Title
Market Structure and Behavioral Frictions: Demand and Supply Perspectives

Author
Rubinstein, Alon Yehoshua

Publication Date
2023

Peer reviewed|Thesis/dissertation

Market Structure and Behavioral Frictions: Demand and Supply Perspectives

by

Alon Y. Rubinstein

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Prof. Stefano DellaVigna, Co-chair
Prof. Benjamin Handel, Co-chair
Prof. Kei Kawai

Spring 2023

Market Structure and Behavioral Frictions: Demand and Supply Perspectives

Abstract

Market Structure and Behavioral Frictions: Demand and Supply Perspectives

by

Alon Y. Rubinstein

Doctor of Philosophy in Economics

University of California, Berkeley

Prof. Stefano DellaVigna, Co-chair

Prof. Benjamin Handel, Co-chair

In my research, I use modern industrial organization theory and econometrics to study the impact of behavioral frictions among buyers and sellers in different market environments on market outcomes and welfare. Following the seminal works of Becker (1957) and Arrow (1972), it is common knowledge that market competition mitigates such non-cognitive frictions as employer taste-based prejudice. Hence, it is not surprising that today, almost half a century after Kahneman and Tversky (1979) introduced their Prospect Theory, many economists believe that market forces —the invisible hands of sophisticated profit-maximizing suppliers—mitigate supply-side cognitive-based frictions as well. In my research, I challenge these popular views and show that market forces do not necessarily compete away the effects of behavioral cognitive or non-cognitive frictions in firm pricing and employment.

The dissertation consists of three chapters: (i) "Behavioral Professionals: Evidence From the Commercial Auto Insurance Industry," (ii) "Price and Prejudice: Customer Taste-Based Discrimination and Competition," and (iii) "The Lemons Gap: Demand For Insurance of Quality Uncertain Goods."

In Chapter 1, I present my work "Behavioral Professionals: Evidence From the Commercial Auto Insurance Industry." A cornerstone of the IO study of *selection markets* is that competition disciplines sellers to customize coverage and premiums optimally. But is this the case? Using data from one of the largest Israeli commercial auto insurance providers, an affiliate of a multinational insurance company, I find there is too little adjustment in the intensive margin. Premiums barely change with expected costs as projected by pre-determined factors (vehicle age) and signals (claim history). At the same time, I find there is too much adjustment in the extensive margin, with an excessive denial of insurance in response to recent claims.

Using unique grading documents, I integrate the insurer's subjective risk assessment into the study of insurance markets. I find that the insurer's risk assessment outweighs recent claims and misevaluates vehicle age. Structural model estimates suggest that insurers enjoy incumbency advantages over their own customers, and clients are rationally inattentive to competitors' pricing unless they are faced with a price increase. Both channels allow sub-optimal behavior to persist. Finally, I find that supply-side behavioral frictions, which result in excessive denial, mainly harm *disadvantaged customers*—single-fleet clients of old vehicles—and diminish with the client's fleet size.

In Chapter 2, I present my work "Price and Prejudice: Customer Taste-Based Discrimination and Competition." This work investigates the effect of competition on the incidence of tastes for discrimination. The model shows that monopolistic sellers discipline discriminatory buyers by taxing their taste for discrimination. In equilibrium, monopolistic sellers hire a lower share of White workers and pay them a lower premium than sellers in a competitive market. These results are tested in the context of the US banking deregulation that affected product market competition, and I quantify its impact on customer-driven labor market discrimination. Using Census/ACS data from 1960 to 2010, O*NET measures of job requirements, and GSS measures of discriminatory attitudes against Black by state, I find that the Black-White wage and employment gaps increased following bank deregulation in jobs requiring intensive contact with clients, especially in states with high measures of prejudicial preferences.

In Chapter 3, I present my work "The Lemons Gap: Demand For Insurance of Quality Uncertain Goods." This work studies the difference between insuring a quality uncertain good and a monetary loss. I integrate key insights from the pre-owned market into the analysis of the demand for insurance. I find that adverse selection in the resale market results in a missing insurance market. There's a gap between the insured vehicle and the resale market's quality, especially for new vehicles. As a result, clients over-insure their quality uncertain goods, yet demand drops over the vehicle life cycle. The partial compensation further amplifies over-insurance patterns driven by behavioral attributes. The gap results in time trends. As the vehicle ages, demand drops, the insurance market is more adversely selected, and moral hazard increases. The incomplete compensation can result in context-dependence demand for insurance, customers' over-insurance limited risk, in general, and for durable goods, in particular.

Acknowledgments

Completing a doctoral dissertation is an arduous journey, and I am humbled and grateful for the support and guidance of many individuals who have contributed to my success. It is my pleasure to acknowledge their contributions and extend my heartfelt thanks.

First and foremost, I would like to thank my advisors, Prof. Stefano DellaVigna and Prof. Benjamin Handel, for their constant support, invaluable mentorship, and advice throughout this process. Their dedication and guidance have been truly instrumental in shaping this work. I will always cherish their training and guidance.

I would also like to express my gratitude to Prof. Ned Augeblick, Prof. Kei Kawai, and Prof. Steve Tadelis. Their insightful comments, feedback, and expertise have been pivotal in helping me complete this dissertation and will serve me throughout my academic career. I would also like to thank Prof. David Card, Prof. Dmitry Taubinsky, and Prof. Jonathan Kolstad, and Prof. David Card for their fruitful comments and suggestions.

Lastly, I would like to thank my parents, Emanuela and Yona, and my brothers, Eran and Yuval, who have provided me with unwavering support, encouragement, and understanding throughout the demanding journey of completing a doctoral degree.

Contents

List of Figures

List of Tables

Chapter 1

Behavioral Professionals: Evidence From the Commercial Auto Insurance Industry

1.1 Introduction

Perhaps the key feature of *selection markets*, such as the private market for insurance, is that consumers vary not only in their willingness to pay but also in how costly they are to the seller. Therefore, insurance providers care about both the *quantity* of policies they sell and the *quality* of the clients they cover. Market forces can fail to achieve efficiency if buyers know better than sellers how risky they truly are (Akerlof, 1970, Rothschild and Stiglitz, 1976).

Thus, it is not surprising that much attention in the IO literature is devoted to studying the demand side in insurance markets and quantifying the implications of selective sorting on policy and welfare (e.g., Cutler and Reber, 1998, Chiappori and Salanie, 2000, Cardon and Hendel, 2001, Cohen and Einav, 2007, Fang, Keane, and Silverman, 2008, Cutler, Finkelstein, and McGarry, 2008, Carlin and Town, 2009, Lustig, 2010, Einav, Finkelstein, and Cullen, 2010a, Einav et al., 2013, Starc, 2014, Finkelstein and Poterba, 2014, Hackmann, Kolstad, and Kowalski, 2015, Handel, Hendel, and Whinston, 2015a, Cabral, 2016, Cabral, Geruso, and Mahoney, 2018 and Einav, Finkelstein, and Tebaldi, 2019).[1] A cornerstone of these studies is that sellers customize coverage and premiums optimally. The view is that while market forces might fail to discipline buyers, competition disciplines similar, equally informed

[1]See handbook chapters by Einav, Finkelstein, and Mahoney (2021) and Handel and Ho (2021) describing studies on IO of selection markets, in general, and insurance markets, in particular.

insurers; providers offer optimal coverage and premiums. In reality, however, insurers are not identical and might not be equally informed. For example, customers know their insurer better than other providers, and insurers know their clients better than their competitors. Due to these relational asymmetries, market forces might fail to eliminate mispriced risk, suggesting that supply-side frictions may affect premiums and coverage.

In this paper, I offer an alternative perspective on the insurance market. As in recent studies, I allow for behavioral frictions in demand for insurance. However, in contrast to the literature on selection markets, I recognize that insurers might fail to assess risk accurately and to customize coverage plans and prices optimally.

With this perspective in mind, this paper studies the pricing and coverage behavior in private insurance markets. Specifically, I address four main questions. First, do large sellers in private insurance markets customize offers and prices as the IO literature predicts? If not, does it reflect their biased beliefs? Third, why do market forces fail to compete away sellers with biased beliefs? Fourth, what are the welfare implications of supply-side behavioral frictions, and do these vary if clients are covered by individual or fleet base contracts?

To address these questions, I study the Israeli commercial auto insurance market. This market provides an excellent laboratory to study supply-side frictions for two reasons. First, we expect professional buyers to choose carefully between insurance plans and discipline sellers to customize offers accordingly. Second, the market is limited to unregulated property coverage; insurers can charge any price and deny coverage without constraints. I use comprehensive data from one of the largest commercial auto insurance providers in Israel, an affiliate of a large multinational insurance company. The data includes *all* the information available to the insurer: (i) premiums, coverage, and claim expenses by policy and client; and (ii) internal policy pre-renewal assessments, known as the *"Go—No Go" grades.* Furthermore, I obtain data on the market competitors' premiums by generating fictitious policy applications. These datasets allow me to portray the gap between premiums and expected costs by pre-determined factors and claim history, identify the gap between objective and subjective risk assessment, and quantify its impact on coverage, pricing, profits, and welfare.

I start by providing evidence of the gap between premiums charged by the insurer and the cost of providing coverage as a function of pre-determined and stochastic factors. I find there is too little adjustment in the intensive margin. The insurer barely adjusts premiums per value with determinants predicting higher expected cost per value, such as vehicle age and claim history. Consequently, the insurer profits by providing coverage to new vehicles and clients with favorable past performance.

Interestingly, I find that the insurer's adjustment of premiums per value regarding claim history is based solely on recent claims while putting no emphasis on augmented past performance, despite aggregate claim history serving as a predictive signal of future claims. In contrast, recent performance has no additional predictive power.

Next, I study whether these pricing patterns are specific to this particular insurer or apply to other market competitors. Specifically, I generate fictitious policy applications using an Israeli insurance agency and examine how the premiums vary by vehicle characteristics and claim history. I find that the market-wide price patterns are comparable to those of the insurer. Moreover, the analysis of market-wide premiums for new policies indicates a substantial adjustment on the extensive margin. It is impossible to generate a premium offer for a new policy if a customer has been involved in at least two claim events in the three preceding years.

After providing evidence indicating that both the insurer and its competitors do not adjust premiums based on a customer's observable characteristics, I turn to the internal grading documents. Despite their richness, the observed prices are insufficient to identify the insurer's beliefs, as both supply and demand factors determine equilibrium premiums. I exploit the variation in the "Go—No Go" grades and policies' observable characteristics to identify the impact of pre-determined and stochastic factors in determining the insurer's subjective risk assessment. Internal grading recommends no change in premiums—"Go"—for most policies, ignoring the predictive power of vehicle age and claim history on costs. The lack of recommended price adjustment spills over to coverage. Internal grading data recommends denying comprehensive coverage to almost half of the "No Go" graded policies rather than increasing their premiums. This is especially relevant for old vehicles and costly clients. This strategy also reflects a biased risk assessment as signaled by recent and augmented claim history. The "Go—No Go" grades are too sensitive to recent claims with almost no predictive power of future costs, conditional on the augmented history of claims.

Internal grading and premiums reflect demand and supply factors. To distinguish between these forces, I develop and estimate a structural model that allows customized prices and coverage to reflect the insurer's subjective risk assessment and commonly used demand and supply factors. In its simplified version, the model consists of two periods. In the first period, customers self-sort to sellers. In the second period, a wedge emerges between their insurer and other providers. Customers decide whether to renew their policies or search for an outside offer. The decision to renew depends not only on their private information and search costs but also on the supply side, that is, their insurer's private information, subjective risk assessment, and customized offers.

I take advantage of the panel structure and internal grading to identify and quantify a client's willingness to pay and the insurer's subjective risk assessment.

Regarding the willingness to pay, a key concern is that premiums are subject to strategic considerations. I use the panel structure of my data, which follows many clients with large fleets over multiple coverage periods, to identify an external source of variation in premiums and estimate clients' willingness to pay. The across-client variation permits conditioning out the client-specific effect on premiums. The within-client variation allows for identifying exogenous shocks in price adjustments by using predicted —rather than actual—adjustments in premiums for those who renew their coverage (Bundorf, Levin, and Mahoney, 2012) and those who do not (Crawford, Pavanini, and Schivardi, 2018).

As for the insurer's subjective risk assessment, this is identified by decomposing expected profits into premiums and expected costs using a two-step procedure. First, I take advantage of the informational symmetry between the insurer and the econometrician to nonparametrically identify the expected profits for each policy by inversion of the share of recommendations (Berry, 1994). Then, using policies for which no change in premiums is recommended ("Go"), I identify the insurer's beliefs of expected costs by subtracting the previous year's premiums from expected profits. To further account for possible latent strategic considerations, I focus on nonfleet clients.

Three main demand side findings emerge. First, customers adversely select to renew coverage. Second, new customers are adversely selected; they cost more, conditional on observables. Last, customers are rationally inattentive to premiums unless they incur a price increase. Both the adverse selection of new customers and the rational inattention of renewing consumers point to asymmetries between incumbent insurer-insuree pairs and others that allow the insurers room for error in customized prices and coverage.

In terms of the supply side, the insurer's subjective risk assessment, two main findings emerge. First, the insurer gives more weight to recent claims without predictive power of future costs. The law of large numbers implies that demand exhibits increasing returns to fleet size, as large fleets are less likely to be affected by the insurer's biased risk assessment. Second, the insurer erroneously evaluates common predetermined factors such as vehicle age.

Using the estimated demand and supply parameters, I analyze the impact of the insurer's biased risk assessment on coverage, premiums, profits, and welfare using a set of counterfactuals. I find that supply-side frictions mainly harm *disadvantaged customers*—single-fleet clients of old vehicles. A profit-maximizing firm does not deny coverage as informational asymmetries between the customer and the insurer are modest. In contrast, the insurer denies coverage to old vehicles and clients with poor recent performance, which results in lower profits. Furthermore, the clients face a substantial reclassification risk, diminishing with customers' fleet size; volatility in recent performance drops with the number of insured vehicles. As a result, customers benefit from purchasing coverage as a group.

Finally, do insurers adjust premiums once they learn those might be mispriced? During my study, I shared my preliminary stylized findings that point to possible mispricing by vehicle age with the managerial team. To assess the impact of information on pricing, I compare premiums by vehicle age over the covered period. I find almost no change in premiums and profits in consecutive years before the managerial team learned about my findings. In contrast, I find a moderate increase in premiums between periods once they were informed of my findings. The adjustment of prices upon learning is consistent with my findings that much of the mispricing reflects the insurer's biased assessment of risk rather than strategic considerations.

This paper contributes to the literature on firm behavior. A growing body of research documents that large suppliers in nonselection markets customize prices too little based on observable demand factors (Orbach and Einav, 2007, McMillan, 2007, Cho and Rust, 2010, Shiller and Waldfogel, 2011, Cavallo, Neiman, and Rigobon, 2014, DellaVigna and Gentzkow, 2019). My paper shows that sellers in selection markets fail to customize prices also on expected cost, a cornerstone of the study of selection markets.

This paper also adds to the literature on reclassification risk and market unraveling (Cutler and Reber, 1998, Hendel and Lizzeri, 2003, Koch, 2014, Finkelstein, McGarry, and Sufi, 2005, Handel, Hendel, and Whinston, 2015a, Hendren, 2017, Fleitas, Gowrisankaran, and Sasso, 2020, Ghili et al., 2021, Cuesta and Sepúlveda, 2021). IO theory attributes both aspects to asymmetric information and regulation. I find that insurers do not adjust premiums and excessively deny customers, despite modest asymmetric information and no regulations. For individual customers, the insurer amplifies welfare loss from the reclassification of risk. This relates to group insurance (Bundorf, Levin, and Mahoney, 2012, Tilipman, 2022), as client's size dilutes insurers' misevaluation of risk and the consequences of overdenial.

The paper also relates to the literature on imperfect competition in selection markets (Veiga and Weyl, 2016, Mahoney and Weyl, 2017, Lester et al., 2019, Cuesta and Sepúlveda, 2021, Tebaldi, 2022). I show two channels generating market power: information asymmetries among insurers (Jin and Vasserman, 2020) and behavioral demand-side frictions (Sydnor, 2010, Abaluck and Gruber, 2011, Barseghyan et al., 2013, Handel, 2013, Handel and Kolstad, 2015, Spinnewijn, 2017, Bhargava, Loewenstein, and Sydnor, 2017, Brot-Goldberg et al., 2017, Ho, Hogan, and Scott Morton, 2017, Handel, Kolstad, and Spinnewijn, 2019a, Gottlieb and Smetters, 2021).[2] Im-

[2]Due to informational asymmetries across insurance providers, a perfectly competitive outcome is implausible even when considering a frictionless economy with homogeneous products.

perfect competition allows imperfect behavior by insurers, which might drastically change the welfare consequences of imperfect competition.

Last, my paper also contributes to the literature on managerial practices, highlighting the impact of monitoring, feedback, and on-the-job training (Bloom and Van Reenen, 2007 and Bloom et al., 2013). My findings indicate that these elements improve profits even among professional sellers in a big-data industry.

The remainder of the paper is organized as follows. In section 2, I describe the setting. Section 3 provides descriptive statistics regarding the data I exploit in the empirical analysis. In Section 4, I provide evidence of the gap between insurer pricing and expected cost. Section 5 presents evidence of the gap between objective and subjective expected costs using the "Go—No Go" grades. Section 6 then develops and estimates the demand for policy renewal, and section 7 develops and estimates the insurer's subjective costs and supply of insurance. In section 8, I conduct a counterfactual analysis to study the implications of supply-side frictions. Finally, Section 9 concludes.

1.2 Setting

In this paper, I take advantage of proprietary for the years 2013 to 2020 obtained from a large Israeli company (with an annual average revenue during the sample period of $37.5 million in 2020 terms) operating in the commercial auto-insurance market to examine the relationship between insurer pricing, perceived costs, and customer's realized costs.[3] The provided dataset includes all data the insurer has from 2013 to 2020. In the empirical application, I take advantage of the information symmetry between the insurer and the econometrician in terms of the determinants of costs and pricing. In this section, I characterize in detail the insurer's affiliation with an international insurance company, its portfolio (in terms of both customers and products), data, and business operations.

Vertical Relationship

The insurance company operates under a unique vertical relationship, compared with the standard market structure in the insurance literature in general and in the auto insurance literature in particular. The Israeli insurance company is an affiliate of a large international insurance company (henceforth, IIC) with asset value of over

[3]Throughout the paper, I use and report monetary values in nominal New Israeli Shekels (ILS) to avoid creating artificial variation in the data. Annual inflation between 2013 and 2020 ranged from 0.84% to -0.63%, and the value of 1 ILS ranged from $0.26 to $0.29.

\$50 billion as of December 2020. The IIC provides capacity, which allows the Israeli insurer to sell policies. This is a result of regulation in Israel, which sets reserve requirements per premiums charged to avoid the failure of insurers to repay claims.[4] In terms of the division of cost and revenue, the IIC pays all claim damages (net of deductibles), while the Israeli insurer pays all additional operational costs. Revenue is split between the IIC and the Israeli insurer based on yearly agreed-upon shares.

A possible concern is that the distorted incentives might lead the Israeli insurer to oversupply insurance, as claims are paid by the IIC. Thus, the IIC provides guidance on pricing by setting a lower bound on premiums charged, conditional on vehicle and client observable characteristics. As a result of repeated interactions with the IIC, the Israeli insurer puts emphasis on the portfolio's return, as negative outcomes often lead to a lower share of revenue in succeeding years. Throughout the paper, I consider the joint profits of both the Israeli insurer and the IIC from operations in this market as a whole.

Insurer Portfolio

The insurer provides three types of coverage: (i) third-party Coverage, which only covers the cost of damage to third-party property; (ii) comprehensive coverage, which includes all damages to a vehicle in addition to damages covered by third-party coverage; and (ii) partial coverage, which covers the same types of damage as comprehensive coverage, excluding theft. None of these types of policies cover bodily injuries to the policyholder (or to third parties). Israeli regulations mandate that all vehicle owners purchase special coverage for bodily injury through a separate, heavily regulated policy. The vast majority of the insurer's portfolio—over 87%—consists of comprehensive coverage policies (see Appendix Figure A.2, Panel A).

The insurers provide coverage to commercial vehicles, including trucks, buses, mini-buses, trailers, and heavy equipment (e.g., tractors, bulldozers, cranes). Approximately two-thirds of the insurer's portfolio consists of trucks (see Appendix Figure A.2, Panel B). Therefore, I focus mainly on comprehensive coverage policies for trucks throughout the empirical application.

The relationship between the insurer and the clients differs in four ways from the common relationship in previous studies on insurance markets. First, clients typically own a fleet of vehicles. The insurer's clients are quite diverse in terms of their fleet size. More than 10% of policies are of a single-vehicle client, a quarter of policies are of clients who insure a fleet of fewer than five vehicles, and more than a

[4]Throughout the sample years, the capacity constraint was not binding. Therefore, I consider the opportunity cost of providing insurance to a different customer to be zero.

quarter of the policies are of clients who insure at least 100 vehicles (see Appendix Figure A.1).

Furthermore, unlike some markets in which firms offer "take-it-or-leave-it" prices, equilibrium premiums are the result of bargaining between the insurer and the client (especially when they are fleet owners). Equilibrium price-setting has both favorable and unfavorable consequences. On the one hand, premiums are endogenous and might be correlated with unobservable (to the insurer and the econometrician) components of the demand for insurance. On the other hand, exogenous variation in pricing as a result of a firm experiment or pilot does not allow examination of its behavior relative to profit-maximizing behavior, as these prices are, by definition, nonoptimal off-equilibrium premiums.

As in other markets, premiums are prorated. Yet, unlike standard insurance markets, 8 comprehensive and partial coverage policies are priced in terms of *premium per value*. For instance, a premium per value of 4% implies that the customer pays 4,000 ILS for a 6-month policy for a vehicle valued at 200,000 ILS.[5]

Finally, unlike the common setting in the selection market literature, there is no regulation of pricing and coverage provided by the insurer. The lack of regulation is consequential; the insurer can provide a customer any coverage at any premium, and deny coverage if it wishes to do so. Since premiums are not regulated, IO theory suggests that customer denial can only be explained by excessive adverse selection (Akerlof, 1970).[6]

Business Operations: "Go-No Go"

The Israeli insurer, operating since the 1950s, employs hundreds of workers. These include employees in the analytical team, overseen by the Chief Operating Officer (henceforth, COO), and underwriters, who are in contact with the customers, either directly or through their agents. Over the sample period, the insurer sold approximately 175,000 policies to over 13,000 different customers. Due to the high volume of customers, and the differentiated occupational requirements and skills of employees, the firm operates in an orderly, systematic structure.

On a monthly basis, the COO provides the underwriters a document, which is named "Go—No Go". The document includes a grade for each policy that is about to end (usually a month or two before the end of the policy coverage contract). The grading system is defined as follows: a "Go" grade implies that the analytical team recommends renewing the customer's policy at the same premium per value and

[5] Vehicle values are usually determined by the Levi Itzhak vehicle price list, which is the standard practice by both commercial and noncommercial auto insurance markets.

[6] An alternative explanation could be high operational costs.

terms (deductibles). A "No-Go" grade implies that the analytical team recommends nonrenewal of the customer's policy at the same premium per value. Typically, a "No-Go" grade will include a recommendation on how to continue the relationship with the customer, if at all. There are four common recommendations: (i) renew the policy without increasing premiums (i.e., increase deductibles), (ii) increase the premium per value of the policy, (iii) do not provide comprehensive coverage (i.e., third-party only), and (iv) do not provide any coverage (deny). An example is provided in Appendix Figure A.3.

At first glance, the complexity of the firm's operations might seem disadvantageous. Yet, this complexity provides additional information which otherwise could have been obtained only by a survey of the analytical team employees. I take advantage of the insurer grading in the empirical application to extract insurer beliefs. Observed prices are insufficient to identify the insurer's beliefs, as they are determined in equilibrium by both supply and demand forces; an insurer might increase prices either because of high expected costs or due to high demand for insurance. Additional assumptions are required to differentiate between the two. In this paper, I identify the insurer's beliefs by exploiting the variation in insurer grading without any structural assumption regarding insurer behavior.

1.3 Summary Statistics

In this section, I describe the data and provide descriptive evidence of insurer pricing and realized costs (I refer to the "Go—No Go" grades in section 1.5). As mentioned earlier, the main dataset includes all the insurer's data from 2013 to 2020. The data include (i) contract characteristics (premium, coverage type, deductibles, duration, and an indicator on whether a driver under the age of 24 is allowed to operate vehicles), (ii) vehicle characteristics (vehicle value, vehicle age, vehicle type, vehicle weight, and vehicle model), (iii) customer characteristics (claim history, zip code, and fleet size), (iv) costs (commission and claim damages), and (v) identifying information (policy id number, vehicle license number, and client id number).[7] Using the identifying information, it is possible to track the clients and their vehicles over time.

In the empirical analysis, I focus on comprehensive insurance policies. Column 1 in Table 2.1 presents the summary statistics of comprehensive coverage policies

[7]In this market, insurance policies are not tied to a specific driver but rather to a specific vehicle. In general, each insured vehicle can be operated by any driver over the age of 24 with a valid license to operate a vehicle of that class. A client can extend the policy coverage for young drivers (between the age of 21 to 24), which in general increases the premium charged.

for trucks (Column 1 in Appendix Table A.1 presents the summary statistics for all vehicles).[8] The sample consists of 51,684 policies. Claims are reported for roughly a quarter of the policies. The insurer enjoys a mean profit of 1,587 ILS and a profit margin of 16%, as the mean premium and costs are 9,938 (or 3.33% of the vehicle value) and 8,361 ILS, respectively. From the insurer's perspective, the portfolio's performance (a profit margin of 16%) is satisfactory. The insurer's positive performance is of key importance, as the systematic mispricing of risk cannot be concluded by observing a failing firm. The profitable performance supports the notion that this particular insurer is not *competed away* from the market. This is complemented by the fact that the insurer is one of the largest insurers in the Israeli commercial auto-insurance market and is affiliated with one of the largest multinational insurance companies.

In this paper, I examine whether the insurer assesses risk correctly in both the intensive and the extensive margins. I examine the intensive margin by estimating the gap between the actual pricing and the expected cost as a function of a customers' observable factors. In the case of a recurring customer, the insurer should adjust premiums or possibly even deny insuring the customer based on changes in the observable characteristics. The evolution of observable factors can be divided into two groups: predetermined changes (i.e., vehicle age) and stochastic shocks (i.e., claim history).

In Figure 1.1, I provide a first glance at the relationship between the premium, costs, and vehicle age. Figure 1.1, presents the mean premium per value and the cost per value by vehicle age (0 to 10) for comprehensive insurance policies for trucks. As is apparent from the figure, the premiums do not adjust optimally over vehicle age, as the cost-per-value curve is, on average, a counter-clockwise rotation of the premium-per-value curve. Therefore, the insurer generates its profits by providing insurance for new vehicles and incurs losses on old vehicles.[9]

Column 2 in Table 2.1 presents the summary statistics for a subset of the comprehensive coverage policies for trucks: those with a vehicle age of six years or above.[10] This subset represents about one-fifth of this sample.[11] In general, the share of poli-

[8]In the empirical framework, I mainly focus on trucks, as the premiums charged by the market competitors are available for trucks only. I also repeat the entire empirical analysis for all vehicle as well. The results, which are consistent with those analyzing insurance for trucks, are reported in the appendix.

[9]The nominal relationship between mean premium, cost, and vehicle age is of a similar nature and is reported in Appendix Figure A.4.

[10]Column 2 in Appendix Table A.1 presents the summary statistics for the equivalent subset of all vehicles.

[11]Appendix Figure A.5 depicts the distribution of policies by vehicle age.

cies involved in a claim is not higher than that of the entire sample (Column 1). Yet, the mean damage (i.e., cost of claims) per value is substantially higher, by more than 60%. The mean damage is 14% lower, yet the vehicle value depreciates by 47%. This suggests that, conditional on a claim, the expected damage is not proportional to vehicle value. Unlike the mean damage per value, the mean premium per value increases by only 20%. As a result, providing comprehensive coverage policies to old vehicles (six years and above) on average does not benefit the insurer but instead generates losses. Consequently, the insurer's satisfactory profit margin is derived by mostly providing coverage to relatively new trucks.

Regarding past performance, Column 3 in Table 2.1 presents the summary statistics for a subset of the comprehensive coverage policies for trucks: those that reported a claim in the previous period.[12] This subsample consists of 16% of the whole sample and does not differ substantially from that of the entire sample in terms of premium per value (3.56% relative to 3.33%), vehicle age (5.00 relative to 4.18), or vehicle value (275,954 relative to 298,659).[13] Yet, policies with past realized claims incur higher costs. Among policies with a reported claim in the previous period, 34.82% report a claim also in the current period (relative to 23.98%), and the mean damage is 9,345 ILS, which is substantially higher than the entire sample (6,794 ILS). Consequently, on average, the insurer exhibits losses for providing coverage in policies with a reported claim in the previous period.

Thus far, I have divided the sample based on whether a policy incurred a claim in the previous period. In the insurance market in general, and with this insurer (and its competitors) in particular, it is customary to measure the client's performance based on the aggregate loss ratio. The aggregate loss ratio is defined as the ratio of damages (i.e., net cost of claims) to revenue (i.e., premiums) with respect to all of the customer's past policies. In Panel A of Figure 1.2, I divide the sample into four groups based on the level of loss ratio (at the start of the policy) and depict the relationship between the premium per value, cost per value, and the client's aggregate loss ratio. The mean cost per value increases with the loss ratio, which suggests a positive and persistent relationship between past and future performance, consistent with the

[12] It should be noted that "reported claim" does not necessarily imply that the customer reported the claim, as third parties usually report claims on customers that generate third-party damages. Furthermore, throughout the analysis, I do not consider the "at-fault" side. A reported claim is defined as an event in which the insurer exhibits costs as a consequence of providing coverage to the client. In addition, Column 3 in Appendix Table A.1 represents the same segmentation of all vehicles.

[13] Since the sample begins in 2013, I am unable to observe reported claims in the previous period. I take this into account in the empirical framework and omit the 2013 policies, or policies of new clients, when conducting comparative statics regarding past performance. About one-fifth of the policies with at least one year of documented history reported a claim in the previous period.

relationship between past and future claims. Yet, the mean premium per value does not adjust accordingly. It is flat in both relative and absolute terms. As with claims at the policy level, on average, the insurer exhibits losses when providing coverage to customers with poor past performance, while it enjoys profits by providing insurance to customers that performed well in the past, as they also tend to perform well in the future. As with vehicle age and claims, the insurer's portfolio is profitable as the overwhelming majority of the insurer customers are beneficial —with a loss ratio of under 1, greater than 82% (Panel B of Figure 1.2). This composition is not exogenously determined; the insurer denies customers with poor performance at a higher rate (see section 1.5).

To summarize, the statistics presented in this section raise two opposing findings. On the one hand, the insurer is profitable. The profit margin of comprehensive coverage policies for all vehicles is 19.65%, and 15.97% in the case of trucks. As mentioned above, these margins are satisfactory from the insurer's perspective. On the other hand, the data presented suggest that the insurer can do better. Both predetermined (age) and stochastic factors (past performance) are correlated with a higher cost per value in the future, yet, the relationship between those factors and premium per value is quite flat. On average, the insurer's profits are generated by a specific segment of customers and vehicles. New vehicles and customers with adequate past performance are beneficial, while old vehicles and customers with poor past performance are costly to insure.

1.4 Insurer Pricing and Costs

In this section, I estimate the gap between actual pricing and the expected cost of policies as a function of predetermined and stochastic factors. This section is ordered as follows. First, I analyze the relationship between premium per value, cost per value, and vehicle age. Then, I examine the predictive power of past performance and its relationship with current costs, premiums, and profits. I differentiate between recent and overall past performance by considering the aggregate loss ratio and the previous year's loss ratio. I study whether the recent claim history is more predictive of future costs and the relationship of both with pricing. In addition, I examine the competitors' pricing schemes to assess whether the documented insurer behavior is unique or similar to market-wide patterns. To further establish that the insurer misprices risk, I conduct a few robustness tests to rule out alternative channels. Finally, I provide evidence of firm learning by examining insurer pricing after providing information on the flat pricing scheme over the vehicle life cycle.

Predetermined Changes

In this part, I examine whether the insurer adjusts premiums according to the evolution of costs by vehicle age. Specifically, I quantify the relationship between premium per value, cost per value, and vehicle age. I do so by estimating the following fixed-effect models.

$$\frac{\text{Premium}_{\ell t}}{\text{Value}_{\ell t}} = \sum_{a=1}^{A} \beta_a^p \{\text{Vehicle Age}_{\ell t} = a\} + \eta_\ell^p + \varepsilon_{\ell t}^p$$
$$\frac{\text{Cost}_{\ell t}}{\text{Value}_{\ell t}} = \sum_{a=1}^{A} \beta_a^c \{\text{Vehicle Age}_{\ell t} = a\} + \eta_\ell^c + \varepsilon_{\ell t}^c,$$

where ℓ and t index the license (vehicle) and period, respectively. $\{\text{Vehicle Age}_{\ell t} = a\}$ is an indicator of a specific vehicle age, a, and $(\eta_\ell^p, \eta_\ell^c)$ are vehicle fixed effects (with respect to pricing and realized costs). I estimate the premium and cost-per-value trends over the vehicle life cycle using a saturated model with fixed effects at the license level. That is, I identify and quantify the trends using within-license variation in the premium and cost per value.

The results are depicted in Figure 1.3. The patterns are consistent with the summary statistics provided in the previous section. The premium per value does not change substantially over the vehicle's life cycle; it increases by less than 0.5 p.p over the first seven years of the vehicle's age, and by less than 1 p.p over the first ten years. In contrast, costs increase considerably over the vehicle's life cycle. The cost per value increases by more than 3 p.p over the first seven years of the vehicle's age, and by more than 5 p.p over the first ten years. These patterns are inconsistent with *perfect* insurer behavior; the optimal pricing strategy suggests that premiums should adjust according to changes in the expected cost of providing insurance. Since the expected cost per value increases with vehicle age, so should the premium per value; however, the observed premiums per value are quite flat.[14]

The result described in Figure 1.3 suggests that there is limited variation in pricing within vehicles, but does not imply limited variation in premium per value *across* vehicles. Appendix Figure A.7 shows that this is not the case, however. There is substantial variation in premium per value, as expected when (i) equilibrium premiums are determined in a bargaining process between the insurer and the client

[14] A variant of the model in terms of nominal ILS (instead of per value) is conducted as well. Results are presented in Appendix Figure A.6, Panel A. Furthermore, a replication of the model with regard to the entire sample (all vehicle types) is reported in Appendix Figure A.6, Panel B. The estimated patterns in both variants are consistent with this figure.

and (ii) there is substantial heterogeneity in bargaining power, possibly due to the considerable variation in clients' fleet size.

A possible explanation for the lack of price variation over the vehicle's life cycle is related to client characteristics. As noted earlier, a substantial portion of clients purchase insurance coverage for multiple vehicles. It could be the case that the observed flatness in premiums is artificial. When providing insurance coverage to a large fleet, premiums per value are not expected to change if the client's vehicle age distribution does not vary over time, as fleet owners purchase new vehicles to replace the old ones. Consequently, the lack of variation over vehicle age does not necessarily reflect a lack of adjustment in pricing, as fleet price adjustments might not be necessary. Cross-subsidization within a fleet is an alternative mechanism that can explain the documented trends. An optimally behaving insurer might not change premiums. The cross-subsidization results in clients artificially overpaying to insure new vehicles and underpaying to insure old ones.

I test whether the observed lack of price adjustment is solely driven by fleet cross-subsidization. I do so by examining nonfleet customers. I re-estimate the models considering only nonfleet customers. The results are reported in Figure 1.4. The reported premium per value and cost per value are similar in spirit to those in Figure 1.3. Although the standard errors are larger relative to the entire sample (as expected when considering a smaller sample), the patterns are quite similar. The premium per value barely changes over the vehicle's life cycle, increasing by less than 1 p.p over the first ten years. In contrast, the cost per value increases substantially over the vehicle's life cycle by more than 6 p.p over the first ten years. Therefore, although fleet cross-subsidization might be a complementing factor, it is certainly not the sole determinant generating the patterns in the data.

Cost Shocks

After documenting price misadjustments with respect to a deterministic factor, I examine whether the insurer adjusts premiums optimally when faced with a stochastic cost. Specifically, I quantify the relationship between different current period outcomes and previous period performance.[15] I do so by estimating the following model.

$$Y_{jt} = \beta^c \{\text{Claim}_{jt-1} \geq 1\} + X_{jt}\delta + \varepsilon_{jt},$$

[15] Therefore, I exclude new policies in the following empirical analysis.

where j and t index policy and period, respectively.[16] $\{\text{Claim}_{jt-1} \geq 1\}$ is an indicator for whether policy j was involved in at least one claim event in the previous period. Y_{jt} indicates the current period's four outcomes in question. Specifically, (i) $\{\text{Claim}_{jt} \geq 1\}$, an indicator for whether at least one claim was reported in the current period, (ii) current period damage (net claim cost) per value, (iii) change in premium per value (that is, the ratio of current to previous premium per value, minus one), and (iv) current period policy loss ratio (damage over premium). I control for vehicle characteristics (value, age, weight, type, and a young driver indicator).[17]

The results are reported in Table 1.2, Panel A.[18] Column 1 reports the relationship between the previous and current claim outcomes. The probability of at least one claim in the current period is 12 p.p higher if a claim was reported in the previous period.[19] The results demonstrate that claim history serves as a persistent signal of current performance (t-stat = 15.22), even when considering a relatively naive measure of past performance.[20] Column 2 reports the estimated model with regards to damage per value. Consistent with the findings in Column 1, the damage per value ratio is on average 1.7 p.p higher, compared with a policy that was not involved in a claim event in the previous period.

Column 3 reports the estimated model of the change in premium per value. If the insurer adjusts premiums correctly, the standard model predicts that under optimal pricing, the premiums should increase with claim history, as past performance serves as a predictive signal of current claims. However, the results suggest that this is not the case. Premiums do not significantly differ, and the coefficient is of the wrong sign; the premium per value drops by -0.1 percent when a claim is reported in the previous period. As a result, the policy is less profitable. As expected, given the results on damages and premiums, the loss ratio associated with a reported claim in the previous period is 41.8 p.p point higher (column 4).

Similar to the examination of predetermined changes, fleet cross-subsidization may be an alternative mechanism giving rise to artificial noncorrelation between the

---

[16] I conduct the analysis at the policy level and not the license level (as before). Not doing so would result in a selected sample. Intuitively, an insurance policy that covered a vehicle that incurred a total loss claim in the previous period would be renewed (if at all) in the current period with respect to coverage of a different vehicle.

[17] I do not control for vehicle value when considering the outcome variables damage per value or change in premium per value.

[18] A replication of the model with regard to the entire sample (all vehicle types) in reported in Appendix Table A.2, Panel A. The results are similar.

[19] This is an interpretation based on the linear probability modeling assumption. The relationship is robust with regard to other specifications, such as a logistic and probit model (not reported).

[20] In the next part, I consider the customer's performance (aggregate loss ratio), which takes into account both the cost of damages of all of the customer's policies.

policy's current premium and past performance; clients overpay for policies with good performance and underpay for policies with poor performance. Furthermore, an additional channel that might explain the lack of correlation is that the insurer adjusts prices based on the customer's overall performance—with regard to all policies—and does not assess each policy separately. Regarding vehicle age, I test whether large fleets might give rise to the observed noncorrelation by re-estimating the model when considering solely nonfleet customers. The results are reported in Table 1.2, Panel B. The relationship between current costs (claim indicator and damage per value) and past claims is similar to those reported for the entire sample. In contrast, the results reported in Column 3 indicate that premiums are adjusted based on past performance. When considering nonfleet policies, the correlation between past claims and current premiums is significant and positive. Policies that incurred a claim in the previous period face a 1.7 percent increase in premiums, compared with policies that were not involved in a claim event. Despite the price increase following poor performance, the relationship between current period loss ratio and past performance suggests that the price adjustment is inadequate. The loss ratio associated with a reported claim in the previous period is 54.7 p.p higher. These results illustrate that the insurer is aware of the persistence of claim history, yet does not adjust premiums sufficiently. Thus, fleet cross-subsidization cannot explain the observed patterns.

Recent vs. Older Claim History

In the previous subsection, I study how the insurer adjusts prices with regards to previous period claims. In this part, I examine how the insurer adjusts prices when considering both new information and past signals. In particular, I consider the following two signals: (i) the client's aggregate loss ratio over time and (ii) the client's previous year loss ratio.[21] Optimal insurer behavior implies that premiums should adjust with respect to each of these signals based on their relative predictive power: the signal-to-noise ratio. I quantify the relationship between the two signals and different current period outcomes using the following regression model:

$$Y_{jt} = \alpha \text{Aggregate LR}_{jt} + \beta \text{Prev. Yr. LR}_{jt} + X_{jt}\delta + \varepsilon_{jt},$$

where j and t index policy and period, respectively. The two explanatory variables, Aggregate LR_{jt} and Prev. Yr. LR_{jt} are the two signals: the clients' aggregate loss ratio over time and the clients' previous year loss ratio, respectively. It is important to note the previous year's performance is reflected in both variables. Yet, Aggregate LR_{jt} weighs previous data equally, without considering the recency of

[21] Similar to previously, I exclude policies of customers with less than one year of observed history.

previous period outcomes. Y_{jt} indicates the four outcomes of the current period in question.

The results are reported in Table 1.3. Columns 1 and 2 describe the relationship between the client's aggregate loss ratio over time and the client's previous year loss ratio with current period damages: the probability of at least one claim in the current period and current period damage per value, both at the policy level. The customer's aggregate loss ratio serves as a predictive signal of future claims. The aggregate loss at the start of the policy is positively correlated (statistically significant) with (i) the indicator of whether the policy incurs a claim in the current period (column 1: t-st=5.95) and (ii) the policy's damage per value (column 2: t-st=4.23). In contrast, the customer's previous year loss ratio is not correlated with either cost variable. The coefficient is either in the *wrong* sign, in the case of the indicator of at least one claim in the current period (column 1:-0.0004), or substantially smaller in order of magnitude relative to the aggregate loss ratio, as is the case when considering the damage per value (column 2: 0.0003 relative to 0.0055); in both cases, the coefficients are not statistically significant. The lack of additional information does not imply that recent performance is not informative, but rather that it is not more informative than older claim history. Consequently, a *perfect* insurer should only consider the aggregate loss ratio when adjusting premiums.

In Column 3, I report the estimated model, which quantifies the relationship between the two variables measuring the previous loss ratio and the change in premium per value. The result indicates a deviation from optimal pricing. Unlike an optimally price-setting perfect insurer, the insurer does not increase premiums when facing a high aggregate loss-ratio client. Furthermore, the insurer reacts to negative results, but considers the wrong signal. Premiums increase when considering a customer with poor performance in the previous year, controlling for aggregate performance over time, despite (i) recent performance not incorporating any additional information relative to aggregate claim history, and (ii) the insurer not adjusting premiums based on the more informative signal—the aggregate loss ratio. Column 4 reports the estimated model with regard to the policy's loss ratio. Consistent with misadjustments, policies of clients with higher aggregate loss ratios are associated with adverse results. Yet, conditional on aggregate loss ratios, policies of clients with higher aggregate loss ratios in previous year are not associated with these results. The insurer overreacts to recent noisy shocks (previous year loss ratio) and underweighs the predictive power of the augmented claim history (aggregate loss ratio).

Market Behavior

Thus far, I have provided evidence that the insurer misprices risk. Premiums insufficiently adjust to predetermined changes and stochastic shocks. A possible concern is the external validity of these results. The findings are based on the pricing data of one insurer (although it is one of the largest insurers in the market and affiliated with a large multinational insurance company). Observing that one insurance company systematically misprices risk does not imply that the *market* is *imperfect*. If the competitors price risk correctly, we expect that in the long run, an imperfect insurer would be competed out of the market.[22]

Unlike many cases where it is difficult to observe the pricing of all firms in the industry, I am able to extract prices for a large number of trucks at different ages and in different conditions (e.g., claims) for all major competitors. I address this issue by examining the market competitors' pricing schemes. I do so with data from the Israeli insurance agency *Orlan Insurance Agency, Ltd. (henceforth, Orlan) (1994)*. As part of its business operations, Orlan has ties to the largest insurance companies in Israel (including the insurer from which I obtained the data). To provide competitive insurance premiums to its clients, Orlan's agents can compare premiums (and coverage terms) for new policies across all insurers (in contact with Orlan), as a function of their characteristics. Orlan's agents access the data using the Orlanet Calculator (henceforth, calculator), which provides information regarding offered pricing and terms from each insurance provider.[23]

Using the calculator, I generate fictitious offers for comprehensive insurance policies for 2,041 distinct trucks model-value-year triads insured between January and March 2020.[24] I use standard insurance coverage and vehicle characteristics as additional inputs necessary to generate an offer.[25] I generate two observations for each distinct vehicle model-year-value triad: (i) no claims in the last 3 years and (ii) one claim in the last 3 years, which occurred last year. I focus on the four largest insurers in this market: the insurer that provided me the data (denoted as "the insurer"), and its three main competitors (denoted as "rival 1", "rival 2", and "rival 3").

[22]This statement is true if (i) customers treat insurance coverage as an homogeneous good, (ii) customer search does not incur any costs and (iii) incumbent insurers do not possess an informational advantage over their competitors.

[23]Orlan state that agents should not use the provided dataset in order to price renewing policies, but rather use the calculator solely for new policies.

[24]The data generating process was conducted in the beginning of March 2020, before Israel began enforcing social distancing and other rules to limit the spread of COVID-19.

[25]I.e., vehicles without heavy equipment, default driver characteristics (any driver over the age of 24, excluding individuals with a criminal record or a revoked license), and no additional coverage (e.g., extensive legal defense, riots, earthquakes).

I examine the external validity of my findings in two steps. First, I assess the validity of the calculator. I do so by conducting a within-insurer comparison between the offered premium for coverage through the insurer (using the calculator) and the actual premium charged, to verify that the calculator offers' premiums match the data provided by the insurer.[26] After verification of the calculator's validity, I conduct an across-insurer comparison of the premium offers and examine pricing trends of vehicle age and claim history across the market.[27]

I use the calculator's generated offers to conduct a comparison between the market's insurance providers. I examine the market premium trends of both predetermined changes and stochastic shocks. Figure 1.5 graphs the premium per value trend throughout the age distribution for the four insurers.[28] The results demonstrate that not only is the insurer imperfect, but rather the *market* is *imperfect*. The price trend of the insurer and rivals 1 and 2 is remarkably similar; no trend in premium per value almost throughout the entire age distribution.[29] In contrast, rival 3 raises the premium per value across the age distribution, which suggests that rival 3 adjusts premiums similarly to an optimally pricing firm. Yet, rival 3 is charging higher premiums relative to other competitors. The minimal premium per value charged by any insurance provider (not restricted to the four insurers) is quite flat throughout the age distribution.

In terms of claim history, Table 1.4 presents the relationship between the premium per value offered by each of the four insurers and the minimum premium per value in the market as a function of claim history. The results suggest that, conditional on offering coverage, the insurer is more sensitive to the previous claim history than

[26]The offered premiums do not need to be identical to those provided for a few reasons. Mainly, the calculator suits nonfleet truck owners as it does not take into account customer's fleet size.

[27]Panel A in Appendix Figure A.9 depicts the within comparison. The correlation is very high (R^2=0.90). Policies with a higher offered premium (using the calculator) are on average charged a higher price in practice (using the insurer's dataset). Yet, the coefficient is not 1 (0.8). Panel B in Appendix Figure A.9 graphs the difference between the premiums generated by the calculator relative to the data. The graph indicates that the calculator premiums trend over the vehicle life cycle is steeper than the one shown in my data. This implies that the results in this part might overemphasize the steepness trend in market premiums over vehicle life cycles. As I show, this is not a concern as the across-insurer comparison suggests that the market premiums per value are quite flat with regard to vehicle age.

[28]The sample consists of 876 observations, as some of the rival insurers do not offer coverage to vehicles of specific types and weights.

[29]A possible concern is that the similarities in pricing suggest that the insurer and the other two rivals coordinate premiums. I examine this issue by considering the variation in pricing of the three insurers across different observations. The results are provided in Appendix Figure A.10. The figure demonstrates that, although the trends over vehicle age are similar, the premiums differ substantially within the truck model-value-year triads.

rivals 1 and 2, as they offer the *same* premium per value, regardless of whether the customer reported a claim in the previous year. Such behavior diverges from optimal pricing insurers; the claim history serves as a precise signal for future performance (see Tables 1.2 and 1.3). Thus, the competition also appears to diverge from optimal pricing behavior.

Furthermore, it is impossible to generate a policy offer from rival 3—which is the only insurer to substantially increase premiums per value over the vehicle life cycle—for a customer with a single claim last year (and none in the two years before that). The lack of provided coverage is market wide when considering a customer with at least two claims in the last three years.[30] This pattern suggests that the insurer's adjustment to customer's risk is mostly on the *extensive margin* (i.e., whether to provide coverage at all) and less so on the intensive margin (i.e., increasing premiums); a pattern which the insurance literature has attributed to adverse selection. I document a similar pattern by the insurer in section 1.5.

Robustness

The findings in this section indicate that the insurer (and its market competitors) misprice risk for a significant segment of its customers. Specifically, the insurer does not adequately adjust premiums when considering changes in predetermined characteristics (vehicle age) or stochastic factors (claim history). A potential alternative explanation to the observed patterns in the data is fleet cross-subsidization. As I show, the documented patterns hold when considering nonfleet customers, suggesting that cross-subsidization is not the only mechanism generating the documented patterns. Furthermore, the analysis of market premiums complements this notion, as the offered premiums are generated for nonfleet customers. In this part, I examine two additional alternative explanations: (i) weak predictive power and (ii) dynamic pricing strategies.

Out-of-Sample Prediction

The results so far suggest that the insurer (and market) misprices risk and underestimates the value of both vehicle age and claim history. In this part, I quantify the predictive power of vehicle characteristics and past performance, as weak predictive power suggests that disregarding the information might be optimal.

The analysis is conducted in two steps. First, I use data on comprehensive insurance policies from 2014 to 2018 to estimate a cost function. The cost function

[30] Regulation in Israel limit insurers acquirement of information by allowing them to require new customers to provide information with regard to claim history from the last three years.

is constructed using a regression analysis without any customer or license fixed effects. Although it is possible to generate a more precise cost function, I use a simple regression analysis method to demonstrate that even simple methods can generate beneficial cost estimates for out-of-sample observations. The regression model includes observable vehicle characteristics (age, value, weight, class, and type) and claim history (aggregate loss ratio).

Using the estimated cost function, I divide the 2019–2020 sample into 25 groups based on projected damage per value. Appendix Figure A.11 reports the relationship between the mean predicted damage per value for each of the 25 groups, the premium per value charged, and the actual damage per value. As the figure illustrates, the relationship between the predicted and actual damage per value is almost one-to-one, which illustrates that vehicle and customer observable covariates serve as predictive signals of future claims. The figure further demonstrates the lack of adjustment in premium per value. Consistent with earlier findings, the correlation between premium per value and damage per value is less than one.

Dynamic Complementary

Thus far, I have considered a static framework in analyzing the insurer's mispricing of risk. Yet, the relationship between the insurer and its customers is not static. It could be the case the insurer is aware of the incurred losses yet continues to provide coverage since it believes the customer is profitable in the long run. Excluding coverage or increasing premiums from a costly customer today might lead the insurer to not enjoy future profits when the customer becomes profitable.

In this part, I examine whether dynamic considerations can explain the flat pricing patterns (relative to expected cost) using a relatively simple method. First, using data from 2014 to 2015, I divide the insurer's clients into two groups: the first consists of customers with both an average loss ratio of at least 2 and an average vehicle age of at least 5, with 166 customers in all; and the second consists of all other customers who purchased at least one policy from 2014 to 2015 and are not part of the first group, with 3,170 customers. After classifying the costly clients, I examine how the profits are affected when these customers are dropped. Specifically, I analyze the insurer's ex-post performance if it did not provide any coverage to those customers from 2016 to 2020.

In Appendix Table A.3 I provide summary statistics on the policies of both types of customers from 2016 to 2020. The classified "drop" group consists of 1.5% of the relevant policies. These policies exhibit an average loss ratio of 1. The insurer incurs losses from providing them coverage. Their average profit margin is approximately

–14%. Hence, dynamic complementary cannot be the sole channel generating these patterns.

Learning

Would the insurer change its policy if it was informed of these patterns? Is the insurer aware of these patterns? During my interaction with the insurer, I provided findings similar to those presented in Figures 1.1 and 1.3. In this part, I examine how the insurer reacts when informed that premiums do not sufficiently adjust over the vehicle life cycle.

An analysis of the firm pricing scheme before and after I presented the findings to the insurer serves two purposes. First, this analysis provides an additional robustness test, as it further examines whether the documented pricing scheme reflects optimal behavior *from the insurer's perspective.* If not fully adjusting premiums over the vehicle life cycle is the optimal pricing scheme, then the information is already incorporated into the firm's decision-making. Therefore, information on flat premiums per value should not cause the insurer to react. However, if premiums adjust, then it must be the case that the insurer was not fully aware of the trends, and, thus, did not price optimally. Second, this part is of interest as pre- and post-price trend analysis allows examination of managerial practice in general, and firm learning, on-the-job training, monitoring, and feedback, in particular. I examine how the insurer reacted to the provided information using the following event study model:

$$\frac{\text{Premium}_{\ell t}}{\text{Value}_{\ell t}} = \sum_{a=1}^{A} \Big(\sum_{t=-2}^{-1} \beta_a^t \{\text{Vehicle Age}_{\ell t} \in G_a\} \Big) \\ + \alpha_a \{t > 0\} \{\text{Vehicle Age}_{\ell t} \in G_a\} + \eta_\ell + \varepsilon_{\ell t},$$

where ℓ and t index the license and period, respectively. $\{\text{Vehicle Age}_{\ell t} \in G_a\}$ is an indicator of a specific vehicle age group, G_a (sample is divided to five groups: 0-1,2-4,5-7,8-10, and >10). β_a^t represents the offered premium per value, by age group, before information was given to the insurer, while α_a indicates the premium per value after information was given to the insurer. As before, the analysis includes a license fixed effect.

The results are depicted in Figure 1.6. The patterns for the two years before the event indicate that there is no pretrend. The premium per value does not change substantially over the vehicle's life cycle; it increases by less than 0.2 p.p over the first ten years of the vehicle's life, and by less than 1 p.p over the entire vehicle life cycle. Interestingly, the premium per value after I provided the information differs from beforehand. The insurer increases premium per value by approximately 0.5 p.p over the first ten years of the vehicle's age, and by more than 1 p.p over the entire vehicle life cycle. Yet, the increase in price trend over vehicle age is inadequate. The premium per value trend over the vehicle life cycle is still flatter than that of the cost per value; The insurer should have increased the premium per value for age groups 8–10, and for those >10 by an additional 1 p.p.

The results are mixed. On the one hand, the significant change in premium per value indicates that the insurer mispriced risk beforehand and that the flat pricing scheme was inadequate. On the other hand, the relatively small adjustment indicates that even when the insurer is aware of the patterns, it does not fully adjust premiums. As I further demonstrate in section 1.5, this does not indicate that the insurer misprices risk, but rather that it prefers not to implement its beliefs on the intensive margin. Instead, the insurer prefers to exclude costly policies.

1.5 Insurer Grading

In section 1.4, I show that the insurer (and market competitors) does not adjust premiums optimally when faced with predetermined changes and stochastic shocks. In spite of their richness, the observed prices are insufficient to identify the insurer's beliefs for two reasons. First, the observed premiums are determined in equilibrium by both supply and demand factors. The insurer might not raise the premium charged either because it believes that the expected cost of providing insurance did not change, or because an increase in premiums substantially increases the probability the customer will not renew the policy. In addition, the analysis of premiums and costs is based on a selected sample; the subsample of customers who choose to renew the policy and the policies the insurer selects to renew. Separation might not be exogenous. The sample potentially consists only of those customers who are perceived as profitable by the insurer.

I use the insurer's internal grading documents, the "Go—No Go" grading (see section 1.2) to extrapolate the insurer's beliefs, as "Go" grades are assigned to profitable policies, regardless of demand factors. Furthermore, the "Go—No Go" documents grade all customers, regardless of their decision to renew.

In this section, I describe the "Go—No Go" grading data, provide descriptive evidence on the relationship between insurer grading and observable predetermined and stochastic factors, and estimate how these factors affect insurer's grading.

Summary Statistics

The "Go—No Go" grading documents are provided on a monthly basis by the COO to the insurer's underwriters, who are in contact with the customers. The decoded documents consist of 14,288 grades of policies, providing comprehensive or partial insurance to all vehicle types. The dataset consists of a subsample of the insurer's customers. Summary statistics are provided in Appendix Table A.4. In general, the customers for which I observe grades are more than half of the insurer's portfolio. These customers are different from those with no grade. They are substantially smaller in terms of fleet size (average size of the fleet is 6.09), they are charged a lower premium per value (average of 3.03% relative to 2.32%), they are more profitable (average profit margin of 24.42% relative to 13.01%), and they have a lower loss ratio (average loss ratio of 59.98% relative to 77.10%)

Table 1.5 presents summary statistics on both graded policies and the distribution of insurer grading. In general, the sample consists of profitable policies (the average loss ratio is below 50%). Among the sample, 87% of the policies are given a "Go" grade (column 2). These policies are quite similar in observable characteristics to policies that face a term adjustment or rejection (premium per value, vehicle age, and vehicle value). The substantial difference is the previous year performance; 12% of the policies reported a claim during the policy duration, and the average loss ratio is 29%, while policies that face an increase in deductibles, premiums or were denied incurred losses in the previous year (i.e., loss ratio above 100% in the previous year). These statistics indicate that the overwhelming majority of policies are profitable and do not require adjustments in pricing. The rest of the portfolio requires further consideration.

Among the policies that received a "No-Go" grade, 35% required only an increase in deductibles (column 3), and 21% of the policies required an increase in premiums (column 4). For the remaining 44%, the analytical team recommends not providing comprehensive coverage; either to provide only third-party coverage (14% - Column 5) or to provide no coverage at all (30% - column 6). These statistics clearly demonstrate that a substantial portion of the insurer's adjustment is on the extensive margin, by not providing comprehensive coverage instead of increasing prices.

The statistics also demonstrate that the insurer's grades are based on shocks and predetermined changes. In general, the insurer appreciates the importance of performance and vehicle age. The average loss ratio of policies that the analytical team recommends increasing deductibles or premiums is above 100% and the average loss ratio of policies that the analytical team recommends denying is over 300%. As for vehicle age, the analytical team recommends offering only third-party coverage to old vehicles, regardless of their performance. The average loss ratio among the policies that the analytical team recommends only third-party coverage is 40%—profitable policies—yet the average vehicle age is close to 10 years. Overall, the insurer's beliefs are not entirely misguided: it is aware of the importance of taking into account vehicle age and loss ratio when renewing a policy.

Regression Analysis

In this part, I examine how the insurer grades customers based on their observable characteristics. Specifically, I examine how a "Go" grade is assigned using the following probit model.

$$GO_{ij} = \sum_{a=1}^{A} \beta_a \{\text{Vehicle Age}_{ij} \in a\} + \alpha_{lr} \text{Aggregate LR}_{ij} + \alpha_{recent} \text{Prev. Yr. LR}_{ij} + \delta X_{ij} + \eta_i + \varepsilon_{ij},$$

where i and j index the customer and the policy, respectively. GO_{ij} is an indicator of whether the policy is assigned a "Go" grade. The variables Vehicle Age_{ij}, Aggregate LR_{ij}, and Prev. Yr. LR_{ij} denote vehicle age, the client's aggregate loss ratio over time, and the client's previous year loss ratio, respectively. Furthermore, I consider η_i to denote an unobserved random error at the client level. I divide the sample into four groups based on vehicle age: (i) up to 1 year old, (ii) age 2 to 4 years old, (iii) age 5 to 7 years old, and (iv) age 8 years and above. Controls include observable characteristics, client fleet size, vehicle weight, vehicle type, and underage driver indicator.

The results are presented in Table 1.6. In Column 1, I examine how vehicle age determines a "Go" grade assignment. The insurer discontinuously reacts to vehicle age. Vehicle age affects grading mainly when the vehicle age at least 8 years. This is in contrast to the findings in section 4, which suggest that damage per value increases continuously throughout the vehicle life cycle. Therefore, the results indicate that the insurer erroneously reacts to predetermined factors using simplistic rule-of-thumb rules. Regarding claim history, the insurer considers past performance in grading. The results in Column 2 demonstrate that a client's aggregate loss ratio over-time is significantly correlated with the "Go" assignment. Interestingly, the client's previous year loss ratio has an additional effect on the "Go" assignment, despite it not having any predictive power of future claims, conditional on augmented claim history. This is consistent with the relationship with premiums as well (as documented in section 4). These patterns do not change when estimating the effect of all variables jointly (column 3). To further examine whether the recent year loss ratio has a substantial and persistent role in grading, I re-estimate the model while considering a subsample of policies: those with at least five years of data observed by the insurer on the client. The results are presented in Column 4.

Even though considerable information on the customer is available, the insurer does not underweigh the importance of the previous year loss ratio. This indicates that the recency bias persists and does not diminish with the insurer's data history on the client.

To conclude, the results in this section suggest that the insurer does not change terms for the overwhelming majority of policies, while extensively denying comprehensive coverage from a substantial portion of policies requiring an adjustment in terms. In general, the insurer is aware of the impact of vehicle age and claim history in predicting future performance. However, it mistakenly regards vehicle age in a discontinuous fashion and outweighs the importance of recent performance.

1.6 Demand for Renewal

In this section, I develop and estimate a model of customer policy renewal and cost realization. The goal of this section is to identify and quantify (i) the cost of comprehensive provided coverage, (ii) customers' willingness to pay for policy renewal, and (iii) customers' private information on risk and its relation to demand.

Model Setup

At the end of the policy contract, a customer can choose whether to renew a comprehensive coverage policy. The customer's net utility from policy renewal equals the difference between the utility from renewing the policy with the insurer, u^1, and the outside option, u^0:

$$U = \underbrace{\upsilon(d) - \alpha p}_{u^1} - \underbrace{\max\{\upsilon(d) - \tilde{\alpha}\tilde{p} - k + \lambda\mathcal{I}, 0\}}_{u^0}, \tag{1.1}$$

where $\upsilon(d)$ is the utility from comprehensive coverage, which depends on the policy's expected damage, d. $\upsilon'(d) > 0$ implies that consumers are adversely selected. p is the premium per value charged by the insurer, while $\tilde{p}$ is the expected premium per value charged by the market competitors. k reflects the costs incurred by the customer when searching for offers among market competitors. $\mathcal{I}$ is a dummy variable that equals 1 if the customer faces a price increase. $\lambda\mathcal{I}$ represents a customer's rational inattention. Intuitively, the consumer becomes aware of the outside option and might decide to search and reoptimize when faced with a price increase.[31] Rational inattention is modeled as an additional component of search cost. The standard search cost generates power for the insurer market. This is true for rational inattention as well. Unlike the standard search cost, rationally inattentive consumers who become attentive following a price increase represent a demand-side friction that might drive a profit-maximizing insurer to provide a uniform pricing scheme.

Econometric Model

I specify the net utility from policy renewal as a linear function of customer and vehicle observable characteristics, premium per value, the previous period premium per value, and expected damage per value.

$$U_{ijt} = -\alpha p_{ijt} - \lambda \mathcal{I}_{ijt} + \beta x_{ijt} + \gamma d_{ijt} + \psi_{ij} + \omega_i + \varepsilon_{ijt},$$

where i, j, and t index customer, vehicle, and period, respectively. p_{ijt} is the premium per value. $\lambda \mathcal{I}_{ijt} = \lambda \times \{p_{ijt} > p_{ijt-1}\}$ represents customer sensitivity to price increase. x_{ijt} are vehicle and client observable characteristics. d_{ijt} is the policy's expected damage per value. γ represents customer selection based on expected damage. Specifically, $\gamma > 0$ implies adverse selection, as customers with a higher utility from renewing their comprehensive insurance policy tend to cost more. ω_i and ψ_{ij} represent the customer-level and license-level unobserved demand components which are fixed over time, respectively. ψ_{ij} is the unobserved license-level demand component, which might be correlated with premiums, while ω_i is an exogenous unobserved client-level demand component; ω_i is normally distributed, $\omega_i \sim N(0, \sigma^2_\omega)$. ε_{ijt} represents variation in unobserved demand factors across the client's vehicles and over time, which follows a logistic distribution.

Equation 1.1 suggests that the net utility from policy renewal is non-linear. Some customers search among market competitors, while others decide not to purchase comprehensive insurance at all. Data limitations do not permit determining whether the customer decides to purchase a policy from a competitor. Furthermore, I do not obtain competitor's pricing scheme for all types of vehicles (specifically, all non-truck vehicles). Nevertheless, competitor-offered premiums are a function of observable characteristics, all observed by the econometrician. Therefore, βx_{ijt} takes into account search costs k, utility from renewal $\upsilon(d)$, and competitor's pricing scheme $\tilde{p}$.

[31] In this model, I consider rational inattention, but unlike Ho, Hogan, and Scott Morton (2017), I do not estimate the increase in premiums necessary to cause the customers to be attentive. Instead, I set that increase to zero, which match both the uniform pricing in the data and the recommendation regarding a "Go" grade assignment.

Regarding the expected damage per value, I adopt a specification similar in spirit to Einav et al., 2013. Damage per value follows a pseudo-Poisson distribution according to the following exponential expected value:

$$d_{ijt} = \exp(\delta x_{ijt} + \nu_i), \tag{1.2}$$

where ν_i represents customers' private information regarding risk, which is constant over time.[32] ν_i is normally distributed, $\nu_i \sim N(0, \sigma_\nu^2)$, independent of observable characteristics x_{ijt}. This is an appropriate modeling fit as the damage per value is heavily skewed, conditional on claim, and characterized by a significant number of small-scale claims in terms of damage per value (see Appendix Figure A.12). The unobserved constant structure implies that the damage per value varies over time, yet the unobserved component is fixed, as in Einav et al., 2013. This modeling assumption allows taking into account policies that are not up for renewal (for instance, ones that exhibited a total loss event) when estimating Equation 1.2.

The logistic distribution assumption implies that the probability of renewal, denoted by $R_{ijt} = 1$, is described using the standard logit model.

$$\Pr(R_{ijt} = 1 | X_{ijt}, d_{ijt}, \omega_i, \psi_{ij}) = \frac{\exp(-\alpha p_{ijt} - \lambda \mathcal{I}_{ijt} + \beta x_{ijt} + \gamma d_{ijt} + \psi_{ij} + \omega_i)}{1 + \exp(-\alpha p_{ijt} - \lambda \mathcal{I}_{ijt} + \beta x_{ijt} + \gamma d_{ijt} + \psi_{ij} + \omega_i)}$$

Identification

I discuss how variation in the data identifies the model. Identification of damage per value is established in the literature.[33] The effect of observable characteristics on expected damage per value—δ—is identified by variation in the realized damage per value across clients' and vehicles' characteristics, as in Cohen and Einav (2007), Bundorf, Levin, and Mahoney, 2012. Identification of heterogeneity in clients' private information of cost—σ_ν—is established by the within-client correlation across the client's different policies, as in Einav et al. (2013).

[32] Cohen and Einav (2007) use a Poisson distribution to fit the policy's number of claims. This implies that the claim process is both state-independent and independent of conditional damage. Since the main goal is to estimate the relationship between consumer demand and the insurer's cost of providing coverage, I use a pseudo-Poisson distribution which accommodates both (i) the possibility of no damages at all, which occurs quite frequently, and (ii) possible dependence between the number of claims and the conditional damage of claims. In practice, I take into account the duration of each policy by estimating a pro-rated variant of Equation 1.2: $d_{ijt} = \exp(\delta x_{ijt} + \nu_i) \times \tau_{ijt}$, where τ_{ijt} is the duration of the policy. See Appendix B for extensive discussion of the modeling assumptions and estimation process.

[33] See Einav, Finkelstein, and Levin (2010a).

The main identification challenge is to estimate consumers' willingness to pay with observational data. I observe adjustments in premiums between consecutive coverage periods for a non-random sample of vehicles. Premiums are determined in equilibrium by both supply and demand forces. I address this issue in two steps. First, premiums might be correlated with the vehicle-specific unobserved component, ψ_{ij}. I take advantage of the population of clients, that contains a large number of fleets with multiple vehicles over multiple coverage periods and treat license-specific unobserved determinants of premiums as fixed rather than random. In the spirit of Crawford, Pavanini, and Schivardi (2018), I estimate the following (log) premium equation.

$$\log(P_{ijt}) = \Omega X_{ijt} + F_{ij} + \zeta_{ijt}, \tag{1.3}$$

where F_{ij} reflects the license-specific determinant of the log premium charged by the insurer. I use the license fixed-effect estimated in the premium model above as a proxy for the constant demand unobservable, $\psi_{ij} = \rho F_{ij}$, to control for the license-specific constant-term which endogenously sets the premium.

Another concern is that the change in premium charged over time is correlated with unobserved demand factors that vary within client or vehicle, over time. I deal with this challenge using the fact that premiums are explained remarkably well by the premium equation above—Equation 1.3—R^2=0.98. The panel structure allows me to use predicted, rather than actual, premiums and impute premiums for those who renew their policy, as in Bundorf, Levin, and Mahoney, 2012, as well as for those who do not, as in Crawford, Pavanini, and Schivardi, 2018. Specifically, I consider the following renewal probability:

$$\Pr(R_{ijt} = 1 | X_{ijt}, d_{ijt}, \hat{p}_{ijt}, \hat{f}_{ijt}, \omega_i) = \frac{\exp(-\alpha \hat{p}_{ijt} - \lambda \hat{\mathcal{I}}_{ijt} + \beta x_{ijt} + \gamma d_{ijt} + \rho \hat{f}_{ijt} + \omega_i)}{1 + \exp(-\alpha \hat{p}_{ijt} - \lambda \hat{\mathcal{I}}_{ijt} + \beta x_{ijt} + \gamma d_{ijt} + \rho \hat{f}_{ijt} + \omega_i)}, \tag{1.4}$$

where $\hat{f}_{ijt} = \frac{\hat{F}_{ij}}{V_{ijt}}$ and $\hat{p}_{ijt} = \frac{\hat{\Omega} X_{ijt}}{V_{ijt}}$ are the predicted premium license-specific constant term and time-varying term, respectively. Identification of both price sensitivity, α, and sensitivity to price increase, λ, are established using variation in the predicted time-varying premium component, $\hat{p}_{ijt}$, and its relationship to the previous period predicted value term, $\hat{\mathcal{I}}_{ijt} = \{\hat{p}_{ijt} - \hat{p}_{ijt-1} > 0\}$. As with cost, the effect of observable characteristics on renewal—β—is identified by variation across clients' and vehicles' characteristics, and heterogeneity in clients' demand component (which is independent of premiums)—σ_ω—is established by within-client correlation across the client's different policies, as in Einav et al. (2013). Last, the selection parameter, γ, is identified using within-client variation in observable characteristics that determine the damage per value and explain renewal. To illustrate this, consider a vehicle characterized by a low mean expected damage per value at the current period and a high mean expected damage per value at the next period. Adverse selection implies that the willingness to pay for insurance will increase with private information regarding cost; and the increase in damage per value between periods for costly clients with a high probability of a claim event will be higher than customers with low probability of a claim.

Estimation proceeds in three steps. Initially, since some policies are not renewed, I do not observe their vehicle value. I predict the (log) vehicle value as a function of the previous period value and vehicle type, and treat it as data; both for those that renewed and those that did not. The new vehicle value is well explained by the preceding one (see Appendix Table A.6). Then, I generate the predicted premium per value and predicted increase in premium per value using the estimates of Equation 1.3 - (see Appendix Table A.7). Finally, I jointly estimate Equations 1.2 and 1.4 via Maximum Simulated Likelihood, similar to the approach in Train (2009). I estimate the parameters in the policy renewal and damage equations maximum likelihood using the observable client and vehicle characteristics x, predicted premium per value $\hat{p}$, dummy indicator for an increase in predicted premium per value $\mathcal{I}$, and realizations regarding damage per value and policy renewal. To estimate σ, I use 200 Halton draws for each client: 100 with respect to the unobserved demand component and 100 for private information regarding cost (as in Train (2000)). I then exploit the normal distribution density to derive the likelihood function. See Appendix B for more details.

Results

The estimation results are provided in Table 1.7.[34] The cost estimates regarding vehicle age and claim history are similar to those documented in the reduced-form analysis. Expected damage per value increases sharply with vehicle age, and the aggregate loss ratio also plays a significant role in predicting future claims.

In contrast, the previous year loss ratio does not provide any additional information, conditional on augmented data. In addition, new customers are adversely selected; customers who joined the insurer in the last year tend to cost more, conditional on observables. These results demonstrate that the incumbent insurer enjoys a comparative information advantage, as past performance is not shared across the market and competitors observe only recent claims.

[34]The estimated coefficient fit both the damage per value distribution and renewal probability quite well. See Appendix Figure A.13.

As for demand, the estimation of the structural model indicates that customers adversely select to renew policies. Private information regarding customer cost of coverage is associated with a higher renewal probability. However, it should be pointed out that asymmetric information is quite modest.

Price elasticity estimates suggest that customers are sensitive to premiums charged. Moreover, they are very sensitive to an increase in premiums per value (relative to previous period). The coefficient suggests that an increase in premium per value is equivalent to an additional increase of 0.4 p.p in premium per value, which is about 10 percent of the average premium per value charged. The demand-side friction both generates the insurer market and incentivizes a uniform pricing scheme.

1.7 Insurer Pricing and Cost

In this part, I develop and estimate a model of insurer policy assessment and supply of coverage. The goal is to identify and quantify the insurer's subjective expected cost of providing coverage and the process of choosing between increasing premiums and rejection, when an adjustment in terms is perceived as necessary.

Model Setup

When the policy contract is about to end, the insurer has three options with regard to policy i: (i) "Go" - the analytical team recommends renewing the policy with the same premium as previous period; (ii) "Adjust" - the analytical team recommends renewing the policy but not in the current terms, by either increasing premiums and/or deductibles; and (iii) "Reject" - the analytical team recommends not to provide comprehensive coverage. The profit margin, denoted by Π_i under each alternative, denoted by A_i, is defined as follows:

$$\Pi_i = \begin{cases} \Pi_i^g = 1 - \frac{d_i}{p_{-1i}} & \text{if } A_i = \text{Go} \\ \Pi_i^a = q_i(\Delta p_i) \times \left(1 - \frac{d_i}{\tau(\Delta p_i) \times (p_{-1i} + \Delta p_i)}\right) & \text{if } A_i = \text{Adjust} \\ 0 & \text{if } A_i = \text{Reject}, \end{cases}$$

where d_i is the insurer's perceived damage per value of policy i. p_{-1i} is policy i's previous period premium per value. Δp_i is the optimal increase in premium per value by the insurer when adjusting prices. $q_i(\Delta p_i) \in [0, 1]$ is the renewal probability in case of an increase in premium of $\Delta p_i > 0$.[35] $\tau(\Delta p_i) \in [0,1]$ represents exerted costs incurred by the insurer in case of an adjustment in terms.

The recommendation is intended to maximize the expected profit margin. Therefore, the optimal strategy is equivalent to selecting the recommendation that maximizes π_i, defined as:

$$\pi_i = \begin{cases} \pi_i^g = \log(p_{-1i}) - \log(d_i) & \text{if } A_i = \text{Go} \\ \pi_i^a = -\log(z_i) & \text{if } A_i = \text{Adjust} \\ 0 & \text{if } A_i = \text{Reject}, \end{cases} \tag{1.5}$$

where $z_i = 1 - q_i(\Delta p_i) \times (1 - \frac{d_i}{\tau(\Delta p_i) \times (p_{-1i} + \Delta p_i)})$.

Before describing the econometric model, I illustrate the implications of a high adjustment cost on the supply of insurance. Adjustment costs affect customers as an increase in $\tau(\Delta p_i)$ reduces π_i^a. Therefore, adjustment cost both reduce the probability of a price increase, and increase the probability of denial of coverage. These outcomes jointly determine the effect of supply-side frictions on consumer surplus. When π_i^g is sufficiently large, consumers would benefit from these frictions. In contrast, customers who the insurer perceives as costly, i.e., low π_i^g, are less likely to be offered coverage. Furthermore, supply-side frictions might negatively affect low-cost customers as well; that is, a good customer might be involved in a claim event. As a result, customers with high volatility in outcomes are prone to denial, while customers with modest volatility are not.

Econometric Model

The insurer's decision can be re-expressed as determined by the customer's observable factors $X_i = (p_i, x_i)$.

$$\pi_i = \begin{cases} \pi_i^g = \underbrace{\log(p_{-1i}) - \log(d(X_i))}_{\pi^g(X_i)} + \varepsilon_i^g & \text{if } A_i = \text{Go} \\ \pi_i^a = \underbrace{\log(z(X_i))}_{\pi^a(X_i)} + \varepsilon_i^a & \text{if } A_i = \text{Adjust} \\ 0 & \text{if } A_i = \text{Reject}. \end{cases} \tag{1.6}$$

Identification of the model requires independence between the observed and unobserved factors. Formally:

$$\varepsilon \perp X.$$

[35] This implies that I am normalizing the probability of renewal to be one in case a "Go" grade is given. Allowing the renewal probability to be lower than one does not change the analysis.

This assumption holds if the insurer does not have any informational advantage relative to the econometrician, i.e., no omitted variables. The rationale for the identification strategy of the structural model is based on the fact that I observe *all* of the information documented by the insurer, which implies informational symmetry. The main challenge in identifying π is that unobserved demand factors might be correlated with the client's and vehicle's observable characteristics. These include undocumented "soft information" (Crawford, Pavanini, and Schivardi (2018)) and strategic factors. Specifically, (i) the insurer might know more about the customer's willingness to pay than the econometrician—as reflected by client's and vehicle's observable characteristics (including premiums)—and (ii) strategic considerations, which include fleet cross-subsidization and possible marketing incentives, as providing coverage to a large fleet might serve as an advertisement that attracts new customers. As in the reduced-form analysis, I address these challenges by focusing on non-fleet customers.

As a result, $\pi^g(X)$ and $\pi^a(X)$ are nonparameterically identified (Berry (1994)). Furthermore, since a "Go" grade assignment implies a recommendation to renew a policy with the same premium per value, identification of $\pi^g(X)$ permits identification of the insurer's risk assessment, both in terms of level (group) and slope (selection).

$$\begin{aligned}\pi^g(X) &= \log(p_{-1}) - \log(d(x,p)) \\ &= \log(p_{-1}) - \underbrace{\log(d(x,\bar{p}(x)))}_{\text{Group cost}} - \Big(\underbrace{\log(d(x,p)) - \log(d(x,\bar{p}(x)))}_{\text{Selection}}\Big),\end{aligned}$$

where $\bar{p}(x)$ is the average premium per value charged by customers of observable characteristics x. Given the model specification, the perceived cost function is identified using variation in characteristics across groups and within variation in premiums. Variation in group observable characteristics, x and $\bar{p}(x)$, identifies the insurer's perceived cost of providing comprehensive coverage to the mean customer of group x. Variation within groups in premium charged identifies the perceived selection by the insurer. The one-to-one relationship between $\pi^g(X)$ and $\log(d(x,p))$ demonstrates that identification of the perceived cost function depends on the significant proportion of policies assigned a "Go" grade, for which demand forces do not play a role. If all policies are assigned either an "Adjust" or "Reject" grades, it is not possible to identify the perceived cost. To illustrate this, consider two policies: one that is assigned an "Adjust" grade, and the other a "Reject" grade. I cannot identify whether the rejected policy is denied coverage due to supply forces, i.e., a higher cost of providing coverage, or demand, i.e., a lower willingness to pay. The same holds for two policies with different recommendations regarding the magnitude of the increase in premiums. The offset of demand forces when a "Go" grade is assigned allows for extraction of the insurer's subjective risk assessment.

While the data is rich enough to extract the insurer's perceived costs, it is insufficient to separately identify demand forces and supply frictions determining denial. To illustrate this, consider that for any price increase from p to $\tilde{p}$, the insurer's profit margin, including adjustment cost, is $q(\tilde{p}-p) \times (1 - \frac{d}{\tau(\tilde{p}-p)\times\tilde{p}})$. Identification of adjustment costs, $\tau(\cdot)$, is only possible when the renewal probability is (or approaches) one. Yet, both data limitations, specifically as the number of recommended adjustments are quite scarce and consumers are sensitive to a price increase, do not allow quantifying both channels using, for instance, an identification at infinity approach.[36]

Despite the non-parametric identification, the estimation is made using some parametric assumption. Similar to demand estimation, I assume that the expected damage per value of policy i is defined using an exponential term.

$$d_i = \exp(\tilde{\delta} x_i + \tilde{\rho}(p_i - \bar{p}(x_i))$$

An examination of both $\tilde{\delta}$ and δ permits a comparison between the objective and perceived determinants of damage per value. As a result, $\pi^g(X)$ is expressed using a linear term, while $\pi^a(X_i)$ is estimated using a linear function. Furthermore, I assume $\varepsilon_i^g = \sigma(\epsilon_i^g - \tilde{\epsilon}_i)$ and $\varepsilon_i^a = \sigma(\epsilon_i^a - \tilde{\epsilon}_i)$. $(\epsilon^g, \epsilon^a, \tilde{\epsilon})$ are i.i.d and follow a Type-I extreme value distribution. The model is estimated using a multinominal logit model.[37]

$$\begin{aligned}
\pi_i^g &= \frac{1}{\sigma}(\log(p_{-1}) - \tilde{\delta} x_i + \tilde{\rho}(p_i - \bar{p}(x_i)) + \epsilon_i^g \\
\pi_i^a &= \frac{1}{\sigma}(\beta_0 \log(p_{-1}) + \beta_x x_i + \beta_p \bar{p}(x_i)) + \epsilon_i^a \\
\pi_i^r &= \epsilon_i^r
\end{aligned}$$

Results

The results are presented in Table 1.8. In the first column, I examine how vehicle age determines a "Go" grade assignment. The insurer discontinuously evaluates vehicle age. The insurer does not consider vehicle age, as long as it is less than 8 years, while not assigning a "Go" grade for a substantial portion of old vehicles, defined as age 8 and above. This result is consistent with reduced-form analysis

[36] The recommended increase in premiums is quite discrete as well. Distribution of recommended price increases can be found in Appendix Figure A.14.

[37] This implies that identification of the subjective cost function components is independent of $\pi^a(X_i)$'s function form.

findings of the "Go" grade assignment (see Table 1.6), while it is in contrast to the estimation results of the cost function in Table 1.7, which suggests that the cost of providing coverage increases throughout the vehicle life cycle.

With regard to past performance, the insurer perceives claim history as a predictive signal regarding future claims. The probability of a "Go" grade assignment is less for policies with a reported claim in the previous period and a higher aggregate loss ratio. As documented in Table 1.6, the insurer places additional emphasis on the previous year's performance, despite the fact that estimation results of the cost function (see Table 1.7) indicate that this is not a predictive signal. The overweighing of the previous year loss ratio affects customers heterogeneously, depending on their fleet size. Single-fleet customers are exposed to substantial volatility in their performance. A good driver might incur high damages. Large fleets are less exposed to this risk; as fleet size increases, the probability of extreme events diminishes, suggesting that the insurer assessment and pricing scheme might be advantageous for large fleets, yet disadvantageous for a single-fleet customer.

Lastly, the coefficient regarding selection $\tilde{\delta}$ indicates that the insurer perceives substantial adverse selection. Customers paying more than the average premium paid, based on its observable characteristics, have a lower probability of their policy being assigned a "Go" grade, although cost estimates suggest modest private information.

In Figure 1.7, I present a comparison between the insurer's subjective determinants of cost, $\tilde{\delta}$, and the objective cost function δ. The figure indicates that the insurer underweighs factors such as new customer indicator and vehicle age groups below 8 (age 2–4 and 5–7). In contrast, the insurer overweighs other factors, including the previous year's loss ratio (although it has no predictive power), the aggregate loss ratio, especially that of a loss-generating clients (i.e., loss ratio above one), and vehicle age, if it is at least 8 years old.

To summarize, the insurer is aware of the importance of vehicle age and claim history as determinants of future performance. Yet, it misevaluates vehicle age in a discontinuous fashion and overweighs the importance of recent performance. Two key conclusions emerge. First, the results of this analysis further demonstrate the importance of considering not only the intensive margin but also the extensive margin. Second, the insurer's biased assessment harms disadvantaged customers; that is, those that purchase comprehensive coverage for a single, old vehicle.

1.8 Counterfactual Analysis

Using the structural estimates of the demand for policy renewal and the insurer's behavior implied by the "Go—No Go" grading, I conduct a few counterfactuals to assess the implications of supply-side behavioral frictions. The analysis is conducted under the assumption that market competitors do not respond to changes in the insurer's behavior, as I do not estimate the cross-substitution patterns between the market competitors' pricing schemes and net utility from policy renewal. The counterfactual analysis is conducted by drawing 200 Halton values: 100 values for the private information regarding cost and 100 values for unobserved demand factors.

I start by examining the premium per value charged for providing coverage for the average truck owner, who was charged a premium per value of 3.5 p.p in the previous period. The results are presented in Panel A of Figure 1.8. I present three different pricing schemes. The blue curve is the optimal pricing scheme, that is, the profit-maximizing premiums. The red curve is the optimal pricing scheme, while restricting the premiums per value does not change throughout the possible states. The green curve is implied by a structural estimation of the insurer's behavior (see Table 1.8). The results in Figure 1.8, Panel A indicate that asymmetric information is quite modest, as the average truck owner is not denied coverage by either an unrestricted or uniform-price restricted profit-maximizing insurer throughout the vehicle life cycle. As with expected damage per value, a profit-maximizing insurer increases premiums per value with vehicle age. In addition, a profit-maximizing insurer does not increase premium per value for a new vehicle. This demonstrates how rationally inattentive consumers might lead profit-maximizing insurers to provide a somewhat uniform pricing scheme.

The insurer behavior implied by the "Go—No Go" grades differs substantially from that of a profit-maximizing insurer. The insurer does not change premiums for vehicles up to the age of 8, and then rejects the policy. The rejection occurs despite the limited selection, suggesting that the insurer excessively denies customers, relative to a rational insurer facing customers with private information of cost. It is not adverse selection that generates rejection, but rather firm practices; under-adjustment of the intensive margin spills over to the extensive margin. The insurer forgoes profits. In particular, the insurer could have increased profits by 7 percent if it had acted as a profit-maximizing firm.

In Panel B of Figure 1.8, I examine the pricing scheme for a single-fleet truck owner following different claim realizations. I consider the observable characteristics of the average customer, characterized by a loss ratio of 70 percent. Note that a single-vehicle customer with an expected loss ratio of 70 percent, who is a profitable customer on average, exhibits a loss ratio of 200 percent or above at a probability of 11 percent. That probability diminishes substantially when considering a fleet of vehicles with the same loss ratio.

As in the case of a new vehicle, optimal pricing does not change following a significant positive realization (loss ratio below 50 percent). The premium per value does change substantially, however, after a negative realization. The premium almost doubles after an outcome of loss ratio of 200 percent, or above. As mentioned above, asymmetric information is quite modest, as the average single-fleet truck owner is not denied coverage by an unrestricted or uniform-price restricted profit-maximizing insurer following any claim realization.

As with vehicle age, the insurer behavior implied by the "Go—No Go" grades extensively deviates from that of a profit-maximizing insurer. The insurer does not change premiums for a loss ratio below 200 percent, yet rejects the policy when it exceeds that loss ratio, despite the limited selection, suggesting that the insurer excessively denies customer, as with old vehicles. The insurer forgoes profits of 16 percent by deviating from profit-maximizing behavior. Furthermore, the net consumer surplus from facing a behavioral insurer, relative to a profit-maximizing firm, is negative. This is not surprising as the probability of denial by the behavioral insurer is higher than the probability that only the profit-maximizing firm would increase premiums. The insurer's biased assessment regarding the recent claim history and the lack of adjustment on the intensive margin harms disadvantage customers—those that purchase comprehensive coverage for a single vehicle. The probability of facing a denial drops substantially with fleet size, suggesting the demand is increasing return to scale. Identical single-fleet customers benefit from purchasing coverage as a whole, independent of price bargaining or risk pooling incentives.

1.9 Conclusion

A cornerstone in the research on risk and insurance is that providers price correctly. In this paper, I inquire whether this is the case. Using data from the one of the largest Israeli commercial auto insurance providers, I find there is too little adjustment in the intensive margin. Premiums barely change with expected costs as projected by predetermined factors (vehicle age) and signals (claim history). Furthermore, I find there is too much adjustment in the extensive margin; that is, an excessive denial of insurance following a negative realization. Using unique grading documents, I integrate the insurer's subjective risk assessment into the study of selection markets, in general, and insurance markets, in particular. I find that the insurer's risk assessment overweighs recent claims and misevaluates vehicle age. Structural model estimates suggest that insurers enjoy incumbency advantages over their own customers, and clients are rationally inattentive to competitors' pricing unless they are faced with a price increase. Both channels allow suboptimal behavior to persist. Finally, I find that supply-side behavioral frictions, which result in excessive denial, diminish with a client's fleet size. This implies that disadvantaged single-vehicle owners are harmed by supply-side frictions, while purchasing insurance coverage as a whole dilutes those losses and might even generate benefits. The spillover of the lack of intensive-margin adjustment on the extensive margin raises important concerns regarding policy intervention.[38]

The results in this paper document the importance of considering both behavioral frictions in selection market analysis and implementing IO structural tools in behavioral economics. The insurer's subjective beliefs regarding the cost of providing coverage differ from an objective assessment. Moreover, the implied behavior suggests overadjustment on the extensive margin and underadjustment on the intensive margin, relative to that implied by state-of-the-art IO analysis. With regard to behavioral economics, examining solely premiums might be quite misleading. When only considering the intensive margin, one might erroneously conclude that the insurer does not take into account observable characteristics when assessing risk. The IO setting, which considers both the intensive and the extensive margins, is essential for identifying and quantifying the effect of biased beliefs.

[38]See Einav, Finkelstein, and Mahoney (2021) regarding the equity and fairness of price variation.

1.10 Figures

Figure 1.1: Summary statistics of premium and cost per value by vehicle age

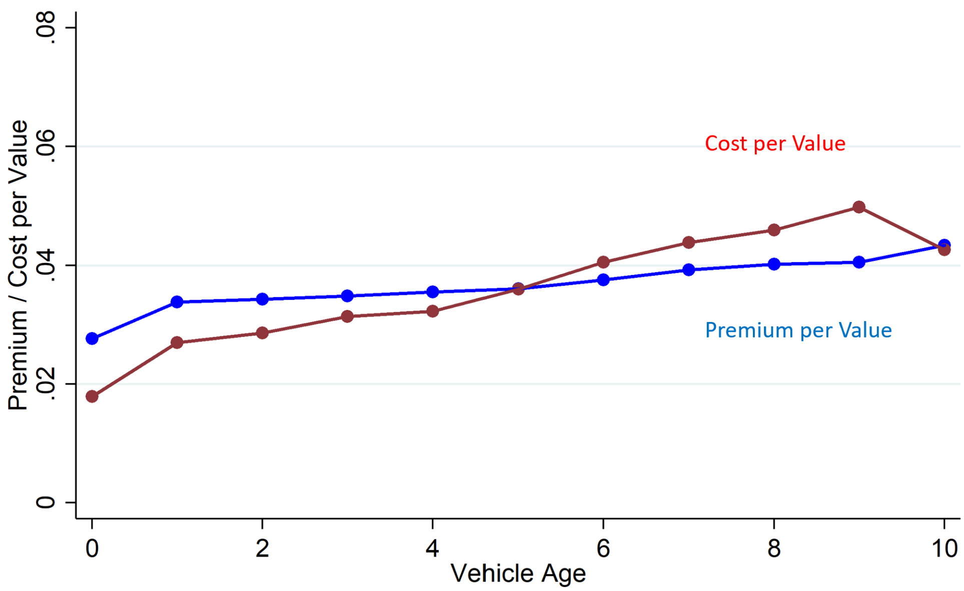

Notes: The figure describes the relationship between premium per value, cost per value, and vehicle age. The vertical axis depicts the mean premium per value (in blue) and costs per value (in red) of comprehensive coverage policies for trucks from 2013 to 2020. The horizontal axis depicts the vehicle's age. No controls are added. Both variables are standardized to an annual term policy. Premiums, costs, and vehicle values are measured in New Israeli Shekel (ILS).

Figure 1.2: Summary statistics of client loss ratio

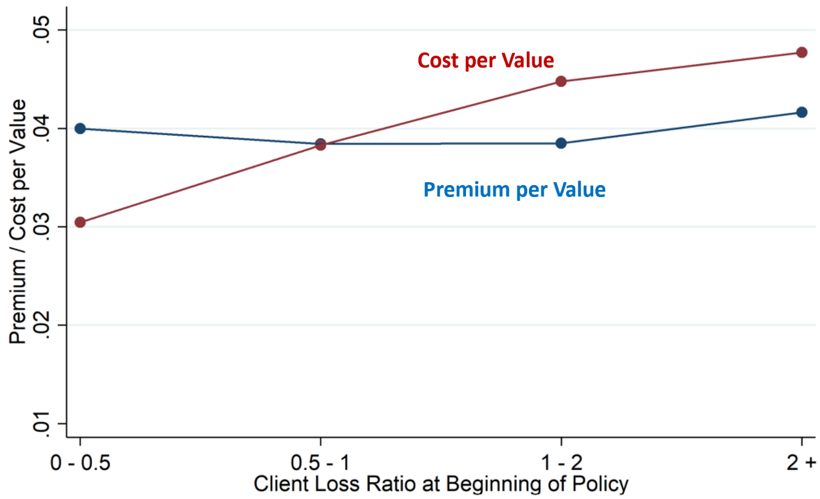

Panel A: Premium and cost per value by client loss ratio

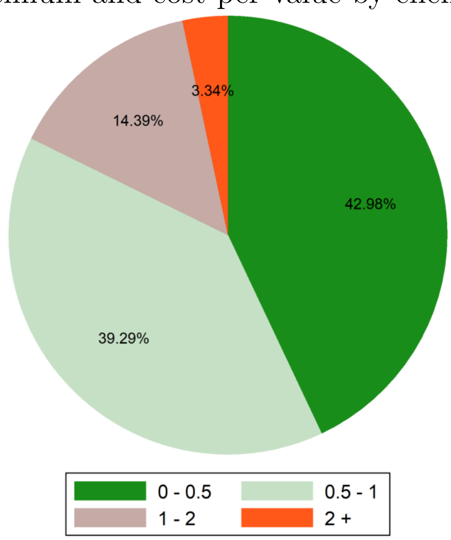

Panel B: Distribution of policies by client loss ratio

Notes: The figure describes premium per value and cost per value by client loss ratio and the sample distribution of client loss ratio. In Panel A, the vertical axis depicts the mean premium per value (in blue) and cost per value (in red) of comprehensive coverage policies for trucks from 2014 to 2020. The horizontal axis depicts the vehicle's age. No controls are added. Both variables are standardized to an annual term policy. Premiums, costs, and vehicle values are measured in New Israeli Shekel (ILS). The sample is divided into four groups: "0–0.5" client loss ratio (aggregate damage over aggregate premium) is up to 0.5, "0.5–1" client loss ratio is above 0.5 and below 1, "1–2" client loss ratio is above 1 and below 2. and "2+" client loss ratio is above 2. Panel B reports the distribution of the sample among the four groups of clients past performance.

Figure 1.3: Premium and cost per value by vehicle age

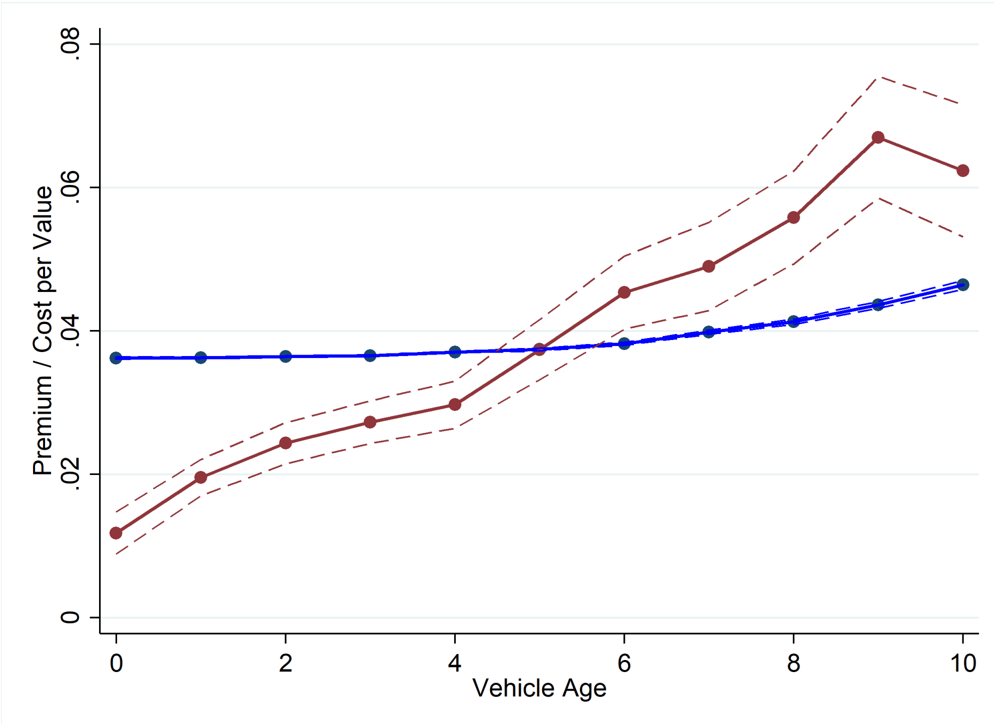

Notes: The figure reports the estimation results of a fixed-effect (license level) saturated regression of premium per vehicle value (in blue) and cost per vehicle value (in red) on vehicle age. Each vehicle age, from 0 to 10, has a unique coefficient. The vertical axis depicts the two dependent variables. The horizontal axis depicts vehicle age. The solid lines represent the regression coefficients. The dashed lines depict the 95% confidence interval. The confidence interval is constructed using robust standard errors clustered at the client level. The sample includes comprehensive insurance coverage policies for trucks from 2013 through 2020. Premiums and costs are normalized to an annual policy length. Premiums, costs, and vehicle values are measured in New Israeli Shekel (ILS).

Figure 1.4: Premium and cost per value by vehicle age for nonfleet clients

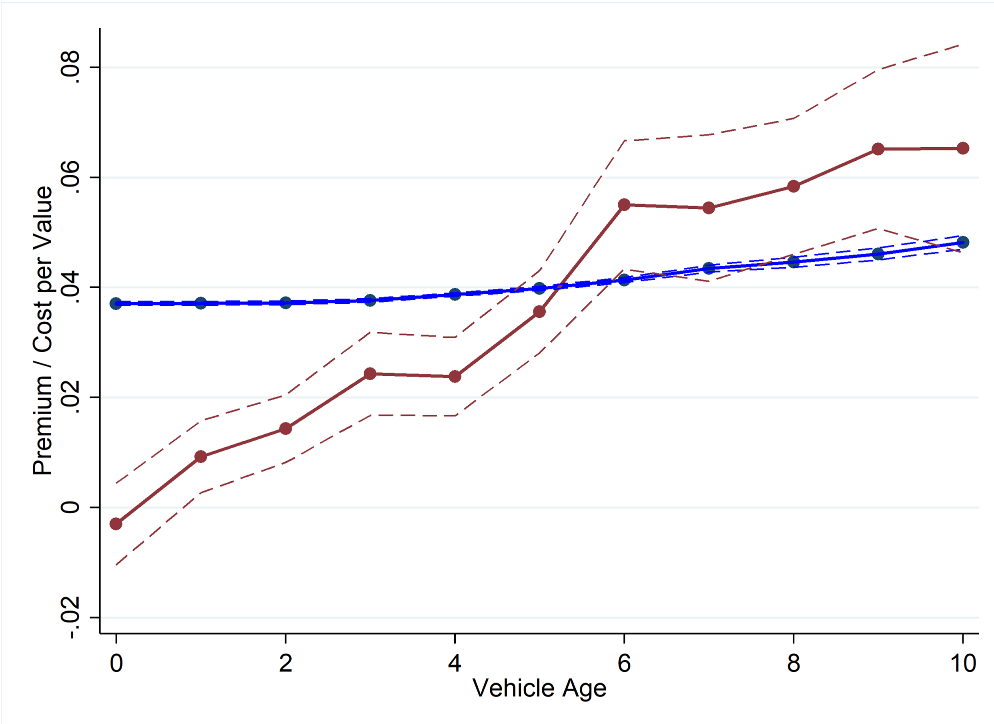

Notes: The figure reports a robustness regression estimation results of the model presented in Figure 1.3, for non-fleet customers, as defined by the insurer. The number of vehicles insured via any type of coverage by client in a given year is less than five.

Figure 1.5: Market-wide premiums per value by vehicle age

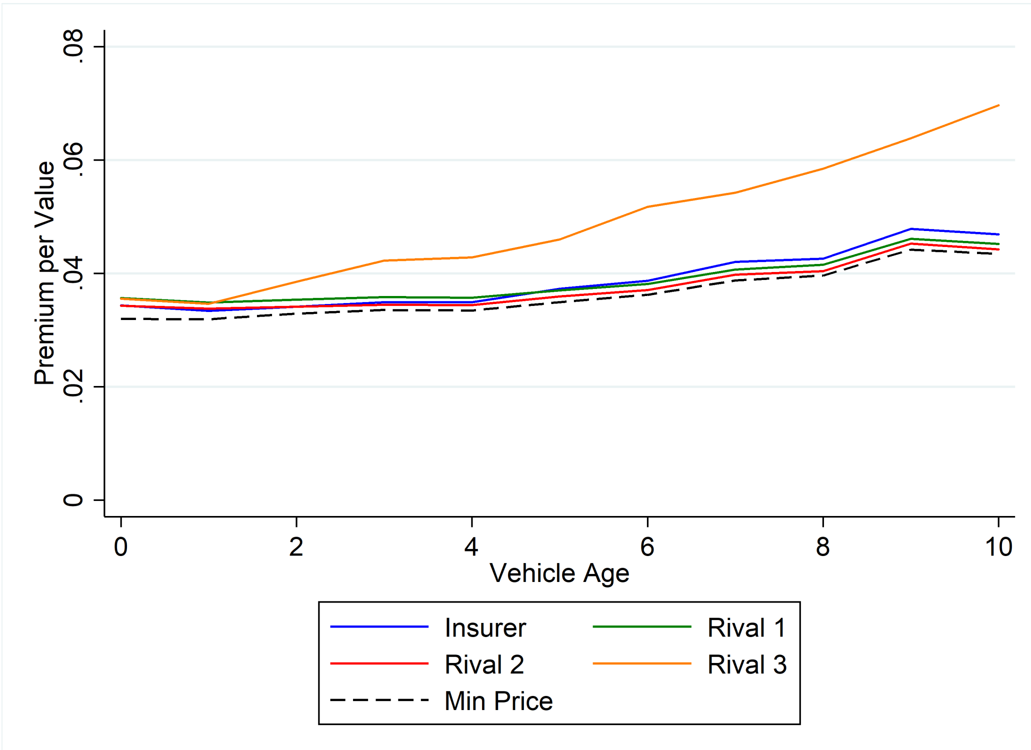

Notes: The figure reports the relationship between market premiums per value and vehicle age. Market-wide premiums are collected via fictitious policy offers generating via Orlan insurance agency's platform (Orlanet Calculator). The sample consists of 876 distinct vehicle model-age values for the top four insurers in the market (the insurer that provided the data and its three main competitors), without any reported claim in the last three years. The horizontal axis depicts premiums per value. The vertical axis depicts vehicle age. The curves are the coefficients of a saturated regression of premiums per value on vehicle age. The dashed line depicts the minimum premium per value in the market (not restricted to the four insurers). Premiums and values are measured in New Israeli Shekel (ILS).

Figure 1.6: Premium adjustment based on new information

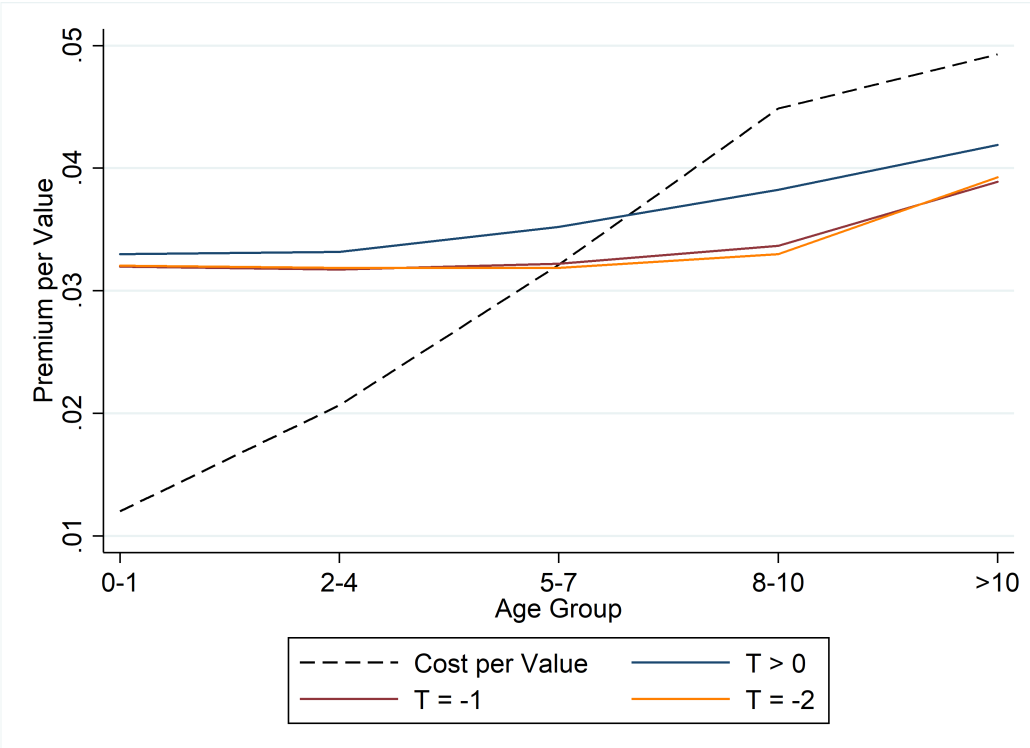

Notes: The figure reports the estimation results of a fixed-effect regression of premium per vehicle value and cost per vehicle value on vehicle age. Observations are divided into five groups based on vehicle age. $T = 0$ indicates the timing at which the insurer was provided information regarding the misadjustment in pricing over the vehicle life cycle. The vertical axis depicts both the premium and cost per value variables. The horizontal axis depicts vehicle age. The red and orange lines represent the estimated premium per value before information was given. The blue line represents the estimated premium per value after information was given. The black dashed line represents the cost per value. The sample includes comprehensive insurance coverage policies for trucks from 2013 through 2020. Premiums and costs are normalized to an annual policy length. Premiums, costs, and vehicle values are measured in New Israeli Shekel (ILS).

Figure 1.7: Comparison of objective and subjective cost estimates

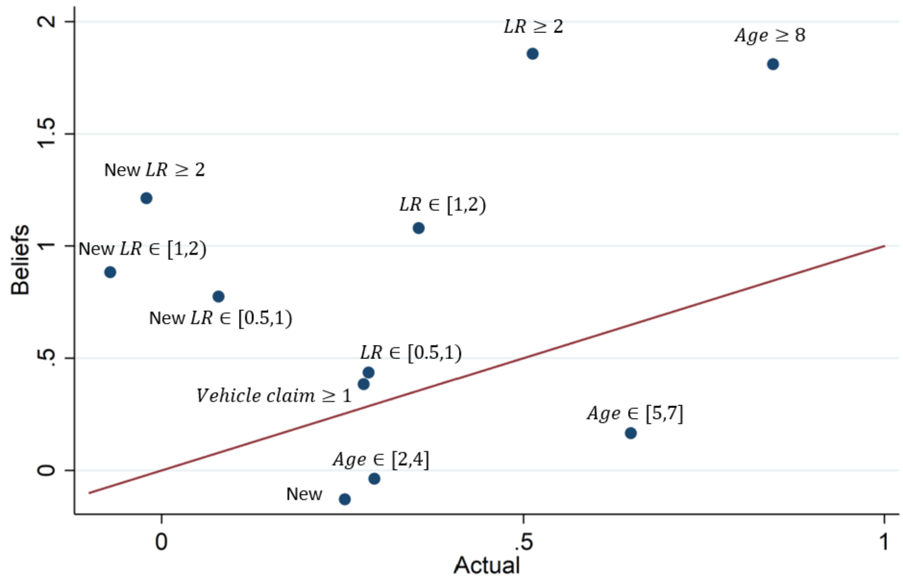

Notes: The figure reports the estimation results of both the objective cost function, as reported in Table 1.7 and the insurer's perceived cost, as reported in Table 1.8. The vertical axis measures the coefficient of the subjective cost components, while the horizontal axis measures the coefficient of the objective cost components.

Figure 1.8: Counterfactual analysis

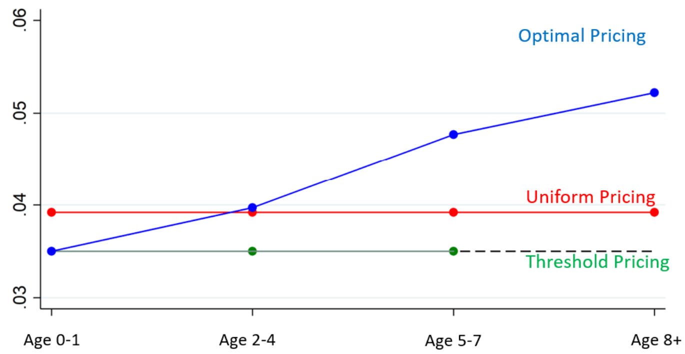

Panel A: Vehicle age

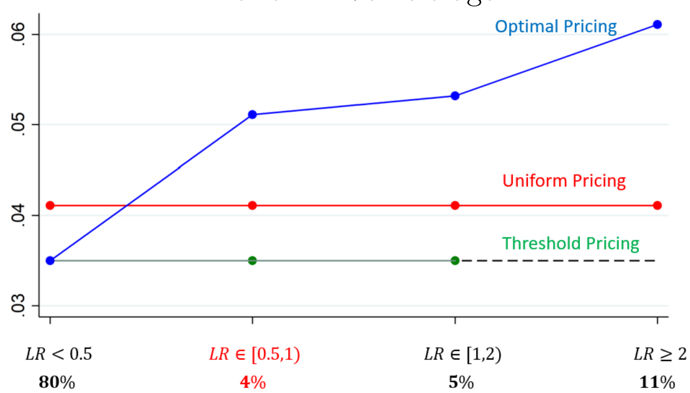

Panel B: Claim

Notes: The figures above provide two counterfactuals. In Panel A, I examine the trend in premium per value (measured on the vertical axis) over the vehicle life cycle (measured on the horizontal axis). The counterfactual analysis is conducted for the average truck owner, who paid 3.5 p.p premium per value. In Panel B, I examine the trend in premium per value (measured on the vertical axis) over different realizations of current year loss ratio (measured on the horizontal axis). The counterfactual analysis is conducted for a single vehicle truck owner with an average loss ratio (0.7) who paid 3.5 p.p premium per value. The blue curve indicates optimal pricing, the red curve indicates optimal pricing conditional on uniform pricing, and the green curve indicates the pricing strategy based on the behavior implied by the "Go—No Go" grades, as presented in Table 1.8.

1.11 Tables

Table 1.1: Summary statistics of comprehensive coverage policies for trucks

	(1) All	(2) Vehicle Age$\geq$ 6	(3) Claim$_{t-1}$ $\geq$1
Policies	51,684	15,506	8,358
Share	100%	30.00%	16.17%
Weighted Share (by Premium)	100%	18.86%	15.96%
Mean Premium	9,938	6,246	9,811
At least 1 claim	23.98%	23.78%	34.82%
Mean Damage	6,794	5,861	9,345
Mean Commission	1,557	1,012	1,577
Mean Profit	1,587	-627	-1,111
Profit Margin	15.97%	-10.04%	-11.32%
Mean Vehicle Age	4.18	8.93	5.00
Mean Vehicle Value	298,659	160,383	275,954
Mean Premium per Value	3.33%	4.01%	3.56%

Notes: The table reports summary statistics of comprehensive coverage policies for trucks between 2013 and 2020. The first column reports statistics for all policies, the second column describes the statistics for policies with a vehicle age of six or above, and the third column describes the statistics for policies with at least one claim in the previous period (Claim$_{t-1}$ $\geq$1). Profit margin is defined as mean profit (=premium$-$damage$-$commission) over mean premium. Mean premium per value is defined as premium over vehicle value. Mean damage is the mean damage of customers' claims (net of deductibles). Vehicle value, premium, commission, paid claims, and profit are measured in New Israeli Shekel (ILS). Vehicle age is measured in years. I exclude from the sample observation with an error , a change in vehicle within the policy, a change in coverage terms over the policy, and policies that did not end or that lasted for less than 30 days (without a claim).

Table 1.2: Policy outcomes and past performance

	Panel A: Entire Sample			
	(1)	(2)	(3)	(4)
	$\text{Claim}_t \geq 1$	$\frac{\text{Damage}_t}{\text{Value}_t}$	$\%\Delta\frac{\text{Premium}_t}{\text{Value}_t}$	Loss Ratio_t
$\text{Claim}_{t-1} \geq 1$	0.122***	0.017***	-0.001	0.418***
	(0.008)	(0.002)	(0.002)	(0.049)
log(Value)	Y	N	N	Y
Vehicle Age - 2^{nd} order	Y	Y	Y	Y
Vehicle Weight Class	Y	Y	Y	Y
Driver Underage Indicator	Y	Y	Y	Y
Observations	32,870	32,870	32,870	32,870
R-squared	0.022	0.009	0.034	0.006
	Panel B: Non-Fleet Policies			
	(1)	(2)	(3)	(4)
	$\text{Claim}_t \geq 1$	$\frac{\text{Damage}_t}{\text{Value}_t}$	$\%\Delta\frac{\text{Premium}_t}{\text{Value}_t}$	Loss Ratio_t
$\text{Claim}_{t-1} \geq 1$	0.118***	0.024***	0.017***	0.547***
	(0.013)	(0.004)	(0.005)	(0.101)
log(Value)	Y	N	N	Y
Vehicle Age - 2^{nd} order	Y	Y	Y	Y
Vehicle Weight Class	Y	Y	Y	Y
Driver Underage Indicator	Y	Y	Y	Y
Observations	8,372	8,367	8,367	8,367
R-squared	0.023	0.009	0.014	0.006

Notes: The table reports the relationship between previous claim history and current outcomes. Panel A's sample includes all non-new comprehensive insurance policies for trucks. Panel B's sample includes only comprehensive insurance policies for trucks of nonfleet clients (number of vehicles insured via any type of coverage by client in a given year is less than five year). $\text{Claim}_{t-1} \geq 1$ is an indicator that equals one if at least one claim has been reported with regard to the policy in the previous period. $\text{Claim}_t \geq 1$ is an indicator that equals one if at least one claim has been reported with regard to the policy in the current period. $\frac{\text{Damage}_t}{\text{Value}_t}$ denotes damage (net claim expenses) per value in the current period, $\%\Delta\frac{\text{Premium}_t}{\text{Value}_t}$ denotes the percent change in premium per value in the current period relative to the previous period ($\frac{p_t - p_{t-1}}{p_{t-1}}$, where p is the premium per value) and Loss Ratio_t denotes the current period policy's loss ratio measured as damage over premiums. Controls include (log) vehicle value, vehicle age (2^{nd} order), vehicle type, vehicle weight class, and driver underage indicator. Vehicle value, premium, damage, and loss ratio are measured in New Israeli Shekel (ILS). Vehicle age is measured in years. Robust standard errors, clustered at the client level, are reported in parentheses. *** p<0.01, ** p<0.05, *p<0.1.

Table 1.3: Policy outcomes and past performance: recent vs. older

	Panel A: Entire Sample			
	(1)	(2)	(3)	(4)
	$\text{Claim}_t \geq 1$	$\frac{\text{Damage}_t}{\text{Value}_t}$	$\%\Delta\frac{\text{Premium}_t}{\text{Value}_t}$	Loss Ratio_t
Client's Agg. Loss Ratio	0.036***	0.006***	-0.002	0.148***
	(0.006)	(0.001)	(0.002)	(0.031)
Client's Prev. Yr. Loss Ratio	-0.000	0.000	0.003**	0.007
	(0.002)	(0.001)	(0.002)	(0.015)
log(Value)	Y	N	N	Y
Vehicle Age - 2^{nd} order	Y	Y	Y	Y
Vehicle Type	Y	Y	Y	Y
Vehicle Weight Class	Y	Y	Y	Y
Driver Underage Indicator	Y	Y	Y	Y
Observations	35,765	35,765	35,765	35,765
R-squared	0.009	0.007	0.009	0.004

Notes: The table reports the relationship between previous claim history and current outcomes. The sample includes all comprehensive insurance policies for trucks from 2014 to 2020, with at least one year of performance history. Client's aggregate loss ratio is the client's ratio of total damages (starting 2013) per total revenue (starting 2013). The client's previous year loss ratio is the client's ratio of previous year's total damages (net claim expenses) over the previous year's total revenue (paid premiums). $\text{Claim}_t \geq 1$ is an indicator that equals one if at least one claim has been reported with regard to the policy at the current period. $\frac{\text{Damage}_t}{\text{Value}_t}$ denotes damage (net claim expenses) per value at the current period, $\%\Delta\frac{\text{Premium}_t}{\text{Value}_t}$ denotes the percent change in premium per value at the current period, relative to the previous period ($\frac{p_t - p_{t-1}}{p_{t-1}}$, where p is the premium per value) and Loss Ratio_t denotes the current period policy's loss ratio, measured as damage over premiums. Controls include (log) vehicle value, vehicle age (2^{nd} order), vehicle type, vehicle weight class, and driver underage indicator. Vehicle value, premium, damage, and loss ratio are measured in New Israeli Shekel (ILS). Vehicle age is measured in years. Robust standard errors, clustered at client level, are reported in parentheses. *** $p<0.01$, ** $p<0.05$, *$p<0.1$.

Table 1.4: Market premiums by claim history

	Dependent Variable: Premium per Value (in percent)				
	(1)	(2)	(3)	(4)	(5)
	Insurer	Rival 1	Rival 2	Rival 3	Min. Price
Constant	3.79***	3.80***	3.69***	4.76***	3.58***
	(0.03)	(0.02)	(0.02)	(0.05)	(0.02)
1 Claim Last Yr (3 yrs)	0.28***	0.00	0.00	·	0.00
	(0.00)	(0.00)	(0.00)		(0.00)
$\geq$ 2 Claims Last 3 Yrs	·	·	·	·	·
Observations	1,752	1,752	1,752	876	1,752
R-squared	0.03	0.00	0.00	0.00	0.00

Notes: The table reports the relationship between market premiums per value and claim history. Market-wide premiums are collected via fictitious policy offers generating via Orlan insurance agency's platform (Orlanet Calculator). The sample consists of 876 distinct vehicle model-age - values for the top four insurers in the market (the insurer that provided the data and its three main competitors). For each one, I generate two observations: one without any claim in the last three years and one with one claim in the last three years, which occurred last year. Note that the Orlanet Calculator does not generate policy offers for the case of at least two claims in the last three years. Columns 1 through 4 present the relationship between claim history and premium per value offered by the top four insurers in the market, while column 5 describes the relationship with regard to the minimum premium per value in the market (not restricted to the four insures). Premiums and values are measured in New Israeli Shekel (ILS). Robust standard errors are reported in parentheses. *** p<0.01, ** p<0.05, *p<0.1.

Table 1.5: Summary statistics of the Go—No Go grading

	(1) All	(2) Go	(3) Inc. Ded.	(4) Inc. Prem.	(5) TP Only	(6) Deny
Policies	14,288	12,478	625	386	258	541
Share	100%	87.33%	4.37%	2.70%	1.81%	3.79%
Premium	9,028	8,965	8,762	10,809	7,251	10,368
Policies With Claim	15.39%	12.29%	36.32%	37.82%	20.54%	44.18%
Damage	4,320	2,558	9,635	12,208	2,924	33,845
% Loss Ratio	47.85%	28.53%	109.96%	112.94%	40.32%	326.45%
Vehicle Age	4.16	4.02	4.39	4.12	9.80	4.35
Vehicle Value	260,601	261,628	242,907	305,377	167,203	269,947

Notes: The table reports summary statistics of insurer grading for all comprehensive and partial coverage policies between 2018 and 2020. The first column reports statistics for all policies, the second column describe the statistics for policies that received a Go grade. Columns 3 through 6 describe the statistics for policies that received a No-Go grade. Column 3 describe the statistics for policies that the operational team recommends a change in terms without increasing premiums (increase deductibles), column 4 describes the statistics for policies that the operational team recommends a price increase, column 5 describes the statistics for policies that the operational team recommends to offer only third-party coverage (i.e., not to provide comprehensive coverage), and column 6 describes the statistics for policies that the operational team recommends denying. Variables are defined as in Table 2.1.

Table 1.6: Probability of a Go grade

	Probit Model. Dependent Variable: $Go = 1$			
	(1)	(2)	(3)	(4)
Age 2-4	-0.023		-0.020	-0.024
	(0.014)		(0.014)	(0.017)
Age 5-7	-0.026*		-0.023*	-0.026
	(0.014)		(0.013)	(0.017)
Age 8+	-0.093***		-0.086***	-0.094***
	(0.017)		(0.016)	(0.020)
Client's Aggregate Loss Ratio		-0.032***	-0.032***	-0.072***
		(0.010)	(0.010)	(0.014)
Client's Prev. Yr. Loss Ratio		-0.025***	-0.024***	-0.026***
		(0.005)	(0.005)	(0.007)
Fleet Size	Y	Y	Y	Y
Vehicle Type	Y	Y	Y	Y
Vehicle Weight Class	Y	Y	Y	Y
Driver Underage Indicator	Y	Y	Y	Y
Sample	All	All	All	History ≥ 5
Observations	14,288	14,288	14,288	9,586
Pseudo R-squared	0.03	0.12	0.13	0.12

Notes: The table reports the relationship between the previous claim history, vehicle, and assignment of a Go grade. The sample includes all comprehensive and partial coverage policies with an assigned insurer grading between 2018 and 2020. The dependent variable is equal if the operational team assigned the policy with a Go grade. The explanatory variables consist of three age variables: (i) Age 2-4, a dummy variable that equals one if the vehicle age is between 2 and 4, (i) Age 5-7, a dummy variable that equals one if the vehicle age is between 5 and 7 and (i) Age 8+, a dummy variable that equals one if the vehicle age is 8, or above. There are two explanatory variables with regard to claim history: (i) client's aggregate loss ratio, which is the client's ratio of total damages (starting in 2013) per total revenue (starting in 2013) and (ii) client's previous year loss ratio is the client's ratio of previous year total damages (net claim expenses) over the previous year total revenue (paid premiums). Both are measured in New Israeli Shekel (ILS). Controls include fleet size of client, which is defined as the number of vehicles insured by the client at a given year, vehicle type, vehicle weight class and an indicator for permitted underage driver. The estimation is conducted using a probit model. Coefficients reported are marginal effect at mean. Columns 1 through 3 include the entire sample, while column 4 includes the sum sample of policies with at least 5 years of recorded history (starting in 2013). History is measured as the sum of the years each of the client's policies are observed. Robust standard errors, clustered at client level, are reported in parentheses. *** p<0.01, ** p<0.05, *p<0.1.

Table 1.7: Structural estimation - demand side

	Damage per Value		Renew Comp. Coverage	
Premium per Value			-162.991***	(8.834)
License Avg. Premium per Value			123.330***	(7.288)
Vehicle Value (in 100,000 ILS)			0.124***	(0.013)
Price Increase			-0.718***	(0.032)
Age groups:				
0-1	(omitted)		(omitted)	
2-4	0.219***	(0.070)	-0.594***	(0.037)
5-7	0.522***	(0.073)	-1.188***	(0.045)
≥ 8	0.725***	(0.074)	-1.423***	(0.049)
Client's Aggregate Loss Ratio:				
$[0.5, 1)$	0.348***	(0.070)	-0.237***	(0.049)
$[1, 2)$	0.436***	(0.089)	-0.681***	(0.066)
≥ 2	0.600***	(0.147)	-0.723***	(0.109)
Client's Prev. Yr. Loss Ratio:				
$[0.5, 1)$	0.108	(0.069)	0.230***	(0.034)
$[1, 2)$	-0.046	(0.085)	0.334***	(0.042)
≥ 2	-0.005	(0.132)	0.418***	(0.072)
$(\text{Claim}_{i\ell t-1}) \geq 1$	0.311***	(0.058)	0.652***	(0.031)
Comp. Coverage	0.202*	(0.112)	0.332***	(0.056)
Fleet Size	-0.0004**	(0.0001)	-0.0019**	(0.0002)
Underage Driver	0.056	(0.076)	0.081*	(0.044)
Joined last yr.	0.196***	(0.071)	0.433**	(0.204)
History (in 1,000 yrs)	-0.017	(0.011)	0.001	(0.013)
Selection	1.079***	(0.284)		
Client unobs. s.e.	0.061	(0.056)	1.741	(0.041)
Observations	91,603		73,171	
Log Likelihood	-8696		-33756	

Notes: The table reports the results of the structural estimation of the demand for insurance. The left column presents the main estimates of damage per value. The right panel presents the main estimates of renewal. Joined last year is an indicator that equals to 1 if the client purchased its first policy from the insurer in the last year. Client's aggregate loss ratio is the client's ratio of total damages (starting in 2013) per total revenue (starting in 2013). Client's previous year loss ratio is the client's ratio of previous year total damages (net claim expenses) over the previous year total revenue (paid premiums). $(\text{Claim}_{i\ell t-1}) \geq 1$ is an indicator that equals one if at least one claim has been reported with regard to the license in the previous period. Price increase is an indicator that equals one if the policy faces an increase in premium per value. Selection measures the relationship between the unobserved client-level damage component and the unobserved client-level demand component. Premium, damage, and vehicle value are measured in New Israeli Shekel (ILS). Client unobserved s.e. measures the magnitude of heterogeneity (standard errors) in unobserved client-level damage and demand components. *** p<0.01, ** p<0.05, *p<0.1.

Table 1.8: Structural estimation - supply side

	Go		Adjust Terms		Deny
$\log(p_i)$	1.223***	(0.461)	1.549***	(0.568)	
$\log(p_i) - \log(\bar{p}(x))$	-1.555***	(0.531)	-2.566***	(0.652)	
Age groups:					
0-1	(omitted)		(omitted)		
2-4	-0.015	(0.209)	0.074	(0.245)	
5-7	-0.282	(0.244)	0.427	(0.280)	
≥ 8	-2.297***	(0.267)	-2.457***	(0.351)	
Client's Aggregate Loss Ratio:					
$[0.5, 1)$	-0.552**	(0.220)	0.258	(0.267)	
$[1, 2)$	-1.393***	(0.204)	-0.205	(0.258)	
≥ 2	-2.476***	(0.220)	-0.984***	(0.283)	
Client's Prev. Yr. Loss Ratio:					
$[0.5, 1)$	-0.964***	(0.286)	0.341	(0.335)	
$[1, 2)$	-1.211***	(0.261)	0.232	(0.309)	
≥ 2	-1.544***	(0.230)	-0.098	(0.285)	
$(\text{Claim}_{it-1}) \geq 1$	-0.466***	(0.175)	0.326	(0.209)	
Comp. Coverage	0.051	(0.206)	0.115	(0.270)	
Fleet Size	0.035	(0.070)	0.184**	(0.086)	
Underage Driver	-0.501*	(0.272)	-0.694**	(0.338)	
New Client	0.196	(0.171)	0.433**	(0.204)	
History	-0.017	(0.011)	0.001	(0.013)	
Constant	8.943***	(1.659)	4.818**	(2.054)	0
Observations	6,347				
Log Likelihood	-1954.2				
Pseudo R-squared	0.266				

Notes: The table reports the relationship between the policy's observable characteristics and insurer grading. The sample includes all comprehensive and partial coverage policies with an assigned insurer grading between 2018 and 2020 for nonfleet policies (i.e., policies for clients with a fleet size below 5 during the relevant year). The insurer's alternatives are (i) "Go", which means the operational team recommends renewing policy with the same terms, (ii) "Adjust", which means the operational team recommends offering a policy with increased premiums or deductibles, and (iii) "Deny", which means the operational team recommends denying comprehensive coverage. The explanatory variables consist of vehicle age variables, client's aggregate loss ratio, and client's previous year loss ratio, as defined in Table 1.6. $\log(p_i)$ is the policy's log premium per value and $\log(p_i) - \log(\bar{p}(x))$ is the difference between the policy's log premium per value and the log average premium per value paid for a policy with the same observable characteristics. Additional explanatory variables include, $(\text{Claim}_{it-1}) \geq 1$, an indicator for a claim event at the policy level in the previous period, comprehensive coverage dummy variable, underage driver indicator, fleet size, history of client with insurer, and new client indicator (joined last year). Estimation includes controls for vehicle type and weight class and year. History and vehicle age are measured in years. Premiums, damages, and vehicle values are measured in New Israeli Shekel (ILS). Estimation is conducted using a multinominal logistic regression model. Analytical asymptotic standard errors are reported in parentheses. *** p<0.01, ** p<0.05, *p<0.1.

Chapter 2

Price and Prejudice: Customer Taste-Based Discrimination and Competition

2.1 Introduction

Today, almost 250 years after Adam Smith, 1776 's *The Wealth of Nations*, there is a broad consensus among economists that a competitive market setting is fundamental for spurring well-being. Competition over demanding buyers drive firms to lower production costs, innovate, offer greater variety and a higher quality of goods, and improve services. The welfare gains exceed the product market as competition *disciplines* discriminatory sellers to practice equity and fairness in the labor market (Becker, 1957) while *accommodating* buyers' demand for discrimination (Arrow, 1972, Cain, 1986).

Inspired by Becker 's market framing of taste-based discrimination in his seminal work *The Economics of Discrimination* (1957), this paper offers a different perspective on the impact of competition in the product market on labor market discrimination. Using theory and data, I show conceptually and document empirically that competition in the product market intensifies customer taste-based discrimination in the labor market. The overall impact of product market competition on labor market discrimination is ambiguous since it combines two offsetting forces—a *disciplining* effect on the seller and an *accommodating* effect on buyers.

At the heart of my framework is the idea that a monopolistic seller jointly determines the quantity of production and the share of White workers. The monopolistic seller tax customer discrimination—with a mark-up—when the marginal willingness to pay diminishes faster with the share of White employees. This occurs, for example, when White buyers have higher purchasing power than Black buyers.

I study the equilibrium implications of competition in the product market on discrimination in the labor market by integrating this demand-side property into a Dixit-Stiglitz two-sector general equilibrium model, as in Spence (1977) and Dixit (1979). Sellers produce and offer bundled products of goods and services. They hire Black and White workers who interact with buyers.

In this case, the monopolistic seller mark-up each component, the quantities, and the share of White employees leading to the known Spence (1975) distortion.[1] A competitive market mirrors the customers' preferences, whereas a monopoly *disciplines* discriminatory buyers by taxing their desire to replace Black workers with equally productive White workers. They charge more and provide less discrimination. Consequently, monopolistic sellers hire a more balanced racial composition of labor and pay White workers lower wage premiums than sellers in a competitive product market.

The canonical model of discrimination suggests that competition diminishes customers' taste-based discrimination by reallocating workers between equally paid jobs (Arrow, 1972, Cain, 1986). The impact of product market competition on labor market discrimination against Black workers depends not only on customers' prejudicial preferences but also on the share of the service sector in the overall economy. If the cost of segregation—the Blacks' opportunity cost from avoiding the service sector—is too high, the lower labor demand is also reflected in wage differential.

Altogether, competition in the product market might intensify taste-based discrimination in the labor market. It depends on the shrinkage of employer-driven discrimination and the amplification of customer-driven prejudice. Product market competition is advantageous when customers' prejudicial preferences and interaction with employees are minimal.

Guided by this set of testable implications, I next turn to the data. I use the massive deregulation of the Banking sector in the US to identify shocks to product market competition and quantify its impact on customer taste-based discrimination. The banking sector provides an excellent laboratory to assess the impact of market structure on labor market taste-based discrimination for four main reasons.

First, the extensive deregulation across the US banking sector, which ended local banking monopolies, identifies local shocks to competition in the banking sector. US states relaxed branching restrictions between the mid-1970s and the mid-1990s by allowing banks to merge within and across states. The timing of inter-state and intra-state deregulation was not related to local labor market conditions, the wage

[1] See also Sheshinski (1976).

premium in the banking sector, and within-sector inequalities between Black and White workers (Black and Strahan, 2001, Levine, Levkov, and Rubinstein, 2008, Beck, Levine, and Levkov, 2010).

Second, the banking sector provides financial services using workers in a wide range of specialized occupations. Some jobs, such as financial managers, bank tellers, and customer service representatives, require intense contact with clients. Other jobs, such as computer software developers, bookkeepers, and clerks, do not interact with clients. This feature provides an opportunity to estimate the impact of intensified competition in the banking sector on the Black-White gaps in employment and wages, comparing occupations requiring substantial contact with clients to other occupations, inside and outside the banking sector.

Third, the banking sector provides similar services in all states. Nevertheless, not all states share the same prejudicial preferences. I use heterogeneity in taste for discrimination to assess whether the impact of bank deregulation on racial gaps in employment and wages varies with clients' discriminatory attitudes toward Blacks. Fourth, the purchasing power of White buyers is disproportionally higher than that of their Black peers.

I use data from four primary sources. First, I use micro-data from the decennial US Censuses from 1960 through 2000 and the 2010-2012 3-year ACS to create a representative sample of the Black and White prime-aged population in the US during this period. Second, I construct post-intra-state and inter-state deregulation indicators, as in Black and Strahan (2001). Third, I use O*NET data on occupational task characteristics to measure the intensity of the job's requirement contact with customers, following Deming (2017). Finally, I complete my data with measures of racial prejudice by state using the GSS data, following Charles and Guryan (2008).

I use the combined data to estimate the impact of bank deregulation on wages and employment by the occupation's required customer contact. My wage and employment sample focuses on prime-aged males who are full-time and full-year workers.

I begin with a detailed assessment of the change in the Black-White wage gap by replicating the Black and Strahan (2001) analysis for different sub-sets of the data, constructed based on the occupational requirement to deal with customers. I find that in low customer contact occupations, mainly exposed to employer discrimination, the Black-White wage gap dropped following intra-state deregulation that increased competition, as the Becker (1957) model predicts and consistent with the findings in Black and Strahan (2001). In contrast, the Black-White wage gap substantially increased in high customer contact occupations. The amplified gaps demonstrate that competition's detriment reinforcement of buyers' taste for discrimination outweighs the beneficial shrinkage of prejudiced, inefficient sellers.

I then examine how the banking sector' Black-White wage gap evolved differentially by exploiting variation in the customer index values. Following intra-state deregulation, the increase in the Black-White wage gap is 11 percent higher when considering occupations that require more contact with customers, with one standard deviation higher customer index value.

Furthermore, I consider how the documented systematic change in the Black-White wage gap by customer index varies with the state's prejudicial tastes. I divide the states into three groups based on the prejudice measure of the "marginal" White individual, as in Charles and Guryan (2008). The results indicate that states with strong prejudicial tastes are driving the increase in the Black-White wage gap in customer-contact occupations. In the high prejudice index states, the increase in the Black-White wage gap is 25 percent higher when considering occupations with one standard deviation higher customer index value, while only an insignificant effect of 6 percent in the rest of the states.

Lastly, I examine whether the competition shock affected the extensive margin: employment. Using within banking variation, I find that the relative proportion of Blacks in occupations requiring above-average customer contact dropped by 3.2 p.p following intra-state deregulation. The drop is especially sharp among occupations with meaningful interaction with customers, financial managers, and specialists. Similar to wages, the proportional decrease in Black employment among customer-contact occupations is more substantial in states with strong prejudicial tastes.

This paper contributes to the literature on discrimination and competition, both theoretically and empirically. Theoretically, several papers illustrate how discrimination can persist in perfectly competitive product markets; alternative channels include imperfect information (Borjas and Bronars, 1989, Black, 1995, Rosén, 1997), adjustment costs (Lang, Manove, and Dickens, 2005), imperfect substitution of labor (Kahn, 1991), and employer's will to bear the cost of discrimination (Goldberg, 1982).[2] I show that discrimination persists in perfectly competitive markets and that competition intensifies discrimination in frictionless markets when prejudiced customers drive discrimination.

Empirically, following Becker's seminal work, a large body of research studies the overall effect of competition on labor market taste-based discrimination. Findings are mixed; some document a substantial impact of product market competition on labor market gender and racial gaps (Ashenfelter and Hannan, 1986, Black and Strahan, 2001, Peoples and Talley, 2001, Black and Brainerd, 2004, Levine, Levkov, and Rubinstein, 2008, Heyman, Svaleryd, and Vlachos, 2013, Weber and Zulehner, 2014, Hirata and Soares, 2020). Others document limited or no effect of competitive forces on labor market gaps in outcomes (Hellerstein, Neumark, and Troske, 1997, Coleman,

[2]See Altonji and Blank (1999) and Lang and Lehmann, 2012 for an extensive review.

2004, Berson, 2012, Cooke, Fernandes, and Ferreira, 2019, Aneja and Krishnamurthy, 2022). I show that examining the impact of competitive forces on labor outcomes at the aggregate might be misleading, as the effect of product market competition on labor market discrimination depends on the factors driving the differential outcomes: *employers* or *customers*.[3]

This paper contributes to the literature on customer taste-based discrimination. Several papers examine the implications of prejudiced importance by evaluating the price or transaction volumes of items based on the products' racial identity (Nardinelli and Simon, 1990, List, 2004) or the seller's racial identity (Doleac and Stein, 2013, Ayres, Banaji, and Jolls, 2015, Bar and Zussman, 2017). Others examined the labor market implications in terms of earnings and employment (Kahn and Sherer, 1988, Holzer and Ihlanfeldt, 1998, Bertrand and Mullainathan, 2004, Leonard, Levine, and Giuliano, 2010, Combes et al., 2016, Hurst, Rubinstein, and Shimizu, 2021, Kline, Rose, and Walters, 2022).[4] My paper adds to this literature as I emphasize that discrimination depends both on the marginal buyer (Heckman (1998)) and the product market structure. The growing share of the service sector (Autor and Dorn (2013)) points out the increasing relevancy of customer taste-based discrimination and anti-discriminatory policies (Donohue and Heckman (1991)) and effectively the labor market discrimination that Black workers face over time (Bayer and Charles (2018)).

The rest of the paper is organized as follows. Section 2 presents a stylized model demonstrating the implications of customer heterogeneity in taste for discrimination on the aggregate demand. Section 3 presents the Dixit-Stiglitz two-sector general equilibrium model. Section 4 provides a background on the US banking deregulation. Section 5 describes the data. Section 6 describes the empirical analysis and reports the result. Section 7 concludes.

2.2 Heterogeneity in Taste for Discrimination

In this part, I introduce a stylized model of customer heterogeneity. This model allows examining the implications of customers' heterogeneous tastes for discrimination on the firm's aggregate demand. Individuals vary in terms of their utility from the good and their taste for discrimination. This model deviates from previous studies, allowing heterogeneity in the customers' prejudicial preferences. I incorporate heterogeneity in Becker's taste-based discrimination model. Heterogeneity in prejudicial preferences plays a substantial role. I show that the employee's group

[3] I do not consider the interaction among employees of different race (Onuchic and Ray, 2021).
[4] See Lang and Kahn-Lang Spitzer (2020) for an extensive review.

affiliation shapes the aggregate demand in terms of level and, more importantly, slope.

Consider a unit mass population of individuals interested in purchasing one unit of a good. When purchasing the product, the consumers interact (and are randomly matched) with the firm's employees.[5] There are two types of employees providing service: White and Black. Customers have a taste for discrimination. Customers are willing to pay more for service provided by an employee of the preferred group. The customers differ in two dimensions. They vary in terms of their value from the product and their prejudicial preferences.

Previous studies model the consumers' prejudicial tastes as observationally equivalent to a homogeneous taste for White employees. I deviate from previous studies; taste-based discrimination varies across customers. Intuitively, some customers prefer White employees to serve them, while others prefer Black employees. The differences in taste raise a puzzle. Why do Black employees face discrimination? If White and Black customers prefer to face an employee of their group affiliation, it is unclear why one group of employees would be favored over the other. Becker has attributed racial labor market discrimination to Blacks' relatively weak economic influence.

"A necessary condition for effective discrimination against N is that N is an economic minority...a necessary and sufficient condition is that N be more of an economic minority than a numerical majority...analysis of discrimination in competitive free-enterprise societies also uses a minority-majority framework, but the concept of economic minorities is somewhat more important here than the numerical ones." Becker (1971).

I address this issue by assuming that White customers have a higher purchasing power than Blacks. Formally, let $\theta_i(s)$ denote individual i's (log) willingness to pay. The willingness to pay for the product depends on the firm's composition of employees by group affiliation. Let s denote the share of White employees among the firm's employees.[6] The willingness to pay, $\theta_i(s)$, is defined as follows:

$$\theta_i(s) = \underbrace{\nu_i + \rho\alpha_i}_{\tilde{\nu}_i} + \alpha_i s. \tag{2.1}$$

$\tilde{\nu}_i$ represents i's value of the good (when served by a Black employee). $\rho \geq 0$ reflects that, on average, customers with a higher demand for the good are willing to pay more for White service. ν_i is the fraction of willingness to pay unrelated to service; $\nu\alpha$. ν_i is distributed according to a continuous and strictly positive density, F_ν, defined on $\tilde{\nu}_i \in (\nu_{min}, \nu_{max})$. α_i represents the individual's net utility from facing a White

[5] Consumers cannot pay a premium to face an employee of their choice.

[6] All consumers are utility-maximizers and risk-neutral.

employee. α_i is distributed according to a continuous and strictly positive density, F_α, defined on $\alpha_i \in (\alpha_{min}, \alpha_{max})$, where $\alpha_{min} < 0$ and $\alpha_{max} > 0$.[7] The consumer's willingness to pay is the sum of the positive benefits from both attributes—good and service—as in Rosen's (1974) hedonic price framework.

The variation in purchasing power is not limited to level differences. The differences in willingness to pay between the economic majority and minority are starker as the share of White employees increases. Intuitively, customers with a higher willingness to pay prefer, on average, White service and are willing to pay even more as the share of White employees increases. In contrast, on average, customers who are willing to pay less prefer Black service and are willing to pay less when the share of White employees increases.

Proposition 1. *Let $F_\alpha(\cdot|Q, s)$ denote the cumulative distribution of preferences for White service among the $Q \in (0, 1)$ consumers with the highest willingness to pay, given s. The infra-marginal customers' distribution of prejudicial preferences given s first-order stochastically dominates that of $s' < s$. I.e., $F_\alpha(\alpha|Q, s) \leq F_\alpha(\alpha|Q, s')$, $\forall \alpha \in (\alpha_{min}, \alpha_{max})$, $s' < s \in (0, 1]$, $Q \in (0, 1)$, and strict inequality for some $\alpha \in (\alpha_{min}, \alpha_{max})$.*[8]

Customer heterogeneity in taste for discrimination affects the sorting of individuals along the demand curve. When considering a homogeneous taste for discrimination, the marginal customer's benefit from White service is identical to that of the infra-marginal customers. As a result, the firm's employment decision is unrelated to the market structure. In contrast, a heterogeneous taste for White service results in a clear difference between the two, which systematically changes with the firm's employment share of White employees. When facing a White employee, customers with a strong taste for discrimination towards Black employees have a higher demand than those with mild prejudicial preferences or even those who prefer Black employees.

Consequently, the aggregate demand is *prejudicially selected*. As the share of White employees increases, demand sorting is based more on the benefit of White service. Customers with a strong taste for discrimination towards Black employees increase their demand for the product more than customers with a mild taste or preference for Black service. Therefore, as the firm's share of White employees increases, the inframarginal customers' increase in willingness to pay is, on average, higher than that of the marginal customer.

---

[7]The results in this part hold for a generalized utility function as well; $\theta_i(s) = \tilde{\nu}_i + \alpha_i g(s)$, where $g(0) = 0$ and $g'(s) > 0$. Furthermore, the results hold for $\alpha_{min} = 0$.

[8]Proofs of all propositions can be found in the appendix.

Proposition 2. $\exists\, q_b \in (0,1)$ *such that* $\forall s' > s$*,* $P_{s'}(Q) < P_s(Q) + \min\{(s'-s)[\alpha], 0\}$ *,* $\forall Q > q_b$ *and* $\exists\, q_t \in (0, q_b]$ *such that* $\forall s' > s$*,* $P_{s'}(Q) > P_s(Q) + \max\{(s'-s)[\alpha], 0\}, \forall Q < q_t$*.*

Proposition 2 has two implications concerning the aggregate demand curve. First, as the share of White employees increases, the aggregate demand curve shifts upwards if, on average, customers prefer White service. By proposition 2, the effect of increasing White employment is not limited to a level effect. To illustrate, consider the "top" and the "bottom" regions of the aggregate demand curve. The "top" region consists of customers with stronger preferences for White employees than the average customer. Therefore, the increase in the share of White employees raises their willingness to pay by more than that of the average customer. Demand sorting further implies that their increase in willingness to pay is, on average, more significant than the rest of the customers, who are willing to pay less. In contrast, the "bottom" region consists of customers who prefer Black service. As the share of White employees increases, their willingness to pay drops. Since the "Bottom" region is dominated by customers that prefer Black service, demand resorting implies that the willingness to pay drops, on average, by more than for customers with a higher willingness to pay (for some, the demand increases). Therefore, customer resorting implies that, on average, the demand curve is steeper as White employment increases. It could be the case that the demand is not well-behaved throughout the customers' distribution of willingness to pay. To establish that the demand curve is *always* steeper as White employment increases, a monotonicity assumption is required; $\forall s' > s$, $P_{s'}(Q) - P_s(Q)$ monotonic in Q. This assumption holds, for example, when ν and α are normally distributed.

In this part, I have shown that as the firm hires a higher share of White employees, the aggregate demand curve shifts and becomes steeper. Since the share of White employees is endogenous, the firm takes advantage of the reshaping properties of the demand curve for its advantage. This part of the analysis is standard with the analysis of product quality. When the difference between the high White share and low White service demand curves decreases with production—i.e., when the demand is *prejudicially selected*the monopolist would under-provide White service (given quantity of production), this is a result of the Spence distortion (spence1975monopoly, sheshinski1976price); When hiring, monopolist considers the prejudicial preferences of the marginal buyer instead of considering the average benefit from White service of the infra-marginal customers. Next, I consider the firm's behavior in a general equilibrium setting where hiring and production are endogenously set.

2.3 General Equilibrium Model

In this part, I integrate heterogeneity in customers' taste for discrimination into a Dixit-Stiglitz-variant framework. I characterize the relative labor demand for employees by group affiliation, the implied labor market equilibrium wage, and employment gaps under different product market structures and customers' preferences. Last, I conduct a comparative static analysis to determine the implications of product market competition on labor market outcomes when considering both prejudiced customers and discriminatory employers.

Setting

In this part, I outline the model's setting. In this general equilibrium model, wages, employment, production, and prices are determined in equilibrium by both the product market and the labor market clearing conditions. The model consists of three agents: (i) customers that are interested in purchasing the offered goods, (ii) profit-maximizing firms that decide whether to enter the market and, if so, select quantities of production and composition of employed labor, and (iii) labor force, differentiated by group affiliation and productivity, that select employment to maximize wage.

Customers

Consider the service sector in the spirit of Dixit and Stiglitz, 1977. Potentially infinite firms can enter the market and produce their unique good. The goods are provided to the customers by the firms' employees. Customers' utility depends on both the goods and the services provided. The representative consumer's benefit from the service sector is defined as follows:

$$\mathcal{V} = \frac{1}{1-\eta}\Big(\int_0^n u(x_i, s_i)di\Big)^{1-\eta} + m, \quad \text{s.t.} \int_0^n x_i P_i di + m \leq I, \tag{2.2}$$

where n represents the number of offered products in the sector. $u(x_i, s_i)$ denotes the (direct) utility from product i. It is a function of both quantity, x_i, and the share of White employees providing service, s_i. P_i denotes product i's price. I focus on the case of product-market substitutes; $\eta \in (0, 1)$.[9] The rest of the economy is aggregated and is represented by a numerairé m. Last, I denotes the consumer's income.

Dixit and Stiglitz, 1977 model consumer preferences for quantity using an iso-elastic utility function. I incorporate heterogeneity in customers' non-price attributes

[9]All of the results in this paper that do not address the monopoly's solution hold for any η.

(i.e., taste for discrimination), as introduced by Spence (1975):

$$u(x,s) = A(s)x^{\sigma(s)} \tag{2.3}$$

where $A(s)$ denotes the level benefit from White service, and $\sigma(s)$ represents the heterogeneous component; the sorting effect on an increase in the share of White employees. This modeling approach, as done in Spence (1975, 1977) and Dixit, 1979, incorporates preferences for both quantity and White service to a general equilibrium model of monopolistic competition à la Dixit and Stiglitz, 1977 while maintaining the standard constant price elasticity assumption.

The definitions of both the representative consumer's benefit from the service sector (equation 2.2) and the iso-elastic utility (equation 2.3) imply that product i's inverse demand function is given by:

$$P_i(x_i, s_i) = \frac{\partial \mathcal{V}}{\partial x_i} = \lambda^{-\eta}\sigma(s)A(s)x^{\sigma(s)-1}, \tag{2.4}$$

where $\lambda = \int_0^n u(x_i, s_i)di$ denotes the consumers surplus (from the sector's products).

Throughout the analysis, I focus on the following case:

$$\sigma'(s) < 0.$$

This assumption has a couple of implications. First, consistent with the stylized model of customer heterogeneity in prejudicial tastes, an increase in the share of White employees generates a steeper aggregate demand curve. The aggregate demand is less price sensitive when customers face a higher share of White employees. In particular, the price elasticity equals $\frac{1}{1-\sigma(s)}$.[10] Furthermore, consistent with the stylized model of heterogeneity in consumer preferences, the Spence distortion $\frac{\frac{1}{x}\int_0^x P_i(q,s_i)dq}{P_i(x,s_i)} = \frac{1}{\sigma(s)}$ increases with s; an increase in the share of White employees increases the dispersion of the customers' willingness to pay distribution. Therefore, the ratio of the infra-marginal customers' average willingness to pay to that of the marginal consumer increases with White service. Lastly, given a constant marginal cost of production, the firm's optimal pricing scheme is a proportional mark-up over the cost of size $\frac{1}{\sigma(s)}$. Improving non-price attributes—, i.e., increasing the share of White employees—enables the firms to extract a higher mark-up from its customers.

[10] I assume that n is sufficiently large and accordingly neglect the effect of a change in $P_i(x_i, s_i)$ on λ.

Firms

Potentially infinite firms can enter the market and produce their unique product. For simplicity, I consider the case of symmetric production functions, as in Spence (1975,1977), Dixit and Stiglitz (1977), and Dixit (1979). A firm can produce a unit of goods at a constant rate of δ units of labor. The cost of production, which is denoted by $c(x, s)$, include a fixed cost of f units of good, which limits entry:

$$c(x, s) = \delta(x + f)(\omega^b + s\Delta\omega), \tag{2.5}$$

where ω^i denotes the wage by group affiliation, and $\Delta = \omega^{\mathcal{A}} - \omega^{\mathcal{B}}$ denotes the wage differential.

Labor Force

Previous studies consider labor mobility as a channel for discriminated groups to bypass the customers' prejudicial preferences (becker1971economics, arrow1972some, cain1986economic). Consequently, customer discrimination has been considered "irrelevant" (arrow1972some) due to the ability of the disfavored labor force to work in occupations without contact with customers. The underlying assumption is that the labor supply is perfectly elastic; it is not costly for the discriminated labor force to avoid prejudice by forgoing jobs in the service sector while only considering working in jobs without consumer contact. However, workers differ in productivity; some have better service skills than others. A non-perfectly elastic labor supply implies that prejudicial preferences are reflected not only in employment (i.e., segregation) but also in wages.

I consider a two-sector economy: (i) a service sector—the one described above—in which employees are in contact with customers, and (ii) a manufacturing sector, in which employees are not in contact with customers. I assume that the manufacturing sector is perfectly competitive (employees are paid based on their productivity), and workers are perfect substitutes (as in the service sector).

The labor force consists of two groups: type White, which is a share ρ of the labor force, and type $\mathcal{B}$. The labor force is heterogeneous in terms of productivity in the manufacturing sector.[11] Yet, the productivity distributions of the labor forces across groups are identical.

Workers select their occupation based on the offered salary. Thus, the group-i workers that are the least productive in the manufacturing sector are the ones employed in the service sector. Let $\tilde{\omega}^i(m)$ denote the m^{th} least productive worker among

[11]The results in this part hold if the labor force is heterogeneous in production in both sectors.

the group-i labor force. Since productivity distributions are identical across group affiliations, $\tilde{\omega}^i(m_i) = \tilde{\omega}(m)$. As with regards to customers and firms, I impose that the distribution of productivity is scale-neutral; specifically, the distribution of manufacturing productivity (that is, the service sector labor supply) exhibits constant elasticity:

$$\tilde{\omega}(m) = Km^{\kappa}, \quad K > 0, \kappa > 0. \tag{2.6}$$

Labor Demand

In this part, I characterize the labor demand under different market structures, in general, and the relative labor demand for White workers, in particular. Specifically, I consider three different scenarios: (i) social planner, (ii) fully-collusive monopoly, and (iii) monopolistic competition. Although non-symmetric equilibria might exist, I restrict my attention to symmetric ones; $x_i = x$, $s_i = s$, $\forall i$, and $\lambda = nu(x, s)$.[12] In what follows, I characterize the optimal share of White employees, given wages and entry costs, under different market structures.

Social Planner

I start by characterizing the customer's optimal magnitude of production and the share of White employees, given wages and entry costs. The welfare-maximizing scenario serves as a benchmark when analyzing the behavior of a profit-maximizing firm. It is, in spirit, the differentiated-good variate of a perfectly competitive sector; all firms enjoy zero profits, and firms select the magnitude of production and labor composition to maximize consumer surplus.

Social Optimum is achieved by maximizing the representative consumer's surplus, subject to the zero-profit condition, denoted by $\mathcal{V}^c$. By equations 2.2 and 2.5:

$$\max_{x,s,n} \mathcal{V}^c = \frac{1}{1-\eta}(nu(x,s))^{1-\eta} - nc(x,s)$$
$$\max_{s,\lambda} \mathcal{V}^c = \frac{1}{1-\eta}\lambda^{1-\eta} - \lambda M(s).$$

where $M(s)$ denotes the minimum cost per utils, conditional on s,

$$M(s) = \min_x \frac{c(x,s)}{u(x,s)}. \tag{2.7}$$

[12]By symmetry and concavity of $u(x, s)$ and $c(x, s)$, the social optimum will be symmetric.

The objective function demonstrates that production and employment are selected to minimize the cost per utils. That is, to maximize the infra-marginal customer's utility relative to its costs. Therefore, the share of White employees that maximizes consumer surplus (given wages), denoted by $s^o(\omega)$, is characterized by the following first-order condition:

$$\hat{M}(s^o(\omega)) \equiv \frac{M'(s^o(\omega))}{M(s^o(\omega))} = 0. \tag{2.8}$$

An implication of equation 2.8 is that the welfare maximizing share of White employees is stationary and convex. I shall assume that the property holds throughout the relevant range of s. The monotonicity assumption permits to conduct comparative statics regarding different market structures.

Assumption. $\hat{M}(s) \equiv \frac{M'(s)}{M(s)}$ *monotonically increases in* $s \in [0, 1]$.[13]

Monopoly

In this part, I characterize the behavior of a fully-collusive monopoly. The monopoly decides on the optimal number of products sold in the sector, in addition to setting the quantity of production and the share of White employees. The firm's objective is to maximize total profits from the sector's goods. Let Π denote the sector's total profits. By symmetry and the inverse demand function (equation 2.4), Π is defined as follows:

$$\max_{x,s,n} \Pi = nx \underbrace{\lambda^{-\eta}\sigma(s)A(s)x^{\sigma(s)-1}}_{p} - nc(x, s)$$

$$\max_{s,\lambda} \Pi = \sigma(s)\lambda^{1-\eta} - \lambda M(s).$$

First-order conditions imply that the monopoly's optimal share of White employees (conditional on wages), denoted by $s^m(\omega)$, is characterized as follows:

$$\hat{M}(s^m(\omega)) - \frac{1}{1-\eta}\hat{\sigma}(s^m(\omega)) \equiv \frac{M'(s^m(\omega))}{M(s^m(\omega))} = \frac{1}{1-\eta}\frac{\sigma'(s^m(\omega))}{\sigma(s^m(\omega))} = 0, \tag{2.9}$$

while the necessary second-order conditions imply that the left-hand-side of equation 2.9 decreases in s. $\sigma'(s) < 0$ and the convexity of $\hat{M}(s)$) permit a comparison between $s^m(\omega)$ and $s^o(\omega)$, by equations 2.8 and 2.9.

[13]This assumption can be weakened by considering the second-order conditions at specific ranges of s.

Corollary 1. *Conditional on wages, the monopoly under-provides White service;* $s^m(\omega) < s^o(\omega)$.

Corollary 1 is driven by two different forces: (i) Spence distortion and (ii) product cross-substitution. To exemplify the Spence distortion, consider the aggregate demand of customers with heterogeneous tastes, as in the stylized model. The customer's optimal share of White employees (and quantity) is selected to minimize the ratio of the cost of production, $c(x, s)$, to the average utility of the infra-marginal customers $u(x, s)$. The monopoly does not consider the infra-marginal customers but solely considers the marginal customer. Due to the ordering of customers based, on average, on the taste for discrimination, the marginal customer's willingness to pay for service by a White employee is lower.

Consequently, the monopoly's relative labor demand for White employees is lower than under "perfect competition." Furthermore, employment is affected by cross-substitution patterns. Hiring more White employees to provide service for product i increases the demand for product i and decreases the demand for the sector's other products. The monopoly takes into account the negative externalities. Therefore, the relative labor demand for White employees further drops.

Equation 2.9 illustrates the importance of the heterogeneity in customer's taste for discrimination. If all customers have an identical bias against Black employees—$\sigma'(s) = 0$—the monopoly will hire the same share of White employees as a social planner. The difference in tastes gives rise to distortions driven by the ability of firms to influence both price and non-price attributes in their favor.

Monopolistic Competition

In this part, I characterize the behavior of profit-maximizing firms in a monopolistically-competitive setting. Unlike the case of a monopoly, the firms do not select how many products are sold (or how many firms are participating) in the sector, but rather each firm takes the competitive setting as given. As n is sufficiently large, the firms do not consider the externalities of their behavior on other firms (and thus surplus). The firm selects price and employment by group affiliation, given the number of products and the competitors' pricing and labor composition, to maximize its objective function, profits, denoted by π.[14]

$$\max_{x,s} \pi = x \underbrace{\lambda^{-\eta}\sigma(s)A(s)x^{\sigma(s)-1}}_{P} - c(x, s) = \lambda^{-\eta}\sigma(s)u(x, s) - c(x, s)$$

[14]By symmetry $\pi_i = \pi$.

Since entry is free, the marginal firm breaks even. By symmetry, all firms enjoy zero profits. The first order conditions and free entry condition are:

$$\frac{\partial \pi}{\partial x} = 0 \Rightarrow \sigma(s)\lambda^{-\eta}u_x(x,s) - c_x(x,s) = 0 \tag{2.10}$$
$$\frac{\partial \pi}{\partial s} = 0 \Rightarrow \sigma(s)\lambda^{-\eta}u_s(x,s) + \sigma'(s)\lambda^{-\eta}u(x,s) - c_s(x,s) = 0 \tag{2.11}$$
$$\pi = 0 \Rightarrow \sigma(s)\lambda^{-\eta}u(x,s) - c(x,s) = 0. \tag{2.12}$$

Similar to both the monopoly's and the social planner's behavior, firms set production to minimize the cost in utils (by equations 2.10 and 2.12). The threat of rival entry drives firms to produce efficiently. By plugging in equation 2.12 in 2.11, by definition of $M(s)$, and using the envelope theorem ($\hat{M}(s) \equiv \frac{M'(s)}{M(s)} = \frac{c_s(x,s)}{c(x,s)} - \frac{u_s(x,s)}{u(x,s)}$) the firm's optimal share of White employees (given wages), denoted by $s^c(\omega)$, is characterized as follows:

$$\hat{M}(s^c(\omega)) - \hat{\sigma}(s^c(\omega)) \equiv \frac{M'(s^c(\omega))}{M(s^c(\omega))} = \frac{\sigma'(s^c(\omega))}{\sigma(s^c(\omega))} = 0. \tag{2.13}$$

As in the case of a monopoly, by $\sigma'(s) < 0$, firms operating in a monopolistic competitive setting decide to under-provide White service, $s^c(\omega) < s^o(\omega)$. The Spence distortion—the divergence between the infra-marginal customers and the marginal customer—results in profit-maximizing firms under-providing service. However, unlike a fully-collusive monopoly, the firm does not consider cross-substitution patterns. Consequently, it has a higher labor demand for White employees.

Corollary 2. *Conditional on wages, a firm operating in a monopolistically-competitive firm under-provides White, yet hires a higher share of White employees relative to a monopoly;* $s^m(\omega) < s^c(\omega) < s^o(\omega)$.

Equilibrium Employment and Wages

In this part, I examine the effect of competition, through its impact on the relative demand for White employees, on equilibrium labor market outcomes: relative wages and employment. Equilibrium relative wage and employment are set such that labor market clearing condition holds:

$$Z^j(s|\tilde{\Delta}, f) = 0$$
$$\frac{s}{1-s} = \frac{\rho}{1-\rho}(\tilde{\Delta})^{\frac{1}{\kappa}},$$

where $\tilde{\Delta} = \frac{\Delta\omega}{\omega^b}$ denotes the wage gap and $Z^j(s|\tilde{\Delta}, f)$ is the sector's excess demand from White employees, which depends on the market structure. Specifically,

$$Z^j(s|\tilde{\Delta}, \omega^b, f) = \begin{cases} -\hat{M}(s|\tilde{\Delta}, f) & \text{for social planner,} \\ \hat{\sigma}(s) - \hat{M}(s|\tilde{\Delta}, f) & \text{for a monopolistically-competitive setting,} \\ \frac{1}{1-\eta}\hat{\sigma}(s) - \hat{M}(s|\tilde{\Delta}, f) & \text{for monopoly.} \end{cases}$$

I conduct a few comparative statics to assess how the equilibrium gap in wages and employment differs by market structure. To do so, I characterize $M(\cdot)$. By definition of $M(\cdot)$, the optimal level of production given the share of White employees—$x(s)$—is given by:

$$x(s) = \arg\min_x \frac{\delta(x+f)\omega^b(1+s(\tilde{\Delta}-1))}{A(s)x^{\sigma(s)}} = \frac{f\sigma(s)}{1-\sigma(s)}. \tag{2.14}$$

Consequently,

$$M(s|\tilde{\Delta}, \omega^b, f) = \frac{\delta\omega^b(1+s(\tilde{\Delta}-1))f^{1-\sigma(s)}}{(1-\sigma(s))^{1-\sigma(s)}\sigma(s)^{\sigma(s)}A(s)}. \tag{2.15}$$

Equation 2.15 demonstrates that the (log) minimum cost per utils of providing White service increases with the wage gap, $\tilde{\Delta}$:

$$\frac{\partial}{\partial\tilde{\Delta}}\hat{M}(s|\tilde{\Delta}, f) = \frac{\tilde{\Delta}-1}{(1+s(\tilde{\Delta})-1)} > 0.$$

The relationship between the wage gap and the cost of production is straightforward—an increase in the wage gap results in a higher marginal cost of providing White service.[15] Since the relative labor supply of White employees increases with $\tilde{\Delta}$, a unique equilibrium labor allocation exists. By corollary 1, the excess labor demand increases with competition; $Z^m(s|\tilde{\Delta}, \omega^b, f) < Z^c(s|\tilde{\Delta}, \omega^b, f) < Z^o(s|\tilde{\Delta}, \omega^b, f)$. Since labor supply is not perfectly elastic, the labor market clearing condition implies that the relatively lower labor demand for White employees is reflected in both employment and wages.

Corollary 3. *In equilibrium, the wage and employment gaps increase with the market's competitiveness; the service sector's wage and employment gaps are lower when a monopoly operates relative to a social planner. Furthermore, a monopolistically-competitive setting implies that wage and employment gaps are between both scenarios;* $s^m < s^c < s^o$, *and* $\tilde{\Delta}^m < \tilde{\Delta}^c < \tilde{\Delta}^o$.

[15] It should be noted that in any equilibrium (regardless of f), the wage gap is positive, $\omega^w > \omega^b$. Otherwise, firms hire only White employees. Therefore, the labor market clearing condition implies $\omega^w > \omega^b$.

After establishing the relationship between different market structures and labor outcomes, I examine the impact of entry costs on wage and employment gaps in a monopolistically-competitive setting. Following equation 2.15, the minimum cost per utils of providing more White service increases:

$$\frac{\partial}{\partial f}\hat{M}(s|\tilde{\Delta}, f) = \frac{-\sigma'(s)}{f} > 0.$$

Intuitively, as more units are lost due to production, fewer firms operate in the sector, and each one produces more units of the good (equation 2.14). Although customers prefer White service, the benefit from increasing the share of White employees drops with production, as $\sigma'(s) < 0$. Hence, as entry cost and production increase, the firm faces a marginal customer with a lower, on average, taste for discrimination. Therefore, the marginal return for providing more White service drops with entry barriers. Since labor demand drops, the inelastic labor supply implies that the sector's equilibrium wage and employment gaps also drop.

Corollary 4. *Let $s^m(f)$ and $\tilde{\Delta}(f)$ denote the equilibrium share of White employees in the service sector and the equilibrium wage gap, respectively, as a function of entry barriers f. As entry costs increase, both the share of White employees and the wage gap drop; $s^m(f) > s^m(f')$ and $\tilde{\Delta}(f) > \tilde{\Delta}(f')$, $\forall f' > f$.*

Model Predictions

In this part, I highlight the model predictions I take to the data. To do so, I consider the realistic case of both prejudiced customers and employers. The standard analysis of the economics of discrimination mainly considers employer taste-based discrimination. The relationship between competition in the product market and discrimination in the labor market depends on the driving force of prejudice. As previous studies indicate, when prejudice is employer-driven, competition mitigates discrimination (Becker (1957)) and is possibly eliminated (Arrow (1972)). In contrast, as I showed earlier, competition exacerbates discrimination when prejudice is customer-driven.

I consider 4 (2-by-2) scenarios. I characterize the wage gap for two market structures: monopoly and monopolistic competition. I examine the wage gap for each market structure when customer taste-based discrimination is not present (i.e., when customers do not interact with employees). I show that the relationship between competition and discrimination is ambiguous when employers and customers are prejudiced. Nevertheless, the drop in the wage gap is unambiguously more significant in non-customer-contact occupations.

In the spirit of Becker, 1957 and Arrow, 1972, consider the case of a fully-collusive monopoly with a taste for discrimination against Blacks. The employer's taste for discrimination can be expressed as if the effective cost of hiring Black employees is $\omega_b(1+d)$, where $d > 0$.[16] The monopolistic market structure allows the employer's prejudiced behavior to persist, as it can over-charge customers to fund the inefficiencies that arise from its taste-based discrimination. In contrast, prejudiced employers are driven out of the market in a competitive economy. Thus, if prejudice is driven solely by employers, the sector's wage gap is d under a monopoly market structure; $\frac{\omega^w}{\omega^b} = 1+d$. However, the wage gap does not persist in a competitive market; $\frac{\omega^w}{\omega^b} = 1$.

The relationship between the product's market structure and the within-sector wage gap is ambiguous when customers are prejudiced. Using equations 2.9 and 2.15, the fully-collusive monopoly's demand for White employees, given wages, is characterized as follows:

$$\frac{\tilde{M}'(s^m|\omega^w, \omega^b, d)}{\tilde{M}(s^m|\omega^w, \omega^b, d)} = \frac{1}{1-\eta}\frac{\sigma'(s^m)}{\sigma(s^m)}, \tag{2.16}$$

where

$$\tilde{M}(s|\omega^w, \omega^b, d) = \frac{\delta\omega^b(1+d)(1+s\frac{(\omega^w-\omega^b(1+d))}{\omega^b(1+d)})f^{1-\sigma(s)}}{(1-\sigma(s))^{1-\sigma(s)}\sigma(s)^{\sigma(s)}A(s)}. \tag{2.17}$$

The monopoly's behavior is similar to before. The only difference is that the monopoly up-weighs the cost of hiring a Black employee by a factor of $1+d$ (that is, $\tilde{M}(s|\omega^w, \omega^b, 0) = M(s|\omega^w, \omega^b)$). As a result, it is unclear whether the monopoly's labor demand for White employees is higher or lower relative to the sector's labor demand for White employees under a monopolistically-competitive market structure. Using equations 2.13 and 2.16, I can characterize the relationship between the two. A monopoly hires a higher (lower) share of White employees if

$$(1-\eta)\Big|\frac{\tilde{M}'(s|\omega^w, \omega^b, d)}{\tilde{M}(s|\omega^w, \omega^b, d)}\Big| > (<)\Big|\frac{M'(s|\omega^w, \omega^b)}{M(s|\omega^w, \omega^b)}\Big|. \tag{2.18}$$

Equation 2.18 highlights the two forces generating a divergence between monopolistic and competitive labor decisions. $\eta \in (0,1)$ illustrates that a monopoly has a lower demand for White employees as it considers the negative externalities of the cross-substitution patterns. In contrast, for $d > 0$, $\big|\frac{\tilde{M}'(s|\omega^w,\omega^b,d)}{\tilde{M}(s|\omega^w,\omega^b,d)}\big| > \big|\frac{M'(s|\omega^w,\omega^b)}{M(s|\omega^w,\omega^b)}\big|$, which identifies the monopoly's preference to hire more White employees.

[16] This functional form holds without loss of generality.

Wage gap	No	Prejudice Customers
Monopoly	$1+d$	r^m
Competition	1	r^c
Change	$\frac{1}{1+d}$	$\frac{r^c}{r^m} > \frac{1}{1+d}$

Although it is unclear whether the wage gap increases or decreases following a competition shock, when the customers are prejudiced, it is possible to compare the impact of competition on the wage gap when customers drive discrimination. As illustrated in the matrix above, when customers are not prejudiced (or when there is no contact between customers and employees), the wage gap drops by a factor of $\frac{1}{1+d}$ as the sector's discriminatory preferences are eliminated by competition. However, when the market structure changes from monopolistic to competitive, the wage gap drops less (and might even increase) when customers are prejudiced.

Proposition 3. *Let r^m and r^* denote the wage gap under monopolistic and competitive market structures. Competition in the product market decreases the sector's wage gap by more when the customer's prejudice does not play a role;* $\frac{1}{1+d} < \frac{r^c}{r^m}$.

Intuitively, the change in the wage gap following a competition shock depends on two factors: the monopoly's and the customer's taste for discrimination. The two factors are independent; when considering the implications of a prejudicial employer, the unexplained difference between the monopolistic and competitive wage gap is solely due to the customer's prejudicial tastes. To see this, notice that when scaling the wage of White employees by a factor of $1 + d$, the residual difference is driven by η. By equation 2.17, $\tilde{M}(s|\omega^w(1+d), \omega^b, d) = (1+d)M(s|\omega^w, \omega^b, d)$. Therefore, $|\frac{\tilde{M}'(s|\omega^w(1+d),\omega^b,d)}{\tilde{M}(s|\omega^w(1+d),\omega^b,d)}| = |\frac{M'(s|\omega^w,\omega^b)}{M(s|\omega^w,\omega^b)}|$. Since the monopoly takes into account the negative externalities of providing White service for one product on the rest of its commodity—$\eta \in (0,1)$—competition in the product market decreases the sector's wage gap by less when customer's taste-based discrimination plays a role (and might even increase).

2.4 US Banking Deregulation

In this part, I present a brief overview of the US banking deregulation process, which I exploit in the empirical analysis. States began limiting the banks' within-state behavior by introducing branching regulations in the nineteenth century. The regulation favored small, weakly capitalized, inefficient banks, as it limited competition.[17] Up until 1970, only 12 states permitted unrestricted state-wide branching, while some states (the "Unit banking" states) imposed significant limitations by prohibiting any branching.[18]

Since early 1970s, states relaxed branching restrictions. In general, intra-state deregulation occurred in three stages. First, states permitted the formation of multi-bank holding companies (MBHC). Despite joint ownership, the MBHC's banks were not allowed to operate jointly. Later, banks were allowed to branch only via mergers and acquisitions (M&A). Finally, states allowed unrestricted (*de novo*) intra-state branching.

Deregulation impacted the banking sector's market structure, profitability, and efficiency. The most critical stage was branching deregulation via M&A (Kroszner and Strahan (1999)). Among all the deregulation stages, Black and Strahan, 2001 find that M&A branching deregulation had the most substantial impact on employees' wages. It also has the most considerable effect on the industry's gender wage gap. M&A branching deregulation evolved substantially over the late twentieth century. In 1969, only 12 states permitted M&A branching, while all states allowed it by 1999 (see table 2.1).

Unlike the case of intra-state M&A branch deregulation, Black and Strahan, 2001 do not document a significant impact of inter-state deregulation on wages, in general, and the gender wage gap, in particular. They state two possible reasons: (i) inter-state deregulation followed M&A intra-state deregulation and thus did not have a substantial effect on the banking industry, and (ii) inter-state deregulation occurred through a relatively short period, limiting the option to statistically differentiate it from time-trend fixed effects. Table 2.1 presents the year of intra-state M&A branch and inter-state deregulation for each state.

Several factors drove the significant deregulation of the banking industry. Kane, 1996 suggest that the wave of bank and thrift failures in the 1980s increased public awareness of the advantages that large, diversified banks have over small, local banks. Kroszner and Strahan, 1999 suggest that new technologies in both deposit and loans changed the legislators' political incentives from protecting the small, local banks to serving large, multi-state banks.

The Black and Strahan, 2001 identification strategy uses variation in the timing of inter-state and intra-state deregulation across states as a competition shock to the banking industry, which is exogenous to the state's labor market conditions. They

[17] Economides, Hubbard, and Palia, 1996 show that the 1927 McFadden Act, which gave states the authority to regulate national banks' branching power, was favored by states characterized by many weakly capitalized, small banks.

[18] The "Unit banking" states: AR, CO, FL, IL, IA, KS, MN, MO, MT, NE, ND, OK, TX, WI, WV, and WY.

support this assumption by stating that there is no correlation at the state-level between the timing of deregulation and (i) unionization in the banking industry or overall unionization rates and (ii) the level of banking wages.

2.5 Data

In this paper, I take the model's predictions to the data. I do so by empirically examining the relationship between competition in the product market and the Black-White labor wage and employment differentials in general and consumer-contact occupations in particular. I exploit the US banking sector deregulation expansion of state-level restrictions *a la* Black and Strahan, 2001 to study the relationship between product market competition and labor market discrimination. In my model, product market competition differentially impacts labor market outcomes based on the job's interaction with customers, thus adding another dimension to Black and Strahan, 2001. I evaluate the systematic change in the Black-White wage and employment gaps across different occupations based on the job's requirement to deal with customers. I do so by taking advantage of the heterogeneity in the consumer contact level across occupations within the banking sector. Some of the bank's employees constantly interact with customers, while others do not. This property allows me to evaluate the impact of competition on labor market discrimination when prejudice is customer-driven.

This section presents the data and measures used throughout the paper. In addition to the state-year deregulation status (table 2.1), I take advantage of variables collected from three separate sources; US Censuses and the annual American Community Surveys (ACS), Occupational Information Network (O*Net), and the General Social Survey (GSS).

Census and American Community Survey

To assess the effect of increased product market competition on racial labor market discrimination, I use data from the decennial US Censuses from 1960 through 2000 and the 2010-2012 3-year ACS.[19] I restrict the sample to contain 26 to 55-year-old, Black and White, US-born men who do not live in a group quarter.[20] I further exclude individuals currently unemployed, self-employed, working in the military, or

[19] I use the following US censuses samples: 1960 5%, 1970 1% state-form 1 and 2, 1980 5% state, 1990 5% state, and 2000 5%.

[20] I consider observations of females when examining the impact of product market competition on employment as gender substitution patterns might not be identical across races.

out of the labor force. In addition, I drop observations with missing variable values (sex, wage, education, occupation, race). As in Black and Strahan, 2001, I exclude observations (i) from Delaware and South Dakota, as these states experienced a dramatic expansion in their banking sectors during the 1980s, and (ii) from states in which the deregulation occurred in the same year.[21] In addition, I omit observation from 1960 if the deregulation occurred beforehand.[22] I weigh the data using the survey weights provided by the Censuses and the ACS, respectively.

In addition to the variables described above, the US Census and ACS datasets contain information on the individual's occupation (1990 3-digit code), industry (1990 3-digit code), and wage and salary income in the previous year. Using these variables, I can compare the pre and post-deregulation Black-White labor outcome gaps within the banking industry relative to other industries (as in Black and Strahan (2001)). Furthermore, using data on the occupation's work context (see O*Net), I can analyze how the change in the Black-White wage and employment gaps are related to the occupation's required contact with customers.

Before examining how deregulation affected the Black-White wage gap in the banking sector, I study how deregulation affected the banking sector's employees' wages. As in Black and Strahan, 2001, I use the dates reported in table 2.1 to construct an indicator variable for states permitting intra-state branching by M&A (M&A) and an additional one for states permitting inter-state banking (INTER). In addition, since pre-deregulation, the Unit banking states had tighter constraints. In terms of regulation, I construct an additional indicator variable that equals one for the Unit banking state (Unit), as in Black and Strahan (2001).

I regress the log wage and salary income on the indicators mentioned above, interacting with an indicator variable on whether the individual is employed in the banking sector or not. I also control for individual-level characteristics, state-specific and year-specific components of the banking sector's wages with two fixed effects, and race-year-state level dummies. The regression model is similar to the differences-in-differences analysis in Black and Strahan, 2001. The results are reported in table 2.2.

In column 1, I report the regression results of log wage regression on the BANK-M&A interaction. I find a significant 4 percent drop in wages in the banking sector following intra-state branching deregulation. In column 2, I provide the results, allowing differential effects of branching deregulation by Unit state classification. I find that states that began with tighter restrictions on branching had a higher wage

[21] The expansion in Delaware and South Dakota is attributed to liberal usury laws. As a result, credit card operations were moved to these states.

[22] As it is unknown to me if the deregulation occurred before or in 1959.

decline (although not statistically significant). For non-Unit states, I find a non-significant 2 percent drop in wages, while a significant 6 percent drop for the Unit states, following intra-state branching deregulation via M&A.

In columns 3 and 4, I report the estimation results of the models reported in columns 1 and 2, respectively, when including inter-state deregulation. In general, the frequency of the decennial data limits the option to separately identify and quantify the impact of each deregulation. Following intra-state branching deregulation, there was a significant 4 percent drop in wages in the banking sector, while only a non-statistically significant 2 percent drop following permission of inter-state branching deregulation (column 3), consistent with earlier findings on the effects of banking deregulation on competition within the banking industry. The aggregate effect of deregulation (intra-state M&A and inter-state) results in a 6 percent wage drop. When considering the differential effect on interstate M&A, I find that the impact of deregulation is significant for Unit states; a 6 percent drop in wages following intra-state M&A deregulation and an 8 percent drop when considering the combination of both types of deregulation (intra-state M&A and inter-state). This is not the case for non-Unit states: following deregulation, the drop in wages within the banking sector is lower and not statistically significant. The results suggest that wages in the banking sector fell as competition increased. The banks could not continue to share rents with the employees.

Occupational Information Network

In order to characterize each occupation by its level of consumer contact, I use the data collected from O*Net, sponsored by the US Department of Labor/Employment and Training Administration (USDOL/ETA). The O*Net database contains information on occupational worker characteristics (e.g., abilities, interests, skills, knowledge, education) and occupational requirements (work activities, work context, and organizational context). In this paper, I use the 1998 O*Net dataset. Using Deming, 2017 's replication cross-walk, I can match the 1998 O*Net's occupational code to the 1990 census 3-digit occupational code level, which allows me to merge these measures into my dataset.[23]

I focus occupational task measures that are relevant to customer contact: ***"Deal With External Customers"***. Deming, 2017 uses this variable, among others, to generate a measure of occupational tasks that require social skills. Since I focus on the effect of competition on discrimination when customers drive prejudice, I

[23] There are two 1990 census 3-digit occupational codes that are unmatched and thus omitted; professionals, n.e.c., and inspectors, n.e.c.

do not consider social skill requirements in general but focus on the occupation's requirement to deal with customers. The value of the 1998 O*Net version variable ranges from 0 to 5 (higher values describe more contact with external customers). I convert the customer contact measure into a z-score using the unweighted mean and standard deviation across occupations. The units of measure of customer contact are standard deviation differences relative to all other occupations.

Figure 2.1 depicts the distribution of the customer index within the banking sector's employees among occupations comprising at least 1 percent of the banking sector's labor force. The banking industry's occupations are characterized by high consumer contact relative to the rest of the economy (mean = 0.482, and median = 0.474). However, there is substantial variation in the sector's required magnitude of contact with external customers; the occupation's requirement to deal with customers. Four of the 16 occupations with the largest share of employees are characterized by a negative z-score (computer systems analysts, bookkeepers/ accountants, equipment operators, and computer software developers). Six have a z-score higher than one (bank tellers, financial services sales, customer service reps, chief executive/ public administrators, marketing /advertising /public relations, and managers/ administrators, n.e.c.). The substantial variation in customer contact is a critical component of the analysis. It provides an opportunity to isolate the effect of competition on the Black-White wage and employment gaps when prejudicial customers induce discrimination.

General Social Survey

I further examine whether the rise in the banking sector's Black-White wage and employment gap following the increase in competition is higher in states with strong discriminatory attitudes toward Blacks . I generate a prejudice index for each state, in the spirit of Charles and Guryan, 2008, using the GSS datasets.

The GSS is a nationally representative survey of US adults conducted since 1972. Among the questions asked, the GSS collects responses from survey questions about matters strongly related to the respondents' racially prejudicial tastes. Following Charles and Guryan, 2008, I use data from multiple waves (1972,1977,1982,1985,1988-1991,1993,1994 and 1996) for White individuals over the age of 18. Over this period, respondents answered 26 different questions relating to some aspect of racial feelings. Some of these questions focus on governmental intervention and race; others are not asked consistently over time. As in Charles and Guryan, 2008, I do not focus on these questions. Instead, I take advantage of four questions unrelated to government policy, which are asked consistently in each wave. (i) Do you think there should be laws against marriages between Blacks and Whites? (ii) If your party nominated a

Black for president, would you vote for him if he were qualified for the job? (iii) Aggregation of three questions: whether you would object to sending your kids to a school that had few/half/most Black students? (iv) Agree? White people have the right to keep Black people out of their neighborhoods, and Blacks should respect that right. I code the prejudice measure from each question as follows: (i) $d^1_{it} = 1$ if yes (0 if no), (ii) $d^2_{it} = 1$ if no (0 of yes), (iii) $d^3_{it} = k$, where k equals the number of questions the respondent agreed, and (iv) d^4_{it} equals four if agree strongly, three if agree slightly, two if disagree slightly, and one if disagree strongly.

I closely follow Charles and Guryan, 2008 and use the response to the four questions above to generate an individual prejudice measure as follows:

$$D_{it} = \frac{1}{4}\sum_{k=1}^{4} \frac{d^k_{it} - \bar{d}^k_{77}}{\sigma^k_{72}},$$

where $\bar{d}^k_{77}$ is the 1977 sample average measure for question k, and σ^k_{72} is the 1972 sample standard deviation with regards to question k.[24]

Lastly, I construct a measure of the aggregate prejudice level, as in Charles and Guryan, 2008; the prejudice level of the "marginal" White discriminator. The measurement is constructed in three stages. First, I regress D_{it} on a full set of year dummies, $D_{itr} = \alpha_t + \varepsilon_{itr}$. Then, I calculate the cumulative distribution function of the regional's prejudice index using the residuals from the previous regression, $F_r(\varepsilon)$.[25] Finally, I use the US censuses between 1970 and 2000 to calculate the average share of Blacks among the total labor force in each state, p_s. $F_r^{-1}(p_s)$ is the states prejudice index.

I divide and classify the states into three groups: low, medium, and high levels of prejudice. The low-level prejudice states include 17 states with the minimum prejudice index value. Many states share this value. Some are a part of regions with low levels of prejudiced tastes. More importantly, the share of Blacks among the local labor force in these states is meager. The set of prejudiced states with high discriminatory attitudes towards Blacks includes the ten states with the highest prejudicial index. The medium group includes the rest of the states. Figure 2.2 presents the state classification by the three groups. The high-prejudice states are

[24] By normalizing by the standard deviation in the first year and not the overall standard deviation—as done in Charles and Guryan (2008)—I avoid a mechanical relationship between trends in responses and the weight the question receives in the overall aggregate. Charles and Guryan, 2008 chose 1977 as the normalization year because it was the year in which the largest number of prejudice questions were asked.

[25] The restrictive GSS data I use does not contain the individual's state, but rather the respondent's region.

generally from regions with high prejudicial tastes, while low-prejudice states are mainly from regions with low prejudicial tastes; New England, Pacific, and Mountain. However, the relationship is not one-to-one. For instance, Kentucky, which is part of a highly prejudicial region, is not part of the highly prejudicial states as the share of the Black labor force is low (6.8%). Iowa, which is not part of a low prejudicial region, is classified as a low prejudicial state due to its low share of the labor force (1.3%).

2.6 Empirical Analysis

In this section, I test the model's predictions. I examine how the banking sector's Black-White labor outcomes changed following the competition shock; deregulation in the US banking sector. Specifically, I examine how the Black-White wage and employment gaps systematically changed depending on the occupation's customer contact characteristics. This analysis is composed of two parts. Previous studies suggest that the Black-White gaps would drop in all occupations following deregulation that intensifies competition. My model predicts that the Black-White gaps would drop in occupations with no or limited contact with customers. Nevertheless, the Black-White wage and employment gaps would rise in service occupations characterized by substantial contact with customers.

First, I study the change in the Black-White wage gap by conducting the Black and Strahan, 2001 analysis for different occupations based on the jobs required to deal with customers. I implement customer contact in the Black and Strahan, 2001 model by including the customer index as an explanatory variable. I then examine the change in the Black-White employment gap following deregulation. Lastly, I examine how systematic change in the Black-White wage gap interacts with state-level prejudicial preferences.

Black and Strahan, 2001 Model

I first test how the banking sector' Black-White wage gap evolved following deregulation, in the spirit of Black and Strahan, 2001. Specifically, I estimate the following wage equation:

$$\begin{aligned}\log(\text{Wage})_{ist} &= \alpha_0 + \alpha_1 \text{BANK}_{ist}\text{M\&A}_{ist} + \beta \text{BANK}_{ist}\text{M\&A}_{ist}\text{Black}_{ist} \\ &+ \delta_t X_{ist} + \tilde{\delta}_t X_{ist}\text{Black}_{ist} + \nu_{st} + \tilde{\nu}_{st}\text{Black}_{ist} + \gamma_s \text{BANK}_{ist} \\ &+ \tilde{\gamma}_s \text{BANK}_{ist}\text{Black}_{ist} + \kappa_t \text{BANK}_{ist} + \tilde{\kappa}_t \text{BANK}_{ist}\text{Black}_{ist} + \varepsilon_{ist}. \quad (2.19)\end{aligned}$$

Black_{ist} is a dummy variable that equals 1 if individual i is Black. BANK_{ist} is a dummy variable that equals 1 if individual i is employed at the banking sector. $\text{M\&A}_{st}$ is a dummy variable that equals one if the state permits intra-state branching via M&A at time t. X_{ist} includes a set of covariates; six age and education dummies, each interacted with year and race, a race-state-year, race-bank-state, and race-bank-year fixed-effects.

The coefficient β on the interaction term between Black, BANK, and M&A generates the triple difference coefficient, which provides an estimate of how the banking sector's Black-White wage gap changes, relative to the general economy, following the intra-state branching deregulation.

To study how competition affects the Black-White wage gap when customer-driven prejudice might be in play, I estimate equation 2.19 for three separate groups: occupations with low, medium, and high customer index. I divide my sample to three thirds (by person sample weight); occupations are classified as low contact (low CI) if their customer index is in the bottom third (≤ -0.356), high contact (high CI) if their customer index is in the top third (≥ 0.980), and medium contact (medium CI) otherwise.

The results are reported in figure 2.3. The top panel reports the estimation result of equation 2.19 for the entire sample. When examining the population at the aggregate, the Black-White wage gap in the banking sector dropped by 4.6 percent (although not statistically significant) following the intra-state branching deregulation. However, as the depicted results suggest, analyzing the change at the aggregate might be misleading. Estimation of the model presented in equation 2.19 for three separate groups suggests that the Black-White wage gap dropped by more than 23 percent following intra-state deregulation in low customer contact occupations. In contrast, in high customer contact occupations, the Black-White wage gap increased by 14 percent following M&A deregulation. Although the increase is not statistically significant, the estimate juxtaposes previous studies.

The center and bottom panels of figure 2.3 report the estimation results of an extension of model 2.19, which permits differential effects for non-Unit and Unit states, respectively. For non-Unit states, the change in the Black-White wage gap following deregulation is not statistically significant but similar in sign to previous estimates. In the aggregate, the Black-White wage gap decreased by more than 4 percent. For low customer contact occupations, the Black-White wage gap dropped and increased when considering high customer contact jobs. Similar to the change in wages (see table 2.2), the change in the Black-White wage gap following deregulation is more substantial for the Unit states. For low customer contact occupations, the Black-White wage gap dropped and significantly increased for high customer contact jobs—by more than 27 percent. A comparison between column 2 and 4 suggest that the change in the wage gap following a positive competition shock depends on whether prejudice is solely employer-driven or customer discriminatory preferences play a role. Furthermore, the results presented in column 4 suggest that in high-customer occupations, the importance of customer-driven prejudice outweighs that of prejudicial employers, as competition increases the wage gap when prejudice is customer-driven.

To further demonstrate the non-linear relationship between the Black-White wage gap and deregulation across occupations, I estimate equation 2.19 among different subsets of the sample, depending on the customer index. Specifically, I estimate the model for observation with a customer index higher than multiple lower bounds. I consider different thresholds, from -1.2 to 1.2, in steps of 0.1. The results depicted in figure 2.4 demonstrate the non-linear relationship. The high customer contact jobs drive the rise in the Black-White wage gap following M&A deregulation. As done earlier, I re-estimate the model, allowing for a differential effect for Unit and non-Unit states. The results indicate, both in magnitude and statistical significance, that the Black-White wage gap rose substantially after deregulation when the competition shock was more dominant. For Unit states, the Black-White wage gap increased substantially following deregulation when considering all occupations with customer contact ≥ 0.3. Consistent with earlier findings, the non-linear trends indicate that high customer contact occupations play the primary role.

Interaction with Customer Index

In the first part of the analysis, I show how the change in the Black-White wage gap systematically differs across occupations by constructing sub-samples based on the job's characterized contact with consumers and analyzing the change in the Black-White wage gap for each one. In this part, I test how the banking sector' Black-White wage gap evolved (following deregulation) differentially by estimating the following wage equation for the entire sample:

$$\begin{aligned}\log(\text{Wage})_{ist} &= \alpha_0 + \alpha_1 CI_{ist} + \alpha_1 \text{BANK}_{ist} CI_{ist} + \alpha_3 \text{BANK}_{ist} \text{M\&A}_{ist} \\ &+ \alpha_4 \text{M\&A}_{ist} \text{CI}_{ist} + \alpha_5 \text{Black}_{ist} \text{CI}_{ist} + \alpha_6 \text{Black}_{ist} \text{M\&A}_{ist} \text{CI}_{ist} \\ &+ \alpha_7 \text{Black}_{ist} \text{BANK}_{ist} \text{M\&A}_{ist} + \alpha_8 \text{Black}_{ist} \text{BANK}_{ist} \text{CI}_{ist} \\ &+ \alpha_8 \text{BANK}_{ist} \text{M\&A}_{ist} \text{CI}_{ist} + \beta \text{BANK}_{ist} \text{M\&A}_{ist} \text{Black}_{ist} \text{CI}_{ist} \\ &+ \delta_t X_{ist} + \tilde{\delta}_t X_{ist} \text{Black}_{ist} + \nu_{st} + \tilde{\nu}_{st} \text{Black}_{ist} + \gamma_s \text{BANK}_{ist} \\ &+ \tilde{\gamma}_s \text{BANK}_{ist} \text{Black}_{ist} + \kappa_t \text{BANK}_{ist} + \tilde{\kappa}_t \text{BANK}_{ist} \text{Black}_{ist} + \varepsilon_{ist}. \quad (2.20)\end{aligned}$$

CI_{ist} is the occupation's customer contact index. This model effectively adds an additional dimension to the model presented in equation 2.19. The coefficient of interest is β; it provides an estimate of how the Black-White wage systematically changed following the intra-state branching deregulation across different occupations within the banking sector, relative to the rest of the economy, based on the occupation's required contact with customers.

The results are presented in table 2.4. Consistent with the model presented before, the results in column 1 indicate that the change in the Black-White wage gap is 11.1 percent higher following intra-state M&A deregulation when considering occupations with one standard deviation higher customer index. In column 2, I estimate the model when including intra-state M&A and inter-state branching deregulation. Both coefficients are negative but not statistically significant. Nevertheless, when considering the total effect of deregulation (inter-state and intra-state M&A), the change in the Black-White wage gap is 11.7 percent higher following both types of deregulation when considering occupations with one standard deviation higher customer index (row F). Columns 3 and 4 report an extension of the estimation models reported in columns 1 and 2, which permit differential effects for Unit and non-Unit states. As reported in figure 2.3, the effect of deregulation is substantial when considering the Unit states (although the difference is non-statistically significant). When considering solely intra-state branching via M&A (column 3), the change in the Black-White wage gap following deregulation is 9.7 percent higher (and not statistically significant) in non-Unit states, while 16 percent higher in Unit states when considering occupations with one standard deviation higher customer index. The results are similar in spirit when taking into account inter-state branching deregulation as well; When considering the total effect of deregulation (inter-state and intra-state M&A), the change in the Black-White wage gap following deregulation is 12.6 percent higher (and not statistically significant) in non-Unit states, while 16.2 percent higher in Unit states when considering occupations with one standard deviation higher customer index. In columns 5 and 6, I report the estimation results of the model presented in columns 1 and 2, but when solely considering bank employees. The estimates are similar in value and are statistically significant (although only at a 90% level).

The results in this part suggest that, in contrast to previous literature, an increase in competition might exacerbate discrimination when prejudice is customer-driven and outweigh the benefit of competition (via mitigation of employer-driven discrimination).

Interaction with Customer Index and Prejudicial Preferences

So far, I show that following deregulation, which intensified competition, the Black-White wage gap decreased in low customer-contact occupations yet increased in high customer-contact occupations. In this part, I examine the relationship between the change in the Black-White wage gap, the occupation's customer contact characteristics, and taste for discrimination. To do so, I generate a prejudice index for each state, in the spirit of Charles and Guryan, 2008, using the GSS datasets. As mentioned earlier, I divide the states into three groups, based on the Charles and Guryan, 2008 index; High, medium, and low prejudice states. If prejudice is customer-driven, a competition shock should exacerbate discrimination in high-customer-contact occupations relative to low-customer-contact occupations.

The analysis is done in two steps. First, I test how the banking sector' Black-White wage gap evolved (following deregulation) differentially by the state's measure of prejudice by estimating equation 2.20 separately for each subset. The results are presented in table 2.4. Columns 1-3 report the results for the high, medium, and low prejudice states sub-sample, respectively. The results indicate that the Black-White wage gap systematically changes with the states' taste for discrimination. In high prejudice states (column 1), the change in the Black-White wage gap is 24.7 percent higher following intra-state M&A deregulation when considering occupations with one standard deviation higher customer index. In medium (column 2) and low (column 3) prejudice states, the change in the Black-White wage gap is much lower and not statistically significant.

The results in columns 1-3 suggest that the states with a high taste for discrimination are the driving force behind the increase in the Black-White wage gap in customer-contact occupations. To further investigate these patterns, I extend the wage equation (2.20) by allowing differential effects for high prejudice states relative to others. The result is reported in column 4. The difference between highly prejudiced states and the rest is substantial. In high prejudice states (column 1), the change in the Black-White wage gap is 25.7 percent higher following intra-state M&A deregulation when considering occupations with one standard deviation higher customer index, while only 6.2 percent (and not statistically significant) in the other states. In column 5, I report the estimation result of the model presented in column 4, but when solely considering bank employees. The results are consistent with column 4, which considers the entire economy. The difference between highly prejudiced states and the rest of the states is substantial (25.5 percent and statistically significant). In high prejudice states (column 1), the change in the Black-White wage gap is 25.5 percent higher following intra-state M&A deregulation when considering occupations with one standard deviation higher customer index, while only 3.8 percent (and not statistically significant) in the other states. The results in this part indicate that the high prejudice states are the main driving force of the increase in the Black-White wage gap among the banking sector's customer contact occupations. Furthermore, the result in this part serves as a robustness test, ruling out alternative explanations for the increase in the Black-White wage gap in occupations that require customer contact.

Employment

Previously, I documented that the Black-White wage gap increases (in relative and absolute terms) among intense customer-contact occupations. In this part, I examine whether employment decreased as well. The model predictions suggest that the relative White labor demand increases following a competition shock. Similar to Black and Strahan, 2001, I examine whether the proportion of Blacks decreased in customer-contact occupations among the banking sector's employees. In this part, I also consider observations of females, as the impact of product market competition on male employment might not be identical across races. Specifically, I test how the banking sector' Black-White employment gap evolved (following deregulation) differentially by estimating the following equation for all banking sector's employees:

$$\begin{aligned}\text{Black}_{ist} = \alpha_0 + \beta \text{M\&A}_{ist}\{\text{CI}_{ist} > 0\} + \delta_t X_{ist} + \tilde{\delta}_t X_{ist}\text{Female}_{ist} \\ + \nu_{st} + \tilde{\nu}_{st}\text{Female}_{ist} + \Psi_t(\text{Female}_{ist}, \text{Customer}_{ist}) + \varepsilon_{ist}. \qquad (2.21)\end{aligned}$$

$\{\text{CI}_{ist} > 0\}$ is an indicator as to whether the individual is employed within the banking sector in an occupation that requires above-average customer contact. As before, X_{ist} includes a set of covariates; six age and education dummies. In this model, each interacted with year and gender. The model further includes a gender-state-year fixed-effects. I use a saturated control function to take into account non-linear changes in the banking sector's occupational composition and the gender composition of the banking sector by customer contact over time.

The results are presented in table 2.5. Consistent with the model, the results in column 1 indicate the banking sector's Black-White employment gap—relative to below-average customer contact occupations—increased by 3.2 p.p among occupations that require above-average customer contact. Similar to the wage gap, the results in column 2 indicate that the banking sector's Black-White employment gap increased, relative to below-average customer contact occupations, by 5.8 p.p among occupations that require above-average customer contact in high prejudice states, while only 2.2 p.p in the rest of the states.

In columns 3 and 4, I consider the two main occupations among male employees: financial managers and financial specialists. Together, they constitute about 40 percent of the sample's male banking sector employees. Both occupations require above-average customer contact, yet some occupations require more. Nevertheless, I focus on these occupations since (i) customers might care about interacting with financial managers and specialists more than other, higher customer contact occupations, such as bank tellers or customer service representatives, and (ii) the customers interacting with such employees might be disproportionally wealthy and White. The results in both columns suggest that interaction with financial managers and specialists is more important than other occupations that require above-average customer contact. The results in column 3 indicate the banking sector's Black-White employment gap is 4 p.p higher among financial managers and specialists, while only 2.6 p.p higher regarding the rest of the occupations that require above-average customer contact. These differences are further amplified when considering high-prejudice states. Among the high prejudice states, the banking sector's Black-White employment gap is 7.5 p.p higher among financial managers and specialists, while only 2.8 p.p higher regarding the rest of the occupations that require above-average customer contact.

The result in this part demonstrates that the competition shock increased the relative labor demand for White employees, especially in occupations with meaningful contact with customers and in states with high discriminatory tastes towards Blacks.

2.7 Conclusion

In his seminal work, Becker, 1957 shows how non-cognitive behavioral frictions can lead to racial wage gaps in the labor market. Product market competition disciplines prejudiced firms to adjust or exit the market. Consequently, much attention is focused on identifying and quantifying the effects of (i) discriminatory attitudes against Blacks on their labor market outcomes and (ii) product market competition as a disciplinary anti-discrimination policy tool, ignoring whether workers are subject to employer or customer taste-based discrimination. However, whether minorities are subject to employer or customer discrimination matters, conceptually and empirically.

In this paper, I show conceptually and find empirically that *intensified* competition in the product market *facilitates* customer taste-based discrimination in the labor market. Theoretically, I incorporate customer heterogeneity in taste for discrimination into a two-sector Dixit-Stiglitz-variant general equilibrium model, in each both product and labor market outcomes are determined. I find that a competitive product market does not tax customers' demand for discrimination as much as less competitive markets reflected in higher demand for White workers and higher wage gaps among equally productive workers.

Product market on discrimination in the labor market by integrating this demand-side property into a Dixit-Stiglitz two-sector general equilibrium model, as in Spence (1977) and Dixit (1979). Sellers produce and offer bundled products of goods and services. They hire Black and White workers who interact with the buyers.

Empirically, using U.S. census/ACS data from 1960 to 2010, data on bank deregulation (black2001division), and data on occupations' required customer contact (deming2017growing), I find that competition in the banking sector *increased* the black-white wage gap in client intensive occupations, especially in states with high discriminatory tastes (charles2008prejudice). The overall results suggest that customer taste-based discrimination is of significant importance that might over-weigh employers' prejudicial preferences, especially as the service sector is rapidly increasing (Autor and Dorn, 2013).

2.8 Figures

Figure 2.1: Characterization of banking sector's occupation by customer contact

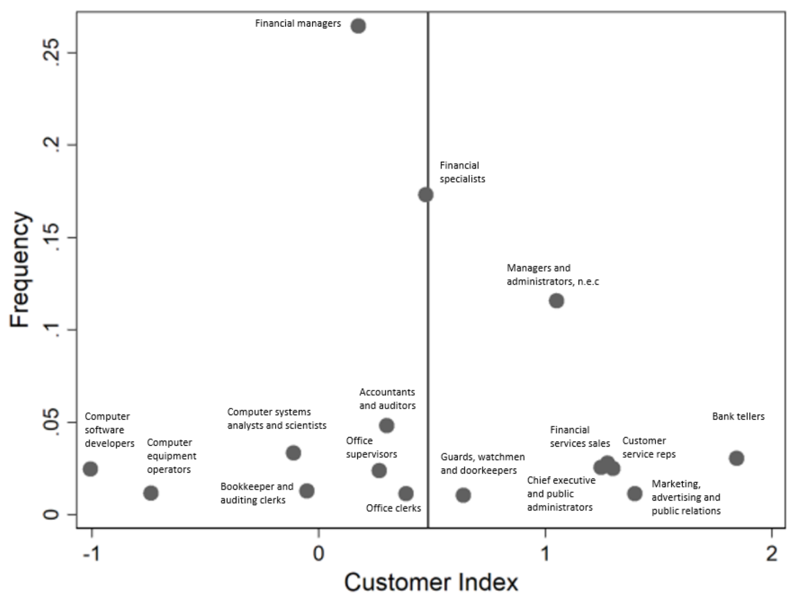

Notes: The figure describes the distribution of customer index for the banking sector's occupations. The horizontal axis denotes the occupation's consumer contact z-score (henceforth, customer index); a higher customer index implies more required contact with customers. The vertical axis measures the share of banking employees working in each occupation. The Black line denotes the banking sector's mean customer index (0.482). In this figure, I present occupations with a share of at least one percent of the banking industry's labor force among the sample.

Figure 2.2: Classification of states by prejudice index

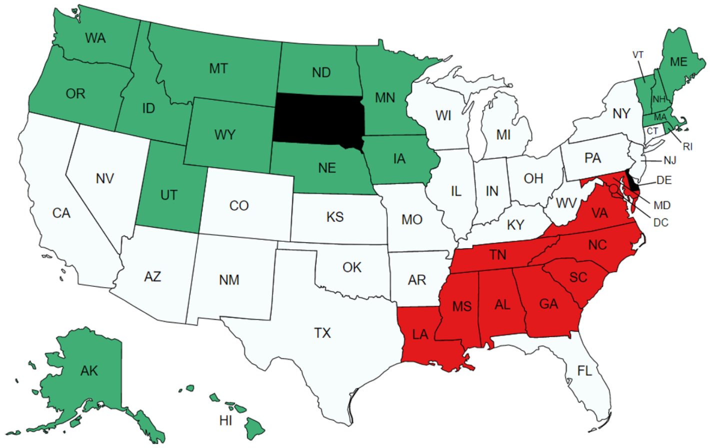

Notes: Figure 2.2 presents the classification of state by the three groups, based on the value of discriminatory attitudes towards Black of the "marginal" White individual, as in Charles and Guryan, 2008. The states with low prejudicial attitudes toward blacks are marked in green, while the states in red are the high prejudiced tastes.

Figure 2.3: Change in Black-White wage gap by customer index groups

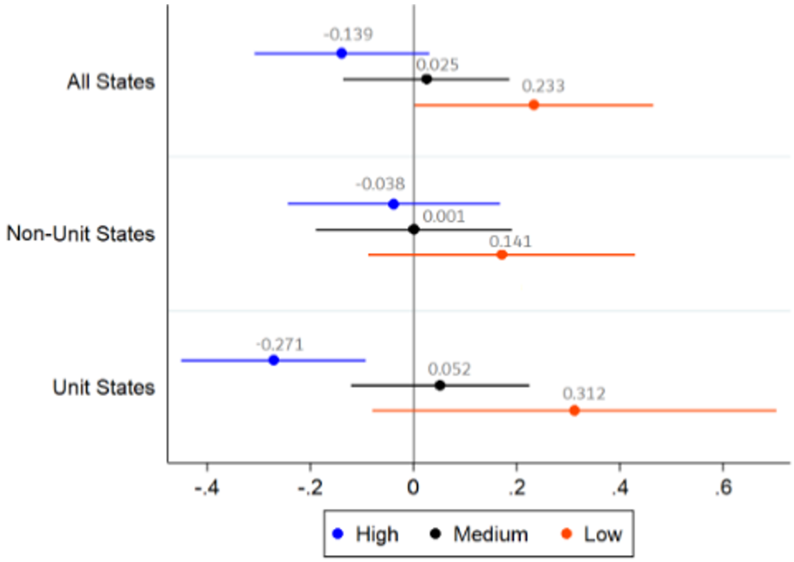

Notes: This figure depicts the estimated change in the banking sector's Black-White wage gap following deregulation by customer contact. Equation 2.19 is estimated for different subsets of the sample: (i) low customer index (i.e., < -0.355), (ii) medium customer index (i.e., $[-0.355, 0.908)$, and (iii) high customer index (i.e., > 0.908). The horizontal axis represents the value of the coefficient of interest; β of equation 2.19. The top panel ("All states") depicts β when considering all states, the center panel ("Non-Unit States") depicts β when considering only non-Unit states, and the bottom panel ("Unit States") depicts β when considering only Unit states. 95% confidence intervals are measured using robust standard errors clustered at the state level.

Figure 2.4: Change in Black-White wage gap by customer index

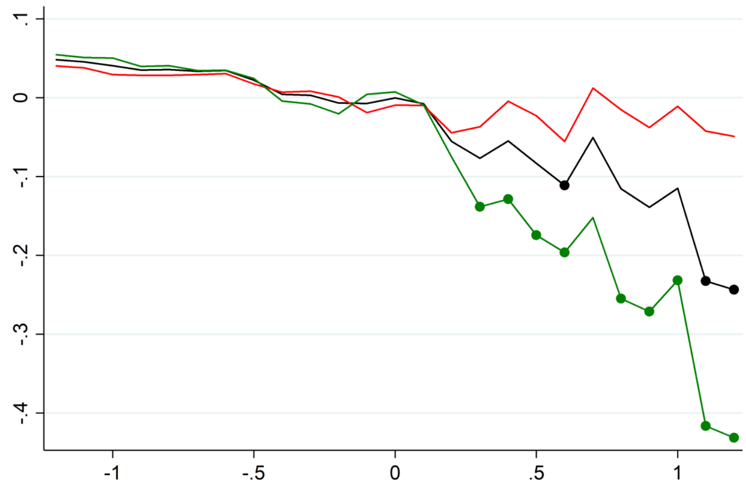

Notes: This figure depicts the estimated change in the banking sector's Black-White wage gap following deregulation by customer contact. Equation 2.19 is estimated for different sample subsets. The different sub-samples include observation with customer index larger than some lower bounds. The horizontal axis represents the different lower bounds, -1.2 to 1.2, in steps of 0.1. The vertical axis represents the value of the coefficient of interest; β of equation 2.19. The Black curve represents the estimated coefficient values, while the Black dots identify statistically significant coefficients at the 90% level. The green and red curves depict the estimation results when considering solely unit (the green curve) and non-unit states (the red curve), respectively. The green and red dots identify statistically significant coefficients at the 90% level. Confidence intervals are measured using robust standard errors clustered at the state level.

2.9 Tables

Table 2.1: Timing of intra-state and inter-state deregulation, by state

State	M&A	Interstate	State	M&A	Interstate
Alabama	1981	1987	Alaska	<1960	1982
Arizona	<1960	1986	Arkansas	1994	1989
California	<1960	1987	Colorado	1991	1988
Connecticut	1980	1983	Delaware	<1960	1988
District of Columbia	<1960	1985	Florida	1988	1985
Georgia	1983	1985	Hawaii	1986	1995
Idaho	<1960	1985	Illinois	1988	1986
Indiana	1989	1986	Iowa	1999	1991
Kansas	1987	1992	Kentucky	1990	1984
Louisiana	1988	1987	Maine	1975	1978
Maryland	<1960	1985	Massachusetts	1984	1983
Michigan	1987	1986	Minnesota	1993	1986
Mississippi	1986	1988	Missouri	1990	1986
Montana	1990	1993	Nebraska	1985	1990
Nevada	<1960	1985	New Hampshire	1987	1987
New Jersey	1977	1986	New Mexico	1991	1989
New York	1976	1982	North Carolina	<1960	1985
North Dakota	1987	1991	Ohio	1979	1985
Oklahoma	1988	1987	Oregon	1985	1986
Pennsylvania	1982	1986	Rhode Island	<1960	1984
South Carolina	<1960	1986	South Dakota	<1960	1988
Tennessee	1985	1985	Texas	1988	1987
Utah	1981	1984	Vermont	1970	1988
Virginia	1978	1985	Washington	1985	1987
West Virginia	1987	1988	Wisconsin	1990	1987
			Wyoming	1988	1987

Notes: The table reports for each state the year in which intra- and inter-state deregulation occurred. Sources: Amel, 1993, Kroszner and Strahan, 1999, and Demyanyk, Ostergaard, and Sørensen, 2007. The M&A column denotes the year at which each state allowed intra-state branching via mergers and acquisitions only. Interstate denotes the year at which unrestricted inter-state branching was permitted.

Table 2.2: The log of total earnings for banking employees and deregulation

	(1)	(2)	(3)	(4)
(A) BANK × M&A	-.041**	-.022	-.040**	-.020
	(.017)	(.020)	(.017)	(.020)
(B) BANK × M&A × UNIT		-.041		.041
		(.028)		(.029)
(C) BANK × INTER			-.020	-.027
			(.027)	(.026)
Total effects:				
(D) A + B		-.062**		-.062**
		(.026)		(.026)
(E) A + C			-.060*	-.047
			(.031)	(.031)
(E) A + B + C				-.088**
				(.037)
Observations	4,970,609	4,970,609	4,970,609	4,970,609
R-square	0.258	0.258	0.258	0.258

Notes: The dependent variable equals the log earnings. BANK is a dummy variable that equals one if the individual is employed in the banking sector. M&A is a dummy variable that equals one if the state permits intra-state branching via M&A. INTER is a dummy variable that equals one if the state permits unrestricted inter-state branching. The UNIT variable would equal one if the state began the sample with a complete prohibition on branching. The estimated models include six age and education dummies, each interacting with year and race, a race-state-year, bank-state, and bank-year fixed-effects. Columns 2 and 4 include the UNIT variable as well. Robust cluster standard errors are clustered at the state level and displayed in parentheses. *** indicates significance at the 99% level. ** indicates significance at the 95% level. * indicates significance at the 90% level.

Table 2.3: The banking sector's Black-White wage gap, deregulation, and customers indicators

	(1)	(2)	(3)	(4)	(5)	(6)
Black × BANK × CI interaction:						
(A) × M&A	-.111**	-.069	-.097	-.066	-.096*	-.064
	(.049)	(.069)	(.060)	(.076)	(.053)	(.072)
(B) × INTER		-.048		-.037		-.038
		(.061)		(.062)		(.063)
(C)× M&A × UNIT			-.063	-.060		
			(.099)	(.099)		
Total effect:						
(D) A + B		-.117**		-.102		-.102*
		(.050)		(.061)		(.055)
(E) A + C			-.160**	-.126		
			(.076)	(.010)		
(F) A + B + C				-.162**		
				(.076)		
Sample	All	All	All	All	Bank Only	Bank Only
Observations	4,970,608	4,970,608	4,970,608	4,970,608	55,263	55,263
R-square	0.264	0.264	0.264	0.264	0.366	0.366

Notes: The dependent variable equals the log earnings. Black is a dummy variable that equals one if the individual is Black. BANK is a dummy variable that equals one if the individual works in the banking sector. M&A is a dummy variable that equals one if the state permits intra-state branching via M&A. INTER is a dummy variable that equals one if the state permits unrestricted inter-state branching. The UNIT interaction equals one if the state began the sample with a complete prohibition on branching. CI denotes customer index value. The estimated models include six age and education dummies, each interacting with year and race, a race-state-year, race-bank-state, and race-bank-year fixed-effects. Robust cluster standard errors are clustered at the state level and displayed in parentheses. *** indicates significance at the 99% level. ** indicates significance at the 95% level. * indicates significance at the 90% level.

Table 2.4: The banking sector's Black-White wage gap, deregulation, customers indicators, and state's discriminatory attitudes

	(1)	(2)	(3)	(4)	(5)
(A) Black × BANK × M&A × CI	-.247***	-.066	.092	-.062	-.038
	(.054)	(.060)	(.083)	(.055)	(.063)
(B) Black × BANK × M&A × CI × HP				-.195**	-.255***
				(.078)	(.071)
(C) Black × BANK × INTER × CI					
(D) Black × BANK × INTER × CI × HP					
Sample	HP Only	MP Only	LP Only	All	Bank Only
Observations	974,753	3,293,879	701,976	4,970,608	55,263
R-square	0.309	0.250	0.220	0.264	0.366

Notes: The dependent variable equals the log earnings. HP, MP, and LP denote high, medium, and low prejudice states, respectively. CI denotes customer index value. The estimated models include six age and education dummies, each interacting with year and race, a race-state-year, race-bank-state, and race-bank-year fixed-effects. Robust cluster standard errors are clustered at the state level and displayed in parentheses. *** indicates significance at the 99% level. ** indicates significance at the 95% level. * indicates significance at the 90% level.

Table 2.5: Employment within banking sector and deregulation by customer contact

	(1)	(2)	(3)	(4)
(A) $\{\text{CI}_{ist} > 0\} \times$ M&A	-.032**	-.022*	-.026**	-.018
	(.012)	(.012)	(.012)	(.012)
(B) $\{\text{CI}_{ist} > 0\} \times$ M&A $\times$ HP		-.037***		-.031**
		(.013)		(.014)
(C) Fin. Manager/Specialist $\times$ M&A			-.014**	-.011*
			(.006)	(.006)
(D) Fin. Manager/Specialist $\times$ M&A $\times$ HP				-.016**
				(.008)
Total effects:				
(E) A + B		-.058***		-.028**
		(.016)		(.013)
(F) Total		-.058***	-.040***	-.075***
		(.016)	(.014)	(.017)
Observations	143,363	143,363	143,363	143,363
R-square	0.127	0.128	0.128	0.129

Notes: The dependent variable is a dummy variable that equals one if the individual is Black. The sample consists of female and male individuals employed in the banking sector. M&A is a dummy variable that equals one if the state permits intra-state branching via M&A. HP is a dummy variable that equals one if the employee's state is classified as high prejudice towards Blacks. Fin. Fin. Fin. Manager/Specialist is an indicator that equals one if the employee's occupation is either a financial manager or financial specialist. $\{\text{CI}_{ist} > 0\}$ is an indicator of whether the individual is employed within the banking sector in an occupation with above-average customer contact. The model includes a set of covariates; six age and education dummies, each interacting with year and gender. Columns 3 and 4 include a financial manager/specialist indicator, interacted with gender and year. The model includes a gender-state-year fixed-effects and a saturated control function of customer index, gender, and year. Robust cluster standard errors are clustered at the state level and displayed in parentheses. *** indicates significance at the 99% level. ** indicates significance at the 95% level. * indicates significance at the 90% level.

Chapter 3

The Lemons Gap: Demand For Insurance of Quality Uncertain Goods

3.1 Introduction

In theory, insurance coverage of durable goods is viewed of as a financial asset (see Schlesinger (2013)). In the case of a loss, the policy holder is completely compensated, once a deductible is met, with a replacement of the same value. Consequently, insuring any item is regarded as insuring a *wealth loss*. The policy is valued by the coverage terms, the premium and the deductible, and the probability of a claim. Specifically, the insured product's characteristics (e.g., house, automobile, or any monetary asset) should not affect the demand for insurance.

In reality, however, the empirical findings suggest that the demand for insurance varies according to context (Barseghyan, Prince, and Teitelbaum, 2011, Einav et al., 2012). In particular, the demand for insurance of a possession is not the same as the demand for insuring a financial asset, and clients over-insure their homes or automobiles relative to its monetary value (Sydnor, 2010, Barseghyan et al., 2013). The findings above present a puzzle for the standard insurance theory. Specifically, it has difficulties explaining (i) why customer's demand for insurance depends on the underlying asset, or (ii) why customers over-insure limited risk, in general, and (iii) why over-insure durable goods, in particular. For example, standard theory would require customers to inflate deductibles by 50 to 250 percent in order to explain automobile over-insurance (Cutler and Zeckhauser, 2004).[1]

[1]See Cutler and Zeckhauser, 2004, and Kunreuther and Pauly, 2006 for an extended review of

Recent work has attributed over-insurance to two aspects that were introduced in the prospect theory (Kahneman and Tversky, 1979, Tversky and Kahneman, 1992); Loss aversion (Kőszegi and Rabin, 2006, Kőszegi and Rabin, 2007) and over-weighting of low probability events (Lattimore, Baker, and Witte, 1992; Prelec, 1998), especially if recall of information is selected and limited (Gennaioli and Shleifer, 2010). Others attribute over-insurance to the fact that insurance and warranties frequently enable speedier replacement of the insured item (Meza and Reyniers, 2020). While these factors might explain over-insurance in general, they cannot explain why over-insurance is context-dependent; In particular why clients over-insure possessions by paying too much to face a lower deductible, relative to an income loss.

In this paper, I offer a new unifying explanation that jointly addresses the three puzzles. I do so by focusing on the difference between the item and its replacement. A key difference between financial assets and durable goods are that the good is valued via the pre-owned market. While all dollar bills are equally useful, this is not true for cars. I integrate the market of "lemons" into the analysis of the demand for insurance. Specifically, I show the importance of the "lemons gap", which is the difference between the insured and replacement values, on the willingness to pay for insurance, selection to comprehensive coverage, and moral hazard. I find that the lemons gap can explain the puzzling findings.

In his seminal paper, Akerlof (1970) identifies the role of adverse selection as a source of inefficiency in trade. Asymmetric information leads to the lemons to be disproportionally sold, which lessens the average quality of the items available in the resale market. Therefore, compensation for a loss of an insured good with a pre-owned replacement is incomplete. This gives rise to the lemons gap. The lemons gap indicates on the missing market. Customers cannot insure their vehicle; Instead, by purchasing an insurance coverage policy for their own car, they are effectively insuring the average car in the resale market.[2]

The lemons gap is embedded in the demand for insurance. Insured durable goods only exist in the resale market. For instance, it is impossible to replace a used vehicle with a brand-new old vehicle. As a result, the insured vehicle's value is based on the transaction prices of vehicles of the same observable characteristics that are sold through the pre-owned market. Yet, the market for lemons imply that full insurance is incomplete; The insured vehicle's quality is higher than the quality of vehicles sold in the resale market. In spite of its broad implications, the insurance literature ignores the lemons gap; It does not make the distinction between insuring quality

insurance anomalies.

[2]Of the same observable characteristics, such as model and year.

certain and uncertain products, but rather assumes that the replacement is identical in value.

In this paper, I close this gap. I model the demand for insurance when clients are compensated with a replacement from the resale market. I analyze the differences between insuring the loss of a quality uncertain good and a monetary loss. I show that the lemons gap affects the consumer's willingness to pay, and how its implications on the insurance market are consistent with the empirical anomalies.

Intuitively, the coverage is incomplete, as clients are partially compensated in case of a loss; The replacement vehicle is of inferior quality relative to the insured client's car. Therefore, the effective economic deductibles are higher than their monetary values, especially when the vehicle is new. This affects the demand for insurance, adverse selection, and moral hazard, both in terms of levels and age-dependent trends.

Initially, I characterize the pre-owned market equilibrium and its evolution over the vehicle life cycle via the canonical adverse selection models of Akerlof (1970) and Wilson (1980). In addition, in order to study the equilibrium *dynamics*, I incorporate *vehicle depreciation* and an *outside option* to the canonical selection models. The outside option provides owners an additional incentive to sell, independent of their vehicle quality. The reservation price drops over the vehicle life cycle as owners are interested in replacing their vehicles with a brand new car. This is similar in spirit to the dynamic model of Hendel and Lizzeri (1999). Two key findings emerge. First, as in Akerlof (1970), the lemons gap exists as the quality of vehicles not sold is higher than the resale market's average. In addition, the lemons gap decreases over the vehicle life cycle. As the vehicle ages, the incentive to sell depends less on the owner's vehicle quality and more on the outside option value. Consequently, the pre-owned market is less adversely selected over the vehicle life cycle as its quality distribution improves.

Motivated by these findings, I next model the owner's demand for insurance. I show that the lemons gap results in over-insurance. Intuitively, although the difference between two deductibles is constant, the value of both economic deductibles, that is the monetary deductible plus the uncompensated loss, are higher when the replacement is an imperfect substitute. As a result, the demand for more coverage is higher when the compensation is partial. My model shows that the lemons gap can explains why the demand for insurance is context-dependent, and why customers over-insure possessions, such as automobiles. I further examine the interaction between the lemons gap and behavioral attributes, such as over-weighting probability of claim and loss aversion. I show that two are positively related; The existence of both the lemons gap and behavioral aspects in the demand for insurance further amplify the willingness to pay for comprehensive coverage.

The lemons gap also has age-dependent implications on the demand for insurance.

As the vehicle ages, both the vehicle depreciates and the resale market's quality improves. This drives a drop of the lemons gap, and thus a decrease in the demand for insurance. The evolution of the lemons gap can explain why lower risk aversion parameters are associated with older cars, relative to new ones (Cohen and Einav, 2007).

The insurance literature devotes much attention to analyzing the implications of asymmetric information in insurance market. The canonical models of asymmetric information have assumed that selection to comprehensive coverage is adverse (Akerlof, 1970; Rothschild and Stiglitz, 1976). Yet, the empirical validity of the *positive correlation property* between demand and cost, introduced in Chiappori and Salanie, 2000, is limited. Some have documented adverse selection (Cohen, 2005; Einav, Finkelstein, and Schrimpf, 2007, Cohen and Einav, 2007), while others found no evidence of selection (Chiappori and Salanie, 2000; Cardon and Hendel, 2001), or even advantageous selection (Finkelstein and McGarry, 2006, Fang, Keane, and Silverman, 2008, Cutler, Finkelstein, and McGarry, 2008).[3] Extension of the canonical models that allows heterogeneity in both risk and risk preferences can explain the advantageous selection phenomena (De Meza and Webb, 2001). Yet, the insurance theory asserts that selection (whether adverse or advantageous) is static.

My model has dynamic implications on adverse selection. The lemons gap dynamics imply that the insurance market is more adversely selected as the vehicle ages. Intuitively, when the vehicle is new, the uncompensated loss is large. As a result, the demand for insurance is particularly high for risk averse individuals. The lemons gap drops over the vehicle life cycle. As a result, the willing to pay for more coverage decreases, especially for risk-averse owners. Therefore, as the vehicle ages, the market is more adversely selected; the demand for insurance depends more on the insured's risk and less on risk preferences.

With regards to moral hazard, the drop in the lemons gap implies that effort decreases (and claims increase) over the vehicle life cycle, as the economical deductible drops. The drop in the uncompensated loss can explain why vehicle owners tend to put less effort in maintaining their old, high-mileage "beater" cars.

These level and age-dependent trends in demand, selection, and moral hazard occur if and only if the insurance market is incomplete - in other words, if the lemons gap exists. If the insurance coverage provides complete insurance (up to a deductible), vehicle owners should not over-insure their vehicles, and the willingness to pay for a lower deductible is constant over time, which implies that both adverse selection and moral hazard are static.

[3] See Cohen and Siegelman (2010), and Chiappori and Salanié (2013) for an extended review of empirical studies of asymmetric information in the insurance market.

This paper stresses out the importance of taking into account quality uncertainty when conducting an empirical analysis of insurance markets. Echoed by theory, the empirical structural models of insurance (see Einav, Finkelstein, and Levin, 2010b) do not take into account the quality uncertainty of goods by imposing that an insurance coverage provides a complete compensation (once a deductible is met). Assuming that the coverage provides full insurance may explain why the risk aversion parameters depend on the identity of the insured product in general, and why the risk aversion parameters for durable goods are too high in particular, especially when new and valuable. Furthermore, the empirical literature assumes that both adverse selection and moral hazard are static. I show that this is not the case. The market is more adversely selected over the vehicle life cycle, and moral hazard is exacerbated.

The remainder of this paper proceeds as follows. Section 2 examines the automobile resale market. Section 3 integrates the key observations from the pre-owned market to the analysis of the demand for insurance. Section 4 concludes.

3.2 The Automobile Market

Setting

In this section, I show that the lemons gap exists and trends as the vehicle ages. I do so by analyzing the trade decisions in the market for used vehicles. I examine the quality distribution of the sold vehicles in the pre-owned market at a given age. Then, I study how it evolves over the vehicle life cycle.

Consider a model of the used vehicle market *a la* Akerlof, 1970 and Wilson, 1980. A static economy that consists of a unit mass of two types of agents: owners of a vehicle and non-owners. Vehicles, which are characterized by the observable variables model and age, differ in quality as well. Unlike vehicle model and age, which the set of different combinations are denoted by $\mathcal{T}$, quality is unknown to the potential buyer, yet known to the owner. q_{im} denotes the quality index of vehicle i of model m. Quality index q is distributed according to a continuously differentiable and strictly positive density, $f_m(q)$, defined on $[q_{0m}, q_{1m}]$ with $q_{0m} > 0$.

Owners

As in Akerlof, 1970 and Wilson, 1980, owners value vehicle quality in a similar way.[4] Vehicle's age plays a role as well. All agents enjoy a higher utility from owning

[4]The results in this paper hold if owners differ in their preferences for quality as long as the owners' distribution of preferences is independent of the vehicle quality distribution.

a newer vehicle relative to an older one of the same quality. The owner's value from vehicle i of model m at age t is defined as $V_{imt} = q_{im}\omega_t$. ω_t reflects the role of vehicle age; ω_t decreases with vehicle age as value depreciates over the vehicle life cycle.

The owners decide whether to sell their vehicle. The owner's utility from selling the vehicle is $P_{mt} + \nu_i$. P_{mt} is the price of the vehicle. ν_i represents the (additional net) utility from the outside option. The utility from the outside option is distributed according to a continuously differentiable and strictly positive density, $g(\nu)$, defined on $[\nu_0, \nu_1]$ with $\nu_0 \geq 0$. $G(\nu) = \int_{\nu_0}^{\nu} g(y)dy$ is the cumulative distribution function of ν. In addition, I assume that vehicle quality and the utility from the outside option are independently distributed.

The outside option is the main difference between this model and the classical models of markets with adverse selection. In Akerlof, 1970 and Wilson, 1980, the incentive for trade arise solely from differences in vehicle valuation. Buyers have a stronger preference for the vehicles than the sellers. ν represents an additional incentive to trade, unrelated to vehicle quality. Owners might consider selling their used vehicle in order to replace it with a brand new car of higher quality. This is similar in spirit to the incentive to trade in Hendel and Lizzeri, 1999.

Owners choose to sell their vehicle if and only if its value is lower than the utility from a sale.

$$V_{imt} = q_{im}\omega_t \leq P_{mt} + \nu_i \Rightarrow u_{imt} \equiv q_{im} - \frac{\nu_i}{\omega_t} \leq p_{mt} \equiv \frac{P_{mt}}{\omega_t}, \tag{3.1}$$

where u_{imt} is owner i's (age-adjusted net) utility from the vehicle, and p_{mt} is the age-adjusted price.

Equation 3.1 demonstrates how the incentive to sell changes over the vehicle life cycle. The decision depends on two factors: (i) quality q, and (ii) the outside option ν. The relative weight of each factor is a function of vehicle age. As the vehicle ages, ω_t drops. As a result, the owner's trade decision depends more on the outside option value and less on the vehicle's quality. Intuitively, the decision to sell a new vehicle depends mostly on its quality, while the selling decision of an old one has more to do with the replacement.

Non-Owners

Non-owners are characterized by a von Neumann-Morgenstern utility function and have the same expectation regarding vehicle quality, as in the standard models of adverse selection (Akerlof, 1970; Wilson, 1980). Non-owner's utility from owning a vehicle of quality q at age t is defined as $\hat{V}_t(q) = \gamma q \omega_t$, where $\gamma > 1$. $\gamma > 1$ reflects the standard gain from trade; Non-owners who are interested in purchasing a vehicle value it more than the owners selling one.

Alike Akerlof, 1970, and unlike Wilson, 1980, I restrict my attention to homogeneous quality valuation. This implies that there exists a sufficient supply of non-owners such that the demand for a particular model does not drop with volume of trade. As I show in the next part, this implies that the vehicle's equilibrium price perfectly depends on its offered quality (adjusted by age) in the resale market. For instance, two different vehicle models of the same age and offered quality (in the pre-owned market) are priced identically. Furthermore, the difference in prices between a high and low quality vehicles drop with vehicle age.[5]

Non-owners prefer to purchase a vehicle that generates them the highest net expected utility.

$$[\hat{V}_{mt}(q) - P_{mt}|P_{mt}] = \gamma\mu_{mt}\omega_t - P_{mt} = \omega_t(\hat{u}_{mt} - p_{mt}), \tag{3.2}$$

where $\hat{u}_{mt} = \gamma\mu_{mt}$ denotes the non-owner's (age-adjusted) expected utility, as μ_t is the non-owners' expected quality of the cars offered for sale, $\mu_t = [q|u_{imt} \geq p_{mt}]$.

Equilibrium

As mentioned earlier, I assume that buyers do not observe the vehicle's quality before purchase. Therefore, the selling prices of all vehicles of the same model and age are identical, as buyers cannot discriminate between vehicles of the same type. Equilibrium price is set to equate between demand and supply. As in asymmetric information models, a price increase affects demand via two channels. First, the increase in price reduces the number of non-owners willing to purchase a vehicle (and increases the number of owners wishing to sell one). In addition, the price increase improves the quality distribution of vehicle offered in the resale market, and thus increases demand. As a result, a price increase unambiguously shifts the supply upwards, yet it is unclear how the demand is affected.

By equation 3.1, owners decide to sell their vehicle if and only if their utility from the vehicle is higher than the price. This implies that the aggregated supply for a given price p, denoted by $S_{mt}(p)$, is equal to:

$$S_{mt}(p) = \Pr(u_{imt} \leq p) = \int_{q_{0m}}^{q_{m1}} \Pr\left(q - \frac{\nu_i}{\omega_t} \leq p\right) f(q)dq =_q [1 - G(\omega_t(q-p))],$$

where the last equality is due to the fact that νq. Thus, supply is continuous and (strictly) increasing in price p (for $S_{mt}(p) \in (0,1)$).

[5] Akerlof, 1970 and Wilson, 1980 analyze the equilibrium price and volume of trade of a single type of vehicle. Th homogeneity assumption allows me to characterize the entire pre-owned market, which consists of multiple (observed) types of vehicles.

Although they are unaware of the specific vehicle quality, non-owners understand the implications of a change in price on the quality distribution of vehicles offered in the market. Specifically, the potential buyers' expectation regarding the quality of the vehicles sold in the pre-owned market (for price p_{mt}) is correct. As with supply, the expected quality of the cars offered for sale, $\mu_{mt}(p)$, is a function of price as well. I assume that $\mu_{mt}(p) = q_{0m}$ when no cars are offered for sale, $S_{mt}(p) = 0$, and that the expected quality equals the average quality in the resale market in the case of positive trade (as in Akerlof, 1970; Wilson, 1980). Since νq,

$$\mu_{mt}(p) = \begin{cases} \frac{1}{S_{mt}(p)}{}_{q_m}[q_m(1 - G(\omega_t(q_m - p))] & \text{if } S_{mt}(p) > 0; \\ q_{0m} & \text{if } S_{mt}(p) = 0. \end{cases}$$

Similar to supply, the equation above indicates that the expected quality is continuous in price p.

In this paper, I perform comparative statics to study the importance of vehicle age on the quality distribution offered in the pre-owned market. In order to analyze the effect of vehicle depreciation, I invert the equilibrium condition.[6] Let $R_{mt}(Q)$ denote the reservation price of the Q^{th} owner (ordered from lowest to highest valuation) of vehicle model m at age t. $R_{mt}(Q)$ is defined as follows:

$$R_{mt}(Q) = \begin{cases} q_{0m} - \frac{\nu_1}{\omega_t} & \text{if } Q = 0; \\ S_{mt}^{-1}(Q) & \text{if } Q \in (0, 1); \\ q_{1m} - \frac{\nu_0}{\omega_t} & \text{if } Q = 1. \end{cases}$$

The reservation price function $R_{mt}(Q)$ is continuous and strictly increasing in Q.

With regards to demand, $\mathcal{W}_{mt}(Q)$ denotes the willingness to pay for vehicle model m of age t when Q vehicles are offered for sale. In this paper I focus on the non-degenerate case in which equilibrium is characterized with partial trade. Full trade does not occur for any vehicle model m at any age t if the non-owners' willingness to pay for the (population) mean vehicle is lower than the maximum reservation price

[6] This is as done in Handel, Kolstad, and Spinnewijn, 2019b. Handel, Kolstad, and Spinnewijn, 2019b examines the welfare implications of an information policy, which reduces frictions between consumers' willingness to pay and their true utility, on equilibrium insurance markets. A reduction in frictions affects equilibrium coverage through both a level effect, and a sorting effect. They invert the equilibrium condition to disentangle between the two; a level shift in demand (which does not affect the insurer's average cost function for a given quantity of coverage policies sold) and changes in consumer sorting (which does affect the insurer's average cost function for a given quantity of coverage policies sold). The logic in this paper is the same. As vehicle ages, all owners have a lower reservation price. Yet, some owners decrease in reservation price is larger than others, depending on the owner's vehicle quality.

among owners, characterized by the highest quality and the lowest outside option value:

$$\gamma[q_m] < q_{1m} - \frac{\nu_0}{\omega_t}, \quad \forall(m,t). \tag{3.3}$$

Since at equilibrium not all vehicles are sold, some of the non-owners do not purchase a car in the resale market. Therefore, in equilibrium, owners are indifferent between purchasing a vehicle, and not. By equation 3.2, equilibrium prices, $P^e_{mt} = p^e_{mt}\omega_t$, are set to equal the non-owners' willingness to pay (as in Akerlof, 1970), $\mathcal{W}_{mt}(Q)$:

$$\mathcal{W}_{mt}(Q) = \gamma\tilde{\mu}_{mt}(Q),$$

where $\tilde{\mu}_{mt}(Q) = \mu_{mt}(R_{mt}(Q))$ denotes the expected quality of the cars (of model m and age t) in the pre-owned market when Q vehicles are offered for sale. Similar to before, $\tilde{\mu}_{mt}(Q)$ is continuous in Q. Therefore, as with the reservation price function, the willingness to pay function $\mathcal{W}_{mt}(Q)$ is continuous in Q.

As mentioned earlier, the assumption of an homogeneous γ implies that equilibrium prices are set such that $P^e_{mt} = \gamma\mu_{mt}(Q^e_{mt}), \forall(m,t) \in \mathcal{T}$. Thus, $P^e_{mt} \geq P^e_{m't'}$ if and only if $\tilde{\mu}_{mt}\omega_t \geq \tilde{\mu}_{m't'}\omega_{t'}$. Furthermore, the difference in prices between a high and low quality vehicles (of those offered for sale) of the same age drops over the vehicle life cycle. Specifically, consider two different vehicle models of the same age, $\tilde{\mu}_{lt}$ and $\tilde{\mu}_{ht} > \tilde{\mu}_{lt}$. Conditional on the offered quality, i.e., $\tilde{\mu}_{mt} = \tilde{\mu}_m$, $P^e_{ht} - P^e_{lt} = \gamma(\tilde{\mu}_h - \tilde{\mu}_l)\omega_t$ is positive and drops with age (as ω_t drops).

Using the definitions above, equilibrium condition implies that that the volume of trade, denoted by Q^e_{mt}, is set to equate between the owner's reservation price and the non-owners' willingness to pay.

$$R_{mt}(Q^e_{mt}) - \mathcal{W}_{mt}(Q^e_{mt}) = 0, \quad \forall(m,t) \in \mathcal{T} \tag{3.4}$$

To show an equilibrium exists, first consider the case of a single type (model-age) of vehicle that exists in the pre-owned market. By equation 3.3, $R_{mt}(1) - \mathcal{W}_{mt}(1) = q_{1m} - \frac{\nu_0}{\omega_t} - \gamma[q_m] > 0$ as not all vehicles are sold in the resale market. In addition, $q_{0m} > 0$ implies that the resale market is characterized with a positive volume of trade; $R_{mt}(0) - \mathcal{W}_{mt}(0) < 0$ as $R_{mt}(0) = q_{0m} - \frac{\nu_1}{\omega_t} < \gamma q_{0m} = \mathcal{W}_{mt}(0)$. By continuity of $R_{mt}(Q)$ and $\mathcal{W}_{mt}(Q)$, there exists an equilibrium quantity of trade $Q^e_{mt} \in (0,1)$ such that the reservation price equals the willingness to pay. Since the equilibrium condition in equation 3.4 implies that non-owners are indifferent between the purchase of any vehicle (as $p^e_{mt} = \mathcal{W}_{mt}(Q^e_{mt}), \forall(m,t) \in \mathcal{T}$), equilibrium exists for the case of multiple types of vehicles (grouped by model and age) as well.

Equation 3.4 defines not only the equilibrium aggregate level of trade but also the extent of trade in equilibrium by quality. In equilibrium, there is partial trade of vehicles. Since (i) buyers have a higher value for vehicle quality relative to the sellers and (ii) buyers' willingness to pay is based on the average vehicle quality offered, in equilibrium, all owners with vehicle quality below the equilibrium average quality, set by the equilibrium (age-adjusted) price p^e_{mt}, prefer to sell their vehicle. Proposition 4 proofs this property.

Proposition 4. *In equilibrium, all owners with vehicle quality lower than the market's expected quality, $q_{im} \leq \mu_{mt}(p^e_{mt})$, sell their vehicle.*

Proof: *See appendix.*

The implications of proposition is illustrated in Figure 3.1. Vehicle quality q is measured on the horizontal axis and the probability of sale on the vertical axis. The curve depicted represents an arbitrary equilibrium trade allocation of vehicles by quality. By proposition 4, all owners with vehicle quality below the equilibrium average quality $\mu_{mt}(p^e_{mt})$ sell their vehicles. Thus, in the non-degenerate case of no full trade the used automobile market is adversely selected. This is in spirit, the Akerlof, 1970 argument, generalized to the case of an additional incentive of trade, independent of vehicle quality. The vehicle quality of owners who decide not to sell their car is higher than the average quality offered in the pre-owned market. This implies that the lemons gap exists; compensating an owner who lost her vehicle via the pre-owned market provides incomplete insurance as the quality of the vehicles sold in the pre-owned market is a deficient alternative. I now examine how the lemons gap evolve over the vehicle life cycle.

Comparative Statics

Equilibrium condition (stated in equation 3.4) allows us to examine how equilibrium allocation change as the vehicle ages, i.e. equilibrium price of a particular vehicle model at age $t+1$ with $\omega_{t+1} < \omega_t$, relative to age t. Equation 3.1 indicates that the decision to sale depends less on quality q and more on the outside option ν as vehicle value depreciates. Therefore, the distribution of vehicles offered for sale by owners is more favorable over the vehicle life cycle. Proposition 5 proofs this statement.

Proposition 5. *Let $F_{mt}(\cdot|Q)$ denote the vehicle model m's quality distribution of the Q vehicles that are supplied at the pre-owned market at age t. The ordering of supply improves over the vehicle life cycle such that, conditional on Q, the quality distribution of supplied vehicles at age $t+1$ first order stochastically dominates that of age t, $F_{mt+1}(\cdot|Q) \leq F_{mt}(\cdot|Q), Q \in [0,1]$ (and strict inequality for some values when $Q \in (0,1)$).*

Proof: *See appendix.*

The implication of proposition 5 is depicted in Figure 3.2. Vehicle quality q is measured on the horizontal axis and the outside option value ν is on the vertical axis. The area within the dashed-line triangle ABC denotes an arbitrary $Q \in (0,1)$ owners with the lowest reservation price at age t, while the area within the solid-line triangle ADE denotes the $Q \in (0,1)$ owners with the lowest reservation price at age $t+1$.

At $t+1$ the market is less adversely selected. For any given $Q \in (0,1)$ offered for sale, the distribution of quality is more favorable at age $t+1$ relative to age t. As the vehicle ages the incentive to sell depends more on the outside option and less on the vehicle's quality. As a result, the sorting of the supply changes; Owners with high quality vehicle - high outside option value have a lower reservation price than some owners with a low quality vehicle - low outside option value, which previously had a lower reservation price.

The evolution of demand, supply and equilibrium outcomes is depicted in figure 3.3. The horizontal axis measure the volume of trade, in terms of the share of pre-owned vehicles. The vertical axis measures the equilibrium age-adjusted price, p^e_{mt}. The dashed curves reflects the owners' reservation price at age t (in gray), and at age $t+1$ (in black), while the solid curves measure the non-owners' willingness to pay at age t (in gray), and at age $t+1$ (in black).

By proposition 5, since the market is less adversely selected, demand increases; For any given level of partial trade $Q \in (0,1)$, the willingness to pay is higher in period $t+1$ relative to t, as $\mathcal{W}_{mt+1}(Q) = \gamma\tilde{\mu}_{mt+1}(Q) > \gamma\tilde{\mu}_{mt}(Q) = \mathcal{W}_t(Q)$. In addition, supply increases as well, as the utility from owning a vehicle drops $u_{imt+1} = q_{im} - \frac{\nu_i}{\omega_{t+1}} < q_{im} - \frac{\nu_i}{\omega_t} = u_{imt}$ (equation 3.1). Therefore, the reservation price drops $R_{mt+1}(Q) < R_{mt}(Q), \forall Q \in [0,1)$. As a result, equilibrium (age-adjusted) price and volume of trade rise over the vehicle life cycle.

Corollary 6. *For any given vehicle model m, equilibrium (age-adjusted) price and quantities sold increase over the vehicle life cycle, $p^e_{mt+1} > p^e_{mt}$, and $Q^e_{mt+1} > Q^e_{mt}$.*

3.3 The Insurance Market

Setting

In this section, I incorporate key insights from the automobile market into the analysis of the demand for insurance. The analysis of the used-vehicle market provides two key observations. First, the resale market is adversely selected (proposition 3.1) as vehicles not sold by the owners are of higher quality, relative to those sold in the pre-owned market. In addition, the used-vehicle market is less adversely selected over the vehicle life cycle. As a result, the age-adjusted price increases with age (corollary 6).

Consider an owner of a vehicle facing a decision on which insurance coverage plan to purchase. I focus on a benchmark case where clients select one of two insurance coverage plans. Either a comprehensive coverage plan, plan 1, which is characterized by premium ρ_1 and deductible d_1, or a partial coverage plan, plan 0, which is characterized by a lower premium $\rho_0 < \rho_1$ and a higher deductible $d_0 > d_1$.[7] The contract is over a short time period τ, where premiums are linearly prorated, as standard in the industry (Cohen and Einav, 2007). There are two possible states. Either no loss, or a total-loss event occurs. In the case of a total-loss event, the insuree is compensated by the market value of the vehicle, which is based on its observable characteristics (model and age), and pays the plan's deductible to the insurer. Individual i's utility from insurance coverage c for a vehicle of model m and age t, $\mathcal{U}^c_{imt}$, is defined as follows:

$$\mathcal{U}^c_{imt} = (1-\pi\tau)(U(W-\rho_c\tau)+\phi(q_{imt}\omega_t))+\pi\tau(U(W-\rho_c\tau-d_c-P_{m^r t^r}+P_{mt})+\phi(\mu_{m^r t^r}\omega_{t^r})),$$

where π is the Poisson risk rate of a total-loss event. U is strictly concave function that reflects the insuree's utility from consumption, while ϕ is a strictly increasing function that measures the client's induced utility from vehicle quality (adjusted by age). W denotes the individual's wealth. ρ_c and d_c denote the insurance coverage plan c's premium and deductible, respectively. P_{mt} denotes the market value of the insuree's vehicle. The (m^r, t^r) indexes denote the model and age of the replacement vehicle purchased by the insuree in the case of a loss. Thus, the expected quality of the replacement vehicle is $\mu_{m^r t^r}\omega_{t^r}$, while $P_{m^r t^r}$ denotes its price.

As evident from the formulation of $\mathcal{U}^c_{imt}$, the identity of the replacement vehicle plays a role. Suppose that the insuree's replacement vehicle (in the case of a loss) is of the same (expected) quality as her vehicle, a standard assumption in the insurance literature.[8] Since $p_{mt} = \mu_{mt}\omega_t, \forall(m,t)$, $\mathcal{U}^c_{imt}$ can be expressed as follows:

$$\begin{aligned}\mathcal{U}^c_{imt} =&(1-\pi\tau)U(W-\rho_c\tau)+\pi\tau U(W-\rho_c\tau-d_c-\gamma(q_{imt}-\mu_{mt})\omega_t)+\phi(q_{imt}\omega_t)\\ =&(1-\pi\tau)U(W-\rho_c\tau)+\pi\tau U(W-\rho_c\tau-d_c-L_{imt})+\phi(q_{imt}\omega_t). \qquad (3.5)\end{aligned}$$

[7] I focus on the non-degenerate case where vehicle value is higher than both deductibles.

[8] The benchmark modeling design of the demand for insurance (see Schlesinger, 2013) is that in the case of a claim event, the insuree's cost is equal to a deductible payment. Thus, the income effect generated by the deductible payment, or any other friction which drive a difference between the insured item and its replacement are not taken into account. In this paper, I do not deviate from previous studies in this dimension, but rather focus solely on the difference in compensation with regards to a loss of a quality uncertain good, relative to a good of observed quality (or income loss).

Equation 3.5 re-expresses $\mathcal{U}^c_{imt}$ to incorporate two key insights from the automobile market analysis in the analysis of the demand for insurance. $L_{imt} = \gamma(q_{imt} - \mu_{mt})\omega_t$ is the lemons gap; The uncompensated loss in the case of a total-loss event. By proposition 3.1, the lemons gap exists. The replacement vehicle serves as a partial substitute as the owners' vehicle quality is better than the average quality at the pre-owned market. In order for the insuree to purchase a vehicle of the same quality, she must forgo some units of consumption as the compensation is partial, $L_{imt} > 0$. In addition, the lemons gap decreases over the vehicle life cycle, $L_{imt+1} < L_{imt}$. This stems from the fact that (i) by corollary 6, the pre-owned market (age-adjusted) price of the insuree's car increases over the vehicle life cycle, and (ii) depreciation ($\omega_{t+1} < \omega_t$) implies that as the vehicle ages, the uncompensated loss from the inadequate substitute is of less importance, as the difference in prices between high and low quality vehicle drops with age.

Demand

In this part, I examine how the demand for insurance of a quality uncertain good differs than that of a good of known quality (or a monetary loss). This is done by examining the net utility from the comprehensive insurance coverage plan, relative to the partial coverage, for both the case of $L_{imt} > 0$ and $L_{imt} = 0$. The insuree buys the comprehensive coverage plan if and only if $\mathcal{U}^1_{imt} \geq \mathcal{U}^0_{imt}$. By equation 3.5, this condition can be re-expressed as:

$$U(W-\rho_0\tau)-U(W-\rho_1\tau) \leq \frac{\pi\tau}{1-\pi\tau}\Big(U(W-\rho_1\tau-d_1-L_{imt})-U(W-\rho_0\tau-d_0-L_{imt})\Big) \tag{3.6}$$

Equation 3.6 highlights differences between the demand for insurance of a perfectly replaceable good ($L_{imt} = 0$) and that of an quality uncertain item ($L_{imt} > 0$). The demand for insurance of a quality uncertain good differs than that of an observed income loss both in terms of level and life cycle trends. Yet, insurance theory as solely focused on the case of $L_{imt} = 0$ (see Schlesinger, 2013).

The demand for a comprehensive insurance plan is higher when the insured item's quality is uncertain. When only partially compensated in a case of a loss, $L_{imt} > 0$, the clients must forgo units of consumption (in addition to those necessary to pay the deductible) in order to replace their vehicle (in the case of a total-loss event). Due to the strict concavity of U, the benefit from lowering the deductible is larger, relative to the case of complete insurance, $L_{imt} = 0$. When insuring an income loss (or when the compensation is sufficient to replace the lost item with a perfect substitute), the economic and monetary values (from both coverages) converge. In contrast, the lemons gap imply that an uncompensated loss exists. This gives rise to the over-insurance phenomena. When considering solely the monetary characteristics of both coverage plans, high risk aversion parameters are necessary in order to explain the high demand for full insurance.

The implications of the lemons gap are not limited to a static level shift in demand. The lemons gap generates life cycle patterns as well. Insurance theory predicts that the demand for insurance is static, as $L_{imt} = 0$. Although the vehicle value drops with age, the monetary terms of the two policies do not change over the vehicle life cycle (when premiums and deductibles are fixed). Therefore, the insuree's willingness to pay does not change over time, as $\mathcal{U}^1_{imt} - \mathcal{U}^0_{imt} = \mathcal{U}^1_{im} - \mathcal{U}^0_{im}, \forall (m,t) \in \mathcal{T}$. This stems from the fact that the sole incentive to purchase a comprehensive insurance plan is the right to pay a lower deductible (by $d_0 - d_1$), in the case of a loss. This is the incentive to take-up comprehensive insurance both in the standard neoclassical models and in the alternative models, mentioned earlier. Behavioral aspects and precautionary savings motives do generate over-insurance. Yet, the demand for insurance does not change over the vehicle life cycle.

However, when the insurance market is incomplete, the uncompensated loss affects the demand for insurance, both in levels and life cycle trends. By equation 3.6, the demand for comprehensive insurance increases with the uncompensated loss L_{imt}. The uncompensated loss does not only up-weight the willingness to pay for insurance, but also generates dynamic trends in the demand for insurance. The lemons gap has a differential impact on the willingness to pay. Corollary 6 indicates that L_{imt} drops over the vehicle life cycle. Since U is strictly concave, the right hand side of equation 3.6 drops with age. Thus, the insuree's willingness to pay for a comprehensive insurance coverage plan drops over the vehicle life cycle as well.

Corollary 7. *If the insured good is perfectly replaceable ($L_{imt} = 0$), the demand for insurance is static and identical to an income loss. Yet, if the insuree is only partially compensated via the pre-owned market value of the good ($L_{imt} > L_{imt+1} > 0$), the demand for insurance is higher thanthat of an income loss, and drops over the vehicle life cycle.*

Behavioral Demand

So far, I considered the demand for insurance of a perfectly rational client. Recent work has highlighted how behavioral aspects, such as loss aversion and over-weighting of low probability events, can resolve the divergence of observed customer behavior from that of a rational frictionless agent. Specifically, these behavioral aspects can explain the existence of the over-insurance phenomena, as they inflate the demand for insurance, relative to its monetary compensation. In this part, I examine the interaction between loss aversion and probability distortion and the lemons gap. I show that the behavioral attributes and the lemons gap positively interact; The increase in the demand for insurance generated by both (i) the lemons gap, and (ii) either the up-weighting of loss probability or due to the individual's will to consume more than the reference level of consumption is larger than the sum of the increase in willingness to pay of both attributes, when occur separately.

In this part, I incorporate two behavioral aspects that were introduced in the prospect theory (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992): Probability distortion (Lattimore, Baker, and Witte, 1992; Prelec, 1998), and loss aversion (Kőszegi and Rabin, 2006, Kőszegi and Rabin, 2007). I embed probability distortion by considering the client's subjective beliefs on the likelihood of a claim, that is the Poisson claim rate, which up-weights the client's objective probability of a loss. This is done by replacing the objective probability $\pi\tau$ with the higher subjective probability $\tilde{\pi}\tau > \pi\tau$. Loss aversion is incorporated as in the "rational expectations" loss aversion model of Kőszegi and Rabin, 2006, which serves as the benchmark implementation of loss aversion in agents choice among risky events, such as selecting the optimal deductibles (Kőszegi and Rabin, 2007; Sydnor, 2010; Barseghyan et al., 2013). In this case, the expected utility from consumption, defined as $\sum_{s\in S}\tilde{\pi}_s\tau U(c_s)$, where $\tilde{\pi}_s\tau$ and c_s represent the subjective probability of, and the level of the consumption at state s, respectively, is replaced by $\sum_{s\in S}\sum_{s'\in S}\tilde{\pi}_s\tau\tilde{\pi}_{s'}\tau\tilde{U}(c_s|\tilde{c}_{s'})$ where $\tilde{c}$ is the individual's reference point, in terms of consumption. Specifically,

$$\tilde{U}(c|\tilde{c}) = \begin{cases} U(c) + \eta(U(c) - U(\tilde{c})), & \text{if} \quad c \geq \tilde{c} \\ U(c) + \lambda\eta(U(c) - U(\tilde{c})), & \text{if} \quad c < \tilde{c}. \end{cases}$$

The gain-loss utility takes the usual two-part linear form. $\eta \geq 0$ measures of the importance of the reference point (gain-loss utility relative to the intrinsic utility from consumption), while $\lambda \geq 1$ measures the magnitude of loss aversion. Kőszegi and Rabin propose that the reference lottery is determined by recent expectations about outcomes. Specifically, when dealing with an insurance choice framework, they suggest the solution concept of "choice-acclimating personal equilibrium". By CPE, the reference lottery is set by the probability and outcomes of each state under the selected coverage plan. Since the individual's vehicle quality is independent of (i) whether a total loss event occurred, and of (ii) the selected coverage, CPE implies that $\mathcal{U}^c_{imt}$ can be re-expressed as follows:

$$\begin{aligned}\mathcal{U}^c_{imt} &= (1-\tilde{\pi}\tau)U(W-\rho_c\tau) + (\tilde{\pi}\tau)U(W-\rho_c\tau-d_c-L_{imt}) \\ &\quad + (\tilde{\pi}\tau)(1-\tilde{\pi}\tau)\Lambda\Big(U(W-\rho_c\tau-d_c-L_{imt}) - U(W-\rho_c\tau)\Big),\end{aligned}$$

where $\Lambda = \eta(\lambda - 1)$. Hence, individual i prefers a comprehensive coverage plan over partial insurance if and only if $\mathcal{U}^1_{imt} \geq \mathcal{U}^0_{imt}$, or:

$$U(W - \rho_0\tau) - U(W - \rho_1\tau) \leq \frac{\tilde{\pi}\tau + \Lambda\tilde{\pi}\tau(1-\tilde{\pi}\tau)}{1 - \tilde{\pi}\tau - \Lambda\tilde{\pi}\tau(1-\tilde{\pi}\tau)} \times \Big(U(W - \rho_1\tau - d_1 - L_{imt}) - U(W - \rho_0\tau - d_0 - L_{imt})\Big) \quad (3.7)$$

It is straightforward to see that a higher Λ (either due to a higher λ or a higher η) generates a higher willingness to pay for a comprehensive coverage plan. Furthermore, over-weighting of low probability events, i.e., $\tilde{\pi}\tau \in (\pi\tau, 0.5)$, results in a higher demand for insurance as well.[9] As stated before, the existence of the lemons gap generates a higher demand for insurance. Specifically, by strict concavity of U, the right-hand side of equation 3.7 increases with L_{imt}. Therefore, even moderate behavioral attributes are sufficient to generate substantial over-insurance patters, relative to insurance of an income loss. To see this, let $\rho_1(\tilde{\pi}\tau, \Lambda, L_{imt})$, defined over $(\tilde{\pi}\tau, \Lambda, L_{imt}) \in [0, 0.5] \times \mathbb{R}_+ \times \mathbb{R}_+$, denotes the willingness to pay for the comprehensive coverage plan. By definition, $\rho_1(\tilde{\pi}\tau, \Lambda, L_{imt})$ is the upper bound for which equation 3.7 holds (with equality). $\rho_1(\tilde{\pi}\tau, \Lambda, L_{imt})$ increases in all three inputs. Moreover, when the right-hand side of equation 3.7 (L_{imt}) is larger, presence of behavioral attributes result in a larger in the demand for comprehensive insurance coverage.

Corollary 8. *$\rho_1(\tilde{\pi}, \Lambda, L_{imt})$ increases in all three inputs. Furthermore, cross-derivative of the lemons gap and the behavioral attributes is positive; $\frac{\partial^2}{\partial L_{imt}\partial\Lambda}\rho_1(\tilde{\pi}, \Lambda, L_{imt}) > 0$ (for $\tilde{\pi}\tau > 0$), and $\frac{\partial^2}{\partial L_{imt}\partial\tilde{\pi}}\rho_1(\tilde{\pi}, \Lambda, L_{imt}) > 0$ (for $\tilde{\pi}\tau < 0.5$).*

Adverse Selection

In previous parts, I study the lemons gap impact on the demand for insurance at the individual level. Corollary 7 highlights that trends in the lemons gap result in a drop in demand over the vehicle life cycle. Yet, the magnitude of the drop in demand is heterogeneous across insurees. In this part, I examine the effect at the aggregate: both in level and slope. Specifically, I study the implications of the lemons gap on coverage selection over the vehicle life cycle.

[9]Throughout the analysis I restrict my attention to low probability events, i.e., $\pi\tau < \tilde{\pi}\tau < 0.5$. This is a reasonable assumption in this setting; Purchasing an insurance coverage can be thought of as a short term commitment, as coverage plans are linearly prorated and can be cancelled at any time - a standard in the insurance industry (Cohen and Einav, 2007). Furthermore, in the unlikely event that the probability of a total loss event is higher than 50%, over-weighting the probability of a claim might result in a decrease in the demand for comprehensive insurance as $\Lambda\tilde{\pi}\tau(1-\tilde{\pi}\tau)$ decreases with $\tilde{\pi}\tau$, for $\tilde{\pi}\tau > 0.5$.

Suppose that rational clients, that own a vehicle of the same model and age, differ in both risk and risk preferences. Let π_i (defined over $[\pi_0, \pi_1]$), and r_i (defined over $[r_0, r_1]$) denote individual i's claim rate, and the coefficient of risk aversion parameter, respectively. I assume that the third derivative of the vNM utility function is not too large.[10] Cohen and Einav, 2007 show that by taking limits with respect to τ (and applying L'Hopital's rule), we can approximate the individual's willingness to pay for the comprehensive insurance. Using a second-order Taylor expansion, individual i's (log) willingness to pay for the comprehensive coverage plan at age t, θ_{it}, is defined as follows:

$$\theta_{it} = \log(\pi_i) + \log(\Delta d) + \log(1 + r_i(L_t + \bar{d})), \tag{3.8}$$

where $\Delta d = d_0 - d_1$, and $\bar{d} = \frac{d_0+d_1}{2}$. The insuree buys the comprehensive coverage plan if and only if $\theta_{it} \geq \log \Delta\rho = \log(\rho_1 - \rho_0)$.

Using equation 3.8, we can identify the role of the lemons gap in terms of trends in selection. Selection is static when the insurance coverage provides full insurance ($L_t = 0$). To illustrate, consider two individuals, denoted by h and l. Suppose h has a higher willingness to pay for comprehensive insurance at age t, $\theta_{ht} \geq \theta_{lt}$. Corollary 7 states that demand is static for $L_t = 0$. Therefore $\theta_{ht+1} = \theta_{ht} \geq \theta_{lt} = \theta_{lt+1}$. Thus, h has a willingness to pay for comprehensive insurance at age $t+1$ as well.

In contrast, when the insurance market is incomplete, the uncompensated loss has an heterogeneous effect on the demand for insurance. By equation 3.8, individuals with different characteristics will sort into the comprehensive coverage plan depending on their risk (π_i) and risk preferences (r_i). Both high risk types and risk averse individuals have a high preference for the comprehensive coverage plan. The relative importance of each component depends on the lemons gap (and the deductibles). The reduction of the lemons gap over the vehicle life cycle affects the sorting of individuals. Over the vehicle life cycle, the importance of risk aversion drops, while the relative importance of risk rises. As a result, the insurance market is more adversely selected as the vehicle ages. Over the vehicle life cycle, low risk averse - high risk individuals are sorted higher on the demand curve for comprehensive insurance coverage.

Proposition 9. *If the insured asset is perfectly replaceable ($L_t = 0$), the ordering of the demand for insurance is static. In contrast, if the insuree is partially compensated via a pre-owned market product ($L_t > L_{t+1} > 0$), the market becomes more adversely selected over the vehicle life cycle.*

[10] The results in this section hold for standard utility functions with a non-zero third derivative, such as the constant absolute risk aversion (CARA) or the constant relative risk aversion (CRRA) utility functions.

Proof: *See appendix.*

Figure 3.4 depicts the conclusion from corollary 7 and proposition 9. The quantity of the comprehensive coverage plans sold (in terms of the share of vehicles) is measured on the horizontal axis and the willingness to pay for a comprehensive insurance coverage plan is on the vertical axis. The solid lines denote an arbitrary demand curve at age t (in gray), $D(t)$, and at age $t+1$ (in black), $D(t+1)$. The dashed curves are the average cost of the insured clients who purchase comprehensive coverage plan at age t (in gray), $AC(t)$, and at age $t+1$ (in black), $AC(t+1)$. By corollary 7, the demand curve drops with vehicle age. Furthermore, by proposition 9, the average cost curve shifts upwards (for $Q \in (0,1)$), as the market is more adversely selected over the vehicle life cycle. The gray and black dots denote the perfectly competitive equilibrium at age t and $t+1$, respectively.[11] In equilibrium, the number of comprehensive coverage plans sold drops, and the price of comprehensive coverage rise over the vehicle life cycle.

Moral Hazard

In this part I depart from the analysis of the interaction between the lemons gap and the demand for insurance. Instead, I examine the lemons gap role in providing incentive to reduce endogenous risk. In particular, I study the interaction between the replacement value and moral hazard. As before, I show that the nature of the relationship depends on whether the replacement product is an imperfect substitute (if $L_{imt} > 0$), and if so, on the vehicle's age as well ($L_{imt} > L_{imt+1} > 0$).

Consider a client that purchased insurance coverage c (both at t and $t+1$). Unlike before, risk is endogenous. By exerting effort, the client can reduce the probability of a total-loss event. Yet, effort is costly. I reformulate equation 3.5 to implement moral hazard in the individual i's expected utility from insurance coverage plan c at

[11] These are the equilibrium price and quantities when we consider (i) at least 2 identical risk neutral insurance providers who take the offered coverage plans (deductibles) as given, and (ii) the demand curve crosses the marginal cost curve at most once. In this scenario, equilibrium price (and quantity) is characterized by the lowest price, subject to zero profits, i.e. the minimum price that equals the average cost of consumers who purchase the comprehensive coverage plan, given the price (see Einav, Finkelstein, and Cullen, 2010b). That is, the product of the probability of a total loss and the difference in deductibles (between the comprehensive and partial plan). The notion of equilibrium presented here, as in Einav, Finkelstein, and Cullen, 2010b, considers a case in which competition occurs only over the comprehensive insurance coverage plan. Equilibrium characterization differs (and might not exists) when considering competition over the two coverage plans (see Handel, Hendel, and Whinston, 2015b).

vehicle age t.

$$\mathcal{U}^c_{imt}(e) = (1-\pi(e)\tau)U(W-\rho_c\tau)+\pi(e)\tau U(W-\rho_c\tau-d_c-L_{imt})+\phi(q_{imt}\omega_t)-\tau C(e), \tag{3.9}$$

where e (defined on $[0,\infty)$) is the exerted level of effort. Effort decreases the claim probability, $\pi'(e) < 0$, at a diminishing rate, $\pi''(e) > 0$. $C(e)$ denotes the cost of effort, which is increasing, $C'(e) > 0$, and convex $C''(e) \geq 0$. I also assume that (i) $\pi\tau > 0$, and (ii) the marginal cost is zero when no effort is exerted, $C'(0) = 0$. These assumptions insure a unique interior solution exists.

The optimal level of effort at age t, e^*_t, is defined by the first order condition of equation 3.9.

$$\frac{\partial}{\partial e}\mathcal{U}^c_{imt}(e) = 0 \Rightarrow -\pi'(e^*_t)(U(W-\rho_c\tau) - U(W-\rho_c\tau-d_c-L_{imt})) = C'(e^*_t) \tag{3.10}$$

Equation 3.10 highlights the relevant components in setting the optimal level of effort. As theory predicts, effort is higher when the deductible is higher. Moreover, the lemons gap's role is similar to that of the deductible. A replacement vehicle which is an inadequate substitute is equivalent to a higher deductible payment. As a result, effort increases with the lemons gap.

With regards to trends over the vehicle life cycle, when the insurance market provides full coverage (up to a deductible), $L_{imt} = 0$, the insuree's optimal level of effort, e^*_t, does not change as the vehicle ages (given the same policy). As with demand and selection, the uncompensated loss affects the optimal level of effort differentially over the vehicle life cycle. When there is an uncompensated loss which decreases with vehicle age, effort drops over the vehicle life cycle (given the same coverage), $e^*_t > e^*_{t+1}$, as the benefit from avoiding a total-loss event, $U(W-\rho_c\tau) - U(W-\rho_c\tau-d_c-L_t)$, drops as the vehicle ages.

Corollary 10. *If the insured asset is perfectly replaceable ($L_t = 0$), exerted effort is static. Yet, if the insuree is partially compensated via a pre-owned market product ($L_{imt} > 0$), the exerted effort is higher. Moreover, exerted effort drops over the vehicle life cycle.*

By corollary 10, the drop in the lemons gap can explain why vehicle owners tend to put less effort in maintaining their old, high-mileage "beater" cars.

3.4 Conclusion

In this paper, I study the implications of insuring a quality uncertain good relative to a item of known value (or an income loss). Specifically, I examine the differences between the item and its replacement. I do so by incorporating valuable insights from the lemons market into the analysis of the demand for insurance. I find that insurance market is incomplete. Due to quality uncertainty, customers cannot insure their vehicle based on its quality. Rather, they insure their own vehicle based on the average quality in the pre-owned market. The adverse selection in the resale market implies that there is a lemons gap, an uncompensated loss, due to the pre-owned market's inferior quality. The lemons gap explains why the customer's demand for comprehensive insurance is context-dependent, and why clients over-insure possessions, such as real estate and automobiles, by paying too much to face a lower deductible, relative to an income loss. This is in contrast to the insurance literature, which views insurance as a financial asset.

The lemons gap is of considerable significance on the analysis of the insurance market. The uncompensated loss affects the demand for insurance both in levels and in trends. The lemons gap increases the economic deductibles. Due to risk aversion, the demand for insurance is higher for quality uncertain goods, relative to insurance of a monetary loss of the same value. In addition, the lemons gap interacts with behavioral attributes such as probability distortion and loss aversion, to further increase the demand for insurance. Thus, even moderate behavioral attributes can generate substantial over-insurance patters.

Moreover, lemons gap has a differential impact on the client's willingness to pay for insurance. As the asset depreciates, the lemons gap drops. This occurs since both the pre-owned market quality distribution improves over the vehicle life cycle, and the uncompensated loss from an inadequate substitute is of less importance. The drop in the lemons gap drives demand down. Yet, the effect is heterogeneous. Selection to comprehensive coverage becomes adversely selected. The lemons gap affects the cost of insurance as well. The insuree's exerted effort (loss probability) drops (rise) with the lemons gap. As a result, moral hazard is exacerbated over the vehicle life cycle.

This paper contributes to both the theoretical and empirical insurance literature. The theoretical literature did not make a distinction between quality certain and uncertain goods, but rather assumes that the value is known to both parties. I demonstrate how the difference between insuring the loss of a quality uncertain good and a monetary loss can have substantial consequences on demand, adverse selection, and moral hazard. Empirically, the cutting-edge insurance literature does not take the uncompensated loss into account. This implies that the empirical structural models in use effectively assume that both adverse selection, and moral hazard are independent of the vehicle's age. Thus, not taking the lemons gap into account may lead to biased results.

Figure 3.1: Probability of Sale by Vehicle Quality

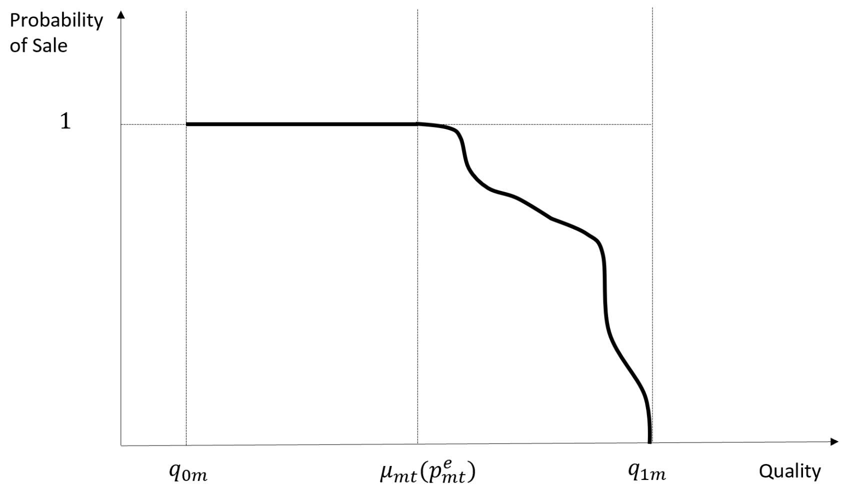

Figure 3.2: Probability of Sale by Vehicle Quality

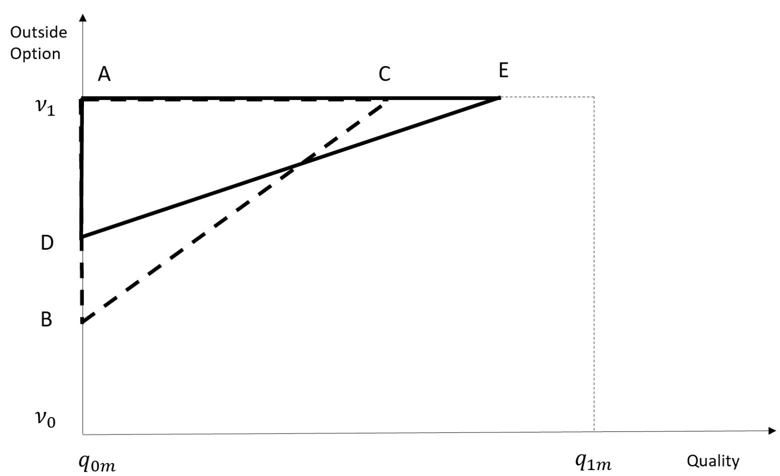

Figure 3.3: Equilibrium Price and Vehicle Age

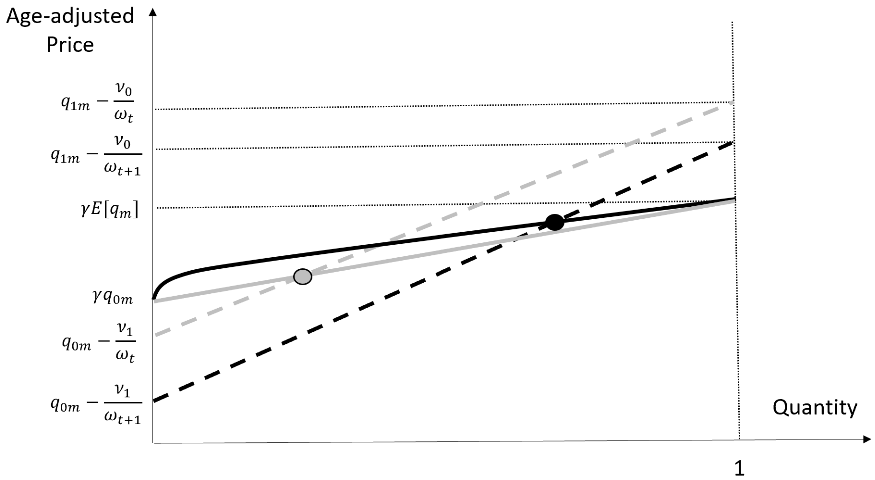

Figure 3.4: Demand and Average Cost Curves by Vehicle Age

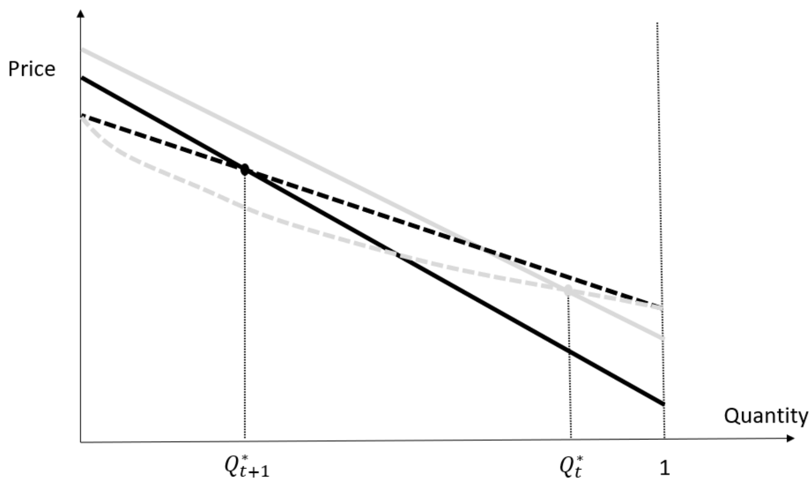

Bibliography

[1] Jason Abaluck and Jonathan Gruber. "Choice inconsistencies among the elderly: evidence from plan choice in the Medicare Part D program". In: *American Economic Review* 101.4 (2011), pp. 1180–1210.

[2] George A Akerlof. "The Market for "Lemons": Quality Uncertainty and the Market Mechanism". In: *The Quarterly Journal of Economics* 84.3 (1970), pp. 488–500.

[3] Joseph G Altonji and Rebecca M Blank. "Race and gender in the labor market". In: *Handbook of labor economics* 3 (1999), pp. 3143–3259.

[4] Dean F Amel. "State laws affecting the geographic expansion of commercial banks". In: (1993).

[5] Abhay Aneja and Prasad Krishnamurthy. "Merger Deregulation, Wages, and Inequality: Evidence from the US Banking Industry". In: *Wages, and Inequality: Evidence from the US Banking Industry (July 24, 2022)* (2022).

[6] Kenneth J Arrow. "Some Mathematical Models of Race in the Labor Market." In: *In Racial Discrimination in Economic Life* (1972), pp. 187–203.

[7] Orley Ashenfelter and Timothy Hannan. "Sex discrimination and product market competition: The case of the banking industry". In: *The Quarterly Journal of Economics* 101.1 (1986), pp. 149–173.

[8] David Autor and David Dorn. "The growth of low-skill service jobs and the polarization of the US labor market". In: *American economic review* 103.5 (2013), pp. 1553–97.

[9] Ian Ayres, Mahzarin Banaji, and Christine Jolls. "Race effects on eBay". In: *The RAND Journal of Economics* 46.4 (2015), pp. 891–917.

[10] Revital Bar and Asaf Zussman. "Customer discrimination: evidence from Israel". In: *Journal of Labor Economics* 35.4 (2017), pp. 1031–1059.

[11] Levon Barseghyan, Jeffrey Prince, and Joshua C Teitelbaum. "Are risk preferences stable across contexts? Evidence from insurance data". In: *American Economic Review* 101.2 (2011), pp. 591–631.

[12] Levon Barseghyan et al. "The nature of risk preferences: Evidence from insurance choices". In: *American economic review* 103.6 (2013), pp. 2499–2529.

[13] Patrick Bayer and Kerwin Kofi Charles. "Divergent paths: A new perspective on earnings differences between black and white men since 1940". In: *The Quarterly Journal of Economics* 133.3 (2018), pp. 1459–1501.

[14] Thorsten Beck, Ross Levine, and Alexey Levkov. "Big bad banks? The winners and losers from bank deregulation in the United States". In: *The Journal of Finance* 65.5 (2010), pp. 1637–1667.

[15] Gary S Becker. "The Economics of Discrimination". In: (1957).

[16] Gary S Becker. "The Economics of Discrimination". In: *University of Chicago Press Economics Books* (1971).

[17] Steven T Berry. "Estimating discrete-choice models of product differentiation". In: *The RAND Journal of Economics* (1994), pp. 242–262.

[18] Clémence Berson. "Does competition induce hiring equity?" In: (2012).

[19] Marianne Bertrand and Sendhil Mullainathan. "Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination". In: *American economic review* 94.4 (2004), pp. 991–1013.

[20] Saurabh Bhargava, George Loewenstein, and Justin Sydnor. "Choose to lose: Health plan choices from a menu with dominated option". In: *The Quarterly Journal of Economics* 132.3 (2017), pp. 1319–1372.

[21] Dan A Black. "Discrimination in an equilibrium search model". In: *Journal of labor Economics* 13.2 (1995), pp. 309–334.

[22] Sandra E Black and Elizabeth Brainerd. "Importing Equality? The Impact of Globalization on Gender Discrimination". In: *Industrial & labor relations review* 57.4 (2004), pp. 540–559.

[23] Sandra E Black and Philip E Strahan. "The division of spoils: rent-sharing and discrimination in a regulated industry". In: *American Economic Review* 91.4 (2001), pp. 814–831.

[24] Nicholas Bloom and John Van Reenen. "Measuring and explaining management practices across firms and countries". In: *The quarterly journal of Economics* 122.4 (2007), pp. 1351–1408.

[25] Nicholas Bloom et al. “Does management matter? Evidence from India”. In: *The Quarterly journal of economics* 128.1 (2013), pp. 1–51.

[26] George J Borjas and Stephen G Bronars. “Consumer discrimination and self-employment”. In: *Journal of political economy* 97.3 (1989), pp. 581–605.

[27] Zarek C Brot-Goldberg et al. “What does a deductible do? The impact of cost-sharing on health care prices, quantities, and spending dynamics”. In: *The Quarterly Journal of Economics* 132.3 (2017), pp. 1261–1318.

[28] M Kate Bundorf, Jonathan Levin, and Neale Mahoney. “Pricing and welfare in health plan choice”. In: *American Economic Review* 102.7 (2012), pp. 3214–48.

[29] Marika Cabral. “Claim timing and ex post adverse selection”. In: *The Review of Economic Studies* 84.1 (2016), pp. 1–44.

[30] Marika Cabral, Michael Geruso, and Neale Mahoney. “Do larger health insurance subsidies benefit patients or producers? Evidence from Medicare Advantage”. In: *American Economic Review* 108.8 (2018), pp. 2048–87.

[31] Glen G Cain. “The economic analysis of labor market discrimination: A survey”. In: *Handbook of labor economics* 1 (1986), pp. 693–785.

[32] James H Cardon and Igal Hendel. “Asymmetric information in health insurance: evidence from the National Medical Expenditure Survey”. In: *RAND Journal of Economics* (2001), pp. 408–427.

[33] Caroline Carlin and Robert Town. “Adverse selection, welfare, and optimal pricing of employer sponsored health plans”. In: *U. Minnesota Working Paper* 1.2 (2009).

[34] Alberto Cavallo, Brent Neiman, and Roberto Rigobon. “Currency unions, product introductions, and the real exchange rate”. In: *The Quarterly Journal of Economics* 129.2 (2014), pp. 529–595.

[35] Kerwin Kofi Charles and Jonathan Guryan. “Prejudice and wages: an empirical assessment of Becker’s The Economics of Discrimination”. In: *Journal of political economy* 116.5 (2008), pp. 773–809.

[36] Pierre-André Chiappori and Bernard Salanié. “Asymmetric information in insurance markets: Predictions and tests”. In: *Handbook of insurance* (2013), pp. 397–422.

[37] Pierre-André Chiappori and Bernard Salanie. “Testing for asymmetric information in insurance markets”. In: *Journal of political Economy* 108.1 (2000), pp. 56–78.

[38] Sungjin Cho and John Rust. "The flat rental puzzle". In: *The Review of Economic Studies* 77.2 (2010), pp. 560–594.

[39] Alma Cohen. "Asymmetric information and learning: Evidence from the automobile insurance market". In: *Review of Economics and statistics* 87.2 (2005), pp. 197–207.

[40] Alma Cohen and Liran Einav. "Estimating risk preferences from deductible choice". In: *American economic review* 97.3 (2007), pp. 745–788.

[41] Alma Cohen and Peter Siegelman. "Testing for adverse selection in insurance markets". In: *Journal of Risk and insurance* 77.1 (2010), pp. 39–84.

[42] Major G Coleman. "Racial discrimination in the workplace: does market structure make a difference?" In: *Industrial Relations: A Journal of Economy and Society* 43.3 (2004), pp. 660–689.

[43] Pierre-Philippe Combes et al. "Customer discrimination and employment outcomes: theory and evidence from the french labor market". In: *Journal of Labor Economics* 34.1 (2016), pp. 107–160.

[44] Dudley Cooke, Ana P Fernandes, and Priscila Ferreira. "Product market competition and gender discrimination". In: *Journal of Economic Behavior & Organization* 157 (2019), pp. 496–522.

[45] Gregory S Crawford, Nicola Pavanini, and Fabiano Schivardi. "Asymmetric information and imperfect competition in lending markets". In: *American Economic Review* 108.7 (2018), pp. 1659–1701.

[46] José Ignacio Cuesta and Alberto Sepúlveda. "Price regulation in credit markets: A trade-off between consumer protection and credit access". In: *Available at SSRN 3282910* (2021).

[47] David M Cutler, Amy Finkelstein, and Kathleen McGarry. "Preference heterogeneity and insurance markets: Explaining a puzzle of insurance". In: *American Economic Review* 98.2 (2008), pp. 157–62.

[48] David M Cutler and Sarah J Reber. "Paying for health insurance: the trade-off between competition and adverse selection". In: *The Quarterly Journal of Economics* 113.2 (1998), pp. 433–466.

[49] David M Cutler and Richard Zeckhauser. "Extending the theory to meet the practice of insurance". In: *Brookings-Wharton Papers on Financial Services* 2004.1 (2004), pp. 1–53.

[50] David De Meza and David C Webb. "Advantageous selection in insurance markets". In: *RAND Journal of Economics* (2001), pp. 249–262.

[51] Stefano DellaVigna and Matthew Gentzkow. "Uniform pricing in us retail chains". In: *The Quarterly Journal of Economics* 134.4 (2019), pp. 2011–2084.

[52] David J Deming. "The growing importance of social skills in the labor market". In: *The Quarterly Journal of Economics* 132.4 (2017), pp. 1593–1640.

[53] Yuliya Demyanyk, Charlotte Ostergaard, and Bent E Sørensen. "US banking deregulation, small businesses, and interstate insurance of personal income". In: *The Journal of Finance* 62.6 (2007), pp. 2763–2801.

[54] Avinash Dixit. "Quality and quantity competition". In: *The Review of Economic Studies* 46.4 (1979), pp. 587–599.

[55] Avinash K Dixit and Joseph E Stiglitz. "Monopolistic competition and optimum product diversity". In: *The American economic review* 67.3 (1977), pp. 297–308.

[56] Jennifer L Doleac and Luke CD Stein. "The visible hand: Race and online market outcomes". In: *The Economic Journal* 123.572 (2013), F469–F492.

[57] John J Donohue and James J Heckman. *Continuous versus episodic change: The impact of civil rights policy on the economic status of blacks*. 1991.

[58] Nicholas Economides, R Glenn Hubbard, and Darius Palia. "The political economy of branching restrictions and deposit insurance: A model of monopolistic competition among small and large banks". In: *The Journal of Law and Economics* 39.2 (1996), pp. 667–704.

[59] Liran Einav, Amy Finkelstein, and Mark R Cullen. "Estimating welfare in insurance markets using variation in prices". In: *The quarterly journal of economics* 125.3 (2010), pp. 877–921.

[60] Liran Einav, Amy Finkelstein, and Mark R Cullen. "Estimating welfare in insurance markets using variation in prices". In: *The Quarterly Journal of Economics* 125.3 (2010), pp. 877–921.

[61] Liran Einav, Amy Finkelstein, and Jonathan Levin. "Beyond testing: Empirical models of insurance markets". In: *Annu. Rev. Econ.* 2.1 (2010), pp. 311–336.

[62] Liran Einav, Amy Finkelstein, and Jonathan Levin. "Beyond testing: Empirical models of insurance markets". In: *Annual review of economics.* 2.2010 (2010), pp. 311–336.

[63] Liran Einav, Amy Finkelstein, and Neale Mahoney. "The IO of selection markets". In: *Handbook of Industrial Organization*. Vol. 5. 1. Elsevier, 2021, pp. 389–426.

[64] Liran Einav, Amy Finkelstein, and Paul Schrimpf. *The welfare cost of asymmetric information: Evidence from the UK annuity market*. 2007.

[65] Liran Einav, Amy Finkelstein, and Pietro Tebaldi. "Market design in regulated health insurance markets: Risk adjustment vs. subsidies". In: *Unpublished mimeo, Stanford University, MIT, and University of Chicago* 7 (2019), p. 32.

[66] Liran Einav et al. "How general are risk preferences? Choices under uncertainty in different domains". In: *American Economic Review* 102.6 (2012), pp. 2606–38.

[67] Liran Einav et al. "Selection on moral hazard in health insurance". In: *American Economic Review* 103.1 (2013), pp. 178–219.

[68] Hanming Fang, Michael P Keane, and Dan Silverman. "Sources of advantageous selection: Evidence from the Medigap insurance market". In: *Journal of political Economy* 116.2 (2008), pp. 303–350.

[69] Amy Finkelstein and Kathleen McGarry. "Multiple dimensions of private information: evidence from the long-term care insurance market". In: *American Economic Review* 96.4 (2006), pp. 938–958.

[70] Amy Finkelstein, Kathleen McGarry, and Amir Sufi. "Dynamic inefficiencies in insurance markets: Evidence from long-term care insurance". In: *American Economic Review* 95.2 (2005), pp. 224–228.

[71] Amy Finkelstein and James Poterba. "Testing for asymmetric information using "unused observables" in insurance markets: Evidence from the UK annuity market". In: *Journal of Risk and Insurance* 81.4 (2014), pp. 709–734.

[72] Sebastian Fleitas, Gautam Gowrisankaran, and Anthony Lo Sasso. "Reclassification Risk in the Small Group Health Insurance Market". In: (2020).

[73] Nicola Gennaioli and Andrei Shleifer. "What comes to mind". In: *The Quarterly journal of economics* 125.4 (2010), pp. 1399–1433.

[74] Soheil Ghili et al. "Optimal long-term health insurance contracts: characterization, computation, and welfare effects". In: (2021).

[75] Matthew S Goldberg. "Discrimination, nepotism, and long-run wage differentials". In: *The quarterly journal of economics* 97.2 (1982), pp. 307–319.

[76] Daniel Gottlieb and Kent Smetters. “Lapse-based insurance”. In: *American Economic Review* 111.8 (2021), pp. 2377–2416.

[77] Martin B Hackmann, Jonathan T Kolstad, and Amanda E Kowalski. “Adverse selection and an individual mandate: When theory meets practice”. In: *American Economic Review* 105.3 (2015), pp. 1030–66.

[78] Ben Handel, Igal Hendel, and Michael D Whinston. “Equilibria in health exchanges: Adverse selection versus reclassification risk”. In: *Econometrica* 83.4 (2015), pp. 1261–1313.

[79] Ben Handel, Igal Hendel, and Michael D Whinston. “Equilibria in health exchanges: Adverse selection versus reclassification risk”. In: *Econometrica* 83.4 (2015), pp. 1261–1313.

[80] Ben Handel and Kate Ho. “The industrial organization of health care markets”. In: *Handbook of Industrial Organization*. Vol. 5. 1. Elsevier, 2021, pp. 521–614.

[81] Benjamin R Handel. “Adverse selection and inertia in health insurance markets: When nudging hurts”. In: *American Economic Review* 103.7 (2013), pp. 2643–82.

[82] Benjamin R Handel and Jonathan T Kolstad. “Health insurance for” humans”: Information frictions, plan choice, and consumer welfare”. In: *American Economic Review* 105.8 (2015), pp. 2449–2500.

[83] Benjamin R Handel, Jonathan T Kolstad, and Johannes Spinnewijn. “Information frictions and adverse selection: Policy interventions in health insurance markets”. In: *Review of Economics and Statistics* 101.2 (2019), pp. 326–340.

[84] Benjamin R Handel, Jonathan T Kolstad, and Johannes Spinnewijn. “Information frictions and adverse selection: Policy interventions in health insurance markets”. In: *Review of Economics and Statistics* 101.2 (2019), pp. 326–340.

[85] James J Heckman. “Detecting discrimination”. In: *Journal of economic perspectives* 12.2 (1998), pp. 101–116.

[86] Judith K Hellerstein, David Neumark, and Kenneth R Troske. *Market forces and sex discrimination*. 1997.

[87] Igal Hendel and Alessandro Lizzeri. “Adverse selection in durable goods markets”. In: *American Economic Review* 89.5 (1999), pp. 1097–1115.

[88] Igal Hendel and Alessandro Lizzeri. “The role of commitment in dynamic contracts: Evidence from life insurance”. In: *The Quarterly journal of economics* 118.1 (2003), pp. 299–328.

[89] Nathaniel Hendren. “Knowledge of future job loss and implications for unemployment insurance”. In: *American Economic Review* 107.7 (2017), pp. 1778–1823.

[90] Fredrik Heyman, Helena Svaleryd, and Jonas Vlachos. “Competition, takeovers, and gender discrimination”. In: *ILR Review* 66.2 (2013), pp. 409–432.

[91] Guilherme Hirata and Rodrigo R Soares. “Competition and the racial wage gap: Evidence from Brazil”. In: *Journal of Development Economics* 146 (2020), p. 102519.

[92] Kate Ho, Joseph Hogan, and Fiona Scott Morton. “The impact of consumer inattention on insurer pricing in the Medicare Part D program”. In: *The RAND Journal of Economics* 48.4 (2017), pp. 877–905.

[93] Harry J Holzer and Keith R Ihlanfeldt. “Customer discrimination and employment outcomes for minority workers”. In: *The Quarterly Journal of Economics* 113.3 (1998), pp. 835–867.

[94] Erik Hurst, Yona Rubinstein, and Kazuatsu Shimizu. *Task-based discrimination*. Tech. rep. National Bureau of Economic Research, 2021.

[95] Yizhou Jin and Shoshana Vasserman. *Buying Data from Consumers: The Impact of Monitoring in US Auto Insurance*. 2020.

[96] Lawrence M Kahn. “Customer discrimination and affirmative action”. In: *Economic Inquiry* 29.3 (1991), pp. 555–571.

[97] Lawrence M Kahn and Peter D Sherer. “Racial differences in professional basketball players’ compensation”. In: *Journal of Labor Economics* 6.1 (1988), pp. 40–61.

[98] Daniel Kahneman and Amos Tversky. “Prospect Theory: An Analysis of Decision under Risk”. In: *Econometrica* 47.2 (1979), pp. 263–292.

[99] Edward J Kane. “De jure interstate banking: Why only now?” In: *Journal of Money, Credit and Banking* 28.2 (1996), pp. 141–161.

[100] Patrick Kline, Evan K Rose, and Christopher R Walters. “Systemic discrimination among large US employers”. In: *The Quarterly Journal of Economics* 137.4 (2022), pp. 1963–2036.

[101] Thomas G Koch. “One pool to insure them all? Age, risk and the price (s) of medical insurance”. In: *International Journal of Industrial Organization* 35 (2014), pp. 1–11.

[102] Botond Kőszegi and Matthew Rabin. "A model of reference-dependent preferences". In: *The Quarterly Journal of Economics* 121.4 (2006), pp. 1133–1165.

[103] Botond Kőszegi and Matthew Rabin. "Reference-dependent risk attitudes". In: *American Economic Review* 97.4 (2007), pp. 1047–1073.

[104] Randall S Kroszner and Philip E Strahan. "What drives deregulation? Economics and politics of the relaxation of bank branching restrictions". In: *The Quarterly Journal of Economics* 114.4 (1999), pp. 1437–1467.

[105] Howard Kunreuther and Mark Pauly. *Insurance decision-making and market behavior*. now publishers Inc, 2006.

[106] Kevin Lang and Ariella Kahn-Lang Spitzer. "Race discrimination: An economic perspective". In: *Journal of Economic Perspectives* 34.2 (2020), pp. 68–89.

[107] Kevin Lang and Jee-Yeon K Lehmann. "Racial discrimination in the labor market: Theory and empirics". In: *Journal of Economic Literature* 50.4 (2012), pp. 959–1006.

[108] Kevin Lang, Michael Manove, and William T Dickens. "Racial discrimination in labor markets with posted wage offers". In: *American Economic Review* 95.4 (2005), pp. 1327–1340.

[109] Pamela K Lattimore, Joanna R Baker, and Ann D Witte. "The influence of probability on risky choice: A parametric examination". In: *Journal of economic behavior & organization* 17.3 (1992), pp. 377–400.

[110] Jonathan S Leonard, David I Levine, and Laura Giuliano. "Customer discrimination". In: *The Review of Economics and Statistics* 92.3 (2010), pp. 670–678.

[111] Benjamin Lester et al. "Screening and adverse selection in frictional markets". In: *Journal of Political Economy* 127.1 (2019), pp. 338–377.

[112] Ross Levine, Alexey Levkov, and Yona Rubinstein. *Racial discrimination and competition*. Tech. rep. National Bureau of Economic Research, 2008.

[113] John A List. "The nature and extent of discrimination in the marketplace: Evidence from the field". In: *The Quarterly Journal of Economics* 119.1 (2004), pp. 49–89.

[114] Josh Lustig. "Measuring welfare losses from adverse selection and imperfect competition in privatized Medicare". In: *Manuscript. Boston University Department of Economics* (2010).

[115] Neale Mahoney and E Glen Weyl. "Imperfect competition in selection markets". In: *Review of Economics and Statistics* 99.4 (2017), pp. 637–651.

[116] Robert Stanton McMillan. "Different flavor, same price: The puzzle of uniform pricing for differentiated products". In: *Same Price: The Puzzle of Uniform Pricing for Differentiated Products (January 17, 2007)* (2007).

[117] David de Meza and Diane Reyniers. "Insuring replaceable possessions: The irrelevance of risk aversion coefficients". In: (2020).

[118] Clark Nardinelli and Curtis Simon. "Customer racial discrimination in the market for memorabilia: The case of baseball". In: *The Quarterly Journal of Economics* 105.3 (1990), pp. 575–595.

[119] Paula Onuchic and Debraj Ray. *Signaling and Discrimination in Collaborative Projects*. Tech. rep. National Bureau of Economic Research, 2021.

[120] Barak Y Orbach and Liran Einav. "Uniform prices for differentiated goods: The case of the movie-theater industry". In: *International Review of Law and Economics* 27.2 (2007), pp. 129–153.

[121] James Peoples and Wayne K Talley. "Black-white earnings differentials: privatization versus deregulation". In: *American Economic Review* 91.2 (2001), pp. 164–168.

[122] Drazen Prelec. "The probability weighting function". In: *Econometrica* (1998), pp. 497–527.

[123] Åsa Rosén. "An equilibrium search-matching model of discrimination". In: *European Economic Review* 41.8 (1997), pp. 1589–1613.

[124] Sherwin Rosen. "Hedonic prices and implicit markets: product differentiation in pure competition". In: *Journal of political economy* 82.1 (1974), pp. 34–55.

[125] Michael Rothschild and Joseph Stiglitz. "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information". In: *The Quarterly Journal of Economics* 90.4 (1976), pp. 629–649.

[126] Harris Schlesinger. "The theory of insurance demand". In: *Handbook of insurance*. Springer, 2013, pp. 167–184.

[127] Eytan Sheshinski. "Price, quality and quantity regulation in monopoly situations". In: *Economica* 43.170 (1976), pp. 127–137.

[128] Ben Shiller and Joel Waldfogel. "Music for a song: an empirical look at uniform pricing and its alternatives". In: *The Journal of Industrial Economics* 59.4 (2011), pp. 630–660.

[129] Adam Smith. "The Wealth of Nations". In: (1776).

[130] A Michael Spence. "Monopoly, quality, and regulation". In: *The Bell Journal of Economics* (1975), pp. 417–429.

[131] Michael Spence. "Nonprice competition". In: *The American Economic Review* 67.1 (1977), pp. 255–259.

[132] Johannes Spinnewijn. "Heterogeneity, demand for insurance, and adverse selection". In: *American Economic Journal: Economic Policy* 9.1 (2017), pp. 308–43.

[133] Amanda Starc. "Insurer pricing and consumer welfare: Evidence from medigap". In: *The RAND Journal of Economics* 45.1 (2014), pp. 198–220.

[134] Justin Sydnor. "(Over) insuring modest risks". In: *American Economic Journal: Applied Economics* 2.4 (2010), pp. 177–99.

[135] Pietro Tebaldi. *Estimating equilibrium in health insurance exchanges: Price competition and subsidy design under the aca*. Tech. rep. National Bureau of Economic Research, 2022.

[136] Nicholas Tilipman. "Employer Incentives and Distortions in Health Insurance Design: Implications for Welfare and Costs". In: *American Economic Review* 112.3 (2022), pp. 998–1037.

[137] Kenneth Train. "Halton sequences for mixed logit". In: (2000).

[138] Kenneth E Train. *Discrete choice methods with simulation*. Cambridge university press, 2009.

[139] Amos Tversky and Daniel Kahneman. "Advances in prospect theory: Cumulative representation of uncertainty". In: *Journal of Risk and uncertainty* 5.4 (1992), pp. 297–323.

[140] André Veiga and E Glen Weyl. "Product design in selection markets". In: *The Quarterly Journal of Economics* 131.2 (2016), pp. 1007–1056.

[141] Andrea Weber and Christine Zulehner. "Competition and gender prejudice: Are discriminatory employers doomed to fail?" In: *Journal of the European Economic Association* 12.2 (2014), pp. 492–521.

[142] Charles Wilson. "The nature of equilibrium in markets with adverse selection". In: *The Bell Journal of Economics* (1980), pp. 108–130.

Appendix A

Appendix of Behavioral Professionals: Evidence From the Commercial Auto Insurance Industry

A.1 Additional Figures and Tables

Figure A.1: Distribution of policies by client's fleet size

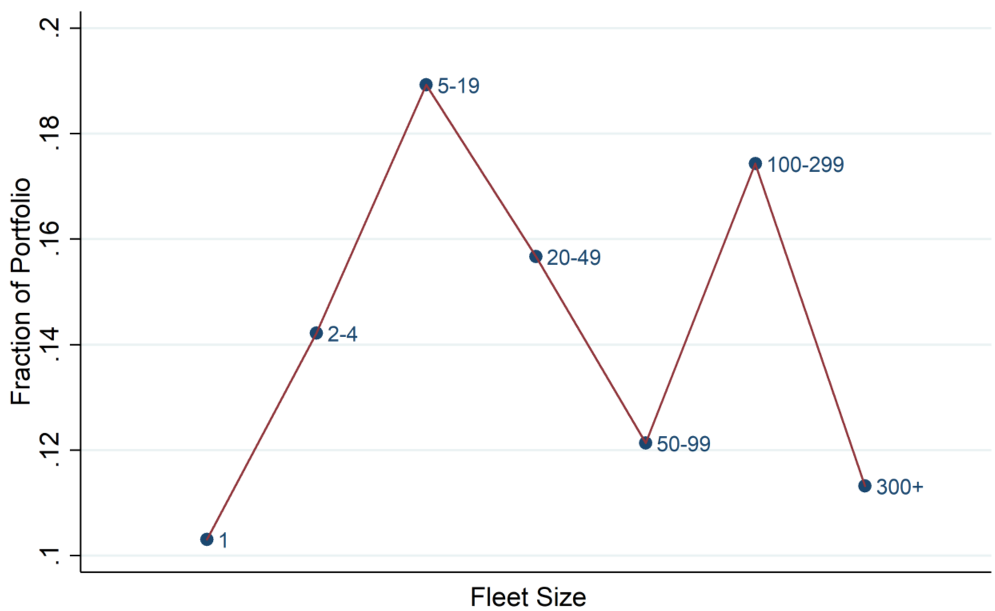

Notes: The figure reports the distribution of all insurance policies from 2013 to 2020 by client's fleet size. Client's fleet size is defined by the number of total insurance policies purchased by the client in a given year. The distribution of policies is weighted by policy premiums, which are measured in New Israeli Shekel (ILS).

Figure A.2: Policies by coverage type and vehicle type

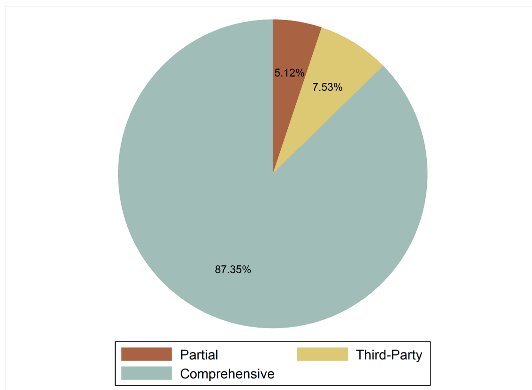

Panel A: Coverage type

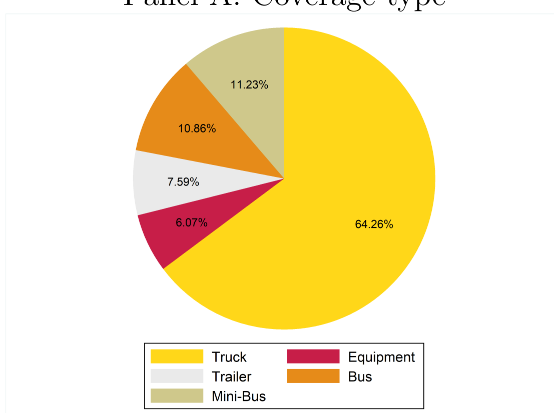

Panel A: Panel B: Vehicle type

Notes: The figure depicts the distribution of policies by coverage type (Panel A) and vehicle type (Panel B). The sample include all insurance policies for all vehicles from 2013 through 2020. Policies are weighted by premiums, measured in New Israeli Shekel (ILS).

Figure A.3: Example of a "Go—No Go" grade document

	הערות	תאריך סוף פוליסה	תאריך תחילת פוליסה	פוליסה
Check	✓	2/29/2020	3/1/2019	390342117419
Same as last year	חידוש ללא שינוי	2/29/2020	3/1/2019	390344197219
Do not decrease	אין לרדת מתעריף\תנאים קיימים	2/29/2020	3/1/2019	390344983119
Increase third party deductible to 7000	יש להעלות אקסס צד ג' ל-7,000 ₪	2/29/2020	3/1/2019	390342188019
Increase premiums by 7.5%	יש להעלות תעריף ב-7.5% מאשתקד	2/29/2020	3/1/2019	390880039719
Offer third party coverage only	צד ג' בלבד	2/29/2020	3/1/2019	390343082519
Due to claims do not renew	לאור מצב תביעות לא ניתן לחדש באמצעותנו	2/29/2020	3/1/2019	390342820819

Notes: The figure reports an example of a "Go—No Go" grade document. The first column (from right) indicates the policy id number. The second column is the date at which the policy began. The third column is the end date of the policy. The fourth column is the "Go—No Go" grade. In the fifth row I provide a translation of a "Go—No Go" grade from Hebrew. The top three rows are policies that received a "Go" grade, while the bottom four rows are policies that received a "No-Go" grade.

Figure A.4: Summary statistics of premium and cost in nominal value by vehicle age

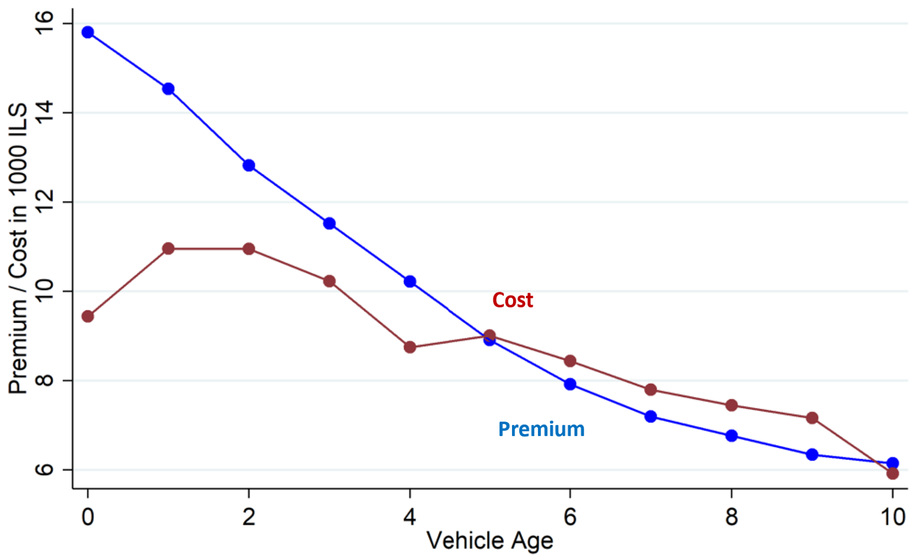

Notes: The figure depicts the premium (in blue) and cost (in red) in nominal values for comprehensive coverage policies for trucks from 2013 to 2020. The vertical axis depicts premiums and costs in 1000 New Israeli Shekel (ILS). No controls are added. Both variables are standardized to an annual term policy. Vehicle values are measured in New Israeli Shekel (ILS).

Figure A.5: Distribution of trucks by vehicle age

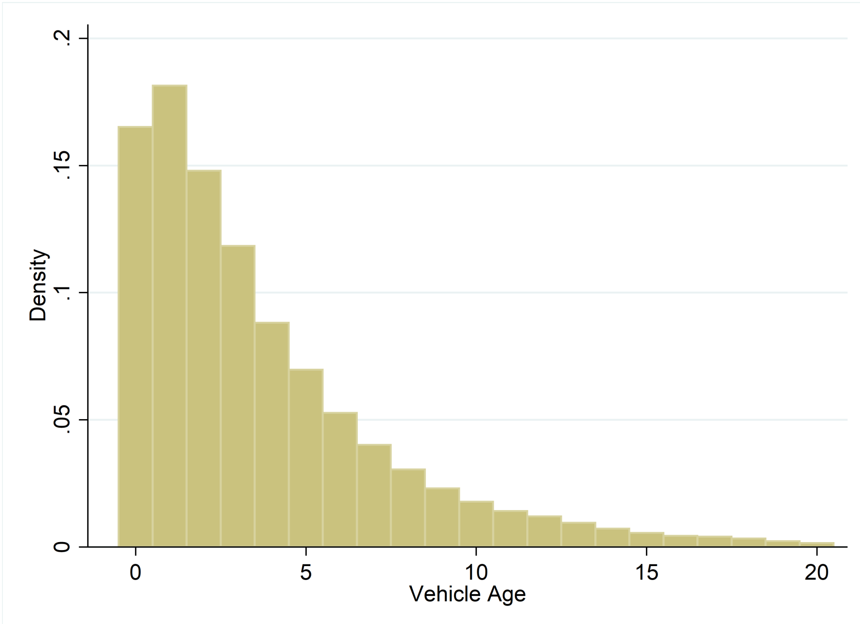

Notes: The figure depicts the distribution of comprehensive coverage policies for trucks by vehicle age. The sample includes insurance policies for trucks from 2013 through 2020. Vehicle age is measured in years.

Figure A.6: Alternative specifications of premium and costs by vehicle age

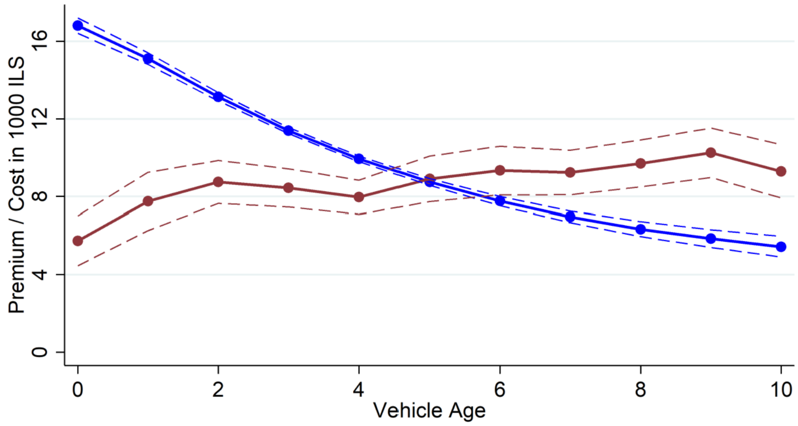

Panel A: Nominal Values

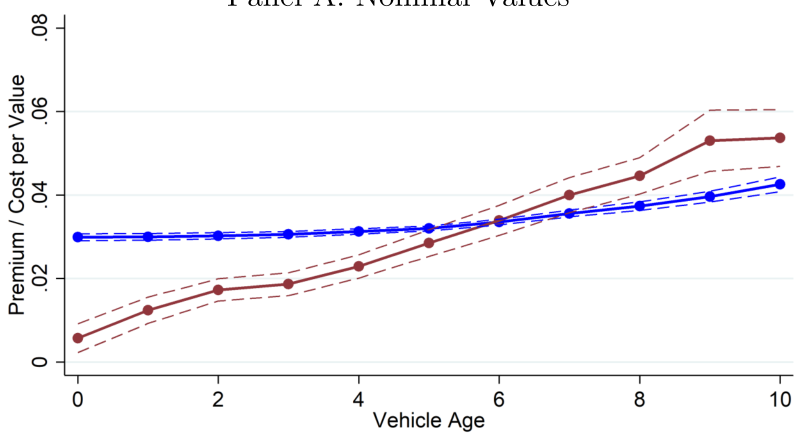

Panel B: All Vehicle Type

Notes: The figure depicts a variation of the analysis conducted presented in Figure 1.3. Panel A depicts premiums (in blue) and costs (in red) in nominal values, instead of normalized by vehicle value. The vertical axis is measured in 1,000 ILS. Panel B depicts a model identical to that of Figure 1.3, but includes all vehicles in the sample, rather than only trucks.

Figure A.7: Distribution of premium per value

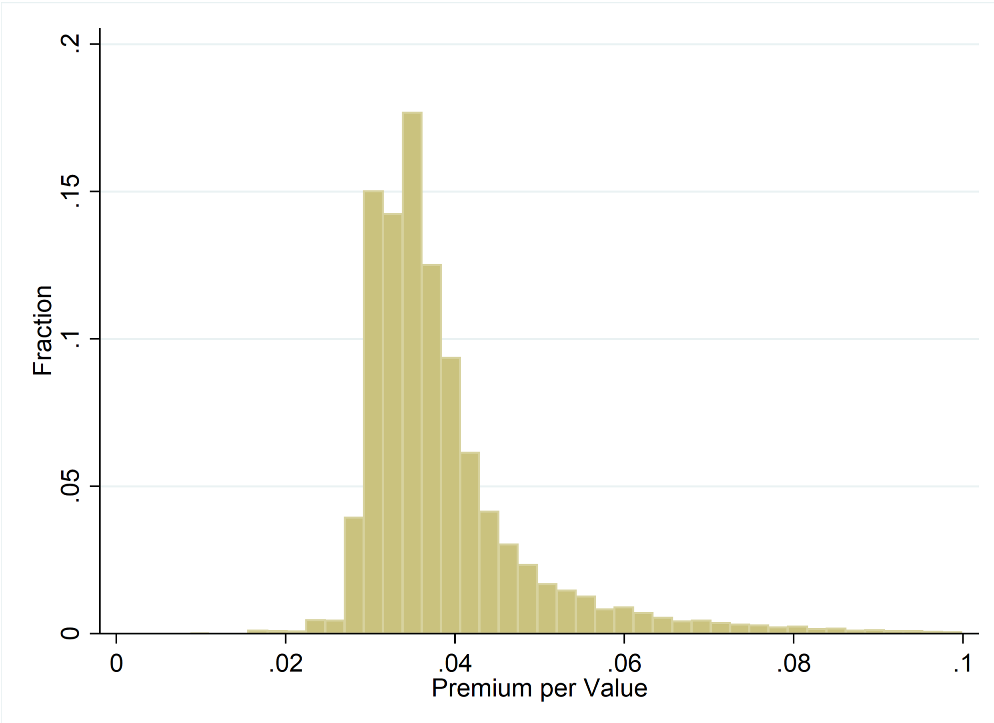

Notes: The figure depicts the distribution of premium per value paid for trucks with comprehensive coverage policies. The sample includes all trucks with comprehensive coverage from 2013 through 2020. Premiums, which are normalized to an annual policy length, and vehicle values are measured in New Israeli Shekel (ILS).

Figure A.8: Orlanet Calculator

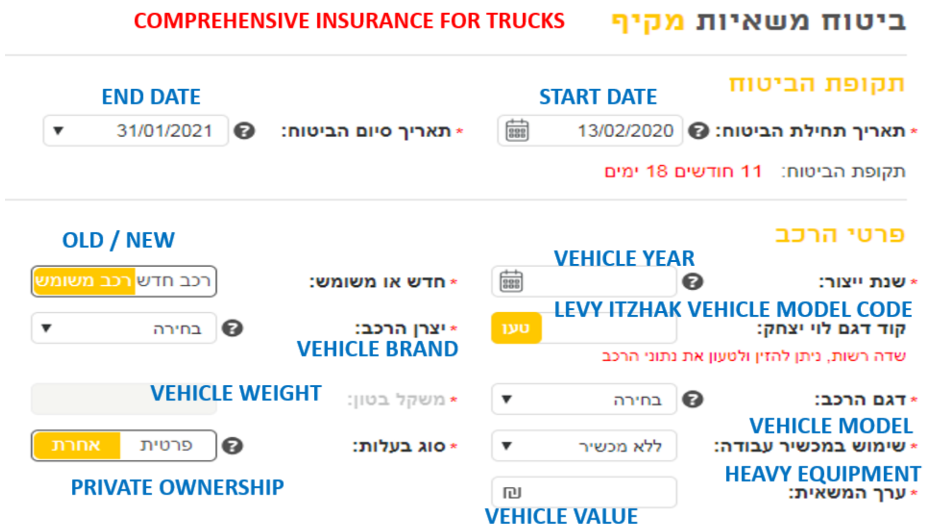

Panel A: Inputs

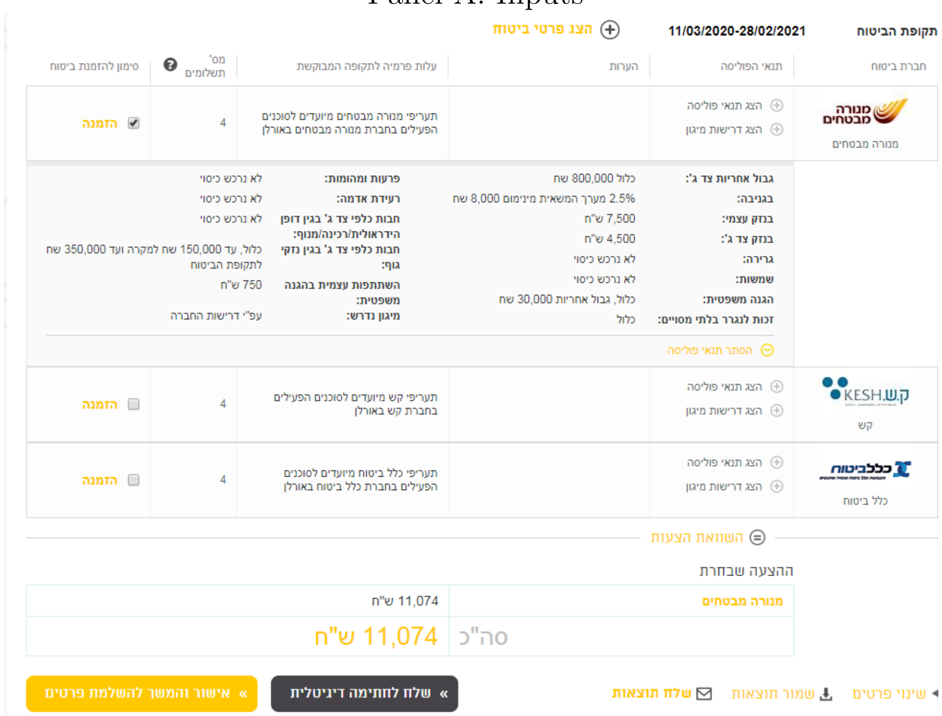

Panel B: Outputs

Notes: The figure depicts the process of generating fictitious comprehensive policy coverage for trucks using the *Orlanet Calculator*. Panel A describe the input process, and panel B illustrates the outputs.

Figure A.9: Orlanet Calculator

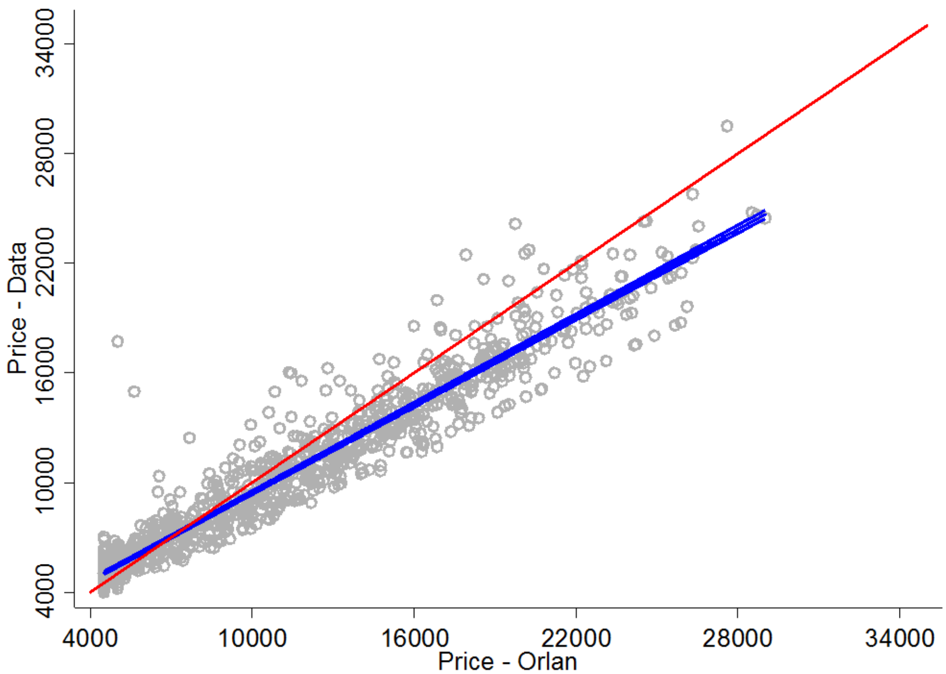

Panel A: Inputs

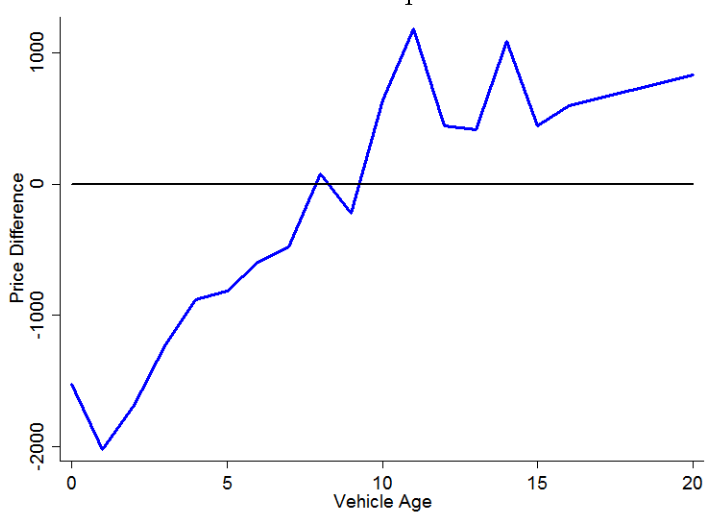

Panel B: Outputs

Notes: The figure in Panel A presents scatter and regression coefficients and 95% confidence interval of a within-insurer comparison in order to validate that the Orlanet Calculator pricing offers match the data provided by the insurer. A total of 2,041 observations are included. For each observation, I calculate the Orlan pricing using the average of both no claims in the last 3 years and one claim in the last 3 years, which occurred last year. The estimated slope equals 0.80 (0.01). R-square = 0.90. The red curve is the 45-degree line. Prices (premiums) are measured in New Israeli Shekel (ILS). The figure in Panel B depicts the mean difference between Orlan premiums and actual premiums by vehicle age.

Figure A.10: Market-wide premiums—Insurer and Rivals 1 and 2

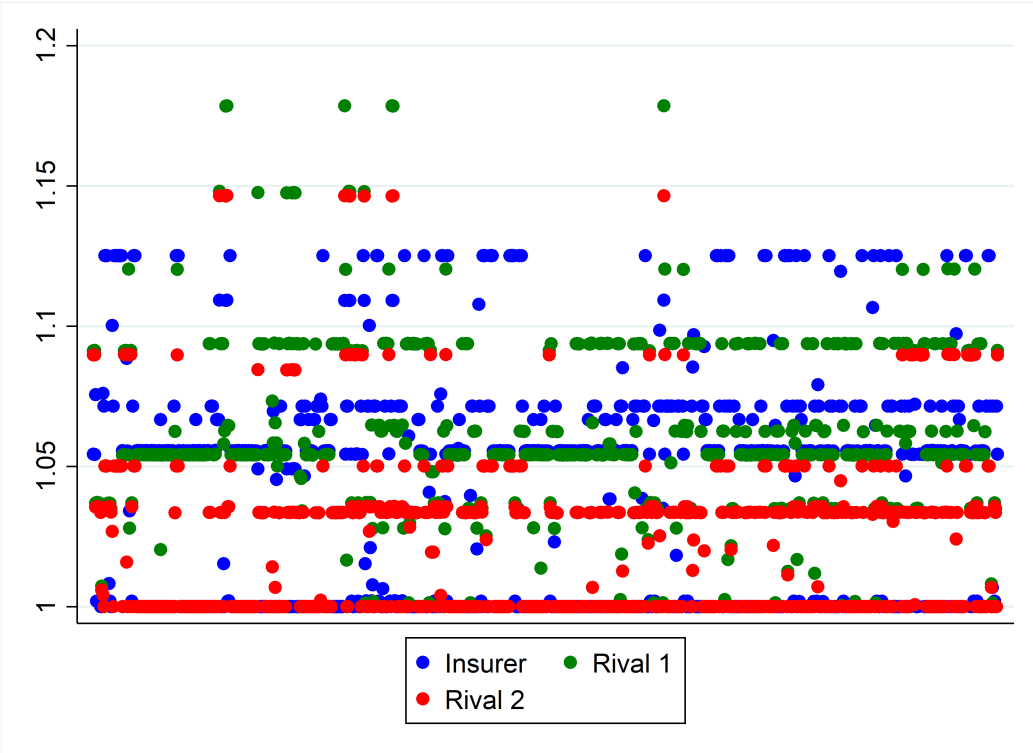

Notes: The figure reports the distribution of premiums offered by the insurer and its two main competitors. Premiums are calculated using the Orlanet Calculator. The horizontal axis depicts all 876 observations with distinct vehicle model-age value characteristics. I use the premium offered for the case of no claim in the last three years. The vertical axis depicts the premiums charged by each insurer, scaled by the lowest premium offered by the three competitors. The lowest premium offered is normalized to one. Prices (premiums) and vehicle values are measured in New Israeli Shekel (ILS).

Figure A.11: Out of sample prediction of costs

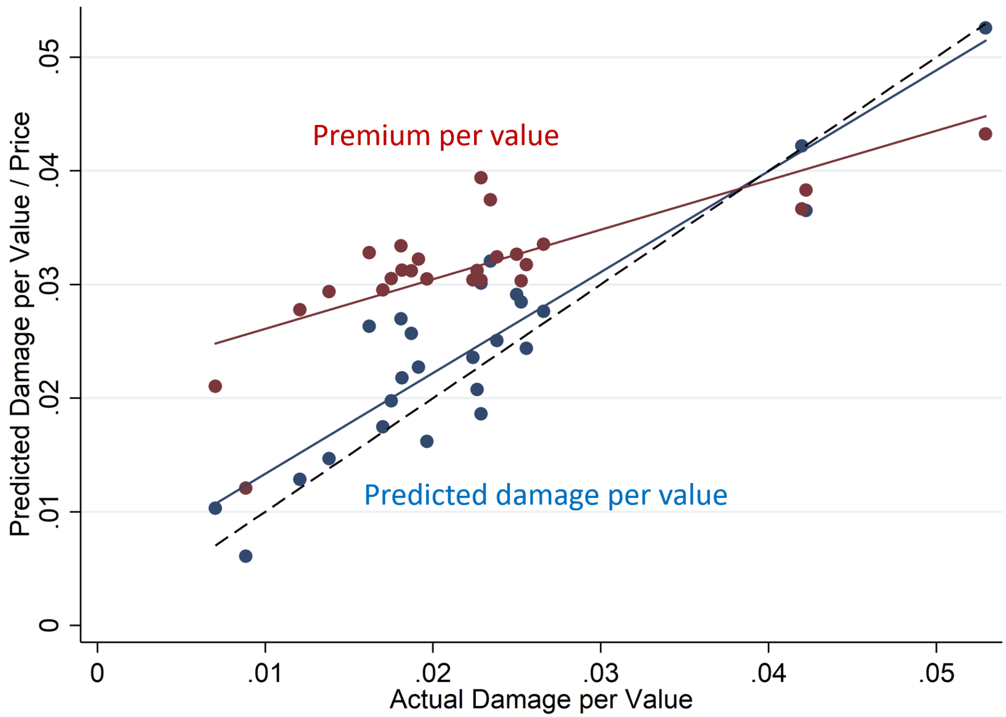

Notes: The figure reports the relationship between premium per value, predicted damage per value, and actual damage per value. Using data on policies from 2014 to 2018, I estimate a cost function (using regression analysis) by the following observable vehicle characteristics (age, value weight class, and type) and claim history (aggregate loss ratio). Based on the cost estimates, I divide my 2019–2020 sample into 25 groups based on projected damage per value. The vertical axis depicts predicted damage per value (in blue) and premium per value (in red). The horizontal axis depicts actual damage per value. The horizontal axis depicts the actual damage per value. The solid blue and red lines represent the regression coefficients of actual damage per value on predicted damage per value and premium per value, respectively. The dashed line is the 45-degree line. Premiums, damages, and vehicle values are measured in New Israeli Shekel (ILS).

Figure A.12: Distribution of conditional damage per value

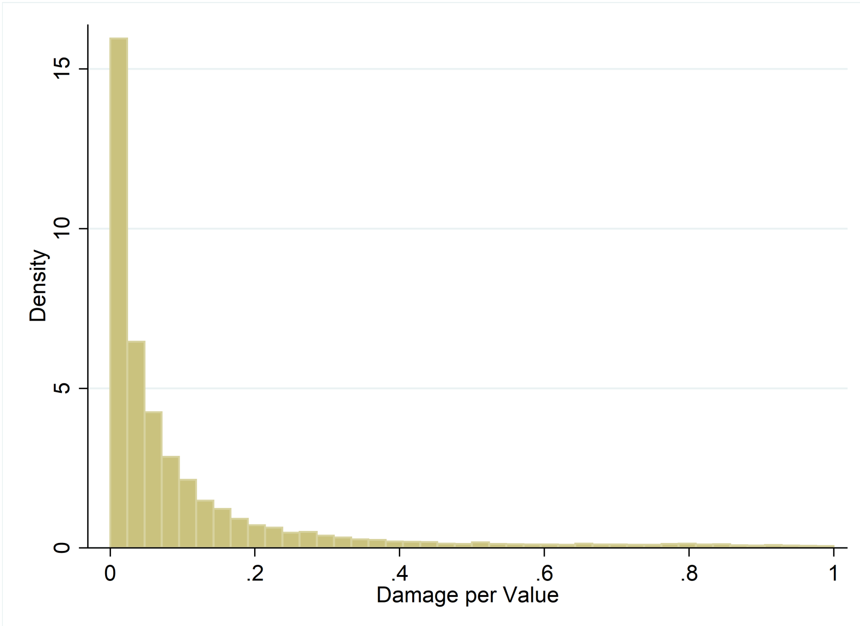

Notes: The figure reports the distribution of conditional damage per value; that is, damage per value if at least one claim occurred for comprehensive and partial coverage policies for all vehicles from 2013 to 2020. Damages and vehicle values are measured in New Israeli Shekel (ILS).

Figure A.13: Model fit

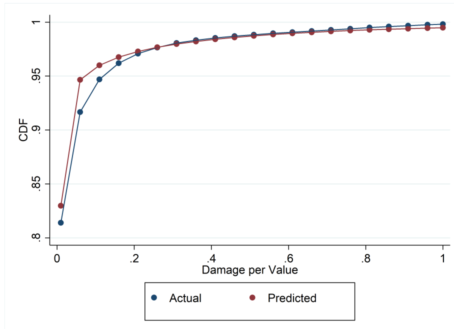

Panel A: Distribution of damage per value

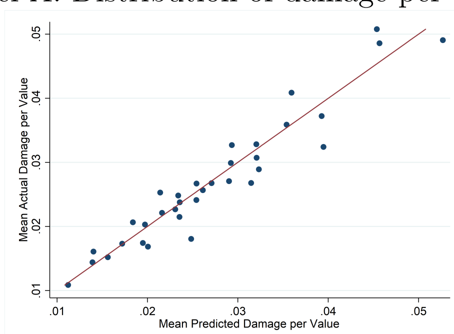

Panel B: Avg. damage per value by groups

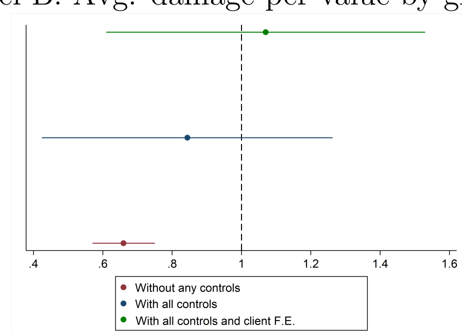

Panel C: Relationship between predicted and actual renewal

Notes: The figure reports the model fit. Panel A depicts the predicted and actual distribution of damage per value. Panel B depicts the relationship between predicted and actual mean damage per value. The sample is divided into groups based on the category variables vehicle age, aggregate client loss ratio, client's previous year loss ratio, and new client indicator. Figure presents group with at least 300 observations. Panel C depicts the relationship between the predicted probability of renewal and the share of realized renewal.

Figure A.14: Distribution of recommended increase in premium per value

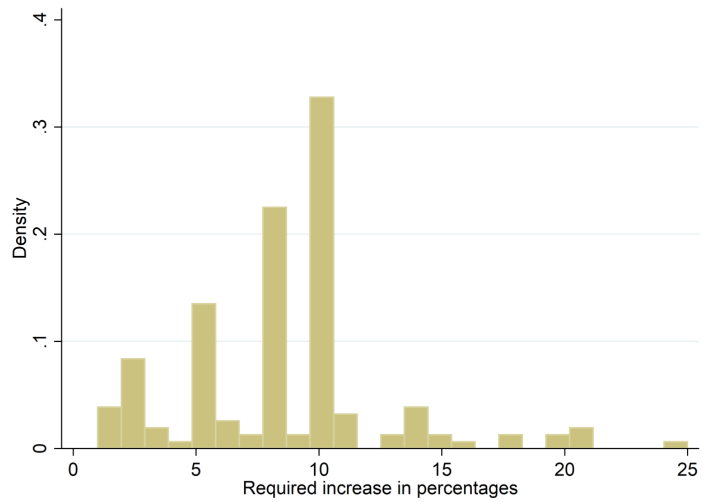

Notes: The figure reports the recommended increase in premium per value, for all policies for which the analytical team recommends a price increase. Premiums and values are measured in New Israeli Shekel (ILS).

Table A.1: Summary statistics of comprehensive coverage policies for all vehicles

	All	Vehicle Age$\geq$ 6	Claim$_{t-1}$ $\geq$1
Policies	102,372	33,378	12,347
Share (by Premium)	100%	32.60%	12.06%
Weighted Share (by Premium)	100%	19.13%	13.95%
Mean Premium	7,463	4,710	8,635
At least 1 claim	18.26%	16.55%	30.63%
Mean Damage	4,998	4,021	8,062
Mean Commission	998	668	1,027
Mean Profit	1,467	21	-1,084
Profit Margin	19.65%	0.45%	-5.26%
Mean Vehicle Age	4.78	10.22	5.08
Mean Vehicle Value	270,456	151,123	279,721
Mean Premium per Value	2.76%	3.12%	3.09%

Notes: The table is a replication of Table 2.1, consisting of comprehensive coverage policies for all vehicle types.

Table A.2: Policy outcomes and past performance - all vehicles

	Panel A: Entire Sample			
	(1)	(2)	(3)	(4)
	$Claim_t \geq 1$	$\frac{Damage_t}{Value_t}$	$\%\Delta \frac{Premium_t}{Value_t}$	$Loss\ Ratio_t$
$Claim_{t-1} \geq 1$	0.103***	0.013***	-0.003	0.350***
	(0.007)	(0.002)	(0.002)	(0.038)
log(Value)	Y	N	N	Y
Vehicle Age - 2^{nd} order	Y	Y	Y	Y
Vehicle Type	Y	Y	Y	Y
Vehicle Weight Class	Y	Y	Y	Y
Driver Underage Indicator	Y	Y	Y	Y
Observations	65,031	65,031	65,031	65,031
R-squared	0.051	0.009	0.023	0.003
	Panel B: Non-Fleet Policies			
	(1)	(2)	(3)	(4)
	$Claim_t \geq 1$	$\frac{Damage_t}{Value_t}$	$\%\Delta \frac{Premium_t}{Value_t}$	$Loss\ Ratio_t$
$Claim_{t-1} \geq 1$	0.104***	0.019***	0.009***	0.463***
	(0.011)	(0.004)	(0.004)	(0.094)
log(Value)	Y	N	N	Y
Vehicle Age - 2^{nd} order	Y	Y	Y	Y
Vehicle Type	Y	Y	Y	Y
Vehicle Weight Class	Y	Y	Y	Y
Driver Underage Indicator	Y	Y	Y	Y
Observations	12,715	12,715	12,715	12,715
R-squared	0.029	0.009	0.029	0.005

Notes: The table reports the results of the estimation model presented in Table 1.2, with regard to all vehicle types.

Table A.3: Summary statistics by customer classification

	Drop	Keep
Policies	74,456	1,102
Customers	3,170	166
Mean Premium	5,843	5,274
Mean Damage	3,658	5,252
Mean Commission	768	752
Mean Profit	1,417	-730
Profit Margin	24.26%	-13.83%
Loss Ratio	62.60%	99.58%

Notes: The table reports summary statistics for all policies from 2016 to 2020. The policies are classified into two groups. "Drop" includes policies of clients that incurred a loss ratio of at least 2 between 2013 and 2015 and their average vehicle age is at least 5. "Keep" includes policies of clients that incurred a loss ratio of at least 2 between 2013 and 2015 or their average vehicle age is below 5. Profit margin is defined as mean profit (premium-damage-commission) over mean premium. Loss ratio is defined as the mean damage of customers' claims (net of deductibles) over mean premiums. Premiums, commissions, damages, and profits are measured in New Israeli Shekel (ILS).

Table A.4: Summary statistics of policies of graded customers vs. rest

	All	Graded Customer	Not
Policies	109,630	55,868	53,762
Share (by Premium)	100%	50.96%	49.04%
Weighted Share (by Premium)	100%	56.27%	43.73%
Mean Premium	7,371	8,138	6,573
Mean Damage	4,973	4,882	5,067
Mean Commission	965	1,269	650
Mean Profit	1,432	1,988	855
Profit Margin	19.43%	24.42%	13.01%
Loss Ratio	67.47%	59.98%	77.10%
% Comprehensive	93.38%	95.20%	91.49%
Mean Vehicle Age	4.83	4.95	4.72
Mean Vehicle Value	275,594	268,580	282,883
Mean Premium per Value	2.67%	3.03%	2.32%

Notes: The table reports summary statistic to all comprehensive and partial coverage policies for all vehicle types between 2013 and 2020. The first column reports statistics for all policies, the second column describes the statistics for a sub-sample of the data consisting of all clients for whom at least one of their policies has a documented grade. Column 3 describes all other policies with regard to the other customers. Profit margin is defined as mean profit (premium-damage-commission) over mean premium. Loss ratio is defined as mean damage of customers' claims (net of deductibles) over mean premiums. Vehicle value, premium, commission, damages, and profits are measured in New Israeli Shekel (ILS). Vehicle age is measured in years. I exclude from the sample observation with an error, change in vehicle within the policy, change in coverage terms over the policy and policies that did not end, or that lasted for less than 30 days (without a claim).

Table A.5: Insurer grading and policy renewal

	(1) $\Delta\frac{\text{Premium}}{\text{Value}}$	(2) ΔDed. TP	(3) $\Delta\frac{\text{Ded. Own}}{\text{Value}}$	(4) Third Party	(5) Renew
Recommendation:					
$\Delta\frac{\text{Premium}}{\text{Value}}$	0.855**	-2.420	0.411*		
	(0.350)	(18.21)	(0.220)		
ΔDed. TP	-0.000	0.716***	0.001		
	(0.000)	(0.116)	(0.001)		
$\Delta\frac{\text{Ded. Own}}{\text{Value}}$	-0.002	0.243	1.085***		
	(0.002)	(0.259)	(0.008)		
Go	0.002	0.129	0.002		
	(0.001)	(0.151)	(0.002)		
Inc. Ded.	0.003**			-0.027*	0.031
	(0.001)			(0.016)	(0.046)
Inc. Prem.		0.082	-0.001	-0.016	-0.022
		(0.142)	(0.002)	(0.016)	(0.051)
TP Only				0.149***	0.006
				(0.049)	(0.042)
Deny				0.066	-0.670***
				(0.057)	(0.021)
V. Age Inc.	Y	Y	Y	Y	Y
Observations	8,652	8,652	8,652	10,080	14,282
R-squared	0.0630	0.1753	0.4894	0.1179	0.0822

Notes: The table reports the relationship between the analytical team's recommendation and the terms of renewal. The dependent variables are the change in premium per value (column 1), the change in deductible with regard to third-party property damage (column 2), the change in deductible with regard to own-property damage, normalized by vehicle value (column 3), an indicator as to whether the policy has been renewed with only third-party coverage (column 4), and whether the policy has been renewed at all (column 5). The explanatory variables, in order, are the recommended change in premium per value, third-party property damage deductible, and own property damage deductible, an indicator as to whether the policy received a "Go" grade, or whether the analytical team recommends increasing deductibles, premium, or to not renew the policy with comprehensive coverage, or at all. Robust standard errors are reported in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A.6: First stage prediction of log value based on previous year log value

Vehicle type	Coeff.	s.e.	R^2	Obs.
Truck	0.895	(0.001)	0.98	39,599
Heavy eq.	0.936	(0.002)	0.97	5,267
Trailer	0.923	(0.002)	0.90	17,058
Bus	0.898	(0.001)	0.99	7,378
Mini-bus	0.836	(0.002)	0.95	9,665
Heavy eq. add-on	0.929	(0.004)	0.92	4,270

Notes: The table reports the results of the first stage estimation: prediction of log value based on previous year log value. The vehicles' values are estimated separately for the five vehicle type groups. Furthermore, the values of heavy equipment add-ons are estimated separately. Vehicle values are measured in 100,000 New Israeli Shekels (ILS). History is measured in 1,000 years.

Table A.7: First stage prediction of log premium

IIC min price	0.496***	(0.056)	Claim Last Yr. (Vehicle)	-0.006***	(0.001)
Aggregate Loss Ratio:			Previous Yr. Loss Ratio:		
< 0.5	(omitted)		< 0.5	(omitted)	
≥ 0.5 and < 1	-0.003**	(0.002)	≥ 0.5 and < 1	0.002	(0.002)
≥ 1 and < 2	-0.004*	(0.002)	≥ 1 and < 2	0.021***	(0.003)
≥ 2	0.003	(0.004)	≥ 2	0.021***	(0.005)
Vehicle Age:					
0-1	(omitted)		log(Value)	0.222***	(0.013)
2-4	-0.017***	(0.001)	log(Total Value)	0.446***	(0.013)
5-7	-0.068***	(0.002)	Partial Coverage	-0.240***	(0.008)
8+	-0.065***	(0.003)	Underage driver	0.017***	(0.004)
			Joined last year	0.004**	(0.002)
History	-0.023	(0.002)	Joined 2 years ago	-0.004**	(0.001)
Observations	101,019				
R-squared	0.978				
Within R-squared	0.662				

Notes: The table reports the results of the first-stage estimation: prediction of log premium per value. The regression includes client and license fixed effects, in addition for year, vehicle type, and vehicle weight dummies. IIC min price is the international insurance company minimum price, based on observable characteristics. The sample excludes third-party coverage policies and 2013 policies. Premiums are measured in New Israeli Shekel (ILS) and vehicle value is measured in 100,000 New Israeli Shekels (ILS). History is measured in 1,000 years. *** p<0.01, ** p<0.05, * p<0.1.

A.2 Estimation Details

Estimation Details

I discuss the joint estimation of the parameters of the demand model: consumers' willingness to pay for policy renewal and expected damage per value of policy. Estimation is conducted in three steps. First, since I do not observe vehicle values for policies that are not renewed, I predict the (log) vehicle value as a function of previous period value and vehicle type, and treat it as data; both for those that renewed and those that did not. Then, I generate the predicted premium per value and predicted increase in premium per value using vehicle and client covariates and license fixed effect. Lastly, I jointly estimate Equations A.2.1 and A.2.2 via Maximum Simulated Likelihood. In this section, I describe in detail the Maximum Simulated Likelihood estimation process.

The parameter of interest are the expected damage per value and policy renewal equations.

$$d_{ijt} = \exp(\delta x_{ijt} + \log(\tau_{ijt}) + \nu_i) \tag{A.2.1}$$

$$\Pr(R_{ijt} = 1 | x_{ijt}, d_{ijt}, \hat{p}_{ijt}, \hat{\mathcal{I}}_{ijt}, \hat{f}_{ijt}, \omega_i) = \frac{\exp(-\alpha\hat{p}_{ijt} - \lambda\hat{\mathcal{I}}_{ijt} + \beta x_{ijt} + \gamma d_{ijt} + \rho\hat{f}_{ijt} + \omega_i)}{1 + \exp(-\alpha\hat{p}_{ijt} - \lambda\hat{\mathcal{I}}_{ijt} + \beta x_{ijt} + \gamma d_{ijt} + \rho\hat{f}_{ijt} + \omega_i)} \tag{A.2.2}$$

where ν_i represents customers' private information regarding risk, which is constant over time, is normally distributed, $\nu_i \sim N(0, \sigma_\nu^2)$. ω_i is an exogenous unobserved client-level demand component, which is constant over time; ω_i is normally distributed, $\omega_i \sim N(0, \sigma_\omega^2)$. ν_i and ω_i are uncorrelated and independent of all other covariates. τ_{ijt} is the duration of the policy, which take into account the duration of each policy by estimating a pro-rated variant of the damage per value equation.

I start by describing the damage per value likelihood. Damage per value follows a pseudo-Poisson distribution defined by parameter (which also equals the expected value) described in Equation A.2.1. This choice has two implications. On the one hand, it accommodates both (i) the possibility of no damages at all, which occurs quite frequently, and (ii) possible dependence between the number of claims and the conditional damage of claims, (iii) implementing pro-rated policies is quite simple.

The log likelihood of observing damage per value D_{ijt} is given by:

$$\begin{aligned}\log(d_{ijt} = D_{ijt} | x_{ijt}, \tau_{ijt}, \nu_i) = {} & D_{ijt}(\delta x_{ijt} + \ln(\tau_{ijt}) + \nu_i) \\ & - \exp(\delta x_{ijt} + \ln(\tau_{ijt}) + \nu_i) - \log(D_{ijt}!). \end{aligned} \tag{A.2.3}$$

While $d_{ijt}!$ is not well-defined when it is a continuous number, it does not possess a challenge in estimation as eliminating the element does not change the maximum function. In order to calculate the implied probabilities for each event, I approximate $d_{ijt}!$ for continuous values using the Gamma alternative of Stirling's formula, which fits small values especially well.

With regards to renewal, the log likelihood of observing a renewal or not, denoted by $\dagger_{ijt} = \{0, 1\}$ is given by Equation A.2.2 and can be re-expressed as:

$$\begin{aligned}
&\log Pr(R_{ijt} = \dagger_{ijt} | x_{ijt}, d_{ijt}, \hat{p}_{ijt}, \hat{\mathcal{I}}_{ijt}, \hat{f}_{ijt}, \omega_i, \nu_i) \\
&= \dagger_{ijt}\Big(- \alpha\hat{p}_{ijt} - \lambda\hat{\mathcal{I}}_{ijt} + \beta x_{ijt} + \gamma \exp(\delta x_{ijt} + \log(\tau_{ijt}) \\
&+ \nu_i) + \rho\hat{f}_{ijt} + \omega_i\Big) \\
&+ \log(1 + \exp(-\alpha\hat{p}_{ijt} - \lambda\hat{\mathcal{I}}_{ijt} + \beta x_{ijt} + \gamma \exp(\delta x_{ijt} + \log(\tau_{ijt}) + \nu_i) + \rho\hat{f}_{ijt} + \omega_i))
\end{aligned}$$

Therefore, the joint log likelihood of observing for client i's damage per value D_{ijt} and renewal or non-renewal, $\dagger_{ijt} = \{0, 1\}$ for each policy can be described as follows:

$$\begin{aligned}
L_i = \sum_{jt \in \mathcal{J}_i} \int\int & \log(d_{ijt} = D_{ijt} | x_{ijt}, \tau_{ijt}, \sigma_\nu \varepsilon_i^d) \\
& + \log Pr(R_{ijt} = \dagger_{ijt} | x_{ijt}, d_{ijt}, \hat{p}_{ijt}, \hat{\mathcal{I}}_{ijt}, \hat{f}_{ijt}, \sigma_\omega \varepsilon_i^r, \sigma_\nu \varepsilon_i^d)\phi(\varepsilon_i^d)\phi(\varepsilon_i^r) d\varepsilon_i^d d\varepsilon_i^r
\end{aligned}$$

where $\mathcal{J}_i$ includes all of the policies (vehicle-period) of client i, and ϕ denotes standard Normal distribution.

In practice, I approximate both integral using 100 Halton draws for each unobserved term. 100 Halton draws achieve greater accuracy in some setting (for instance, mixed logit estimations) than 1,000 pseudo-random draws (Train, 2000). I follow the procedure as described in Train (2000). In order to take into account policies that are not up for renewal, I set the log renewal likelihood to zero.

The estimated parameters are the set of $\theta = (\alpha, \lambda, \beta, \gamma, \rho, \delta, \sigma_{nu}, \sigma_{omega})$ that maximize the likelihood for the N clients:

$$\theta = \arg\max \sum_{i=1}^{N} L_i$$

Appendix B

Appendix of Price and Prejudice: Customer Taste-Based Discrimination and Competition

B.1 Proofs

Proof of Proposition 1

Proposition. *Let $F_\alpha(\cdot|Q,s)$ denote the cumulative distribution of preferences for $\mathcal{A}$ service among the $Q \in (0,1)$ consumers with the highest willingness to pay, given s. The infra-marginal customers' distribution of prejudicial preferences given s first-order stochastically dominates that of $s' < s$. I.e., $F_\alpha(\alpha|Q,s) \leq F_\alpha(\alpha|Q,s')$, $\forall \alpha \in (\alpha_{min}, \alpha_{max})$, $s' < s \in (0,1]$, $Q \in (0,1)$, and strict inequality for some $\alpha \in (\alpha_{min}, \alpha_{max})$.*

Proof. Let $\Delta\theta_i = \theta_i(s) - \theta_i(s') = \alpha_i(s - s')$ denote customer i's increase in willingness to pay when facing a higher share of White employees, $s > s' \in [0,1)$. Let a and b denote two arbitrary clients that their ordering changed as the share of $\mathcal{A}$ employees increases. That is $\theta_a(s') < \theta_b(s')$, and $\theta_a(s) > \theta_b(s)$. Thus, $\Delta\theta_a > \Delta\theta_b$, which imply that $a_i > b_i$. □

Proof of Proposition 2

Proposition. $\ni q_b \in (0,1)$ *such that* $\forall s' > s$, $P_{s'}(Q) < P_s(Q) + \min\{(s'-s)[\alpha], 0\}$ *,* $\forall Q > q_b$ *and* $\ni q_t \in (0, q_b]$ *such that* $\forall s' > s$, $P_{s'}(Q) > P_s(Q) + \max\{(s' - s)[\alpha], 0\}, \forall Q < q_t$.

Proof. Let $F_\theta(\cdot|s)$ denote the distribution of willingness to pay, given s. Consider arbitrary shares of white employment $\{s_0, s_1\} \in [0,1]$. Without loss of generality, $s_1 > s_0$. By definition of $\theta(s)$, $\Delta B \equiv [\theta(s_1)] - [\theta(s_0)] = (s_1 - s_0)$. Let $F_{\theta+\Delta B}(\cdot|s_0)$ denote the distribution of willingness to pay, given s_0, plus the mean difference in benefit from higher white service, $\Delta B = (s_1 - s_0)[\alpha]$. Both $F_{\theta+\Delta B}(\cdot|s_0)$ and $F_\theta(\cdot|s_1)$ have the same mean willingness to pay. Yet, due to the heterogeneity in prejudicial preferences, $F_\theta(\cdot|s_1)$ is a mean preserving spread of (and hence second-order stochastically dominated by) $F_{\theta+\Delta B}(\cdot|s_0)$. Define $T_{d|s}(P) = \int_0^P F_d(p)dp$ as the integral of the cumulative distribution function of d, given s. By second-order stochastic dominance, $T_{\theta|s_1}(p) > T_{\theta+\Delta B|s_0}(p)$, $\forall p \in (\nu_{min}+s_1\alpha_{min}, \nu_{max}+s_1\alpha_{max})$. By definition of $T_{d|s}(\cdot)$, $T_{\theta|s_1}(\nu_{min} + s_1\alpha_{min}) = T_{\theta+\Delta B|s_0}(\nu_{min} + s_1\alpha_{min}) = 0$. Therefore, $\ni p_l \in (0,1)$ such that $F_\theta(p|s_1) > F_{\theta+\Delta B}(p|s_0)$, $\forall p < p_l$. Hence, $\forall p < p_l$, $\ni p'$ such that $F_\theta(p'|s_1) = F_{\theta+\Delta B}(p|s_0)$. Since $F_\theta(\cdot|s)$ is strictly increasing, p' is unique and $p' < p$.

Moreover, by definition of $F_{\theta+\Delta B}(\cdot|s_0)$ and $T_{d|s}(P)$, $T_{\theta+\Delta B|s_0}(\nu_{max} + s_1\alpha_{max}) = \nu_{max}+s_1\alpha_{max}-[\nu+s_0\alpha+[(s_1-s_0)\alpha]] = \nu_{max}+s_1\alpha_{max}-[\nu+s_1\alpha] = T_{\theta|s_1}(\nu_{max}+s_1\alpha_{max})$. Thus, $\ni p_t \in [q_b, 1)$ such that $F_\theta(p|s_1) < F_{\theta+\Delta B}(p|s_0)$, $\forall p > p_t$. This implies that $\forall p > p_t$, $\ni p'$ such that $F_\theta(p'|s_1) = F_{\theta+\Delta B}(p|s_0)$. Since $F_\theta(\cdot|s)$ is strictly increasing, p' is unique and $p' > p$.

Furthermore, since $\alpha_{max} > 0$ and $\alpha_{min} < 0$, $\ni \{\tilde{p}_t, \tilde{p}_l\}$ such that $\max\{\alpha_i : \theta_i(s_1) \leq \tilde{p}_l\} < 0, \forall p \leq \tilde{p}_l$ and $\min\{\alpha_i : \theta_i(s_1) \leq \tilde{p}_t\} > 0, \forall p \geq \tilde{p}_t$. By demand resorting and continuity of preferences, $\forall p \leq \tilde{p}_l, \exists! p' < p$ such that $F_\theta(p'|s_1) = F_\theta(p|s_0))$, and $\forall p \geq \tilde{p}_t, \exists! p' > p$ such that $F_\theta(p'|s_1) = F_\theta(p|s_0))$. □

Proof of Proposition 3

Proposition. *Let r^m and r^c denote the wage gap under monopolistic and competitive market structures. Competition in the product market decreases the sector's wage gap by more when the customer's prejudice does not play a role;* $\frac{1}{1+d} < \frac{r^*}{r^m}$.

Proof. By way of contradiction, suppose $r^m \geq r^*(1+d)$. Since the employer exhibits a cost when increasing the share of white employees, it must be the case that the monopoly's demand for white employees for a given wage gap $\frac{\tilde{\omega}^w}{\tilde{\omega}^b} = r^*(1+d)$ is larger than that of a competitive firm, facing a wage gap of $\frac{\omega^w}{\omega^b} = r^*$.[1] By equations 2.13

[1] I denote the equilibrium wage gap in the case of a monopoly by $\tilde{\omega}^w$ and $\tilde{\omega}^b$ to take into account the fact that the scale of the wages might differ than that of a competitive market structure.

and 2.16, this implies that

$$\frac{(1-\eta)\tilde{M}'(s^*|r^*(1+d)\tilde{\omega}^b,\tilde{\omega}^b,d)}{\tilde{M}(s^*|r^*(1+d)\tilde{\omega}^b,\tilde{\omega}^b,d)} \leq \frac{M(s^*|r^*\omega^b,\omega^b)}{M'(s^*|r^*\omega^b,\omega^b)}. \tag{B.1.1}$$

By definition of $\tilde{M}$ (equation 2.17):

$$\tilde{M}(s|r^*(1+d)\tilde{\omega}^b,\tilde{\omega}^b,d) = \frac{\delta\tilde{\omega}^b(1+d)(1+s\frac{(r^*(1+d)\tilde{\omega}^b-\omega^b(1+d))}{\omega^b(1+d)})f^{1-\sigma(s)}}{(1-\sigma(s))^{1-\sigma(s)}\sigma(s)^{\sigma(s)}A(s)},$$

which implies that $\frac{\tilde{M}'(s^*|r^*(1+d)\tilde{\omega}^b,\tilde{\omega}^b,d)}{\tilde{M}(s^*|r^*(1+d)\tilde{\omega}^b,\tilde{\omega}^b,d)} \leq \frac{M(s^*|r^*\omega^b,\omega^b)}{M'(s^*|r^*\omega^b,\omega^b)}$. Since $\in (0,1)$, equation B.1.1 does not hold. Contradiction. □

Appendix C

Appendix of The Lemons Gap: Demand For Insurance of Quality Uncertain Goods

C.1 Proofs

Proof of Proposition 4

Proposition. *In equilibrium, all owners with vehicle quality lower than the market's expected quality, $q_{im} \leq \mu_{mt}(p^e_{mt})$, sell their vehicle.*

Proof. In equilibrium, by equation 3.1, owners do not sell vehicle iff $q_{im} - \frac{\nu_i}{\omega_t} > p_{mt}$. By way of contradiction, suppose owner j with vehicle quality $q_{jm} \leq \mu_{mt}(p^e_{mt})$ is unwilling to sell at equilibrium price p^e_t. Thus, $q_{jm} \geq q_{jm} - \frac{\nu_j}{\omega_t} > p^e_{mt}$. With regards to the demand side, by definition, the marginal buyer is indifferent between purchasing the vehicle, or not. Equation 3.2 implies that $\gamma\mu_{mt}(p^e_t) = p^e_{mt}$. Since markets clear in equilibrium and $\gamma > 1$, $\mu_{mt}(p^e_{mt}) < p^e_{mt}$, which imply that $q_{jm} < p^e_{mt}$. Contradiction. □

Proof of Proposition 5

Proposition. *Let $F_{mt}(\cdot|Q)$ denote the vehicle model m's quality distribution of the Q vehicles that are supplied at the pre-owned market age t. The ordering of supply improves over the vehicle life cycle such that, conditional on Q, the quality distribution of supplied vehicles at age $t+1$ first order stochastically dominates that of* *age t,* $F_{mt+1}(\cdot|Q) \leq F_{mt}(\cdot|Q), Q \in [0,1]$ *(and strict inequality for some values when* $Q \in (0,1)$*).*

Proof. $\Delta u_i \equiv u_{imt+1} - u_{imt} = \nu_i\left(\frac{1}{\omega_t} - \frac{1}{\omega_{t+1}}\right) < 0$ denotes owner's i change in utility between age t and $t+1$. Let a, b be 2 owners such that at age t owner b has a lower reservation price $u_{amt} \geq u_{bmt}$ and at age $t+1$ owner a has a lower reservation price $u_{amt+1} < u_{bmt+1}$. Thus, $\Delta u_b - \Delta u_a = (\nu_b - \nu_a)\left(\frac{1}{\omega_t} - \frac{1}{\omega_{t+1}}\right) > 0$. Since $\omega_{t+1} < \omega_t$, $\nu_b < \nu_a$. Furthermore, by $u_{amt} \geq u_{bmt}$, $q_a - \frac{\nu_a}{\omega_t} \geq q_b - \frac{\nu_b}{\omega_t} \Rightarrow q_a - q_b \geq (\nu_a - \nu_b)\frac{1}{\omega_t}$. $\nu_a > \nu_b$ implies $q_a > q_b$. □

Proof of Proposition 9

Proposition. *If the insured asset is perfectly replaceable (*$L_t = 0$*), the ordering of the demand for insurance is static. In contrast, if the insuree is partially compensated via a second-hand market product (*$L_t > L_{t+1} > 0$*), the market becomes more adversely selected over the vehicle life cycle.*

Proof. $\Delta\theta_i \equiv \theta_{it+1} - \theta_{it} = \log(1 + r(L_{t+1} + \bar{d})) - \log(1 + r(L_t + \bar{d})) < 0$ denotes insuree's i change in (log) willingness to pay for the comprehensive coverage plan between age t and $t+1$. Let a and b denote two arbitrary clients such that at age t owner a has a lower willingness to pay $\theta_{bt} \geq \theta_{at}$ and at age $t+1$ owner b has a lower willingness to pay $\theta_{at+1} > \theta_{bt+1}$. Thus, $\Delta\theta_b < \Delta\theta_a \Rightarrow \frac{1+r_b(L_{t+1}+\bar{d})}{1+r_b(L_t+\bar{d})} < \frac{1+r_a(L_{t+1}+\bar{d})}{1+r_a(L_t+\bar{d})} \Rightarrow r_a(L_t - L_{t+1}) < r_b(L_t - L_{t+1})$. Since $L_{t+1} < L_t$, $r_a < r_b$. Furthermore, by equation 3.8, $\theta_{at+1} > \theta_{bt+1}$ implies that $\frac{\lambda_b}{\lambda_a} < \frac{1+r_a(L_{t+1}+\bar{d})}{1+r_b(L_{t+1}+\bar{d})}$. $r_b > r_a$ implies $\lambda_a > \lambda_b$. □