy 2001 Donohue and Levitt 2001, Antonovics and Knight 2009, and Price and Wolfers 2010, among others). These studies compare the treatment by different groups of potential discriminators, which

are distinguished by the same observable characteristic in which discriminatory behavior is suspected. Group bias would imply that people treat members of their own group more favorably than they treat other people. In my setting, I test if the employee's gender matters to explain gender-based differences in prices. If statistical discrimination alone explained differences in treatment, then everything else equal, treatment differences should be independent of the group characteristics. For instance, woman employees should not help other woman customers get better deals. I test this hypothesis in a simple regression where I include a dummy for perceived woman employees and its interaction with a woman customer dummy. Table 3 presents the results. I find that there is no significant change in price-gaps when a woman employee gives a quote. This provides no support to a taste-based motive driven price-gap. In addition, the observed price differences across treatment groups are robust to the inclusion of these controls.

Finally, I explore whether the main results are associated with observed shop types. Is it a common practice for dealerships and franchises to price customers uniformly? Presumably, independent shops have more scope to update estimates more frequently. Table 3 shows these results. As expected, non-independent shops (franchises and dealerships) provide higher estimates, men receive increased estimates of about 100 dollars higher than those from independent shops. Perhaps unexpectedly, however, the relative gender difference increases about 19 dollars more among non-independent shops. The reason for this rather surprising result is not obvious, but I rule out some explanations. First, there is no differential reply rates among non-independent shops. Second, I can also rule out that this increase in price gaps is driven by systematic differences in the composition of quotes. That is, franchises and dealerships do not include relatively more additional replacement parts in the price estimates provided to women than those provided to men. Consistent with List (2004) study, where more experienced recruited sellers provided higher prices to minorities who on expectation were inexperienced at a sports card market, one could consider that

Table 3: Effects by shop characteristics

	Dependent variable: Price estimates		
	(1)	(2)	(3)
Women	9.585** (4.326)	10.040** (4.317)	7.949* (4.410)
Women employee	8.373 (5.414)	9.162* (5.430)	8.741 (5.435)
Women employee \times Women	-3.474 (7.335)	-1.941 (7.370)	-1.137 (7.380)
Non independent shop	106.299*** (4.454)	109.703*** (4.549)	100.049*** (6.426)
Non independent shop \times Women			18.557** (8.955)
Controls	State FE, mail items, DOW, week FE	+CZ FE	
Women employee (%)	17.5	17.5	17.5
Non-independent	11	11	11

Notes: This table reports coefficients from regressions of total price estimates on shop characteristics (an indicator for perceived women employees, and indicator for non-independent shops) and their interaction with an indicator for woman customers, an indicator for woman customers, customer types and their interaction with woman customer (with baseline customer type as the omitted category). Column 1 shop characteristics refers to non-independent shop types, and includes state fixed effects, other email varying items: car year, subject and script number controls, adds day of week and week fixed effects - indicating when the email was sent to a shop. Column 2 adds Commuting Zone fixed effects with an additional dummy for shops not matched to a Commuting Zone. Robust standard errors in parenthesis, with significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

dealerships and franchises are be more *experienced* and have larger market power and are, thereby, able to discriminate more. This brings into consideration how market competition levels could impact gender price gaps.

4.3 Do increases in competition levels help close price-gaps?

The relationship between gender-price discrimination and competition is theoretically ambiguous. On the one hand, a monopolist with market power can perfectly price discriminate across consumers. Opposite to this, when markets are perfectly competitive and firms have no market power, then the law of one price prevails and price discrimination cannot exist. From these extreme cases, it would follow that the degree of price discrimination should decrease as markets become less concentrated. However, this theoretical prediction is not supported in intermediate cases. To the contrary, theoretical studies suggest that price discrimination actually can increase when moving from a monopoly to imperfect competition (Borenstein1985, Borenstein1994, and Holmes 1989).

Thus, the relationship between price discrimination and competition becomes an empirical question. This field experiment provides an ideal setting to study whether gender-based discrimination becomes more or less pervasive in more competitive markets. I combine experimental results of the effects of gender on prices with non-exogenous variation in market structure. Taking advantage of rich detailed data on the location of each car repair shop and the variation of geographic characteristics across the country, I define measures of competition based on the density of nearby competitors.

In my setting, I define competition levels by the number of competitors in each market and each market for a radiator repair by 1-km radius circles centered around each shop. Table 4 presents the main results, including several specifications to gain insights into the robustness of the results. Each column regresses price estimates on the number of competitors, the number of competitors interacted with a dummy for perceived women customers, a dummy for women customers, customer type controls, a dummy for non-independent nearby competitors to control for additional characteristics of the type of competitors, as well as the standard main controls from the experimental design (mail varying items, time, state and commuting zone fixed effects).

Overall, the results suggest that there is a robust negative association between gender

based price discrimination and increases in the level of competition. When there are no nearby competitors, the gender price gap for baseline customers increases to 15 dollars. This is a significant increase with respect to the 9 dollar gap shown in table 2. Each additional competitor is associated with a decrease in average prices. Importantly, it is associated with an additional significant decrease in estimates provided to women. Women receive estimates of roughly 1.5 fewer dollars with an additional competitor nearby. Taken at face value, this estimates imply that price gaps would disappear with ten shops.

Table 4: Effects on price gaps with competition

	Dependent variable: Price estimates		
	(1)	(2)	(3)
Competitors	-0.736** (0.344)	-0.246 (0.340)	-0.623* (0.336)
Competitors × Women	-1.466*** (0.473)	-1.448*** (0.463)	-1.470*** (0.459)
Women	15.379*** (4.577)	14.900*** (4.458)	15.274*** (4.507)
Controls	State FE, mail items, DOW, week FE	+ Shop type	+ CZ FE
Mean Competitors	3.88	3.88	3.88

Notes: This reports coefficients from a regression of total price estimates on the number of competitors within a 1-km radius, its interaction with a woman indicator, a woman indicator, and each customer type control, and their interaction with woman customer, with baseline customer type as the omitted categories. Column 1 includes state fixed effects, other email varying items: car year, subject and script number controls, and day of week and week fixed effects - indicating when the email was sent to a shop. Column 2 includes a non-independent shop control (non-independent shops: dealership and franchises). Column 3 adds CZ fixed effects, with a dummy for observations not linked to a CZ. Robust standard errors in parenthesis, with significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

This result suggests that while consumer heterogeneity may lead shops to price discriminate, the extent of competition also matters. This supports the hypothesis that price

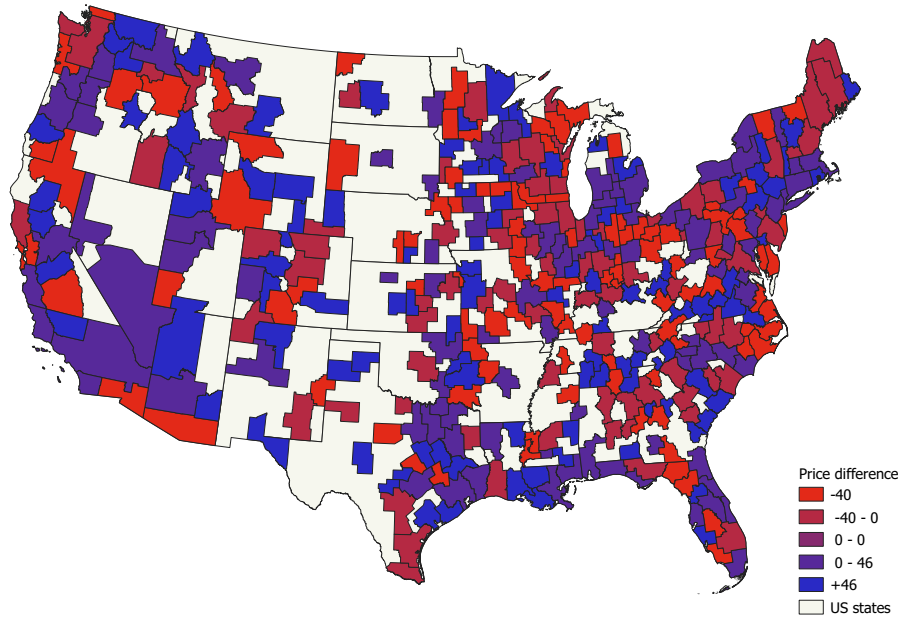
discrimination is negatively associated with the degree of competition. Unlike the previous exercises comparing prices changes by employee gender and customer types, these results are consistent with taste-based discrimination. Discrimination against women is less present in markets with more competition, as sellers for whom it is more costly to interact with women will be competed out by sellers who do not discriminate (Becker, 1957). Nevertheless, this result can also be rationalized by a statistical discrimination framework, in which sellers are allowed to discriminate based on observable customer characteristics as well as the level of market competition. Appendix 7 illustrates this mechanism in the context of a simple model in which statistical discrimination is allowed. In my setting, the degree in which shops vary prices with perceived gender also depends on the ease with which customers can substitute stores. As it becomes easier for customers to switch to a competitor shop, increasing prices based on gender becomes less profitable.

As an additional exercise, I show the results are robust to alternative samples and market definitions in Tables A.9 and A.10. First, I verify that the observed pattern is not driven by outliers and influential points by running a robust regression, which re-weights observations based on how well behaved they are (Fox, 1997, see Column 2, Table A.9. Second, I show that the results are qualitatively similar if I restrict the sample of competitor shops to those that at least have an email or website in their YP listing. These could be considered the shops with "online" presence that would compete for customers using email communication. The results are also similar when using the unrestricted sample of YP listings, independently of being identified as a mechanic car shop. Finally, I show that results are robust to alternative definitions of the threshold used to compute the competition measure. Column 1-5 re-run the regression in Table A.10 by using radiuses thresholds that increase by 0.5 km (1, 1.5, 2, 2.5, and 3 km, respectively). Results are qualitatively similar, but quantitatively the results gets dampened with larger distances. I explore this further by defining mutually exclusive rings of these same circumferences and

regressing prices on them. Column 6 shows that the results are driven by the innermost ring.

Another potential venue to explore is how discrimination varies with other geographic characteristics. For instance, are women more discriminated against in areas where more single parents live? Are more ethnically diverse places associated with more equal prices? Or maybe gender price gaps are higher in places where commuting to work takes longer. A first glance at this suggests that geographic characteristics are not strongly correlated with gender price gaps. Figure 6 shows the price differential for each commuting zone, which has no evident pattern to other commonly studied characteristics. The price differential is obtained from a regression of commuting zone fixed effects interacted with a woman dummy, controlling for commuting zone. Figure A.7 shows this differential aggregated at the state level.

Figure 6: Price differential against women by commuting zone



Notes: This map presents coefficient estimates from a regression of price estimates on commuting zone fixed effects interacted with a woman dummy, controlling for commuting zone fixed effects. The sample is restricted to commuting zones with at least two estimates by gender. Larger coefficients (blue) represent higher prices for women relative to men. Bins have equal counts of observations.

5 Conclusion

In this study, I present evidence that women customers receive worse estimates than their male counterparts when no information other than gender is provided to shops. Women in the *baseline* group receive price estimates of roughly 486 dollars - about 9 dollars more than men. This price gap remains once we include controls for other experimental design features or geographic characteristics. Furthermore, this gender price gap is not driven by

differential response rates or changes in the composition detail of the provided quotes. These effects, however, disappear once additional information about the customer is revealed. Once customers additionally signal to be in the *uninformed*, *search*, or *income* groups, they receive price estimates that do not vary significantly by perceived gender. Women benefit when they mention they are searching for other price quotes, since they receive estimates about 11 dollars lower, while men quotes do not vary with this signal. Thus, this signal effectively closes the gap by lowering female price quotes. On the other hand, signaling high income also closes this gender price gap, but by hurting men. Men whose implicit income is high receive estimates close to 490 dollars, similar to women estimates in the baseline group.

Overall, these results suggest that discrimination patterns are explained by a statistical discrimination model. Shops priors are that absent additional information, women have a higher willingness to pay but once additional information is revealed, gender becomes less relevant. This results may be consistent with shops beliefs that women have higher search costs than men, thereby helping women get more competitive prices once they signal the opposite, and hurting men by getting higher prices once they signal higher search costs.

This paper also examines the extent to which discrimination patterns vary with shop and market characteristics. I provide further evidence suggesting that discrimination arises through a signal extraction problem, rather than simple distaste towards women. I find that results are not affected by employee's gender, which indicates that employee's preferences for a specific customer are not driving the results. Consistent with previous literature on discrimination in car sales, I find that dealerships and franchises discriminate against women, and interestingly, this bias is greater than that of independent shops. Importantly, the disadvantages faced by women are reduced when markets are more competitive.

As with all such correspondence studies, there is an important consideration to these findings. Ultimately, the value we are most interested in is the final price paid once a repair

is done, and I observe an intermediate outcome, the price estimate. Car repair shops have both an incentive to provide low price estimates to attract a new customer, but also to provide an estimate close to the final price to avoid surprising customers with higher quotes at the risk of losing them. This paper setting makes it difficult to provide lower estimates than final, as they are written and not given over a phone call. Furthermore, a discrepancy between estimates and final prices would be problematic if we thought that shops increased prices by more to men than to women, vanishing - or reversing - the discrimination price gaps against women. On the other hand, discrimination against women would increase if we thought that shops increased final prices by more to women. One may speculate that the results are a lower bound on the extent of price discrimination with respect to a setting in which only in-person visits occur. The internet facilitates information search and has been found to be particularly beneficial for individuals whose characteristics disadvantage them in negotiating (Morton, Zettelmeyer and Silva-Risso, 2003).

In addition, there are important considerations about the generality of this paper findings. The representativeness of this paper results should be interpreted with caution. While observed estimates come from shops throughout the US, they represent less than 10 percent of all repair shops in the US. Only 40 percent of the YP listings have an available email address, and close to one out of five emails sent replied with an estimate. This motivates further comparisons between shops, and an in-depth exploration of the relation between discrimination patterns and geographic characteristics.

Overall, this paper shows that gender-based price discrimination still exists despite increased progress toward gender equity and equality. Such discrimination patterns seem to arise through shops' prior beliefs about men and women willingness to pay rather than a distaste for serving women customers. Gender becomes less relevant once additional information on search costs or household income is observed. Furthermore, revealing that customers are searching for other price quotes helps them get more competitive prices and

closes the gender-based premium, and the gender premium also decreases with increased competition. In this setting, competition may be a powerful tool to reduce gender-based discrimination.

References

- Abadie, Alberto, Susan Athey, Guido W Imbens, and Jeffrey Wooldridge,** “When Should You Adjust Standard Errors for Clustering ?,” *National Bureau of Economic Research*, 2017, No w24003.
- Ahmed, Ali M and Mats Hammarstedt,** “Discrimination in the rental housing market : A field experiment on the Internet,” *Journal of Urban Economics*, 2008, 64 (2), 362–372.
- Aigner, Dennis J and Glen G Cain,** “Statistical theories of discrimination in labor markets,” *ILR Review*, 1977, 30 (2), 175–187.
- Anderson, Simon P and André De Palma,** “Price dispersion and consumer reservation prices,” *Journal of economics & management strategy*, 2005, 14 (1), 61–91.
- Antonovics, Kate L and Brian G Knight,** “A New Look at Racial Profiling : Evidence from the Boston Police Department,” *The Review of Economics and Statistics*, 2009, 91 (1), 163–177.
- Anwar, Shamena and Hanming Fang,** “An alternative test of racial prejudice in motor vehicle searches : Theory and evidence,” *American Economic Review*, 2006, 96 (1), 127–151.
- Arrow, Kenneth,** “The Theory of Discrimination,” *Discrimination in labor markets*, 1973, pp. 3–33.
- Ayres, Ian and Peter Siegelman,** “Race and gender discrimination in bargaining for a new car,” *The American Economic Review*, 1995, pp. 304–321.

- Barron, John M, Beck A Taylor, and John R Umbeck**, “Number of sellers , average prices , and price dispersion,” *International Journal of Industrial Organization*, 2004, *22*, 1041–1066.
- Becker, Gary S**, “The Economics of Discrimination,” *University of Chicago Press*, 1957.
- Bertrand, Marianne and Esther Duflo**, “Field Experiments on Discrimination,” in “Handbook of Economic Field Experiments,” Vol. 1, North-Holland, 2017, pp. 309–393.
- **and Sendhil Mullainathan**, “Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination,” *The American Economic Review*, 2004, *94* (1996), 991–1013.
- Bessendorf, Anna**, “From cradle to cane : The cost of being a female consumer. A study of gender pricing in New York City,” Technical Report 2015.
- Borenstein, Severin and Nancy Rose**, “Competition and price dispersion in the US airline industry,” *Journal of Political Economy*, 1994, *102* (4), 653–683.
- Bosch, Mariano, M Angeles Carnero, and Lidia Farré**, “Information and discrimination in the rental housing market : Evidence from a field experiment,” *Regional Science and Urban Economics*, 2010, *40* (1), 11–19.
- Busse, Meghan R, Ayelet Israeli, and Florian Zettelmeyer**, “Repairing the damage : The effect of price knowledge and gender on auto repair price quotes,” *Journal of Marketing Research*, 2017, *54* (1), 75–95.
- Castillo, Marco, Ragan Petrie, Maximo Torero, and Lise Vesterlund**, “Gender differences in bargaining outcomes : A field experiment on discrimination,” *Journal of Public Economics*, 2013, *99*, 35–48.

- Christos, Lauren**, *ABC of women workers' rights and gender equality*, geneva, in ed. 2009.
- Darby, Michael R. and Edi Karni**, "Free competition and the optimal amount of fraud," *The Journal of law and economics*, 1973, 16 (1), 67–88.
- Doleac, Jennifer and Luke Stein**, "The visible hand: Race and online market outcomes," *The Economic Journal*, 2013, 123 (572), F469–F492.
- Donohue, John J and Steven D Levitt**, "The impact of race on policing and arrests," *The Journal of Law and Economics*, 2001, 44 (2), 367–394.
- Dulleck, Uwe and Rudolf Kerschbamer**, "On doctors, mechanics , and computer specialists : The economics of credence goods," *Journal of Economic literature*, 2006, 44 (1), 5–42.
- , – , and **Matthias Sutter**, "The economics of credence goods : An experiment on the role of liability , verifiability , reputation , and competition," *American Economic Review*, 2011, 101 (2), 526–555.
- Edelman, Benjamin, Michael Luca, and Dan Svirsky**, "Racial discrimination in the sharing economy: Evidence from a field experiment," *American Economic Journal: Applied Economics*, 2017, 9 (2), 1–22.
- Ewens, Michael, Bryan Tomlin, and Liang Choon Wang**, "Statistical discrimination or prejudice? A large sample field experiment," *Review of Economics and Statistics*, 2014, 96 (March), 119–134.
- Ferrel, O.C, D Kapelianis, L Ferrel, and L Rowland**, "Expectations and attitudes toward gender-based price discrimination," *Journal of Business Ethics*, 2016, pp. 1–18.

- Fershtman, Chaim and Uri Gneezy**, “Discrimination in a segmented society: An experimental approach,” *The Quarterly Journal of Economics*, 2001, 116 (1), 351–377.
- Fox, John**, *Applied regression analysis, linear models, and related methods*, sage publi ed. 1997.
- Gerardi, Kristopher S. and Adam Hale Shapiro**, “Does competition reduce price dispersion ? New Evidence from the airline industry,” *Journal of Political Economy*, 2009, 117 (1), 1–37.
- Gneezy, Uri, John List, and Michael K Price**, “Toward an understanding of why people discriminate: Evidence from a series of natural field experiments,” *National Bureau of Economic Research*, 2012.
- Guryan, Jonathan and Kerwin Kofi Charles**, “Taste-based or statistical discrimination: the economics of discrimination returns to its roots,” *The Economic Journal*, 2013, 123 (572), F417–F432.
- Hanson, Andrew and Zackary Hawley**, “Do landlords discriminate in the rental housing market ? Evidence from an internet field experiment in US cities,” *Journal of Urban Economics*, 2011, 70 (2-3), 99–114.
- Heckman, James J**, “Detecting Discrimination,” *Journal of Economic Perspectives*, 1998, 12 (2), 101–116.
- Holmes, Thomas J**, “The effects of third-degree price discrimination in oligopoly,” *The American Economic Review*, 1989, 79 (1), 244–250.
- Jowell, Roger and Patricia Prescott-Clarke**, “Racial discrimination and white-collar workers in Britain,” *Race*, 1970, 11 (4), 397–417.

- Kerschbamer, Rudolf and Matthias Sutter**, “The economics of credence goods—a survey of recent lab and field experiments,” *CEsifo Economic Studies*, 2017, 63 (1), 1–23.
- Lee, David S**, “Training , wages , and sample selection : Estimating sharp bounds on treatment effects,” *The Review of Economic Studies*, 2009, 76 (3), 1071–1102.
- Levitt, Steven D. and Stephen J. Dubner**, *Freakonomics. Revised and Extended Edition* 2006.
- List, John**, “The nature and extent of discrimination in the marketplace: Evidence from the field,” *The Quarterly Journal of Economics*, 2004, 119 (1), 49–89.
- Morton, Fiona Scott, Florian Zettelmeyer, and Jorge Silva-Risso**, “Consumer information and discrimination : Does the internet affect the pricing of new cars to women and minorities ?,” *Quantitative marketing and Economics*, 2003, 1 (1), 65–92.
- Nunley, John M, Mark F Owens, and R Stephen Howard**, “The effects of information and competition on racial discrimination: Evidence from a field experiment,” *Journal of Economic Behavior & Organization*, 2011, 80 (3), 670–679.
- Phelps, Edmund S**, “The Statistical Theory of Racism and Sexism,” *The American Economic Review*, 1972, pp. 659–661.
- Pope, Devin G and Justin R Sydnor**, “What ’s in a picture ? Evidence of discrimination from prosper.com,” *Journal of Human resources*, 2011, 46 (1), 53–92.
- Price, Joseph and Justin Wolfers**, “Racial discrimination among NBA referees,” *The Quarterly journal of economics*, 2010, 124 (4), 1859–1887.
- Ravina, Enrichetta**, “Beauty , personal characteristics and trust in credit markets,” 2007.

Riach, Peter A and Judith Rich, “Testing for racial discrimination in the labour market labour market,” *Cambridge Journal of Economics*, 1991, 15 (3), 239–256.

– **and** –, “Field experiments of discrimination in the market place,” *The economic journal*, 2002, 112 (1974), 480–518.

Schneider, Henry S., “Agency Problems and Reputation in Expert Services: Evidence from Auto Repair,” *The Journal of Industrial Economics*, 2012, 60 (3), 406–433.

Shepard, Andrea, “Price discrimination and retail configuration,” *Journal of Political Economy*, 1991, 99 (1), 30–53.

Siegelman, Peter and James J Heckman, “Clear and convincing evidence: Measurement of discrimination in America,” in “The Urban Institute Audit Studies : Their Methods and Findings,” Vol. 187, Urban Institute Press., 1993, p. 258.

Stavins, Joanna, “Price discrimination in the airline market : The effect of market concentration,” *Review of Economics and Statistics*, 2001, 83 (1), 200–202.

Stole, Lars A, “Price discrimination and competition,” in “Handbook of industrial organization,” Vol. 3 2007, pp. 2221–2299.

6 Appendix: Experimental design details

6.A Randomization: Components and implementation steps

This project used a simple randomization method to assign an email account (perceived gender variation) and the email content (customer type variation) to each shop. The randomized component is visualized below.

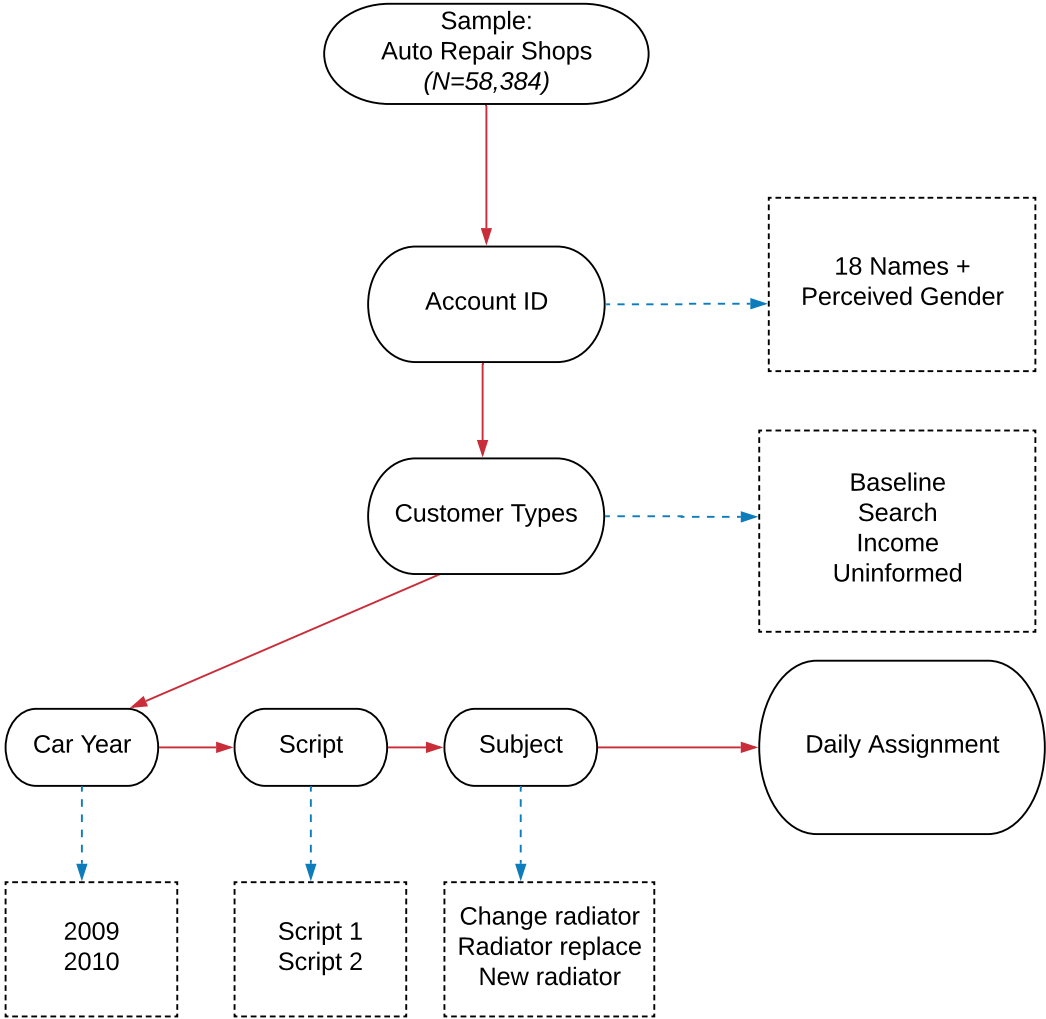
First, I randomized gender and customer types. I randomly assigned one of the eighteen account IDs to each shop, matching the same number of emails to each account ID, and therefore, assigning half of the sample to women. Then, I randomly assigned each of the four customer type treatments within each gender. Thus, 18. percent of shops will be contacted by woman customers using standard scripts, 16 percent will be contacted by low-search woman customers, 8 percent by high-educated women and 8 percent by uninformed women. Equivalent shares are assigned to each man-type combination.

Then, I randomized additional components of each email; the car year, script design and email subject. Each of these variations is included to avoid sending a second identical email to each shop. The car years are 2009 and 2010, both belong to the same car generation. Within each customer type, there are two scrips; script 1 and script 2. The content is the same, but they are reordered and slightly differently worded. Across types, each script number mirrors one another. That is, the high-education type script 1 is the same as the standard script 1 except for the signature, which adds the Ph.D. title. All the subjects transmit the need for a radiator replacement, these are “change radiator,” “radiator replace” and “new radiator”.

Finally, the order and date in which each shop is contacted is also randomly assigned. The assignment of shops to treatments was done in weekly waves while sending a similar number of emails by day. Along the process, some shops were identified as not repairing

Honda radiators, thereby dropped from the sample. For instance, collision shops were dropped as they do not perform repairs unless caused by an accident.

Figure A.1: Randomization Steps

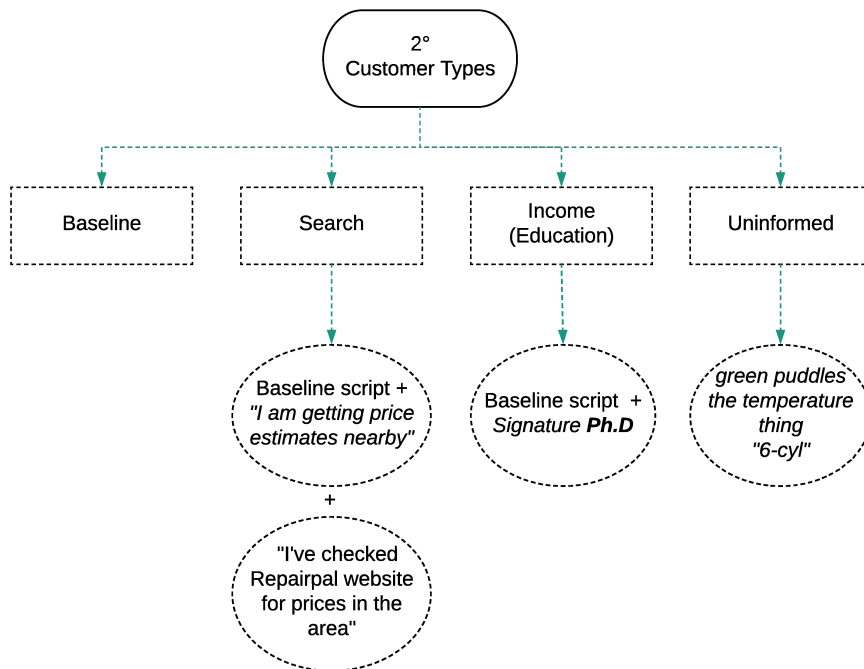


Note: This figure shows each of the randomized components used to contact each facility. The red arrows indicate new reshuffles in the sample. The blue dotted arrows indicate the components that were randomly assigned.

6.B Customer gender and type scripts

This section provides details of customer type manipulation trait. Figure A.2 summarizes the main changes added to the baseline script. Complete scripted details used for each customer type are also presented. In each script, the terms $\{\text{FILL_CITY}\}$, $\{\text{FILL_YEAR}\}$ and $\{\text{FILL_NAME}\}$ are included. At the time of sending the script, those terms are automatically replaced with the corresponding shop's city, car year and user account name. As can be observed, script 1 is identical across type's except for the information revealing content. Likewise, script 2 is identical across types. Script 1 and Script 2 within types aim to convene the same information and structure, in a slightly reorganized and re-worded mail.

Figure A.2: Customer type variations to Baseline type script



Note: This figure shows each of the main variations to the baseline script used to signal a low-search cost, a high-income, and a uninformed type customer.

Standard script 1

“Hi,

Do you replace radiators? Would it be possible to know what is the total estimate to replace mine and when I could take it? I live in `#{FILL_CITY}`, near your service center.

I have a 6 cylinder `#{FILL_YEAR}` automatic, sedan Honda Accord. It has been leaking and overheating. I already had some leaks fixed, now I'd like to replace the radiator.

Thank you!

Best,

`#{FILL_NAME}`”

Standard script 2

“Hi,

I need to replace my car's radiator. It had cracks fixed before, but now it is overheating and leaking coolant again. I'd like to directly get a new radiator. The car is a `#{FILL_YEAR}` 6-cylinder, automatic, sedan Honda Accord.

Would you please tell me what is the total replacement cost? When can I stop by? I live close to your shop, here in `#{FILL_CITY}`.

Thanks!

Best,

`#{FILL_NAME}`”

Low-Search script 1 and 1b

“Hi,

Do you replace radiators? Would it be possible to know what is the total estimate to replace mine and when I could take it? I live in \${FILL_CITY}, near your service center.

I have a 6 cylinder \${FILL_YEAR} automatic, sedan Honda Accord. It has been leaking and overheating. I already had some leaks fixed, now I'd like to replace the radiator.

“I am getting price estimates nearby”

or *“I've checked Repairpal website for price references in the area. I will try to get a price estimate nearby.”*

Thank you!

Best,

\${FILL_NAME}”

Low-search script 2

“Hi,

I need to replace my car's radiator. It had cracks fixed before, but now it is overheating and leaking coolant again. I'd like to directly get a new radiator. The car is a \${FILL_YEAR} 6-cylinder, automatic, sedan Honda Accord.

Would you please tell me what is the total replacement cost? When can I stop by? I live close to your shop, here in \${FILL_CITY}.

“I am getting price estimates nearby”

or “I’ve checked Repairpal website for price references in the area. I will try to get a price estimate nearby.”

Thanks!

Best,

#{FILL_NAME}”

High-Education script 1

“Hi,

Do you replace radiators? Would it be possible to know what is the total estimate to replace mine and when I could take it? I live in #{FILL_CITY}, near your service center.

I have a 6 cylinder #{FILL_YEAR} automatic, sedan Honda Accord. It has been leaking and overheating. I already had some leaks fixed, now I’d like to replace the radiator.

Thank you!

Best,

#{FILL_NAME}, *Ph.D.*”

Standard script 2

“Hi,

I need to replace my car’s radiator. It had cracks fixed before, but now it is overheating and leaking coolant again. I’d like to directly get a new radiator.

The car is a #{FILL_YEAR} 6-cylinder, automatic, sedan Honda Accord.

Would you please tell me what is the total replacement cost? When can I stop by? I live close to your shop, here in \${FILL_CITY}.

Thanks!

Best,

\${FILL_NAME}, *Ph.D.*”

Uninformed script 1

“Hi,

Do you replace radiators? Would you let me know what is the total cost to replace mine and when I could take my car? I live in \${FILL_CITY}, near your service center.

My car is leaving *these green liquid puddles* and the *temperature thing* is going up a lot. I had the radiator fixed before but I think it is time to replace it. The car is a Honda Accord. Not sure if it helps, it is an automatic, 4-door, "6-cyl." car from \${FILL_YEAR}.

Thank you!

Best,

\${FILL_NAME}”

Uninformed script 2

“Hi,

My car is leaving *these green liquid puddles* and the *temperature thing* is rising too much. I had the radiator fixed before, now I think I should get a new one

placed. The car is a Honda Accord. In case it is useful, it is an automatic, 4-door and "6-cyl." car from \${FILL_YEAR}.

Would you let me know what is the total cost estimate to replace my radiator?

What would be a good time to take it? I live close to your shop, here in \${FILL_CITY}.

Thanks!

Best,

\${FILL_NAME}"

6.C Reply rules and template scripts

The following guidelines are used to reply to shops, regardless of the email account and customer type used to contact each shop.

The main two rules are to reply only once to shop's follow-up emails, and second, to use an available reply template whenever possible. These templates are listed below and anticipate shop's common requests, such as whether the customer has a phone or would be able to take the car for an inspection. Each request reply has two alternative templates, which are rotated on a weekly basis. If further car information is asked, the first template "car information" should be used, adding either the information asked. The car is 6-cylinder, Honda Accord, sub-model EX, 4 door, automatic car.

Another rule used was to try to reply to every email. So, I replied "thanks" in each of the following cases: when a shop mentions it is their policy not to give estimates without seeing the car, when they suggest contacting other person in the shop or, if they reply the following week of the experiment. This last rule is designed to avoid sending the script corresponding to the new week to a shop contacted in the previous week. Finally, I replied "thanks" when a shop provides a quote. If the shop mentions that the part should be ordered or offers to

set an appointment, I replied “thanks, I will let you know” in order to avoid any confusion. There are some cases in which these shop emails remained unanswered, this should in no way affect this paper results, as no quote is given.

Finally, when an employee mentioned being out of office, I left this email unanswered, as I only wanted to recruit estimated from employees while they are working.

As will be noted in the reply scripts, they have the text `#{name}` and `#{sender_name}`. Each of these is manually replaced. `#{name}` refers to the employee’s name, and is highlighted in yellow so I minimized the chance of skipping this text replacement step. `#{sender_name}` is replaced by the account id being used. The email setting from each account had a different background theme, the name of the user in the gchat window, and I distributed accounts in each browser - allowing me to also familiarize which account is being used based on the explorer (chrome, explorer and Firefox). Additionally, every account allowed the “undo” option after sending an email for 20 seconds.

Reply scripts 1

Additional car information

Hi `#{name}`,

It’s a 3.5-liter car. Just in case, it also has AC.

Thank you!

`#{sender_name}`

VIN

Hi `#{name}`,

I am sorry, I won’t be close to my car until tomorrow. If it helps towards an approximate estimate, it is a 3.5-liter car, sedan, and with AC.

Thank you!

Best,

\${sender_name}

VIN + CALL

Dear \${name},

I am sorry, I am not close to my car now. If it helps towards an approximate estimate, it is a 3.5-liter car, sedan, and with AC.

If it's OK, I'd prefer to talk by over email.

Thank you!

Best regards,

\${sender_name}

Call

Hi \${name},

If it's OK, I'd prefer to talk over email. In case it is helpful, it is a 3.5-liter car, with air conditioner. In case it is the radiator, I would like to have a ballpark figure, if possible.

Thank you!

Best regards,

\${sender_name}

Inspect

Hi \${name},

I realize the final estimate may change upon inspection, but I was hoping to have a sense of total repair costs before bringing it in (in case it is the radiator).

If it is, I would not want to fix it but to have a new one installed.

Just in case, the engine is 3.5 liters and it is a car with air conditioner. Thank you, I appreciate all your help!

Best,

`\${sender_name}`

Call + inspect

Hi `\${name}`,

If it is OK, I'd like to talk over email before taking my car. I realize the final estimate may change upon inspection, but I was hoping to have a sense of total repair costs before bringing it in (in case it is the radiator). If it is, I would not want to fix it but to have a new one installed.

Just in case, the engine is 3.5 liters and it is a car with air conditioner. Thank you, I appreciate all your help!

Best,

`\${sender_name}`

Info on service, no price

Dear `\${name}`,

I realize the final estimate may change upon inspection, but I was hoping to have a sense of total repair costs before bringing it in (in case it is the radiator).

If it is, I would not want to fix it but to have a new one installed. Just in case, the engine is 3.5 liters and it is a car with air conditioner.

Thank you, I appreciate all your help!

Best,

`{sender_name}`

Are you bringing the radiator?

Oh no, I would just bring the car in. In case it helps, the engine is 3.5 liters and it is a car with air conditioner.

Thank you!

`{sender_name}`

Previous repairs

Hi `{name}`,

I had these same car issues a while ago. I had it repaired, not replaced. I think some cracks were closed, but I cannot say for sure. That is why, if the radiator is having problems I will just want a new one. That is why I am asking for an approximate estimate for this, to have an idea beforehand. In case it helps, the car has air conditioner and it's 3.5 liters.

Thank you!

Best,

`{sender_name}`

Still need repair?

Yes, thank you. In case it helps, the car has air conditioner and it's 3.5 liters. Just want to have a sense of approximate costs if it is the radiator that has issues (would like a new one, not repairing mine).

Best regards,

\${sender_name}

Others – Radiator type:

I have no preference between radiator brands. I would just want something that I won't have to replace again soon.

Thank you!

\${sender_name}

Reply Scripts 2

Additional car information

Hi \${name},

Thanks for your reply! It is a 3.5-liter car, with AC.

Best,

\${sender_name}

VIN

Dear \${name},

Thanks for your reply. Unfortunately, I cannot check the VIN until tomorrow night. Is there any other information I could give you now instead of waiting for the VIN number?

Just in case, the car has AC and it is 3.5 liters.

Best,

\${sender_name}

VIN + CALL

Hi \${name},

Thanks for your reply. If you would not mind, I would rather correspond by email. Unfortunately, I cannot check the VIN until tomorrow night. Is there any other information I could give you now instead of waiting for the VIN number?

Just in case, the car has AC and it is 3.5 liters.

Best,

\${sender_name}

Call

Hi \${name},

Thanks for your reply. If you would not mind, I would rather correspond by email. I'd appreciate having an approximate estimate if it is indeed the radiator.

Just in case, the car has AC, 3.5 liters.

Best,

\${sender_name}

Inspect

Dear \${name},

If possible, could you let me know approximately how much it would cost? I just want to have a ballpark figure if it is indeed the radiator. Of course, I understand that it is an estimate before an inspection. But, if it is the radiator that has issues, I want to get a new one.

If it helps, the car has AC, 3.5 liters.

Thanks for your reply!

Best,

\${sender_name}

Call + inspect

Dear \${name},

Thanks for your reply. If you would not mind, I would rather correspond by email first. If possible, could you let me know approximately how much it would cost? I just want to have a ballpark figure if it is indeed the radiator. Of course, I understand that it is an estimate before an inspection. But, if it is the radiator that has issues, I want to get a new one.

If it helps, the car has AC, 3.5 liters.

Best,

\${sender_name}

Info on service, no price

Hi \${name},

If possible, could you tell me a ballpark figure in case it is the radiator? I wanted to have an idea in case it is the radiator. I would like to get a new one and not fix the one I have. Of course, I understand that the final estimate may vary after an inspection. Let me know if there is any other information I could give you.

The car has AC, and it's 3.5-liter.

Best,

#{sender_name}

Are you bringing the radiator?

No, I would just take my car. Just in case, it is a 3.5-liter car, with AC.

Best,

#{sender_name}

Previous repairs

Dear #{name},

Thanks for your reply. I am not 100% certain on what specific repair I had on the radiator. I know it had some cracks that were leaking, and they added something to close them (but this may not be entirely accurate). I had this done a while ago but now I have the same issues. So, if it is indeed the radiator that has issues, I'd like to have an idea of how much it would cost to replace it instead of fixing it. I want to get a new one.

Just in case, it is a 3.5-liter car, with AC.

Best,

#{sender_name}

Still need repair?

Thank you for your reply. Yes, I still need to get my car fixed. Let me know if you could give me a ballpark figure in case it is the radiator. I don't want it fixed but I'd like getting a new one.

Just in case, the car has AC, and it is 3.5-liter.

Best,

`#{sender_name}`

Others – Radiator type:

I'd want one that I won't need to fix again in the near future, but I am indifferent between brands.

Thanks!

`#{sender_name}`

6.D Data collection rules

With respect to the data collection, I created Gmail labels each week to group emails into *invalid email*, *not do service*, *do service and give quote*, *do service and not give quote*. These are used to estimate reply rates by email reply content. I automated the weekly download of every email thread into a spreadsheet, with some variables automatically filled in based on the email's text. These were later reviewed and completed manually. Three research assistants, unaware of this project hypotheses, helped with this last task. As much as possible, they reviewed *give quote label* emails, while I reviewed the remaining labels. Nevertheless, all collected information is not subject to an individual's interpretation and is easily verifiable (i.e., a shop either gives a price or not, mentions a discount or not, etc.).

The most relevant pre-specified rules used to translate email replies into a spreadsheet are the following: keep one observation for each contacted shop, the reported time and date of each reply corresponds to the first reply received by each shop. Identify if a reply is sent before obtaining a quote from a shop, record the employee first name used to sign the email, and if missing, the name of the email account. Keep the first observed employee name when more than one employee reply to an email. If available, keep track of the price composition


and which parts are explicitly not included (if any), the parts are labor, parts, radiator, coolant (includes and indicator for coolant flush services), hoses and clamps, thermostat, other (includes shop fees and other offers), taxes, quantity of estimated labor and estimated coolant. Also keep track of the radiator brand offered in the estimate (in particular, if it is a factory Honda part). Whenever more than one estimate is sent by the same shop, but in separated emails, keep the first estimate and add a comment to identify these cases.

6.E Design validation

This section shows further validations of the experimental design implementation. First, I plot the geographic distribution of auto-repair shops assigned to be contacted by each treatment group. Figure A.3 shows the assignment of shops to men and women separately, and further dividing the assignment of shops by treatment group (combining each customer type and gender). As can be observed, both maps present a similar distribution of geographical locations contacted by each group.

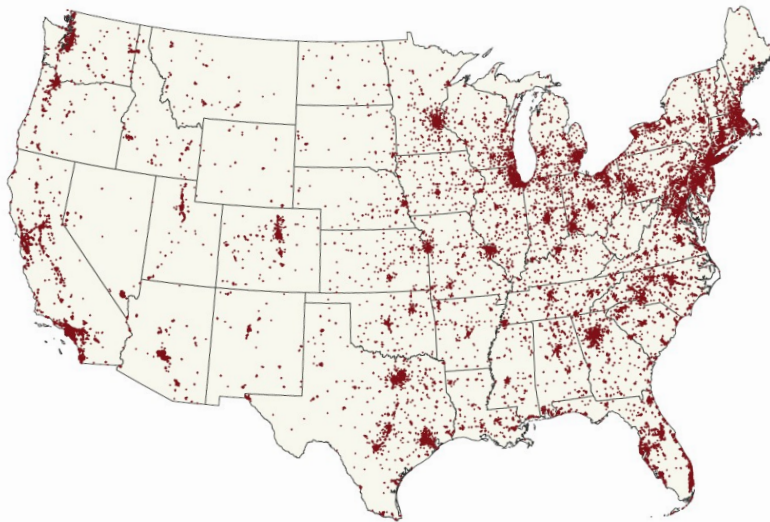
In addition, Figure A.4 shows the distribution of invalid emails by treatment group, which reassuringly shows not significant differential across groups. Table A.5 complements these results showing the reply rates by gender across customer types and within customer type across genders for all invalid email replies, as well as replies of shops not performing Honda radiator replacements, performing the service and restricting to those that perform the service and provide an estimate.

Figure A.3: Geographic distribution of contacted shops



Distribution of Shops assigned to males

A. Men




Distribution of Shops assigned to females

B. Women

Notes: These figures show the location of shops contacted by woman and man fictitious customers

Figure A.4: Distribution of invalid-emails



Notes: This figure shows average invalid emails for women and men by their customer type used to contact each shop: Baseline, uninformed, quote and high-income. The plot shows the 95% confidence intervals, obtained with robust standard errors. The sample mean invalid emails is marked by a grey line within the plot.

7 Theoretical Framework

This section provides a simple theoretical framework to help guide the empirical exercise in Section 4. I posit two dimensions of consumer heterogeneity to capture the main sources of variation in the empirical section.

On the one hand, consumers - indexed by i - differ in their willingness to pay for the car repair service θ_i . This parameter θ_i is assumed to be distributed according to G^k where $k \in K$ is an observable characteristic. In the model, stores are allowed to discriminate based on observable k , which they use to infer θ_i . In the experiment there are 8 observable

characteristics (i.e., $\#K = 8$), which differ by gender and customer type characteristics. While customers within group are heterogeneous, one may expect that some agents, such as those with high-income or high-search costs, to “typically” have a higher willingness to pay. In the model, I say that agent k is “typically” more willing to pay than agent k' if G^k satisfies the monotone likelihood ratio property (MLRP) with respect to $G^{k'}$. This implies that the sensitivity of demand of agents k relative to k' is relatively larger when the price is relatively high (i.e., for $p > p'$, $\frac{g^k(p)}{g^{k'}(p)} > \frac{g^k(p')}{g^{k'}(p')}$). As I will show later, this implies that if $G^{females}$ satisfies the MLRP with respect to G^{males} , then there will be a positive gender gap.¹⁴ Note that, in principle, one can enrich the set of observable characteristics by interacting these with geographic characteristics. For example, if women living in area A earn much less than men in area A , then one would expect the G^k s to be closer (or even reversed) to one another, leading to a lower (or opposite) price gap.

On the other hand, I assume that consumers differ in their ability to substitute the current shop with another potential shop. Each consumer i is randomly assigned to an initial shop with equal probability and faces a switching cost λ_i of shopping in an alternative store. This parameter λ_i is assumed to be distributed according to F^l where $l \in L$ is an observable characteristic. Again, stores in the model are allowed to discriminate based on the observable (l), which they use to infer λ_i . The counterpart of λ_i in the experiment is the extent of competition in any given market, i.e., the ease with which customers may substitute stores. For example, if there are many stores within a 1 kilometer radius, then one may expect this λ_i to be typically lower. In this sense, the correct interpretation is that the consumer lives near this store (or happens to “passively” shop in it as in Anderson and De Palma, 2005), and finds it costly to visit another store.¹⁵ Importantly, in the model λ_i is not an informational cost. Indeed, in the equilibrium I describe below there is a pure

¹⁴First-order stochastic dominance is implied but not enough to guarantee this result.

¹⁵The model presented here is a two-store model. I conjecture that having many stores will behave similarly to lowering switching costs in the context of the two-store model.

strategies equilibrium so agents fully anticipate the distribution of prices at other stores in the Nash equilibrium.¹⁶ However, the credible threat of switching stores is key for the properties of the Nash equilibrium.

Consumer i utility function is given by

$$u_i = \begin{cases} \theta_i - p_s & \text{if } i \text{ buys in the closest shop} \\ \theta_i - p_{s'} - \lambda_i & \text{if } i \text{ buys in a different shop} \end{cases}$$

There are 2 stores. Each store is assigned 1/2 “initial” consumers of each type $\{k, l\}$. Stores may set price differently across types. For simplicity, I assume stores face a constant marginal cost, making the problem with respect to each type $\{k, l\}$ separate from one another. Both shops have the same marginal cost, which I set without loss of generality to 0.

Consider the problem of shop j selling to a customer of type $\{k, l\}$. The shop’s profits are given by

$$\pi_j^{kl} = \underbrace{p_j^{kl} \left(1 - G^k(p_j^{kl})\right)}_{\text{initial customers}} - \underbrace{1_{p_j^{kl} > p_{-j}^{kl}} p_j^{kl} \left(F(p_j^{kl} - p_{-j}^{kl})\right)}_{\text{customers lost if price is high}} + \underbrace{1_{p_j^{kl} < p_{-j}^{kl}} p_j^{kl} \int_0^{p_{-j}^{kl} - p_j^{kl}} f^l(\lambda) (1 - G^k(p_j^{kl} + \lambda)) d\lambda}_{\text{customers gained if price is low}}$$

¹⁶The interaction with informational frictions is an interesting avenue for research, but it is outside the scope of this paper. I decided to focus on the physical cost for two reasons. First of all, given the rise of the internet, the cost of sending emails to figure out quotes is likely to be small relative to the cost of actually going to the store. Second, as I show below, this “physical” cost is very well-behaved in the sense that the Nash equilibrium features a number of desirable properties, while the predictions regarding informational frictions and search costs depend significantly on details of the model and may give rise to a number of counter-intuitive predictions, such as the Diamond paradox (for a survey see (Stole, 2007)).

When $p_j^{kl} \rightarrow p_{-j}^{kl}$, one can show the first-order condition converges to^{17,18}

$$\frac{1 - G^k(p_j^{kl})}{g^k(p_j^{kl})} (1 - p_j^{kl} f^l(0)) - p_j^{kl} = 0. \quad (1)$$

If consumers were immobile, i.e., if $f^l(0) \rightarrow 0$, then p_j^{kl} is the monopoly price. Each p_j^{kl} varies with customer observable characteristics, as stores engage in statistical discrimination to maximize profits.¹⁹ Proposition 7 shows that this is still the case when switching costs are nontrivial. By contrast, if $f^l(0) \rightarrow \infty$, $p_j^{kl} \rightarrow 0$. This is the perfect competition outcome: Stores cannot discriminate based on observable characteristics since they would lose all their customers to the other shop. Proposition 7 shows that this intuition is more general: the price increases monotonically with switching costs. One may then wonder if statistical discrimination and switching costs are “complements”, i.e., whether statistical discrimination is stronger when switching costs are larger. While the previous results imply this is obviously true in the $f^l(0) \rightarrow 0$ and $f^l(0) \rightarrow \infty$ cases, a stronger result can be shown when G is uniform. In a symmetric Nash equilibrium, the following is true:

(a) Suppose $G^k(\cdot)$ satisfies the monotone likelihood ratio property with respect to $G^{k'}(\cdot)$.

Then, everything else equal, agents of type k are charged a higher price, i.e., $p^{kl} > p^{k'l}$.

(b) Suppose agents of type l have a higher switching cost, i.e., that $f^l(0) < f^{l'}(0)$. Then, everything else equal, agents of type l are charged a higher price, i.e., $p^{kl} > p^{k'l}$.

(c) Suppose $G^k(x) = \frac{\bar{\theta}^k - x}{\bar{\theta}^k - \underline{\theta}^k}$, i.e., G^k is uniform. Then, agents with high willingness to pay are discriminated against more strongly when switching costs increase, i.e., if $\bar{\theta}^k > \bar{\theta}^{k'}$ and $f^l(0) < f^{l'}(0)$, then $p^{kl} - p^{k'l} > p^{k'l} - p^{k'l'}$.

¹⁷Despite the indicator functions, the objective is smooth at p_{-j} since you only lose/gain an infinitesimal number of consumers from an infinitesimal change in price. This smoothness precludes many of the standard problems of inexistence of pure strategy equilibria in this kind of models.


¹⁸One naturally also needs to impose a condition on f to show this (local) solution to the first-order conditions is also a global maximum. I assume this is the case in the exposition.

¹⁹There is no taste-based discrimination in the model. One could potentially incorporate this as a difference in the marginal cost of serving different kinds of customers.

Proof. (a) and (b) follows from a straightforward application of the implicit function theorem on 1. (c) follows from computing the cross-derivative using the results in (a) or (b). □

Figure A.5 illustrates the results. I assume $f = (\bar{\lambda}^l)^{-1}$ and vary $\bar{\lambda}^l$ in the x axis. The solid line shows the equilibrium price where agents have a high willingness to pay ($\bar{\theta}^k = 1$) while the dashed line corresponds to a case with a lower willingness to pay ($\bar{\theta}^k = 1/2$). As switching costs increase, stores compete less with one another and are able to extract more surplus from agents with a high willingness to pay.

Figure A.5: Price Discrimination and Competition



8 Replies and price estimates descriptive statistics

Figure A.6: Distribution of shops providing estimates

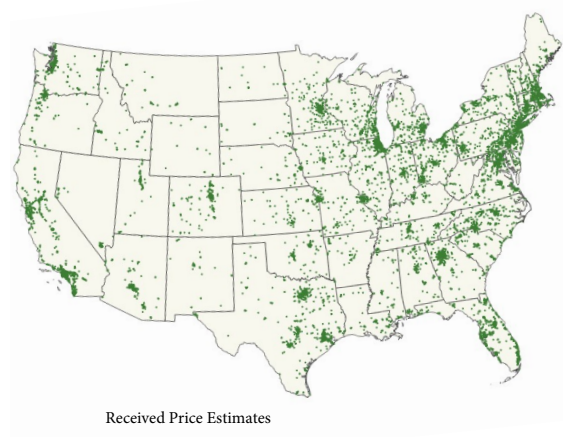



Figure A.7: Price differential against women by state



Notes: This map presents coefficient estimates from a regression of price estimates on state fixed effects interacted with a woman dummy, controlling for state fixed effects. Larger coefficients (blue) represent higher prices for women relative to men. Bins have equal counts of observations. In this sample, shops from 32 states quote higher estimates to women.

9 Price differences by treatment group: Robustness

This section re-does the main analysis in a simple regression framework. The OLS estimation is:

$$p_{jk} = \alpha_0 + \beta_1 Women_k + \beta_2 Type_k + \beta_3 (Type_k \times Women_k) + \omega_s + \theta X_{jk} + \varepsilon_{jk} \quad (2)$$


where p_{jk} is the estimate that shop j provides to k customer (gender and type customer combination), $Women_k$ is an indicator variable equal to 1 when a woman customer contacts shop j , $Type_k$ is a vector with indicator variables for each customer type. The omitted category is always the baseline type. And, $Type_k \times Women_k$ is the vector of customer type interactions with the woman indicator. ω_s are state fixed effects, X_{jk} correspond to controls of experimental design items that vary including the timing in which the emails were sent. That is, day of the week and week of the experiment in which shop j is contacted, and script number, car year and email subject covariates. Additional specifications include a control for non-independent shops and commuting zone fixed effects. Finally, ε_{jk} is the unobserved component, which has mean zero and is uncorrelated with our treatment groups.²⁰

Table A.6 shows these results. As pre-specified in the analysis plan, column 1 controls for state fixed effects, additional varying email items (car year, subject controls and script version used), and day of week and week fixed effects indicating when each shop is contacted. Column 2 adds a non-independent shop type control (i.e., franchises and dealerships) since their price levels are usually higher, and column 3 adds Commuting Zone (CZ) fixed effects and an additional indicator dummy for the shops that were not matched to any commuting zone. Figure A.8 complements these last results, directly plots adjusted point estimates for each treatment group with respect to men in the baseline group. Over-

²⁰I do not cluster standard errors given this field experiment's sampling design and experimental design Abadie et al. (2017)

all, the regression-adjusted results are consistent with raw price differences. Indeed, the gender-based premium in the baseline group remains significant across specifications and increases to an average of 9.6 dollars. The price decrease for women in the quote search group is marginally significant, while the price increase for men in the income group is statistically significant and varies between an average of 11.2 to 12.2 dollars with respect to men in the baseline group. Reassuringly, results are robust to the inclusion of controls, and none of the additional experimental design features is relevant to explain variation in prices.

Figure A.8: Effects of treatment groups on price



Notes: This figure shows average price gaps for each gender and customer type relative to male-baseline customers. The coefficient estimates are obtained from a specification with indicator variables for each customer type interacted with an indicator for female, an indicator for man customers, respectively. The omitted category is male-baseline customer, marked by a triangle with no confidence intervals in the plot. Additional controls include state and commuting zone fixed effects, with a dummy variable for observations not linked to a commuting zone, car year, subject and script number, day of the week and week fixed effects, and a non-independent shop indicator (i.e., franchise, dealership). The scatter shows each average price differential with their 95% confidence intervals. A null price gap estimate ($y\text{-axis}=0$) is marked by a grey line within the plot.

Table A.6: Impact of gender and customer type on prices: Robustness

	Dependent variable: Price estimates		
	(1)	(2)	(3)
Women	10.026** (4.236)	9.506** (4.115)	9.689** (4.176)
Search	-0.816 (4.555)	-0.175 (4.440)	0.389 (4.531)
Search \times Women	-9.821 (6.396)	-11.317* (6.188)	-11.101* (6.291)
Income	12.772** (5.298)	12.248** (5.173)	13.999*** (5.351)
Income \times Women	-13.792* (7.560)	-12.719* (7.349)	-13.856* (7.586)
Uninformed	-1.568 (5.450)	-3.169 (5.267)	-1.988 (5.367)
Uninformed \times Women	-6.012 (7.601)	-5.000 (7.321)	-5.597 (7.500)
Subject: Change radiator	-3.044 (3.206)	-2.481 (3.100)	-1.756 (3.175)
Subject: Radiator replace	-2.284 (3.302)	-0.116 (3.205)	0.245 (3.264)
Car year 2010	1.117 (2.590)	1.402 (2.509)	1.555 (2.551)
Script 1	-3.460 (2.609)	-2.477 (2.524)	-2.519 (2.570)
Controls	State FE DOW fe, week FE	+ Shop type	+ CZ FE
Observations	10,323	10,323	10,323

Notes: This table reports coefficients from a regression of total price estimates on an indicator for woman customer, and indicator for each customer type, and their interaction with woman customer, with baseline customer type as the omitted categories. Column (1) includes state fixed effects, other email varying items: car year, subject and script number controls, and day of week and week fixed effects - indicating when the email was sent to a shop. Column (2) includes a non-independent shop control (non-independent shops: dealership and franchises). Column (3) adds CZ fixed effects, with a dummy for observations not linked to a CZ. Robust standard errors in parenthesis, with significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Selection bias correction


Table A.7: Robustness: Selection in replies

	Give quote (1)	Two-step correction		Non-parametric bounds	
		Price estimate (2)	Inverse MR (3)	Lower bound (4)	Upper bound (5)
Baseline: Women vs men	0.008 (0.005)	8.522* (5.039)	-31.508 (62.763)	4.781 (14.732)	15.415 (9.896)
Search: Women vs men	0.016*** (0.006)	-0.129 (5.985)	25.571 (44.465)	-21.027 (14.153)	14.877 (11.814)
Women: Search vs baseline	-0.012** (0.006)	-12.240** (5.138)	6.421 (55.102)	-21.912* (10.623)	1.557 (22.536)
Men: Search vs baseline	-0.020*** (0.005)	0.738 (5.936)	-17.923 (61.157)	-20.430* (10.055)	24.778* (12.203)
Men: Income vs baseline	0.008 (0.007)	12.111* (6.682)	-36.880 (60.486)	3.726 (16.228)	18.134 (14.312)
Women: Uninformed vs baseline	-0.012* (0.007)	-8.984 (5.892)	13.613 (58.443)	-20.615* (10.716)	9.187 (13.736)

Notes: This table reports the difference in the rate of shops providing quotes (Column 1), the corrected price estimate differential using Heckman’s two-step correction method (Column 2), and the associated inverse Mills ratio coefficient (Column 3), for each comparison group that showed a significant difference in the rate of price estimates provided. The coefficient obtained in Column 2 controls for state fixed effects, email varying items, a dummy for non-independent shops, and day of week and week fixed effects. Standard errors in parenthesis, with significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Alternative quote definitions

Figure A.9: Robustness: Alternative quote definitions



Notes: This figure shows average price gaps differential for baseline customers using alternative quote definitions. The coefficient estimates are obtained from a specification with indicator variables for each customer type interacted with an indicator for female, an indicator for man customers, respectively. The omitted category is male-baseline customer, marked by a triangle with no confidence intervals in the plot. Additional controls include state and commuting zone fixed effects, with a dummy variable for observations not linked to a commuting zone, car year, subject and script number, day of the week and week fixed effects, and a non-independent shop indicator (i.e., franchise, dealership). The scatter shows each average price differential with their 95% confidence intervals. A null price gap estimate (y-axis= 0) is marked by a grey line within the plot.

Table A.8: Robustness: Alternative quote definitions

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Prices	No disc.	Min P	Max P	Match P	Invalid P	Robust
Women	9.689** (4.176)	9.891** (4.186)	8.747** (4.180)	10.635** (4.202)	10.096** (4.202)	9.984** (4.177)	10.745*** (3.979)
Search	0.389 (4.531)	0.402 (4.549)	-0.475 (4.516)	1.268 (4.581)	0.108 (4.570)	3.265 (5.284)	1.970 (4.313)
Search \times Women	-11.101* (6.291)	-10.846* (6.321)	-9.416 (6.279)	-12.805** (6.349)	-12.365* (6.348)	-14.107** (6.960)	-13.228** (5.994)
Income	13.999*** (5.351)	14.407*** (5.371)	13.215** (5.340)	14.795*** (5.398)	14.377*** (5.381)	14.156*** (5.357)	17.624*** (5.150)
Income \times Women	-13.856* (7.586)	-14.684* (7.601)	-12.785* (7.575)	-14.823* (7.654)	-13.535* (7.623)	-14.229* (7.585)	-17.905** (7.227)
Uninformed	-1.988 (5.367)	-1.943 (5.387)	-2.774 (5.364)	-1.195 (5.408)	-0.983 (5.395)	-2.005 (5.368)	-0.855 (5.157)
Uninformed \times Women	-5.597 (7.500)	-5.996 (7.525)	-4.966 (7.509)	-6.261 (7.563)	-6.602 (7.552)	-5.321 (7.509)	-5.405 (7.288)
Changed P	0	116	240	240	56	119	
Observations	10323	10323	10323	10323	10328	10333	10324

Notes: This table reports coefficients from a regression of total price estimates on an indicator for woman customer, and indicator for each customer type, and their interaction with woman customer, with baseline customer type as the omitted categories. Columns (1)-(6) redefine total quote measure. Column (1) uses the main pre-specified quote definition. Column (2) excludes discounts from quotes. Column (3) and (4) use minimum and maximum quote whenever a range of prices is provided, Column (5) selects minimum price provided by other shop to a customer during that week whenever a shop offers to price match quotes, Column (6) includes invalid quotes. Column (7) uses the main quote definition in a robust regression. Robust standard errors in parenthesis, with significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Alternative competition measures

Table A.9: Robustness: Alternative competitors definitions

	Main sample (1)	Robust regression (2)	Online shops (3)	Unrestricted sample (4)
Competitors	-0.623* (0.336)	-0.545* (0.329)	0.086 (0.488)	-0.618** (0.253)
Competitors × Women	-1.470*** (0.459)	-1.390*** (0.440)	-2.079*** (0.669)	-0.963*** (0.342)
Women	15.274*** (4.507)	15.867*** (4.318)	14.948*** (4.459)	15.403*** (4.623)
Controls	Yes	Yes	Yes	Yes
Observations	10,323	10,313	10,323	10,323
Av. Competitors	3.88	3.88	2.57	5.00

Notes: This table reports coefficients from a regression of total price estimates on market concentration measures and their interaction with woman customers. Each column used an alternative definition to measure competitors within a 1-km radius. Column 1 repeats the benchmark specification, with competitors defined as all mechanic shops within a 1-km radius. Column 2 uses the same measure and performs a robust regression - which weights observations differently based on how well behaved these are. Column 3 counts nearby competitors that have a listed website or email address available in the YP. Column 4 includes the unrestricted sample of listed car repair shops in the YP, regardless of being identified as mechanic shops. All specifications control for woman indicator, each customer type and their interaction with women indicator with baseline customer type as the omitted category, state fixed effects, other email varying items: car year, subject and script number controls, and day of week and week fixed effects - indicating when the email was sent to a shop, non-independent shop control, whether a non-independent shop is within the market concentration measure, CZ fixed effects, with a dummy for observations not linked to a CZ. Robust standard errors in parenthesis, with significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.10: Price gaps with alternative distance radius: Robustness

	1-km (1)	1.5-km (2)	2-km (3)	2.5-km (4)	3-km (5)	1-km (6)
Competitors	-0.623* (0.336)	-0.546** (0.233)	-0.625*** (0.179)	-0.623*** (0.146)	-0.588*** (0.117)	-0.064 (0.378)
Competitors×Women	-1.470*** (0.459)	-0.967*** (0.304)	-0.682*** (0.223)	-0.455*** (0.175)	-0.335** (0.141)	-1.356** (0.532)
Women	15.274*** (4.507)	16.239*** (4.574)	16.283*** (4.637)	15.728*** (4.659)	15.300*** (4.679)	15.578*** (4.736)
Main controls	Yes	Yes	Yes	Yes	Yes	Yes
Added controls						Yes

Notes: This table reports coefficients from a regression of total price estimates on number of competitors and its interaction with a dummy for female. Each column defines the number of competitors within radius thresholds increasing by 0.5 km, starting with a 1 km radius in column (1), and ending with a 3-km radius in column (6). Column (1) defines competitors as the count within a 1-km radius, and the specification includes controls for competitors in the 1.5, 2, 2.5 and 3 km outer rings. All specifications control for woman indicator, a woman indicator, and each customer type control, and their interaction with woman customer, with baseline customer type as the omitted categories, state fixed effects, other email varying items: car year, subject and script number controls, and day of week and week fixed effects - indicating when the email was sent to a shop, non-independent shop control, whether a non-independent shop is within the market concentration measure, CZ fixed effects, with a dummy for observations not linked to a CZ. Robust standard errors in parenthesis, with significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

10 Are effects driven by other service offers?

Table A.11: Effect of additional part and service offers

	(1)	(2)	(3)	(4)	(5)
	No detail	Add parts	Fees	Warranty	Offers
Women	-0.026** (0.013)	0.001 (0.012)	0.009 (0.016)	0.010 (0.009)	-0.001 (0.007)
Search	-0.005 (0.015)	-0.004 (0.013)	0.019 (0.017)	0.021** (0.010)	0.001 (0.007)
Search × Women	0.018 (0.020)	-0.000 (0.017)	-0.006 (0.024)	-0.011 (0.014)	0.006 (0.010)
Income	-0.033** (0.017)	-0.000 (0.015)	0.052** (0.020)	0.004 (0.012)	-0.005 (0.008)
Income × Women	0.025 (0.024)	0.013 (0.022)	-0.032 (0.029)	-0.007 (0.017)	-0.002 (0.011)
Uninformed	0.051*** (0.018)	-0.014 (0.015)	-0.028 (0.020)	-0.005 (0.011)	0.017* (0.010)
Uninformed × Women	-0.041 (0.025)	0.007 (0.021)	0.003 (0.028)	-0.009 (0.016)	-0.003 (0.014)
Controls	Yes	Yes	Yes	Yes	Yes
Av Mean	0.21	0.15	0.35	0.08	0.04

Notes: This table reports coefficients from a regression of not giving detail of total quote components (Column 1), including additional car parts in quote (radiator cap, radiator hoses, thermostat or others) (Column 2), include tax and fees in quote (Column 3), providing a warranty (Column 4), and providing additional offers such as discount, a free inspection, price-matching, and financing options or a shuttle (Column 3) on an indicator for woman customer, and indicator for customer types (excluding baseline type) and female, and indicators for each customer type, with baseline customer type as the omitted categories, commuting zone and state fixed effects, email varying items: car year, subject and script number controls, and adds day of week and week fixed effects, and a dummy for non-independent shop (dealership and franchises). Every variable takes a value equal to 0 when the shop does not mention its inclusion. Robust standard errors in in parenthesis, with significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.5: Email replies: Impact of gender and gender-type on replies

	Invalid (1)	Not repair (2)	Repair (3)	Give quote (4)	Give quote (5)
A. Men customers					
Search	0.006 (0.006)	-0.003 (0.003)	-0.034*** (0.006)	-0.020*** (0.005)	-0.020*** (0.005)
Income	-0.002 (0.007)	-0.002 (0.004)	-0.002 (0.008)	0.009 (0.007)	0.009 (0.007)
Uninformed	-0.002 (0.007)	-0.004 (0.004)	0.010 (0.008)	-0.004 (0.007)	-0.005 (0.007)
Constant	0.189*** (0.004)	0.055*** (0.002)	0.252*** (0.004)	0.181*** (0.004)	
A. Women customers					
Search	0.001 (0.006)	-0.006* (0.003)	-0.023*** (0.006)	-0.012** (0.006)	-0.013** (0.006)
Income	0.000 (0.007)	-0.003 (0.004)	-0.000 (0.008)	0.004 (0.007)	0.004 (0.007)
Uninformed	-0.004 (0.007)	0.003 (0.004)	0.006 (0.008)	-0.013* (0.007)	-0.012* (0.007)
Constant	0.186*** (0.004)	0.056*** (0.002)	0.258*** (0.004)	0.189*** (0.004)	
C. Gender differential by customer type					
Search×Women	-0.007 (0.006)	-0.002 (0.003)	0.018*** (0.006)	0.016*** (0.006)	0.015*** (0.006)
Income×Women	0.000 (0.008)	-0.001 (0.005)	0.007 (0.009)	0.003 (0.008)	0.003 (0.008)
Uninformed×Women	-0.004 (0.008)	0.007 (0.005)	0.002 (0.009)	-0.001 (0.008)	0.000 (0.008)
Baseline×Women	-0.002 (0.005)	0.000 (0.003)	0.006 (0.006)	0.009 (0.005)	0.008 (0.005)
Main Controls					Yes

Notes: This table reports coefficients from a regressions of invalid email reply rates (Column 1), not perform service shop reply rates (Column 2), do service reply rates (Column 3), and do service and give price estimate reply rates (Columns 4 and 5) on customer types separately for each gender (Panel A presents results for men, Panel B for women) with baseline type customers as the omitted category. Panel C includes both genders in the sample and adds the interaction coefficients between each customer type and women customers. Column 5 controls for state fixed effects, other email varying items: car year, subject and script number controls, and day of week and week fixed effects - indicating when the email was sent to a shop, and non-independent shops. Robust standard errors in parenthesis, with significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.