The Impact of Monitoring Technologies on Contracts and Employee Behavior: Experimental Evidence from Kenya's Matatu Industry

Erin M. Kelley, Gregory Lane, and David Schönholzer*

July 31, 2018

Abstract

Asymmetric information makes designing contracts challenging. Agency theory suggests that additional information about employees can always help employers improve contracts, but we have little empirical evidence documenting how this subsequently impacts firm operations in developing countries. This project uses a randomized control trial to investigate whether monitoring technologies can ease labor contracting frictions and help firms design a better contract that increases profits and encourages business growth. To this end, we designed a monitoring device that delivers real-time vehicle location and daily productivity and safety statistics to 250 minibuses operating across 9 major commuter routes in Nairobi. These privately owned minibuses are the primary providers of public transportation services in Nairobi, a dynamic that we see in many large urban centers across low-income countries. Owners and drivers in both treatment and control groups submitted daily data on business outcomes, which we used to track changes in contracts and drivers behavior. We find that providing information to vehicle owners allows them to modify the terms of their contracts with drivers. As a result, we find that employees exert more effort, decrease risky behavior that damages the vehicle, and under-report revenue by less, leading to an overall increase in firm profitability and making it easier for firms to expand. Finally we investigate whether these gains to the company come at the expense of commuters and passengers safety, which are already at risk in this industry where accidents are common. While we do not find any evidence that conditions deteriorate, they also do not improve despite the detailed information on safe driving we provide. Only by incentivizing drivers through an additional cash-treatment do we see improvements in safety. This suggests that the introduction of monitoring technologies into informal transportation sectors may need to be coupled with stronger regulation if safety standards are to be improved.

*We are grateful to Jeremy Magruder and Elisabeth Sadoulet for numerous discussions and feedback. We also thank Fiona Burlig, Kyle Emerick, Jeremy Gordon, Nick Li, John Loeser, Matt Pecenco, and Alain de Janvry for helpful comments and suggestions. We gratefully acknowledge financial support from DIL and PEDL. We thank Jeremy Gordon for designing an amazing tracking device. We are also extremely grateful to the entire team at EchoMobile (Boris, James, Vitalis, Langat, Wambururu, and Kevin) for exceptional field management and implementation.

1 Introduction

It is well established that firms have to design and enforce contracts that properly incentivize employees to behave in a way that maximizes firm profits. Under asymmetric information this becomes more challenging because firms cannot observe their employees behavior. Firms are constrained to adopt second-best contracts that reduce their profits and leaves rents for workers. In theory, firms can overcome this friction by investing in monitoring technologies that reveal the performance of their workers more accurately. As the these technologies become more affordable and more accessible, it becomes increasingly important to measure the extent to which they mitigate moral hazard in labor contracting, and how this can improve firm operations. This is particularly meaningful to establish in low-income countries because their economies are dominated by small firms that struggle to grow beyond a few employees. Moreover we expect that designing efficient contracts in the presence of asymmetric information is even more challenging in these environments because employees lack of capital means that firms are often held liable for damages to company assets, and weak legal institutions prevent companies from credibly sanctioning bad behavior.

The primary objective of this paper is to determine the impact of moral hazard on labor contracting, productivity, and firm profits in developing countries and the extent to which improved monitoring eases these frictions. Specifically, we investigate whether providing employers with real-time visibility into their employees' behavior generates changes to the contract, and whether these adjustments result in higher earnings for either party. We study this in the context of Kenya's public transportation system, which functions like many other informal transit systems worldwide (including Tanzania's dala dala's, Haiti's tap tap's and India's rickshaws among others). This setting is an important one to study because the sector is economically meaningfully in terms of the number of individuals it employs, and the amount of income it generates. In Kenya, estimates suggest that the industry employs over 500,000 people and contributes significantly to the country's GDP. Most importantly however, this context allows us to overcome major data constraints that have limited previous research in the space. Namely, we are able to collect detailed information on the contract terms set by the employer and the actions of the employee (their choice of effort and lying). We also introduce exogenous variation into the costs of monitoring in order to observe changes to the contract.

Studying the impact of monitoring on labor contracting is also interesting in this setting because *private* firms are responsible for providing crucial *public* services. Kenya's informal transportation system is dominated by small-scale entrepreneurs that own a few minibuses ("matatus") that run on designated routes. Estimates suggest that over 3,000 individual firms are responsible for managing the 15-20,000 minibuses that operate across Nairobi's 135 major commuter routes. These matatus are the only reliable form of public transportation and serve the majority of Nairobi's 4 million commuters daily. In this setting, it becomes important to document whether improved monitoring affects the quality of the public service they provide. To the extent that employees use the additional information to incentivize employees to maximize production, we might be concerned that the quality of their service could deteriorate. In the public transportation industry, the way transit vehicles operate directly affects commuters and passenger safety. While firms care about limiting dangerous driving behavior that can lead to accidents and scare passengers away, they are primarily driven to maximize revenue which could exacerbate unsafe driving practices. Our context is unique in that it allows us to collect detailed data on these externalities (speeding, dangerous driving maneuvers and accidents) and determine whether monitoring affects them or not. Because the safety externalities produced by Kenya's matatus are severe, understanding the implications of these technologies is important for institutions that want to promote better road safety.

The contracting environment we study is not unique to Kenya: the dynamics we outline here are prevalent in many other transportation industries, in agriculture and in the service industry. Two main dynamics characterize the space. First firms cannot observe the amount of revenue the driver collects (unobserved output) nor the amount of effort the driver invests (unobserved effort). The presence of unobserved output creates an additional friction that exacerbates the negative impacts traditionally associated with moral hazard in effort. Moreover, drivers in this setting are from relatively poor households and they cannot afford to walk away without pay on days when total revenue is low, nor can they pay for repairs when the vehicle is damaged (limited liability). Drivers are known to run away from accidents so they cannot be held accountable by the owner or the police. Together, these two features create an incentive for the driver to underreport total revenue and take shortcuts that may damage the vehicle in order to maximize their own utility. In light of these constraints, firms have overwhelmingly opted for a fixed rent contract (locally referred to as a "target" contract). The owners specify an amount that the driver must deliver by the end of the day net of fuel expenses. This contract induces the agent to invest effort because they are the residual claimant on revenue generated above the rental price. Moreover, the contract limits instances of under-reporting to days when the driver is unlikely to meet the agreed upon "target" price and walk away with their reservation wage. However, due to the limited liability constraint, the supply of effort will be less than the first-best outcome because the driver does not internalize the effect of his choice of over the full range of output.

Within this environment we introduce a new monitoring technology that allows owners to observe how productive and careful the drivers are being throughout the day. We developed our own device because available alternatives on the market were either too costly, or not sophisticated enough. Our device records and transmits via a mobile app the location of the vehicle, the number of kilometers driven, and the number of hours the ignition was on. While the owner will not know the number of passengers that boarded the vehicle, they can use this information to monitor driver's operations throughout the day and gain a more precise estimate for total daily revenue. We recruited 250 matatus in Nairobi to participate in the study, and 125 were randomly selected to be part of the treatment group. We only included owners that had 1 minibus, and managed their own vehicle in order to focus our analysis on the classic principal-agent relationship. The monitoring device was fitted to all the matatus in our sample, but the information produced by the device was only shared with treatment owners via mobile app.

In order to understand the impact of the device, we extend the classical principal agent model with unobservable effort to include unobservable output, which is a common feature in real world contracting relationships. Drivers choose the amount of effort and risk they invest and the amount of revenue they report to the owner in order to maximize their utility. Owners decide where to set the "target price". We model the technology as impacting the cost of under-reporting by increasing the probability the owner catches the driver lying and punishes them, thereby reducing the drivers' incentive to lie about the total amount of revenue they collected. The standard predictions from a model with unobserved effort are borne out in our model: contracts should improve when asymmetric information is reduced by inducing higher effort from the agent. However, our model also yields additional predictions on reporting behavior. We hypothesize that we should always see under-reporting of revenue (shading) below a certain threshold value, and this should decrease in the presence of higher monitoring.

We have suggestive evidence that owners use the information to restructure how they interact with their drivers, and improve upon the contract. We see that owners are more likely to reprimand drivers in earlier months as they use the information to correct bad practices. We do not see any impacts on the number of firings however - likely because the study's time span was too short. It also appears that receiving information from the monitoring device allows owners to modify the contract by decreasing the target they demand from drivers. According to our model owners recognize that the technology reduces driver utility by making lying more costly and compensate them by reducing the target. Empirically, we find that owners steadily reduce the target throughout the study period. By the last month of the study the target is approximately 4.1 percent below where it started. While the downward trend is prominent, the coefficient does not become statistically significant, likely because contract change of this magnitude in an environment where norms are so engrained takes time.

We find that drivers respond to this change as theory would predict. Namely, they invest more effort in driving the minibus, they drive in a way that is less damaging to the vehicle (less risky), they lie less, and they are more likely to make the target. We see effort, measured by the number of hours supplied, increase by 12 percentage points. When breaking up the effects along their extensive and intensive margins, we see that the increase in effort is primarily being driven by an increase in the probability drivers are on the road (by 6-8 percentage points). We also see drivers taking better care of the vehicle. We find damages decrease significantly, which we hypothesize comes from fewer instances of driving on alternate-routes that are bumpy (the distribution of vertical and lateral acceleration shift towards less "erratic" behavior in the treatment group). Finally, we see drivers lying significantly less, stemming from a reduction in the amount they lie when they report revenue below the target, and from making the target more often (when they make the target their is no incentive to lie because they get to keep everything they earn).

Overall we find that these changes lead to an increase in firm profitability. The gains in firm profits more than offset the cost of the device, suggesting that a tracking device like the one we designed for this study would be a worthwhile investment provided it was available on the market. Owners also report that monitoring their drivers has become significantly easier, and they trust their drivers more, thereby reducing some of the management costs that owners struggle with. Together these results suggest that treatment owners may have an easier time expanding their business than control owners. At endline, we asked owners about the number of vehicles in their fleet and we find that treatment owners have 0.154 more vehicles than control owners, suggesting that the introduction of monitoring technologies constitutes another important way small firms in developing countries can overcome barriers to growth.

This first set of results suggest that monitoring technologies can lead to changes in contracting terms, more effort, less risk, and larger profits for owners, all in ways consistent with agency theory. However, the presence of monitoring devices can also have effects outside of the firm. In public transportation systems, monitoring technologies are often used to check and limit instances of unsafe driving. Kenya's matatu sector is notorious for it's poor safety standards: drivers often swerve, stop suddenly, and turn sharply in order to collect more passengers. Our monitoring technology records these instances and conveys them to owners through a separate tab in the mobile app. A priori, it is not clear how the owners would use the safety information we provided. If owners internalized the dangers associated with unsafe driving, they might use the information to improve driving quality. However, to the extent that owners care only about revenue, and effort and social safety are substitutes, these outcomes could worsen over the study period, producing a negative externality for Nairobi's commuters. Despite all the of safety information we provided, the frequency of unsafe driving events flagged by the device do not change significantly, and instances of speeding remain the same. It follows that the gains to the firm do not come at the cost to commuters. However, this also suggest that external intervention may be necessary if safety standards are to be improved in an environment where accidents are common.

We test the efficacy of one such intervention by providing small cash incentives to drivers conditional on safe driving. This experiment is designed to mirror the actions that a regulatory body could potentially take in this setting (South Africa's Ethekwini municipality is testing one such intervention in the coming months). Our objective is to determine the effectiveness of an intervention that encourages the employees (drivers) rather than the employers (owners) to internalize the externalities generated by the business. Half of the drivers in our sample were randomly selected to receive 600 schillings at the beginning of each day, but incurred a penalty of 100 KES for every safety violation they committed (more than 6 violations resulted in a payout of 0). We find that the cash treatment meaningfully reduced safety violations committed by drivers, confirming that third party intervention is necessary to address these firm externalities. However, these effects did not persist after the removal of the cash incentives, suggesting that further action or permanent regulation is needed to induce long-lasting change.

This paper contributes to four different literatures. First, the paper speaks to the vast theoretical work on principal agent relationships and contract formation. The model we develop extends the standard principle agent models with unobserved effort (Laffont and Martimort (2002)) to include unobserved output. This feature is prevalent in many informal transportation systems worldwide. It also characterizes relationships in agriculture where absentee landlords cannot supervise their tenants; in the service industry where employers cannot record the number of services provided by their employees; and in businesses where inventory is difficult to monitor. To our knowledge there is only one other paper that documents the implications of moral hazard in effort and output (de Janvry and Sadoulet (2007)). The paper derives the optimal contract within the set of possible landlord-tenant sharecropping contracts, and finds that it is second best in that it allows for residual moral hazard in both effort and output reporting. In contrast, our model takes the fixed rent contract as given and investigates how the introduction of monitoring technologies allows the principal to change the parameters of the contract.

Second, this paper documents profit gains among companies that adopt monitoring technologies. Empirical evidence on the impacts of monitoring is limited because shirking behavior is hard to detect by design, firms' decision to monitor is often not random, and data on firms' operations are difficult to obtain. A handful of papers overcome these limitations by randomizing the introduction of additional monitoring technologies and/or exploiting the roll-out of new monitoring schemes within an industry. For the most part these papers are concentrated in developed countries, and point to the efficacy of monitoring technologies. Nagin et al. (2002) find that employees cheat less when the probability of being detected through random audits increases. Hubbard (2000), Hubbard (2003), Baker and Hubbard (2003), and Baker and Hubbard (2004) use detailed survey data to examine how tracking devices have changed the way the trucking industry operates in the United States.[1]

While these papers document important findings, we have reason to believe the impacts of monitoring could be different in a developing country. First, management quality is different (Bloom et al. 2016), and employers might not use the information as effectively. Second, employers face additional frictions that may limit their ability to use the information: contracts are not binding, and law enforcement is weak. There is only one other paper to our knowledge that investigates the impact of monitoring on employee behavior in a developing country. de Rochambeau (2018) studies the use of GPS devices by managers in Liberia's long range trucking industry. She finds that low-performing drivers increase their efforts dramatically in response to being monitored, while high-performing drivers react by taking less good care of the asset, and breaking company rules more often.[2] Our paper focuses on a different contracting environment where output is unobserved and owners must rent the asset out to the drivers. This more closely resembles informal public transportation systems within urban settings across the developing world. Our results stress the

[1]Hubbard (2000) finds that the returns to monitoring are higher when drivers operate on long-haul routes with infrequent stops (because they are more likely to speed and take longer breaks to recover the time - which isn't in the firm's best interest); Hubbard (2003) documents higher capacity utilization (loaded miles per period in use) among vehicles with on board computers (OBC); Baker and Hubbard (2003) find that shippers are more likely to use their own vehicles when they have access to monitoring technologies; Baker and Hubbard (2004) demonstrate that drivers working for companies with OBCs are less likely to own their own trucks because they can with the introduction of OBC's because the owner can ensure the driver preservers the truck value without selling them the asset.

[2]Note that there are additional studies that document the impacts of monitoring in a developing country context. Duflo, Hanna, and Ryan (2012) find that teacher absenteeism in India decreases when their attendance is carefully monitored, while Björkman and Svensson (2009) demonstrate that community health workers exert more effort when their performance is scrutinized by the community. While these papers shed light on individuals' response to incentives, they do not study the employer-employee relationship within the firm.

impact of the technology on employers' ability to change the parameters of the contract, and induce behavior that generates higher profits for the company. This suggests that monitoring can benefit firms despite the other frictions that prevail in these informal industries.

Third, our conclusions identify another channel for boosting firm productivity, namely the effective management of company employees through monitoring. These results add new insights to a growing literature documenting the challenges that firms face in developing countries. The work focuses primarily on small firms because of the important contributions they make to these countries' economies. According to the World Bank they contribute up to 60% of total employment and up to 40% of national income (GDP) in many developing nations (World Bank). In Kenya, where this project takes place, 80% of jobs are created in the informal sector, which is dominated by small firms. While researchers have documented a number of challenges facing firms, more recent work has focused on the frictions generated by poor management practices. A large survey of small firms across 7 different countries highlights large variation within countries in terms of the business practices used by small firms, and demonstrates strong correlations between better management practices and firm profits/survival (McKenzie and Woodruff (2016)). Researchers have investigated the effectiveness of different training programs designed to improve firms' management practices. Generally they find that training has small (0.1 or 0.2 standard deviations) or insignificant effects on firm level outcomes (Berge, Bjorvatn, and Tungodden (2014), de Mel, McKenzie, and Woodruff (2014), Valvidia (2012)). In contrast, our results suggest that monitoring technologies may constitute a new policy instrument for helping firms grow in the developing world. They represent a particularly promising solution because they are rapidly spreading throughout low income countries and becoming increasingly affordable.

Finally, our results suggest that the quality of service provided by private firms does not deteriorate with the introduction of monitoring. This is consistent with some of the existing work on the provision of incentive schemes (Duflo, Hanna, and Ryan (2012)). Equally important to note, however, is that the quality of service does not improve, despite the breadth of information we provided to owners about unsafe driving maneuvers. The only way we could successfully induce better driving in this environment was to incentivize drivers along this dimension. This is an important issue for policymakers who struggle to promote safe driving in urban hubs across the developing world.[3] In recent years, international institutions have provided funding, knowledge and technical assistance to build systems to reduce the number of traffic injuries and deaths worldwide (World Bank (2014)). These efforts are inherently hard to evaluate because the investments are multi-faceted and typically rolled out across an entire city. One exception is a program that was launched by Habyarimana and Jack (2015) in Kenya, which placed stickers inside Nairobi's matatus to encourage passengers to complain to their drivers about unsafe driving. They find that the intervention reduced insurance claims (filed when there is an accident) by 25-30 percent. Our paper takes a different approach by implementing an intervention that a regulatory body could poten-

[3]According to the Global Status Report on Road Safety, 1.24 million people are killed in traffic accidents each year and 90 percent of these deaths occur in low- and middle-income countries (LMICs) World Bank (2014)

tially enforce.[4] We find that the program results in improved road safety conditions by encouraging drivers to reduce their speeds and instances of sharp breaking. This suggests that investments in technologies that monitor and punish unsafe driving can be effective so long as they are properly enforced.

The remainder of this paper is organized as follows. Section 2 discusses Kenya's public transportation system, the prevalence of moral hazard, and the scope for monitoring. Section 3 reviews how the experiment was implemented in the field, data collection, and basic characteristics of the sample. We present a simple theoretical framework in Section 5. The purpose of the model is to generate predictions about how monitoring should affect the principal-agent relationship with unobservable output and unobservable effort. Section 6 discusses each our results. We then provide an overview and discuss the implications of the findings in the final section.

2 Context

2.1 Nairobi's Matatus

Nairobi's transportation system was developed after independence in 1963 (Mutongi (2017)). Small-scale entrepreneurs responded to the growing demand for mobility by retrofitting old vehicles and transporting passengers from the suburbs to the urban center. The buses were labelled "matatus", meaning three in Kikuyu, in reference to the early ticket price in KES of a matatu ride. These private businesses were legalized in 1973, but remained largely unregulated until 2003 when the government passed the Michuki rules, requiring that buses install speed governors, safety belts, and ensure that all drivers exhibit valid licenses (Michuki (2003)). To date these regulations are rarely enforced. In 2010 the Ministry of Transport issued a new directive to further formalize the industry, and eliminate the presence of gangs that were becoming increasingly active in the space. This required that all minibus owners form or join transport Savings and Credit Cooperatives (SACCOs) or Transport companies (McCormick et al. (2013)). To this day any industry newcomer must first register with a SACCO or company before they can put their vehicle on the road. Transport companies are much rarer in Nairobi and manage buses on behalf of individual investors. SACCOs on the other hand leave the daily management of the vehicle to the owner, but facilitate centralized organizational activities including scheduling, resolving internal disputes between owners, ensuring compliance with the National Transport and Safety Authority (NTSA) regulations, and providing financial services to owners and drivers.

This informal network of buses constitutes the only dependable transit system in Nairobi, and the city comes to a near standstill on days where drivers strike. To this day it remains almost entirely locally owned: private entrepreneurs purchase 14 or 33 seat minibuses, which they register with an existing SACCO. Rough estimates suggest that 15,000 to 20,000 buses currently circulate throughout the city, swerving on and off the road to collect passengers along their designated route. The presence of severe competition within a route explains the dangerous and reckless driving that

[4] A municipality in South Africa recently adopted this approach

prevails throughout the industry. According to the World Health Organization's Global Status Report on Road Safety, approximately 3,000-13,000 people die annually from traffic incidents where at least 30% of cases involve matatus (WHO (2015)). Conditions have not improved measurably in recent years. However, in an effort to combat negative stereotypes, matatu owners are increasingly investing in the comfort of their vehicle, the aesthetic (colorful interior and exterior), the quality of the "experience" (helping passengers on and off the bus), and the perks (TV's) (Reed (2018)). The more attractive and comfortable vehicles can charge up to twice the price of regular ones. Matatu fares vary between 0.5 and 1.5 USD for travel inside the city center, and between 1 to 5 USD for trips to the outskirts.

Finally it is worth emphasizing that informal transportation systems like this one are not unique to Kenya. In many low income countries governments cannot afford the investments required to build centralized bus or train networks that are regulated by the transport authority. In these environments, small private operators (including minibuses, vans, taxis, station wagons, three-wheelers) become responsible for all commuter traffic by default (McCormick et al. (2013)). These features are present for *peseros* in Mexico, *jeepney's* in the Philippines, *tuk-tuks* in Indonesia, *rickshaws* in India, *dala-dala's* in Tanzania, among others (Mutongi (2017)). It follows that the findings established here will have broader implications beyond Kenya.

2.2 Driver and owners

The relationship between matatu owners and their drivers is a principal agent relationship with asymmetric information. Owners rent their vehicles to a driver for an agreed upon "target price" (henceforth referred to as the 'target'), which is usually approximately 3000 KES (30 USD). Unlike the taxi systems in many high-income countries, the driver is expected to deliver this amount at the *end* of the day once all the fares have been collected. This is primarily because drivers have limited capital and cannot afford to pay the amount up front. Drivers are the residual claimants in this contract and keep everything they earn above the target. The owner is not allowed to revise the terms of the contract and claim more revenue if the driver had a good day. In the event that the driver cannot make the target, they must provide a justification to the owner. If they fail to make the target too many times they will be fired. Drivers can choose the number of hours they work, their driving style, and the amount of revenue they deliver to the owner. Owners cannot observe their driver's actions throughout the day, and must resort to costly interventions to check in on their drivers. This includes phone calls, dropping by the terminus of the route and staging someone at various stops to monitor whether the bus drives by.

This contract structure appears to be the only viable alternative in the industry. A fixed wage payment is unattractive to most owners because drivers face incentives to undersupply effort when they cannot be monitored.[5] The few SACCOs that have adopted this payment scheme have also hired full time managers who supervise the drivers closely. The traditional sharecropping

[5] Driver's also mentioned they disliked this remuneration scheme because it eliminates the large windfall days they receive when they are the residual claimant

model is also non-existent because owners do not observe the amount of revenue generated in fares, which means that drivers can consistently underreport the amount they collect. The cost of under-reporting is low because drivers can easily hide undisclosed revenue.

2.3 Device (Hardware and Software)

Monitoring technologies are becoming widely available in many developing countries, including Kenya. The majority of long-range bus companies that travel between the country's main cities are equipped with tracking devices. Moreover, some banks in Nairobi recently announced that they would only issue loans for minibuses whose location could be tracked with a device. Despite their availability, most medium range buses and inner-city public transportation vehicles are not yet using them. When asked why, most vehicle owners cite the high cost of sophisticated tracking systems (approximately 600 dollars for the tracker and additional monthly installments for system access), or the lack of detailed information provided by the cheaper alternatives.

To fill this need, the research team created a new monitoring system for city buses that is considerably cheaper, more flexible and more powerful than traditional tracking devices. The physical tracking units were procured for 125$ from a company in the United States (CalAmp). They feature GPS, internal back-up battery packs, 3-axis accelerometer for motion sense, tilt and impact detection. The device was designed to capture and transmit the information we required, including the 95th percentile and average forward/backward/lateral/vertical acceleration, as well as the 95th percentile and average forward/backward jerk. The device was also calibrated to generate alerts for every instance of vehicle speeding, over-acceleration, sharp braking and sharp turning. These safety alerts were calculated by an internal algorithm built into the CalAmp device with threshold parameters as inputs, using the full sequence of acceleration and speed data to identify unsafe driving actions. Further processing of the CalAmp system data on the server provided additional measures of interest including the total number of kilometers traveled that day, the total time the matatu was running, and a safety index (from aggregating the day's safety alerts). Finally, an API call was generated each time the owner used the app to request data from the server. These calls were recorded in the database and provided a measure of the owner 's usage of the app. In this way, we could track which types of information the owner found most valuable and how often the owner requested this information.

The data captured by the CalAmp device was transmitted to owners via a mobile application that was specifically designed to present information simply. The app (referred to as "Smart-Matatu") provided information in three ways. The first tab was a map of Nairobi and presented the real-time location of the vehicle. The owner could also use the map to find out where the matatu had been throughout the day. By entering a specific date and time interval into the phone, the app would draw out the exact routes traveled by the matatu over this time period. The second tab displayed all the safety alerts that were captured by the device. The owner could click on the safety event to find out when and where it occurred on the map. The final tab conveyed a summary of the driver's productivity and safety. The productivity section of this page listed the total mileage covered, and the duration the ignition was turned on that day. The safety section of this page provided the owner with an overview of the number of safety violations that occurred that day, as well as the driver's daily safety rating relative to all other drivers on the road that day (where a thumbs up appears for scores of 60% and above, a sideways thumb for scores between 40% and 60%, and a thumbs downs for scores of 40% and below).

3 Experimental Design

3.1 Sample Recruitment

We conducted an extensive recruitment drive in late 2015 by contacting 61 SACCOs that were operating on various routes across the city. We organized several large meetings with matatus owners in each SACCO, presenting the study's goals and methodology. We also continuously relied on referrals from the matatu owner's we were interacting to increase awareness about the project. All matatu owners were informed at the time of recruitment that a monitoring device would be placed in their vehicle free of charge and that they would be required to *provide* daily information about their business operations. We also mentioned that a random subset of owners would be selected to *receive* information from the tracker via a smartphone app for a six month time period, while others would have to wait 6 months before gaining access to the information for a shorter 2 month period. Recruitment lasted for 4 months in order to secure enough participants.

Owners who expressed interest in the study during the recruitment drive were contacted again over the phone to confirm their willingness to participate in the experiment, and to check that they met the three study requirements. First, owners were required to manage their matatu themselves, as opposed to hiring a third party manager. Second, owners could only own a single matatu at the time of recruitment. Third, the owners could not be the primary driver of the vehicle. According to an exploratory survey we conducted in the pre-pilot phase of the experiment, approximately one quarter of matatu owners in the general population met these three criteria (if we allowed owners to posses two or three matatus, over 50% of matatu owners satisfied these conditions). We set these constraints to focus the research on the classical principal-agent relationship with one owner and one driver. Failing to meet any one of these requirements would have invalidated our attempts to understand contract formation and principal's use of monitoring information.

3.2 Installations

The first installation took place in November 2016, and continued until April 2017. The field team managed by EchoMobile scheduled a time to meet every owner and install the device in their vehicle. Every owner in the study was offered a one-time payment of 5000 KES (50 USD) to compensate them for the time spent off-road to perform the installation of the device; and a new Android phone (worth approximately 80 USD) to ensure they could access the SmartMatatu app. The installation process was rather complex, requiring a team of 3 staff (an enumerator, a field manager, and an engineer) to meet the owner and driver of the matatu at a predetermined time and location. While the mechanic worked on fitting the device in the matatu, the field manager took the owner aside to re-explain the purpose of research project and the tracking devices' functionality. For owners in the treatment group, the manager conducted the additional app training. While the fleet manager explained these processes to the owner, the enumerator administered the baseline survey to the driver in a separate location outside of the owner's earshot so that the driver felt comfortable answering the questions honestly. Once the field manager finished the training with the owner, and the enumerator finished administering the survey to drivers, they switched. The field manager then took the driver aside to explain that they would receive a daily SMS to elicit information about the day's operations and to emphasize that all the data they shared would be kept confidential. Meanwhile, the enumerator conducted a 20-minute baseline survey with the owner. This whole installation process took approximately 1 hour to complete. The field manager shared his contact information with the owner and the driver so they could contact him with any further questions they had.

3.3 Treatment Assignment

The first treatment arm is referred to as the "information treatment". Owners in our sample were randomly allocated to a treatment and a control group. Owners in the treatment group were provided with free access to the data produced by the monitoring device immediately after installation. Owners in the control group were informed that they would receive the same access six months after the device was installed. During the device installations our field manager spent an additional 30 minutes with treatment owners explaining the types of data that would visible on the SmartMatatu app. A small survey was administered with the owners at the end of training to make sure they knew how to find all the information contained in the app.

Four months after the information treatment was launched we introduced a second treatment arm, referred to as the "safety" treatment. We selected half of the treatment drivers and half of the control drivers and offered them cash incentives to drive safely. This arm was designed to simulate the role of a functioning regulatory system and monetize the tradeoff between revenue and safety that drivers face. The cash incentive drivers were again randomly split into two groups: a one-month treatment group and a two-month treatment group (which we did not inform drivers of). This was done so we could study whether any changes in driving behavior that might be induced by the incentives persisted once they were removed. The specific incentive amount they received was determined by a safety rating, calculated daily for each driver in the following way. We analyzed two weeks of data for each driver (dropping days with less than 30km), tracking 1) the number of alerts of each type k (speeding, heavy braking, sharp turning and over-acceleration), and 2) the number of hours worked. For each driver, day and alert type we computed the rate of violations by dividing the number of alerts of type k on a given day t for driver i by the number of hours worked a given day t for driver i. For each driver i, we then construct a distribution of these rates for each alert type k and found the percentile that day's alert rate falls in. We then calculated the weighted average percentile for driver i on day t, by adding the alert rates for each type, applying weights of 1/3 for over-speeding and breaking, and 1/6 for over-acceleration and turning. The average lies between 1 and 100, and for each driver on each day, we then assess the cutoff below which they fall and disburse their incentives accordingly.

4 Data Collection

We collected data from three different sources. The first data set is a panel of daily responses from owners and drivers which we gathered through the app and SMS surveys, respectively. Next the enumerators conducted 8 monthly surveys, beginning with the baseline, followed by 6 monthly surveys and wrapping pup with the endline. Finally the GPS tracker collected a wealth of data that we use to measure safety violations committed throughout the day.

4.1 Non-system application variables

The SmartMatatu app was also designed to collect information from owners. Collecting accurate data can be very challenging in these settings, and this system was created to improve the quality of the data we received. Owners in the study were reminded daily via a push notification to report on that day's business activities through a form located on the app. They were asked to submit data on: the "target" amount assigned to their driver at the beginning day; the amount the driver delivered to the owner; any repair costs incurred and how much of those repair costs were paid by the owner directly; an overall satisfaction score for their driver's performance (bad, neutral, good); and whether the driver left employment that day (either voluntary or involuntary). Once the report was successfully submitted, owners received 40KES in airtime. We collected similar information from drivers through SMS surveys (because the drivers were not provided with a smartphones). At the end of every work day around 10pm, drivers would receive a text message asking whether they were ready to respond to the survey. Once they agreed individual text messages were sent to the driver asking for: the total revenue the matatu collected from fares that day; the amount they spent on fuel; and their "take home salary" (their residual income after expenses and paying the owner). Once the driver responded to all the questions, they were sent 40KES to incentivize consistent reporting.

We developed a set of processes for checking and validating the daily data we receive from owners and drivers. Echomobile wrote code to check for anomalies including outliers and entries that did not make sense and/or suggested the owner/driver may not have fully understood how to answer the question. A team of enumerators would then follow up with owners and drivers over the phone about each one of these entries. In cases owners and drivers were able to justify their responses, the enumerators would record their justifications in an excel spreadsheet. The necessary changes were made when the data needed to be corrected.

4.2 Monthly Surveys

We conducted eight rounds of surveys. We first administered the baseline surveys during the tracker installation. The *owner* baseline survey collected detailed information regarding basic demographics, employment history, features of the matatu, and their relationship with the current driver. Similarly the *driver* baseline asked about driver demographics, experience as a driver, unemployment spells, and their relationship with the current owner. For both owners and drivers we measured cognitive ability through raven's matrices. We also used games to gauge drivers' risk aversion and driver/owner propensity to trust one another. To measure risk we asked respondents whether they would prefer to receive 500 KES for certain or play a lottery to win 1500 KES. The game was repeated multiple times, with increasingly favorable lottery odds. The trust game on the other hand presented owners with 500KES and asked them to select a certain amount to be placed back in an envelope. They were informed that this amount would be tripled and delivered to a random matatu driver who was then going to decide how much to keep for himself and how much to return to the owner. The amount they chose to place in the envelope was recorded in the survey. When playing the game with drivers, we first presented them with an envelope containing 900KES. This amount was standardized across all drivers to ensure they faced the same choice. The drivers were informed about the owner's decision and how this amount was then tripled. The drivers were asked to return however much they wanted to the same owner.

We proceeded with 5 monthly follow up surveys. The monthly surveys were administered with three purposes in mind. First, they provided an opportunity for enumerators to follow up regularly with matatu owners and drivers and address any questions they might have about the device. Second, they were used to remind both parties to continue to submitting the daily reports in the SmartMatatu app. Finally, they were designed to track changes in the owner-driver relationship (how much monitoring had taken place, how satisfied the owner and driver were with the others' management/driving respectively), and any large expenses that may have occurred during month. As owners and drivers reached the 6-month mark, we conducted an endline survey to measure changes in key outcomes and assess the impact of the SmartMatatu app's information stream and the cash incentives on owners and drivers respectively.

4.3 Tracking data

The CalAmp tracking device transmitted high frequency data on forward/backward/lateral/vertical acceleration, jerk, location and a timestamp. For analysis purposes we aggregate this data to the day level. The tracker subsequently fed the raw data into an algorithm that computed the number of safety events that occurred in a 30 second time frame. Thresholds were calibrated for the Kenyan roads so as to avoid recoding an unreasonable number of safety violations and losing credibility among owners. These events included instances of speeding, over-acceleration, sharp braking, and sharp turns. The data was then further aggregate on the backend to produce daily reports on the number of safety violations, which is what we use for our analysis.

5 A Principal-Agent model with unobserved output

The model was built to reflect the actual contracting structure owners and drivers use, and to generate key predictions about the impacts of monitoring on contracts and employee behavior. The owner (principal) sets a target amount T at the beginning of the day, which they expect the driver (agent) to deliver by the end of the day. The target is chosen to maximize owner expected profits. The driver chooses their level of effort e to maximize their own returns, and the days events unfold. The driver earns revenue q from passenger fares, and decides how much to report to the owner $\tilde{q}$. We refer to the difference between q and $\tilde{q}$ as the "shade" amount. The model is solved using backward induction.

5.1 Step 1: Driver's optimal reporting choice

The driver can choose to report above the target ($\tilde{q} > T$) or below the target ($\tilde{q} < T$). The driver incurs a punishment from the owner on days where they report below the target. We define $\hat{q}$ to be the owner's belief about true q, where $\hat{q} = q - \sigma$, and $\sigma \sim U\left(\frac{-1}{\alpha}, \frac{1}{\alpha}\right)$ represents the precision of the owner's estimate of true output. The owner's signal of true output is more precise the larger α is. It follows that a higher α means the owner is more likely to detect underreporting. The punishment the driver incurs is a function of 1) the difference between what the owner expected revenue to be and the reported amount (owners are less upset on days where the driver reports below the target and they know for a fact that conditions were difficult), and 2) the probability this difference is greater than zero:

$$E(\hat{q} - \tilde{q}) \cdot \Pr(\hat{q} - \tilde{q} > 0) = \frac{\alpha}{2}(q - \tilde{q})^2 + \frac{1}{2}(q - \tilde{q})$$

The driver chooses the amount he reports to the owner ($\tilde{q}$) to maximize his utility (which we assume to be linear). He can either choose to report above or below the target. When the agent chooses to report above the target, there is no incentive to lie because the agent keeps everything above the target and the owner can't renegotiate the terms of the contract. Therefore they report truthfully $\tilde{q} = q$ and their payoff is $q - T$. However, if they decide to report below the target, they need to choose $\tilde{q}$ to maximize their utility:

$$\max_{\tilde{q}} \qquad U^D = (q - \tilde{q}) - \frac{\alpha}{2}(q - \tilde{q})^2 - \frac{1}{2}(q - \tilde{q})$$

Solving for $\tilde{q}$ yields:

$$\tilde{q} = q - \frac{1}{2\alpha}$$

Where

$$\frac{\partial \tilde{q}}{\partial \alpha} = \frac{1}{4\alpha^2} > 0$$

Proposition 1 *The optimal amount to report is $\frac{1}{2\alpha}$ below true revenue. In other words drivers shade by a constant amount.*

The reasons for underreporting are twofold. When the driver fails to make the target altogether $q < T$, they are expected to hand everything to the owner and walk away with 0. Clearly this is suboptimal and they should lie in order to walk away with some money. When the amount of revenue is slightly above the target $q > T$, they also face an incentive to lie in order to walk away with slightly more income than what they otherwise would if they reported truthfully and had to hand over the target.

Proposition 2 *The amount of shading increases if the probability of detection rises.*

Next we need to determine the point at which the driver is indifferent between reporting above the target (and telling the truth) and reporting below the target (and shading). When I tell the truth i.e $\tilde{q} = q$, I get utility:

$$(q - T)$$

When I lie "optimally" i.e $\tilde{q} = q - \frac{1}{2\alpha}$, I get utility:

$$(q - \tilde{q}) - \frac{\alpha}{2}(q - \tilde{q})^2 - \frac{1}{2}(q - \tilde{q})$$

Setting the two utilities equal and solving:

$$q^* = \frac{1}{4\alpha} + T$$

Proposition 3 *As the probability of detection increases, the revenue required to truthfully report and make the target (q^*) will fall. In other words drivers will make the target more often in a window of revenue around q^*.*

Proposition 4 *As the owner decreases the target, the revenue required to truthfully report and make the target (q^*) will decrease. In other words drivers will make the target more often in a window of revenue around q^*.*

5.2 Step 2: Driver's optimal choice of effort

The driver chooses effort to maximize his utility

$$\max_{e,r} \quad \underbrace{E[(q - T) \mid q \geq q^*] \cdot Pr(q \geq q^*)}_{Truth} + \underbrace{E[(q - \tilde{q}) - \alpha(q - \tilde{q})^2 \mid q < q^*] \cdot Pr(q < q^*)}_{Shade} - h(e, r)$$

where $q = e + r \cdot \varepsilon$

Which yields the following F.O.C with respect to e and r, respectively :

$$1 - F_\varepsilon\left(\frac{q^* - e}{r}\right) - h'_e = 0$$
$$\int_{\frac{q^*-e}{r}}^{\infty} \varepsilon f_\varepsilon(\varepsilon) d\varepsilon - h'_r = 0$$

Applying the implicit function theorem, we can derive the following two propositions:

Proposition 5 $\frac{\partial r}{\partial \alpha} < 0$ & $\frac{\partial e}{\partial \alpha} > 0$

This says that as the probability of detection increases, the driver will move away from behavior that the owner dislikes: namely risky driving behavior and supplying less effort. Supplying low effort results in low output, and now that the owner can monitor the driver more closely, the driver can no longer shade as much and walk away with the same take-home pay. This creates an incentive to increases their effort, and output. Similarly with higher α, the driver is going to have to truthfully report more often (q^* falls), which increases the instances when they have to bear the full cost of a bad outcome. Being exposed to greater downside risk makes risky behavior less attractive.

Proposition 6 $\frac{\partial r}{\partial T} > 0$ & $\frac{\partial e}{\partial T} < 0$

This says that as the owner increases the target, T, the driver will decrease effort and increase risk. If the owner increases the target, the driver is less likely to make it, which reduces the returns to effort. Conversely, raising the target encourages the driver to take on more risk because they are the residual claimant on days where they make above the target, but they do not incur the downside risk when they don't.

5.3 Step 3: Owner's choice of the target

Constrained case

The owner chooses T to maximize his utility:

$$\max_T \qquad T \cdot Pr(q \geq q^*) + E[\tilde{q} \mid q < q^*] \cdot Pr(q < q^*) \qquad \text{s.t}$$
$$E[(q-T)|q \geq q^*] \cdot Pr(q \geq q^*) + E[(q - \tilde{q}(q)) - \alpha(q - \tilde{q}(q))^2 \mid q < q^*] \cdot Pr(q < q^*) - h(e^*, r^*) > R$$

Which yields the following F.O.C with respect to T and λ, respectively :

$$\frac{\partial}{\partial T} = 1 - F_\varepsilon(\cdot) + \frac{\partial e}{\partial T}F_\varepsilon(\cdot) + \frac{\partial r}{\partial T}\int_0^{\frac{q^*-e^*}{r*}} \varepsilon f_\varepsilon(\varepsilon)d\varepsilon - \frac{\partial}{\partial T}\left(\frac{q^*-e^*}{r^*}\right)\left(\frac{1}{4\alpha}\right)f_\varepsilon(\cdot)+$$
$$\lambda\left[\underbrace{-(1-F_\varepsilon(\cdot)) + \frac{\partial e}{\partial T}(1-F_\varepsilon(\cdot)) + \frac{\partial r}{\partial T}\int_{\frac{q^*-e^*}{r*}}^{\infty} \varepsilon f_\varepsilon(\varepsilon)d\varepsilon}_{C'} - \underbrace{h'\left(\frac{\partial e}{\partial T}+\frac{\partial r}{\partial T}\right)}_{D'}\right] = 0$$
$$\frac{\partial}{\partial \alpha} = \int_{\frac{q^*-e^*}{r*}}^{\infty} (e + r\varepsilon - T)\, f_\varepsilon(\varepsilon)d\varepsilon + \int_0^{\frac{q^*-e^*}{r*}} \frac{1}{4\alpha}f_\varepsilon(\varepsilon)d\varepsilon - h(e,r) = 0$$

Applying the implicit function theorem, we can derive following proposition:

Proposition 7 $\frac{\partial T}{\partial \alpha} < 0$

This says that as the probability of detection increases, owners will reduce the target. We know that the owner is unambiguously better off with higher α: at every point below the target, the driver is offering more than he once did, and as q^* falls, the probability the owner will receive the target increases. However, the driver is unambiguously worse off because they can lie by less. Because the constraint binds, the owner needs to reduce the target to ensure the driver makes their reservation wage.

Unconstrained case

The owner chooses T to maximize his utility:

$$\max_T \qquad T \cdot Pr(q \geq q^*) + E[\tilde{q} \mid q < q^*] \cdot Pr(q < q^*)$$

Which yields the following F.O.C with respect to T

$$\underbrace{(1-F_\varepsilon(\cdot))}_{A'} + \underbrace{\frac{\partial e}{\partial T}F_\varepsilon(\cdot) + \frac{\partial r}{\partial T}\left(\int_0^{\frac{q^*-e^*}{r*}} \varepsilon f_\varepsilon(\varepsilon)d\varepsilon\right) - \frac{1}{4\alpha}f_\varepsilon(\cdot)\frac{\partial}{\partial T}\left(\frac{q^*-e}{r}\right)}_{B'}$$

Applying the implicit function theorem, we can derive following proposition:

Proposition 8 $\frac{\partial T}{\partial \alpha} > 0$

Without the driver's constraint binding, the owner can then increase the target because 1) when the driver shade's, they shade by less so it's less costly when they do, b) they can capture the higher revenue on days where the drivers make above the target.

6 Results

6.1 Baseline Characteristics

We work exclusively with matatu owners with one vehicle, which they do not operate themselves. They are approximately 40 years of age, and have completed 11 years of education. These small-scale entrepreneurs have spent an average of 8 years in the matatu industry, owning a vehicle for the past 4 years. While it is possible to have a salaried job and manage a matatu at the same time, only 20% of our sample juggles these two responsibilities. Typically owners have worked with their current driver for the past 2 years. Drivers have very similar profiles, which is to be expected because many owners were previously driving matatus themselves. They are a few years younger (35 years old on average), with slightly lower levels of education (8 years on average). They have worked in the industry for over a decade, driving a vehicle for the past 7-8 years (many start as conductors - the person in charge of collecting fares). They have worked with 5 different owners, averaging 1.5 years with each one. Both driver and owner characteristics are balanced across treatment and control groups (Table 1 and Table 2).

6.2 Usage

We monitor owners usage of the device through the API calls generated by logging into to the app, and requesting different pieces of information (historical location, updating of the summary page, checking where the safety violations occurred on the map). Figure 4 calculates the share of owners that made at least one API call during the week. We find high rates of take-up: 80% of owners are checking the app at least once a week towards the beginning of the study. This share decreases but stabilizes at about 70% as the study progresses. There is also a large share of owners using the app daily: 60% check it once a day at the beginning of the study, and 40% continue their daily usage after 6 months. This suggests that owners are actively engaging with the device throughout the study. Typically adoption rates of new technologies are much lower than what we observe here, suggesting the device fulfilled unmet market demand.

We also check whether owners are internalizing the information we provided. At endline we asked owners to state the revenue earned, the number of kilometers driven, the fuel costs, and the extent of off-road driving on the most recent day their vehicle was active. Owners had the option of answering "don't know". We find that owners in the control group are less likely to know about the the number of kilometers driven and the instances of off-route driving, both of which are registered in the app. We do not find any difference between the treatment and control groups regarding knowledge about revenue and fuel costs which cannot be garnered from the device. As a final test, we also asked owners to rate how challenging it was to monitor their employees on a scale from 1 (not hard) to 5 (very hard). Having access to the information reduces the reported difficulty level by just over 2 points (moving the overall rank from hard in the control group to easy in the treatment group). We do not find any significant changes in traditional monitoring behaviors (checking-in with the driver over the phone, at the stage, through a third party).

6.3 Information Treatment Arm

To study the treatment effect of information on contracts, productivity and safety, over the 6 month time frame we run the following regression model:

$$y_{irmd} = \alpha_d + \tau_r + \sum_{m=1}^{6} D_{im}\beta_m + X_i\gamma + \varepsilon_{irmd}$$

where y_{irmd} is a daily contract/productivity/safety outcome for owner i on route r, on day d, in month m, α_d is a day fixed effect, τ_r is an assigned route fixed effect, D_{im} is a treatment indicator equal to 1 if the owner is in the information group in month m (allowing examination of the treatment effect as it evolves over the six months of the study), X_i is a set of firm-level baseline specific controls included for precision, and ε_{irmd} is an error term.[6] We cluster our standard errors at the matatu level. Note that the study offered the information to control owners in months 7 and 8 (as compensation for participating in the study). As a result all the regressions only include data before month 7. We cluster our standard errors at the matatu level.

We also have an endline survey that asked owners about their perceptions of the driver's performance, their monitoring strategies, and their firm's size. To study the impact of our device on these outcomes we run the following regression model:

$$y_{ir} = \alpha_d + \tau_r + D_i\beta + X_i\gamma + \varepsilon_{ir}$$

where y_{ir} is an endline outcome for owner i on route r, τ_r is an assigned route fixed effect, D_i is a treatment indicator equal to 1 if the owner is in the information group), X_i is a set of firm-level baseline controls included for precision, and ε_{ir} is an error term.

6.3.1 Contracts

We first investigate whether access to the tracking information changes owner interactions with their drivers, and the *terms* of the contract. While the intervention could have also changed the *type* of contract they offered their drivers (fixed wage or sharecropping), extensive interviews with owners suggested this was unlikely to occur. The fixed wage contract is unpopular among owners and drivers, and the sharecropping model is difficult to implement when revenue is unobserved and can be easily withheld by the drivers. Moreover social norms are engrained in this industry, and a change of this magnitude would be unexpected in a 6 month time frame. We further confirmed in our endline survey that none of the owners in our sample had switched over to other contracting types. However, we do find changes in owner-driver interactions. We asked drivers to report the frequency with which they were contacted and criticized by the owner that month. Formal reprimands are not frequent but they are used by owners to signal their displeasure with the drivers behavior. Figure 5 suggests that the number of reprimands is marginally higher in the

[6] Controls include the matatu's age and number of features, as well as owner's age, education, gender, tenure in the industry, their raven score and the number of other jobs they have.

treatment group at the beginning of the study period, consistent with the idea that owners use the information to correct behavior early on. The frequency increases by approximately 20-30 % (off of a control mean of 0.5) in months 1-4 before decreasing significantly in month 6. We also investigated whether owners took more extreme actions and fired their drivers more frequently. While the trend in Figure 6 suggests that the number of firings increased in the second month of the study and decreased thereafter, this result should be interpreted carefully because there are so few firings in our data (17 in total).

Owners can also change the terms of the contract, namely the target. The target for 14 seater buses is usually set at 3000 KES. Discussions with owners confirm this is an industry standard that only fluctuates with good reason (they know that demand will be high or low that day because the weather or road conditions have changed). Charging much more would alienate drivers, and charging any less would cut into firm revenues. Figure 7 depicts the estimated treatment effect on the owner's daily target across the 6 months of the study. There are no significant changes in the first month, likely because owners were still learning how to use the app and experiment with ways to improve their business operations. In subsequent months however, we see the target steadily declining. By month 6, the target amount is 135 KES below the control group, representing a 4.5% decrease (0.2 standard deviation decrease). While the result is not statistically significant (likely because we are underpowered), the downward trend is clear. This steady reduction suggests that the information allows managers to re-optimize the terms of their employees contracts. Taking this result back to the model, it suggests that the drivers are operating at their participation constraints. Recall that if the participation constraint did *not* bind, we would have expected the target to increase in response to the information treatment. However, when the constraint binds, the owner needs to decrease the target in order to compensate the driver for their lost information rents. This decrease reduces owners revenue on days where the driver makes the target. As a result it is only profitable for the owner to do so if they are compensated in other ways, namely higher revenue shares on days the driver doesn't make the target, fewer damages to the vehicle or an increase in the frequency the driver makes the share. We turn to these results next.

6.3.2 Productivity

Drivers choose how much effort to supply, how recklessly they will drive, and the amount of revenue they disclose to the owner (which is either the target amount, or some amount below). We proxy driver effort by the number of hours the tracking device was on (the device powers on and off with the matatu). When the device is installed in the matatu, drivers know that they are more likely to get caught lying, which makes undersupplying effort more costly. As the model demonstrates, this should induce drivers to invest more effort throughout the day. In parallel, owners have decreased the target which means that becoming the residual claimant is more achievable. For both of these reasons we expect effort levels to rise. This prediction is borne out in our data: Figure 8 demonstrates the upward trend in effort that we anticipated. The number of hours the tracking device was on increases by 1.4 hours in month 2 and rises steadily until the end of the study. By month 6, effort levels (proxied by hours worked) increase by 2.1 hours in the treatment group. This represents a 15% increase in drivers' labor supply. With more hours on the road, we also see the number of kilometers increase by approximately 12 kilometers. We further investigate whether this increase in effort is occurring on the extensive or intensive margin. Figure 11 demonstrates that the number of days the device is on increases by 6 percentage points in month 5 and 8 percentage points in month 6. This represents an additional 1-2 days of work per month.[7]

Without the monitoring technology, drivers can choose how recklessly to drive with few repercussions. With the technology, however, owners can choose the amount of risk they want drivers to take because they observe more about their driving behavior. We hypothesize that owners will prefer less risk than what the drivers would optimally choose. With less risk, damages to the vehicle should be reduced. Figure 12 confirms this hypothesis in the data. We see damages decreasing substantially throughout the entire 6 month period. In month 2 daily repair costs for treatment owners are reduced by 100 KES, and continue falling until month 6 where they are 250 KES less than what control owners are incurring on average. This represents a 50% decrease in daily repair costs, which is significant, as these repairs constitute a major business expense for owners. To confirm that this result stems from less risky driving behavior, we investigate whether the distributions of lateral and vertical acceleration differ across the treatment and control groups. We might expect these measures to be different if drivers operate along different routes for example. Drivers often take shortcuts on bumpy roads that are notoriously damaging to matatus. These shortcuts are appealing to the driver because they help them get to the city center more quickly, and avoid traffic jams where they sit ideally without picking up any passengers. Drivers are also off the hook for any damages to the vehicle that result from taking these routes, which mean it is essentially costless to do so because the owners cannot observe them. We find suggestive evidence that driving behavior has changed. In the case of lateral acceleration (tilting from side to side) we see the distribution in the treatment group tightens around 0 (Figure 14). We can reject equality of these distribution functions by applying a K-S test which returns a p-value of 0.000. Similarly the distribution of vertical acceleration (bumping up and down) for the treatment group has more mass around gravity (normal driving).

It is also important to rule out any alternative explanations for these effects on repair costs. Specifically, it could be the case that drivers tend to inflate these repair-costs and the device reduces their incentive to do so because they are more likely to be caught in the lie. This cannot be the case for larger repairs, however, because these are incurred by the owner directly and/or will be validated with the mechanic. We therefore create an indicator for whether the repair costs exceed 1000 KES (80th percentile). Figure 13 demonstrates that the probability of incurring a large repair

[7] While this could be driven by fewer days in the repair shop, we have a few pieces of evidence that suggest this is not the case. First, we ask owners whether the vehicle is supposed to be on the road each day and we see no differences between treatment and control owners throughout the study. If the vehicle was being repaired more often we should see owners reporting more days *off* the road (they always know when it is the shop). Similarly, the driver does not report fewer days on the road, or taking fewer days off. Simultaneously we know that drivers lie about taking side-jobs and whether they are driving the vehicle. We hypothesize that treatment drivers are less able to lie about taking the matatu out on the road now that the owner can monitor them.

cost decreases significantly (7-8 percentage points). This implies that the decrease in the repair costs that we observe cannot be entirely driven by driver inflation.

Next we investigate drivers' reporting behavior and how this changes with the introduction of monitoring. According to the model we should see drivers under-reporting ("shading") below some optimum q^*, at which point they will start truthfully reporting and providing the owner with the target amount. The model also predicts that drivers will under-report by a constant amount below q^*. This is consistent with the idea that drivers have some reservation wage they do not want to fall below. Figure 16 confirms this prediction in the data. The graph summarizes values of shading at unique/binned values of net revenue above the target. We see that divers continuously shade approximately 700 KES (7 USD) until the net revenue they generate exceeds the target by approximately 500-1000 KES.[8] Next, we predict that the monitoring technology should decrease the amount of shading we see in the data. Owners can use the device to more accurately approximate actual revenue, and are therefore more likely to detect when the driver is underreporting. Drivers should respond to this increase in the cost of under-reporting by lying less everywhere below the threshold, and reducing the revenue threshold for truth telling. Figure 17 depicts the same shading graph split across treatment and control groups, to which we apply a non-parametric smoothing function. We see that for net revenue below q^* (which falls somewhere between 500-1000 KES), the treatment group shades less than the control group. Next, we regress the shade amount on treatment status for different possible q^* (between 500-1000), where the regression only considers data below q^* because this is where the model predicts shading will occur.[9] The results in Table 5 confirm that the amount driver shades falls by approximately 70-100 KES depending on the exact location of q^*.

According to the model this reduction in shading stems from q^* shifting down (the driver requires less revenue before they start truthfully reporting), and the $\tilde{q}$ falling (when they under-report, they do so by less). We investigate whether both of these behaviors are indeed happening in reality. We do so by imposing a step function in a regression of shade amount on treatment. In other words we allow shading below q^* and impose zero shading above, allowing this effect to differ by treatment status, just like the model would predict. We run this regression for every reasonable value of q^*. We then plot two outcomes in Figure 18. The dots represent the estimated shading amount in the treatment and control groups across different choices of q^*. We can see that the treatment groups shade by approximately 50-70 schillings less than the control group regardless of the q^* we impose on the model. Next, we plot the Mean Squared Error (MSE) of our regressions (dotted lines) to determine which q^* minimizes the MSE for the treatment and control groups respectively. The vertical lines represent the optimal q^* using this metric. This demonstrates that our best guess of q^* in the treatment group is 150 schillings below our best guess of q^* in the control group. This

[8] To get an accurate measure of shading we want to know the share of joint revenue that the driver withholds. In other words we need to know the income that the owner took home *and* and the salary of the driver. We therefore use net revenue above target on the x-axis, defined as owner income + driver salary − target

[9] The regression includes the standard controls and fixed effects. The regression also excludes data from month 1 because we know that owners were unfamiliar with the device in that first month. The magnitude of the results stay the same when we include month 1 but we lose some precision from the noise this month introduces.

confirms that both factors explain the overall reduction in shading behavior that we observed in the more flexible regression specification above.

Finally we look at the probability that drivers make the target, which depends on q^*. We just saw that $q*$ shifts down because of the higher penalty associated with under-reporting. From the model we also know that any reductions in the target (for which we have suggestive evidence) will also reduce the revenue required to truthfully report (q^*). A lower q^* implies that drivers should make the target more often on days where revenue is close to q^*. To investigate this prediction in the data, we first apply our standard regression specification to determine whether we see a significant increase in the rate at which treatment drivers make the target. Figure 19 suggests that from month 3 onwards, the rate at which drivers make the target increases by 11 percentage points off of a base of 44 percent (significant in month 3 only). It is not altogether surprising that the result is slightly weaker because the analysis considers the full range of revenue rather than focusing on days when drivers are close to q^* (i.e. close to making the target), which is where the model predicts we should see these effects. To investigate this further we calculate the average revenue above target on a route-month in the control group to get a sense of the usual revenue above target generated for a day.[10] We then take each drivers reported revenue above target and subtract the average expected amount.[11] We are left with the daily deviation from expected revenue above target, which we plot in Figure 20. The revenue above target measure has an approximate mean of 4,000 schillings. As such, -2000 KES on the graph implies that drivers only have 2,000 KES in revenue to cover their salary and their costs for that day. This results in a take-home pay of 500 to 1000 schillings, which is right where we expect q^* to be. Figure 20 demonstrates that the probability of making the target increases significantly at this point, which is exactly what we would expect. This represents a meaningful increase in "compliance" with the terms of the contract.

6.3.3 Company Performance and Employee Welfare

We now turn to investigating the impact of the monitoring device on firm performance. Specifically we are interested in determining whether the information we supplied allows companies to generate higher profits and ultimately expand their operations by adding more vehicles to their fleet. Company profits are measured by subtracting costs (repairs and driver salary) from total revenue. We documented substantial reductions in repair costs and, assuming drivers are at their reservation wage, we expect their salary to stay the same (Figure 23 confirms this is true). The model predicts that the impact on revenue, however, is ambiguous: better monitoring should increase driver effort, but reduce the amount of risk-taking behavior they engage in (which we also confirmed in our data). Depending on which of these effects dominates revenue could increase or decrease. Figure 21 illustrates that revenue does not change substantially throughout the study. Taken together, decreasing costs and stable revenues suggest that firm profits will increase. Figure 22 demonstrates

[10] We use gross revenue below average for this outcome instead of net revenue like we did for the shading amount because it only depends on drivers reporting, which means we have more data to work.

[11] This is akin to including route fixed effects, because we know that a certain level of revenue above target will be good on certain routes but not on others

a similar trend to what we've observed to date: profits increase continuously starting month 3, and peak at month 5. Specifically, treatment owners see their daily profits rise by approximately 12% in month 4 and 5 (440 KES). Taking the average gains over the study period and extrapolating to the full year (assuming the matatu operates 25 days a month), we can expect a 120,000 KES (1200 USD) increase in annual firm profits. It is worth mentioning that this profit measure does not take into account any additional gains from having to spend less time and effort monitoring the driver. The device cost 125 USD (including shipping to Kenya), which means that it would take less than 3 months for the investment to become cost-effective for the owner. This return on investment (ROI) suggests that these devices are likely to be welfare improving for owners in the short and long run. One of the reasons we have not see more matatu owners adopting them, however, is because they currently do not exist in this form on the market. The options are either much more expensive (approximately 600 USD and monthly installments), or have more limited capacity. Without having tested their efficacy, owners are hesitant to make the investment. The impact of this intervention is comparable to some of the more successful business training programs documented in the literature. The cost of these trainings range from 20 to 740 dollars and last a few weeks at most. Our technology has the added benefit of requiring a single up-front payment for continued use. Moreover it requires relatively little coordination and training. Finally our profit gains are in line with other estimates in the literature.

Are treatment firms also more likely to grow their business than control firms? We measure firm growth by the number of vehicles that owners have in their fleet at endline. A simple regression of this outcome on treatment with the standard controls reveals that treatment owners have 0.145 more vehicles in their fleet on average than control owners (Table 7). This represents an 11 percentage point increase in fleet size. While treatment owners were also more likely to make changes to their matatu's interior, this result is not statistically significant. We hypothesize that the monitoring device introduced a number of changes that encouraged treatment owners' to grow their businesses more actively. First, profits increased. Second, our results suggest that owners started trusting their drivers more. Table 6 presents four different measures of owner's perceptions of their drivers at endline. We see owner's sending an additional 30 KES to drivers in the trust game the enumerators administered (Column 1). Moreover, treatment owner's assessment of whether their driver's skills have improved, increases by 0.6 points (where they could be assigned a -1 for worse driving, 0 for no change, and 1 for better driving). Finally, treatment owners are more likely to report that their drivers have become more truthful (Column 3). We suspect that greater trust in their drivers abilities/honesty, combined with a reduction in the amount the drivers shade, makes the process of managing of the company easier.

Finally it is worth investigating whether these gains to the company come at the expense of their employees. While it is difficult to measure welfare, we consider three main outcomes that could impact driver's well-being: the amount of effort they supply, their salary and their relationship with the owner. We know the amount of effort they supply increases, and the amount they shade decreases. While their salary per hour remains unchanged, they are potentially worse off for working more hours. However, throughout the course of the study we did not receive any complaints from drivers, despite contacting them regularly to conduct our surveys. To investigate further, we created a small survey that we administered to drivers via SMS 6 months after the original study concluded (at this point we had given control owners 2 months with the information as well so no distinction can be drawn between info treatment and control drivers). Sixty percent of drivers responded (distributed evenly across treatment and control) with very positive experiences about the device: 27% said it improved their relationship with the driver (70% said nothing changed), 65% said it made their job easier (26 % said nothing changed), 96% said they preferred driving with the tracker, and 65% said it changed the way they drove. We do not want to lean too much on this qualitative evidence, but it does suggest that the drivers benefitted from the device as well. Some of the open ended questions reveal that drivers felt a greater sense of security with the device in their car, and they felt it increased owner's trust in their work.

6.3.4 Externalities

The device conveyed information to owners about productivity and safety. To the extent that owners contract explicitly over safety we might expect owners to set higher safety standards for their drivers. However, if owners care only about revenue, and increased effort comes at the expenses of safety, we might expect instances of unsafe driving to increase. This would lead to socially suboptimal behavior by the drivers. The device collected five pieces of information that correlate with safe driving: maximum speeds, speeding over 80km, acceleration, sharp breaking and sharp turning. We do not see increases in maximum or average speeds as the study progresses. Similarly instances of over-acceleration and sharp breaking do not change. We see no effect on sharp-turns or instances of speeding above 80km (which is difficult in Nairobi to begin with). Finally we tracked the number of accidents throughout the project. There are 41 accidents in total throughout the 6 month period, of varying degrees of severity. While the number of accidents trends upwards in months 4 and 5, it is difficult to conclude that accidents increase significantly. Overall the evidence points towards safety standards staying the same, despite the emphasis we placed on safety across all tabs in the app. While this confirms that owners can incentivize optimal levels of effort without further compromising passenger safety, we cannot expect owners to internalize the negative externalities produced by unsafe driving.

6.4 Cash Treatment Arm

Finally, we tested the impact of an intervention that incentivizes drivers to take safety into account. Drivers were offered 600 shillings at the beginning of the day, and incurred a penalty for each safety violation they incurred. The experiment was designed to mimic an intervention that a regulatory body could feasibly implement. We find that the cash treatment has no discernible effect on average speed, over-acceleration, and sharp turning. However, we detect large decreases in the instances of speeding and sharp braking. The number of sharp braking alerts deceases by 0.13 events per day, a 17% decrease relative to the control group. Likewise, the number of sharp braking events decreases by 0.24 per day, representing a 35% decrease. These results suggest that drivers can be incentivized to take safety into account. However the incentives must come from a third party, as owners are unlikely to induce similar changes in driving behavior.

In Table 9 we examine driving behavior among the group of drivers whose cash incentives were removed after the first month. The goal of this exercise is to examine whether the behavioral changes induced by the cash treatment persist after the incentives are removed. We see that the number of speeding events rebounds almost completely to pre-treatment levels, while the number of sharp braking events remains lower but insignificant. Overall, it appears that the behavioral effects of the cash treatment arm wear off after the removal of the incentives. This suggests that inducing better driving habits for a short time period may not be sufficient to see longer run improvement in safety outcomes.

7 Concluding Remarks

In this paper we design a monitoring technology tailored to the minibus industry in Nairobi. The device provides real time information about the productivity and safety of the driver to the owner of the minibus. We find that the monitoring technology eases labor contracting frictions by improving the contract that owners offer their drivers. The drivers respond by supplying more effort, driving in ways that are less damaging to the vehicle, under-reporting revenue by less and making the target more often. This results in higher profits for the firm. Treatment owners also report greater trust in their drivers, and find it less difficult to monitor them, which may explain why their businesses grow faster during the study. Despite the breadth of information we supplied on safety, we do not see drivers improving their performance along this margin unless they are explicitly incentivized to do so with small cash grants. While this suggests that gains to the company do not come at the expense of the quality of service they provide, the technology does not remedy the externality the industry produces to being with.

These results are important for firms, and for policy makers working to improve road safety conditions in urban hubs. We know firms struggle to grow in developing countries for a number of reasons, and this paper identifies another important barrier that a relatively low-cost intervention can help overcome. Monitoring is typically difficult in small firms because they cannot hire dedicated staff to oversee employee performance, and it takes time away from regular business operations. Monitoring technologies are becoming more accessible in these environments, however, and we demonstrate how they can benefit companies using rich data on contracts and employee behavior collected daily over 6 months from employers and employees alike.

We do not find that safety standards improve when information from the device is conveyed to owners. However, when the drivers are incentivized to drive more safely we see instances of speeding and sharp breaking fall. This suggests that simply introducing monitoring technologies, without further regulation, might not achieve the desired effects for governments trying to improve road safety. Local transport authorities in Nairobi and South Africa have already started to discuss ways of introducing remote tracking solutions throughout the transportation industry to help monitor and record the behavior of the drivers on the road. Our research suggests that while this will improve firm operations, more targeted interventions requiring regulatory oversight will be necessary if these devices are to induce safer driving.

This analysis highlights the need for further research estimating the long term impacts of these technologies on firm operations. Our study lasted 6 months, but we hypothesize that we would have seen greater changes in the terms of the contract, and in the *type* of contract being offered had we continued for an additional year. Finally, future research should also investigate how this information can be used to induce longer-lasting improvements in safe driving.

References

Baker, George P. and Thomas N. Hubbard. 2003. "Make Versus Buy in Trucking: Asset Ownership, Job Design, and Information." *American Economic Review* 93 (3):551–572.

———. 2004. "Contractibility and Asset Ownership: On-Board Computers and Governance in U. S. Trucking." *The Quarterly Journal of Economics* 119 (4):1443–1479.

Berge, Lars Ivar Oppedal, Kjetil Bjorvatn, and Bertil Tungodden. 2014. "Human and Financial Capital for Microenterprise Development: Evidence from a Field and Lab Experiment." *Management Science* 61 (4):707–722.

Björkman, Martina and Jakob Svensson. 2009. "Power to the People: Evidence from a Randomized Field Experiment on Community-Based Monitoring in Uganda." *The Quarterly Journal of Economics* 124 (2):735–769.

de Janvry, Alain and Elisabeth Sadoulet. 2007. "Optimal Share Contracts with Moral Hazard on Effort and in Output Reporting: Managing the Double Laffer Curve Effect." *Oxford Economic Papers* 59 (2):253–274.

de Mel, Suresh, David McKenzie, and Christopher Woodruff. 2014. "Business Training and Female Enterprise Start-up, Growth, and Dynamics: Experimental Evidence from Sri Lanka." *Journal of Development Economics* 106:199–210.

de Rochambeau, Golvine. 2018. "Monitoring and Intrinsic Motivation: Evidence from Liberia's Trucking Firms." Working Paper, Columbia University.

Duflo, Esther, Rema Hanna, and Stephen P. Ryan. 2012. "Incentives Work: Getting Teachers to Come to School." *American Economic Review* 102 (4):1241–1278.

Habyarimana, James and William Jack. 2015. "Results of a Large-Scale Randomized Behavior Change Intervention on Road Safety in Kenya." *Proceedings of the National Academy of Sciences of the United States of America* 112 (34):E4661–4670.

Hubbard, Thomas N. 2000. "The Demand for Monitoring Technologies: The Case of Trucking." *The Quarterly Journal of Economics* 115 (2):533–560.

———. 2003. "Information, Decisions, and Productivity: On-Board Computers and Capacity Utilization in Trucking." *American Economic Review* 93 (4):1328–1353.

Laffont, Jean-Jacques and David Martimort. 2002. *The Theory of Incentives*. Princeton University Press.

McCormick, Dorothy, Winnie Mitullah, Preston Chitere, Risper Orero, and Marilyn Ommeh. 2013. "Paratransit Business Strategies: A Bird's-Eye View of Matatus in Nairobi." *Journal of Public Transportation* 16 (2).

McKenzie, David and Christopher Woodruff. 2016. "Business Practices in Small Firms in Developing Countries." *Management Science* 63 (9):2967–2981.

Michuki, John. 2003. "The Traffic Act."

Mutongi, Kenda. 2017. *Matatu, A History of Popular Transportation in Nairobi*. The University of Chicago Press.

Nagin, Daniel S., James B. Rebitzer, Seth Sanders, and Lowell J. Taylor. 2002. "Monitoring, Motivation, and Management: The Determinants of Opportunistic Behavior in a Field Experiment." *American Economic Review* 92 (4):850–873.

Reed, Rachel C. 2018. "Transportation Turned Performance Art: Nairobi's Matatu Crews." *New York Times* .

Valvidia, Martin. 2012. "Training or Technical Assistance for Female Entrepreneurship? Evidence from a Field Experiment in Peru." Tech. rep., GRADE.

WHO, Geneva. 2015. "Global Status Report on Road Safety." Tech. rep., World Health Organization, Geneva.

World Bank, DC. 2014. "Global Road Safety Facility (GRSF) Strategic Plan 2013-2020." Tech. rep., World Bank.

Tables

Table 1: Balance across information treatment (owners)

Variable	Control	Treatment	Difference
Install date (days since July 1, 2016)	183.8	183.7	0.16
			(9.04)
Owner age	36.6	37.2	-0.69
			(0.90)
Owner gender	0.18	0.17	0.0031
			(0.044)
Owner highest level of education	2.97	3	-0.026
			(0.10)
Owner is employed in salaried job	0.23	0.24	-0.011
			(0.049)
Years the owner is in matatu industry	8.07	8.09	-0.022
			(0.76)
Years owner has owned matatus	4.92	4.68	0.24
			(0.49)
Number of drivers hired for this matatu	1.28	1.32	-0.049
			(0.096)
Number of other drivers hired in the past	1.81	2.01	-0.20
			(0.22)
Amount given in trust game	123.7	131.6	-7.87
			(11.5)
Owner Raven's score	4.57	4.59	-0.018
			(0.18)
Driver rating: owner's fairness	8.13	8.33	-0.20
			(0.17)

The data are limited to the 300 owners. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels.

Table 2: Balance across cash treatment (drivers)

Variable	Control	Treatment	Difference
Driver age	34.9	36.9	-2.00
			(0.82)**
Driver highest level of education	2.47	2.47	0.0039
			(0.079)
Driver experience	7.14	8.30	-1.16
			(0.66)*
Driver industry tenure	9.92	11.7	-1.82
			(0.73)**
Weeks unemployed before current job	2.93	2.24	0.69
			(0.67)
Number of vehicles driven for before current	5.93	4.97	0.97
			(0.51)*
Number of conductors	1.23	1.13	0.096
			(0.052)*
Number of past accidents	0.90	0.87	0.034
			(0.13)
Number of months the driver has been employed	16.1	14.3	1.87
			(2.30)
Owner rating: driver's honesty	7.68	7.62	0.067
			(0.17)
Owner rating: how hard driver works	8.26	8.13	0.13
			(0.17)
Owner rating: driver's safety	8.33	8.23	0.10
			(0.18)
Owner rating: driver's performance overall	8.02	8.01	0.013
			(0.16)
Driver days working for owner	453.4	517.0	-63.7
			(65.7)
Driver Raven's score	4.23	4.23	0.0021
			(0.17)
Revenue at baseline	7744.8	7746.2	-1.45
			(207.3)
Baseline target	3203.3	3218.1	-14.8
			(81.9)

The data are limited to the 300 owners. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels.

Table 3: Knowledge gathered through the device

	(1) Know Km	(2) Know Off-route	(3) Know Revenue
Info Treatment	0.268***	0.451***	0.039
	(0.068)	(0.065)	(0.072)
Control Mean of DV	0.47	0.40	0.61
Controls	X	X	X
Matatu N	187	187	187

Each of the variables is a binary indicator for whether the owner knew the number of kilometers, instances of off-route driving and revenue generated by the vehicle. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels. Note these questions were added to the end-line survey after the first wave of endlines had already been completed, which is why we only have 187 observations (balanced across treatment and control)

Table 4: Monitoring through the device

	(1) Difficulty Monitor	(2) Monitoring Change	(3) Check (Phone)	(4) Check (Stage)	(5) Check (Third Party)
Info Treatment	-1.845***	-0.721***	0.966	0.184	-0.116
	(0.156)	(0.053)	(0.895)	(0.383)	(0.257)
Control Mean of DV	4.02	-0.01	7.01	1.95	0.95
Controls	X	X	X	X	X
Matatu N	190	190	190	190	190

These variables capture monitoring behaviors by the owner. Difficulty monitoring is an indicator from 1 to 5 for the level of difficulty associated with monitoring (5 = very hard). Change in monitoring captures whether owners are spending less time monitoring ($= -1$), more time monitoring ($= 1$), or see no change. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels. Note these questions were added to the end-line survey after the first wave of endlines had already been completed, which is why we only have 190 observations (balanced across treatment and control)

Table 5: Shading Behavior

	(1) 500	(2) 600	(3) 700	(4) 800	(5) 900	(6) 1000	(7) 1100
Treatment	-72.437^{**}	-88.290^{**}	-90.520^{**}	-109.736^{**}	-105.451^{**}	-96.244^{**}	-96.353^{**}
	(33.094)	(37.538)	(45.409)	(48.017)	(48.498)	(45.776)	(47.415)
Observations	3,378	3,822	4,503	5,339	5,866	6,820	7,101

Table 6: Perceptions of trust

	(1) Trust Amount	(2) Better Driving	(3) More Honest	(4) Performance Rating
Info Treatment	33.796**	0.626***	0.708***	0.112
	(15.123)	(0.057)	(0.052)	(0.174)
Control Mean of DV	151.61	0.04	0.04	7.21
Controls	X	X	X	X
Matatu N	244	190	190	246

Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels.

Table 7: Business decisions

	(1) Number Vehicles	(2) New Interior
Info Treatment	0.145*	0.074
	(0.078)	(0.057)
Control Mean of DV	1.22	0.21
Controls	X	X
Matatu N	246	240

Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels.

Table 8: Effect of cash (immediate)

	(1) Average speed	(2) Maximum speed	(3) Speeding	(4) Sharp braking	(5) Overacceleration	(6) Sharp turning
Cash Treatment	-0.099	-0.214	-0.250**	-0.131*	-0.009	0.036
	(0.247)	(0.874)	(0.113)	(0.074)	(0.015)	(0.034)
Mileage in km	0.007	0.022	0.001	0.001	0.000	0.000
	(0.005)	(0.014)	(0.001)	(0.001)	(0.000)	(0.000)
Control Mean of DV	15.89	52.64	0.70	0.77	0.08	0.39
Controls	X	X	X	X	X	X
Matatu FE	X	X	X	X	X	X
Day FE	X	X	X	X	X	X
Route FE	X	X	X	X	X	X
Matatu N	213	213	213	213	213	213
Matatu-Day N	39,072	39,072	39,072	39,072	39,072	39,072
R-squared	0.53	0.38	0.42	0.45	0.22	0.38

The data are limited to the 300 owners. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels.

Table 9: Effect of no cash (ongoing)

	(1) Average speed	(2) Maximum speed	(3) Speeding	(4) Sharp braking	(5) Overacceleration	(6) Sharp turning
Cash Treatment	-0.120	-0.409	-0.229**	-0.140**	-0.014	0.000
	(0.220)	(0.756)	(0.100)	(0.058)	(0.013)	(0.030)
One Month Post Treat	-0.039	0.072	-0.059	-0.115	-0.003	-0.015
	(0.260)	(0.971)	(0.138)	(0.089)	(0.012)	(0.030)
Mileage in km	0.008	0.024	0.002	0.001	0.000	0.001
	(0.005)	(0.015)	(0.001)	(0.001)	(0.000)	(0.000)
Control Mean of DV	15.89	52.64	0.70	0.77	0.08	0.39
Controls	X	X	X	X	X	X
Matatu FE	X	X	X	X	X	X
Day FE	X	X	X	X	X	X
Route FE	X	X	X	X	X	X
Matatu N	213	213	213	213	213	213
Matatu-Day N	42,405	42,405	42,405	42,405	42,405	42,405
R-squared	0.53	0.38	0.43	0.45	0.22	0.39

The data are limited to the 300 owners. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels.

Figures

Figure 1: Location

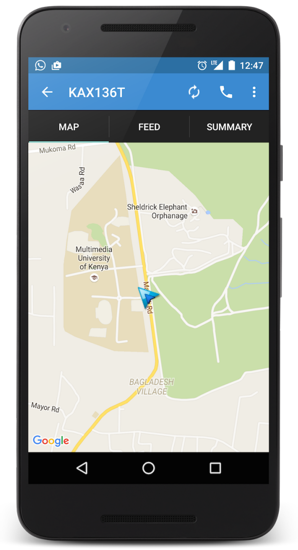

(a) Map Viewer

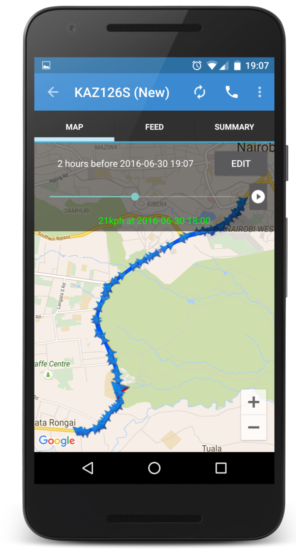

(b) Historical Map Viewer

Figure 2: Productivity and Safety

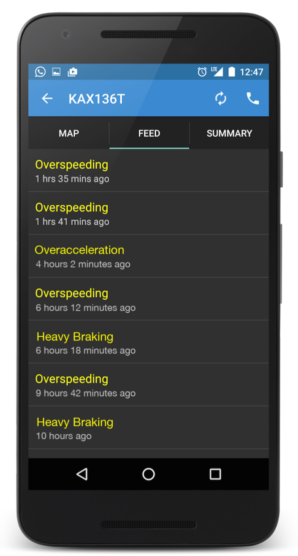

(a) Safety Feed

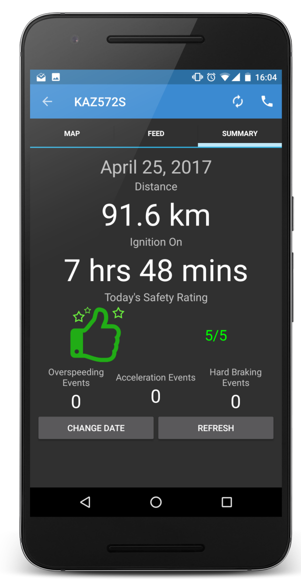

(b) Productivity Summary Viewer

Figure 3: Report Complete

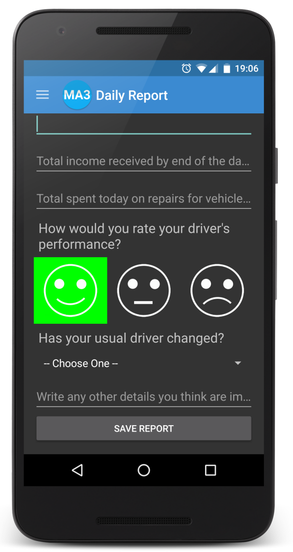

(a) Report

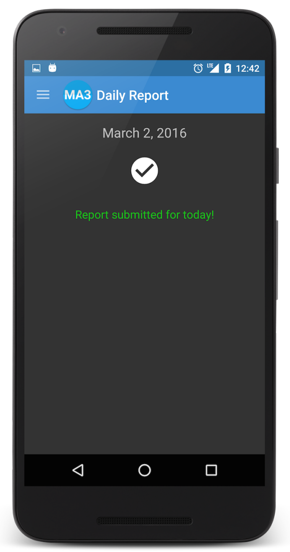

(b) Productivity Summary Viewer

Figure 4: Report Complete

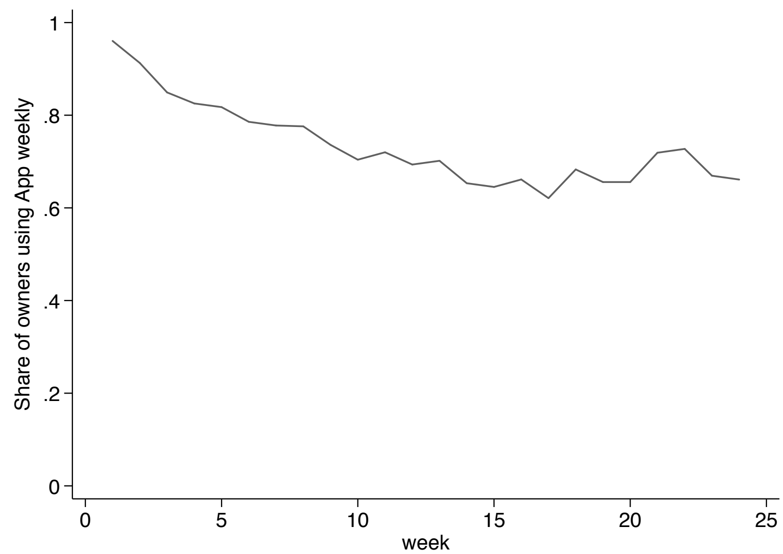

(a) Report

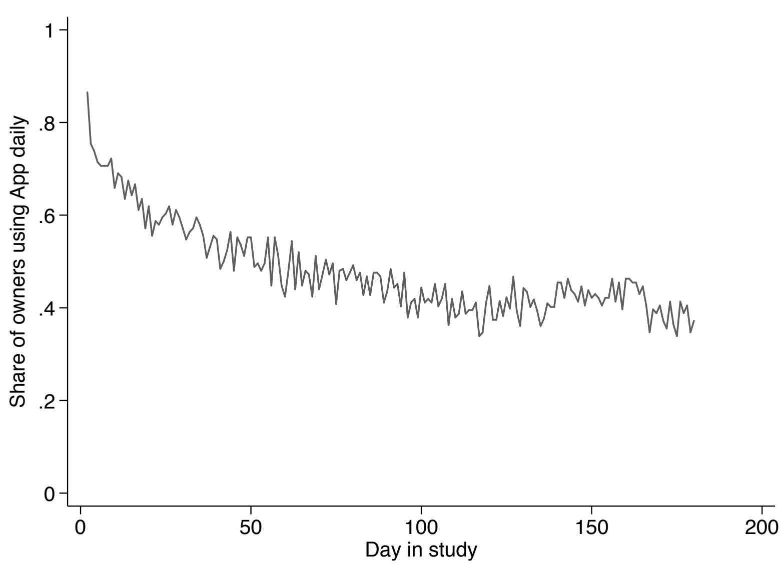

(b) Productivity Summary Viewer

Figure 5: Reprimands

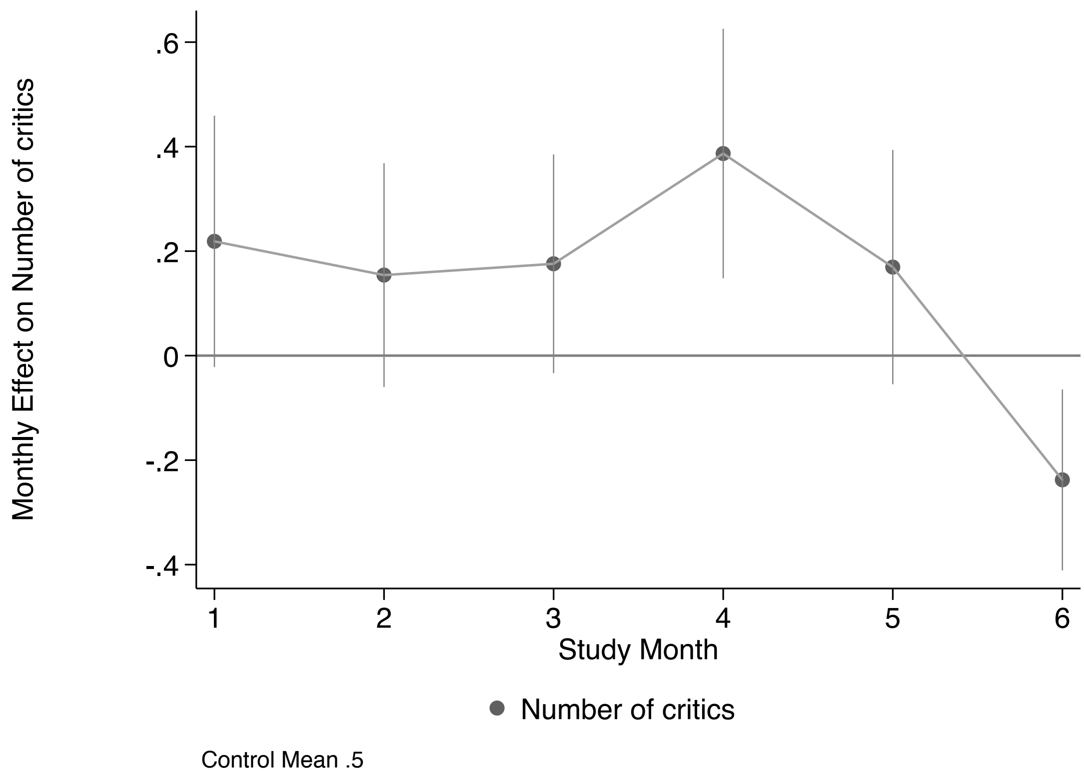

Notes:

Figure 6: Firing

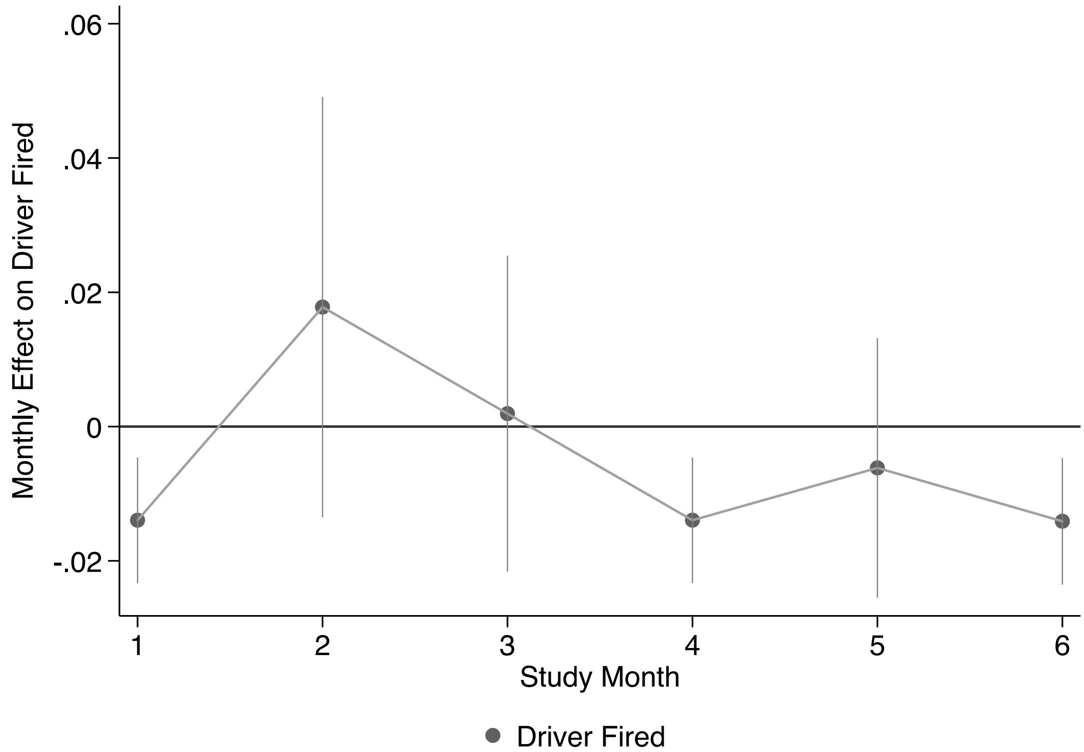

Notes:

Figure 7: Prediction 1 → Target

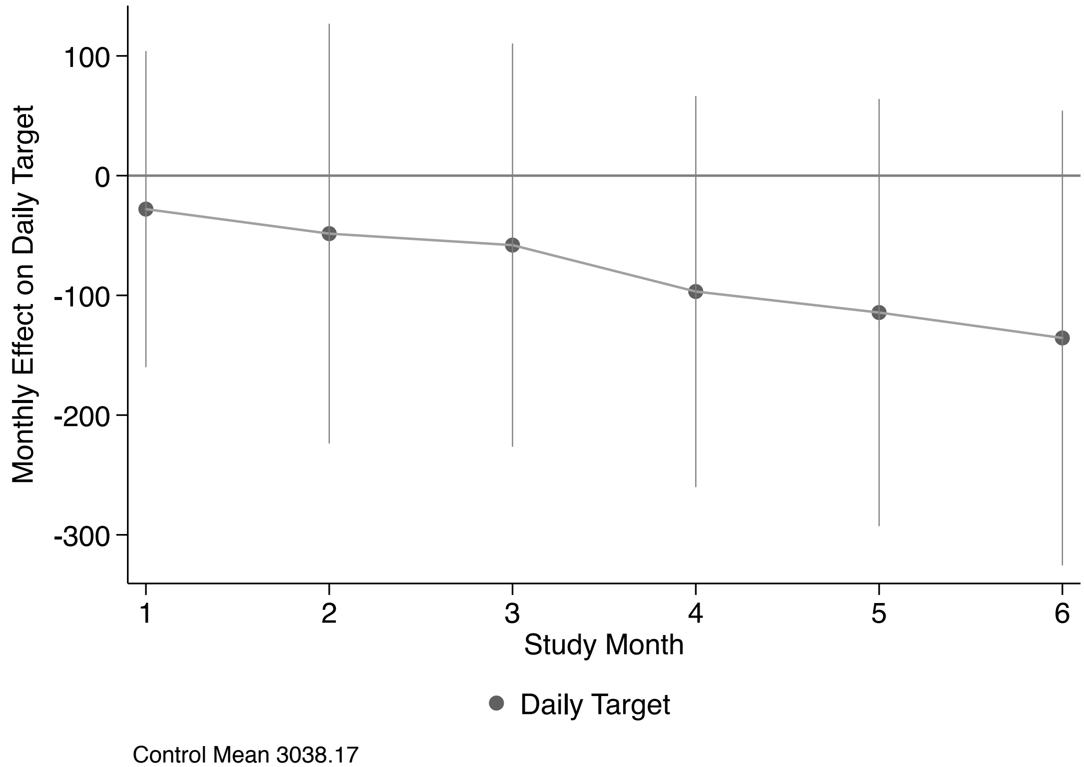

Notes:

Figure 8: Prediction 2 → Effort

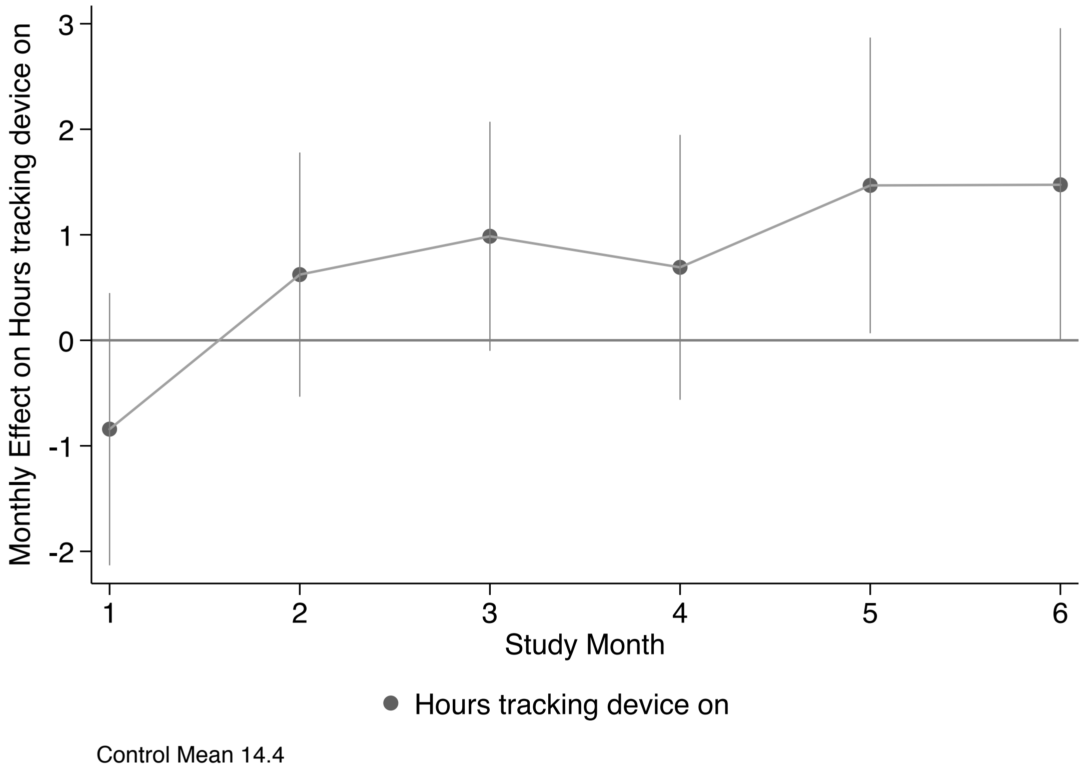

Notes:

Figure 9: Prediction 2 → Effort

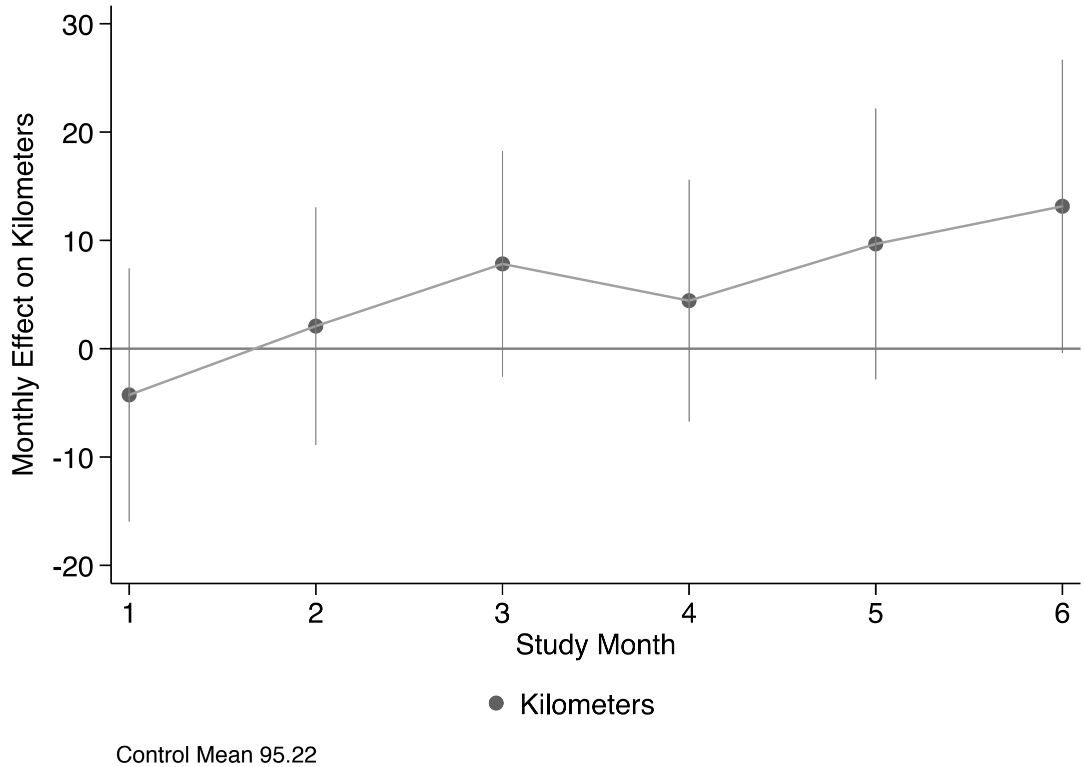

Notes:

Figure 10: Prediction 2 → Effort - Extensive Margin

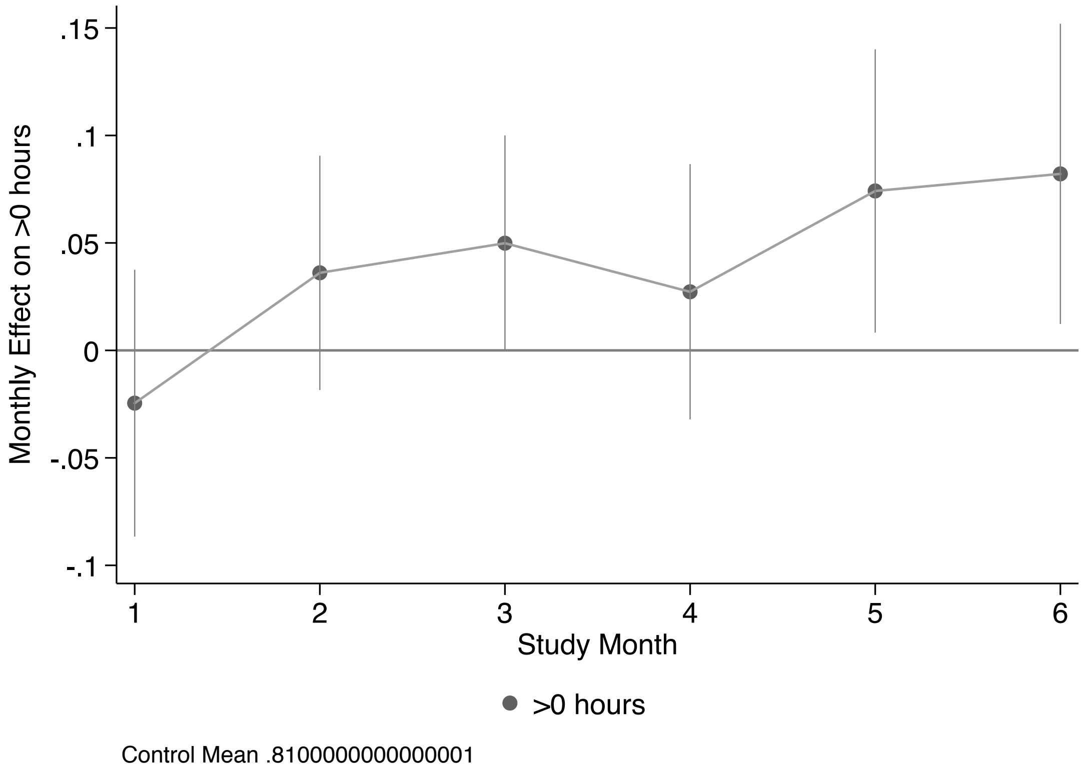

Notes:

Figure 11: Prediction 2 → Effort - Extensive Margin

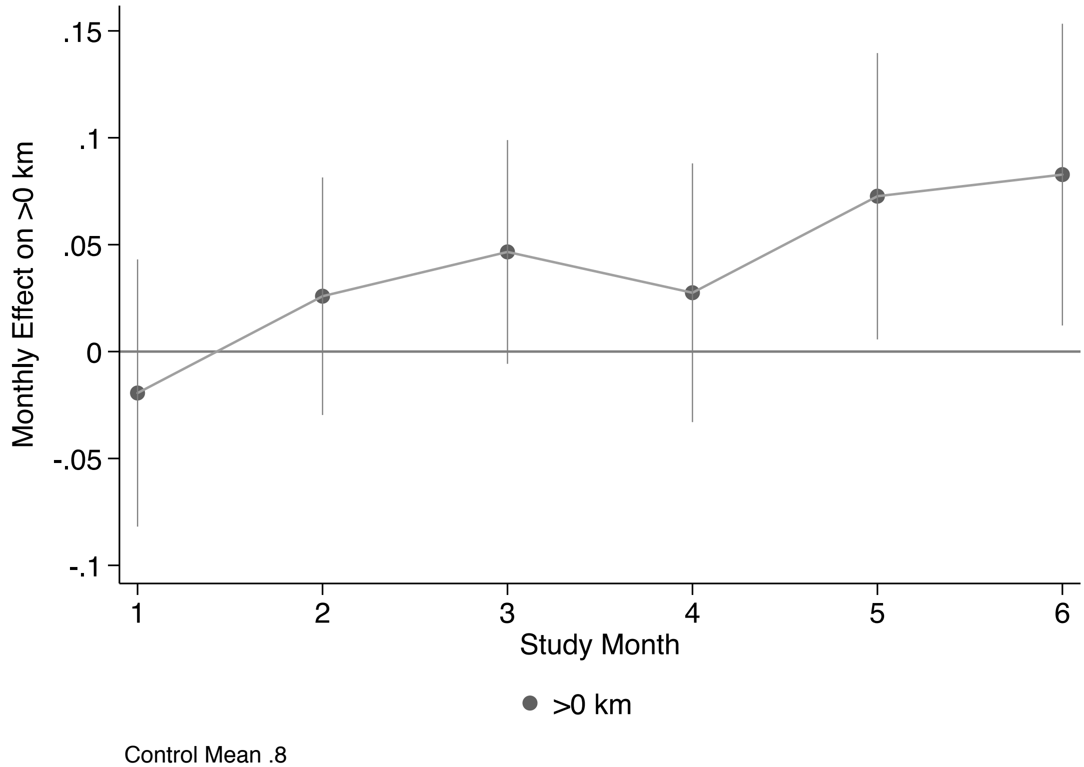

Notes:

Figure 12: Prediction 5 → Fewer Damages

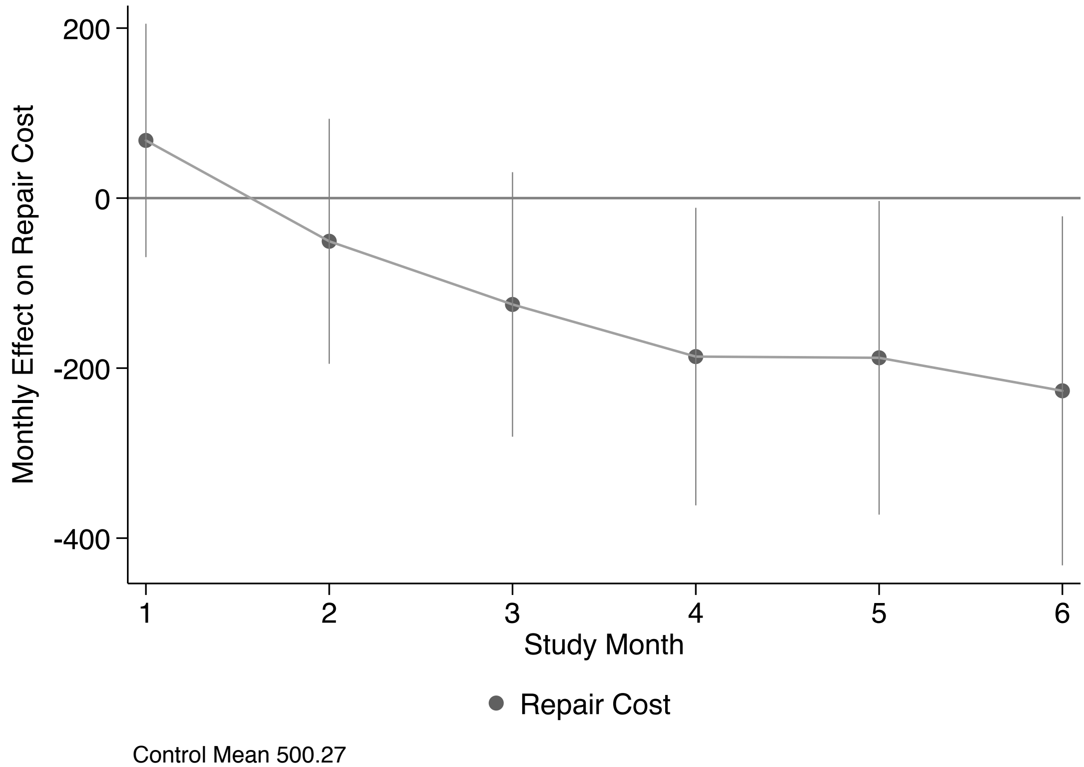

Notes:

Figure 13: Prediction 5 → Fewer Damages

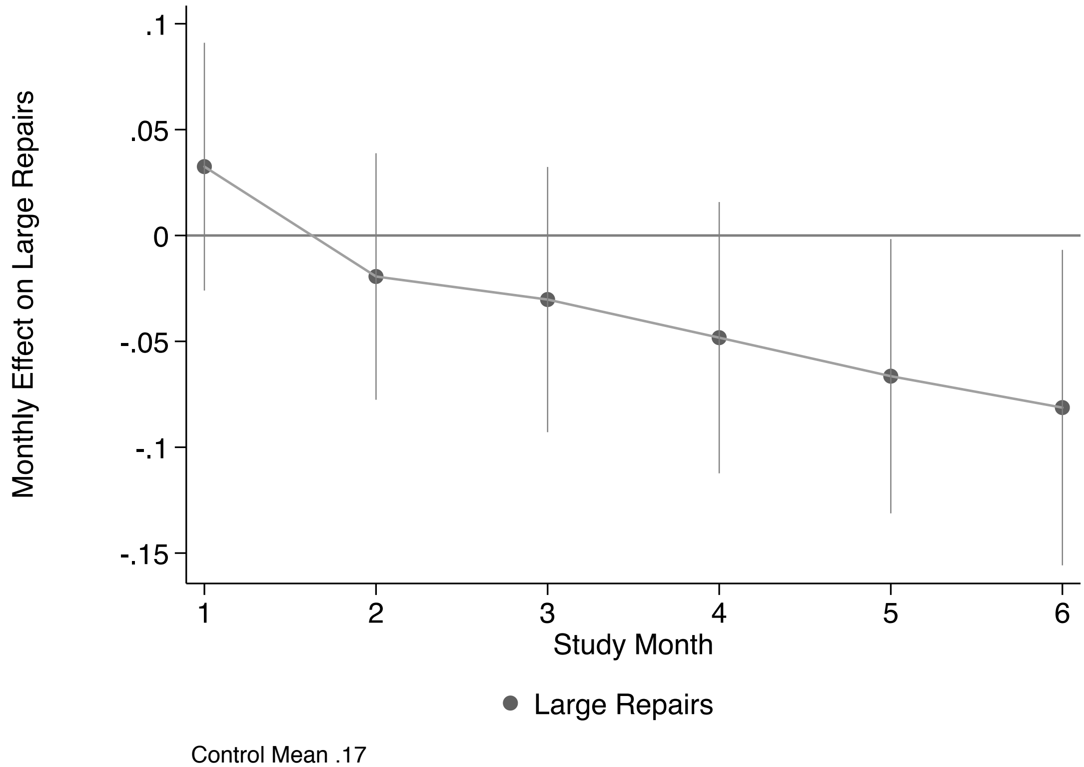

Notes:

Figure 14: Prediction 5 → Fewer Damages

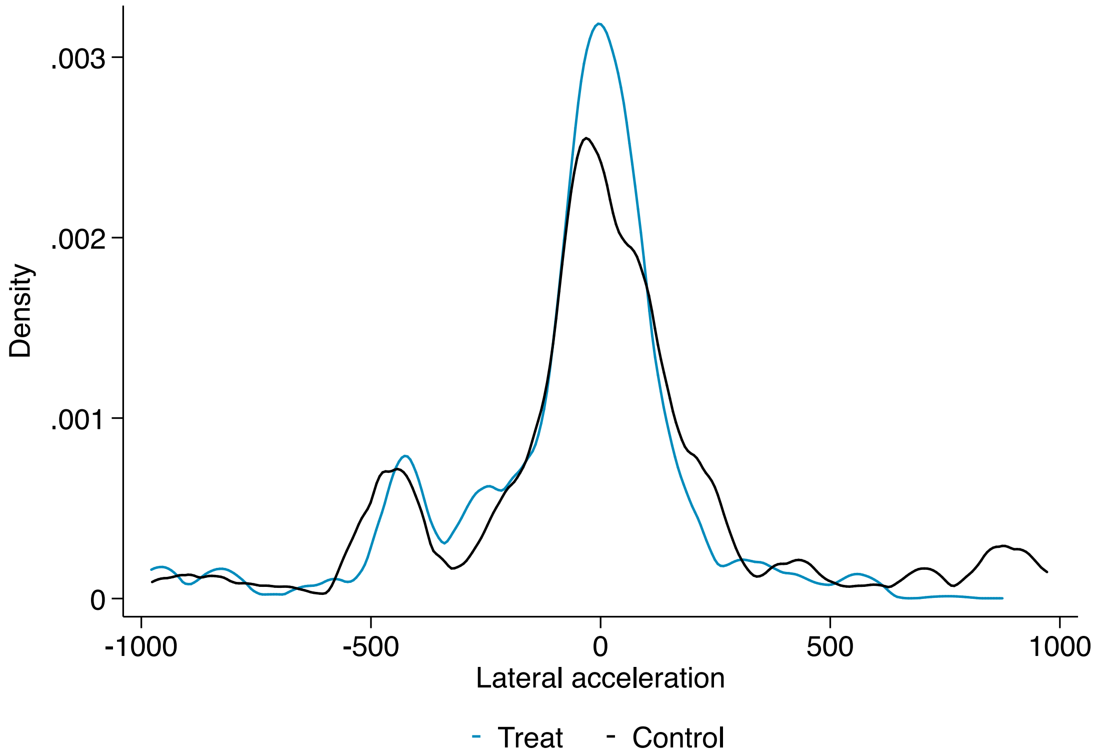

Notes:

Figure 15: Prediction 5 → Fewer Damages

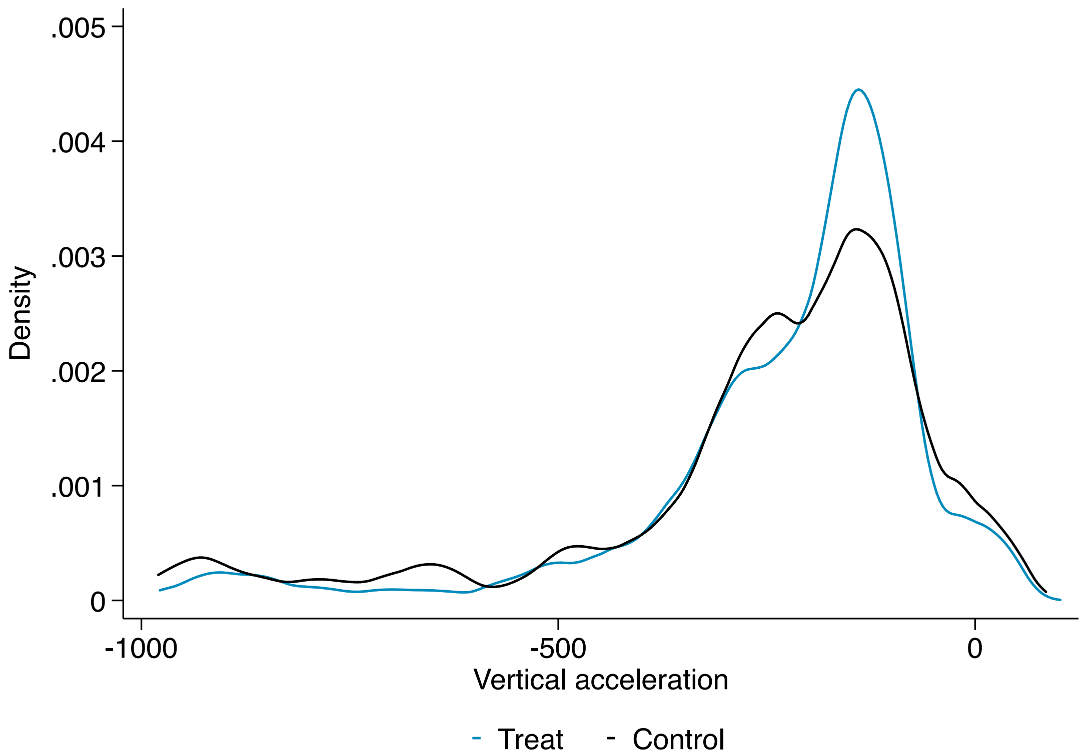

Notes: It was necessary to downsample the 3-axis acceleration data reported by the tracking devices in order to keep data storage manageable. To do this, several summary statistics were calculated for each 30-second window of acceleration data received. For the vertical component of acceleration, which included both positive and negative readings, the system stored the vector with the maximum magnitude for the window, regardless of direction. To analyze this variable, it was therefore necessary to take the absolute value, resulting in a measure equivalent to the maximum vertical component magnitude within each window. The distribution is centered at -200 rather than 0 (gravity) because of some combination of a non exact calibration and the asymmetry of suspension resulting in asymmetrical acceleration

Figure 16: Prediction 6 $\rightarrow$ Constant Shading

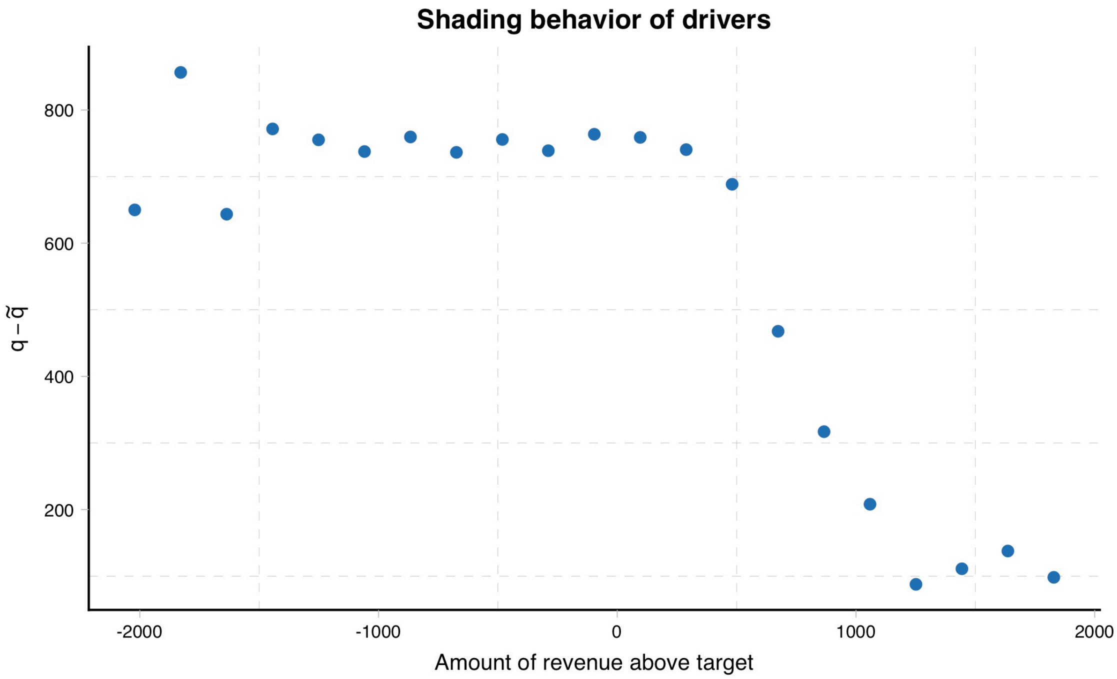

Notes:

Figure 17: Prediction 6 → Less Shading

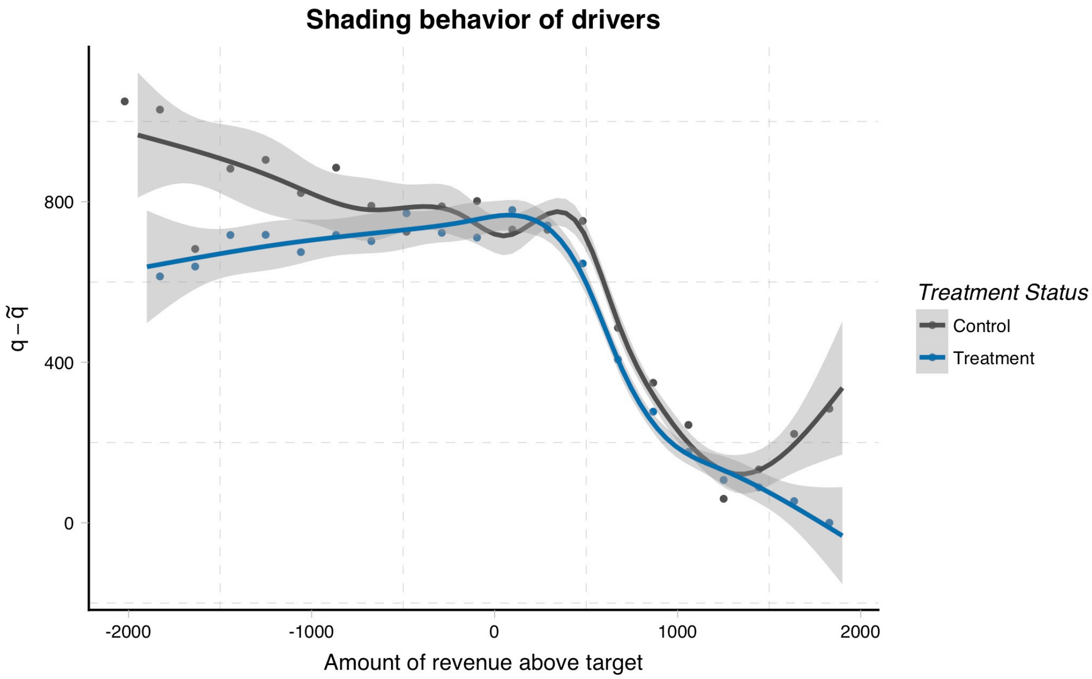

Notes:

Figure 18: Prediction 6 → Less Shading

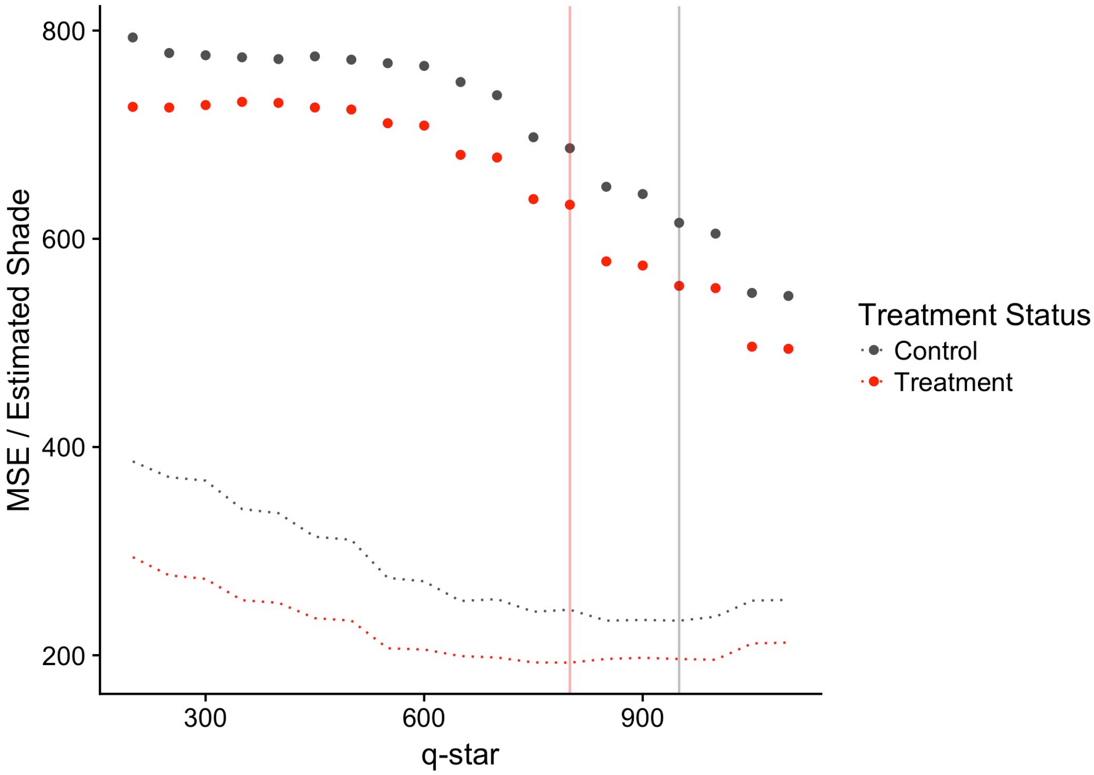

Notes:

Figure 19: Prediction 3 → Achieving Target

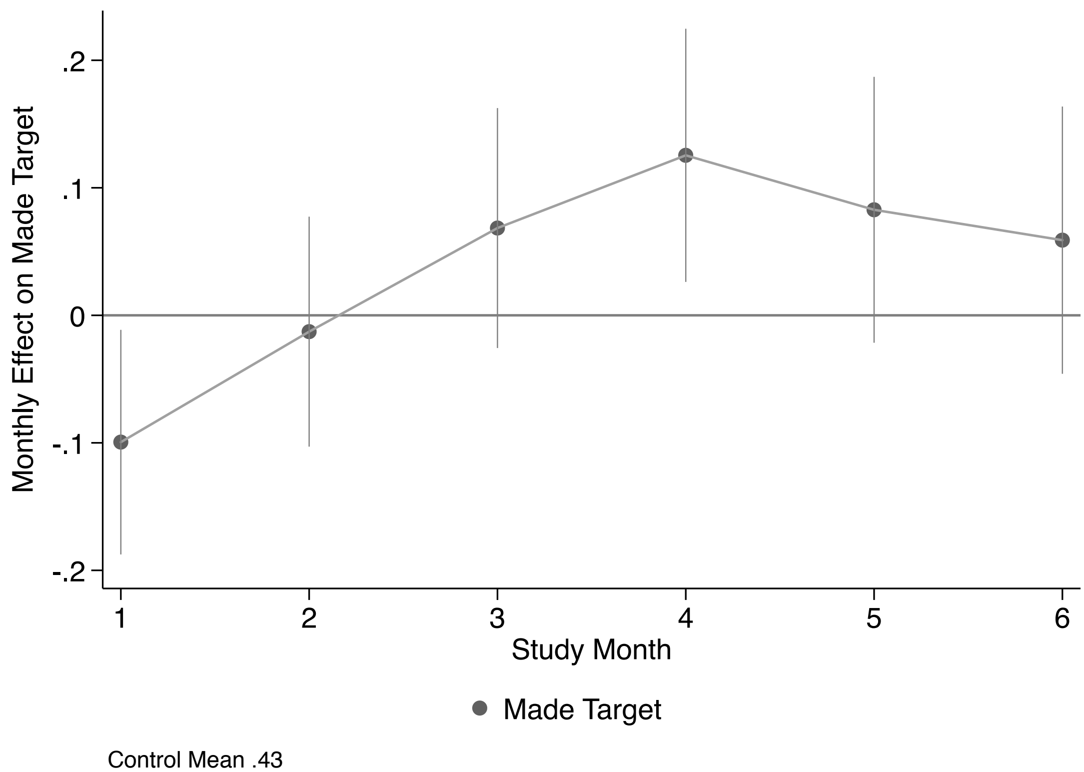

Notes:

Figure 20: Prediction 3 $\rightarrow$ Achieving Target

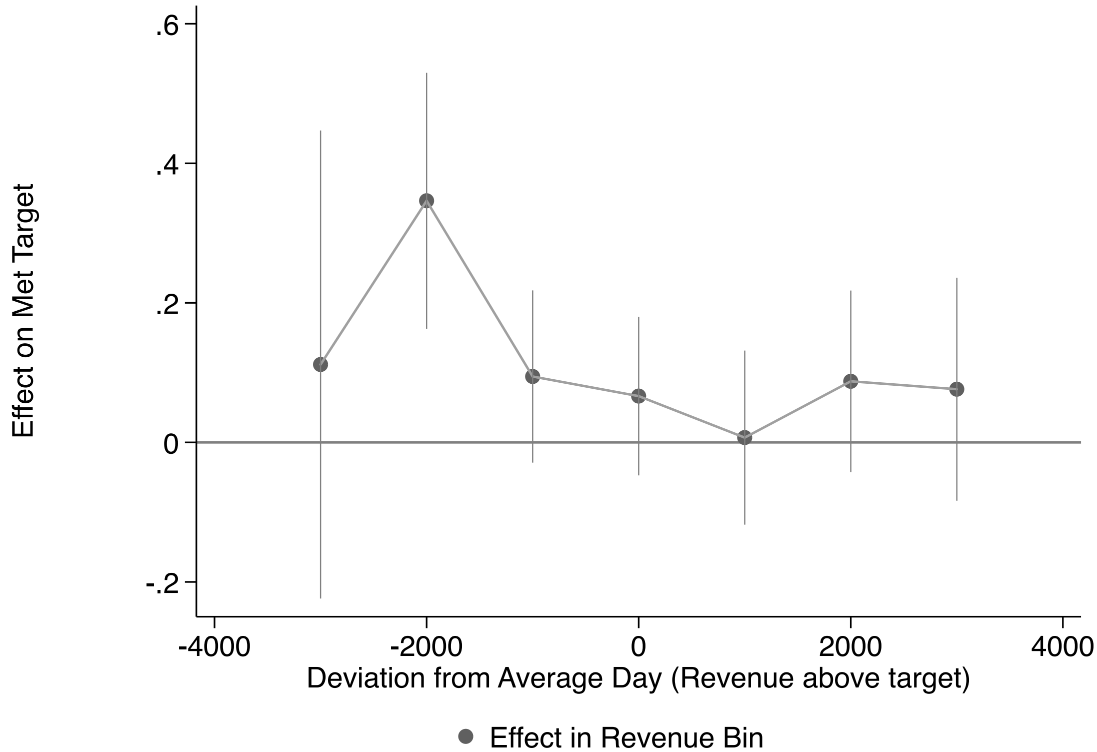

Notes:

Figure 21: Prediction 4 → More revenue

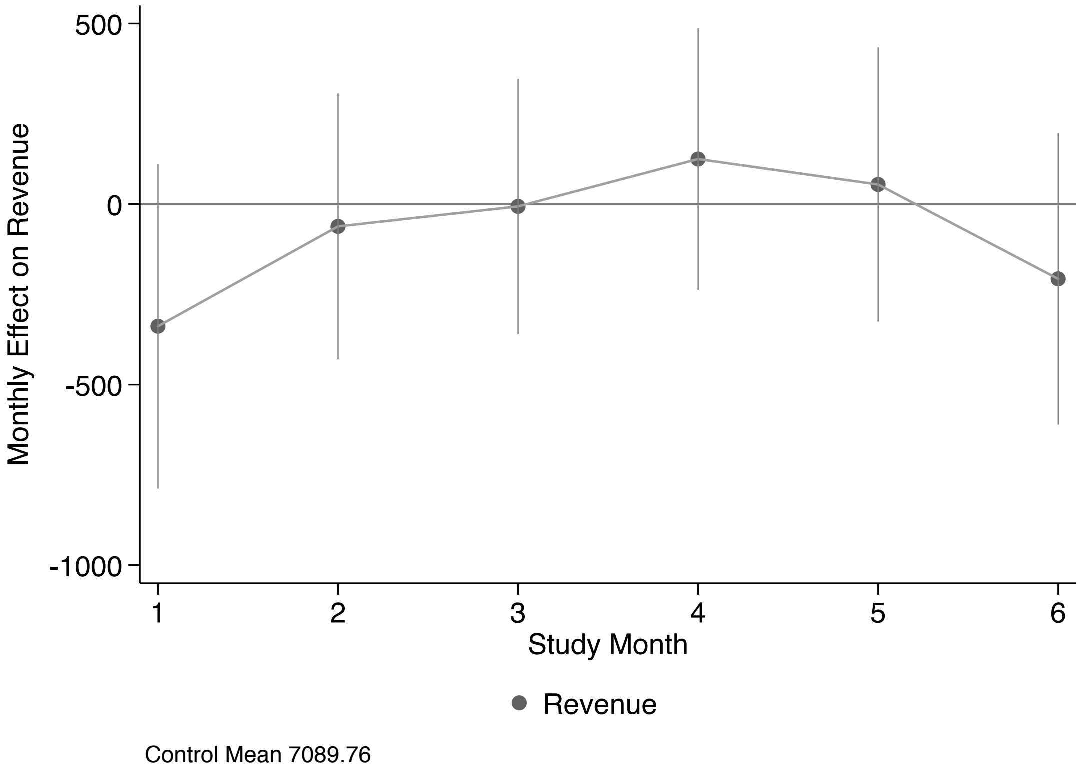

Notes:

Figure 22: Prediction 7 → Higher Profits

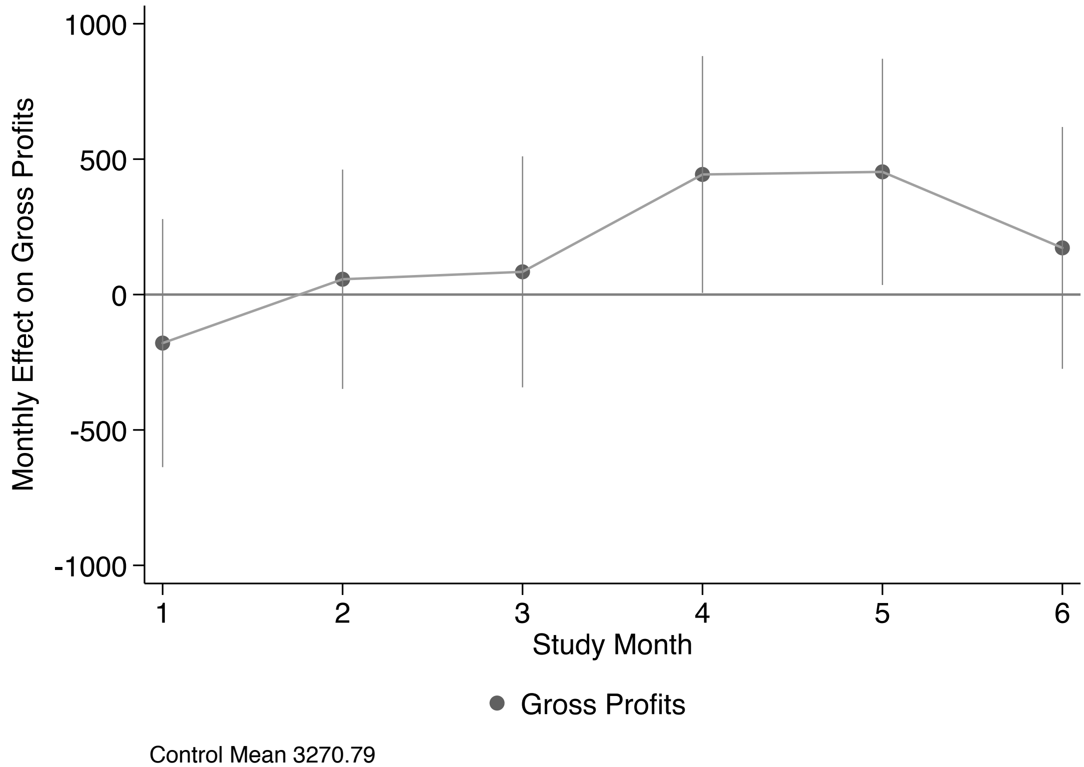

Notes:

Figure 23: Salary

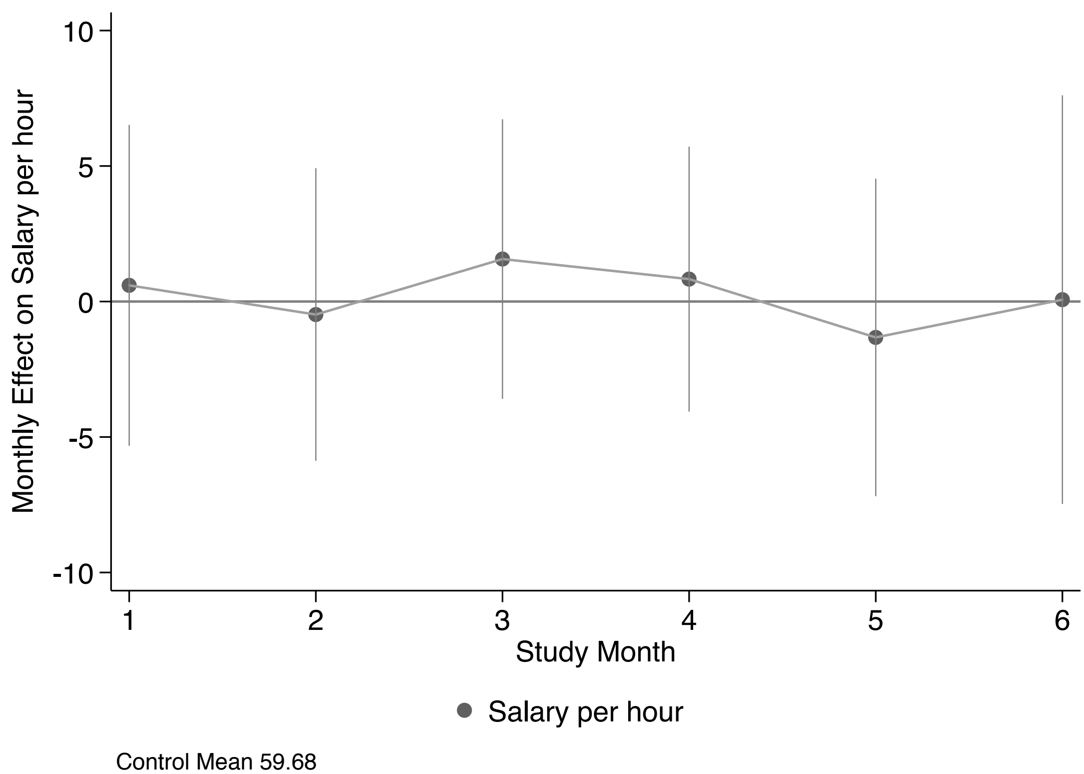

Notes:

Figure 24: Safety: Maximum Speeds

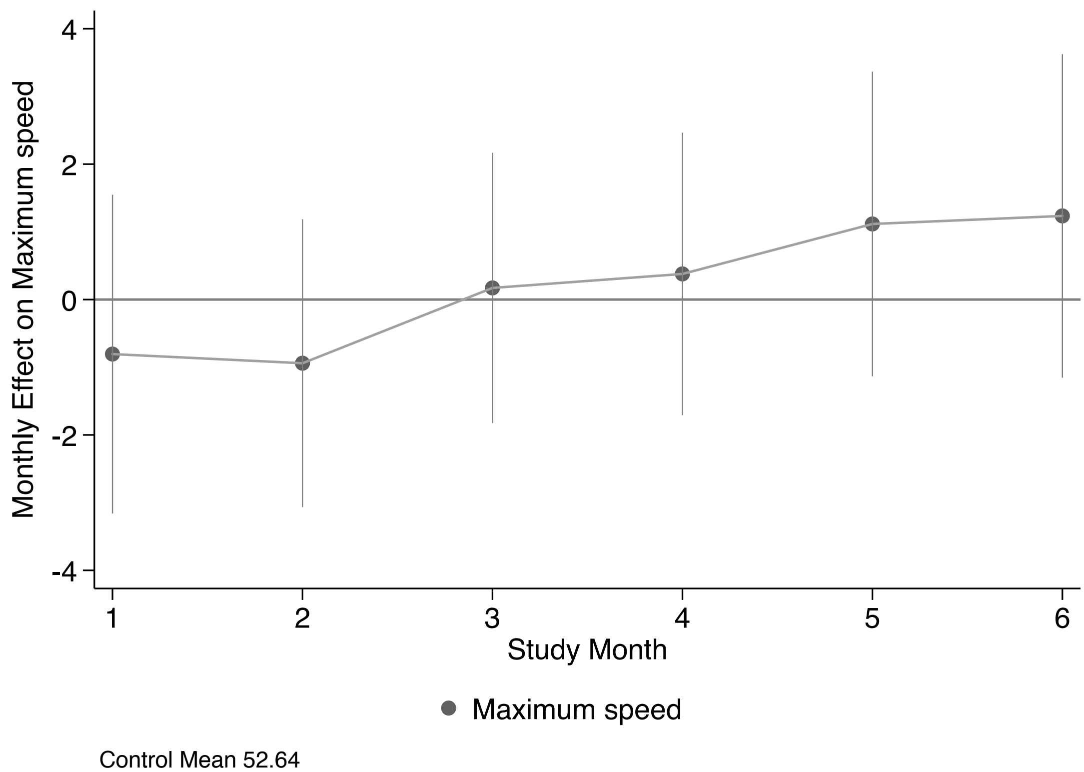

Notes:

Figure 25: Safety: Speeding above 80km

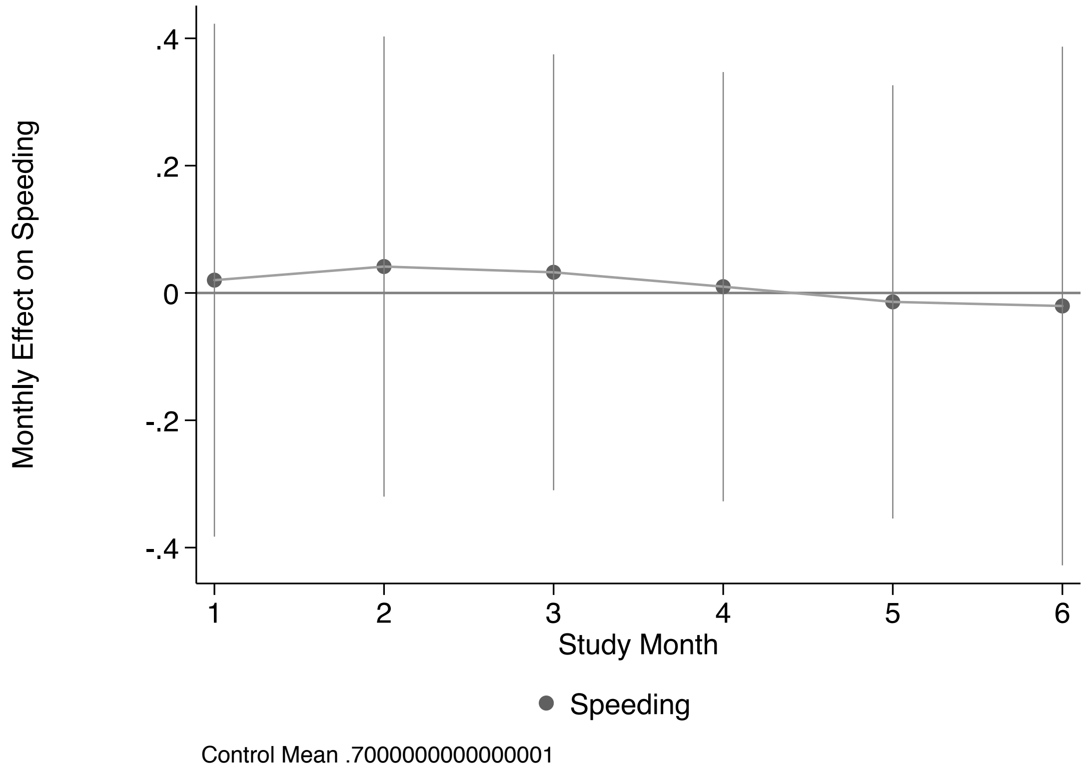

Notes:

Figure 26: Safety: Overacceleration

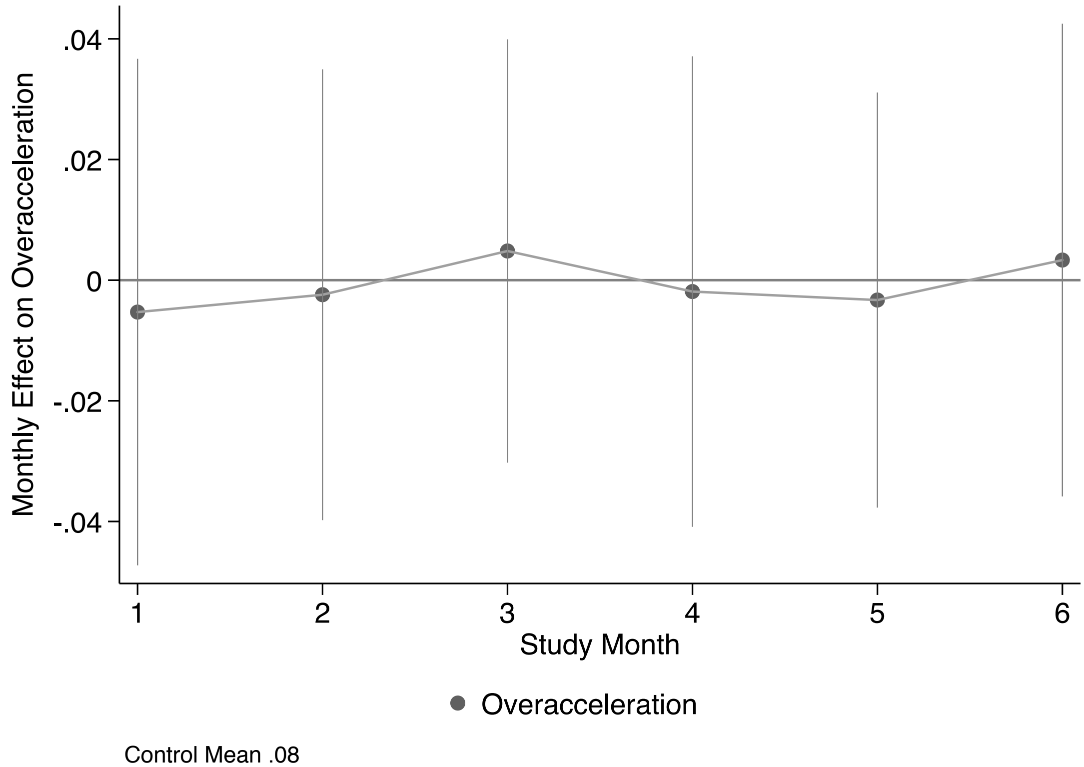

Notes:

Figure 27: Safety: Braking

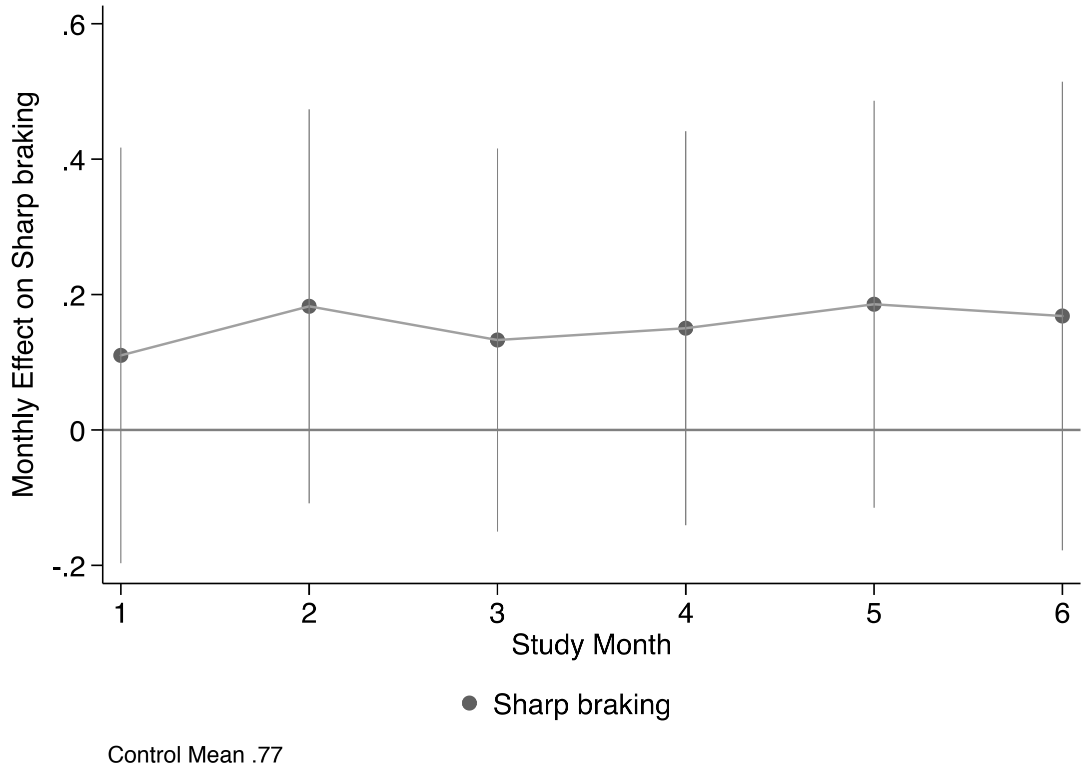

Notes:

Figure 28: Safety: Turning

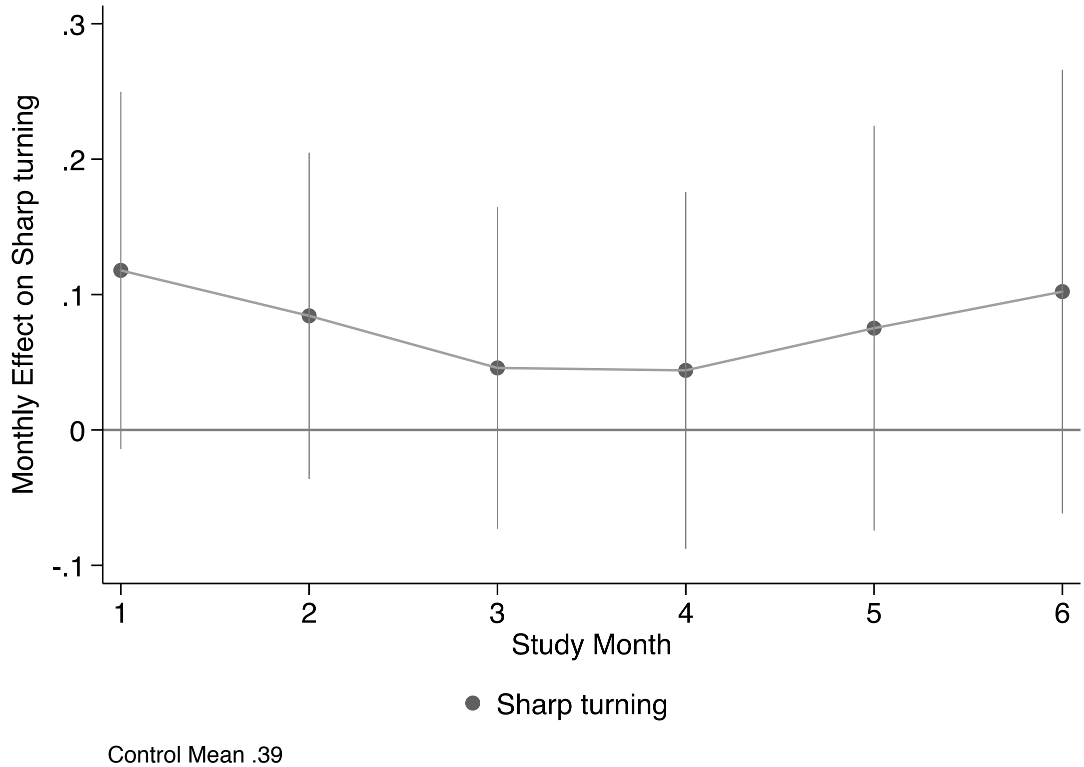

Notes:

Figure 29: Safety: Accident

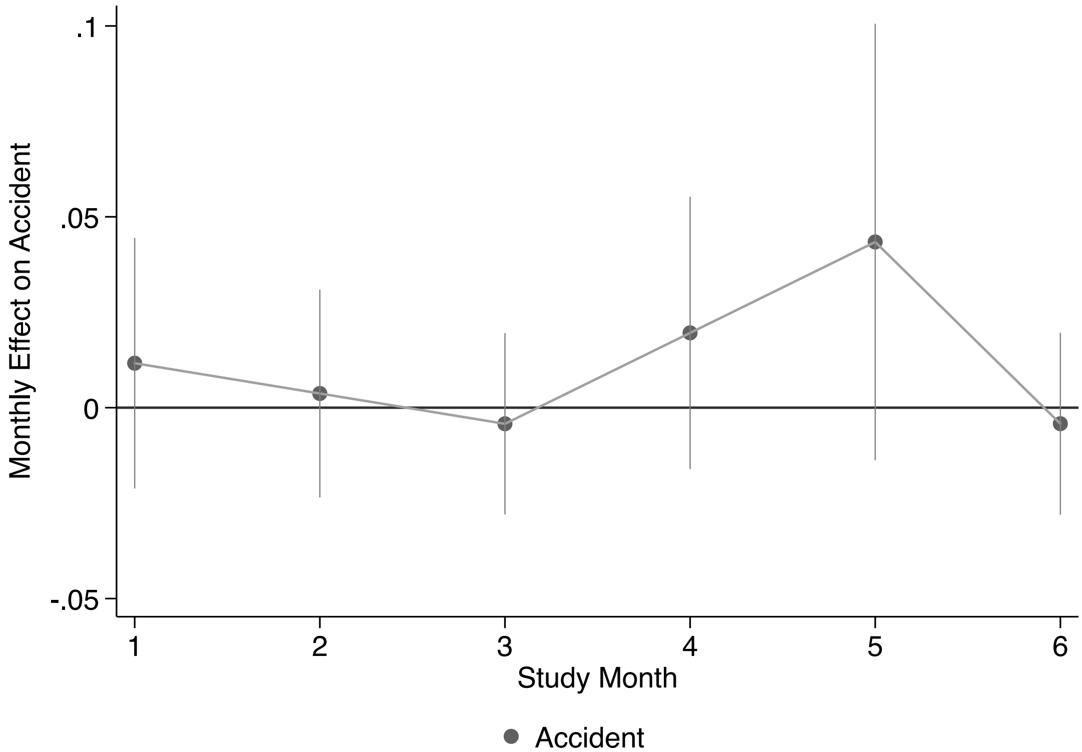

Notes: