Re-examining Moral Hazard under Inattention:

New Evidence from Behavioral Data in Auto Insurance

Yizhou Jin*

August 15, 2023

Abstract

This paper uses novel sensor data to study drivers' risky phone use behavior. The results challenge the conventional wisdom of moral hazard in insurance. We first identify handheld phone use behavior ("HPU") and quantify its causal impact on accident likelihood ("riskiness") using exhaustive fixed-effect models. We then find HPU to be risky but insensitive to both insurance coverage changes and weather shocks that increase its riskiness. This contradicts the prevailing theoretical prediction and empirical studies that have thus far relied on claims data alone. On the other hand, an experiment with a one-time text-message warning led to a persistent 15% HPU reduction. Drivers' inattention to risk thus limits moral hazard.

Keywords: Moral Hazard, Inattention, Auto Insurance, Tort Law, Risky Driving, Telematics

JEL Codes: L14, K13, G22, I12

*Jin: Rotman School of Management and Department of Economics, University of Toronto, Toronto, ON M5S3E6, Canada; contact: yizhou.jin@utoronto.ca +1 (650) 339-3386. The first version of the paper was titled "How to Prevent Traffic Accidents: Moral Hazard, Inattention, and Behavioral Data" and was dated November 18, 2020. Thomas Yu provided excellent research assistance. I thank Ben Handel, Jonathan Kolstad, Ariel Pakes, Steve Tadelis, Tom Baker, Kevin He, Jing Li, and anonymous referees from EC'23 for their valuable comments and helpful suggestions.

Introduction

The idea of moral hazard—that risk-sharing such as providing insurance coverage would incentivize excessive risk-taking and create negative externalities—has been a central concern since the genesis of insurance economics (Ducat 1865; Arrow 1968; Shavell 1979; Baker 1996). It has also profoundly influenced the design of tort law and liability insurance systems, particularly with regard to motor accidents.[1] Without moral hazard, accident liability would have no "deterrence effect" on risky driving, and different systems become little more than risk and monetary transfers between insurers, victims, and injurers (James Jr 1947; Shavell 2018).

However, empirical tests of moral hazard often focus *not* on individuals' risk-taking behaviors but on the cost to insurers alone. Not only does this falsely equate insurance utilization with the underlying risk—particularly salient when the average coverage is as low as auto insurance in the U.S.[2]—it also ignores the externalities borne by non-insurers such as car-accident victims who suffer losses that are almost never fully compensated.[3]

Two factors contribute to this empirical distortion. First, the insurance sector and actuarial scholars often do not distinguish insurers' exposure to risk from the risk itself. Second, studies that do seek to separate moral hazard in risk-taking behaviors—often referred to as "ex-ante" moral hazard—do not observe those behaviors. Instead, they must make inferences based on insurance claims or police reports (Levitt and Porter 2001; Weisburd 2015; Jeziorski et al. 2017).[4] This data limitation necessitates strong assumptions on accident reporting and behavioral rationality that have been challenged by recent evidence on the effectiveness of simple nudges (Choudhary et al. 2019).

Advances in data and AI technology have resolved the behavioral data bottleneck in practice: sensor data have proliferated with the ubiquity of smartphones, which combines with machine learning to allow accurate identification of risky driving behaviors.

In this paper, we collaborate with a transportation company that provides on-demand car rides and self-insures against traffic accidents its drivers incur. Using drivers' smart-

[1]While liability insurance can diminish accident deterrence as "the potential tortfeasor relaxes in his efforts to take due care in order to avoid harm," others "did not believe in the idea of prevention" and argued that low coverage benefited insurers at the expense of accident victims (Wagner 2006).

[2]Mandatory minimum auto liability coverage in U.S. states range from $25 to 100 thousand, compared to GBP 1 million in the U.K., EUR 6 million in Italy, and CAD 10 million in British Columbia.

[3]U.S. auto insurers paid $196 billion in claims in 2011 while the total cost of accidents exceeded $1 trillion adjusted for inflation (Blincoe et al. 2015).

[4]Similarly, Spenkuch (2012) points out that "pure measures of self-protection are rarely available" in the context of health insurance, which contributes to why "'ex-ante moral hazard' has received very little attention in empirical work from the literature" (Einav and Finkelstein 2018).

phone sensor data, we can accurately identify handheld phone use ("HPU") behavior. We use this behavioral data to re-examine the degree of moral hazard and the role of inattention.

Our main result is that moral hazard in risky driving is severely limited by inattention. This challenges the conventional wisdom and empirical findings based on claims data; it also highlights a broader source of inefficiency in auto insurance contracts: features that limit insurance utilization (deductibles, for example) or punish accidents harm risk-sharing without inducing safer driving. Specifically, we find HPU to be risky but insensitive to both insurance coverage reduction and weather shocks that increase its riskiness. However, a one-time text-message warning on HPU led to a persistent 15% HPU reduction.

We first introduce our behavioral data and provide background information. HPU is identified from high-frequency smartphone sensor records based on AI algorithms. We demonstrate that HPU significantly increases the likelihood of accidents and safety-related passenger complaints. This is based on a linear regression model with exhaustive ride-level fix effects. We address potential reverse causality issues by removing all HPU records post accidents in affected rides. Moreover, across-driver heterogeneity accounts for 68% of all HPU variation based on a simple variance decomposition. This suggests that a given driver's HPU may be more stable than implied by full rationality.

Our main analyses begin by laying out a simple rational framework of HPU to define moral hazard formally and to understand the relationship between ex-ante and ex-post moral hazard. We then empirically test it using exogenous weather events and insurance coverage change. Overall, we find evidence rejecting the rational benchmark. First, precipitation strongly heightens the riskiness of HPU but generates minute deterrence effects. This holds true for both across- and within-trip precipitation changes. Second, we study state-level regulatory changes that reduced drivers' auto insurance coverage. We use a synthetic control approach and find no detectable HPU change in response to the insurance reduction. Further, we detect no differential HPU response when comparing drivers that were differentially impacted by the insurance reduction. We thus argue that moral hazard is weak; exposing drivers to more risk does little to discipline and deter their risky behavior.

We then conduct an experiment to explore the role of inattention and the effect of nudges. We sent a simple text message to a random subset of drivers with high HPU. The message brings attention to their behavior. It also informs them of a small added incentive: potential suspension when they receive passenger complaints about unsafe driving. The treatment effect is large and persistent: HPU drops by a third on the first day, stabilizing at 15% after a month. Meanwhile, no change is detected in driving hours or in other types of risky behavior, such as harsh acceleration. To isolate the role of inattention, we note that the experiment imposes no added incentive for non-passenger trips (drivers en route to pick up passengers). Yet we find a large drop in HPU nonetheless. Among passenger trips, treatment effects are larger with precipitation and among drivers that stand to lose more in suspension. This indicates that attention triggers moral hazard as nudges catalyze drivers' responsiveness to risk.

Related Literature This paper primarily contributes to the empirical literature on moral hazard by introducing new data that directly reveals individuals' "ex-ante" risk-taking behaviors. This enriches the literature in four ways.

First, our findings provide important context to understand previous studies showing that claim probabilities respond to incentives in auto insurance. Chiappori and Salanie (2000) and Abbring et al. (2003) first provide evidence of moral hazard using correlations between insurance claims and coverage choices. To overcome the confounding effect of adverse selection, Weisburd (2015) uses an instrumental variable approach with coverage assignment in employer-sponsored plans. Alternatively, Abbring et al. (2008), Jeziorski et al. (2017), and Jin and Vasserman (2019) turn to panel datasets and measure how claim probabilities change in response to mechanical shifts in a driver's risk class and monitoring status. However, many accidents are unreported, especially with deductibles, premium reclassification, and low liability coverage. This is often addressed by looking at insurance claims that are less prone to under-reporting, such as multi-vehicle collision and large liability claims. Our results challenge the effectiveness of this approach and suggest that claim probabilities may measure extensive-margin insurance utilization more than it does accident frequency.

Second, by focusing on accident prevention rather than insurance utilization, our examination of moral hazard is consistent with its broader "hidden-action" interpretation. This highlights its central role not only in managing insurance spending—which is a relatively recent concern (Finkelstein 2014)—but in the design of tort law and insurance more generally.[5] "At the beginning of the nineteenth century, liability insurance would have been unthinkable" and "considered as *immoral*" exactly because it resolved the potential tortfeasors from the consequences of accidents and discouraged exercising care

[5]The historical origin of moral hazard emphasizes perverse incentives for behaviors that heighten the probability of risk (Rowell and Connelly 2012). For example, "heavy insurance also increases the moral hazard, by developing motive for crime;" (Ducat 1865) "...an individual may choose to invest less preventive effort than otherwise to avoid the insured outcome." (Shavell 1979)

to avoid them (Knapp 1983). The subsequent expansion of both liability and no-fault insurance was accompanied by a diminished view of the ability of tort punishment to deter accidents (Wagner 2006). In the second half of the 20th century, Scandinavian lawmakers championed a tort system with strict liability as well as high liability insurance, arguing that "there is no empirical evidence available regarding the comparative merits of either [the Scandinavian or traditional tort liability] system for the prevention of accidents" (Hellner 2001), leading to a dramatic difference in auto insurance mandates between Europe and the U.S. that persists to this day.[6]

Third, we highlight the empirical relevance of inattention and bounded rationality in mitigating moral hazard. Our findings are consistent with the dual-process theory in the psychology of attention (Miller 1956; Stanovich 2009; Kahneman 2003; Gabaix 2019). Drivers must process a myriad of navigation, earning, and risk factors at any given time, but they can only consciously monitor few. This results in inattention to risk, which amplifies risky behavior and reduces drivers' sensitivity to risk, both of which are validated by our empirical tests and are consistent with lab experiments on risky driving (Strayer and Johnston 2001). Moreover, in Section 2, we discuss how ex-ante and ex-post moral hazard are related. Our results are thus close in spirit to Brot-Goldberg et al. (2017), which reveals that consumers' heuristic-based price response drives ex-post moral hazard in health insurance utilization.

Lastly, in the context of digitization, we are related to two papers comparing Uber to traditional taxis. Liu et al. (2018) find that trip monitoring and the rating mechanism limit moral hazard in detours. Athey et al. (2019) study similar risky behaviors as us. They find that nudges, rider preference, and Uber's rating system can partially explain reduced risky driving compared to taxis. Our work covers similar channels, but the main focus is instead on driver preference, accident impact, and insurance design.

1 Background, Data, and Stylized Facts

1.1 Data

Telematics and Behavior Identification Our telematics data is derived from drivers' cell phone sensor records, the collection of which is necessary for the transportation company to dispatch its drivers to passengers. Table 1 summarizes most key variables

[6]See footnote 2.

in our analysis.[7] The average driver is 43 years old and male. The average trip is 4.4 miles long and lasts for 16.8 minutes. The handheld phone use (*"HPU"*) frequency in the average trip is 13.9 meters per mile.

Most phone movements are not HPU. Many drivers keep their phones in cup holders or on dashboards. Disruptive vehicle movements such as going through speed bumps or collision accidents can also cause large sensor data fluctuations. As a result, the telematics event identification requires that the phone orientation be aligned with the car movement, that the car be traveling at nonzero speed, and that phone movement instability be detected for a nontrivial time duration. Sophisticated machine-learning techniques and validation datasets are used to improve event detection. Although our data usage agreement does not permit disclosure of proprietary algorithms, our algorithm achieves more than 99% precision and over 75% recall in a validation test using independently labeled events. The focus on high precision is to minimize potential liabilities associated with unnecessary or even harmful interventions. More details on the machine learning algorithm are provided in Appendix B.

Our analysis focuses on HPU for three reasons. First and foremost, our HPU identification is much more accurate and precise than other risky driving behaviors. This is because the detection of HPU is almost entirely unaffected by environmental factors such as speed limits and weather events (Figure A6(c)). In contrast, behaviors such as driving-under-the-influence ("DUI") or speeding require us to accurately identify the vehicle and driver movement in the context of instantaneous environmental constraints. Relatedly, our training and validation datasets are proprietary and labeled based on test drives and video footage of test drivers, and the labeling of HPU is the most unambiguous and error-free. Lastly, HPU has been widely studied in lab environments and, due to its rising frequency over the last two decades, frequently targeted for safety campaigns and intervention (Strayer and Johnston 2001; Wilson and Stimpson 2010; Llerena et al. 2015; Megat-Johari et al. 2023).

Precipitation Almost a third of all trips in our data experience precipitation, partially reflecting increased passenger demand in such weather conditions. We collect detailed and real-time data on precipitation type and intensity from Darksky's public API. We merge our datasets based on the geohash12 codes and timestamps at the start and the end of each trip. Table 1 includes summary statistics of our precipitation data.

Insurance The transportation company self-insures its drivers and offers no choice in coverage. The insurance cost is passed onto the drivers as a per-mile charge that is

[7] We cannot provide accident statistics per data usage agreement.

actuarially fair based on historical claim records. All drivers enjoy close-to-unlimited coverage for third-party liability but face a high deductible for repairing property damage to their own cars in at-fault accidents. There is additional coverage for drivers' own medical payments or when the at-fault counter-parties of an accident fail to pay their liability in a timely manner. This coverage varies by state and time, which we exploit in Section 3.2.

1.2 Road Accidents and Cellphone Use

Handheld cellphone use while driving is outlawed in 19 U.S. states while texting while driving is outlawed in all states (NHTSA 2019b). However, the law is poorly enforced. Accidents without DUI or speeding are defaulted as simple negligence, and even severe accidents rarely lead to trials. As a result, only 3% of first-respondent police reports for traffic fatalities involve definitive mentioning of cellphone use (NHTSA 2019a). In other regions in which camera and phone records are collected for police reports, the proportion of phone-use-caused road fatalities is far higher: 29.6% in Shanghai in 2019,[8] for example. This severe verification and under-reporting issue make police-report-based risk inference invalid (Levitt and Porter 2001). It has also prompted some fleet owners to adopt telematics or in-vehicle surveillance solutions in some commercial settings (Hubbard 2003). Therefore, although cellphone use is widely accepted as a major risk factor in road accidents, it has not received similar legal emphasis as other factors such as DUI and speeding.

1.3 Riskiness of handheld phone use

This subsection presents results from several observational studies that reveal the riskiness of HPU. The ideal test would be to conduct an experiment that changes HPU for a randomized subset of drivers. However, the extreme rarity of accidents means that tens of millions of trips must be in the treatment and control groups to achieve statistical significance.

We thus rely on observational studies. Specifically, we focus on two channels: accidents and passenger complaints. To obtain a causal estimate, we perform four procedures. First, we add trip- and driver-level control variables. This includes all available variables on passenger demand (number of passengers, request frequency at the time

[8] https://www.chinadaily.com.cn/china/2014-11/14/content_18914543.htm.

Table 1: Summary Statistics of Select Data Variables

Variable		Mean	Std. Dev.	Pctl(10)	Median	Pctl(90)
Driver-level						
Driver characteristics	Age (years)	43.4	12.3	27.9	42.1	60.8
	Is female	0.13	0.34	0.00	0.00	1.00
	Past rides completed	734.4	530.7	98.0	640.0	1505.0
Vehicle characteristics	Vehicle age (years)	5.6	3.3	2.0	5.0	10.0
	Is SUV	0.12	0.33	0.00	0.00	1.00
	Is luxury	0.05	0.22	0.00	0.00	0.00
Trip-level						
Trip characteristics	Distance (miles)	4.4	3.8	1.4	3.2	8.6
	Duration (minutes)	16.8	8.5	8.3	14.8	27.8
	Percent of drive-time with passenger	0.72	0.15	0.50	0.74	0.89
	Is weekend	0.29	0.45	0.00	0.00	1.00
	Is night-time	0.41	0.49	0.00	0.00	1.00
	Population density in geohash6 of pickup (1/sqmi/1k)	17.5	21.4	1.4	9.0	44.9
	Requests in same hour and geohash6 of pickup	30.6	48.7	0.0	9.0	93.0
Precipitation	Is raining	0.27	0.45	0.00	0.00	1.00
	Average rain intensity (in/hr/sqmi/1k)	10.6	21.7	0.5	1.9	27.8
	Change in rain intensity (in/hr/sqmi/1k)	-0.0	4.7	-0.5	0.0	0.5
	Abs. val. change in rain intensity (in/hr/sqmi/1k)	1.6	6.1	0.0	0.0	2.8
	Is snowing	0.05	0.22	0.00	0.00	0.00
	Average snow intensity (in/hr/sqmi/1k)	4.4	6.9	0.4	1.4	10.4
	Change in snow intensity (in/hr/sqmi/1k)	0.0	1.2	-0.1	0.0	0.1
	Abs. val. change in snow intensity (in/hr/sqmi/1k)	0.4	1.5	0.0	0.0	0.8
Risky behavior	Meters of handheld phone use per mile	13.9	64.4	0.0	0.0	0.0
	Harsh acceleration incidents per 1k miles	65.7	206.6	0.0	0.0	232.1
	Safety complaints per 1k miles	1.1	9.4	0.0	0.0	0.0

Notes: This table provides summary statistics of key data variables based on our observational data, which comes from a random sub-sample of drivers and trips from April 2019 to May 2020. We filter out uncommon trips defined by the start-geohash5-end-geohash5 combination having fewer than 10 occurrences. Some variables, such as driver earnings and accidents, are hidden to protect the company's private information. Risky behavior is recorded only when the vehicle is moving. Our sample size is more than 10 million trips.

and location of trip start, etc.), company actions, environmental factors such as precipitation, various time and location fixed effects (hour-of-day, day-of-week, month-year, start-end-geohash5), as well as other types of risky behavior detected from phone sensors. Second, for trips with accidents, we remove all post-accident HPU records. This limits reverse causal effects: accidents may lead to phone use, inflating the causal effect of interests. Then, we add driver fixed effects to arrive at our main specification, which corresponds to our investigation of moral hazard and behavioral improvement in latter sections:

$$\frac{\mathbf{1}_{\text{acc},it}}{\text{miles}_{it}} = p_b b_{it} + x'_{it}\beta + FEs^{i,\text{loc}_t,\text{time}_t} + \varepsilon_{it} \tag{1}$$

Here, b is HPU/mile, x are observable covariates, and FEs are the fixed effects. We find that a 1% increase in HPU/mile leads to a 0.06% increase in accident/mile and a 0.02% increase in passenger safety complaints. This magnitude is large: holding phones for mere 2.3 more meters per mile leads to a one percent increase in accident rates. The results reported here are in elasticity.

Fact 1 *Handheld phone use increases accident likelihood.*

Table A1 shows the cost estimates for each of the steps above. For robustness, we further absorb driver-month-location-pair average to control for driver-specific unobserved location and time factors (row 5 of Table A1). This specification dramatically reduces the number of effective observations since most driver-month-location-pairs in our dataset have only a few observations with even fewer variations in HPU or accident.

1.4 Behavioral Stability

We regress trip-level HPU frequency on exhaustive trip-level covariates, as well as driver, location (geohash6-pair of start and ending locations), and time (hour-of-week and week-year) fixed effects to understand the relative importance of across- vs. within-driver behavioral differences. The exercise suggests that across-driver variation accounts for 68% of the overall behavioral variance. This surprising stability of HPU within drivers may contradict rationality given a myriad of environmental risk factors. Figure A1 further breaks down the incremental effect of including different interactive fixed effects on the predictability of the model.

$$b_{it} = x'_{it}\beta + FEs^{i,\mathrm{loc}_t,\mathrm{time}_t} + \varepsilon_{it} \tag{2}$$

Fact 2 *Across-driver differences account for a majority of the total variation in HPU behavior.*

2 Conceptual Framework

In this section, we lay out a simple rational expectation benchmark to explain handheld phone use ('HPU') behavior. The framework formalizes moral hazard and reveals two testable implications. We also discuss the relationship between ex-ante and ex-post moral hazard.

2.1 Rational Benchmark

HPU, denoted as b, influences consumption h by causing accidents with Bernoulli probability $\lambda(b, x)$ and by providing private benefit $g(b; \beta)$ (Equation 3). Here, x represents environmental factors and influences HPU's marginal contribution to accident likelihood, while β governs the HPU elasticity with respect to incentives. When accidents occur, consumers suffer loss ℓ based on their insurance y, which the transportation company sets exogenously.

Each unit of HPU carries minute incremental accident risk. We thus rely on the Arrow-Pratt approximation to model its utility impact, denoting $\frac{\gamma}{2}$ as the absolute risk aversion coefficient. This leads to the utility characterization in Equation 4. Note that this equation is simply the private benefit of b minus the certainty equivalent (CE) of the accident risk associated with b.

$$h(b|\Omega) = -\lambda(b,x)\ell(y) + g(b;\beta) \tag{3}$$

$$\begin{aligned} u(b|\Omega;\theta) &= \mathbb{E}\left[h\right] - \frac{\gamma}{2}\mathbb{E}\left[h^2\right] \\ &= g(b;\beta) - \overbrace{\underbrace{\lambda(b,x;p)}_{\text{risk occurrence}}\ \underbrace{\ell(y)\left[1+\frac{\gamma}{2}\ell(y)\right]}_{\text{loss after insurance}}}^{CE} \end{aligned} \tag{4}$$

$$\text{where } \Omega = \{x,y\} \text{ and } \theta = \{\beta,\gamma\}$$

This model assumes that (1) the arrival of accidents is a Bernoulli trial with probability λ. (2) accident loss ℓ is deterministic in drivers' expectation conditional on y. (3) $g(b)$ is increasing and concave so that the private benefit of conducting risky behaviors features diminishing returns. (4) The Arrow-Pratt approximation, which is appropriate for small risks, can also be viewed as the "negligible third-derivative" approach *a la* Barseghyan et al. 2018.

2.2 Behavior Characterization and Moral Hazard

The first-order condition can be represented by Equations 5, in which we define the implicit price of HPU as the certainty equivalent of the additional risk exposure caused by each unit of b. Given assumption (3) above, we can adopt a power utility form without loss of generality (Equation 6), which resembles how cellphone usage is modeled in Nevo et al. (2016). This further gives rise to a simple characterization of b (Equation 7).

$$b_{it}^{*} = \left[\underbrace{-\frac{\partial CE(b, y; x, \gamma)}{\partial b}}_{\text{implicit price of } b} \right]^{\beta} \tag{5}$$

$$g(b;\beta) := b^{1+\frac{1}{\beta}} / (1 + \frac{1}{\beta}) \tag{6}$$

$$\log b^{*} = \underbrace{\beta}_{\text{HPU elasticity}} \cdot \left[\underbrace{\log \lambda_b(b;x)}_{\substack{\text{HPU's marginal} \\ \text{contribution on} \\ \text{accident rate}}} + \underbrace{\log(y)}_{\text{deductibles}} + \underbrace{\log(1 + \frac{\gamma}{2} y)}_{\text{risk aversion}} \right] \tag{7}$$

We have now expressed the ex-ante moral hazard effect as a log-linear "demand curve," in which consumers' risk-taking behavior b responds to the "implicit price" of b according to an elasticity term β that denotes how much they value b independently from the risk. Assuming individuals are risk-averse and that the private benefit of HPU is positive, we can generate the following testable implications ("TI").

TI 1 Environmental Sensitivity *Risky behavior reduces in environments that heighten its riskiness (or its marginal contribution on accident rate).*

TI 2 Moral Hazard *More insurance coverage encourages risky behavior.*

2.3 Ex-ante vs. Ex-post Moral Hazard

A prominent empirical literature investigates moral hazard in health insurance utilization (Einav and Finkelstein 2018). Many attribute this to Arrow's argument that health insurance may increase the demand for medical care (Arrow 1963). However, as in Mark Pauly's critique, "the response of seeking more medical care with insurance than in its absence is a result not of moral perfidy, but of rational economic behavior." In response, Arrow highlighted that the central issue with moral hazard was that "the optimality of complete insurance is no longer valid when the method of insurance influences the demand for the services provided by the insurance policy" (Pauly 1968; Arrow 1968).

Our "demand-curve" characterization above demonstrates that ex-ante and ex-post moral hazard are fundamentally rooted in the same behavioral rationality assumption on some behavior b and that b has negative externalities on others. Within the health insurance context, recent evidence suggests that ex-post moral hazard is driven by coarse individual response to price heuristics (Brot-Goldberg et al. 2017). To the extent that individuals are even less salient about the implicit price of yet-to-be-realized risks compared to explicit ones when seeking claim reimbursements, we may expect β to deviate more from the rational benchmark in ex-ante moral hazard. However, the direction of deviation can vary based on the sources of bounded rationality: Prospect Theory and non-linear probability weighting would predict over-exaggeration of small risks and overly-cautious behaviors (Tversky and Kahneman 1992; Barseghyan et al. 2013), while inattention to risk exposure would suggest the reverse.

3 Reduced-Form Evidence

This section conducts reduced-from tests for TI1 (environmental sensitivity) and TI2 (moral hazard). Here, we take advantage of repeated and high-frequency observations for the same driver as well as plausibly exogenous changes to drivers' risk exposure due to weather shocks and to state-level insurance mandate changes. Our results reject strong environmental sensitivity and moral hazard.

3.1 Behavioral Response to Weather Shocks

Precipitation significantly increases accident likelihood. Despite slower traffic, the overall risk of a fatal crash increases by 34% during active precipitation (Stevens et al. 2019). Precipitation also makes HPU riskier. Adding an interaction term between HPU frequency and precipitation indicator to our HPU riskiness analysis (Equation 1 and Table A1 row 3), we find that the average precipitation event increases the marginal accident impact of HPU by 38.7% (10.9%).

The heightened riskiness of HPU should deter HPU based on TI1 (environmental sensitivity). We test this hypothesis by directly regressing trip-level HPU frequency on the precipitation indicator, average precipitation level of a trip, as well as changes in precipitation level within a trip (from beginning to end). The main specification includes trip-level controls specified in Table 1 as outlined in Figure 1 note.

$$b_{it} = \begin{pmatrix} \kappa_{\text{rain}} \\ \kappa_{\text{snow}} \end{pmatrix}_{it} \times (\underbrace{Inch_t}_{\substack{\text{across-trip} \\ \text{precip change}}}, \underbrace{\Delta Inch_t}_{\substack{\text{within-trip} \\ \text{precip change}}}) + x'_{it}\beta + FEs^{i,\text{loc}_t,\text{time}_t} + \varepsilon_{it} \tag{8}$$

Figure 1 reports the regression results and the percentage change in HPU induced by different precipitation events. In the left panel, different colors (top to bottom) represent the HPU impact of across-trip differences in average precipitation, measured in thousandth inches per hour. Along the X-axis, we plot the effect of within-trip precipitation change on HPU. This complements the first estimate by mitigating the effect of potential selection bias: drivers may be more likely to stop driving as precipitation increases.

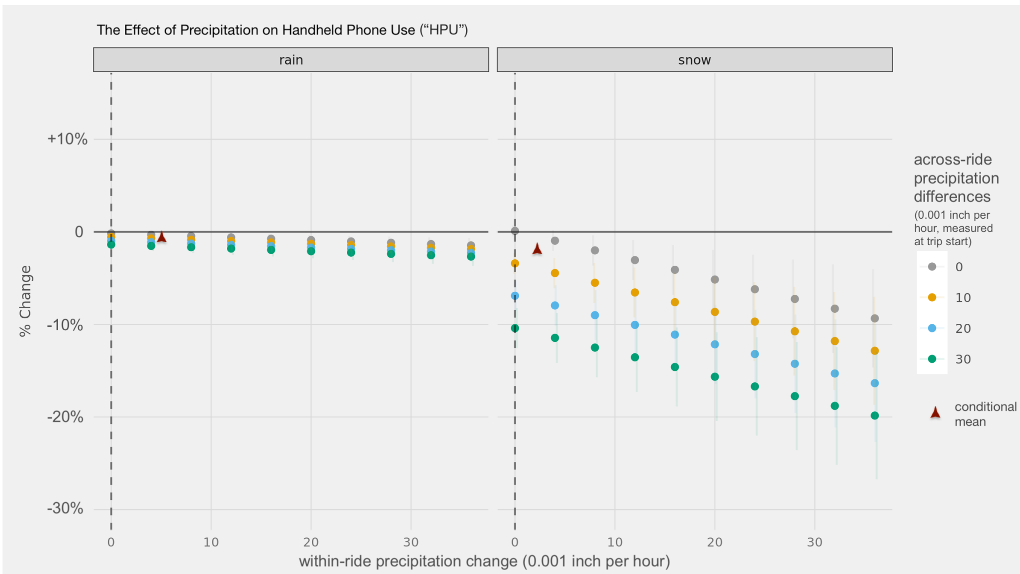

Figure 1: HPU impact of within-trip precipitation change

Notes: This graph reports the result of Equation 8: the effect of precipitation on HPU. It shows the impact (on HPU) of both within- (x-axis/left-to-right) and across-trip (color/top-to-bottom) precipitation change, for rain and snow weather (left vs. right panels). Controls include demand and company/dispatch factors such as ride type and passenger request density at trip start time and location (geohash6), as well as driver, time, and start-geohash5-end-geohash5 fixed effect.

The estimates are precise, and precipitation has an economically small impact on HPU. This is especially so since the mean and median rain precipitation levels are only 11.4e-3 and 1.6e-3 inches/hour. Snow precipitation events elicit larger HPU responses. However, the HPU response is still economically small: the mean and median snow precipitation levels are 4.6e-3 and 1.2e-3 inches/hour.

The main assumption is that precipitation does not directly alter the private benefit of phone use while driving. It may be violated if, for example, the private benefits drivers derive from text-and-drive scale with precipitation both across and within trips. Figure A2 reports the results from specification robustness checks for our results here.

3.2 Behavioral Response to Insurance Changes: A Natural Experiment

As discussed in Section 1.1, drivers at the transportation company are essentially exempt from any liabilities associated with road accidents. However, protection against losses incurred by the drivers themselves is not as well-insured. Drivers use their own cars for trips and still face high deductibles for repair damages. Drivers can also potentially incur high out-of-pocket expenditures from their medical treatments, which are only partially covered. TI2 (moral hazard) implies that the more exposed drivers are to these losses, the less they should exhibit risky behavior such as HPU.

In this subsection, we directly test for moral hazard (TI2) with changes in insurance mandates across several states. Specifically, coverage limits on own medical payments were reduced from a quarter million to the new state mandatory minimum levels, which in some states are no coverage at all. This, combined with differences in the frequency and utilization of this claim type across states, led to significant variation in the impact on drivers' expected loss in accidents (Figure A3). Based on the tailing twelve-month claim records at the time of the mandate change, drivers in most states saw a small impact ($< \$0.01$ per trip), but those in three states experienced changes greater than $1 per trip (Figure 2(c)). This variation is largely driven by tort differences across states, which shift the burden on victims' own medical coverage versus the liability coverage of the injurers. This insurance change was unanticipated; drivers were informed via email the night before the change.

We use a two-step synthetic control approach to estimate the treatment effect of this insurance change (Abadie et al. 2010). In the first step, we estimate residual HPU progression across different states by regressing HPU on trip-level controls and fixed effects. It is similar to Equation 2, with the addition of state and calendar-day interactive fixed effects, which is plotted in Figure 2(a) with colors representing treatment status, lines for states, and X-axis for date relative to insurance change. We see that almost all states experienced a drop in HPU on the day of treatment regardless of insurance change. This is most likely a result of the treatment timing being on the first of the month. It also highlights the importance of unobserved time trends.

In the second step, we calculate the treatment effect using a generalized synthetic control approach proposed by Xu (2017) and the package developed by Xu et al. (2017). For each treated state, we impute a counterfactual "synthetic-control" state using all control-group states based on a linear interactive fixed effects model *a la* Bai (2009). The fixed effects allow flexible interactions between state-specific and time-specific (days-since-treatment) intercepts, and the degree of interaction is selected based on cross-validation in the pre-period. We bootstrap standard errors across the two steps. In the first step, for each state, we bootstrap (sub-sampling with replacement within that state) 100 daily HPU residual estimates (i.e. 100 lines as in Figure 2(a)). In the second step, we also bootstrap 100 times, each using the corresponding bootstrap from the first step.

Figure 2(b) plots the treated-state average against that of the synthetic control estimates. The mean synthetic counterfactual closely tracks real data pre- and post-treatment, which verifies the validity of the synthetic control while indicating a lack of treatment effect.

To further test moral hazard and HPU response to risk, we separate treated states into three groups based on the cost impact of the insurance change reported in Figure 2(c): high-impact states include the three states with per-trip expected loss of over $1 to drivers; low-impact states include those with less than $0.1 expected loss to drivers; the remaining states are included in the medium-impact group. Figure 2(d) reports the average ATT over all states (on the left) as well as across the three cost-impact groups. We find a precise null effect across all groups. Meanwhile, the average cost impact among the high-impact group is over ten times that of the low-impact group. However, this higher expected loss impact did not further deter HPU.

Taken together, our reduced-form tests demonstrate that drivers' risky HPU behavior is insensitive to changes in their accident risk exposure. Precipitation raises HPU riskiness by more than 30% on average, but it only induces less than a 1% reduction in average HPU frequency. In addition, significant insurance coverage reduction also fails to elicit any deterrence effect on HPU. This implies that drivers have a very small effective HPU elasticity (β), which can either reflect an extraordinarily high preference for HPU or, perhaps more plausibly, a lack of attention to the risk exposure associated with their HPU behavior. We investigate the role of inattention in the next section.

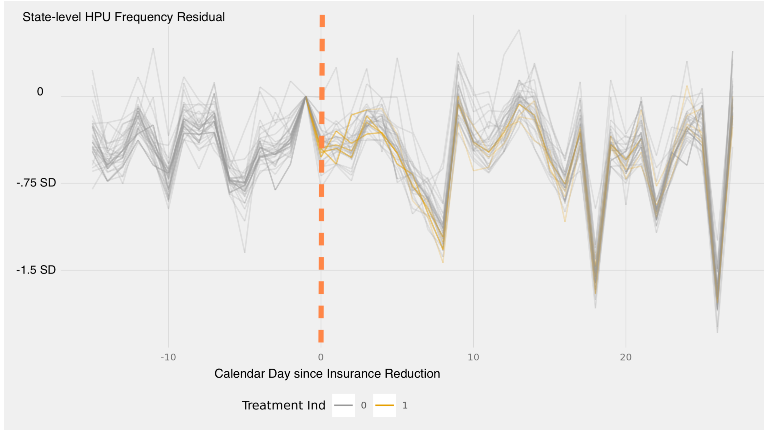

(a) HPU Progression by State Pre and Post Insurance Change

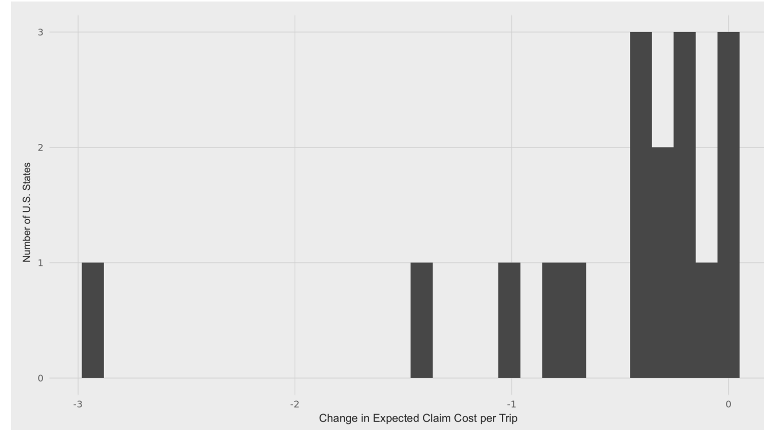

(c) Histogram of Treatment Intensity by State

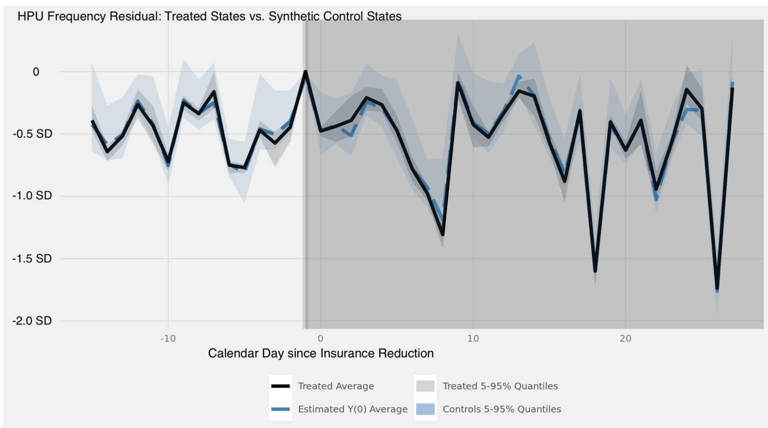

(b) Treated Data vs. Synthetic Control Counterfactual

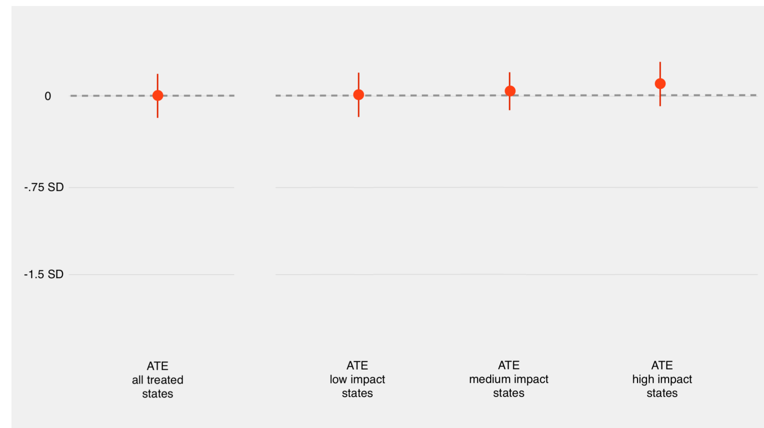

(d) ATT by Treatment Intensity

Figure 2: Behavioral Response to Insurance Change

Notes: These graphs present the synthetic control results on the effect of insurance coverage reduction on handheld phone use frequency. In (a), we plot the calendar-day fixed effects for each U.S. state—the "residual HPU" after controlling for trip-level controls (including location/time fixed effects) as well as driver fixed effects—with colors representing treatment status. "1.5SD" denotes 1.5 times the standard deviation of HPU frequency. In (b), we plot the average residual HPU frequency among the treated states (in black) against that of the synthetic control state (dotted blue line). The synthetic control state is constructed using folding cross-validation in the pre-period. In (c), we plot a histogram of the expected insurance cost savings per trip from the coverage reduction across states, which corresponds to increased risk exposure for drivers in those states. In (d), we trisect treated states into three groups based on the cost impact plotted in (c). We then report the average ATT among each of the three groups as well as across all states. We bootstrap standard error by sampling with replacement among the respective treated groups and among control states.

4 Experimental Evidence

Our findings in the last section reveal a general behavioral insensitivity to risk and a lack of moral hazard. We now conduct an experiment with a simple nudge. It sheds light on the role of inattention HPU and HPU response to incentives when drivers do pay attention.

4.1 Experimental Design

Treatment: we sent the following warning to drivers via a one-time text message:
"Our app shows you may be holding your phone while driving. Passenger reports of unsafe driving, like handheld phone use, can lead to suspension."

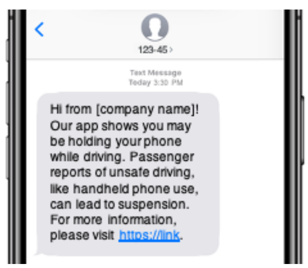

Figure 3: Illustration of the Text-message Intervention

This intervention is designed (1) to raise drivers' *attention* to the riskiness of HPU and (2) to increase the financial risk associated with HPU that passenger complaints may lead to suspension. The text message ends with a link to an FAQ page containing legal primers, explanations of how telematics data are collected as well as details about suspension, including investigation procedures to verify passenger safety complaints.
Target population: drivers with prior passenger complaints and are in the top 5% percentile HPU over the two-week period prior to intervention. We focus on drivers who work at least 20 hours a week on average over the last quarter.
Logistics: we pre-select a 30-min window randomly (3-3:30 pm EST). On the intervention day, all targeted drivers logging on and starting a shift during this period are in our experimental sample. To ensure safety, we automatically exclude drivers who are traveling at non-trivial speeds when they log on. There are 467 treated drivers, to whom we sent the treatment text message right when they log on. The main control group consists of 512 hold-out drivers, including 25 drivers who were randomized into treatment but were unreachable by text message due to engineering issues (connectivity

or device issues). We continue to include trip-level controls and fixed effects as in the reduced-from section. Standard errors are clustered at the driver level.

4.2 Average Treatment Effects

Consistent with the reduced-form analyses, we estimate treatment effects with the regression below, in which τ represents the calendar day relative to the intervention time.[9]

$$b_{it} = \alpha_t + \delta_\tau D_{it} + x'_{it}\beta + FEs^{i,\text{loc}_t,\text{time}_t} + \varepsilon_{it} \tag{9}$$

The average treatment effect is large and persistent. Figure 4(a) shows that HPU reduces by 1/3 in the 24-hour window after the intervention. This likely includes a Hawthorne effect: drivers behave better when they first realize that they are being watched. However, this effect is typically short-lived,[10] which holds up in our case: the first week sees an average treatment effect of 21%, which reduces to stabilize around 15% in subsequent weeks: 14% for weeks 2 and 3; 16% for week 4.

4.3 Inattention and Heterogeneous Treatment Effect

This subsection argues that driver attention is the reason why our experiment generates seemingly contrasting results from our reduced-form tests. First, the intervention generates no detectable change in other types of risky behavior. Figure 4(c) reports the treatment effect on harsh acceleration. The intervention increases the financial risk of all behavior that caused safety complaints, but only specifically called drivers' attention to HPU. It is therefore plausible that driver attention played a key role in the big treatment effect on HPU.

Moreover, the additional treatment incentives only apply when there are passengers in the car. In reality, drivers spend a good amount of time going to pick up assigned passengers (Table 1). During these trips, they are paid similarly and are exposed to the same amount of accident risk as trips with passengers.[11] Figure 4(d) reports the differ-

[9] Including driver fixed effects resembles a block design and puts higher regression weights on drivers with higher variance in the treatment variable, which are those who did not change the number of trips driven before and after treatment. We do not detect a statistically significant treatment effect on driving time (Figure 4(b)).

[10] Friedman and Gokul (2014) provides an empirical survey of the Hawthorne effect:"Of note is that the Hawthorne effect was rather short-lived... teachers habituate quite rapidly to video observation and return to 'normal' levels of practice within a day or so after the introduction of the camera."

[11] Passenger injuries and loss productivity are fully insured by the company.

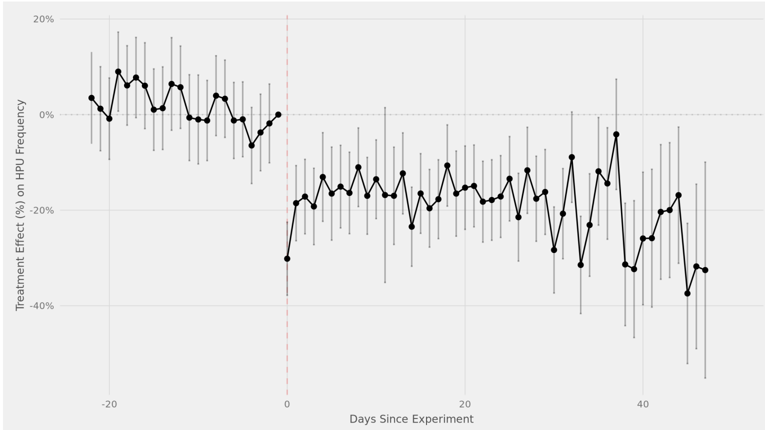

(a) Daily % ATE on HPU Frequency

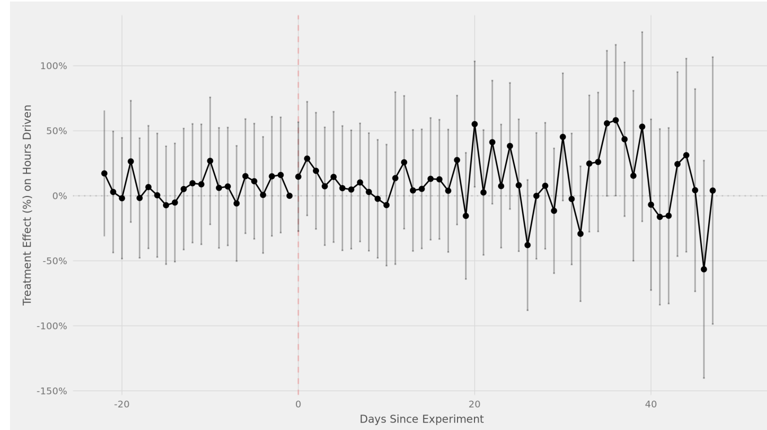

(b) Daily % ATE on Hours Driven

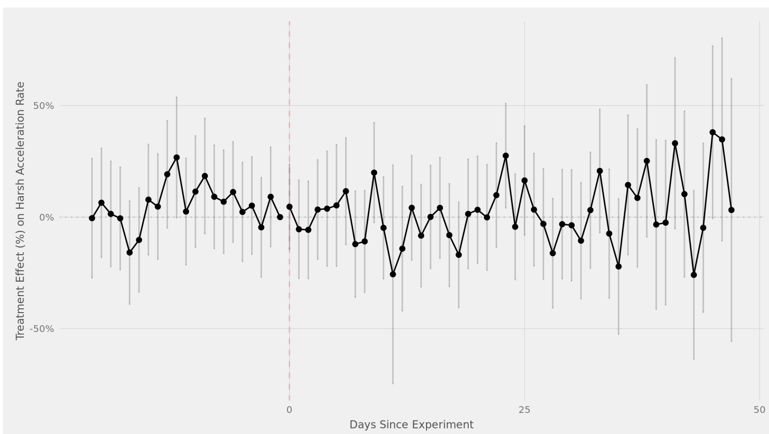

(c) Daily % ATE on Harsh Acceleration Frequency

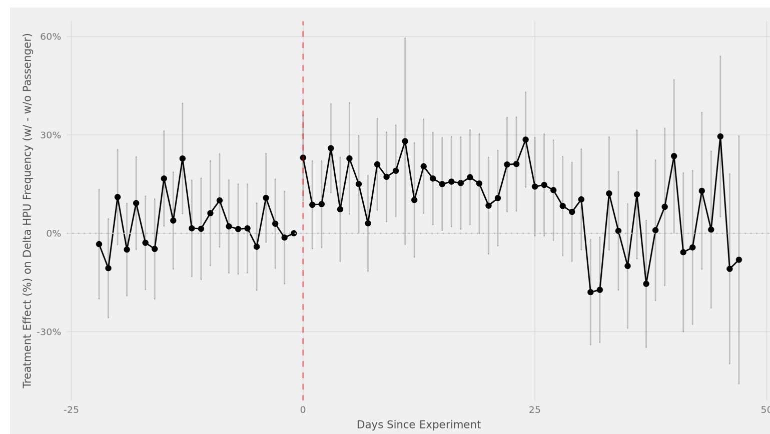

(d) Differential TE for Trip w/ and w/o Passengers

Figure 4: Direct Intervention: Experimental Results

Notes: Figures (a), (b), (c) plot the estimates on the interaction between daily fixed effects and the treatment indicator, converted to percentage terms based on the pre-treatment average level of the outcome variable. Figure (d) further interacts this interaction term with an indicator of whether the trip has passengers and therefore is subject to the treatment incentives. (a) and (d) use the frequency of HPU as the outcome variable, controls for demand and platform factors such as ride type and trip request density, as well as driver, time-of-day, and weekday fixed effects. (c) reports the same estimates as Figure (a), replacing the outcome variable with the frequency of harsh acceleration. (b) is on the driver-day level and thus uses the total driving hour as the outcome variable, with averaged trip-level controls.

-ential treatment effects using a "difference-in-differences" estimator, which cancels out any incentive effect to isolate the attention effect. Surprisingly, the treatment effect is larger for non-passenger trips with no added incentives. This is due to higher pre-treatment HPU, so a smaller proportional treatment effect (among non-passenger trips) translates into a larger drop in HPU level. Quantile regressions reveal that the proportional treatment effect is almost identical across drivers with different pre-treatment levels of HPU (Figures A5(a) and A5(b)). Higher HPU drivers thus saw a much higher decrease in absolute levels.

Interestingly, Figure A5(c) shows that the treatment effect is smaller for drivers with past risk experience: those who had accidents or "near-misses" (as defined in Shum and Xin (2019)), but the effect is strong in snow and rain rides. This order is reversed when calculating the attention effects alone (Figure A5(d)). Combined, we conclude that HPU response to incentives is triggered when drivers are nudged to pay attention to the behavior and its riskiness.

Conclusions

This paper uses novel data on risky driving behavior to directly quantify the magnitude of moral hazard. Focusing on handheld phone use (HPU) while driving, the central finding is that inattention to risk is prevalent, which severely limits moral hazard.

The focus on accident prevention improves our empirical understanding of the expansive role of moral hazard in insurance design and tort law, its interaction with bounded rationality, and the connections between ex-ante and ex-post moral hazard. Our results also have direct policy implications: instead of raising accident punishment or erecting harsher limits and deductibles–essentially reducing insurance–to deter risky driving, raising attention to or directly pricing on such behavior is far more effective.

Our results do not imply that moral hazard does not exist. First, our experiment demonstrates that increased attention can trigger moral hazard. In addition, as discussed in Section 2, the magnitude of moral hazard is highly context-dependent and varies with the behavior of interest. Moreover, the cause of accident risk is diverse, and handheld phone use is but one risk factor. Nonetheless, as sensor data proliferate, our study points to a real possibility that insurers can identify all significant risky driving behaviors and move towards full insurance and a "first-best" contract (Holmström 1979; Jin and Yu 2021).

References

Abadie, Alberto et al. (2010). “Synthetic control methods for comparative case studies: Estimating the effect of Californiaâs tobacco control program”. In: *Journal of the American statistical Association* 105.490, pp. 493–505.

Abbring, Jaap H et al. (2003). “Adverse selection and moral hazard in insurance: Can dynamic data help to distinguish?” In: *Journal of the European Economic Association* 1.2-3, pp. 512–521.

Abbring, Jaap H et al. (2008). “Better safe than sorry? Ex ante and ex post moral hazard in dynamic insurance data”. In:

Arrow, Kenneth J (1963). “Uncertainty and the Welfare Economics of Medical Care”. In: *American Economic Review* 53.5, pp. 941–973.

— (1968). “The economics of moral hazard: further comment”. In: *The American economic review* 58.3, pp. 537–539.

Athey, Susan et al. (2019). “Service quality in the gig Economy: Empirical evidence about driving quality at Uber”. In: *Available at SSRN*.

Bai, Jushan (2009). “Panel data models with interactive fixed effects”. In: *Econometrica* 77.4, pp. 1229–1279.

Baker, Tom (1996). “On the genealogy of moral hazard”. In: *Tex. L. Rev.* 75, p. 237.

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

— (2018). “Estimating risk preferences in the field”. In: *Journal of Economic Literature* 56.2, pp. 501–64.

Blincoe, Lawrence et al. (2015). *The economic and societal impact of motor vehicle crashes, 2010 (Revised)*. Tech. rep.

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

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

Choudhary, Vivek et al. (2019). “Nudging Drivers to Safety: Evidence from a Field Experiment”. In: *INSEAD Working Paper*.

Ducat, Arthur Charles (1865). *The practice of fire underwriting*. T. Jones.

Einav, Liran and Amy Finkelstein (2018). “Moral hazard in health insurance: What we know and how we know it”. In: *Journal of the European Economic Association* 16.4, pp. 957–982.

Finkelstein, Amy (2014). “Moral hazard in health insurance: Developments since Arrow (1963)”. In: *Moral Hazard in Health Insurance*. Columbia University Press, pp. 13–44.

Friedman, Jed and Brindal Gokul (2014). “Quantifying the Hawthorne Effect”. In: *Development Impact, World Bank*.

Gabaix, Xavier (2019). “Behavioral inattention”. In: *Handbook of Behavioral Economics: Applications and Foundations 1*. Vol. 2. Elsevier, pp. 261–343.

Hassan, Mohammed Mehedi et al. (2018). “A robust human activity recognition system using smartphone sensors and deep learning”. In: *Future Generation Computer Systems* 81, pp. 307–313.

Hellner, Jan (2001). “Compensation for personal injuries in Sweden–A Reconsidered View”. In: *Scandinavian studies in law* 41.249,263,267,269,276.

Holmström, Bengt (1979). “Moral hazard and observability”. In: *The Bell journal of economics*, pp. 74–91.

Hubbard, Thomas N (2003). “Information, decisions, and productivity: On-board computers and capacity utilization in trucking”. In: *American Economic Review* 93.4, pp. 1328–1353.

James Jr, Fleming (1947). “Accident Liability Reconsidered: The Impact Liability Insurance”. In: *Yale Lj* 57, p. 549.

Jeziorski, Przemyslaw et al. (2017). “Adverse selection and moral hazard in the dynamic model of auto insurance”. In: *Working Paper*.

Jin, Yizhou and Shoshana Vasserman (2019). “Buying data from consumers: The impact of monitoring programs in us auto insurance”. In: *Working Paper. Harvard University*.

Jin, Yizhou and Thomas Yu (2021). “How to Prevent Traffic Accidents: Moral Hazard, Inattention, and Behavioral Data”. In: *UC Berkeley Working Paper*.

Kahneman, Daniel (2003). “Maps of bounded rationality: Psychology for behavioral economics”. In: *American economic review* 93.5, pp. 1449–1475.

Knapp, V. (1983). *International Encyclopedia of Comparative Law*. International encyclopedia of comparative law. Vol. 11, Torts v. 11, pt. 1. Mohr.

Levitt, Steven D and Jack Porter (2001). “How dangerous are drinking drivers?” In: *Journal of political Economy* 109.6, pp. 1198–1237.

Liu, Meng et al. (2018). *Do digital platforms reduce moral hazard? The case of Uber and taxis*. Tech. rep. National Bureau of Economic Research.

Llerena, Luis E et al. (2015). “An evidence-based review: distracted driver”. In: *Journal of trauma and acute care surgery* 78.1, pp. 147–152.

Megat-Johari, Nusayba et al. (2023). “Evaluation of enforcement and messaging campaign focused on reducing cell phone-related distracted driving”. In: *Transportation research record* 2677.1, pp. 1741–1752.

Miller, George A (1956). “The magical number seven, plus or minus two: Some limits on our capacity for processing information.” In: *Psychological review* 63.2, p. 81.

Nevo, Aviv et al. (2016). “Usage-based pricing and demand for residential broadband”. In: *Econometrica* 84.2, pp. 411–443.

NHTSA (2019a). *Driver Electronic Device Use in 2017 — NHTSA — Department of Transportation*.

— (2019b). *Driver Electronic Device Use in 2018 — NHTSA — Department of Transportation*.

Pauly, Mark V (1968). “The economics of moral hazard: comment”. In: *The american economic review* 58.3, pp. 531–537.

Rowell, David and Luke B Connelly (2012). “A history of the term âmoral hazardâ”. In: *Journal of Risk and Insurance* 79.4, pp. 1051–1075.

Shavell, Steven (1979). “Risk sharing and incentives in the principal and agent relationship”. In: *The Bell Journal of Economics*, pp. 55–73.

— (2018). “On liability and insurance”. In: *Economics and Liability for Environmental Problems*. Routledge, pp. 151–163.

Shum, Matthew and Yi Xin (2019). “Time-Varying Risk Aversion: Evidence from Near-Miss Accidents”. In: *Available at SSRN 3455798*.

Skog, Isaac and Peter Handel (2009). “In-car positioning and navigation technologiesâA survey”. In: *IEEE Transactions on Intelligent Transportation Systems* 10.1, pp. 4–21.

Spenkuch, Jörg L (2012). “Moral hazard and selection among the poor: Evidence from a randomized experiment”. In: *Journal of Health Economics* 31.1, pp. 72–85.

Stanovich, Keith E (2009). “Distinguishing the reflective, algorithmic, and autonomous minds: Is it time for a tri-process theory”. In: *In two minds: Dual processes and beyond*.

Stevens, Scott E et al. (2019). “Precipitation and fatal motor vehicle crashes: continental analysis with high-resolution radar data”. In: *Bulletin of the American Meteorological Society* 100.8, pp. 1453–1461.

Strayer, David L and William A Johnston (2001). “Driven to distraction: Dual-task studies of simulated driving and conversing on a cellular telephone”. In: *Psychological science* 12.6, pp. 462–466.

Tversky, Amos and Daniel Kahneman (1992). "Advances in prospect theory: Cumulative representation of uncertainty". In: *Journal of Risk and uncertainty* 5, pp. 297–323.

Wagner, Gerhard (2006). "Tort law and liability insurance". In: *The Geneva Papers on Risk and Insurance-Issues and Practice* 31, pp. 277–292.

Weisburd, Sarit (2015). "Identifying moral hazard in car insurance contracts". In: *Review of Economics and Statistics* 97.2, pp. 301–313.

Wilson, Fernando A and Jim P Stimpson (2010). "Trends in fatalities from distracted driving in the United States, 1999 to 2008". In: *American journal of public health* 100.11, pp. 2213–2219.

Xu, Yiqing (2017). "Generalized synthetic control method: Causal inference with interactive fixed effects models". In: *Political Analysis* 25.1, pp. 57–76.

Xu, Yiqing et al. (2017). *Package gsynth*.

A Appendix: Additional Figures and Tables [For Online Publication]

Table A1: Robustness - HPU Risk Model

	Specification	Accident	Safety Complaints
1	OLS	-0.074	-0.404
		(0.017)	(0.052)
2	1 + trip & driver controls + location/time fixed effects	0.034	0.032
		(0.019)	(0.007)
3	2 + post-accident filter	0.056	0.033
		(0.019)	(0.007)
4	3 + Driver fixed effects	0.060	0.021
		(0.021)	(0.008)
5	3 + Driver-month-location-pair fixed effects	0.042	0.010
		(0.021)	(0.008)

Notes: This table reports estimates from our observational estimates of the riskiness of HPU. The results are presented in elasticity terms. 'OLS' shows the coefficient of regressing accident or safety complaints directly on trip-level HPU. Location fixed effects use the pair of start and finish geohash5 codes for each trip. Time fixed effects include hour-of-day, day-of-week, month-year dummies. The post-accident filter removes all HPU records that occur after accidents for all accident trips.

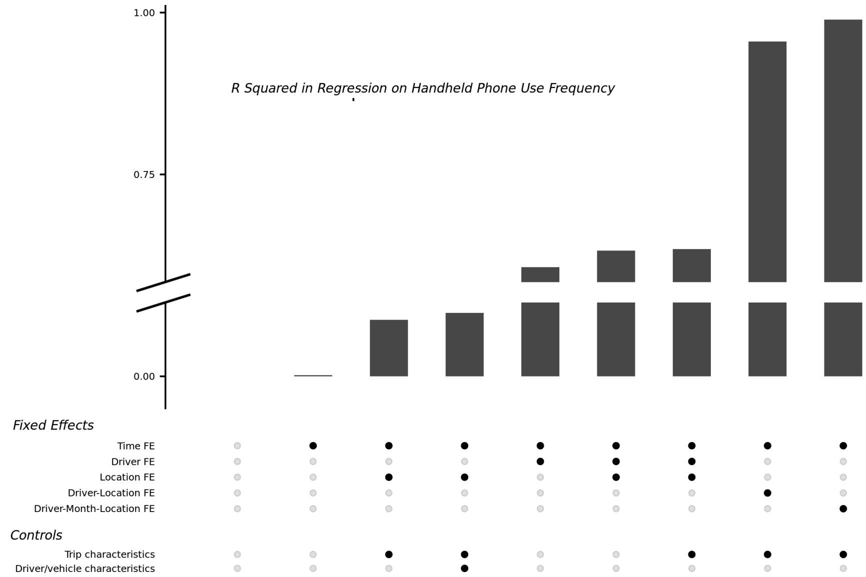

Figure A1: Robustness Checks - R Square of HPU Regressions

Notes: This graph is a robustness check for our regression in Equation 2. It also helps visualize the marginal explanatory power of including various covariates and/or fixed effect.

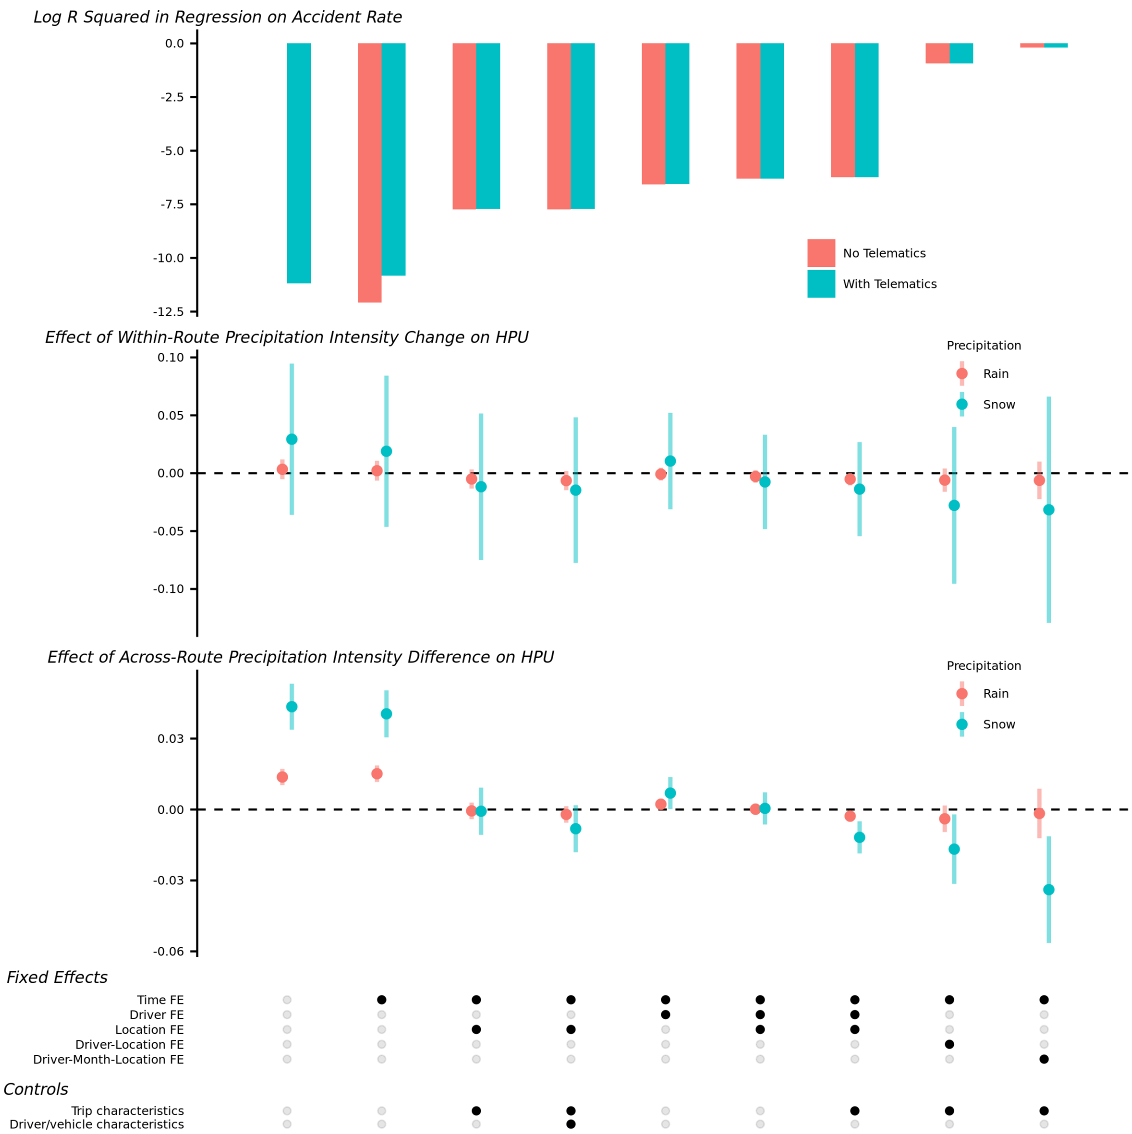

Figure A2: Robustness Checks - Reduced-Form Analyses

Notes: This graph presents robustness check for our regressions in Equation 1 (top panel) and in Equation 8/Figure 1. It helps visualize the marginal explanatory power and impact of including various covariates and/or fixed effect. The mid panel corresponds to the impact on HPU from across-trip precipitation differences. The bottom panel shows the impact of within-trip precipitation change.

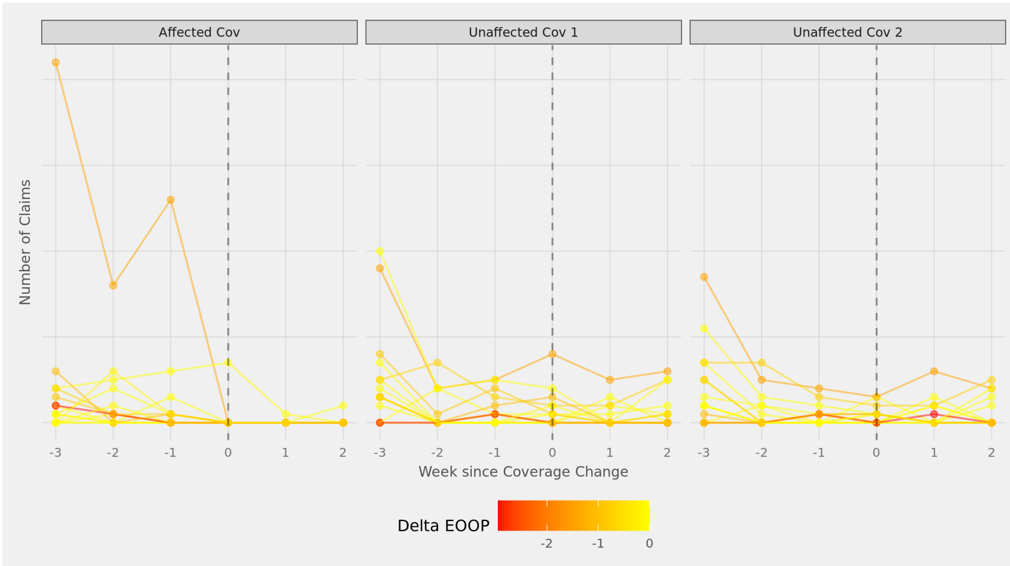

Figure A3: Insurance Change: Claims Progression by State and by Coverage

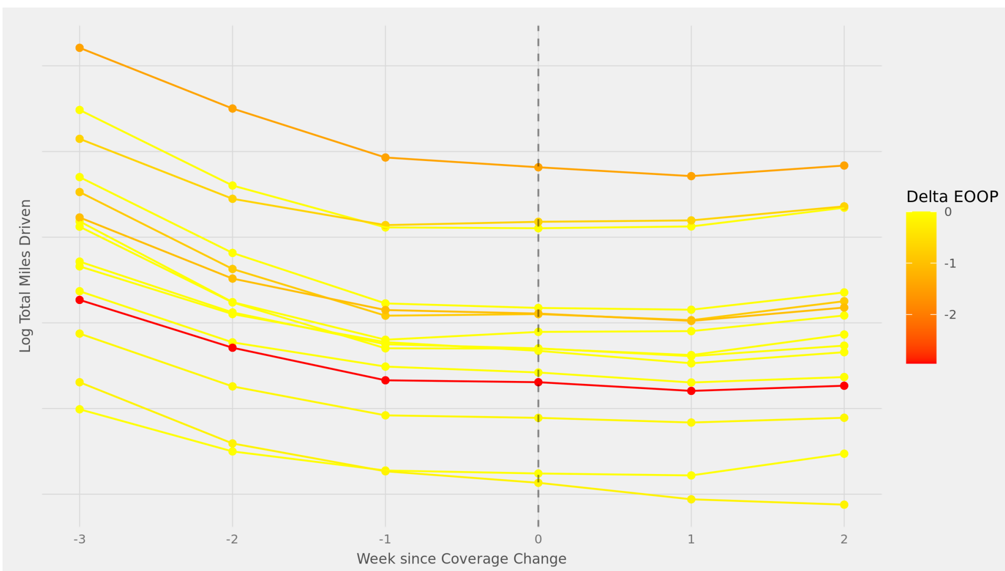

Figure A4: Insurance Change: Miles Driven Progression by State

Notes: These graphs describe the insurance mandate change analyzed in the reduced form section, in which the own medical spending coverage was removed at week zero (Fig A3). Different lines represent different states, while their colors represent the treatment intensity: Delta EOOP means the change in expected out-of-pocket expenditure ("loss") per trip due to the insurance change. States/lines that are colored yellow were unaffected by the insurance change. Both claim counts and miles driven suffer time trends due to macroeconomic and public health events but both have largely stabilized two weeks before the insurance change.

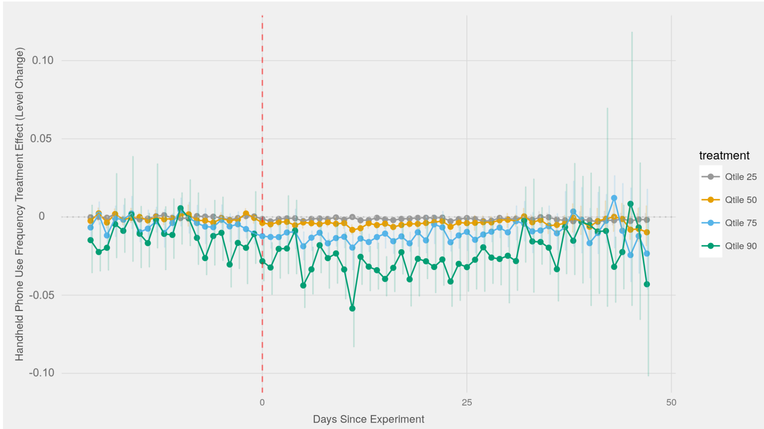

(a) HPU Quantile Treatment Effects

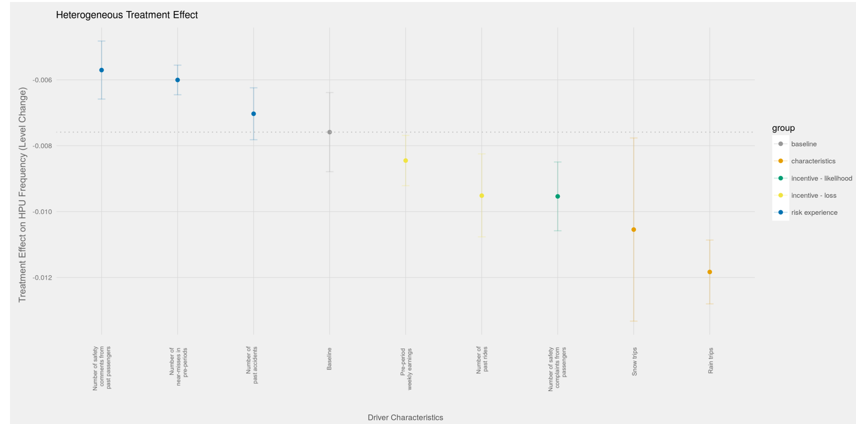

(c) Het. Treatment Effects by Driver Char.

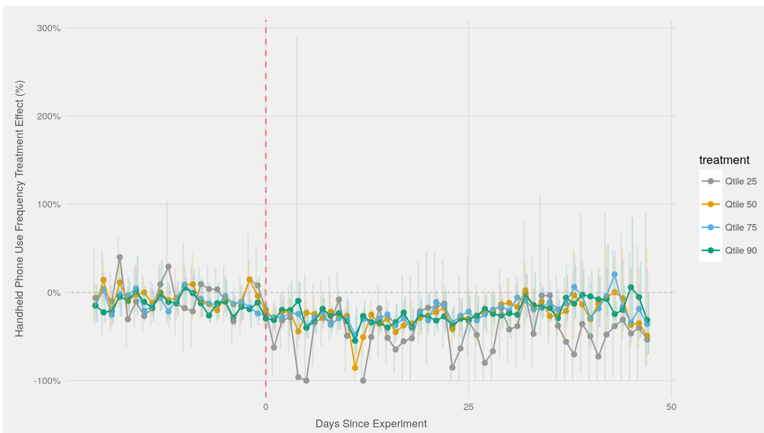

(b) HPU % Quantile Treatment Effects

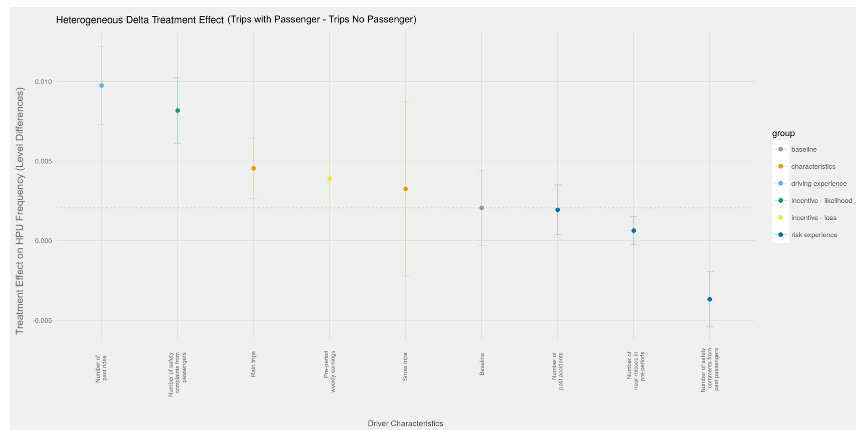

(d) Het. Differential TE by Driver Char. (w/ - w/o Passenger)

Figure A5: Heterogeneous Treatment Effects

Notes: These figures report heterogeneous treatment effects in our direct intervention experiment. Quantile treatment effects in handheld phone use frequency (HPU) are reported in levels in (a) (5% means a 0.05 drop in HPU frequency) and in percentage terms in (b) by dividing the average pre-treatment levels of HPU for each quartile. (c) and (d) reports report heterogeneous treatment effects on HPU frequency by observable driver features. In (c), the coefficients reported are the ones on the interaction between treatment indicator, post treatment indicator, and the corresponding driver feature. Controls are the same as the main regression, including driver fixed effects. (d) adds passenger indicator to the interaction term. The coefficient reports how the treatment effect differ with and without passengers. In (c) and (d), driver characters are selected to represent key factors in the structural model. Specifically, "incentives" includes pre-treatment daily earning and existing safety complaints that increase suspension likelihood. "Risk experience" includes past accidents and passenger safety comments, as well as near-misses identified by severe harsh braking events.

B Data and Algorithm: From Sensors to Handheld Phone Use Behavior ("HPU") [For Online Publication]

Our behavioral data is constructed from raw telematics data in three stages. First, we collect off the smartphone 25Hz three-dimensional streams of accelerations, rotations, and force of gravity based on the phone's accelerometer and gyroscope.[12] We also collect 1Hz streams of GPS location records. These raw sensor records are serialized using Google's Protocol Buffer, compressed down to a very small data package over a small time period, and transmitted to our data center in real-time together with other preexisting data collected for navigation and for normal app functioning.

In the second stage, to process these raw signals, we use a hybrid heuristic- and ML-based approach. We first use a standard low-pass filter to attenuate excessive noise and further regularize the sensor data with a gravity removal procedure outlined in Hassan et al. (2018). We then orient the phone with respect to the vehicle based on GPS movement patterns, which reveal the resting position of the phone within the vehicle (Skog and Handel 2009). These steps make telematics sensor data interpretable from both the phone's and the vehicle's perspectives. For example, Figure A6(a) plots the sensor movements during a right turn in a test drive, where the acceleration and gyration of the phone, adjusted for its GPS speed, allow us to pin down the orientation of the vehicle (top right panel).

In the third stage, we construct sensor data features and identify behaviors using an in-house machine-learning algorithm. We first identify candidate "events," which are defined as abnormal movements in at least one of the sensor data streams. Operationally, we aggregate each sensor data stream over small time intervals and calculate moments of the data. We then detect events using pre-specified moment thresholds. The combination of detected events over several data streams would identify behaviors such as HPU.

There are many types of phone movements that are not HPU (Figure A6(b)). Many drivers keep their phones in cup-holders or on dashboards. Disruptive vehicle movements such as going through speed bumps or collision accidents can also cause large sensor data fluctuations. As a result, the telematics event identification requires that the phone orientation be aligned with the car movement, that the car be traveling at nonzero speed, and that phone movement instability be detected for a nontrivial time duration.

Nonetheless, compared to other risky driving behaviors such as speeding, HPU is

[12] An accelerometer is a compact device designed to measure non-gravitational acceleration. A gyroscope is a device that uses Earth's gravity to help determine orientation.

particularly easy to identify. This is because the detection of HPU is almost entirely unaffected by environmental factors such as speed limits and weather events. Intuitively, HPU identification primarily relies on abnormal gyration of the smartphone itself: events associated with abnormally large variance in gryo sensors tend to indicate either HPU or harsh turns, but HPU is associated with shorter durations of gyro sensor abnormality, which results in higher kurtosis, as well as lower variance in GPS speed (Figure A6(c)). Our training and validation datasets are proprietary and labeled based on test drives and video footage of test drivers. This allows accurate and synchronized labeling of HPU and the corresponding sensor movements. The resulting algorithm achieves over 99% precision at industry-leading recall rates.[13]

[13]The focus on precision is due to limit liabilities associated with intervention based on false-positive identification of HPU. We briefly discuss the implications of our results on optimal precision-recall choice in the conclusion section.

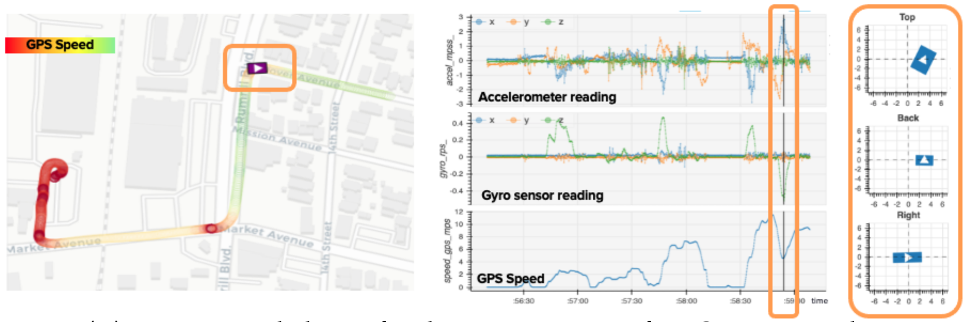

(a) Interpretability of Telematics Data after Steps 1 and 2

(b) HPU versus some notable non-HPU behaviors

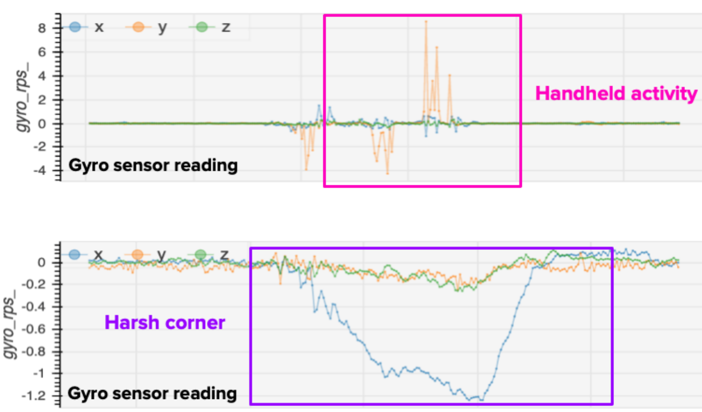

(c) Event and HPU Identification in Step 3

Figure A6: Illustration of HPU Identification from Sensor Data

Notes: These figures correspond to the steps we take to identify risky handheld-phone-use ("HPU") behaviors from telematics sensor data from smartphones. (a) shows the telematics records corresponding to a right turn in our test drive after data collection and cleaning in Steps 1 and 2. (b) illustrates the types of behavior we categorize as HPU versus non-HPU behaviors. (c) illustrates the intuitive identification of HPU using gyro-sensor abnormalities, and how HPU can be distinguished from harsh turns and cornering, which also produce gyro-sensor abnormalities but with lower kurtosis.