NBER WORKING PAPER SERIES

BUYING DATA FROM CONSUMERS: THE IMPACT OF MONITORING PROGRAMS IN U.S. AUTO INSURANCE

Yizhou Jin
Shoshana Vasserman

Working Paper 29096
http://www.nber.org/papers/w29096

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2021

We thank our advisors Ariel Pakes, Nathan Hendren, Robin Lee, Dennis Yao, Leemore Dafny, and Elie Tamer; our data providers Quadrant Information Services and an unnamed auto insurer; Harvard and the Geneva Association for financial support; Jie Bai, Liran Einav, Ashvin Gandhi, Nir Hak, Ben Handel, Oliver Hart, Kevin He, Panle Barwick, Ginger Jin, Myrto Kalouptsidi, Scott Kominers, Jonathan Kolstad, Jing Li, Alex MacKay, James Savage, Steve Tadelis, Andrew Sweeting, Chad Syverson, John Wells, Thomas Wollmann, and various seminar participants for valuable comments. Ability to publish is not contingent on results (data usage agreement contact carolina_harvey@harvard.edu). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Buying Data from Consumers: The Impact of Monitoring Programs in U.S. Auto Insurance
Yizhou Jin and Shoshana Vasserman
NBER Working Paper No. 29096
July 2021
JEL No. L0

ABSTRACT

New technologies have enabled firms to elicit granular behavioral data from consumers in exchange for lower prices and better experiences. This data can mitigate asymmetric information and moral hazard, but it may also increase firms' market power if kept proprietary. We study a voluntary monitoring program by a major U.S. auto insurer, in which drivers accept short-term tracking in exchange for potential discounts on future premiums. Using a proprietary dataset matched with competitor price menus, we document that safer drivers self-select into monitoring, and those who opt in become yet 30% safer while monitored. Using an equilibrium model of consumer choice and firm pricing for insurance and monitoring, we find that the monitoring program generates large profit and welfare gains. However, large demand frictions hurt monitoring adoption, forcing the firm to offer large discounts to induce opt-in while preventing the unmonitored pool from unraveling given the competitive environment. A counterfactual policy requiring the firm to make monitoring data public would thus further reduce the firm's incentive to elicit monitoring data, leading to less monitoring and lower consumer welfare in equilibrium.

Yizhou Jin
2200 Piedmont Ave, F555
Haas School of Business
Berkeley, CA 94608
jyz@berkeley.edu

Shoshana Vasserman
Stanford Graduate School of Business
655 Knight Way
Stanford, CA 94305
and NBER
svass@stanford.edu

Additional appendices are available at http://www.nber.org/data-appendix/w29096

New technologies have enabled individuals to easily and credibly document their behavior. Firms can elicit such data to offer tailored products and personalized pricing, which may in turn benefit consumers and give rise to voluntary data sharing.

The potential benefits of behavioral data are especially compelling in the auto insurance industry. Granular driving data allows insurers to identify drivers with lower accident risk more precisely than their current algorithms, which pool consumers based on demographic features and sparse accident records alone (Cohen and Einav 2007). This benefits both the insurer, who can more easily target safe drivers, and drivers, who may receive price discounts by demonstrating safe driving behavior.

However, consumers must volunteer to share their data in order for these benefits to manifest. Even if a consumer qualifies for safe-driving discounts, she may opt not to share her data due to risk aversion or disutility from being monitored (Handel, Hendel, and Whinston 2015). Furthermore, firms typically retain property rights over the data that they collect. This incentivizes more data collection, but it also generates a competitive advantage that may lead to higher markups in the future (Jin and Wagman 2021).

In this paper, we develop an empirical framework to examine the trade-offs for consumer welfare and firm profit that result from consumer data sharing in the context of an auto-insurance monitoring program ("pay-how-you-drive") in the United States. New customers are invited to plug a simple device into their cars, which tracks and reports their driving behavior for up to six months (Figure A.1). In exchange, the insurer uses the data to better assess accident risk and adjust future premiums. Unlike traditional pricing factors such as age or accident records, monitoring data is not available to other firms. In 2017, insurers serving over 60% of the $267 billion U.S. auto insurance industry offered monitoring options.[1] Similar programs have been introduced in other industries, such as life insurance and consumer lending (Figure A.2).[2] However, despite their growing relevance, empirical evidence on the effects of consumer data sharing mechanisms and their impact on welfare across different sides of the market is sparse.

We construct a novel dataset that combines proprietary individual-level data from a major U.S. auto insurer (hereinafter referred to as "the firm") and price menus offered by its competitors. Our panel covers 22 states and spans from 2012 to 2016, during which the firm's

[1] 2017 annual report of the National Association of Insurance Commissioners.

[2] The Vitality program from life insurer John Hancock tracks and rewards exercise and health-related behaviors. Ant Financial incentivizes users to conduct more personal finance transactions in exchange for borrowing discounts.

monitoring program—the first major program in the industry—was introduced. For each consumer in our panel, we observe demographic characteristics, price menus from top competing insurers, insurance contracts purchased, and realized insurance claims. For each consumer who opts into the monitoring program, we observe a monitoring score reflecting her safe driving performance and the corresponding premium adjustments. Taken together, our analysis uses a panel dataset of over 1 million consumers and 50 million insurance quotes.

We first establish two key facts about the monitoring program as it is observed in our data. First, the monitoring program induces safer driving—a moral hazard effect. Monitoring only occurs in the first period of insurance for new customers who opt in. Using a difference-in-differences estimator to capture within-consumer across-period variation in insurance claims, we find that consumers who opt in to monitoring become 30% safer, on average, while they are monitored. However, this incentive effect only explains 64% of the risk differences between consumers in the monitored and unmonitored groups, and monitoring scores remain highly predictive of risk after the monitoring device is removed. This suggests that monitoring reveals persistent risk differences across consumers that were previously unknown to the firm, leading to advantageous selection into the program.

These reduced-form analyses allow us to separately identify the effects of moral hazard and asymmetric information, extending a rich empirical literature that has thus far relied on correlations between plan choice and claims (Chiappori and Salanie 2000; Cohen and Einav 2007; Jeziorski, Krasnokutskaya, and Ceccarini 2019). However, in order to quantify the impact of the monitoring program—or that of prospective data regulations on consumer welfare—we need to assess how consumer risk and preferences contribute to insurance choices and to the degree of asymmetric information in equilibrium (Einav, Finkelstein, and Levin 2010). To do this, we develop an equilibrium model of insurance and monitoring. On the demand side, we estimate a structural model that captures complex correlations between consumers' insurance-plan choice, their monitoring opt-in decision, as well as the cost needed to insure them. Consumers have heterogeneous private risk types and preference factors—risk aversion, inertia costs, and disutility from being monitored—and choose an insurance plan from a personalized price menu with options offered by the focal firm and its competitors. When consumers engage with the firm for the first time, they can opt in to monitoring and receive an upfront discount. At the end of each period, accidents arrive based on consumers' risk types. In addition, if a consumer is monitored, a score is realized from a noisy distribution that is correlated with her risk type. In expectation, better scores lead to higher discounts in future periods. Consumers are thus incentivized to drive more safely when monitored.

On the supply side, we model firm pricing before and after the monitoring period to reflect the trade-off between eliciting more consumer data and profiting from these data in a competitive market. By offering a high upfront discount, the firm can encourage more consumers to opt in to monitoring. In order for this to be profitable, however, the firm must extract some of the surplus generated from monitoring in renewal periods. In this sense, the amount of monitoring data that is created (due to equilibrium monitoring participation) is endogenous to the firm's dynamic pricing strategy. In addition, the firm's ability to profitably "invest-and-harvest" through the monitoring program is tempered by competition as it must attract and retain consumers in the first place.

Taken together, our model captures several forces that drive consumers' monitoring opt-in choice. First, consumers anticipate that their risk will be lower during the monitoring period if they opt in. Some may also expect higher future discounts and thus stand to gain more from monitoring. But the expected benefit of such future discounts is moderated by elevated reclassification risk due to noise in the monitoring process. Finally, consumers may incur various unobserved costs due to monitoring, including privacy loss, increased effort while driving, and additional decision-making. We quantify the combined effect of these costs with a heterogeneous monitoring disutility term.

Our estimates suggest that the average consumer suffers a $93 disutility from being monitored. This contributes to the low opt-in rate observed in the data, but it simultaneously prevents unraveling in the unmonitored consumer pool *a la* Milgrom (1981). Competition further prevents the firm from surcharging unmonitored consumers to induce unraveling. Nonetheless, we find that monitoring disutility is lower for safer drivers, enhancing advantageous selection beyond what is implied by financial risk and rewards alone. Meanwhile, the average consumer forgoes $284 in financial gain annually by not exploiting outside options from competitors. This suggests that the market may remain imperfectly competitive even with perfect information on consumer risk. Finally, consumers are only moderately risk-averse. The precision of the monitoring score thus has little effect on monitoring demand.

Overall, compared to a counterfactual scenario with no monitoring, both consumers and the firm benefit from the monitoring program as it is: total annual surplus increases by $13.3 (1.7% of premium), 64% of which is due to risk reduction during monitoring. But monitoring take-up is low—due, in part, to demand frictions, and to competition from cheap plans offered by other insurers. As a result, enforcing a counterfactual ban on proprietary data would strongly diminish the firm's incentive to elicit the data in the first place.[3] This would reduce the amount of information that is revealed in equilibrium, hurting both firm profits and consumer welfare.

The paper proceeds as follows. Section 1 describes our data and provides background information on auto insurance and monitoring. Section 2 conducts reduced-form tests to evaluate the effects of moral hazard and selection that are induced by the monitoring program. Section 3 presents our structural model, identification arguments, and estimation procedures to recover key demand and cost parameters. Section 4 discusses estimation results and counterfactual simulations for welfare analyses. Section 5 proposes a model of monitoring pricing and studies the equilibrium implications for optimal pricing and a ban on proprietary data. Section 6 revisits our results in relations to the literature and concludes.

1 Background and Data

In this section, we provide background information on U.S. auto insurance and the monitoring program we study. We also describe our datasets.

1.1 Auto Insurance

Auto insurers in the U.S. collected $267 billion dollars of premiums in 2017.[4] There are two main categories of insurance: liability and property. Property insurance covers damage to one's own car in an accident, regardless of fault. Liability insurance covers injury and property liability associated with an at-fault accident. In all states we study, liability insurance is mandatory, with the minimum required coverage ranging from $25,000 to $100,000.[5]

Pricing Insurance prices are heavily regulated. Firms collect large amounts of consumer information in risk-rating, and are required to publish filings that detail their pricing rules.[6]

[3] This is similar to enforcing data portability across firms, such as Article 20 of the General Data Protection Regulation (De Hert, Papakonstantinou, Malgieri, Beslay, and Sanchez 2018).

[4] Calculated as premiums from property annual statements plus state funds by the National Association of Insurance Commissioners.

[5] All states that we study follow an "at-fault" tort system and mandate liability insurance. In reality, liability insurance is specified by three coverage limits. For example, 20/40/10 means that, in an accident, the insurer covers liability for bodily injuries up to $40,000 overall, but no more than $20,000 per victim; it also covers liability for property damage (cars or other infrastructure) for up to $10,000. We quote the highest number here.

[6] Except in Wyoming, which is not in our dataset.

In general, a pricing rule can be summarized as follows, where the price (p) of a policy is:[7]

$$\begin{aligned} p = \text{base rate} \times \text{driver factor} \times \text{vehicle factor} \times \text{location factor} \\ \times \text{tier factor} \times \text{coverage factor} + \text{loading factor} \end{aligned} \tag{1}$$

Within each firm, price varies by observable characteristics, coverage choice, and time. Base rates vary only by state and time. Driver, vehicle, and location factors include age, vehicle model, and zipcode-level population density, etc. This information is verified and cross-referenced among various public and industry databases. Tier factors incorporate information from claim and credit databases, which include accident, traffic violation (DUI, speeding, etc.), and financial (delinquency, bankruptcy, etc.) records.[8] In addition, choosing a higher coverage scales prices by a positive factor. Finally, firms charge a loading factor that includes markups and overhead for operational and marketing expenditures. Pricing regulations varies by state and time, but the primary focus of regulation is to deter third-degree price discrimination: excessively high prices that hurt affordability, and excessively low prices that raise insurers' default risk. In general, "price optimization" across consumer segments—beyond risk-rating through the various factors in Equation 1, which regulators can verify with historical claims data and reserving assumptions—is generally not allowed, and is explicitly outlawed in 15 states.

Timing and Dynamics For each consumer segment, insurers determine and publicly file *two* prices (Figure 1a): an initial price and a renewal price. New customers must report characteristics at time $t = 0$. This facilitates risk rating, based on which the firm generates personalized initial price menu. Consumers make coverage choices based on the menu offered to them, or leave for another firm. There is no long-term commitment in U.S. auto insurance. Each period at the firm lasts for six months, at the end of which consumers decide to stay or leave given the renewal price menu provided at the end of month five. If an auto accident occurs (Figure 1b), the insured files a claim immediately and, upon evaluation and adjustment by the insurer, gets reimbursed and pays out-of-pocket accordingly. Meanwhile, the claim is recorded in industry databases, leading the consumer to face a claim surcharge at renewal or higher prices when switching to other firms.

[7]We focus on single-driver-single-vehicle policies and liability coverage only. See Appendix H for more details.

[8]See Appendix H, Figures H.7 and H.8

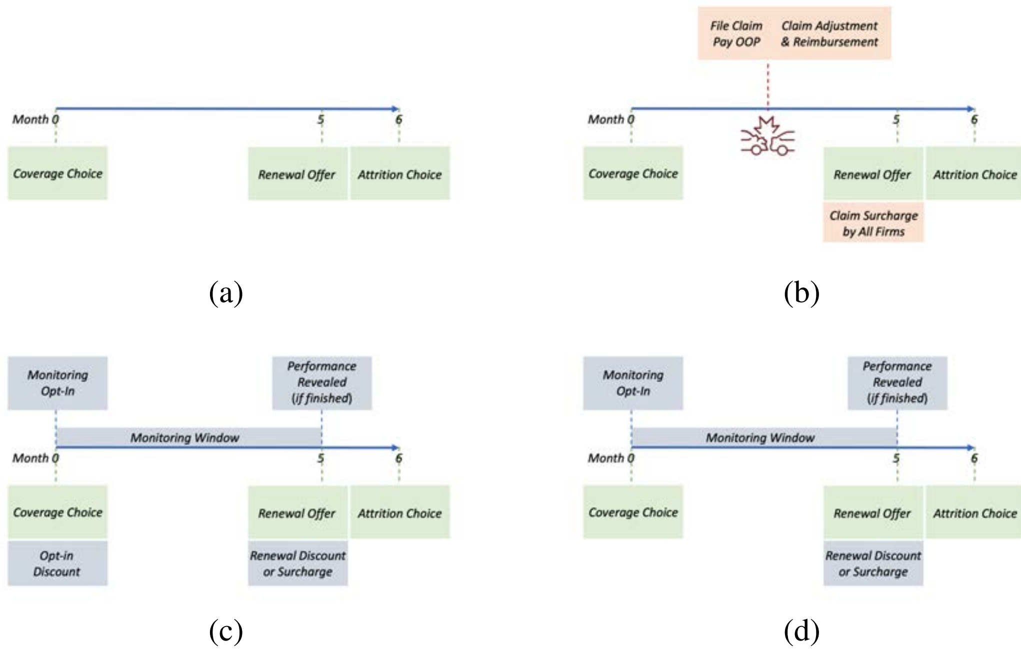

Figure 1: Timing Illustration of Auto Insurance and Monitoring Program

Dataset 1 - Panel data from an auto insurer Our first dataset comes from a national auto insurer in the U.S. It is a panel that spans 2012 to 2016, and covers 22 states. For tractability, we focus only on *single-driver-single-vehicle* insurance policies sold online or via phone so that the consumer of the insurance policy is also the driver being insured. We observe more than 1 million consumer-drivers for an average duration of 1.86 years (3.73 periods)[9]. The date range spans periods pre- and post-introduction of monitoring.

At the beginning of each period, we observe each consumer's observable characteristics[10] as well as the price menu offered to her, which includes a detailed breakdown of all available coverage options offered by the firm. We also see the consumer's coverage choice. For simplicity, we limit our attention to *liability coverage* (limits). Not only is this the most expensive coverage for the average consumer, but its mandatory nature makes it a focal point of firms' strategy and of the allocative benefit provided by the monitoring program. Liability arises in accidents involving two or more parties, in which the policy holder is at

[9]The panel is right-censored, but the censoring is plausibly uninformative.

[10]Main observables include consumer gender, age, martial status, education, out-of-state status, home-ownership, vehicle model, year, and financing, license and vehicle history, violation and accident records, credit history, prior insurance history, and zip code population density. See Table A.3 for a list of observables used in our estimation procedure.

least partially at-fault. As such, our focus also mitigates concerns about under-reporting.[11] During renewals, consumers who had recently filed a claim face a surcharge that ranges from 10% to 50% (Figure A.4).[12] Otherwise, the average consumer experiences close to no price change in renewal periods.[13] Overall, 5% to 20% of consumers leave the firm after each period.[14]

Table 1(a) presents summary statistics of prices, coverage levels, and claims. The average consumer is 33 years old, drives a 2006 vehicle, lives in a zip code with an average annual income of $142,000, and has 0.3 recorded accidents in the past 5 years. Per six-month period, she pays $380 in liability premium and files 0.05 liability claims (1 in ten years). We also observe her assigned risk class, given by her algorithmically generated personalized premium excluding the coverage factor, markups, and fees.

Dataset 2 - Price menus of competitors based on price filings To understand competition, we need to account for consumers' outside options. To do this, we augment our main dataset with price menus from the firm's main competitors. Our data includes quotes from all liability coverage options offered by the firm's top five competitors in each state (based on price filings), harnessed using Quadrant Information Services' proprietary software. For each consumer in our dataset, we obtain a complete set of competing prices through a precise match of her characteristics, including the state and the calendar day of insurance purchase.[15] Table 1(b) compares the quotes for the five most common liability coverage options across competitors in a representative U.S. state. Due to the large size of price menus, we end up with millions of quotes per state. Our reduced-form analysis and cost model estimation make use of the full dataset because they only depend on the firm's proprietary data. However, as demand estimation and counterfactuals require us to account for competitor quotes, we restrict our sample to three adjacent mid-western states, which cover 283,000 consumers and over 50 million quotes, for these analyses.

Looking ahead, observing competitor prices enables us to estimate consumers' inertia to switching firms. This allows us to accurately model the role of price competition in generating the status quo market equilibrium, and to predict how consumer retention would equilibrate in counterfactual information environments.

[11] In contrast, claim filing for single-car accidents is almost entirely discretionary.

[12] The surcharge varies only based on existing claims and traffic violation records.

[13] The first renewal sees some one-time discounts being removed, such as those for online processing.

[14] This number varies across U.S. states and calendar time, and it tends to be lower for older policies (higher period numbers).

[15] We match based on observable characteristics including those in Table A.3, violation records, zipcode, vehicle make and model.

(a) Premium, Coverage and Claims (6-month Period)

Statistic	Mean	St. Dev.	Min	Median	Max
Total premium ($)	632	364	69	548	22,544
Liability premium ($)	380	208	32	336	10,177
Risk class ($)	255	172	50	212	9,724
Total claim ($)	323	2,822	0	0	544,814
Claim count	0.18	0.67	0	0	12
Liability claim ($)	164	2,209	0	0	513,311
Liability claim count	0.05	0.32	0	0	7
Liability coverage ($000)	126	119	25	60	500
Liability coverage (index)	2.10	1.15	1	2	8
Mandatory minimum ind.	0.36	0.48	0	0	1
Renewal count	1.76	2.01	0	1	9
Calendar year (index)	2.66	1.38	0	3	5

Notes: Risk class is the pre-markups-pre-fees premium for liability coverage. Coverage index ranks coverage options in ascending order and sets the mandatory minimum in each state as 1.

(b) By Coverage (a representative U.S. State)

Liability coverage ($000)	40	50	100	300	500
Quotes ($)	335.14	343.43	382.03	422.13	500.48
- Competitor 1	482.68	506.11	564.34	626.81	730.56
- Competitor 2	263.14	279.15	314.46	347.69	405.22
- Competitor 3	319.42	348.97	388.48	428.64	464.36
- Competitor 4	511.24	567.58	613.74	682.87	790.83
- Competitor 5	421.84	363.96	403.64	433.17	497.79
Share within firm (%)	19	39	20	19	3
Liability claim ($)	154.98	155.54	154.16	143.43	107.54
Liability claim count	0.05	0.05	0.04	0.03	0.03

Notes: This table reports the average quotes and claims of the Firm and its top 5 competitors by market share. We focus on one U.S. state to avoid pooling across states with different coverage options. In this state, the mandatory minimum and the most popular coverage changed from $40,000 to $50,000 during the research window.

Table 1: Summary Statistics

1.2 Monitoring Program

Our research focuses on the firm's voluntary monitoring program for new customers.[16] The monitoring process is summarized in Figures 1c and 1d. When consumers first arrive, they choose whether to opt in to monitoring immediately before seeing the coverage price menu. All consumers are provided with information on the kinds of driving behavior that are tracked and rewarded—high mileage, driving at night, high speed, and harsh braking—but the exact mapping to discounts is opaque. In addition, the firm offers an opt-in discount on the first-period premium independent of performance. It also spells out the mean and range of the renewal discount that would be applied to all subsequent (renewal) periods.[17]

Consumers who opt in receive a simple device via mail within a week. They then have until the end of month five to accumulate some 100-150 days of monitored driving. If completed, the firm evaluates their performance and includes an appropriate renewal discount when giving renewal quotes.[18] If an accident occurs, monitoring data do not influence claim reporting, handling, or future premium adjustment. Monitoring continues after any disruptions from the accident.

During the monitoring period, monitored drivers receive real-time feedback on their performance. Key statistics of recorded trips are posted online. The firm also offers active reminders through media such as text messages, mobile app push notifications, and beeping from the monitoring device when punishable behaviors are records.

Nevertheless, monitoring data is *proprietary*. We verify this by confirming that the firm's monitoring information does not appear anywhere in its competitors' price filings. More generally, other firms face many practical hurdles in getting and using monitoring information. First, verifying monitoring outcomes with consumers alone is difficult, labor-intensive, and may be subject to legal liability.[19] More importantly, firms may have different preexisting risk assessments, underlying costs, and markups for serving the same type of consumers.

The proprietary nature of monitoring data also prevents us from observing the details of com-

[16] In our setting, monitoring was offered as a one-time program for new customers only.

[17] The average opt-in discount is 4.6% in our estimation dataset. We cannot disclose the renewal discount range exactly to avoid identifying our data provider, but it centers around 7% and spans zero (-15% to 40%, for example).

[18] 27% of drivers who start monitoring do not finish. Our main analysis treats these drivers as unmonitored and jointly consider consumers' decision to *start and finish* monitoring. This is because 97% of non-finishers drop out during a two-month grace period (no penalty) in which the firm communicates projected renewal discounts to monitored drivers (afterwards, dropping out results in the maximum amount of renewal surcharge). We thus consider non-finishers as if they have reversed their opt-in decision after forming the correct belief about their monitoring outcome. Non-finishers also have similar risk profile as other opt-out consumers: in **??**, if we separate non-finishers from the opt-out group, their average claim count is only 4% below the latter.

[19] The privacy policy agreed upon at monitoring opt-in prevents the firm from sharing personally identifiable data.

peting monitoring programs. Public filings do contain some information on these programs, but it is difficult to interpret reliably. For instance, Reimers and Shiller (2019) uses public filings to construct a dataset of the introduction timeline of U.S. auto-insurance monitoring programs. However, the program that our paper analyzes is shown to have been introduced up to four years before what our proprietary data suggests across U.S. states. The discrepancy is largely due to the various trials and R&D efforts employed by the firm to fine-tune the structure and technology of its monitoring program. Furthermore, monitoring generally takes up a small fraction of the market during our research window—estimated to be under 4% in 2014 by Ptolemus Consulting (2016). In fact, until the second half of 2016, our firm was the only one offering monitoring in the three states used in our estimation sample. Our empirical results, which assume that competing firms do not have their own monitoring programs, are thus an accurate reflection of the competitive environment under study.

Dataset 3 - Monitoring Our data on the firm's monitoring program includes its pricing schedule, consumers' opt-in choices, and realized monitoring scores and renewal discounts for those that opt in. Across calendar time and states, the average monitoring finish rates are around $10-20\%$ (Figure B.1). The firm's monitoring pricing is discussed in detail in Section 5 as well as in Appendix B. Importantly, consumers face the same baseline price (before monitoring-related discounts are applied) whether they choose to opt in to monitoring or not. We empirically validate this pricing policy in B, in which we show that prices without monitoring discounts are exactly the same across the opt-in and opt-out consumer groups.

Monitored drivers' performance is summarized by a one-dimensional *monitoring score*, the distribution of which is plotted in Figure 2(a). The more punishable behavior recorded for a given monitored driver, the *higher* the score that she receives. We treat this score as the output of a fixed monitoring technology that provides additional information on consumers' *future* accident risk, given observations of driving behavior during the monitoring period. Figure 3 plots the average claim count in period two based on monitoring choice and outcome in period one. Compared to unmonitored drivers, those who finished monitoring are 22% safer. Among finishers, the quintile of their monitoring score strongly predicts their second-period risk, which ranges from 60% better to 40% worse than the opt-out pool.

Consumers that opt in to monitoring face the same renewal choices as others, except that their renewal quotes include performance-based monitoring discounts or surcharges. Figure 2(b) compares the distribution of first-renewal price changes: normalizing the average price change for an unmonitored driver to be one, the average monitored driver receives an extra discount of 7%. Moreover, this monitoring renewal discount persists beyond the second period, as shown in Appendix Figure A.3, which is consistent with the firm's upfront communication with consumers about the potential reward of monitoring.

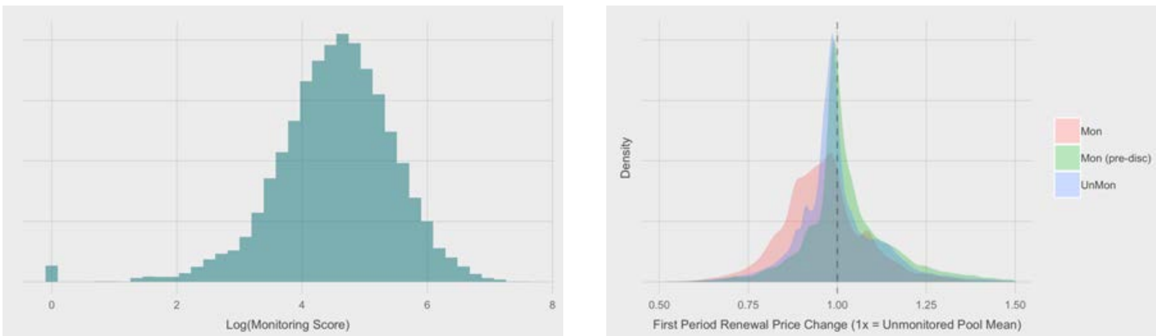

(a) (b)

Figure 2: Monitoring Score and Renewal Discounts

Notes: (a) plots the density of the (natural) log of monitoring score for all monitoring finishers. The lower the score the better. Drivers that received zero score plugged in the device continuously for enough days but did not drive. We ignore these drivers in all subsequent tests. (b) plots the benchmarked (per firm request) distribution of renewal price change at the first renewal, by monitoring group. 1x represents the average renewal price change factor for the unmonitored group. The one-time monitoring opt-in discount (applied on the first period premium) is taken out in order to isolate the renewal discount for monitored drivers. "Mon" and "UnMon" are monitored and unmonitored groups, while "Mon (pre-disc)" is the renewal price change for monitored drivers without the monitoring discount.

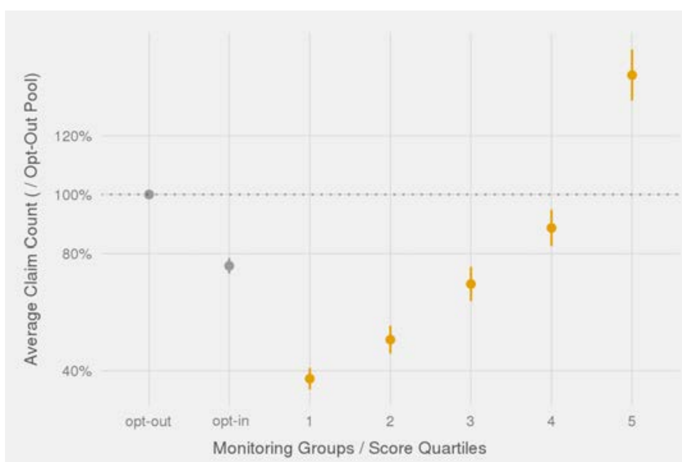

Figure 3: Comparison of subsequent claim cost across monitoring groups

Notes: This is a binned-scatter plot comparing average claim count of the second period ($t = 1$, after monitoring ends) across various monitoring groups. The benchmark is the unmonitored pool, which is the "opt-out" group. Group "opt-in" includes all monitored drivers that finished the program per definition in section 1.2. Groups "1" to "5" breaks down the "finish" group based on the quartile of the drivers' monitoring score. Lower monitoring score means better performance.

2 Reduced-form Evidence

This section introduces and measures the selection and moral hazard effects associated with the monitoring program. Consumers that opt in to monitoring become safer when monitored. Despite this behavioral change, monitoring reveals persistent and previously unobserved risk differences across consumers, which in turn drives advantageous selection into the program.

2.1 Risk Reduction and the Moral Hazard Effect

If monitoring technology is effective, consumers may want to appear safer when monitored. If, in addition, their risk is modifiable, then we should expect them to be riskier in unmonitored periods than in the monitored one–a moral hazard effect.[20]

Since monitoring is temporary, we can directly measure this effect by comparing claim outcomes for the *same* monitored consumers before and after monitoring ends. This exercise requires us to balance our panel. We thus focus on the first three periods (18 months).[21] There may be spurious trends in the claims rate across periods that are irrelevant to monitoring. We thus include exhaustive observable controls and adopt a *difference-in-differences* approach. Among monitored consumers, we take the first difference in claim counts[22] between monitoring and post-monitoring periods. This difference is then benchmarked against its counterpart among unmonitored consumers (the control group).

$$C_{it} = \alpha + \tau m_i + \omega \mathbf{1}_{post,t} + \theta_{mh} m_i \cdot \mathbf{1}_{post,t} + \mathbf{x}'_{it}\beta + \varepsilon_{it} \qquad (2)$$

Here, i,t index the consumer and period in our panel dataset. C denotes the claim count, and m_i is a consumer-specific indicator for whether i has finished monitoring. The vector $\mathbf{x}$ denotes a rich set of observable characteristics that the firm uses in pricing, which includes state fixed effects and third-order polynomials of the calendar year and month, as well as a time-varying risk class measure assigned to each consumer by the firm.[23]

[20] This is studied in Fama (1980) and Holmström (1999). A similar setting is online tracking of consumers' purchase history (Taylor 2004; Fudenberg and Villas-Boas 2006). If consumers know that buying expensive items online labels them as inelastic shoppers and lead to higher future prices, they may refrain from those purchases.

[21] Attrition is about $10-15\%$ per period and our data is right-censored, so balancing the panel eliminates 46% of our data. In our robustness check, we show results with only two periods.

[22] Throughout our reduced-form analyses, we use claim count as our cost proxy. This is because claim severity is extremely noisy and skewed. This is also common practice in the industry, where many risk-rating algorithms are set to predict risk occurrence only. We therefore present our estimates mostly in percentage comparison terms.

[23] See Table A.1 for a list of other main observable characteristics.

Table 2: **Estimates From Incentive Effect Regression**

	dependent variable: claim count (C)								
explanatory variables	(1)	(2)	(3)	(4)	(5)	(6)	Parallel Trend/Placebo		
constant	0.045^{***}	0.002	0.003	0.046^{***}	0.003	0.004	0.001	0.002	0.002
	(0.000)	(0.005)	(0.005)	(0.000)	(0.005)	(0.005)	(0.006)	(0.006)	(0.008)
post monitoring indicator	-0.001^{*}	-0.003^{***}	-0.003^{***}	-0.001^{**}	-0.003^{***}	-0.003^{***}	0.001^{*}	-0.001	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)	(0.001)	(0.002)
monitoring indicator (m)	-0.013^{***}	-0.012^{***}	-0.012^{***}	0.008^{***}	0.005^{***}	0.005^{***}	-0.005^{***}	-0.005^{***}	-0.004^{***}
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)
monitoring duration (z)				-0.026^{***}	-0.020^{***}	-0.020^{***}			
				(0.002)	(0.002)	(0.002)			
interaction ($\mathbf{1}_{post} \times m$)	0.008^{***}	0.009^{***}	0.009^{***}	-0.005^{**}	-0.005^{***}	-0.005^{***}	**-**0.001	0.000	**-**0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)
interaction ($\mathbf{1}_{post} \times z$)				0.015^{***}	0.016^{***}	0.016^{***}			
				(0.002)	(0.002)	(0.002)			
observables controls (x)	N	Y	Y	N	Y	Y	Y	Y	Y
coverage fixed effects	N	N	Y	N	N	Y	Y	Y	Y
implied risk reduction (%)	28.0	29.4	29.5	27.5	29.4	29.6			
pre- / post-periods - "1st diff"	0 / 1-2						1 / 2	2 / 3	3 / 4
treatment / control - "2nd diff"	$t = 0$ finisher / unmonitored			all finishers / unmonitored			$t = 0$ finisher / unmonitored		
number of drivers per period	755,614			809,784			755,614	539,296	397,642

Notes: This table reports results of equation (2). The estimate on the interaction term ($\mathbf{1}_{post} \times m$ or z) measures the "treatment effect" of monitoring ending on claim count across periods. We first balance our panel data to include all drivers who stay till the end of the third semester ($t = 3$). This gives us two renewal semesters ($t \in \{1,2\}$) after the monitoring semester ($t = 0$). We control for a full set of observables, including driver and vehicle characteristics and tiers (past records of violations or claims). It also includes state fixed effects and third-order polynomials of calendar year and month. Continuous observable characteristics are normalized. We report estimates with and without these controls.

Columns (3) and (6) are our main specification. Column (3) focuses on monitored drivers who finished within the first period, while Column (6) introduces additional variation in monitoring duration and timing and looks at all monitoring finishers. Columns (1,2,4,5) show robustness of our estimates to observable and coverage fixed-effect controls. The right-most columns are placebo tests for parallel trends among treatment/control groups after monitoring ends. We first try to detect a similar change from $t = 1$ to $t = 2$. We drop all observations from period 0, and roll the post-period cutoff one period forward, so that $1_{post,t} = 1 \iff t \geq 2$ (changed from $t \geq 1$). Naturally, we look at the future trends of monitored drivers who finished within the first semester and drop other monitored finishers. We find similar results by repeating this test in subsequent periods. As we need to balance panels, number of drivers drop in these tests.

Among consumers who finish the program, the actual duration of monitoring differs. Further, a small fraction of consumers do not finish monitoring until subsequent periods.[24] To make use of this plausibly exogenous variation in monitoring duration, we introduce another specification, using the monitoring duration in the first period, z_i, to indicate treatment intensity. This is calculated as the fraction of days monitored in the first period minus the same fraction in post periods.[25]

Results are reported in Table 2. We find a large moral hazard effect. Column 3 corresponds to the specification in Equation (2) with the addition of insurance coverage fixed effects, which soak up the effect of coverage adjustments between periods. This shows that the average claim count for monitored consumers is 0.009 or 23% lower during the monitoring period, compared to after it. Adjusting for the average monitoring duration of first-period monitoring finishers (142 days), a fully-monitored period would be 29.5% less costly to insure for the same consumer. Adding additional variation in monitoring duration generates similar results (Column 6). Since we do not observe consumer claim progression before they come to the firm, we test for parallel trends between the monitored and unmonitored groups by repeating the baseline specification in subsequent (unmonitored) periods. No differential claim change across periods can be detected (Columns 7-10).

For robustness, we adapt specification 2 slightly to look at the claim progression across monitoring groups on a longer but smaller balanced panel (six periods). The fixed-effect estimates for each policy period, θ_t, are illustrated in Figure 4. The level difference between the grey and the orange lines after period 0 represents a persistent difference in riskiness between opt-in consumers and their opt-out counterparts—a selection effect quantified in the next subsection. There is an additional risk reduction of 0.009 in period 0 among opt-in consumers, which closely tracks the moral hazard effect shown in Table 2 Column (2).

$$C_{it} = \alpha + \tau m_i + \omega_t \mathbf{1}_t + \theta_t m_i \cdot \mathbf{1}_t + \mathbf{x}'_{it}\beta + \varepsilon_{it} \tag{3}$$

Next, we investigate heterogeneity in the moral hazard effect across consumers with different

[24] Based on interviews with managers, among finishers, delays in finishing is predominantly caused by device malfunction or delayed start of monitoring due to mailing issues, etc.

[25] As discussed above, some consumers started monitoring but dropped out without finishing. This would bias our results if claims itself leads to non-finish. Out of more than 10,000 claims we observe among monitored consumers, only 13 occurs within 7 days before or after monitoring drop-out. In Table C.1, we further test the robustness of our results by repeating our main analyses on all consumers who started monitoring. This implies larger moral hazard effect adjusting for monitoring duration. However, if some monitored consumers drop out as they discover that they cannot change their risk, the incentive effect estimate would be contaminated by this selection effect.

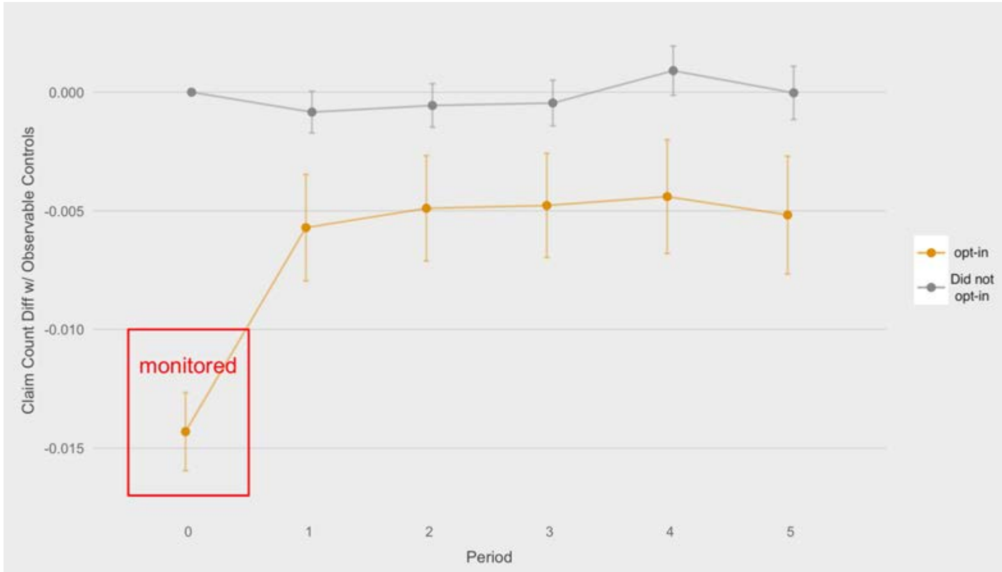

Figure 4: Claim Progression across Monitoring Groups

Notes: This graph reports the fixed effect estimates of eq. (3). The grey line plots ω_t while the orange line plots $\omega_t + \theta_t$, both against insurance periods t. The red box is superimposed ex-post to represent the period when opt-in consumers are monitored. Error-bars report 95% confidence interval.

observable characteristics and different insurance coverage choices:

$$C_{it} = \alpha + [\tau m_i + \omega \mathbf{1}_{post,t} + \theta_{\mathbf{mh}} m_i \cdot \mathbf{1}_{post,t}] \cdot (1, \mathbf{x}_{i0}, y_{i0})' + \mathbf{x}_{it}' \beta + \varepsilon_{it} \tag{4}$$

Here, $(1, \mathbf{x}_{i0}, y_{i0})'$ indicates the initial characteristics and coverage choice of consumer i and is thus time-invariant. This vector is interacted with the difference-in-differences term in C2 to uncover heterogeneity in the moral hazard effect. Results are presented in Figures A.5 and A.6. No systematic heterogeneity is detected.

We discuss two important caveats of our results. First, monitoring mitigates moral hazard because it builds a reputation mechanism by signaling consumers' persistent risk level, *not* because it directly rewards effort during monitoring (Fama 1980; Holmström 1999). The magnitude of risk reduction can be different in the latter setting.[26] However, the dynamic reward in our setting indicates that the average consumer is forward-looking and responds greatly to future incentives. Second, our estimate measures a treatment-on-treated effect. Even though there is little observed heterogeneity in the moral hazard effect among opt-in consumers (Figure A.5 and A.6), those who do not opt in may be less capable of altering their

[26]We are also unable to disentangle the "Hawthrone effect" from consumers' responsiveness to financial incentives in our estimate. Since consumers must be aware of the data collection to be incentivized for it, we consider this effect as part of the incentive effect.

risk. The moral hazard effect that we have identified is thus likely larger than the population average.[27] To avoid external validity concerns, our counterfactual analysis maintains the opt-in structure of the monitoring program and does not extrapolate to scenarios where the monitoring rate is significantly higher than in the data.

2.2 Private Risk and the Selection Effect

Are consumers who choose monitoring safer than those who do not? Table 3 reports the results of regressing claim counts in the first period ($t = 0$) on a monitoring indicator, controlling for the same variables as in Column (3) of Table 2. The coefficient on the monitoring indicator suggests that the risk differential between the two groups is larger than the moral hazard effect discussed in the previous section: the latter only accounts for 64% of the risk differential. This suggests that consumers possess private information on their own risk, driving advantageous selection into monitoring.

Table 3: First-Period Claim Comparison Across Monitoring Groups

	Dependent variable: Claim Count ($t = 0$)
monitoring indicator	-0.014^{***}
	(0.001)
observable controls	Y

Notes: This table reports the results of a regression where the dependent variable is first-period claim count, and the independent variables are the monitoring indicator and observable controls (including a constant). This is done within all first-period finishers of the monitoring program. This variable is consistent with the monitoring indicator in the incentive effect regression (2) (Table 2), so as to facilitate comparison and decomposition. *p<0.1; **p<0.05; ***p<0.01

Selection into monitoring also implies that the technology is effective at capturing previously unobserved differences in drivers' risk types, further allowing the firm to dynamically select safer drivers. The following regression examines both factors. It demonstrates how the average cost to insure a consumer after period 0 varies based on the consumer's period-0 monitoring choice and score.

$$C_{it} = \alpha_t + \theta_{m,t} m_i + \theta_{s,t} s_i + \mathbf{x}_{it}' \beta_{\mathbf{t}} + \varepsilon_{it} \tag{5}$$

[27] We also suppress selection on moral hazard in counterfactuals (Einav, Finkelstein, Ryan, Schrimpf, and Cullen 2013). In equilibrium, the firm assesses the signal that monitored consumers send based on their future claim records when they are no longer monitored, which corresponds to the renewal discount it gives. Thus, risk reduction is compensated only to the extent that it correlates with consumers' unmonitored risk type. If safer consumers' risk levels are also more responsive to incentives, as suggested by a pure effort cost model, selection on the incentive effect can be important. In particular, perfect revelation of a continuum of risk types is possible, as characterized in Mailath (1987), with a monotonicity condition similar to the single-crossing condition. However, consumers likely have multidimensional heterogeneity in reality, so consumers' performance during monitoring may not perfectly reveal their risk types (Frankel and Kartik 2016).

Here, $m = 1$ indicates monitored consumers and s denotes their monitoring scores. The latter is normalized among monitored consumers and set to 0 for others. Figures A.7 and A.8 report $\hat{\theta}_{m,t}$ and $\theta_{s,t}$ for renewal periods $t = 1$ to 5. It shows that a monitored driver who scores one standard deviation above the mean has a 29% higher average claim count in the first renewal. Further, controlling for claims does not alter our estimate much. This suggests that the sparsity of claims greatly limits the informativeness of claim counts on driver risk in the short run, highlighting the value of monitoring data.

A potential concern is that monitored drivers may learn to become safer during the monitoring period. This does not influence our moral hazard estimates (derived from within-consumer behavioral change), which we show is the main source of surplus gain from introducing monitoring in Section 4. However, ignoring learning may exaggerate the degree of advantageous selection. In Appendix C.1, we use consumers' prior speeding violation records (before period 0) to test for learning and cannot reject a no-learning hypothesis.

2.3 Extensive Margin and Renewal Elasticity

Consumer demand for insurance is complex. As we detail in the end of this section, welfare and profit analyses require structural estimation of consumer demand across monitoring, insurance coverage, and insurance firm. Here, we focus on the binary choice of renewal acceptance after monitoring concludes and the differential price sensitivity across the monitored and the unmonitored consumer groups:

$$\mathbf{1}_{it}^{\text{renewed}} = \alpha M_i + \theta \log p_{it}^{\text{renewal}} + \theta_M M_i \cdot \log p_{it}^{\text{renewal}} + \psi s_i + \mathbf{x}_{it}' \beta + \varepsilon_{it} \tag{6}$$

We augment the consumer-specific monitoring indicator m to get M, which admits three discrete monitoring *groups*: unmonitored consumers, monitored consumers that receive discounts, and monitored consumers that receive no discount or a surcharge. We include the same controls $\mathbf{x}$ as Column (3) of Table 2. We also add consumer-specific monitoring score s_i, normalized among monitored consumers and set to 0 for unmonitored ones. The parameters of interest are θ and θ_M, which combines with M to give price elasticities across monitoring groups.

The regression above may produce biased estimates if renewal prices are endogenous to unobserved demand factors or competitor pricing. For robustness, we adopt an instrumental-variable approach to compute price elasticity. We narrow the sample down to 30-day windows around rate revision events within each state, and instrument renewal pricing using Z_{it}, which is a post indicator representing the consumer arriving after her corresponding rate revision takes effect. Our estimation adds rate revision fixed effects to the controls ($\mathbf{x}$) of both the first-stage and the reduced-form regressions (Equations 7 and 6). The exclusion restriction is supported by an event-study argument: consumers arriving right before and after a rate revision event should be similar, and the price differential they face is thus independent from their idiosyncratic unobservable demand factors.

$$\log p_{it}^{\text{renewal}} = \tau M_i + \omega Z_{it} + \omega_M M_i \cdot Z_{it} + \phi s_i + \mathbf{x}_{it}' \gamma + \nu_{it} \tag{7}$$

Figure 5 shows that both OLS and IV regressions produce similar price elasticity estimates across monitoring groups. This suggests that the rigid pricing structure as well as our observable controls and fixed effects helps us avoid price endogeneity conditional on observed pricing factors (including monitoring score). Further, the price regression (Equation (7)) has an adjusted R square of 0.95. After controlling for rate revisions (within-state across calendar-time price variation), the only remaining systematic price variation is across zip codes with similar average household income and population density.

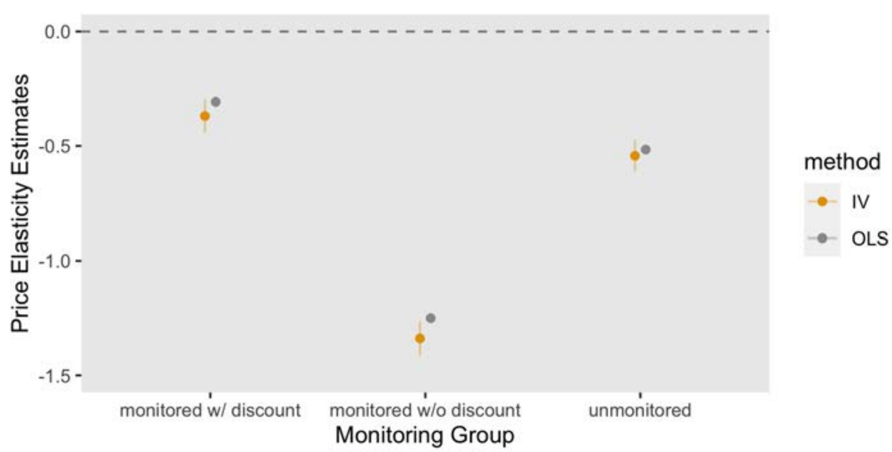

Figure 5: Price Elasticity of Renewal Acceptance by Monitoring Groups

Notes: The dots report the renewal price elasticity estimates across monitoring groups using OLS or IV, which correspond to $\hat{\theta} + \hat{\theta_M} \cdot M$ in Equation (6). They measure the average percentage point increase in renewal likelihood with respect to 1% increase in renewal price. M are indicators for three monitoring groups (from left to right): monitored consumers that receive price discounts, monitored consumers that receive no price discounts or surcharges, unmonitored consumers. Lines reports 95% confidence interval, with standard errors clustered on the consumer level.

During renewal, monitored consumers who are eligible for discounts are less price sensitive than the average unmonitored consumer, while the opposite is true for consumers that receive no discounts or a surcharge. Intuitively, this means that the firm faces a "flatter" residual demand curve when giving discounts to monitored consumers and a "steeper" one when surcharging. For monitored and discount-eligible consumers, which are revealed to be safer than the firm's prior expectation, this suggests an informational rent that the firm can extract based on the monitoring data, potentially enabling higher markups after the monitoring period.[28] On the other hand, higher price elasticity among monitored but discount-ineligible consumers can explain why the firm rarely surcharges monitored consumers.

The need for structural models Taken together, our results in this section show that monitoring produces verifiable information on accident risk, which effectively separates risky consumers from safe ones who are otherwise indistinguishable. This leads to advantageous selection into the program and safer driving when consumers are monitored. We have thus separately identified the effects of moral hazard and asymmetric information, extending the literature that has thus far relied on correlations between plan choice and claims (Chiappori and Salanie 2000; Cohen and Einav 2007; Jeziorski, Krasnokutskaya, and Ceccarini 2019).

As Einav, Finkelstein, and Levin (2010) point out, reduced-form methods are useful in proving existence of moral hazard and selection patterns. However, quantitative welfare and counterfactual analyses require estimating how risk and preferences contribute to consumer choices and the degree of asymmetric information. We achieve in Sections 3 and 5 by specifying a model of consumer utility, risk information, and firm pricing strategy. This serves three important functions. First, monitoring simultaneously affects consumer welfare through multiple channels that cannot be combined or compared without estimating risk preference: accident risk may be reduced; privately safe drivers may anticipate a lower price menu and may increase coverage, but they may face higher premium volatility. Second, consumers make three unordered discrete choices: firm, monitoring opt-in, and insurance coverage. The inter-dependence of these choices and how they correlate with accident risk drive the selection pattern that is critical to market efficiency. Third, the firm faces intense regulatory pressure when introducing the monitoring program, leading to overly conservative pricing.[29] Deriving equilibrium pricing without such restrictions necessitates the modeling of how the firm's information on consumer accident risk evolves with monitoring.

[28]Residual demand elasticity is a function of consumer preferences, frictions (e.g. search and information frictions), competitor pricing, and selection due to remaining consumer private risk even after controlling for their observables and monitoring scores. In latter Sections, we introduce competitor pricing data and explicitly model risk preference, switching inertia, and selection patterns to better understand markups.

[29]For example, as shown in Appendix B, the firm did not raise price for the unmonitored pool, which is primarily to appease some state regulators' blanket aversion to price increase.

3 Cost and Demand Models of Auto Insurance and Monitoring

This section develops a structural model for consumers' accident risk and their demand for insurance and monitoring. In the first period, consumers observe their risk types and make three choices: firm, insurance coverage, and monitoring opt-in. After this, claims are realized; the monitoring scores for opt-in consumers are revealed to the firm. Consumers are then offered the corresponding renewal price for the second period based on the realized claims and the scores.

We describe our model in two parts. First, we characterize a consumer's choice utility upon the realization of her monitoring score and any possible claims ("realized choice utility"). This features risk aversion, path-dependence (choice inertia and disutility for monitoring), and her expectation of future prices. We then describe the data generating processes for claims and monitoring scores in a model of the firm's cost to insure consumers, which features risk heterogeneity, the moral hazard effect, and the monitoring score's signaling precision. We can then unify cost and demand factors with an expected utility framework to capture selection. Lastly, we discuss estimation procedures and sources of identification for key parameters, before demonstrating model fit and validation out-of-sample.

Realized choice utility Aside from consumers' risk types, our choice model highlights three factors. (i) Risk aversion governs both preference for insurance and distaste for future price fluctuations. (ii) Demand frictions: firm-switching inertia leads to imperfect competition among insurers. Consumers' disutility from being monitored accounts for factors such as privacy or effort cost associated with monitoring. These factors also sustain a partial pooling equilibrium, in which only a fraction of the population is monitored. (iii) Future prices contain most of the benefit of monitoring and depends on claims and monitoring score.

Denote consumers, periods and decision menu options ("plans") by i, t, and d, respectively.[30] Plans, $d = \{f, y, m\}$, consist of firm (f), coverage (y), and monitoring (m) choices. Consumer preferences are characterized by a standard von Neumann-Morgenstern utility function u_{idt} with absolute risk aversion, denoted by γ. Each consumer i starts period t with annual income w_{it} and evaluates insurance choices based on their impact on her utility through the consumption term h_{idt}:

[30] Monitoring takes place in the first period ($t = 0$).

$$
\begin{aligned}
u_{idt}(C,s) &= u_\gamma(w_{it} + h_{idt}(C,s)) && (8)\\
h_{idt}(C,s) &= -p_{idt} - \underbrace{\mathbf{1}_{d,t-1}\cdot\psi_{idt}}_{\text{friction}} - \underbrace{e(C,y_d)}_{\text{oop}} - \underbrace{p_{idt}\cdot R_{idt}(C,s)}_{\text{renewal price}} && (9)\\
\text{where } \psi_{idt} &= \underbrace{\mathbf{1}_{d,t-1}\cdot\eta_0}_{\text{baseline inertia}} + \underbrace{\mathbf{1}_{f_d,t-1}\cdot\eta_{it}}_{\text{firm-switching inertia}} + \underbrace{\mathbf{1}_{m_d}\cdot\mathbf{1}_{t=0}\cdot\xi_{it}}_{\text{monitoring disutility}} && (10)
\end{aligned}
$$

Consumption h spans a one-year horizon and consists of two types of components: upfront costs, p and ψ, and stochastic costs, $e(C,y)$ and $R(C,s)$.[31] p_{idt} is the price for plan d in period t. The term ψ_{idt} captures the degree of path-dependence in consumer choice in monetary terms. This includes a cost of overcoming inertia: a baseline inertia η_0 that hinders any choice adjustment (indicated by $\mathbf{1}_{d,t-1} = 1$), and a firm-switching inertia η_{it} that deters consumers from exploiting financially lucrative outside options.[32] It also includes disutility from being monitored, ξ_{it}, which may reflect unobserved factors such as hassle costs and privacy concerns.[33]

Out-of-pocket expenditures ("oop"), e, and renewal prices charged for each plan, R_{idt}, depend on the realization of claims C and the monitoring score s. Consumers pay the portion of accident expenditures that exceeds their plan's coverage limit. Renewal prices are adjusted by multiplying two factors: a baseline factor $R_{0,idt}(s)$ that varies based on monitoring results, and a surcharge for claims, $R_{1,C}$. We model the baseline factor by a Gamma distribution with shape parameter β_R and rate parameter $\alpha_{R,imt(s)}$, the latter of which depends on observables and monitoring opt-in.

Our model maintains the assumption that consumers know their own accident risk. Although future price fluctuations still depress consumer choice, we do not capture the welfare impact of consumers' risk from learning about their own type (Hirshleifer 1978; Handel, Hendel, and Whinston 2015; Hendren 2018). In other words, reclassification risk only comes from uncertainty associated with the firm's inference of consumer accident risk (based on monitoring results), not from consumers' uncertainty about their own risk.

[31] We assume that consumers are myopic beyond a one-year (two-period) horizon. This is the simplest model that captures the different types of costs and benefits of monitoring programs.

[32] These terms capture imperfect competition that supports the observed attrition rate given price dispersion in the data 1(b). Inertia accounts for the search and switching costs as well as potential brand differentiation (Farrell and Klemperer 2007; Honka 2012).

[33] Monitoring is a one-time offering and choice for new customers, so ξ can only incur at $t = 0$. Moreover, our data does not allow us to identify the micro foundation of this disutility term. It may include real costs like privacy and driving effort costs ((Lin 2019)), or systematic misconceptions and salience issues (that goes away in a monitoring mandate). Our counterfactual analyses therefore only focus on local deviations from the current regime.

Claims and monitoring scores Claims arrive according to a Poisson distribution. The rate parameter, λ_{imt}, has a time-varying mean $\mu_{\lambda,imt}$ that depends on monitoring choice m.[34] It also contains an additive error $\varepsilon_{\lambda,i}$ that is individual-specific, persistent over time, and log-normally distributed with spread σ_λ. This error captures unobserved risk differences across consumers. Further, each claim has a stochastic cost ℓ, drawn from an independent Pareto distribution. The monitoring score s is an informative signal of the consumer's risk types. For opt-in consumers, a score is drawn once after the first semester, according to a log-normal distribution with an individual-specific mean $\mu_{s,i}$ and precision σ_s.

At each period t, consumer i chooses d from her feasible choice set D_{it} so as to maximize her expected utility, subject to a random coefficient $\zeta_{idt} \sim \mathcal{N}(0,\sigma_\zeta)$ on plans offered by the monitoring firm $f^\star$, and an independently drawn Type-1 extreme value error ε_{idt}. We evaluate utility using a normalized second-order Taylor approximation of vNM utility around income w: [35]

$$d_{it} = \arg\max_{d\in D_{it}} \{v_{idt} + \varepsilon_{idt}\} \tag{11}$$

$$\text{where } v_{idt} = \mathbb{E}_{C,s}\left[u_{idt}(C,s)\right] = \mathbb{E}\left[h_{idt}\right] - \frac{\gamma}{2}\mathbb{E}\left[h^2_{idt}\right]. \tag{12}$$

Econometric assumptions and heterogeneity Heterogeneity across consumers is captured by a vector of consumer attributes x_{it} and individual random effects.[36]

Our *demand parameters* Θ_d include risk aversion γ, the type I error variance σ, baseline inertia η_0, linear coefficients on consumer attributes for firm-switching inertia, θ_η, and for monitoring disutility, θ_ξ, as well as parameters that characterize (expectation for) renewal pricing $\theta_{\mathbf{R}} = (\theta_{\mathbf{R},\mathbf{0}}, \theta_{\mathbf{R},\mathbf{1}})$:

$$\eta_{it} = (1, x_{it})'\,\theta_\eta$$
$$\xi_{it} = (1, x_{it}, \ln\lambda_{it})'\,\theta_\xi$$
$$\alpha_{R,imt}(s) = \begin{cases} \mathbf{x}^{R\prime}_{it}\theta_{\mathbf{R},\mathbf{0}} & m=0 \\ \left(\mathbf{x}^R_{it}, s\right)'\theta_{\mathbf{R},\mathbf{1}} & m=1 \end{cases}$$

[34] Note that monitoring impacts the claims rate for an individual in the same way no matter what insurance plan she chooses. This is consistent with our findings in section 2, that the moral hazard effect appears invariant to plan choice.

[35] See Cohen and Einav 2007 and Barseghyan, Molinari, O'Donoghue, and Teitelbaum 2013 for further discussion of this approximation. The key underlying assumption is that third- or higher-order derivatives are negligible.

[36] For each type of parameter, we use a set of consumer attributes that is consistent with those used in related actual firm pricing rules: x_{it}, x^R_{it} and x^s_i.

In order to fully capture selection into monitoring, we allow monitoring disutility to vary based not only on observables but also on unobserved risk λ. Without this, given an observable type of consumers, the propensity to opt in to monitoring would be fully determined by the financial rewards (lower accident likelihood and potential monitoring discounts). To the extent that there is heterogeneity—across observable groups and across unobserved risk types—in privacy preference, hassle costs, or in misperception of own risk or of the monitoring program, $\theta_{\xi,\lambda}$ can capture its effect on consumer choices.

Our *cost parameters* Θ_c include linear coefficients on consumer attributes and monitoring status for claim arrival rate, $\theta_\lambda = (\theta_{\lambda,\mathbf{0}}, \theta_{\lambda,m})$, as well as for the monitoring score $\theta_{\mathbf{s}}$. They also include the unobserved risk spread for new and old drivers, $\sigma_{\lambda,\text{new}}$ and $\sigma_{\lambda,\text{old}}$, the monitoring score precision σ_s, as well as the rate and location parameters of the accident loss Pareto distribution (ℓ_0, α_ℓ):

$$\mu_{\lambda,imt} = (1, x_{it})' \theta_{\lambda,\mathbf{0}} + \theta_{\lambda,m} \cdot \mathbf{1}_{m=1} \cdot \mathbf{1}_{t=0} \tag{13}$$

$$\mu_{s,i} = (1, \ln \lambda_i, x_i^s)' \theta_{\mathbf{s}}. \tag{14}$$

For tractability, we abstract away from the structure of effort provision that underlies the moral hazard effect. We assume that the effect is homogeneous across consumers and enters risk in a mechanical and additively-separable fashion via $\theta_{\lambda,m}$.[37] However, our approach allows us to capture the risk information that is revealed in the signaling equilibrium: a consumer's monitoring score contains additional information on her risk when (i) $\theta_{s,\lambda} \neq 0$, (ii) σ_s is finite, and (iii) s is not co-linear with x_i^s.

3.1 Estimation

We estimate our model of consumer cost and insurance demand using a two-step simulated maximum likelihood procedure.[38] First, we estimate the cost parameters Θ_c using the full dataset of claims and monitoring scores. We then estimate the demand parameters Θ_d using menu options, plan choices, and prices, taking the point estimates of the first stage as data.

The Type-1 extreme value distribution of ε_{idt} implies a mixed-logit structure on plan choice

[37] This is supported by our findings in Figure A.5 and A.6. For more careful treatment of moral hazard and risk determination, see Jeziorski, Krasnokutskaya, and Ceccarini (2014).

[38] We adopt the two-step procedure due to computational constraints. This comes at an efficiency cost. Standard errors for the demand estimates are not currently adjusted for two-step estimation.

with choice probabilities:

$$\begin{aligned}\Pr(d_{it}|\Theta_i) &= \Pr(\varepsilon_{idt} - \varepsilon_{id't} > [v_{idt}(\Theta_i) - v_{id't}(\Theta_i)] \ \forall d' \neq d \\ &= \frac{\exp[v_{idt}(\Theta_i)/\sigma]}{\sum_{d'} \exp[v_{id't}(\Theta_i)/\sigma]}\end{aligned} \tag{15}$$

Our model includes random coefficients that enter utility nonlinearly. Private risk, in particular interacts with various observed monitoring and coverage characteristics (renewal price, out-of-pocket expenditure), as well as unobserved demand parameters (risk aversion and monitoring cost). To account for this, we simulate 50 independent draws of private risk (ε_λ) and the zero-mean firm dummy (ζ) for every proposal of Θ_d.[39] We then compute likelihood for choices, claim counts and severities, monitoring scores and renewal prices, and average over the simulated draws.[40]

3.2 Identification

We now provide an informal discussion of the variation in our data that allows us to identify the parameters of our model.

For the cost parameters Θ_c, variation in average claim counts and monitoring scores across observable groups helps identify the associated slope parameters θ_λ and $\theta_{\mathbf{s}}$. Variation in claims between monitored and unmonitored periods and consumers helps identify $\theta_{\lambda,m}$. Given the claim arrival rate of an observable group, the variance in claim counts may deviate from that implied by the Poisson structure and therefore identify the spread of risk across consumers σ_λ. The same quantities in the data, when conditioned on not only observables but also on the monitoring score, help identify σ_s, the precision of the monitoring score signal. The rate parameter characterizing loss severity is identified by observed claim amounts.[41]

Identification of demand parameters Θ_d relies on price and contract space variation. Controlling for the attributes used in firms' pricing rules, the remaining price variation depends on location and calendar time. Specifically, price changes associated with the firm's and its

[39] Increasing the number of draws from 50-200 on a 10,000 sub-sample produces minimal effect on estimates.

[40] The Taylor approximation gives closed-form solutions for the first two moments of out-of-pocket expenditures and renewal prices. We therefore do not simulate losses or monitoring scores for each draw of random coefficients.

[41] Claim amounts are capped above by coverage limits. The Pareto distribution is sufficiently long-tailed so that loss events significantly larger than coverage limits still have non-degenerate support in consumer's expectation.

competitors' rate revisions (back-end changes in pricing rules) as well as cross-zipcode variation that are plausibly exogenous from consumer demand.[42] Notably, the firm altered the monitoring opt-in discount over time, generating a useful source of variation in monitoring incentives.

We also observe variation in consumers' contract space. Specifically, monitoring eligibility differs based on state, time, specific vehicle models, and renewal period. For instance, consumers who arrived before monitoring was introduced in their states or whose vehicles were older than 1995 were not eligible for the monitoring program. Monitoring is also only available to new customers. Meanwhile, mandatory minimum coverage changed in two states within our research window. We use one in our demand estimation and reserve the other for cross-validation (see Table 4).

Our primary concern is in identifying monitoring disutility (ξ) well. Given cost parameters and risk aversion, we can determine the relative attractiveness of the same coverage option with and without monitoring based on objective financial risk and rewards alone. On top of that, the monitoring disutility is pinned down by the actual monitoring share (under various pricing environments). The slope parameter on risk type ($\theta_{\xi,\lambda}$) controls the share of each risk type opting into monitoring. It therefore helps us fit both the share of monitoring and selection on risk.[43]

Another parameter of interest is risk aversion γ. For a given i,t, different γ values imply different gradient of Δv_{idt} across the multiple coverage options we observe in the data.[44] Therefore, conditional on risk parameters, risk aversion can be identified by how the empirical coverage share changes given contract space and pricing environment.[45] In our demand estimation, the Pareto severity parameters can also affect changes in coverage attractiveness. However, we restrict the Pareto distribution to approximate the actual (truncated) claim severity that we observe.

We also need to separately identify baseline inertia (η_0) and consumers' firm-switching iner-

[42]To hone in on this variation, our model include each consumers' assigned risk class in the cost model, and include controls for yearly trends, seasonality, and zipcode characteristics like income and population density in our demand parameters.

[43]Simply raising baseline monitoring cost for all risk types (conditional on observables) enhances selection but also necessarily reduces monitoring share.

[44]This is conditional on the fixed effect for the mandatory minimum plan (ψ_1). The fixed effect adds an additional degree of freedom to more flexibly fit the gradient of willingness-to-pay across coverage options.

[45]Based on the company's pricing rule in Equation 1, the price gradient across coverage options only depends on the actuarial risk class assigned to each consumer and the coverage factor. The latter is heavily regulated and rarely changes empirically. Each state offers an official guidance on the coverage options that auto insurers should offer and the corresponding coverage factors. Firms need to provide actuarial support to deviate from the guidance in order to avoid regulatory scrutiny.

tia (η). Conditional on observables, different levels of these parameters imply unique combinations of the share of consumers who adjust coverage rather than leave the firm at renewal time. We also observe rich variation in the competitive pricing environments conditional on observables. Under a given pricing environment, a choice of inertia parameters implies a corresponding threshold under which consumers would stay with the firm, and another one under which consumers would not adjust choices at all. Our observations of attrition under different environments and price menus thus inform inertia estimates that best rationalize the thresholds that generate them.

3.3 Fit and Validation

We demonstrate that our demand model is flexible enough to produce an accurate fit for four critical moments of the data in Figure 6 and in Table A.5. As Table 4 shows, we match monitoring and coverage shares of the firm well. Further, first-renewal attrition rates, the share of outside option, is also broadly consistent. More importantly, we also accurately fit the expected monitoring score. This demonstrates that the model is capable of capturing selection as well as the effectiveness of the monitoring score. Figure 6 confirms this graphically: we calculate the expected monitoring score for each consumer over all random-coefficient draws. The red line plots the simulated score weighted by the corresponding monitoring choice probability in each draw. The orange line plots the full distribution of expected monitoring scores, had everyone in the data finished monitoring.

Using these estimates, we can calculate the counterfactual (unmonitored) risk type of *monitored* drivers in the first period. When we numerically integrate over private risk ε_λ, we weight each draw by the choice probability of monitoring. This yields the expected risk type in the monitored pool without a moral hazard effect. We then repeat this procedure for the unmonitored pool. The selection effect is the ratio between the two at 21%.[46]

We also cross validate our demand estimates. In particular, one state in our dataset increased its mandatory minimum from \$30,000 to \$50,000. In our demand estimation, we draw from only the pre-change period for this state. The hold-out sample, however, contains all consumers in that state arriving in the post-period. As shown in Table A.6, our model performs well out of sample.

[46] This is similar to the 17% back-of-the-envelope calculation we did in the reduced-form section (Table 3). In Tables A.5 and A.6, we compare our model fit and cross validation to a more basic one that excludes the firm random coefficients ζ and the private monitoring disutility $\theta_{\xi,\lambda}$.

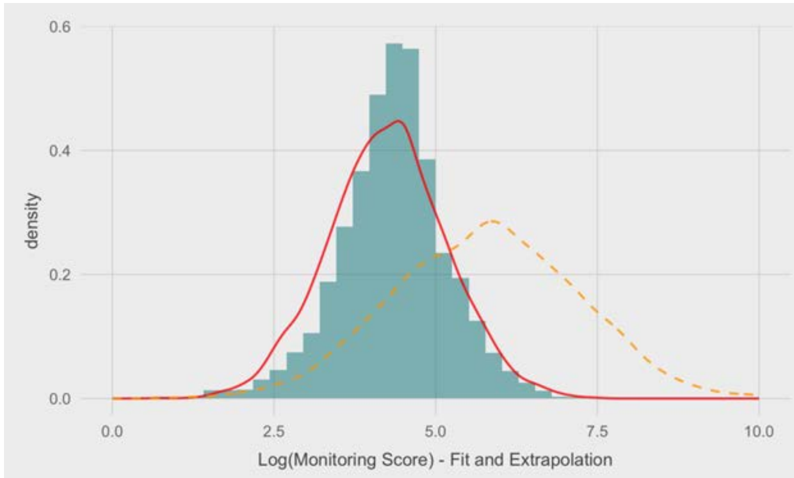

Figure 6: Monitoring Score - Fit and Extrapolation

Notes: The green histogram is the empirical distribution of monitoring score for monitoring finishers in our demand estimation data. The red line plots the fitted distribution as outlined above. The orange dotted line plots the density of the extrapolated distribution of monitoring scores had all drivers finished monitoring.

Table 4: Demand Model Fit and Cross Validation

	Model Fit		**Cross Validation**	
	Fit	**Data**	**Prediction**	**Hold-out Data**
Monitoring share (when eligible)	15.6%	15.3%	17.9%	17.6%
Expected score	4.25	4.30	3.97	4.17
Coverage share				
30K	12.5%	12.7%	-	-
40K	8.2%	8.5%	7.6%	7.2%
50K	49.8%	47.1%	60.5%	58.1%
100K	15.4%	17.0%	17.5%	19.6%
300K	11.9%	12.3%	10.9%	12.8%
500K	2.3%	2.4%	3.6%	2.4%
First renewal attrition	15.6%	15.2%	15.4%	14.7%

Notes: This table reports the fit of our demand model and cross validation results. Our demand estimation data pools across three states with different mandatory minimum. One state changed mandatory minimum from 30K to 50K; estimation data is drawn from only the pre-period of that state to capture monitoring introduction. First renewal attrition rate is benchmarked to data per the firm's request (reporting percent differences, not percentage point differences).

4 Estimation Results and Welfare Calculations

The raw estimates for our model are reported in Tables A.2 and A.3. In this section, we highlight some key results, provide intuition, and conduct welfare calculations.

The magnitude of private risk and the monitoring score's signal precision are presented in the left panel of Table A.2. Compared to Cohen and Einav (2007), we find significantly more unobserved heterogeneity in driving.[47] This can be attributed to our ability to capture information contained in an additional signal of private risk: the monitoring score. New consumers who do not have past claim records see particularly high spreads of private risk. Our estimates also capture the monitoring technology and the firm's renewal prices well. In particular, monitoring score rises with consumer risk, as do renewal prices for monitored consumers (Table A.4).

We find that consumers are not very risk averse in their auto insurance and monitoring choices. Our primary specification assumes homogeneous risk aversion, and the estimate of $\hat{\gamma} = 9.8 \times 10^{-5}$ is broadly consistent with the literature.[48]

Also consistent with prior literature, demand frictions are empirically important. This implies that many consumers who can benefit from monitoring do not participate. In Table 5, we show the empirical distribution of both firm-switching and monitoring costs in the population. The average consumer foregoes \$283 of gain by not choosing an outside option from other firms, which is 36% of annual premium (two periods). Monitoring cost is also large and is heterogeneous across consumers. In particular, the average consumer needs to expect a gain of \$93 to participate in monitoring.

Moreover, monitoring disutility increases with private risk.[49] This further accelerates advantageous selection into monitoring, while suggesting that observed renewal prices alone are not enough to explain the empirical selection pattern. At the same time, we see that older and more educated consumers tend to have lower monitoring costs, as well as those with newer cars and better insurance and traffic records.

The right panel of Table A.2 shows that the baseline inertia cost that consumers need to overcome when adjusting coverage is \$134. This adds to firm-switching and monitoring costs to

[47] Our private risk spread is 0.43 ($\exp(\ln \sigma_{\lambda,\text{old}})$), compared to Cohen and Einav (2007)'s estimate of 0.15.

[48] Figure A.9 benchmarks our risk-aversion parameter against the literature. In the graph, risk aversion is interpreted as the indifference value between inaction and taking a 50-50 bet on gaining \$1000 versus losing that value. Barseghyan, Molinari, O'Donoghue, and Teitelbaum (2013), in particular, differentiate between probability distortion (wrong belief about one's own risk) and risk aversion.

[49] Column (2) of table A.3 in the appendix reports the slope parameter for private risk.

Table 5: Latent Parameters

Statistic	Mean	Std. Dev.	Min	Pctl(25)	Median	Pctl(75)	Max
firm-switching inertia (η_{it}, \$)	284	35	158	265	286	307	407
(% annual premium)	36	5	20	34	37	39	52
monitoring disutility (ξ_{it}, \$)	93	19	10	80	93	105	187
(% annual premium)	12	2	1	10	12	14	24
claim risk (λ_{it})	0.05	0.05	0.001	0.02	0.03	0.06	1.48

Notes: This table reports the distribution of heterogeneous latent parameters in our dataset. We simulate a distribution of private risk and calculate these parameters based on our demand estimates.

further prevent safe drivers from opting in to monitoring. Meanwhile, the average consumer only prefers the mandatory minimum coverage by $26 (fixed effect estimate on that plan), which is low given that the plan commands a market share of nearly 50%. This suggests that the rational amount of coverage for many consumers may be below the mandatory minimum, limiting the extent to which monitoring can affect allocative changes across coverage.

To better demonstrate the relative importance of different demand factors, Appendix G calculates counterfactual demand and firm profit by removing the moral hazard effect, reclassification risk, firm-switching inertia, and monitoring disutility, respectively, from consumer demand. Reclassification risk does not strongly influence demand for monitoring due to low risk aversion. All other forces have much stronger impacts.

Welfare calculation Using our demand estimates, we evaluate the impact that the monitoring program has on consumer welfare and on firm profits in the status quo. To do this, we simulate a counterfactual scenario in which monitoring is not available, and compare it to the status quo baseline with monitoring.[50]

To facilitate this exercise, we make three simplifying assumptions. First, we assume that the focal firm's prices do not change if monitoring is not offered. This is substantiated by the fact that the firm did not change insurance prices when it first introduced monitoring in reality (see Appendix B)—nor did it charge different base prices to consumers who opted

[50]As is standard in counterfactual comparisons, we compare the no-monitoring counterfactual against a simulated baseline with monitoring in order to avoid bias from simulation error.

out after monitoring was introduced (see Appendix B and Figure B.2). Second, as the focal firm does not change its prices, we assume that its competitors' prices do not change either under the counterfactual, and so the competing price menus that we observe would continue being offered to consumers in the counterfactual. Finally, since we do not observe the number of external consumers that are served by each competitor, we maintain a *no-brand-differentiation* assumption so that competitors are equivalent to prospective consumers, other than through their price menus.

For each market in our sample, we calculate each consumer's first-period choice probability—for insurance plans and monitoring in the baseline, or just insurance plans in the no-monitoring counterfactual—given her type and menu offerings. In doing so, we obtain annual (ex-ante) consumer welfare as the utility horizon is over two periods. Similarly, we obtain the firm's expected first-period profit. We then simulate claim and monitoring score realizations (if relevant) in each period, pinning down the firm's second-period information set about consumers and the renewal prices charged. This gives us second-period choice probabilities and the annual profit of the firm. Appendix F contains more details.

We take a certainty equivalent approach in calculating ex-ante welfare. Total surplus is the difference between welfare and total expected cost over two periods. Annual profit is given by observed prices over two periods minus the expected cost of covering consumers in that time. We also account for the resource cost for the firm to administer monitoring. This is set at \$35 per monitored consumer, based on interviews with the program manager and on industry estimates, and includes manufacturing, wireless data transmission, depreciation, inventory, and mailing costs as well as R&D, marketing, and other overheads. We do not directly estimate resource costs since observed monitoring prices are constrained by idiosyncratic preferences by managers and regulators that we cannot observe.

Figure 7 plots the results in per-capita per-year terms. The average consumer gains \$11.6 in certainty equivalent, or 1.5% of premium from the availability of monitoring. Profit increases by \$7.9 per capita, a 23.6% increase. Under our symmetric cost and no-brand-preference assumptions, competitors see a profit decline of \$6.2. This isolates the impact of cream skimming by the monitoring firm because the firm can offer lower prices to some monitored consumers despite charging higher markups. The combined total surplus given monitoring increases by \$13.3 (1.7% of premium) over the no-monitoring scenario.

To disentangle the welfare consequence of the moral hazard effect and allocative changes from mechanical monetary transfers across consumers, we first redo the welfare calculation

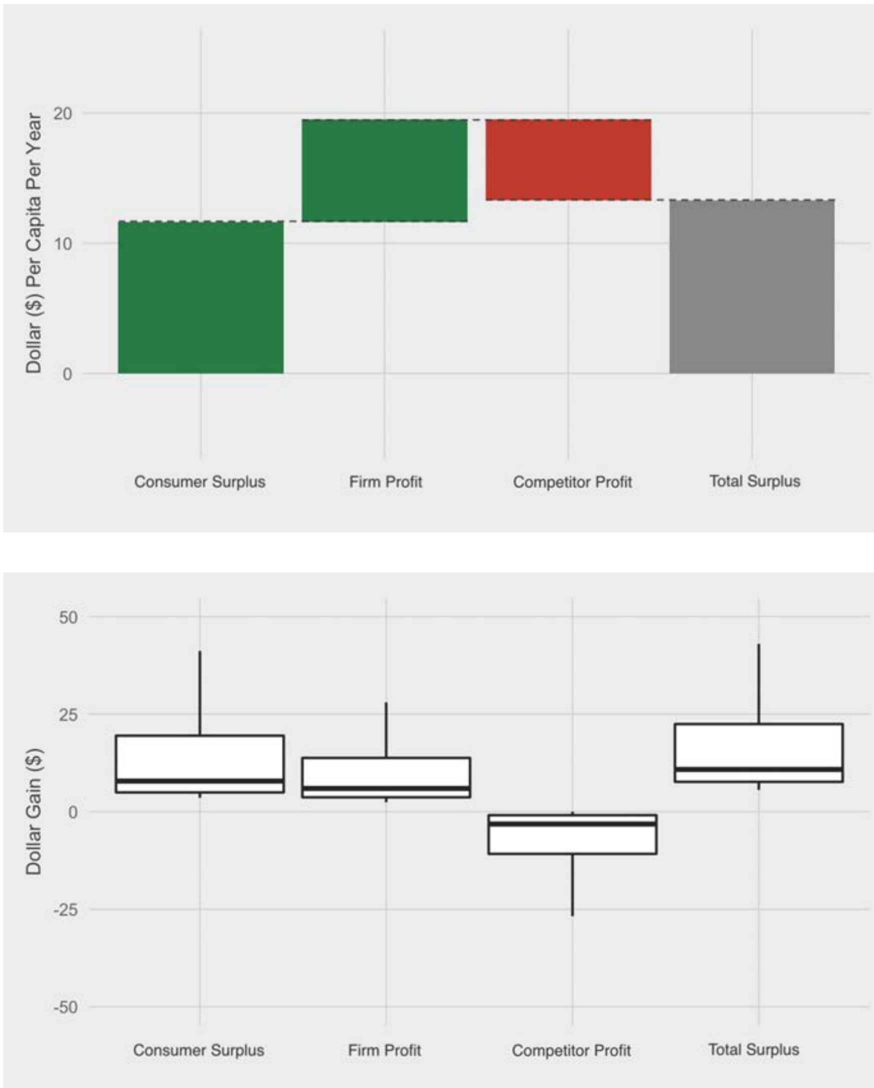

Figure 7: Welfare Calculations

Notes: These figures plot results from our welfare exercise outlined in Section 4. The amount denotes the change moving from a regime where no monitoring is offered to one we observe in the data. We plot the differences in ex-ante certainty equivalent, expected profit (across two-periods) for both the monitoring firm and its competitors, as well as total surplus (welfare minus expected cost). The top graph is a waterfall graph decomposing how the components of total surplus changes. The color green indicates an increase while red indicates a decrease. The box plot show 10/25/50/75/90 percentiles.

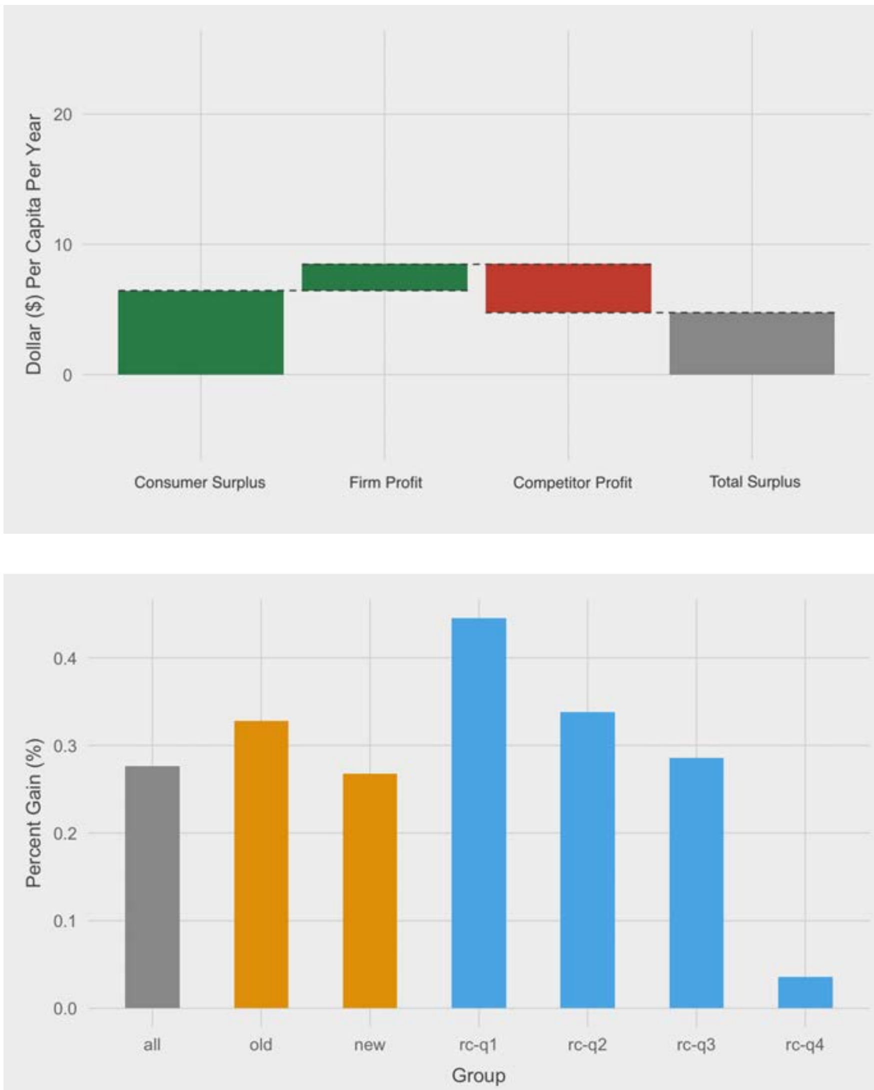

Figure 8: Moral Hazard Effect and Coverage Reallocation

Notes: The top figure plots the same welfare calculation assuming away risk reduction during monitoring based on the incentive effect, per our discussion in the main text. The bottom figure plots average change in coverage amount in percentage across observable groups. "rc-q1" means risk class being in the first quartile at time of choice.

shutting down moral hazard. Consumers' expected utility from monitoring and firms' expected cost for monitored consumers both suffer in this case, reducing the total surplus to \$4.8 per capita. The top panel of Figure 8 plots this effect. The stark decrease suggests that 64% of total surplus gain can be attributed to better driving, implying small allocative efficiency gains. To investigate this further, we look at changes in the quantity of insurance purchased, comparing the observed regime with the no-monitoring one. Because liability insurance is mandatory, this result is entirely due to changes in coverage levels. Overall, insurance coverage increases, but only by 0.28%. Across various observable pools, the safer risk classes stand out despite the fact that they already pay lower premiums. Meanwhile, without risk reduction, overall profit in the industry falls as the monitoring firm offers lower prices to high-performing monitored consumers at the expense of its competitors' profit.

Importantly, as we noted above, our simulations in this section do not allow the availability of monitoring to change baseline firm prices for unmonitored consumers. Any cream-skimming effect would therefore reduce profit in the unmonitored pool as opposed to reduce welfare of unmonitored consumers. While this assumption is true to the data in the status quo, it may not hold in a more general counterfactual setting. In the next section, we propose a model for pricing where the firm can freely surcharge unmonitored consumers.

5 Monitoring Pricing and Equilibrium Implications

In the previous section, we held fixed the price menu that the firm offers to consumers who choose not to opt-in to monitoring. As we note in Section 1.2, this assumption has backing in our data: the introduction of monitoring did not lead to a price increase for consumers who abstained from monitoring. In this section, we argue that the lack of surcharge for unmonitored consumers is an approximate equilibrium outcome driven by the extent of monitoring disutility, switching frictions, and competition in our setting. These factors strongly binds the firm's monitoring pricing decisions in the status quo—so much so that a counterfactual proprietary-data ban (that enforces the portability of consumer data across firms) would have reduced both firm profit and consumer welfare.

True to auto insurers' two-period pricing structure as introduced in Section 1.1, we propose a two-period pricing model focusing on how the introduction of monitoring changes the initial and the renewal prices offered to consumers, as well as how such prices alter the firm's information about consumers' risk types. Specifically, the firm optimally chooses the menu of prices that are offered to each consumer at sign-up depending on whether they opt in to monitoring or not, as well as the renewal discount that is applied after the first period of driving—an "invest-and-harvest" pricing strategy. This model captures a key tension that is raised by our results in Section 4: encouraging consumers to opt in to monitoring increases firm profits, but it requires providing incentives for consumers to overcome their distaste for monitoring and elevated uncertainty.

During the "investment" period, the firm can encourage monitoring opt-in with two main levers: offering discounts for consumers that opt in ("carrot"), or surcharging those that do not ("stick"). However, the firm's ability to wield the latter lever (the "stick") is severely limited by monitoring frictions and competition. Although a surcharge may enhance advantageous selection into monitoring, high monitoring frictions limit adoption and effectively prevent unraveling. As a result, the risk distribution of the unmonitored pool cannot be easily changed, and the benefit from selecting safer consumers may easily be outweighed by ceding market share to other firms.[51]

During the "harvest" period, the firm takes advantage of its informational advantage over competitors and offers monitored consumers a renewal rate that is both more profitable and more appropriate for the consumers' risk types. However, the more that the rate is profitable, the less that consumers stand to benefit from opting in to the program in the first place. The firm thus needs to solve a "rent-sharing" problem with monitored consumers, that is constrained by the "investment"-period prices that encourage consumers to become monitored.

In this section, we synthesize these profit trade-offs and solve the firm's optimal pricing problem given the observed administrative cost of monitoring.[52] We first present the firm's problem holding competitors' prices fixed and compute the optimal opt-in discount, renewal discount, and unmonitored surcharge in the status quo. Next, we allow competitor prices to respond and simulate an equilibrium in which the firm must disclose monitoring data to competitors (enforcing data portability). We find that, given the distribution of demand parameters estimated in Section 3, the firm cannot profit from a large surcharge on unmonitored consumers. Thus, the absence of a surcharge that we observe in practice can be seen as a near-optimal response to the preponderance of high monitoring costs. This leaves the firm

[51] In Appendix Section G, we show that consumers' monitoring disutility and firm-switching inertia are the main factors limiting the adoption of the firm's monitoring program (market share among all consumers in the market). As a result, they also strongly limit the degree to which the monitoring program can "cream-skim" the unmonitored consumer pool.

[52] The prices associated with the monitoring program, as they are observed in the data, are heavily influenced by regulatory pressure as well as managerial and organizational limitations. Since we observe the cost of monitoring to the firm, we can focus on intertemporal profit maximization.

with one main lever for incentivizing monitoring opt-in: high initial discounts. However, we show that the firm's willingness to provide these initial discounts depends on the extent to which it is able to benefit from monitoring ex-post. Thus, while regulations that require public disclosure of data would curb the profitability of the monitoring program, they would also lead to less monitoring data being produced in equilibrium.

5.1 Firm Pricing

In this section, we propose a parsimonious model of auto-insurance pricing that highlights the levers by which the firm can incentivize consumers to opt in to monitoring. Monitoring take-up improves the firm's information on consumers' risk types. This makes the firm's information set endogenous to pricing in our counterfactual simulations. We consider a two-period model in which the firm chooses a set of premium prices to offer each consumer at arrival and at their first renewal, depending on the consumer's observable characteristics and (if relevant) monitoring score. In order to abstract away from regulatory details and conform to the standard approach to program pricing in auto-insurance, we assume that the firm benchmarks its pricing decisions on the existing baseline price rule $p(x)$ for each set of observables x, and differentiates the prices for the monitored and unmonitored pools in each period through multiplicative adjustment factors κ.[53]

In the first period, the firm chooses a multiplicative discount κ_1 for consumers who opt in to monitoring, and a multiplicative surcharge κ_0 for consumers who opt out. That is, if a consumer with observable type x was offered $p(x)$ under the status quo pricing rule, they would be offered $\kappa_1 \cdot p(x)$ if they opt in to monitoring and $\kappa_0 \cdot p(x)$ they opt out. In the second period, the firm chooses a renewal price to offer consumers. For consumers who did not enter monitoring, the firm has no incentive to modify its baseline pricing behavior: it neither has any informational advantage over its competitors, nor does it have reason to nudge them toward a particular plan choice as it did with monitoring in the first period. For consumers who did opt in to monitoring however, the firm does have an informational advantage. A consumer's monitoring score s affords the firm a more precise estimate of the cost to insure them. Thus, for a monitored consumer who is, say, 30% safer than previously

[53] Auto-insurance pricing is highly regulated and subject to a number of external constraints. Firms are required to submit price filings that often span thousands of pages to regulators and face delays in updating them. Adjustments to pricing based on internal programs are typically made through multiplicative factors as with the existing monitoring discount and renewal price. In order to abstract away from regulatory details and conform to the standard approach to program pricing, we compute the firm's baseline price rule $p(x)$ for each set of observables x observed in the data, and parameterize price discrimination between the monitored and unmonitored pools through multiplicative adjustments κ to this price.

expected, the firm may be able to offer a discount that is much smaller than 30% and still be confident that the consumer would not leave for a competitor.

As we discussed in Section 3 (Equation 9), the baseline renewal price offered to a consumer with observables x is given by their first period price $p(x)$ multiplied by a renewal factor $R(C,s) = R^C_{1,C} \cdot R_{0,idt}(s)$ that depends on the consumer's number of claims and monitoring score, which is set to zero ($s = 0$) if the consumer opted out. The wedge between $R(C,s)$ and $R(C,0)$ constitutes the amount of *rent-sharing* between the firm and each monitored consumer that is observed in the data. We model the optimal level of rent-sharing by the choice of a parameter κ_s that scales the existing rent-sharing schedule linearly.

$$p_{m,1}(\vec{\kappa}|x_i,C,s) = p(x_i) \cdot R^C_{1,C} \cdot \begin{cases} 1 & m = 0 \\ \left[1 - \kappa_s \cdot (1 - R_{0,idt(s)})\right] & m = 1 \end{cases}$$

If $\kappa_s = 0$, then the firm keeps all the rent: performance in monitoring has no bearing on renewal pricing. On the other hand, if $\kappa_s > 1$, then the firm shares more rent with consumers than it does in the current regime.

Taken together, the firm's choice of κ_0, κ_1 and κ_s determine both the relative attractiveness of monitoring among the firm's plans, and the firm's competitiveness relative to other firms. A higher surcharge κ_0 encourages consumers to opt in to monitoring and raises profits from unmonitored consumers, but it also discourages unmonitored consumers from staying with the firm. A higher monitoring discount κ_1 lowers direct profits from monitored consumers, but it increases the number of consumers that opt in to monitoring. In this sense, κ_0 and κ_1 constitute the extent to which the firm chooses to *invest* in the information it can learn from monitored consumers. A higher level of rent-sharing κ_s increases the profit that the firm earns from monitored consumers due to its informational advantage but it decreases the probability that those consumers will stay with the firm. In this sense, κ_s constitutes the extent to which the firm is able to *harvest* its investment in the competitive environment.

To choose κ_0, κ_1 and κ_s optimally, the firm maximizes its expected profits across both periods with respect to the anticipated changes in demand and retention across its offerings.[54] We lay out the firm's pricing problem below, suppressing observable notation x for exposition.

[54] As with most factor pricing offered by the firm, we do not allow κ_0, κ_1 and κ_s to vary by coverage choice.

$$
\arg\max_{\vec{\kappa}} \sum_i \int_\lambda \left\{ \sum_{d,m} \underbrace{\Pr\left(f,d,m|\lambda,\mathbf{p_f},\mathbf{p}_{-f};\Theta\right)}_{\text{demand share}} \cdot \left[\underbrace{\left(p^0_{f,d,m}(\vec{\kappa}) - c(\lambda,m) - m\cdot c_m\right)}_{\text{markup}} \right.\right.
$$

$$
\left.\left. + \delta \cdot \mathbb{E}_{C,s|\lambda} \left[\underbrace{\sum_d \Pr\left(f,d,m|\lambda,\mathbf{p_f^1},\mathbf{p}_{-f};\Theta\right)}_{\text{retention rate}} \cdot \underbrace{\left(p^1_{f,d,m}(\vec{\kappa}|C,s) - c(\lambda,0)\right)}_{\text{retention markup}} \right] \right] \right\} g(\lambda) d\lambda
$$

The firm maximizes its total profits across consumers, integrating over the distribution of latent risk types λ for each consumer. For a given choice of $\vec{\kappa}$, the firm expects to earn a first period markup from each consumer consisting of its premium less the costs to insure and monitor the consumer ($c(\lambda,m)$ and $c_m = 35$, respectively) at a chosen plan, multiplied by the consumer's probability of choosing that plan. In the second period, the firm expects to earn a markup governed by the information learned in the first period (claims and the monitoring score if relevant), multiplied by probability that the consumer is retained at a given plan. As $\vec{\kappa}$ changes, the demand and retention shares for each plan adjust according to our model in Section 4.[55] Accordingly, the amount of information revealed by monitoring also adjusts.

5.2 Why is There No Surcharge?

Optimal pricing Under the profit maximizing rule described in Section 5.1, we find that the firm would optimally surcharge the unmonitored pool by 2.7%, but offer a 22.1% up-front discount for opting into monitoring in the first period. This is surprising: as liability insurance is mandatory, our demand model implies that a sufficiently large κ_0 would force all consumers into monitoring if it were not for competition with other firms. Instead, our results suggest that the firm must rely on large discounts κ_1 in order to incentivize consumers to opt-in to monitoring, adding only a modest surcharge for unmonitored consumers. In other words, the firm must make a significant investment from first-period profits in order to elicit the surplus generated by monitoring.

In the renewal period, we find that optimal pricing implies 19.6% less rent-sharing than observed in the data, offering a smaller discount for safe drivers and a smaller surcharge

[55] See Appendix F for more details on the procedures to compute profit under each pricing regime.

for risky ones. This implies that the firm would optimally engage in more aggressive price discrimination conditional on risk. Consistent with our findings in Section 2.3, among monitored consumers, safer ones are only offered a discount by the monitoring firm, and are thus less prone to attrition. Surcharged consumers can avoid the surcharge by switching to a competitor, however, and are thus more price-sensitive.

Overall, the monitoring opt-in rate increases to 4.4% (over the entire enumerated market) under the optimal pricing rule. This constitutes increases in both consumer welfare and market surplus relative to the status quo regime. Although the firm takes a larger share of total surplus, it also generates more surplus by eliciting the creation of more monitoring data in equilibrium. Note that these results assume that the prices offered by the firm's competitors are held fixed. This is because when monitoring data is kept proprietary, changes in the firm's pricing strategy do not have an impact on its competitors' information sets. Competing firms can thus offer only blanket discounts or surcharges in response to changes in the firm's actions. As these blanket price changes would also affect the competitors' existing consumer base—which we do not observe—we do not model it explicitly. However, since the firm's optimal surcharge is already small in the status quo, we expect that competitive equilibrium would be similar. In the next section, we consider a counterfactual in which competing firms' information sets *are* affected by the incentives that the firm offers for monitoring take-up.

5.3 Ban on proprietary data

The optimal pricing regime in Section 5.1 highlights the active investment that the firm must make in order to encourage monitoring take-up. In this section, we consider the equilibrium impact of a ban on proprietary monitoring data ownership. The data produced by the firm's monitoring program is *non-rival* in nature. If the firm were required to disclose it publicly with competitors, those competitors would be able to poach monitored consumers by sharing more rent with them in the renewal period. This diminishes the extent to which the firm can "harvest" its its investment into monitoring take-up.

In order to estimate the equilibrium impact on pricing and take-up for the firm and its competitors, we calculate the fixed point of a Bertrand pricing game that nests each firm's dynamic pricing decisions. Given the information revealed from monitoring take-up at the firm's equilibrium choice of $\vec{\kappa}$, the firm's competitors choose an optimal amount of rent sharing $\kappa_{-f,s}$ to compete for consumers facing renewal. Conversely, the firm's equilibrium

Table 6: Counterfactual Equilibrium Simulations

	Current Regime	**Optimal Pricing**	**Proprietary Data Ban**
Firm Profit	46.5	61.2	49.3
Competitor Profit	149.2	138.2	147.1
Consumer Welfare (CE)	-	+4.7	+2.2
Total Surplus	-	+8.4	+2.9
Monitoring Market Share	3.0%	4.4%	3.4%
Invest			
Unmonitored surcharge	0.0%	2.7%	1.6%
Opt-in discount	4.6%	22.1%	8.3%
Harvest			
Rent-sharing (κ_s)	1	0.80	1.14
Competitor rent-sharing ($\kappa_{s,-f}$)	-	-	1.81

Notes: This table reports results from our counterfactual equilibrium simulations in Section 5. The simulation procedure to calculate welfare, profits, and total surplus is outlined in Section 4. These quantities are reported in dollar per driver per year terms as we translate utility with a certainty equivalent approach. We further enumerate our sample of new customers to the full market by calculating driver weight as in Appendix Section F. The time frame we report is one year (two-period). The level of consumer welfare and total surplus is not identified, so we report only the change in those values in counterfactual regimes compared to the current regime. "Optimal Pricing" represents our equilibrium simulation in Section 5.2. "Data Sharing" represents the equilibrium simulation in Section 5.3, where the monitoring firms is required to share monitoring data to competitors. The "Current Regime" uses monitoring pricing we observe in the data. The rent-sharing parameter (κ_s) is indexed against the one observed in the "Current Regime". Empirically, it is a scalar on top of the firm's existing monitoring renewal schedule. $\kappa_s = 0$ means no rent sharing with consumers (flat pricing schedule regardless of monitoring outcome). $\kappa_s > 1$ means a steeper monitoring discount schedule than observed. This represents more rent-sharing with the consumers.

choice of $\vec{\kappa}$ best responds to its competitors' choice of $\kappa_{-f,s}$.

We make two main assumptions to facilitate this exercise. First, we assume that information sharing is complete and credible. Therefore, firms have symmetric knowledge about the expected cost of monitored consumers, given observables and monitoring scores. Second, the firm's competitors have symmetric profit functions and their action space only consists of setting a single competing rent-sharing schedule $\kappa_{-f,s}$ for monitored consumers. These assumptions abstract away from differences between the competing firms' own knowledge and pricing rules. In this sense, our equilibrium simulation can be thought of as a calibration of how the monitoring firm would expect the market to equilibrate under the proprietary data ban, given all of the information that the firm itself has, but without detailed knowledge of its competitors' proprietary information. Furthermore, the simulation highlights the asymmetric market effect of the policy: while only the monitoring firm invests in eliciting data, the proprietary data ban enables its competitors to collectively "free-ride" by poaching consumers that found to be safer. Finally, while it is possible that competitors may introduce their own monitoring programs, we do not account for this in our simulation, as it would require us to take a stance on a number of operational details that we cannot observe. However, given the rates of monitoring take-up in our data and simulations, we expect that the overall effects on pricing and welfare for our firm's monitoring program would be robust to competitive adoption of a similar monitoring program.

We report the full counterfactual results in Table 6. We find that competitors share 81% more rent than the monitoring firm does under the status quo. Consequently, the firm is forced to share more rent with monitored consumers as well: 14% more compared to the status quo and 43% more compared to the optimal pricing regime. This in turn leads the firm to reduce its opt-in discount to only 8.3%, as well as to reduce the surcharge on the unmonitored pool to 0.8%. Overall, as profit is reallocated across firms, consumer welfare and total surplus decrease slightly compared to the equilibrium without the proprietary data ban (the optimal pricing regime). Ultimately, there is less information in the new equilibrium. The positive impact of curbing ex-post markups is outweighed by the firm's decreased investment in data production. This suggests that a preferable policy for data property rights might instead resemble a *patent* mechanism: short-term property rights that preserve the firm's incentives to produce the data, but long term data sharing that returns control to consumers and curbs markups through competition.

6 Related Literature and Conclusion

Across many markets, firms buy data directly from consumers—with attractive rewards for data-sharing and readily available sensor technologies—and keep what they collect proprietary. The data is then used to mitigate information problems, to gain competitive advantages, and to extract rents from consumers. In our paper, we obtain novel datasets on a voluntary monitoring program in the competitive U.S. auto insurance industry, which give us direct visibility into how proprietary monitoring data is collected and used beyond prior empirical analyses that have relied on simulations (Bordhoff and Noel 2008) and state-firm level aggregates (Hubbard 2000; Reimers and Shiller 2018). We also develop an empirical framework that embeds the firm's data elicitation strategy—and consumers' data-sharing choices—inside the broader product (insurance) market pricing and choice problems. This allows us to account for the value of monitoring data to the firm based on its subsequent use in price discrimination, risk-rating, competitive cream-skimming, and risk reduction.

Our paper builds on a rich literature on measuring the extent and welfare implications of information asymmetries in insurance markets (Chiappori and Salanie 2000; Einav, Finkelstein, and Cullen 2010). Like Cohen and Einav (2007), Barseghyan, Molinari, O'Donoghue, and Teitelbaum (2013), and Jeziorski, Krasnokutskaya, and Ceccarini (2019), we use plan choices and subsequent claims data from an auto insurer to identify consumers' accident risk and risk preferences. Our estimates highlight a similar information asymmetry between the insurer and its consumers, which drives adverse selection. However, our study of the monitoring program—which mirrors many other programs that have since proliferated in the industry—also allows the information environment to no longer be exogenous or fixed, but rather determined endogenously in equilibrium with insurance prices and quantities. To do this, we expand the canonical framework in three ways.

First, monitoring data contains granular information on consumers' accident risk. We incorporate this information by modeling the joint distribution of consumers' monitoring scores with their plan choices and accident claims, which allows us to directly identify a source of private risk (unobserved to the firm pre-monitoring) that drives selection.

Second, the revelation of risk information takes time. When consumers decide whether or not to opt in to monitoring, they anticipate the effect that the monitoring score they receive will have on their renewal price. Our choice model reflects the dynamic nature of monitoring decisions, as well as uncertainty about future premiums that consumers face given their risk types and risk preferences. This allows us to predict how changes in opt-in discounts and renewal pricing alter selection into monitoring (i.e. how many and what kinds of consumers choose to opt in), and consequently, the amount of information that is revealed.

Third, monitoring data is proprietary, and firm pricing must account not only for selection into the program but also for selection into the firm. Accounting for competitor pricing is thus crucial to understanding consumers' outside options and the firm's pricing incentives. Compared to existing studies, we introduce new individual-level competitor pricing data based on price filings; we also expand the canonical model to admit consumer choice along both the intensive and extensive margins.

We also contribute to the empirical literature on the role of data in industrial organization. First, monitoring mirrors more traditional modes of hard information disclosure such as restaurant health score cards (Jin and Leslie 2003) and used-car photos (Lewis 2011). Our study differs conceptually in that consumers have control over information disclosure, albeit heavily influenced by firm pricing. Nonetheless, our findings highlight similar sources of efficiency benefits from richer information as well as the signaling incentives and frictions related to its disclosure. Second, our final counterfactual considers a regulation that would require public disclosure of monitoring data, inspired by recent policy debates regarding data portability and algorithmic transparency (Jin and Wagman 2021). We find that the loss of proprietary control over monitoring data eliminates most of the gains that the firm relies on to recoup the cost of opt-in discounts used to encourage monitoring uptake. This highlights the trade-off between curbing markups and protecting the firm's incentive to produce data in the first place (Posner 1978; Hermalin and Katz 2006).

Finally, our findings provide empirical support for a recent theoretical literature on voluntary disclosure and personalized pricing. Monitoring data is a form of *hard* (verifiable) information that cannot easily be falsified. Although voluntary disclosure of such information by some consumers may lead a monopolist firm (Pram 2020), or one facing horizontal competition (Ali, Lewis, and Vasserman 2019), to infer that others have lower types, consumers can still obtain a Pareto improvement—corresponding to more coverage purchased and hence more risk-sharing in our setting—if there are gains from trade with consumer types that wouldn't be served without the additional information. We also show that consumer control over information sharing—realized through the ex-ante opt-in structure in our setting—mitigates consumer harm from price discrimination. As in Ichihashi (2020) and Montes, Sand-Zantman, and Valletti (2019), we argue that consumer control may thus obviate the need for ex-post regulations on exclusive data ownership.

References

Ali, S Nageeb, Gregory Lewis, and Shoshana Vasserman (2019). "Voluntary Disclosure and Personalized Pricing". In: *NBER Working Paper* w26592.

Barseghyan, Levon, Francesca Molinari, Ted O'Donoghue, and Joshua C. Teitelbaum (2013). "The nature of risk preferences: Evidence from insurance choices". In: *American Economic Review* 103.6, pp. 2499–2529.

Bordhoff, Jason E and Pascal J Noel (2008). *Pay-as-You-Drive Auto Insurance. The Hamilton Project*. Discussion paper 08-09, Brookings Institution, Washington DC.

Chiappori, Pierre Andre and Bernard Salanie (2000). "Testing for Asymmetric Information in Insurance Markets". In: *Journal of Political Economy* 108.1, pp. 56–78.

Cohen, Alma and Liran Einav (2007). “Estimating Risk Preference from Deductible Choice”. In: *American Economic Review* 97.1994, pp. 745–788.

De Hert, Paul, Vagelis Papakonstantinou, Gianclaudio Malgieri, Laurent Beslay, and Ignacio Sanchez (2018). “The right to data portability in the GDPR: Towards user-centric interoperability of digital services”. In: *Computer law & security review* 34.2, pp. 193–203.

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

Einav, Liran, Amy Finkelstein, and Jonathan Levin (2010). “Beyond testing: Empirical models of insurance markets”. In: *Annual Review of Economics* 2.1, pp. 311–336.

Einav, Liran, Amy Finkelstein, Stephen P. Ryan, Paul Schrimpf, and Mark R. Cullen (2013). “Selection on moral hazard in health insurance”. In: *American Economic Review* 103.1, pp. 178–219.

Fama, Eugene F (1980). “Agency Problems and the Theory of the Firm”. In: *Journal of Political Economy* 88.2, pp. 288–307.

Farrell, Joseph and Paul Klemperer (2007). “Chapter 31 Coordination and Lock-In: Competition with Switching Costs and Network Effects”. In: *Handbook of Industrial Organization* 3.06, pp. 1967–2072.

Frankel, Alex and Navin Kartik (2016). “Muddled information”. In: *Working Paper, University of Chicago, Booth School of Business and Columbia University*.

Fudenberg, Drew and J Miguel Villas-Boas (2006). “Behavior-based price discrimination and customer recognition”. In: *Handbook on Economics and Information Systems* 1, pp. 377–436.

Handel, Ben, Igal Hendel, and Michael D. Whinston (2015). “Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk”. In: *Econometrica* 83.4, pp. 1261–1313.

Handel, Benjamin R. (2013). “Adverse Selection and Switching Costs in Health Insurance Markets: When Nudging Hurts”. In: *American Economic Review* No. 17459, pp. 1–48.

Hendren, Nathaniel (2018). *Measuring ex-ante welfare in insurance markets*. Tech. rep. National Bureau of Economic Research.

Hermalin, Benjamin E and Michael L Katz (2006). “Privacy, Property Rights and Efficiency: The Economics of Privacy as Secrecy”. In: *Quantitative Marketing and Economics* 4, pp. 209–239.

Hirshleifer, Jack (1978). “The private and social value of information and the reward to inventive activity”. In: *Uncertainty in Economics*. Elsevier, pp. 541–556.

Holmström, Bengt (1999). “Managerial Incentive Problems : A Dynamic Perspective”. In: *Review of Economic Studies* 66.1, pp. 169–182.

Honka, Elisabeth (2012). “Quantifying Search and Switching Costs in the U.S. Auto Insurance Industry”. In: *The RAND Journal of Economics* 45.4, pp. 847–884.

Hubbard, Thomas (2000). “The Demand for Monitoring Technologies: The Case of Trucking”. In: *The Quarterly Journal of Economics* 115.2, pp. 533–560.

Ichihashi, Shota (2020). “Online privacy and information disclosure by consumers”. In: *American Economic Review* 110.2, pp. 569–95.

Jeziorski, Przemyslaw, Elena Krasnokutskaya, and Olivia Ceccarini (2014). *Adverse Selection and Moral Hazard in the Dynamic Model of Auto Insurance*. Tech. rep. Mimeo.

Jeziorski, Przemysław, Elena Krasnokutskaya, and Olivia Ceccarini (2019). “Skimming from the bottom: Empirical evidence of adverse selection when poaching customers”. In: *Marketing Science*.

Jin, Ginger Zhe and Phillip Leslie (2003). “The effect of information on product quality: Evidence from restaurant hygiene grade cards”. In: *The Quarterly Journal of Economics* 118.2, pp. 409–451.

Jin, Ginger Zhe and Liad Wagman (2021). “Big data at the crossroads of antitrust and consumer protection”. In: *Information Economics and Policy* 54, p. 100865.

Lewis, Gregory (2011). “Asymmetric information, adverse selection and online disclosure: The case of eBay motors”. In: *American Economic Review* 101.4, pp. 1535–46.

Lin, Tesary (2019). “Valuing Intrinsic and Instrumental Preferences for Privacy”. In:

Mailath, George J (1987). “Incentive compatibility in signaling games with a continuum of types”. In: *Econometrica*, pp. 1349–1365.

Milgrom, Paul R (1981). “Good news and bad news: Representation theorems and applications”. In: *The Bell Journal of Economics*, pp. 380–391.

Montes, Rodrigo, Wilfried Sand-Zantman, and Tommaso Valletti (2019). “The value of personal information in online markets with endogenous privacy”. In: *Management Science* 65.3, pp. 1342–1362.

Posner, Richard A. (1978). “The Right of Privacy”. In: *Georgia Law Review* 12.3, p. 393.

Pram, K (2020). “Disclosure, Welfare and Adverse Selection”. In: *University of Nevada, Reno*.

Ptolemus Consulting (2016). *UBI Global Study*. https://www.ptolemus.com/research/ubistudy2016/.

Reimers, Imke and Benjamin Shiller (2018). "Welfare Implications of Proprietary Data Collection: An Application to Telematics in Auto Insurance". In: *Working Paper*.

Reimers, Imke and Benjamin R Shiller (2019). "The impacts of telematics on competition and consumer behavior in insurance". In: *The Journal of Law and Economics* 62.4, pp. 613–632.

Taylor, Curtis R. (2004). "Consumer Privacy and the Market for Customer Information". In: *The RAND Journal of Economics* 35.4, p. 631.

Train, Kenneth E. (2009). *Discrete Choice Methods with Simulation*. Vol. 2, pp. 1–370.

A Additional Figures and Tables

Figure A.1: Examples of Telematics Devices in U.S. Auto Insurance

Notes: These are some examples of the in-vehicle telecommunication (or "telematics") devices used in monitoring programs in U.S. auto insurance. These devices can be easily installed by plugging them into the on-board diagnostics (OBD) port. The OBD-II specification that these monitoring devices rely on has been mandatory for all cars (passenger cars and light trucks) manufactured or to be sold in the U.S. after 1995.

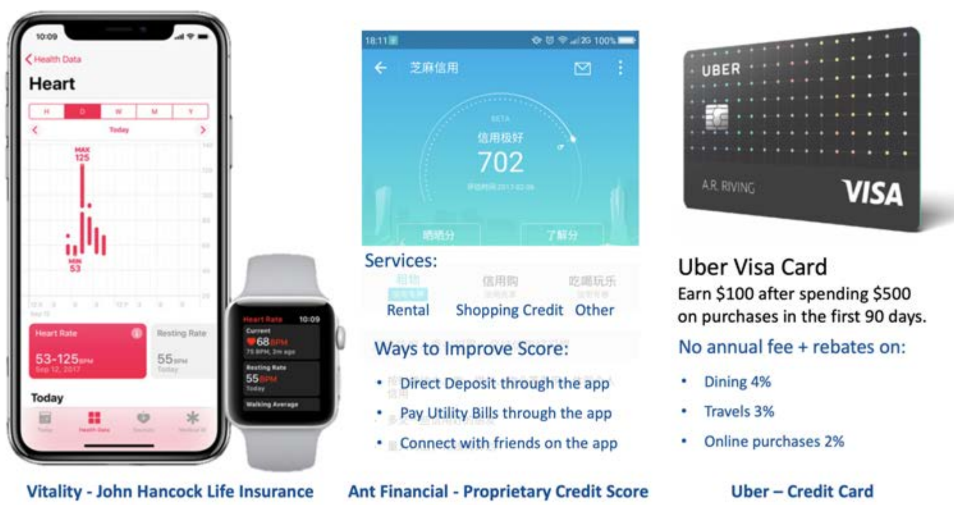

Figure A.2: Other Examples of Direct Transactions of Consumer Data

Notes: Examples of direct transactions of consumer data in other settings. The Vitality program from life insurer John Hancock tracks and rewards exercise and health-related behaviors in exchange for discounts on life insurance premiums. Ant Financial incentivizes users to conduct more personal finance transactions through the platform, such as setting up direct deposit or paying utility bills, in exchange for discounts on various borrowing and rental services. The Uber credit card offered much larger incentives for consumers to use it intensively than the transaction fees charged. One of the plausible business rationales is that the transaction data can be linked back to improve Uber's main businesses in ride sharing and in food delivery.

Table A.1: Summary Statistics on Select Observable Characteristics

Statistic	Mean	St. Dev.	Min	Median	Max
Number of Drivers	1	0	1	1	1
Number of Vehicles	1	0	1	1	1
Calendar month	6.25	3.43	1	6	12
Female Ind.	0.49	0.50	0	0	1
Driver Age	33.42	11.68	15	30	103
Adult Ind.	0.96	0.19	0	1	1
Age <25 Ind.	0.22	0.41	0	0	1
Age <60 Ind.	0.04	0.20	0	0	1
Years of Education	14.46	2.05	9	14	18
College Ind.	0.73	0.44	0	1	1
Post Graduate Ind.	0.41	0.49	0	0	1
Years of License	2.44	1.14	0	3	3
Driver Credit Tier	106	26	0	101	239
Credit Available Ind.	0.96	0.19	0	1	1
Credit Report Ind.	0.83	0.38	0	1	1
Homeowner Ind.	0.17	0.38	0	0	1
Garage Verification Ind.	0.84	0.37	0	1	1
Out-of-state Ind.	0.11	0.32	0	0	1
Population Density Percentile	51	21	0	54	99
Vehicle Model Year	2006	6.05	1928	2007	2018
Vehicle on Lease Ind.	0.51	0.50	0	1	1
Length of Ownership	0.42	0.92	0	0	4
Class C Vehicle indicator	0.89	0.31	0	1	1
ABS Ind.	0.13	0.34	0	0	1
Safe Device Ind.	0.35	0.48	0	0	1
Accident Point	1.53	2.80	0	0	82
At-Fault Accident Count	0.33	0.65	0	0	11
DUI Count	0.05	0.23	0	0	8
Clean Record Ind.	0.64	0.48	0	1	1
Prior Insurance - Some	0.08	0.27	0	0	1
Prior Insurance - Yes	0.57	0.49	0	1	1
Length of Prior Insurance	1.59	1.45	0	2	4
Zipcode AGI ('$000)	142	162	1	114	100,508

Notes: Our data only consist of single-driver-single-vehicle insurance policies. Years of license data is capped at 3 in compliance with regulations that limit risk rating. Zipcode AGI is merged into the dataset by researchers based on zipcode.

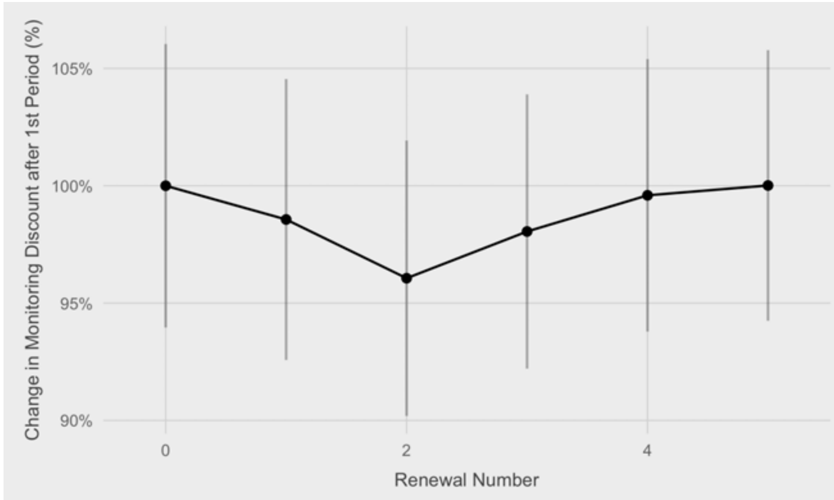

Figure A.3: Persistence of Monitoring Discount

Notes: This graph plots the empirical progression of monitoring discount for all monitoring finishers in one state that stayed with the Firm till at least the end of the 5th periods (so we observe monitoring discount in the renewal quote for the 6th period). The benchmark is monitoring discount in the first renewal quote ($t = 0$). Fluctuations and noises are due to ex-post adjustments. Monitored drivers can report mistakes in their records and have their discount adjusted.

Figure A.4: Renewal Price Claim Surcharge

Notes: This graph plots the empirical claim surcharge function for at-fault accidents. Claim surcharge varies with existing violation points and calendar time. 0.1 means 10% surcharge. This differs from the filed factors because the latter is applied on the base rate only, while this function represents the surcharge percentage on top of overall premium. This is done by regressing renewal price change on violation point last period and current period at-fault claim, controlling for all other observables.

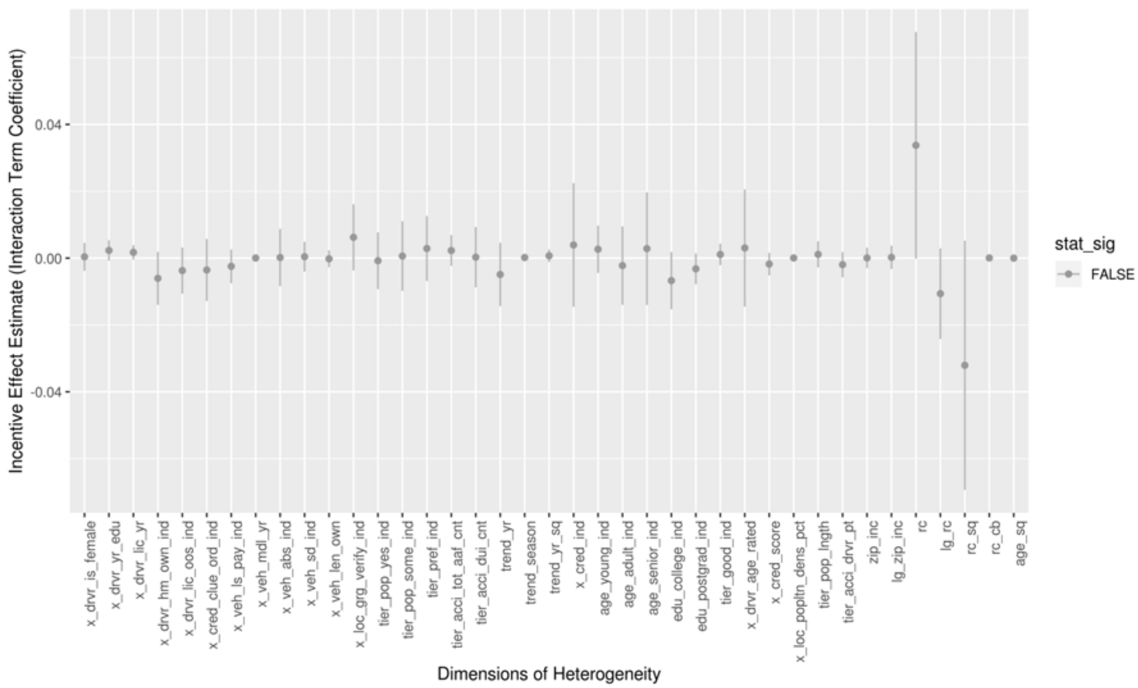

Figure A.5: Heterogeneity in Incentive Effect across Coverage Choice

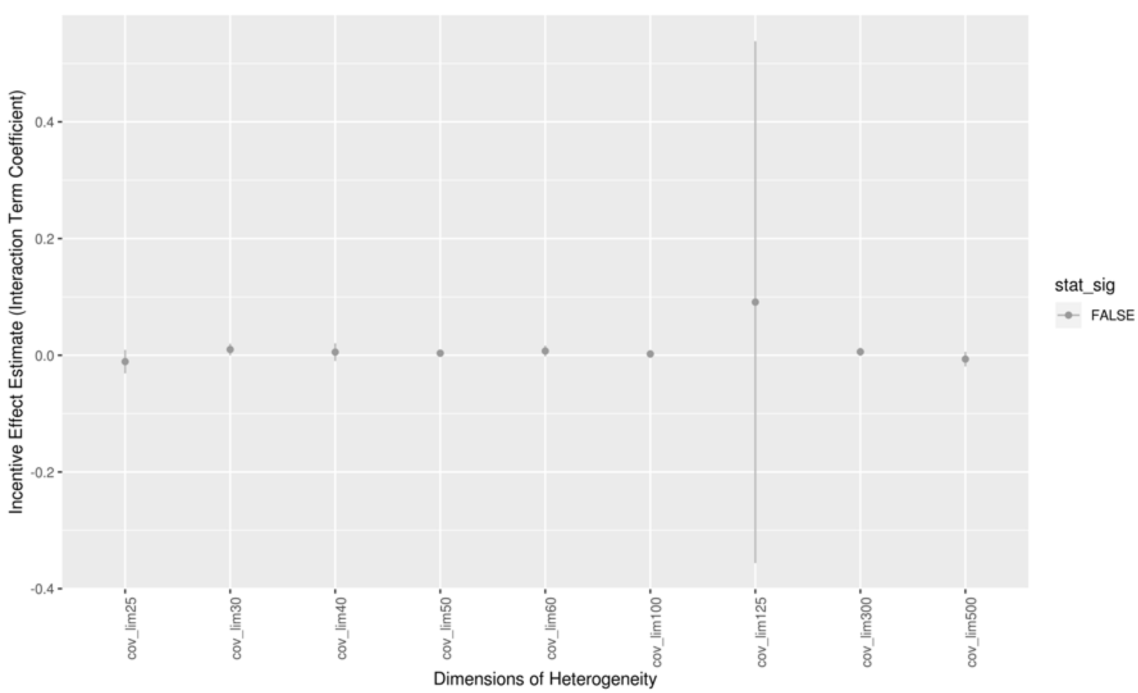

Figure A.6: Heterogeneity in Incentive Effect across Observables

Notes: These figures plot the estimated coefficients $\hat{\theta_{\mathbf{mh,x}}}$ in Equation (4) as well as the corresponding 95% confidence intervals. A positive coefficient means that drivers with higher values (or a 1 in the case of binary variables) in the variables listed in the horizontal axis saw higher claim increase after monitoring, hence have larger incentive effect. 'rc' means risk class.

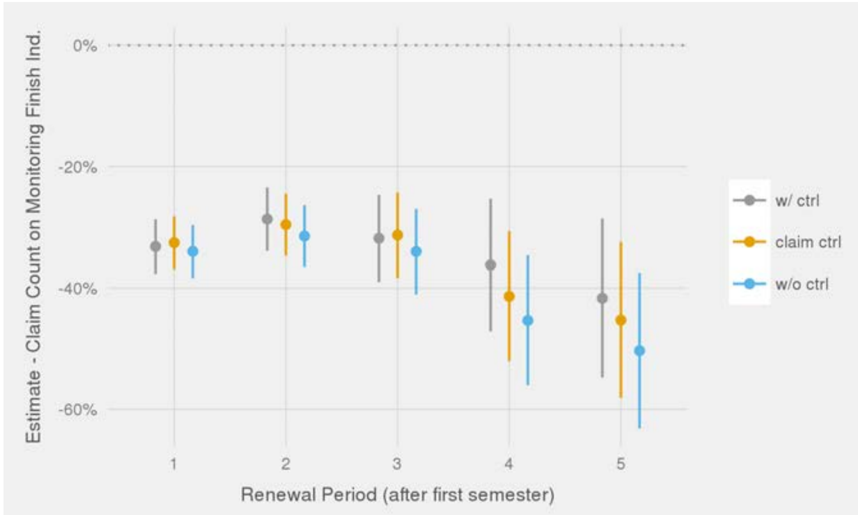

Figure A.7: Estimates - dynamic informativeness of monitoring participation

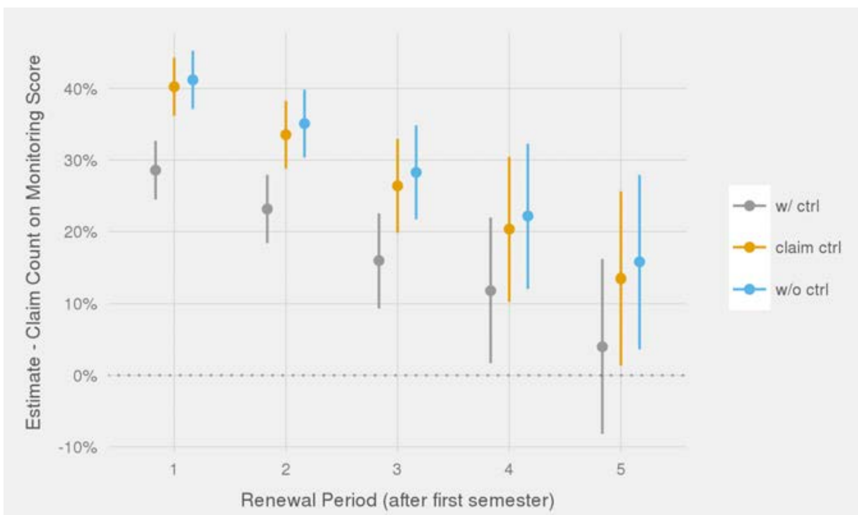

Figure A.8: Estimates - dynamic informativeness of monitoring score

Notes: Figures A.7 and A.8 report the estimate for θ_t and γ_t from regression (5) in percent increase terms. Monitoring participation is an indicator for finishing monitoring. For each $t > 0$, we take all drivers who stayed with the Firm till at least the end of period t. θ_t is the coefficient of claim count of driver i in period t on monitoring score of i, and γ_t is that on monitoring finish indicator of i. Monitoring score is normalized, and defaulted as 0 for unmonitored drivers. So θ_t measures the effect of getting a score one standard deviation above the mean during the monitoring period ($t = 0$). γ_t compares unmonitored drivers with the average monitoring finisher. To further translate these effects into percent increase terms, we divide the estimate of θ_t and γ_t by the average claim count in period $t$of all *monitored* drivers. The horizontal axis represents different regressions for different renewal period $t > 0$. Different colors within each t value represent different specifications of control variables (x_{it}). The grey (left-most) series represents estimates from regressions with the full set of x_{it}; the orange (middle) one includes only claim records revealed since $t = 0$; the blue (right) series includes no control.

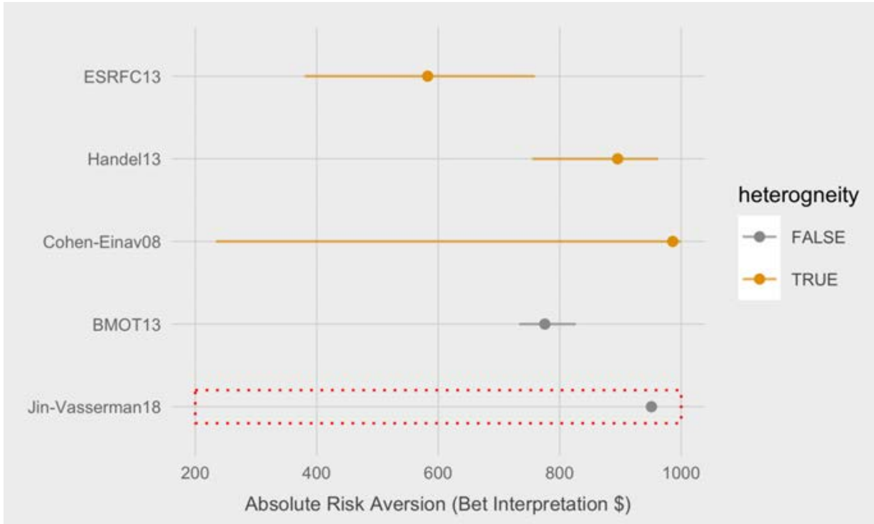

Figure A.9: Risk Aversion Parameter Estimates - Benchmark

Notes: This figure benchmarks our risk aversion parameter estimate to the literature. Heterogeneity indicator means that the author allows risk aversion to vary across people, in which case we plot the range of risk aversion parameters in the population. Otherwise we plot the 95% confidence interval of the homogeneous risk aversion parameter. (Figures A.7 and A.8 on the next page)

Table A.2: Estimates: Homogeneous Parameters

Cost		Score & Pricing		Demand	
$\ln \sigma_{\lambda,\text{new driver}}$	-0.266^{***}	$\ln \sigma_s$	-0.081^{***}	$\ln \gamma$	-9.235^{***}
	(0.060)		(0.007)		(0.089)
$\ln \sigma_{\lambda,\text{old driver}}$	-0.840^{***}	$\beta_{R,\text{new}}$	66.953^{***}	η_0	134.262^{***}
	(0.070)		(0.403)		(2.228)
$\ln \alpha_\ell$	-1.480^{***}	$\beta_{R,\text{monitoring}}$	59.680^{***}	σ_ζ	98.989^{***}
	(0.063)		(0.902)		(2.303)
		$\beta_{R,\text{renw}}$	78.571^{***}	σ	39.213^{***}
			(0.315)		(0.632)

Note: This table reports estimates for homogeneous parameters of our structural model. *Cost*: spread of private risk $\sigma_{\lambda,\text{new driver}}$ and $\sigma_{\lambda,\text{old driver}}$ (new drivers are defined as those licensed in the past three years), claim severity Pareto distribution parameters ℓ_0 and α_ℓ (ℓ_0 is set at $3,000 per discussion in the text). *Score and Pricing*: monitoring score's signal precision σ_s, rate parameters for the renewal price change (R_0) Gamma distribution β_R's. *Demand*: absolute risk aversion coefficient γ, baseline inertia η_0 in dollar term, variance of own firm random coefficient σ_ζ, scale of the logit error σ. *p<0.1; **p<0.05; ***p<0.01 *p<0.1; **p<0.05; ***p<0.01

Table A.3: Estimates: Heterogeneous Latent Parameters

	Log Claim Rate (μ_λ)	Monitoring Disutility (ξ/\$)	Firm-switching Inertia (η/\$)
Intercept	-3.294^{***}	96.773^{***}	228.559^{***}
	(0.080)	(2.813)	(6.213)
Private Risk		25.238^{***}	
		(1.657)	
Monitoring Ind.	0.404^{***}		
	(0.063)		
Monitoring Duration	-0.796^{***}		
	(0.081)		
Driver			
Driver Age	-0.240^{***}	-1.049^{**}	4.526^{***}
	(0.053)	(0.437)	(1.641)
– Square	0.156^{***}	-1.047^{***}	3.816^{**}
	(0.055)	(0.309)	(0.742)
Age < 25	0.081^{**}	0.326	−0.500
	(0.032)	(0.339)	(0.922)
Age > 21	−0.064	−0.059	3.195^{***}
	(0.053)	(0.403)	(0.449)
Age > 60	−0.046	−0.139	−0.275
	(0.068)	(1.689)	(0.340)
Year of Education	0.001	-2.452^{***}	-7.526^{***}
	(0.025)	(0.331)	(0.915)
College Ind.	−0.00001	-0.952^{***}	0.234
	(0.038)	(0.339)	(0.237)
Post Grad Ind.	0.005	−0.728	−1.547
	(0.039)	(1.644)	(1.686)
Female Ind.	0.099^{***}	−0.261	1.007
	(0.021)	(1.643)	(1.686)
Driver License Year	−0.018	−0.016	16.776^{***}
	(0.019)	(0.905)	(0.338)
Home Ownership	−0.020	−0.039	0.058
	(0.038)	(0.447)	(1.653)
Out-of-State License	-0.104^{***}	−0.380	−0.406
	(0.030)	(0.339)	(0.922)
Location			
Garage Verified Ind.	-0.069^{*}	0.008	1.847^{**}
	(0.036)	(0.521)	(0.922)
Population Density	0.076^{***}	0.359	-4.902^{***}

	μ_λ	ξ/\$	η/\$
	(0.015)	(0.419)	(0.445)
Zipcode Income	-0.058^{***}	0.610	-2.936^{*}
	(0.017)	(1.615)	(1.677)
Log Zipcode Income	0.031^{***}	0.284	-0.808
	(0.008)	(2.949)	(1.850)
Vehicle			
Length of Ownership	0.017	-0.918	-0.084
	(0.012)	(0.887)	(0.338)
Vehicle on Lease Ind.	0.092^{***}	-1.058	4.789^{***}
	(0.024)	(1.677)	(0.343)
Model Year	-0.026^{*}	-1.621^{***}	3.211^{***}
	(0.014)	(0.421)	(0.445)
ABS Ind.	-0.058^{*}	0.034	-1.626^{***}
	(0.035)	(0.741)	(0.422)
Airbag Ind.	0.014	0.199	1.225
	(0.021)	(1.644)	(1.686)
Class C Ind.	0.023	0.079	3.843^{**}
	(0.053)	(0.448)	(1.655)
Tier			
Credit Report Ind.	0.044	0.414	1.832^{***}
	(0.035)	(0.429)	(0.448)
Delinq. Score*	-0.016	2.114^{***}	10.959^{***}
	(0.014)	(0.331)	(0.917)
Prior Ins. Length	-0.038^{**}	-2.293	-3.993^{***}
	(0.017)	(1.648)	(0.338)
Has Prior Ins.	-0.067^{*}	-1.183^{***}	-0.759^{*}
	(0.035)	(0.427)	(0.448)
– w/ Lapse	-0.050	0.204	0.001
	(0.043)	(1.686)	(0.620)
Violation Points	-0.032	1.084^{***}	4.333^{***}
	(0.030)	(0.337)	(0.429)
Clean Record Ind.	-0.097^{***}	-0.909	-1.392^{***}
	(0.035)	(0.916)	(0.342)
Total Accident Count	0.115^{***}	0.470	-0.139
	(0.029)	(1.638)	(1.690)
Total DUI Count	-0.233^{***}	0.031	0.326
	(0.065)	(0.922)	(0.536)
Log Risk Class	0.275^{***}		
	(0.046)		
Risk Class	0.042		
	(0.074)		

	μ_λ	ξ/\$	η/\$
– Square	-0.124^{*}		
	(0.073)		
– Cube	0.0002		
	(0.046)		
Seasonality	0.026^{**}	-0.764^{**}	-1.585^{***}
	(0.011)	(0.331)	(0.427)
– Square	0.063	-0.364	-0.519
	(0.046)	(0.340)	(0.430)
Trend Year	0.083^{*}	-1.570	7.417^{***}
	(0.043)	(1.660)	(0.338)
– Square	-0.102^{***}	-1.413	6.199^{***}
	(0.039)	(1.830)	(1.674)

Note: This table reports intercept and slope estimates for heterogeneous latent parameters. Continuous covariates are normalized (except private risk and monitoring duration). Discrete variables are normalized so that the lowest level is zero. "Deliq. (delinquency) Score" is based on records from a credit bureau. Higher scores mean worse records. *p<0.1; **p<0.05; ***p<0.01

Table A.4: Estimates: Renewal Pricing and Monitoring Score

	$\mathbb{E}[R_{0,m=0,t=0}]$	μ_s	$\mathbb{E}[R_{0,m=0,t=1}]$
Intercept	-0.362^{***}	11.367^{***}	-1.131^{***}
	(0.001)	(0.506)	(0.132)
Log Risk Class	-0.413^{***}	-0.384^{**}	-0.080^{***}
	(0.018)	(0.155)	(0.018)
Risk Class	0.367^{***}	-0.077	0.063
	(0.051)	(0.304)	(0.034)
– Square	-0.290^{***}	0.245	-0.155^{***}
	(0.054)	(0.308)	(0.036)
– Cube	-0.229^{***}	-0.039	0.031
	(0.022)	(0.140)	(0.019)
$\ln\lambda$		1.859^{***}	
		(0.094)	
log(Monitoring Score)			0.150^{***}
			(0.005)

Notes: This table reports estimates for the renewal pricing and monitoring score model. Instead of modeling the Gamma shape parameters (α), we use a change-of-variables technique to directly estimate the expected renewal rate. It is modeled with a Sigmoid function between 0.5 (50% cheaper) and 2 (twice as expensive). That is, $\mathbb{E}[R_0] = \sigma(\mathbf{x}'\theta_R) \times 1.5 + 0.5$. We include the appropriate Jacobian adjustments in estimation, and winsorize away extremely large or small renewal price change. *p<0.1; **p<0.05; ***p<0.01

Table A.5: Demand Model Fit

	Basic Specification	**Primary Specification**	**Data**
Monitoring share (when eligible)	17.7%	15.6%	15.3%
Expected score	5.46	4.25	4.30
Selection effect (risk)	6.7%	21.2%	-
Coverage share			
30K	13.7%	12.5%	12.7%
40K	9.1%	8.2%	8.5%
50K	53.2%	49.8%	47.1%
100K	13.0%	15.4%	17.0%
300K	9.3%	11.9%	12.3%
500K	1.8%	2.3%	2.4%
First renewal attrition (indexed)	133.0%	102.9%	100.0%

Notes: This table reports the fit of our demand model as described above. The primary specification is outlined in our econometric model section. Monitoring share is conditional on eligibility. For coverage shares, our demand estimation data pools across three states with different mandatory minimum. One state changed mandatory minimum from 30K to 50K; estimation data is drawn from only the pre-period of that state to capture monitoring introduction. First renewal attrition rate is benchmarked to data per the firm's request (reporting percent differences, not percentage point differences).

Table A.6: Cross Validation

	Basic Specification	**Primary Specification**	**Hold-Out Data**
Monitoring share (when eligible)	21.2%	17.9%	17.6%
Expected score	5.23	3.97	4.17
Selection effect (risk)	5.2%	23.7%	-
Coverage share			
30K	-	-	-
40K	9.4%	7.6%	7.2%
50K	66.3%	60.5%	58.1%
100K	13.4%	17.5%	19.6%
300K	9.7%	10.9%	12.8%
500K	1.3%	3.6%	2.4%
First renewal attrition	132.2%	104.2%	100.0%

Notes: This table reports our cross-validation result. All measures are calculated analogously as Table A.5. For the state that changed mandatory minimum, the hold-out data include all post-period data. For the other two states, the hold-out data include all observations that are not in our demand estimation data.