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

Erin M. Kelley*

Gregory Lane and David Schönholzer

November 30, 2018

[JOB MARKET PAPER]
(Click here for the latest version)

Abstract

Agency theory suggests that moral hazard in employer–employee contracting constrains firm profits. We use a randomized controlled trial to empirically evaluate how information and communication technologies (ICT) can mitigate moral hazard and enable firms to design more efficient contracts which increase profits and engender business growth. Specifically, we study a fleet of 255 minibuses (matatus) in Nairobi, Kenya, where we introduce monitoring devices that track real-time vehicle location, daily productivity, and safety statistics. We randomize whether minibus owners have access to these monitoring data using a novel mobile app. This information allows owners in the treatment group to observe a more precise signal of driver effort, the amount of revenue drivers collected in fares, and the extent to which the driver engages in reckless driving. We find that treated vehicle owners modify the terms of the contract by decreasing the rental price they demand. Drivers respond by working more hours, decreasing behavior that damages the vehicle, and under-reporting revenue by less. These changes improve firm profits and reduce management costs, thereby helping treated firms grow. The device also improves owners' trust in their drivers, which drivers say makes their job easier. Finally, we investigate whether these gains to the company come at the expense of passenger safety, in an environment where accidents are common. While we do not find any evidence that conditions deteriorate, offering detailed information on driving behavior also does not *improve* safety. Only by incentivizing drivers through an additional cash treatment do we detect safety improvements.

1   Introduction

Firms design contracts to ensure their employees exert the profit-maximizing level of effort. In the presence of moral hazard, however, firms cannot condition the terms of the contract on important dimensions of employee behavior, including effort and output. Firms respond by relying on "second-best" self-enforcing incentive contracts or coercive measures to align agents' interests with their own (Hölmstrom, 1979; Grossman and Hart, 1983; Hart and Holmstrom, 1987; Shapiro and Stiglitz, 1984). In theory, firms can overcome these frictions by investing in monitoring technologies that reduce information asymmetries and reveal the performance of their workers more accurately (Harris and Raviv, 1979; Hölmstrom, 1979; Hubbard, 2003). In practice, however, the impact of these technologies on contracts is unclear: the presence of institutional and managerial frictions may limit employers' ability to leverage the additional information needed to change the contract and employees' behavior.

This paper studies the impact of moral hazard on labor contracting, and productivity, and the extent to which improved monitoring eases these frictions. We also investigate whether monitoring technologies subsequently improve firm profits, and worker well-being. We implement a randomized control trial where we introduce a novel monitoring device to a subset of firms operating in Kenya's transit industry. The industry is dominated by thousands of small-scale entrepreneurs who own a few minibuses ("matatus") that run on designated routes. These matatus are the only reliable form of transportation and serve 70% of Nairobi's four million commuters daily.[1] We recruited 255 owners operating along 9 major commuter routes to participate in the study, and we randomly selected 125 to be part of the treatment group. The monitoring device was fitted to all the matatus in our sample, but only transmitted data to minibus owners in the treatment group.[2] 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 does not know the number of passengers that boarded the vehicle, they can use this information to monitor drivers' operations throughout the day and to gain a more precise estimate of total daily revenue.

The contracting environment we study here is not unique to Kenya or transportation, as the dynamics that characterize this space are prevalent in many other settings, including agriculture and the service industry. First, employers (owners) cannot observe the amount of revenue their employees (drivers) collect, nor the amount of effort drivers invests. Second, 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, meaning that in practice they have limited or no liability to the owners. Drivers are known to run away from accidents so they can avoid being held accountable by owners or the police. In light of

---

[1] Similar transit systems are present in Mexico (peseros), the Philippines (jeepney's), Indonesia (tuk-tuks), India (rickshaws), and Tanzania (dala-dala's), among others.

[2] The drivers were present during the installations of the devices, but they were not informed about whether the owner was in the treatment group or not.

1

these constraints, firms have overwhelmingly opted for fixed-rent contracts (locally referred to as a "target" contract) with limited liability. The owner specifies an amount of revenue that the driver must deliver by the end of the day, net of fuel expenses. According to the contract, the driver should deliver the fixed rent ("target") amount if the revenue they collect exceeds the target (keeping any revenue they earn above it). If the earnings are below the target, the contract stipulates that the driver must hand over all of the revenue they earned.

In order to understand the impact of the device, we adapt a standard principal-agent framework to reflect the actual employer-employee relationships within this network. In the absence of a monitoring technology, the model predicts that drivers will engage in a number of behaviors that are sub-optimal from the firm's perspective. First, drivers under-report revenue so they can be sure to walk away with slightly more income than they would otherwise. Note: rampant cheating is kept in check by owners who threaten to punish and ultimately fire drivers who are caught under-reporting. Second, drivers under-supply effort on days when they are unlikely to make the target price. On these particular days, drivers know they will not be the residual claimant and reap the benefits of higher effort. Third, drivers engage in more damaging driving than what owners would optimally choose. Damaging driving refers to the maneuvers drivers make that may damage the vehicle. These actions include driving on the shoulder of the road, or veering off the designated route onto roads that are more bumpy and damaging to the vehicle. The driver engages in such damaging behavior because he reaps the benefits in terms of higher revenue, without bearing any of the downside risk (the limited liability constraint binds).

We model the introduction of the new technology as increasing: 1) the precision of the owner's signal about total revenue and 2) the probability the owner detects damaging driving. This has implications for the owners' choice of contract and drivers' behavior, which we test in our data. According to the model, the monitoring technology reduces drivers' information rents and lowers their utility. Owners recognize this outcome and compensate by reducing the target. Empirically, we have some suggestive evidence that owners steadily reduce the target throughout the study period. By the last month of the study, the target set by treatment owners is approximately 4.1% lower than the target set by control owners.

The model then predicts that drivers' behavior is affected along three key dimensions. First, the drivers' incentive to lie about the total amount of revenue they collected is reduced, which means that we should see lower under-reporting. We confirm this prediction in our data: under-reported revenue falls by approximately 100 shillings (1 USD) per day (a 16% decrease). Second, the model predicts that drivers will increase their effort to compensate for the income they lose from lower under-reporting. In parallel, as the target falls, drivers also have an incentive to increase their effort because they can become the residual claimant more easily. We capture a precise measure of effort through the tracking device, which powers on and off with the matatu, and consequently find that the number of hours the vehicle is on the road increases by 1.4 hours per day (a 9.9% increase). Finally, we expect the driver to reduce instances of damaging driving because they are more likely to be caught by the owner. We proxy damaging driving by the amount of repair costs the owner

incurs, and we find these decrease by 200 shillings (2 USD) per day (46%) by the last month of the study. We have evidence to suggest this outcome comes from fewer instances of driving on alternate routes that are bumpy.

Next, we investigate the effect of these changes on firm profitability. We find that profits increase by 13%, which is primarily driven by lower repair costs. These 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 if it were available on the market. Owners also report that monitoring their drivers has become significantly easier, and they trust their drivers more, consistent with a reduction in management costs. Subsequently, we ask whether this improved profitability and better management fueled business growth. We find that treatment owners have 0.145 more vehicles (11% increase) on average than control owners by the end of the study. This suggests that inadequate monitoring may represent an important barrier to firm growth in low-income countries.

It is also important to investigate whether these benefits to the firm come at the expense of their workers. The introduction of ICT has generated some debate in low-income countries. On the one hand, there is a concern that these technologies concentrate all of the bargaining power in the hands of the employer. We do see some evidence of this as the amount of revenue drivers can under-report falls, and drivers work more (although their salary per hour stays the same). However, proponents of these new technologies suggest that they increase employers' trust in their employees, which makes for better managers (Pierce, Snow, and McAfee, 2015). We have some suggestive evidence of this in our data. In a qualitative survey we conducted six months after the experiment concluded, we find that 65% of drivers said the tracking device made their job easier (26% said nothing changed). This suggests that the effects of new technologies on worker well-being are nuanced, despite the net benefit they represent for firms.

These first results demonstrate how alleviating moral hazard affects operations *within* the firm. However, the presence of monitoring devices can also have effects *outside* of the firm, as profit-maximizing behavior by firms and their employees may impose negative externalities. In public transportation systems, monitoring technologies are often used to check and limit instances of unsafe driving. Kenya's matatu sector is notorious for its poor safety standards: drivers often over-accelerate, speed, stop suddenly, and turn sharply in order to collect more passengers.[3] Our monitoring technology records these four instances, in addition to maximum and average speeds, and conveys them to owners through a separate tab in the mobile app. It is important to note that these traditional measures of unsafe driving need not be perfectly correlated with the damaging driving behavior outlined above. For example, drivers often choose to bypass slow-moving traffic by taking rough, unpaved roads that harm the vehicle but pose no safety danger. Alternatively, driving quickly through crowded pedestrian areas may be unsafe, but it is unlikely to cause significant damage to the vehicle.

A priori, it was not clear how the owners would use the safety information we provided. On the

---

[3]Matatus account for 11% of registered vehicles but 70.2% of passenger casualties (Macharia et al., 2009). Buses in the US account for 1% of registered vehicles and 0.4% of casualties (BTS, 2016).

3

one hand, owners may internalize the dangers associated with unsafe driving because, unlike their drivers, they bear the cost if their vehicle gets into an accident or receives a fine.[4] In that case, they might use the information to induce drivers to limit behaviors that result in these accidents and fines. On the other hand, owners may perceive the risk of accidents to be low and disregard this information. Alternatively, while owners care about damages to the vehicle, these may only be weakly correlated with unsafe driving (as detailed above). If they use the device to incentivize effort only, the number of safety violations could worsen over the study period, producing a negative externality for Nairobi's commuters.

Despite all the safety information we provided, the frequency of unsafe driving events flagged by the device does not change significantly, and instances of speeding remain the same. It follows that the gains to firms do not come at the expense of commuters. However, these results also suggest that external intervention may be necessary to improve safety. We tested the efficacy of one such intervention by providing small cash incentives to drivers conditional on safe driving. This treatment was designed to mirror the actions that a regulatory body could potentially take in this setting.[5] Our objective was to determine the effectiveness of an intervention that encouraged the employees (drivers) rather than the employers (owners) to internalize the negative externalities produced by the business. We find that the cash incentives meaningfully reduce safety violations committed by drivers, confirming that third-party intervention can successfully address these firm externalities. However, these effects do 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 three different literatures. First, the paper speaks to the vast theoretical work on principal-agent relationships and contract formation, which predominantly focuses on deriving the optimal contract subject to various constraints (Hölmstrom, 1979; Grossman and Hart, 1983; Hart and Holmstrom, 1987; Shapiro and Stiglitz, 1984). Our paper, by contrast, empirically demonstrates how contracts change when these constraints are alleviated via monitoring technologies. Measuring the impact of monitoring is challenging because shirking behavior is hard to detect by design, a firm's decision to monitor is often not random, and data on firm operations are difficult to obtain. There are only two other papers to our knowledge that overcome these limitations. Baker and Hubbard (2003, 2004) investigate how the introduction of onboard diagnostic computers (OBCs) change ownership patterns in the U.S trucking industry. Baker and Hubbard (2004) demonstrate that shipping companies respond to the introduction of OBCs by hiring drivers to operate their vehicles (rather than working with drivers who already own their own trucks). Our paper differs from this existing work in a number of ways. We generate exogenous variation in the usage of monitoring technologies by randomizing which companies receive data from a tracking device. We also capture high frequency data on contracts and worker behavior. This allows us

---

[4]The owner may also have limited liability as they are protected by insurance. However insurance does not always cover full repairs, and premiums will rise if the owner continues filing insurance claims. The owner therefore always absorbs much more liability in the event of an accident, which means that owners should prefer less accidents and fines on average than the driver.

[5]South Africa's Ethekwini municipality is testing one such intervention in the coming months (Payet, 2018).

to monitor how different dimensions of the contract, and worker performance, change over time. We can then document the impact on firm profits and worker well-being (salary per hour, hours worked, sense of trust).

These papers are also concentrated in developed countries, and we have reason to believe that the impacts of monitoring could be different in low-income countries. Management quality is different, and employers may not use the information effectively (Bloom et al., 2013; Bloom, Sadun, and Van Reenen, 2017). Employers also face additional frictions that may limit their ability to use the information: law enforcement is weak and limited liability constraints bind. Contrary to Baker and Hubbard (2004), we do not find that ownership patterns change as a result of monitoring, precisely *because* of existing constraints (drivers cannot afford their own vehicles). Nevertheless, we do find that firms use the technology to change the terms of the contract, and to induce their employees to behave in a way that aligns with the firms' best interest. There is only one other paper to our knowledge that investigates the impact of monitoring technologies within firms 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 monitoring technologies crowd out high-performing workers' intrinsic motivation.[6] Our analysis builds on this work by investigating how monitoring technologies alleviate other key dimensions of moral hazard (including damaging driving, and lying about total revenue). We also focus on the impact of these devices on contracts.

Second, our findings add to the literature investigating the barriers to firm growth in low-income countries. Identifying the constraints to firm growth is a question of great policy relevance given the large contribution these firms make to emerging economies. Empirical research on small firm growth has identified three key challenges facing firms: credit constraints, labor-market frictions, and managerial deficits (Bloom et al., 2014). Our paper most closely resembles the work on managerial deficits, which refers to the difficulties firms face managing the day-to-day operations (including financial accounts and inventories), and incentivizing and monitoring workers. Most of the work in this field studies the impact of business training programs (Bloom et al., 2013; Bloom, Sadun, and Van Reenen, 2017; McKenzie and Woodruff, 2016; Berge, Bjorvatn, and Tungodden, 2014; de Mel, McKenzie, and Woodruff, 2014; Valvidia, 2012). These interventions provide information about how to manage aspects of the business that do not involve employees (maintaining business records, separating finances, inventory, controlling for quality, marketing). In contrast, our paper focuses on the role of moral hazard, and how providing information specifically about *employees' behavior* can change firm operations. We find monitoring technologies improve firm profits and reduce management costs, which helps the treated firms grow. As prices fall, these technologies are becoming increasingly prevalent, making their impacts important to understand.

Finally, our results on damaging driving and traditional safety metrics contribute to a growing

---

[6]Note there are additional studies that document the impacts of monitoring in low-income countries - but they do not focus on the employer-employee relationship within the firm. Duflo, Hanna, and Ryan (2012) find that teacher absenteeism in India decreases when their attendance is monitored; Björkman and Svensson (2009) demonstrate that community health workers exert more effort when their performance is scrutinized by the community; and Duflo et al. (2013) find that incentives for third-party auditors can improve their reporting.

empirical literature on policies that promote compliance with government regulation - in this case with safety regulation. In recent years, international institutions have provided funding, knowledge and technical assistance to build systems aimed at reducing the number of traffic injuries and deaths worldwide (World Bank, 2014).[7] These efforts are difficult to evaluate because the investments are multi-faceted and typically rolled out across an entire city. One exception is a program that was launched in Kenya, which placed stickers inside Nairobi's matatus to encourage passengers to complain to their drivers about unsafe driving (Habyarimana and Jack, 2015). They find that the intervention reduced accidents by 25-30%. Our intervention complements their approach by asking whether drivers, in addition to passengers, can be incentivized to improve safe driving. We find that owners with access to the monitoring device do not internalize the negative externalities produced by their drivers. We only document improvements in safety when we directly incentivize drivers. This result suggests that investments in technologies that monitor unsafe driving may be more effective when combined with external incentives.

The remainder of this paper is organized as follows: Section 2 discusses Kenya's transportation system, the prevalence of moral hazard, and the scope for monitoring. Section 3 details the field experiment, and Section 4 reviews the data. We present a simple theoretical framework in Section 5. Section 6 discusses each of our results. We then discuss the implications of the findings and conclude in the final section.

2   Context

2.1   Nairobi's Matatus

Nairobi's transportation system was developed after Kenya's 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 Kenyan Shillings (KES) of a matatu ride (where 100 KES = 1 USD). 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 limiters, 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 sector. This required that all minibus owners form or join transport Savings and Credit Cooperatives (SACCOs) or transport companies licensed to a particular route (McCormick et al., 2013). At present, industry newcomers must first register with a SACCO or transport company before they can put their vehicle on the road. Transport companies are rare 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

---

[7]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)

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 when drivers strike. 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 industry remains almost entirely locally owned: private entrepreneurs purchase 14 or 33 seat minibuses, and hire a driver to operate the vehicle along their SACCO's designated route. The presence of severe competition within a route explains the dangerous driving that 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 in Kenya, and 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). There are no regulations placed on the aesthetic of the vehicle. Nevertheless, 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 and 5 USD for trips to the outskirts.

This transportation industry is appealing to study for a number of reasons. First, it is representative of many other informal transit systems worldwide, including Tanzania's dala dala's, Haiti's tap tap's and India's rickshaws, among others. Moreover, the sector is economically meaningful 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 up to 5% of the country's GDP (Kenya Roads Board, 2007). Most importantly, this context allows us to overcome major data constraints that have limited previous research in the space. Namely, we 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.

2.2 Driver and owners

In this study we work exclusively with small firms that meet three basic criteria. First, the owners of these vehicles manage their matatu themselves, as opposed to hiring a third party manager. Second, the owners are not the primary drivers of the vehicle.[8] These two conditions were designed to focus the research on the classical principal-agent relationship. Finally, we only worked with owners that had a single matatu at the time of recruitment. We chose single owner-driver pairs to remove any

---

[8]Owners do not operate the vehicles themselves for two reasons. First, it allows them to pursue side-jobs that are more lucrative than being a driver. Second, driving is a tough job that individuals like to avoid if they have other options.

dynamics that arise from one driver reporting differently from another, sending competing signals for the owner to parse through. According to an exploratory survey we conducted in the pre-pilot phase of the experiment, approximately 25% of matatu owners in the general population meet these three criteria.[9]

In this industry, owners have settled on a fixed rent contract with limited liability that is negotiated daily. Owners rent their vehicles to a driver every morning for an agreed upon "target price" (henceforth referred to as the 'target'). 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 has had a good day. In the event that the driver cannot make the target, they are supposed to provide the total revenue they earned to the owner. In practice, drivers under-report total revenue to make sure they have some income left over. If they fail to make the target too many times, or they are caught under-reporting too frequently, they will be fired. In addition to choosing the amount of revenue they declare (under-reporting), drivers decide the number of hours they work (effort), their driving style (including driving maneuvers that may damage the vehicle). Owners cannot directly observe these three actions by the driver (under-reporting, effort, damaging driving), and must resort to costly monitoring techniques. 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 negotiation process is repeated daily over the phone (and occasionally in person). Formal documents are not signed because legal recourse is virtually non-existent. Typically owners and drivers have worked with each other for just over two years. The target price is set at approximately 3000 KES (30 USD) every day, and drivers make this target 44% of the time. On days when they do not report making above the target, they under-report revenue by approximately 700 KES. Drivers are typically on the road between 12-14 hours per day, and make approximately 10 trips to and from the city center.

This contract structure appears to be one of the only viable alternatives in this industry. A fixed wage payment is unattractive to most owners because drivers face incentives to undersupply effort when they cannot be monitored. The few SACCOs that have adopted this payment scheme have hired full time managers who supervise the drivers closely. Anecdotal evidence suggests that drivers also dislike this remuneration scheme because it eliminates the large windfall they receive on the best days. Next, a fixed rent contract is impossible to enforce because drivers are poor and hence the limited liability constraint always binds. This leaves the traditional sharecropping model or a fixed rent contract with limited liability. A sharecropping contract in this industry would have to take the form of a profit-sharing agreement where owners and drivers are each responsible for their share of the costs. However, mosts of the costs that the vehicle incurs are beyond the means

---

[9]If we allowed owners to posses two or three matatus, over 50% of matatu owners satisfied these conditions.

of a matatu driver. A typical service fee is 2000 KES, which the driver simply cannot pay upfront. Similarly, in the extreme case that a matatu gets into an accident, drivers are known for running away from the scene. Moreover, a sharecropping model with unobserved output means that drivers can consistently under-report the amount they collect. The cost of under-reporting is low because drivers can easily hide undisclosed revenue. These limitations may explain why owners find this contract structure is less attractive.

A fixed rent contract with limited liability ensures that drivers face incentives to supply effort when they earn more than the target. It also limits under-reporting to days when the driver does not report making the target (there is no incentive to under-report on good days because they are the residual claimant). Note that due to the limited liability constraint, the supply of effort under this contract scheme will be less than the first-best outcome because the driver does not supply optimal effort on days when they do not expect to reach or exceed the target. This contract structure 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.

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 owners' 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 (Figure 1). The first tab was a map of Nairobi and presented the real-time location of the vehicle. By entering a specific date and time interval into the phone, the app would display the exact routes traveled by the matatu over this time period. This first tab provided owners with a more accurate measure of driver effort because they could track where the driver was at any point in time. It also conveyed a more accurate measure of damaging driving because they could see if the driver was operating on bumpy routes. 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 had occurred on the map. It is important to remember that these safety alerts are not necessarily correlated with damaging driving behaviors such as off-route driving. 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. This could be used by owners to estimate total revenue more precisely. The safety section of this page also 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 appeared for scores of 60% and above, a sideways thumb for scores between 40% and 60%, and a thumbs down 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 along various routes across the city. We organized several large meetings with matatu owners in each SACCO, presenting the study's goals and methodology. All owners were informed at the time of recruitment that a monitoring device would be placed in their vehicle free of charge, and 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 two month period. It took approximately four months to recruit enough participants across 9 major commuter routes (Figure 2). 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 (owners had to own a single matatu, which they rented to a driver, and manage the firm's operations themselves).

10

3.2   Installations

The first installation took place in November 2016, and continued until April 2017. The field team, managed by EchoMobile, was able to fit approximately 15 matatus per week with a device (Figure 3). The team scheduled a time to meet each owner individually at a location of their choosing. The owner was compensated for the time their vehicle spent off-road to perform the installation of the device with a one-time payment of 5000 KES (50 USD) 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 three staff (an enumerator, a field manager, and an engineer). While the mechanic worked on fitting the device in the matatu, the field manager took the owner aside to re-explain the purpose of the research project and the tracking device's functionality. For owners in the treatment group, the field manager conducted an additional training on the app. At the same time, 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 the driver, 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 of the data they shared would remain 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 be visible on the SmartMatatu app. A small survey was administered to the owners at the end of their training to make sure they knew how to find all the information contained in the app. Despite this in-depth training, it took owners a few months to feel comfortable navigating the different tabs in the app. We informed treatment and control drivers that a tracker would be placed in their vehicle. We did not mention, however, whether the information would be transferred to the owner. This meant that any subsequent changes we observed in driver behavior could only come from owners using the tracker data, rather than from receiving different information from the enumerators during the installation.

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 then randomly split into two groups: a one-month treatment group and a two-month treatment group. This was done so we could study whether any changes in driving behavior that might be induced by the incentives would persist after 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 (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 by the number of hours worked. For each driver, we then constructed a distribution of these rates for each alert type and found the percentile into which that day's alert rate fell. We then calculated the weighted average percentile for each driver-day by adding the alert rates for each type, applying weights of 1/3 for speeding and braking, and 1/6 for over-acceleration and turning. The average we computed each day lay between 1 and 100. We assessed the cutoff below which they fell and disbursed their incentives accordingly (fewer safety violations resulted in a lower percentile and a higher payout).

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 up with the endline. Finally the GPS tracker collected a wealth of data that we use to measure driving behavior, including safety violations.

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 notification on their phone 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 of the day; the amount the driver delivered to the owner; any repair costs incurred; an overall satisfaction score for their driver's performance (bad, neutral, good); and whether the driver was fired/quit that day. Once the report was successfully submitted, owners received 40 KES via M-Pesa (a mobile phone-based money transfer service). We collected similar information from drivers through SMS surveys (because the drivers were not provided with 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 40 KES via M-Pesa to incentivize consistent reporting.[10]

We developed a set of processes for checking and validating the daily data we received from owners and drivers. Echo Mobile wrote code to check for anomalies including outliers and entries that did not make sense and/or suggested the owner/driver may not have 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 where owners and drivers were able to justify their responses, the enumerators would record their justifications in an excel spreadsheet. In cases when owners and drivers revised their responses, this data was corrected on the server.

4.2 Monthly Surveys

We conducted eight rounds of surveys. We first administered the baseline survey 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 presented owners with 500 KES 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 the 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 900 KES. 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 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 submitting the daily reports in the SmartMatatu app. Finally, they were designed to collect some basic data. As owners and drivers reached the 6-month mark, we conducted an endline survey to measure changes in key outcomes, and to assess the impact of the information treatment and the cash incentives.

---

[10]Note we do not see different submission rates, or differential reporting, between treatment and control - as detailed in appendix Tables 10 and 11).

4.3 Tracking data

The CalAmp tracking device transmitted high frequency data on forward/backward/lateral/vertical acceleration, jerk, location and a timestamp. We use the raw measures of acceleration to investigate changes in driver behavior. Specifically, we look at vertical and lateral acceleration to determine whether the driver is operating on bumpier stretches of road. Furthermore, we use the GPS data to calculate how far each vehicle is from the route they are licensed to be on at any point in time. This provides a measure of how far the driver is deviating from the actual route. Figure 4 depicts the number of times vehicles licensed to route 126 pass through a particular lon-lan cell. The first panel clearly shows what the route should be, and the second panel overlays the designated route to confirm. The figure illustrates that off-route driving is relatively common practice.

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 to avoid capturing 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 aggregated 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 purpose of the model is to generate key predictions about how the monitoring device affects the principal-agent relationship. The owner sets a target in KES that the driver must deliver by the end of the day. The driver chooses the amount of effort and damaging driving they engage in. Damaging (i.e risky) driving refers to the set of behaviors that damage the company asset (which are not necessarily correlated with the unsafe driving metrics we will discuss in the results section). The driver also decides when they will under-report revenue, and by how much - because revenue is unobserved to the owner. This results in five key parameters, and we derive comparative statics for each one. For simplicity we assume that both owners and drivers are risk-neutral.

Given that all of the contracts in this transportation sector follow this fixed form, and they do not change over time, we restrict our attention to contracts of this type (detailed in more depth right below). We then derive comparative statics resulting from the introduction of the monitoring technology, which affects the *parameters* of the contract (not the *type* of contract), and driver behavior. There are three additional departures from the classical setting: 1) the driver selects both effort *and* risk levels, 2) production is unobservable to the principle (allowing the driver to misrepresent production), and 3) the principle observes a signal of production.

5.1 Status Quo

The model is comprised of four steps that correspond to the owner-driver daily interaction.

1. First, the owner sets the target (T), and asks the driver to deliver it by the end of the day.

2. Second, the driver chooses how much effort (e) and damaging driving (r) to engage in. The days' random events unfold (ε), and total revenue is generated. We define total revenue q as:

$$q = e + r\varepsilon$$

where $r \geq 1$ (it is impossible for drivers to avoid risky driving entirely). This production function says that an increase in effort shifts the distribution of output to the right, while an increase in risk reduces the variance of revenue while keeping the mean constant (this is similar to (Ghatak and Pandey, 2000)). We assume drivers choose the amount of effort and risk *before* ε is realized. A driver that chooses to behave in a really risky manner will experience more extreme outcomes depending on the days events. For example, if he chooses to drive off-route a lot (higher risk), he may be caught by the police and impounded (earning very low revenue that day) or he may get away with it and earn high revenue. While $q = q(e, r, \varepsilon)$, I will suppress the arguments for notational convenience throughout the rest of the paper. Note: q is observed to the driver, because they collected passenger fares throughout the day, but it is not unobservable to the owner.

3. Third, the driver chooses how much revenue to report to the owner $(\tilde{q})$.

4. Finally the owner decides whether or not to punish based on 1) whether he detects any damages to the vehicle, and 2) whether he can be sure the driver is under-reporting revenue. Note these punishments are non-monetary (firing, reprimands) because the driver's limited liability constraint binds - they cannot be expected to pay monetary costs.

We solve the model in 4 steps, via backward induction.

Step 1: Owner's punishment

The owner punishes the driver for two types of behaviors. First, they will punish the driver if they detect any damages to the vehicle. The expected punishment is expressed as βr, where β is the probability the owner detects risky driving that damages the vehicle, and r are the damages incurred. Note that without a monitoring technology β is close to zero because the owner has a hard time detecting behavior that damages the vehicle.

Second, the owner will punish based on the driver's reported revenue $(\tilde{q})$. We know the owner (principal) sets a target amount (T) at the beginning of the day, which they want the driver (agent) to deliver by the end of the day. The driver earns revenue (q) from passenger fares, and decides how much to report to the owner $(\tilde{q})$. Owners do not observe actual revenues that are collected by drivers from passenger fares throughout the day (q), but they receive a signal $(\hat{q})$ about what true revenue (q) should be. This signal comes from the information owners gather about driving conditions and driver behavior throughout the day. Absent a monitoring technology, the owners signal will be informed by listening the news, visiting the designated route, or talking to their

friends. We assume the owner's signal is noisy but unbiased, so that $\hat{q}$ is defined as follows:

$$\hat{q} = q - \sigma$$
$$\sigma \sim F\left(0, \frac{1}{\alpha}\right)$$

where α is the precision of the owners' signal about true revenue. Any monitoring technology we introduce will provide more information to the owner about driver behavior. This will increase the precision of the owner's signal about what revenue should be, which gives the driver less leeway to significantly under-report revenue on a particular day. In the case of our monitoring technology specifically, the owner can observe the number of kilometers driven, and where the driver was operating at any point in time. Owners can use this information to estimate the number of trips to and from the city center, which provides a more accurate measure of total daily revenue.

In principle the owner's decision to punish the driver could depend on 1) the level of reported revenue ($\tilde{q}$), and 2) the signal of actual revenue ($\hat{q}$). In theory, punishing based on the level of $\tilde{q}$ is rational because owners may be able to infer that their driver invested low effort or took too much risk from a very low $\tilde{q}$. Alternatively, owners could punish based on ($\hat{q} - \tilde{q}$) because this difference reflects how much the owner thinks the driver is lying. In practice, the most effective punishment will depend on owners ability to infer e and r from reported $\tilde{q}$, which will depend on the variance of ε. If the variance of ε is high, then the owner cannot infer e and r from a low $\tilde{q}$. Anecdotally, it appears the variance of epsilon is large, which means the owner will place more weight on ($\hat{q} - \tilde{q}$) (which owners in our sample confirmed). For ease of exposition we assume that the punishment is quadratic in the difference between the owners signal $\hat{q}$ and the driver's offer $\tilde{q}$ - where the distributional assumptions that rationalize this are presented in the Appendix. Note that because $\hat{q}$ is an unbiased signal of q, and σ is independent of $\tilde{q}$, the punishment can also be expressed as an increasing function of the gap between the true revenue (q) and the offer ($\tilde{q}$) (or in this case, a quadratic of the difference). All comparative statics go through as long as this punishment remains an increasing function of the in gap between $q - \tilde{q}$. The punishment is expressed as:

$$E[punishment] = \frac{\alpha}{4}(q - \tilde{q})^2$$

Step 2: Solve the agent's optimal reporting schedule

The driver needs to choose how much revenue to report to the owner ($\tilde{q}$). Broadly, the driver can choose to report truthfully, or to under-report, and there will be some threshold of revenue beyond which they will choose to truthfully report.

Case 1: The driver chooses to reports above the target ($\tilde{q} > T$)
When the agent chooses to report above the target, they do not face any incentive to lie because they keep everything they earn above the target and the owner cannot renegotiate the terms of the

contract. Therefore they report truthfully $\tilde{q} = q$ and their utility is

$$U^D = \overbrace{q - T}^{salary} - \overbrace{\beta r}^{damages}$$

Case 2: The driver chooses to report below the target $(\tilde{q} < T)$

When the agent chooses to report below the target, they face an incentive to lie in order to increase their take-home pay. Indeed, on days when the driver truly fails to make the target $q < T$, under-reporting revenue means they get to walk away with some money rather than handing it all over to the owner as stipulated by the contract. Even when the amount of revenue is slightly above the target $q > T$, the driver faces some incentive to lie in order to walk away with slightly more income than what they would if they reported truthfully and had to hand over the entire target amount. On the days that the driver decides to report below the target, they know they can be punished by the owner. The driver then needs to choose the amount of revenue to report $(\tilde{q})$ to maximize their utility:

$$\max_{\tilde{q}} \quad U^D = \overbrace{(q - \tilde{q})}^{salary} - \underbrace{\frac{\alpha}{4}(q - \tilde{q})^2}_{punishment} - \overbrace{\beta r}^{damages}$$

Solving for $\tilde{q}$ yields:

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

where

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

This says that the optimal amount for the driver to under-report is a function of the owner's signal. More specifically, it is optimal for the driver to under-report by a constant amount $\left(\frac{2}{\alpha}\right)$. Figure 5 shows that the data confirms this behavior. The graph summarizes values of under-reported revenue at unique/binned values of net revenue above the target.[11] We see that drivers continuously under-report approximately 700 KES (7 USD) until the net revenue they generate exceeds the target by approximately 500-1000 KES. This graph confirms that our model does a good job of predicting the behavior we observe in the data.

Switch point: Next we need to determine the point at which the driver is indifferent between

---

[11]On the y-axis we have under-reported revenue, which we compute from driver and owner's daily surveys. When the driver does not make the target (owner's income is less than the target), any salary the driver reports to us is under-reported revenue because they should have handed this over to the owner. When the driver makes the target (the owner's reported income equals the target), there is no under-reporting and this variable is set to zero. On the x-axis we plot net revenue above target, defined as owner income + driver salary − target. Indeed, to get an accurate measure of under-reporting 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.

reporting above the target (and telling the truth) and reporting below the target (and under-reporting). When the driver tells the truth i.e $\tilde{q} = q$, they get utility:

$$(q - T - \beta r)$$

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

$$(q - \tilde{q}) - \frac{\alpha}{4}(q - \tilde{q})^2 - \beta r = \frac{1}{\alpha} - \beta r$$

Setting the two utilities equal and solving:

$$q - T - \beta r = \frac{1}{\alpha} - \beta r$$
$$q^* = T + \frac{1}{\alpha}$$

Where

$$\frac{\partial q^*}{\partial \alpha} = -\frac{1}{\alpha^2} < 0$$
$$\frac{\partial q^*}{\partial T} = 1 > 0$$

This says that the revenue required to truthfully report, (q^*), is a function of the owner's signal and the target. It is perhaps useful to relate this result that $q^* = T + \frac{1}{\alpha}$ to the optimal reporting behavior $\tilde{q}$. In particular, if the driver observes q at the upper bound of $[T, T + \frac{1}{\alpha}]$, he will report $\tilde{q} = q - \frac{2}{\alpha} = T - \frac{1}{\alpha}$, which is indeed less than the target.

Step 4: Driver's optimal choice of effort

The driver chooses two actions, effort (e) and risk (r), which affect the probability distribution of revenue. Following (Ghatak and Pandey, 2000) we define q as:

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

where ε is a random variable that reflects the idiosyncratic events of the day (weather and traffic). The driver chooses effort and risk to maximize his utility

$$\max_{e,r} \quad \underbrace{E\left[(q - T - \beta r) \mid q \geq q^*\right] \cdot Pr(q \geq q^*)}_{Truth} + \underbrace{E\left[(q - \tilde{q}) - \frac{\alpha}{4}(q - \tilde{q})^2 - \beta r \mid q < q^*\right] \cdot Pr(q < q^*)}_{Under-report} - \underbrace{h(e,r)}_{Cost}$$

Where h(e,r) is the driver's private cost for his actions, and we assume that this function is twice continuously differentiable, monotonically increasing and convex in e and r. The optimization problem yields the following F.O.C with respect to e and r, respectively (all the derivations in the

18

paper can be found in the Appendix):

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

$$\int_{\frac{q^* - e}{r}}^{\infty} \varepsilon f_\varepsilon(\varepsilon) d\varepsilon - h'_r - 2\beta = 0$$

Where

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

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

$$\frac{\partial r}{\partial \beta} < 0 \ \& \ \frac{\partial e}{\partial \beta} > 0$$

This means that the driver's choice of effort and risk is a function of the owner's signal, the target, and the probability the owner detects damaging driving. Intuition for the sign of these partial derivatives will be provided in the next section when we introduce the monitoring technology.

Step 5: Owner's choice of the target

Constrained case

The owner chooses T to maximize his utility:

$$\max_T \quad T \cdot Pr(q \geq q^*) + E[\tilde{q} \mid q < q^*] \cdot Pr(q < q^*) - \gamma(r) \quad \text{s.t}$$

$$E\left[(q - T - \beta r) | q \geq q^*\right] \cdot Pr(q \geq q^*) + E\left[(q - \tilde{q}) - \frac{\alpha}{4}(q - \tilde{q})^2 - \beta r \mid q < q^*\right] \cdot Pr(q < q^*) - h(e^*, r^*) > 0$$

Where $\gamma(r)$ are the costs owners incur for any damages to the vehicle they detect that need repairs. The constraint is the driver's participation constraint, where we assume for simplicity that they have a reservation wage of zero. This optimization problem yields the following F.O.C with respect to T and λ, respectively:

$$\frac{\partial}{\partial T} = 1 - F_\varepsilon(\cdot) + \frac{\partial e}{\partial T} F_\varepsilon(\cdot) + \frac{\partial r}{\partial T} \int_0^{\frac{q^* - e^*}{r^*}} \varepsilon f_\varepsilon(\varepsilon) d\varepsilon - \frac{\partial}{\partial T}\left(\frac{q^* - e^*}{r^*}\right)\left(\frac{1}{\alpha}\right) f_\varepsilon(\cdot) - \gamma'(r)\frac{\partial r}{\partial T} +$$

$$\lambda\left[-(1 - F_\varepsilon(\cdot)) + \frac{\partial e}{\partial T}(1 - F_\varepsilon(\cdot)) + \frac{\partial r}{\partial T}\int_{\frac{q^* - e^*}{r^*}}^{\infty} \varepsilon f_\varepsilon(\varepsilon) d\varepsilon + \frac{\partial r}{\partial T}(-\beta) - h'\left(\frac{\partial e}{\partial T} + \frac{\partial r}{\partial T}\right)\right] = 0$$

$$\frac{\partial}{\partial \lambda} = \int_{\frac{q^* - e^*}{r^*}}^{\infty} (e^* + r^*\varepsilon - T - \beta r^*) f_\varepsilon(\varepsilon) d\varepsilon + \int_0^{\frac{q^* - e^*}{r^*}} \left(\frac{1}{\alpha} - \beta r\right) f_\varepsilon(\varepsilon) d\varepsilon - h(e^*, r^*) = 0$$

Unconstrained case

The owner chooses T to maximize his utility:

$$\max_{T} \quad T \cdot Pr(q \geq q^*) + E[\tilde{q} \mid q < q^*] \cdot Pr(q < q^*) - \gamma(r)$$

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

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

5.2 Monitoring Technology

Introducing the monitoring technology increases α and β. In other words, the precision of the owner's signal about output (α), and the probability of detecting damaging driving (β), increase. In what follows we consider how the owner changes the terms of the contract, and how this subsequently affects the driver's choices of effort, risk and under-reporting. We assume throughout that the driver's participation constraint binds.[12] As α and β increase:

Proposition 1 *(Target)*

$$\frac{\partial T}{\partial \alpha} < 0 \quad and \quad \frac{\partial T}{\partial \beta} < 0$$

This says that as the precision of the owner's signal (α) increases, the owner will reduce the target. We know that the driver's utility will fall as α increases. Because the constraint binds, the owner needs to reduce the target to ensure the driver makes their reservation wage. The same reasoning explains why the owner reduces the target as the probability of detecting damaging driving (β) increases.

Proposition 2 *(Damaging (Risky) Driving)*

$$\frac{dr}{d\alpha} = \overbrace{\frac{\partial r}{\partial \alpha}}^{-} + \overbrace{\frac{\partial r}{\partial T}}^{+}\overbrace{\frac{\partial T}{\partial \alpha}}^{-} < 0$$

$$\frac{dr}{d\beta} = \underbrace{\frac{\partial r}{\partial \beta}}_{-} + \underbrace{\frac{\partial r}{\partial T}}_{+}\underbrace{\frac{\partial T}{\partial \beta}}_{-} < 0$$

This says that risk will unambiguously decrease when α increases. There are two effects at work. As α increases, the driver is going to have to truthfully report more often (q^* falls), which increases the probability they are the residual claimant and must bear the cost of a bad revenue day. Being exposed to greater downside risk makes damaging behavior less attractive, $\frac{\partial r}{\partial \alpha} < 0$. Similarly, as

---

[12]Anecdotally we know that there are a lot of drivers on the market and we think it reasonable to assume that they would have been bargained down to their constraint. Switching costs are not trivial and the owner would prefer less turnover all else equal.

the owner reduces the target in response to a higher α, (q^* falls), and the driver will take on less risk, $\frac{\partial r}{\partial T} > 0$. Next, the driver's immediate response to an increase in probability of detection (β) is to reduce risk, $\frac{\partial r}{\partial \beta} < 0$. Moreover as the owner reduces the target in response to higher β, the driver will further reduce the amount of risk they take.

Proposition 3 *(Effort)*

$$\frac{de}{d\alpha} = \overbrace{\frac{\partial e}{\partial \alpha}}^{+} + \overbrace{\frac{\partial e}{\partial T}}^{-}\overbrace{\frac{\partial T}{\partial \alpha}}^{-} > 0$$

$$\frac{de}{d\beta} = \underbrace{\frac{\partial e}{\partial \beta}}_{+} + \underbrace{\frac{\partial e}{\partial T}}_{-}\underbrace{\frac{\partial T}{\partial \beta}}_{-} > 0$$

This says that effort will unambiguously increase when α increases. Yet again, there are both direct and indirect effects driving this result. As the precision of the owner's signal, α, increases, the driver will supply more effort, $\frac{\partial e}{\partial \alpha} > 0$. Indeed, for every level of output below q^*, the driver can no longer under-report as much as they used to because the owner has a more precise signal of what output should be. Holding all else constant, their revenue will decrease if they do not respond by increasing effort and generating more revenue. Next, we know that the owner responds to an increase in α by lowering the target. As the target decreases the driver is more likely to meet it, which increases the returns to effort and incentivizes the driver to work more, $\frac{\partial e}{\partial T} > 0$. Turning next to the effects of better detection of risky driving (β), we know from above that the driver will have to reduce the amount of risk they take, which reduces the probability of large windfall days. The driver has to compensate by investing more effort, $\frac{\partial e}{\partial \beta} > 0$.

Proposition 4 *(Switch Point)*

$$\frac{dq^*}{d\alpha} = \overbrace{\frac{\partial q^*}{\partial \alpha}}^{-} + \overbrace{\frac{\partial q^*}{\partial T}}^{+}\overbrace{\frac{\partial T}{\partial \alpha}}^{-} < 0$$

As the owner's signal becomes more precise, the revenue required to truthfully report and make the target q^* will decrease, $\frac{\partial q^*}{\partial \alpha} < 0$. Similarly, as the target decreases, the revenue required to truthfully report and make the target q^* will decrease, $\frac{\partial q^*}{\partial T} > 0$. Together these two factors drive the reduction in q^* we anticipate with the introduction of the monitoring technology.

Proposition 5 *(Under-reporting)*

$$\frac{d\tilde{q}}{d\alpha} = \frac{2}{\alpha^2} > 0$$

This says that the optimal amount for the driver to report ($\tilde{q}$) increases with the precision of the owner's signal (α). It follows mechanically that the amount they under-report will fall.

6 Results

We test each one of these predictions in our data. We first investigate how the technology affects the target contract. We then examine drivers response along the key dimensions detailed above, including 1) effort, 2) damaging driving, and 3) reporting behavior (probability of making the target, and under-reporting revenue). Before reviewing whether these behaviors change in the way we would expect from the model, we briefly provide some summary statistics and evidence that owners were interacting with the device. Finally, we conclude this section with an overview of the device's impact on safe driving (which is not highly correlated with the our measure of damaging driving).

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 juggle these two responsibilities. Typically owners have worked with their current driver for the past 2 years. Drivers have very similar profiles - which we expect - 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. 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 Information Treatment Arm

To study the treatment effect of information on contracts, productivity (which includes effort, damaging driving, under-reported revenue, and profits) and safety, over the 6 month time frame we run the following regression model:

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

where y_{ird} is a daily contract/productivity/safety outcome for owner i on route r, on day since installation d; D_{im} is a treatment indicator equal to 1 if the owner is in the information group in month m (which allows us to examine the treatment effect as it evolves over the six months of the study); α_d is a day fixed effect; τ_r is an assigned route fixed effect; X_i is a set of firm-level baseline specific controls included for precision, and ε_{irmd} is an error term.[13] We cluster our standard errors at the firm level. Note that the study offered the information to control owners in months 7 and 8

---

[13]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.

(as compensation for participating in the study). As a result all the regressions only include data before month 7.[14]

We also have an endline survey that asked owners about their use of the device, perceptions of drivers' 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   Usage

First, we monitor owners' usage of the device. We do so by tracking the API calls that are generated every time the owner logs into the app and requests different pieces of information (including historical location, up-to-date summary information, and where the safety violations occurred on the map). Figure 6 calculates the share of owners that made at least one API call during the week. We find high rates of take-up. In the early months of the study approximately 80% of owners are checking the app at least once a week. This share decreases but stabilizes at about 70% as the study progresses. A large share of owners are also using the app daily. In the first few months, 60% of owners check the app once a day, and 40% continue their daily usage after 6 months. This suggests that owners are actively engaging with the device throughout the study.

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 treatment group are 27 percentage points more likely to know about the number of kilometers driven and 45 percentage points more likely to know about the the instances of off-route driving (Table 3 Columns 1 and 2). We do not find any differences between the treatment and control groups regarding knowledge of exact revenue (Column 3). This last result is not altogether surprising. While the device provides a better *estimate* of revenue by revealing the number of kilometers the vehicle has covered, and the location of the vehicle throughout the day, it does not reveal the *exact* revenue because the number of passengers is uncertain. As a final test, we also ask owners to rate how challenging it is 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. In other words, control owners maintain that monitoring is hard while treatment owners reveal that it is easy (Table 4).[15]

---

[14]All of the results from these regressions are presented through figures because it is easier to see the how owners' and drivers' behavior changes over time. Please see the appendix for the specific point estimates from these regressions, and alternate regression specifications (Appendix Tables 12 though 20).

[15]Column 2 of Table 4 also shows that control owners report no change in the time they spend monitoring during

23

We do not, however, find any significant changes in traditional monitoring behaviors (checking-in with the driver over the phone, at the terminal (locally referred to as "stage"), or through a third party).

Finally, we investigate whether there are any changes in how the owners and drivers interact. 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 driver's behavior. Figure 7 suggests that the number of reprimands is marginally higher in the treatment group at the beginning of the study period. This is 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 investigate whether owners take more extreme actions and fire their drivers more frequently. While the trend in Figure 7 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).

6.3.1 Contracts

We first investigate whether access to the tracking information changes the *terms* of the contract. While the intervention could also have 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 limited liability constraints bind, and revenue is unobserved/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 every owner maintained the same type of contract.

As detailed in the model, owners can use the information from the device to change the target they set for drivers (Proposition 1). Absent the technology, 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 8 depicts the estimated treatment effect on the owners' 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 daily target amount is 135 KES (1.35 USD) below the control group, representing a 4.1% decrease (0.2 standard deviation). 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 constraint. When the

---

the study period, while 70% of treatment owners reveal they spend less time monitoring.

constraint binds, the owner needs to decrease the target in order to compensate the driver for their lost information rents. Lowering the target also reduces owners' revenue on days where the driver makes the target. As a result it is only profitable for the owner to reduce the target if the owner is compensated in other ways, namely with a higher share of revenue on days when the driver does not make the target, fewer damages to the vehicle, or an increase in the frequency with which drivers make the target. We turn to these results next.

6.3.2 Productivity

We consider three measures of productivity, which correspond to the choices that drivers make throughout the day. This includes how damaging (risky) they will drive, how much effort to supply, and the amount of revenue they disclose to the owner (which is either the target amount, or some amount below).

1. Damaging (risky) driving: We hypothesize that owners prefer less damaging driving than what the drivers would optimally choose. With the technology, owners can observe driving behavior more accurately, and the probability they detect damaging driving increases. This reduces drivers' incentive to drive in ways that damage the vehicle. Similarly, as the precision of the owner's signal about revenue improves, and they decrease the target, the driver will be exposed to greater downside risk on bad revenue days.[16] This should further reduce their incentive to drive in ways that damage the vehicle (Proposition 2). Figure 9 confirms this hypothesis in the data. We see damages substantially decrease throughout the entire 6 month period. In month 3, daily repair costs for treatment owners are reduced by 125 KES (1.25 USD), and continue falling until month 6, where they are 226 KES (2.26 USD) less than what control owners incur on average (this represents a 46% decrease in daily repair costs). These repair costs represent a major business expense for owners, which makes the impact of the monitoring technology significant.

We want to confirm that this result stems from less damaging driving behavior. One of the greatest sources of damaging driving is off-route driving. Drivers often take shortcuts on bumpy roads that are notoriously damaging to matatus. These shortcuts are appealing to the driver because they help them travel to the city center more quickly, and avoid traffic jams where they sit idly without picking up any passengers. Typically owners cannot observe off-route driving and drivers cannot be expected to pay for vehicle repairs. This means that damaging driving along these alternate routes is costless to drivers but offers the opportunity for large windfall days. When owners have access to the monitoring technology, however, they can inform drivers about how to take better care of the vehicle, and mandate that they stay on their designated routes. To investigate this hypothesis, we compute the distance between each GPS point recorded by the device, and the vehicle's licensed route. Figure 10 demonstrates that treatment drivers are

---

[16]As detailed in the model, the switch point q^* shifts down. This means that there are days when the driver used to under-report, and they now truthfully report. As the residual claimant on these days, they bear the cost of a bad revenue day (which is higher when they make more damaging maneuvers on the road).

on average 400 meters closer to the designated route than control drivers throughout the study period. To confirm that the reduction in off-route driving is responsible for fewer damages, we investigate whether the distributions of lateral and vertical acceleration differ across treatment and control groups. Taking fewer bumpy roads that jostle the vehicle should be visible in the acceleration data. Lateral acceleration measures tilting from side to side, while vertical acceleration captures movement upwards and downwards. We find suggestive evidence that driving behavior has changed. The distribution of lateral acceleration in the treatment group tightens around 0 (less tilting - Figure 11). Similarly, the distribution of vertical acceleration has more mass around gravity (normal driving) for the treatment group. We can reject equality of these distributions across treatment and control by applying a K-S test, which returns a p-value of 0.000 for both measures of acceleration.

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 repair costs, and the device reduces their incentive to do so because they are more likely to be caught in the lie.[17] This cannot be the case for larger repairs, however, because they 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). The second panel in Figure 9 demonstrates that the probability of incurring a large repair cost decreases significantly (7-8 percentage points). This implies that the decrease in the repair costs that we observe cannot be entirely driven by drivers inflating the repair costs. Drivers are also changing *how* they drive as the result of the technology.

2. Effort: Next, we proxy driver effort by the number of hours the tracking device is on (the device powers on and off with the matatu). When the device is installed in the matatu, drivers know that owners have a more precise signal of output, which means they are more likely to get caught if they under-report heavily. This encourages drivers to under-report less, which reduces their take-home pay. As the model demonstrates, this creates an incentive for drivers to invest more effort throughout the day so they can increase total revenue and ensure they maintain similar compensation. In parallel, owners have lowered the target, which means that it is more likely that the driver will become the residual claimant. Finally, the model predicts that drivers will compensate for the reduction in damaging driving by investing more effort. For all of these reasons we expect effort levels to rise (Proposition 3). This prediction is borne out in our data: Figure 12 demonstrates the upward trend in effort that we anticipated. The number of hours the tracking device is on increases by 0.98 hours per day on average in month 3 and rises steadily until the end of the study. By month 6, effort levels increase by 1.47 hours per day on average in the treatment group. This represents a 9.9% increase in drivers' labor supply. This is substantial in an environment where drivers are already working 14 hour days. With more hours on the road, we also see the number of kilometers increase by 12 kilometers per day on average, which corresponds to an extra trip to/from the city

---

[17]Note that drivers pay for some of the smaller repair costs out-of-pocket, if they can be handled quickly by a mechanic during the day. They subsequently report the amount they spent to the owner, which we would then capture in the owner survey.

center (Figure 12).

3. Reporting Behavior: Once the driver chooses how much effort and risk to invest, they need to decide whether or not to truthfully report. The model predicts that the optimal switch point for truthfully reporting (q^*) shifts down because 1) the owner's signal of true revenue becomes more precise, and 2) the owner has lowered the target (Proposition 4). Testing this proposition is difficult because we do not have a direct measure of q^*; we only observe whether or not drivers make the target (when owner reported income is equal to the target). However, a lower q^* implies that drivers should make the target more often, primarily on days when revenue is close to q^* to begin with. To investigate this prediction in the data, we first apply our standard regression specification to determine whether we see a significant change in the probability of making the target.[18] Figure 13 suggests that from month 3 onwards, the rate at which drivers make the daily target increases by 11 percentage points from a base of 44 percent (significant in month 3 only). It is not altogether surprising that this result is slightly weaker because the analysis considers the full range of revenue rather than focusing on days when drivers are close to q^* to begin with (i.e. close to making the target). This 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.[19] We then compute drivers' daily reported revenue above target and subtract the average expected amount. This is akin to including route fixed effects - because we know that a certain level of revenue above target will be acceptable on certain routes but not on others. We are left with a measure of daily deviation from expected revenue above target, which we plot in the second panel of Figure 13. The revenue above target measure has an approximate mean of 4,000 KES. 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 (from plotting the amount drivers under-report in Figure 5). Figure 13 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.

Finally, we anticipate that drivers' reporting behavior ($\tilde{q}$) will change as the monitoring devices are introduced. According to the model, we should see drivers under-reporting below some optimum q^*, at which point they will start truthfully reporting and providing the owner with the target amount. Below this optimum, the model predicts that drivers will under-report by a constant amount. This is consistent with the idea that drivers have some reservation wage they do not want to fall below. We predict that the monitoring technology will decrease the amount of under-reporting we see in the data. Owners can use the device to estimate actual revenue more accurately, and they are more likely to detect when the driver is under-reporting. Drivers should respond by

---

[18]The driver has "made the target" if the owner's reported income is equal to the target. When the owner's reported income is below the target, the driver has not made the target.

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

lying less everywhere below the threshold. Figure 14 depicts under-reporting across treatment and control groups, to which we apply a non-parametric smoothing function. We observe constant under-reporting below some threshold value q^* in both groups (which falls somewhere between 500-1000 KES). Moreover, we observe that the treatment group under-reports less than the control group. To obtain a more precise estimate for the reduction in under-reporting, we regress the amount drivers under-report on treatment status for different possible q^* (between 500-1000). The regression only considers data below q^* because this is where the model predicts shading will occur.[20] The results in Table 5 confirm that the amount drivers under-report falls by approximately 70-100 KES ($\approx$ 1 USD) per day depending on the exact location of q^*. As a final check, taking 900 KES as a benchmark value for q^* and keeping all data below this point, Figure 15 applies our standard regression specification over study months. We see that the amount drivers under-report falls by approximately 100 KES throughout the study (except for month 1).

Without knowing q^* exactly, our estimate of 70-100 KES technically includes both a reduction in $\tilde{q}$ and a drop in q^*. We want to confirm that both of these behaviors are indeed happening in reality. We do so by imposing a step function in a regression of under-reported amount on treatment. In other words, we allow under-reporting below q^* and impose zero shading above. We run this regression for every reasonable value of q^* for treatment and control groups. We then plot two outcomes in the second panel of Figure 14. The dots represent the estimated under-reported amount in the treatment (in red) and control (in black) groups across different choices of q^*. We can see that the treatment groups under-report 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 isolate the q^* that 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 confirms that both factors explain the overall reduction in under-reporting that we observed in the more flexible regression specification above.

6.3.3 Company Performance and Employee Well-being

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 17 confirms this is true). The model predicts that the impact on revenue, however, is ambiguous. Improved monitoring increases driver effort, and reduces under-reporting. However, it also reduces the amount of damaging driving

---

[20]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.

they engage in (which we confirmed in our data). Depending on which of these effects dominates revenue could increase or decrease. Panel 1 of Figure 16 illustrates that revenue does not change substantially throughout the study. Taken together, decreasing costs and stable revenues suggest that firm profits will increase. Panel 2 in Figure 16 demonstrates a similar trend to what we've observed to date: profits increase continuously starting in 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 per day = 4.40 USD per day). Taking the average gains over the study period and extrapolating to the full year (assuming matatus operate 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 capture 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 do 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. It is perhaps worth mentioning that our profit gains are in line with 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 (Bloom et al., 2013; Bloom, Sadun, and Van Reenen, 2017; McKenzie and Woodruff, 2016; Berge, Bjorvatn, and Tungodden, 2014; de Mel, McKenzie, and Woodruff, 2014; Valvidia, 2012). Our technology has the added benefit of requiring a single up-front payment for continued use. Moreover it requires relatively little coordination and training.

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, Column 1). This represents an 11 percent 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 (Table 7, Column 2). There are a number of reasons why the monitoring device could have encouraged treatment owners to grow their businesses more actively. First, profits have increased and under-reporting has decreased. Second, our results suggest that owners started trusting their drivers more. Table 6 presents four different measures of owners' perceptions of their drivers at endline. We see owners sending an additional 30 KES to drivers in the trust game the enumerators administered (Column 1) - a 30% increase. Moreover, treatment owners' assessment of whether their drivers' 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 honest (Column 3). We suspect that greater trust in their drivers' abilities/honesty, combined with a reduction in the amount the drivers under-report, makes the process of managing

the company easier. Together with higher profits, treatment owners may have seen an opportunity to expand that did not exist before.

Finally, it is important to investigate 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 drivers' 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 (Figure 12), and the amount they under-report decreases (Figure 14). While their salary per hour remains unchanged (Figure 17), they are potentially worse off from 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 this further, we administered a small survey 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, and no distinction can be drawn between 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 owner (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. While we do not want to lean too much on this qualitative evidence, 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. They also felt that it increased owners' trust in their work, which reduced their stress levels. In an environment where drivers are constantly being second guessed by their employers, this could represent a meaningful improvement in working conditions.

6.3.4 Externalities

The device conveyed information to owners about productivity and safety. A priori, we thought that owners might care more about safe driving than drivers for two reasons. First, owners are the ones to pay for repairs when the vehicle is damaged, and for fines when the car is ticketed or impounded by the police. Second, drivers can increase their take-home pay if they generate higher revenue in fares by breaking safety regulations to pick up more passengers. It follows that the information we provided should have an effect if unsafe driving is correlated with damages, and the owners want to avoid high repair costs. Alternatively, owners should use the information to reform driving behavior if unsafe driving results in more fines than the owner would optimally choose to incur. The Kenyan government also assumed that owners would reform driving behavior when they mandated that long-range buses be equipped with GPS tracking devices in 2016. However, this measure was only marginally successful, and before rolling out the experiment we were conscious that the safety information we provided may not have the intended effects. This may be because owners down-weight the probability of getting into severe accidents, or fined for unsafe driving. Alternatively, the benefits to unsafe driving in terms of increased revenue may be high for the owners as well. If owners care only about profits, and increased effort comes at the expenses of safety, we might expect instances of unsafe driving to increase. Finally, while owners definitely care

about damages to their vehicle, these damages may only be weakly correlated with unsafe driving. If this is the case, then the information we provided may not change safe driving practices.

The device collected six pieces of information that correlate with safe driving: maximum speed, average speed, speeding over 80km, over-acceleration, sharp braking and sharp turning. We do not see any meaningful increases in maximum or average speeds as the study progresses. Similarly, instances of over-acceleration and speeding above 80km do not change significantly (Figure 19). We see no effect on sharp turns or sharp braking (Figure 20). 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 (Figure 21). 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 highlights that owners can incentivize optimal levels of effort without further compromising passenger safety, we cannot necessarily expect owners to curb unsafe driving with these technologies. This is especially important for local governments in Kenya and South Africa to know, as they continue to take steps to curb unsafe driving by introducing monitoring technologies.

6.4   Cash Treatment Arm

Bearing in mind that owners may not act on the safety information we provided, we tested the impact of an intervention that incentivizes drivers to take safety into account. Drivers were offered 600 KES 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. The variable "One-month post treat" compares drivers who never received cash incentives, to the drivers whose incentives have been 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 is statistically 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 to the owner of the minibus about the productivity and safety of the driver. 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 meeting 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, it also highlights that the technology does not remedy the negative safety externalities of the industry.

These results are important for a number of key stakeholders, including small firms operating in the transportation industry, and 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 is relatively understudied empirically: moral hazard in labor contracting. One solution that can potentially ease this friction is improved monitoring. Monitoring is typically difficult in small firms, however, because they cannot hire dedicated staff to oversee employee performance, and it takes time away from regular business operations. In our paper, we demonstrate that introducing cost-effective monitoring technologies can be a worthwhile investment for companies looking to increase their profits and grow their asset base.

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 braking 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 longer term impacts, and general equilibrium effects, of these technologies on firm operations, and worker outcomes. Our study included 255 matatus and 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 and offered more GPS trackers to owners on the various routes. Similarly, the impacts on driver well-being may have changed if all the matatus on the route were fitted with GPS tracking devices. While our results suggest that benefits accrue to both workers and firms in this context, thinking about who gains and who loses as these technologies become

more pervasive is an important area for future work.

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.

Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts. 2013. "Does Management Matter? Evidence from India." *The Quarterly Journal of Economics* 128 (1):1–51.

Bloom, Nicholas, Gregory Fischer, Imran Rasul, Andreas Rodriguez-Clare, Tavneet Suri, Christopher Udry, Eric Verhoogen, Christopher Woodruff, and Gulia Zane. 2014. "Firm Capabilities and Economic Growth." Evidence Paper, IGC.

Bloom, Nicholas, Raffaella Sadun, and John Van Reenen. 2017. "Management as Technology?" Working Paper, NBER.

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, Michael Greenstone, Rohini Pande, and Nicholas Ryan. 2013. "Truth-Telling by Third-Party Auditors and the Response of Polluting Firms: Experimental Evidence from India." SSRN Scholarly Paper ID 2294736, Social Science Research Network, Rochester, NY.

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

Ghatak, Maitreesh and Priyanka Pandey. 2000. "Contract Choice in Agriculture with Joint Moral Hazard in Effort and Risk." *Journal of Development Economics* 63:303–326.

Grossman, Sanford J. and Oliver D. Hart. 1983. "An Analysis of the Principal-Agent Problem." *Econometrica* 51 (1):7–45.

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.

Harris, Milton and Artur Raviv. 1979. "Optimal Incentive Contracts with Imperfect Information." *Journal of Economic Theory* 20 (2):231–259.

Hart, Oliver and Bengt Holmstrom. 1987. *The Theory of Contracts.* Cambridge: Cambridge University Press, advances in economic theory ed.

Hölmstrom, Bengt. 1979. "Moral Hazard and Observability." *The Bell Journal of Economics* 10 (1):74–91.

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

Kenya Roads Board, Roads. 2007. "Kenyan Transport Sector Details Annex 3.1." Tech. rep., Kenya Roads Board.

Macharia, WM, EK Njeru, F Muli-Musiime, and V Nantulya. 2009. "Severe Road Traffic Injuries in Kenya, Quality of Care and Access." *African Health Sciences* 9 (2):118–124.

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.

Payet, Charmel. 2018. "Green Light for Taxi Incentive Programme." *eThekwini Municipality* .

Pierce, Lamar, Daniel C. Snow, and Andrew McAfee. 2015. "Cleaning House: The Impact of Information Technology Monitoring on Employee Theft and Productivity." *Management Science* 61 (10):2299–2319.

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

Shapiro, Carl and Joseph E. Stiglitz. 1984. "Equilibrium Unemployment as a Worker Discipline Device." *The American Economic Review* 74 (3):433–444.

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)	211.9	212.9	-1.03
			(5.09)
Owner age	36.3	37.3	-0.99
			(0.99)
Owner gender	0.18	0.18	-0.0056
			(0.048)
Owner highest level of education	2.94	2.97	-0.030
			(0.11)
Owner is employed in salaried job	0.21	0.24	-0.030
			(0.052)
Years the owner is in matatu industry	7.71	7.71	-0.0066
			(0.79)
Years owner has owned matatus	4.65	4.47	0.18
			(0.52)
Number of drivers hired for this matatu	1.26	1.37	-0.12
			(0.13)
Number of other drivers hired in the past	1.77	1.94	-0.17
			(0.22)
Amount given in trust game	117.7	126.2	-8.50
			(12.4)
Owner Raven's score	4.51	4.65	-0.14
			(0.19)
Driver rating: owner's fairness	8.11	8.33	-0.23
			(0.18)

The summary statistics are calculated using the owners baseline data. The first column shows the mean in the control group, while the second column show the mean in the treatment group. The final column shows the difference between treatment and control. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels.

Table 2: Balance across cash treatment (drivers)

Variable	Control	Treatment	Difference
Driver age	34.4	37	-2.64
			$(0.88)^{***}$
Driver highest level of education	2.45	2.48	-0.030
			(0.087)
Driver experience	6.95	8.24	-1.29
			$(0.72)^{*}$
Weeks unemployed before current job	2.96	2.28	0.68
			(0.77)
Number of vehicles driven for before current	6.05	4.97	1.08
			$(0.57)^{*}$
Number of past accidents	0.90	0.87	0.035
			(0.13)
Number of months the driver has been employed	15.2	14.3	0.91
			(2.49)
Owner rating: driver's honesty	7.78	7.60	0.18
			(0.18)
Owner rating: how hard driver works	8.29	8.07	0.22
			(0.18)
Owner rating: driver's safety	8.32	8.21	0.11
			(0.18)
Owner rating: driver's performance overall	8.09	8	0.092
			(0.17)
Driver days working for owner	411.7	500.8	-89.2
			(61.2)
Driver Raven's score	4.26	4.28	-0.016
			(0.18)
Revenue at baseline	7744.8	7732.3	12.5
			(207.6)
Baseline target	3113.1	3147.6	-34.5
			(56.6)

The summary statistics are calculated using the driver baseline data. The first column shows the mean in the control group, while the second column show the mean in the treatment group. The final column shows the difference between treatment and control. 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
Route FE	X	X	X
Matatu N	187	187	187

The data are from the owner endline surveys (where 3% of the sample - 9 owners - were unreachable, balanced between treatment and control). Note that the variables in this table were added to the endline survey after the first wave of endlines had already been completed, which is why we only have 187 observations (balanced across treatment and control). Each of the variables is a binary indicator for whether the owner said he knew the exact number of kilometers, instances of off-route driving and revenue generated by the vehicle on the most recent day the vehicle was on road. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels. Standard errors are clustered at the matatu level.

Table 4: Monitoring through the device

	(1) Difficulty Monitor	(2) Monitoring Time	(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
Route FE	X	X	X	X	X
Matatu N	190	190	190	190	190

The data are from the owner endline surveys (where 3% of the sample - 9 owners - were unreachable, balanced between treatment and control). Note that the variables in this table were added to the endline survey after the first wave of endlines had already been completed, which is why we only have 190 observations (balanced across treatment and control). 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). Monitoring time captures whether owners are spending less time monitoring ($= -1$), more time monitoring ($= 1$), or have seen no change ($= 0$) over the last 6 months. The last three columns document the number of times the owner checked up on the driver by phone, at the terminal (stage), and through a third party, respectively. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels. Standard errors are clustered at the matatu level.

	(1) 500	(2) 600	(3) 700	(4) 800	(5) 900	(6) 1000	(7) 1100
Treatment	−63.658**	−79.026**	−80.957*	−97.075**	−93.110*	−81.483*	−81.673*
	(32.366)	(36.311)	(44.088)	(46.778)	(47.664)	(45.250)	(46.821)
Control Mean of DV	767.1	751.3	700.8	750.0	620.3	554.4	751.3
Controls	X	X	X	X	X	X	X
Day FE	X	X	X	X	X	X	X
Route FE	X	X	X	X	X	X	X
Observations	3,378	3,822	4,503	5,339	5,866	6,820	7,101

The data are from the daily surveys we collected from owners and drivers. Under-reporting is the amount of revenue the driver withholds from the owner. According to the contract, the driver must deliver the target to the owner by the end of the day. On days where the driver does not make the target (owner's reported income is below the target), the driver should deliver everything they earned in fares to the owner. On these days, any take-home pay the driver reports to us via the SMS survey is the amount they under-report to the owner. On days when the driver makes the target (owner's reported income is equal to the target), under-reporting is set to zero because they made the target. The model predicts that drivers will continue to under-report until revenue exceeds some threshold q^* of revenue above the target. We cannot determine the exact q^* because we don't observe α for each owner. Therefore we run a regression of the under-reported amount on treatment, and a set of controls, for various possible values of q^*, ranging from 500 to 1000 KES above the target. Note, to get an accurate measure of under-reporting 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 we define q^* to be net revenue above target, defined as owner income + driver salary − target. These regressions are not split out by month and we restrict data to include months 2 onwards (in month 1 there is no behavior change as owners are learning how to use the technology. Note the effects stay the same if month 1 is included, but we lose some precision). Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels. Standard errors are clustered at the matatu level.

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
Route FE	X	X	X	X
Matatu N	244	190	190	246

The data are from the owner endline surveys (where 3% of the sample - 9 owners - were unreachable, balanced between treatment and control). The first column represents the amount of KES that was transferred from the owner to the driver in a game of trust. The owner was given an envelope with 900 KES and told that anything they placed back in the envelop would be tripled and sent to the driver. The driver would then choose how much to send back to the owner. The following two columns ask owners whether their drivers' driving has improved, and whether they have become more honest ($= +1$), or less honest ($= -1$) in the last 6 months. Note that the variables in Column 2 and 3 were added to the endline survey after the first wave of endlines had already been completed, which is why we only have 190 observations (balanced across treatment and control). The final column reflects owners rating of drivers on a scale from 1 to 10 (where 10 is the highest score). Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels. Standard errors are clustered at the matatu level.

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
Route FE	X	X
Matatu N	246	240

The data are from the owner endline surveys (where 3% of the sample - 9 owners- were unreachable, balanced between treatment and control). The first column represents the number of vehicles the owner possesses, while the second column is an indicator $= 1$ if the owner refurbished the interior of his vehicle. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels. Standard errors are clustered at the matatu level.

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	-0.239**	-0.131*	-0.009	0.041
	(0.247)	(0.874)	(0.108)	(0.074)	(0.015)	(0.035)
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	0.69	0.77	0.08	0.40
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-Day N	39,072	39,072	39,072	39,072	39,072	39,072

The data are from the tracking device throughout the study period. Column (1) and (2) capture average and maximum speeds throughout the day, respectively. Column (3)-(6) capture daily alerts for speeding over 80 km/hour, sharp breaking, over-acceleration, and sharp turning, respectively. We control for the number of miles the vehicle was on the road. Cash Treatment = 1 if the driver received cash transfers (whether it be for one month or two). Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels. Standard errors are clustered at the matatu level.

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	-0.220**	-0.140**	-0.014	-0.000
	(0.220)	(0.756)	(0.096)	(0.058)	(0.013)	(0.031)
One Month Post Treat	-0.039	0.072	-0.058	-0.115	-0.003	-0.017
	(0.260)	(0.971)	(0.135)	(0.089)	(0.012)	(0.031)
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	0.69	0.77	0.08	0.40
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-Day N	42,405	42,405	42,405	42,405	42,405	42,405

The data are from the tracking device throughout the study period. Column (1) and (2) capture average and maximum speeds throughout the day, respectively. Column (3)-(6) capture daily alerts for speeding over 80 km/hour, sharp breaking, over-acceleration, and sharp turning, respectively. We control for the number of miles the vehicle was on the road. Cash Treatment = 1 if the driver received cash transfers (whether it be one month or two). One Month Post Cash = 1 for drivers in the 1 month cash treatment group, in the month *after* their cash incentives were stopped. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels. Standard errors are clustered at the matatu level.

Figures

Figure 1: Mobile app "SmartMatatu"

(a) Map Viewer

(b) Historical Map Viewer

(c) Safety Feed

(d) Productivity Summary


(e) Report Submit

(f) Report Complete

Figure 2: Device Location

(a) Designated bus routes in Nairobi (black)

(b) Designated bus routes in Nairobi (black) and routes in our sample (colored)

Figure 3: Study Timeline

Notes: The figure depicts the number of matatus that were fitted with GPS trackers (and hence were added to the study) per week. The first installation took place in November 2016, and continued until April 2017. On average, the field team was able to fit GPS trackers to 15 matatus per week. As a result it took approximately 5 months to finish installations.

Figure 4: Device Location
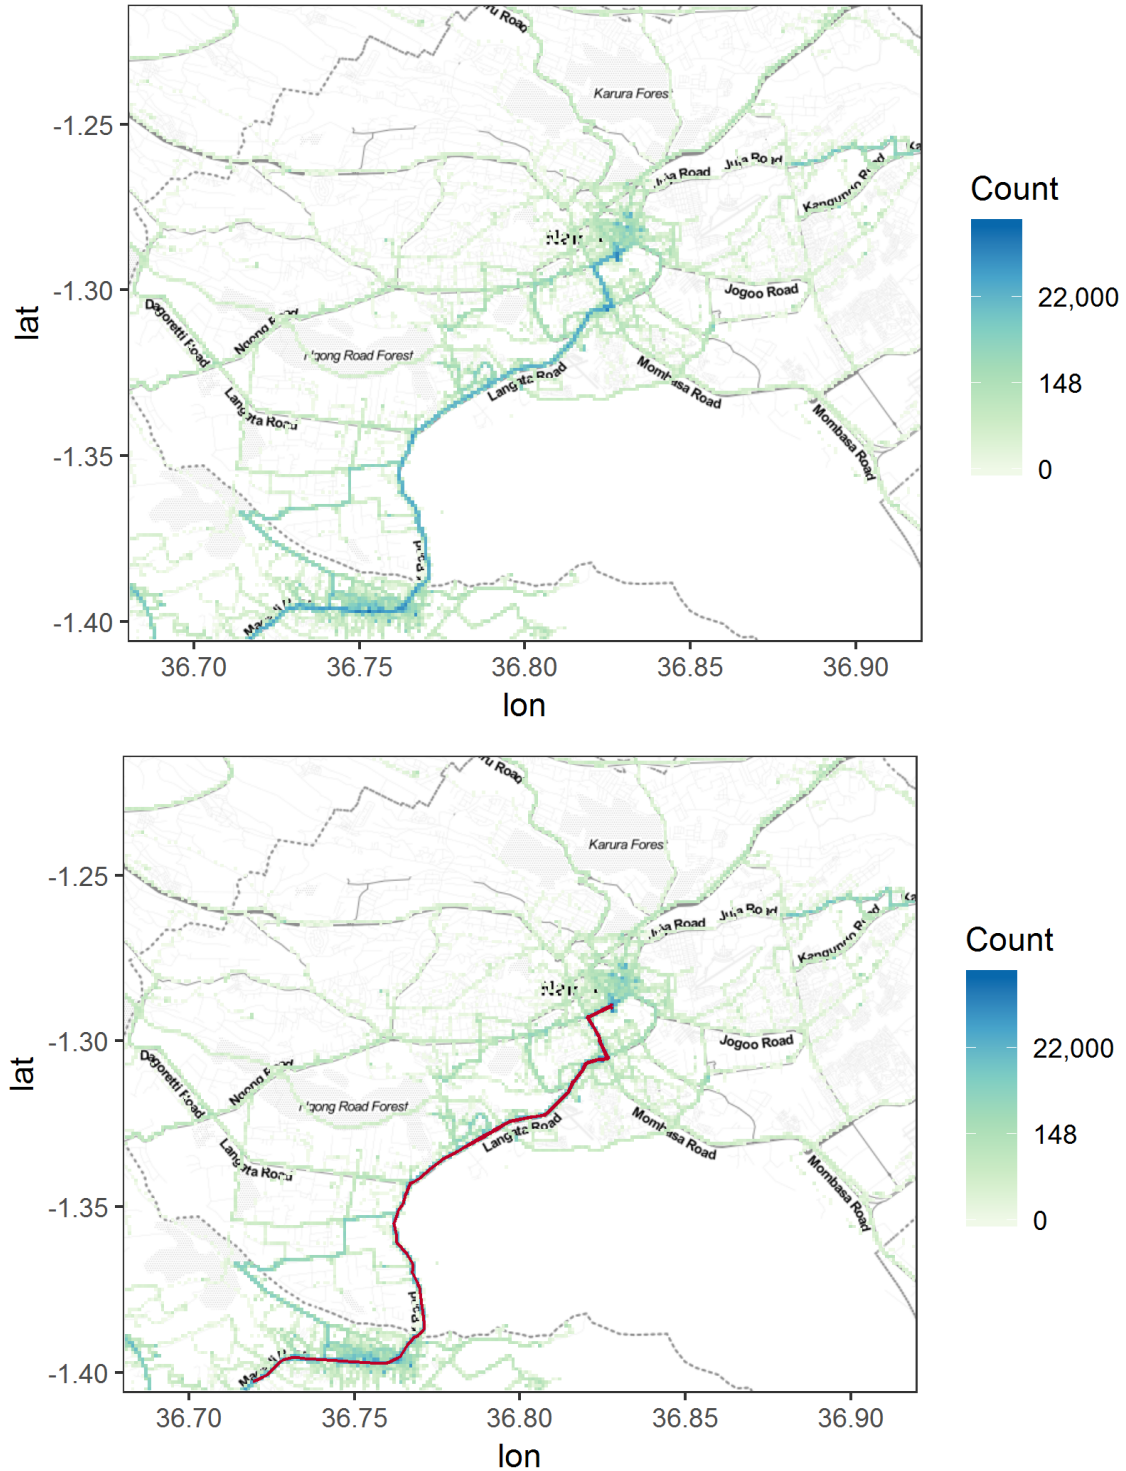
Notes: These maps use data from the trackers that were installed in vehicles licensed to operate on Route 126 (Ongata-Rongai line). Specifically, we count the number of times that vehicles passed through particular longitudinal and latitudinal cells on the map. A deeper shade of blue demonstrates that more vehicles passed through that particular cell. The second panel overlays the designated route that vehicles are supposed to be on (red). Any colored cells outside of the designated route are instances of off-route driving. Some of these are sanctioned by the owner, while others are not.

Figure 5: Model Validation (Constant Shading)

Notes: This figure depicts the amount of under-reporting on the y-axis, and the amount of revenue relative to the target on the x-axis. The data are from the control group. Under-reporting is the amount of revenue the driver withholds from the owner. According to the contract, the driver must deliver the target to the owner by the end of the day. On days where the driver does not make the target (owner's reported income is below the target), they should deliver everything they earned in fares to the owner. On these days, any take-home pay the driver reports to us via the SMS survey, is the amount they under-report to the owner. On days when the driver makes the target (owner's reported income is equal to the target), under-reporting is set to zero because they made the target. Note, to get an accurate measure of under-reporting 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.

Figure 6: Device Usage (API Calls)


(a) Weekly Usage

(b) Daily Usage

Notes: To measure device usage, we capture whether any API calls were made in a day. An API call is generated each time the owner logs into the app throughout the study period. The first panel looks at usage by week, whereas the bottom panel looks at usage per day.

Figure 7: Reprimands and Firing

Notes: Each point on the graphs is the average difference between treatment and control in a particular month for a particular outcome. Note that because installations were rolled out over time, Month 1 (through 6) represent the first month (through the sixth month) since *installation*. The outcome in the first panel is number of instances the owner criticized the driver (driver-reported). This data was collected in the monthly driver surveys. The outcome in the second panel is the number of drivers fired. This data was captured by the daily owner/driver surveys, and then validated by an enumerator who called the owner directly to confirm.

Figure 8: Prediction 1 → Target

Notes: Each point on the graph is the difference between the average daily target reported by treatment and control owners in a particular month. The data was collected through the owner daily surveys. Note that because installations were rolled out over time, Month 1 (through 6) represent the first month (through the sixth month) since *installation*.

Figure 9: Prediction 2 → Damaging (Risky) Driving

Control Mean: 483.48
Last 3 months pooled: -187.62** (89.68)

Control Mean: 0.17
Last 3 months pooled: -0.06* (0.03)

Notes: Each point on the graphs is the average difference between treatment and control in a particular month for a particular outcome. Note that because installations were rolled out over time, Month 1 (through 6) represent the first month (through the sixth month) since *installation*. The outcome in the first panel is the daily repair cost reported by the owner (daily survey). The outcome in the second panel is a binary indicator $= 1$ if the owner's reported repair was "large" (i.e. in the 80th percentile or above).

Figure 10: Prediction 2 → Damaging (Risky) Driving

Control Mean: 815.475
Last 3 months pooled: -371.62** (154.251)

Notes: Each point on the graph is the difference between the average distance to designated route calculated for treatment and control owners in a particular month. "Distance to designated route" is the shortest distance between the GPS tracker data point and the line corresponding to the route the vehicle is supposed to be on. A negative coefficient means the vehicles in the treatment group are closer to the designated route than the vehicles in the control group. Note that because installations were rolled out over time, Month 1 (through 6) represent the first month (through the sixth month) since *installation*.

Figure 11: Prediction 2 → Damaging (Risky) Driving

Notes: We plot the distributions of vertical and lateral accelerations for treatment and control vehicles from month 2 onwards (for consistency with the other pooled regressions, but the results are the same if month 1 is included). These acceleration measures are taken from the device directly. The distributions for vertical and lateral acceleration are centered at -200 and 100, respectively, rather than 0 because of some combination of a non exact calibration and the asymmetry of suspension resulting in asymmetrical acceleration.

Figure 12: Prediction 3 → Effort

Control Mean: 14.79
Last 3 months pooled: 1.12* (0.64)

Control Mean: 96.64
Last 3 months pooled: 8.37 (5.79)

Notes: Each point on the graphs is the average difference between treatment and control in a particular month for a particular outcome. Note that because installations were rolled out over time, Month 1 (through 6) represent the first month (through the sixth month) since *installation*. The outcome in the first panel is number of hours the device was on. The device powers on and off with the vehicle, and thus provides a measure of effort. The outcome in the second panel is the number of kilometers driven. These data points were captured daily by the device.

54

Figure 13: Prediction 5 → Achieving Target

Notes: The outcome in these figures is whether or not the driver made the target (owner income = target). The data is collected from owner daily surveys. In the first panel, each point on the graph is the average difference in the probability of making the target between treatment and control in a particular month. Note that because installations were rolled out over time, Month 1 (through 6) represent the first month (through the sixth month) since *installation*. In the second panel we pool the data from months 2 onwards and look at whether the probability of making the target differs between treatment and control for a particular amount of revenue. The x-axis here depicts revenue relative to an average revenue day on that route (normalized by the target). Note we use gross revenue for this outcome instead of net revenue like we did for the under-reporting amount - because the variable only depends on drivers reporting, which means we have more data to work with.

Figure 14: Prediction 4 → Less Under-reporting
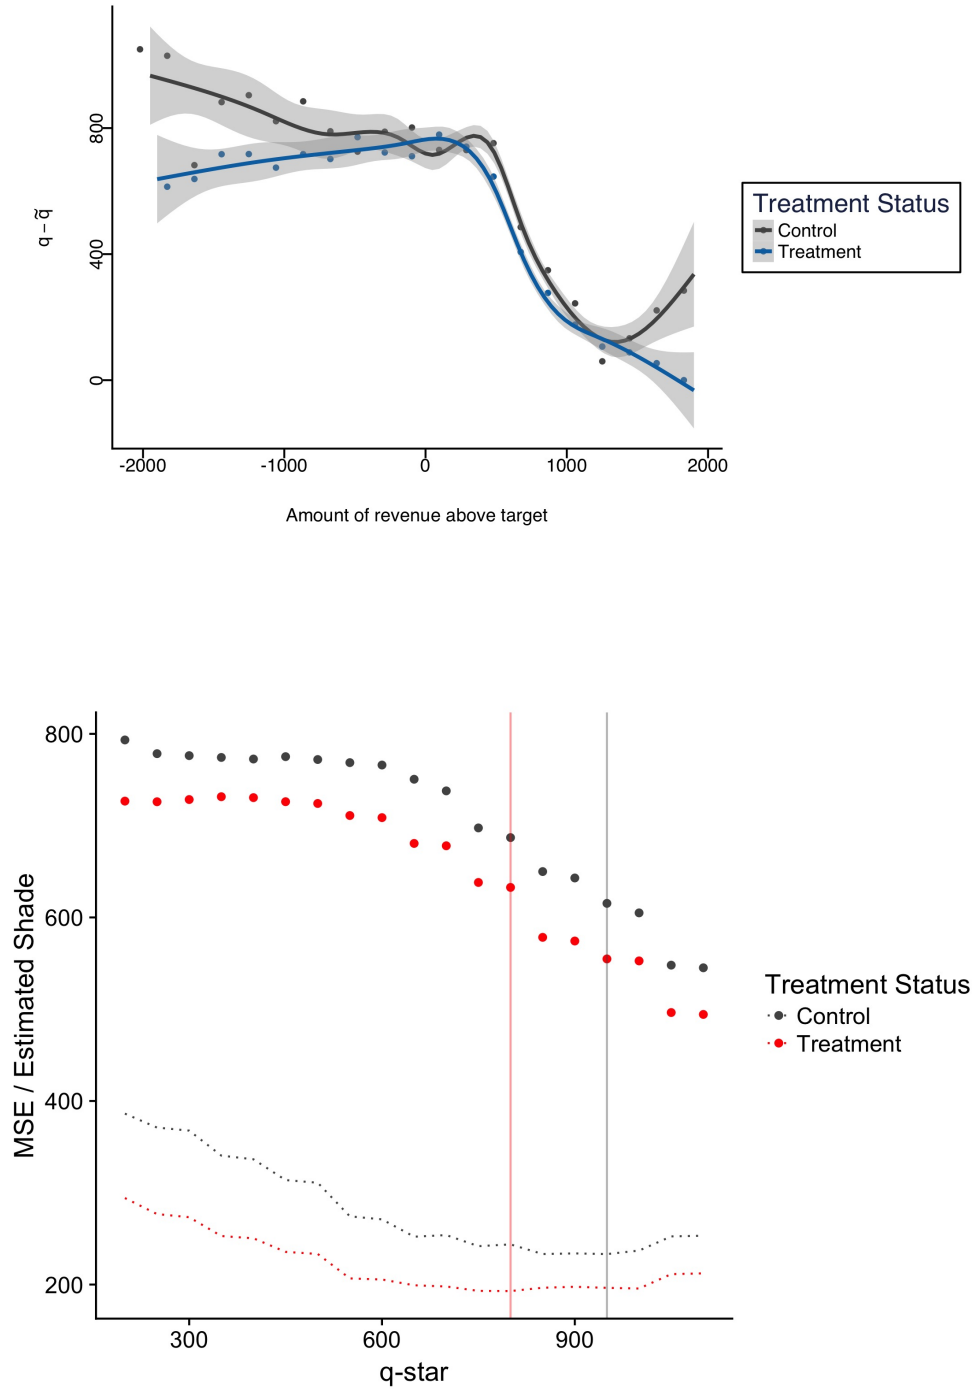
Notes: The outcome in these figures is the amount of revenue drivers under-report. The first panel reproduces Figure 5, but separates treatment drivers from control drivers. We also overlay a non-parametric smoothing function. The second panel imposes the model's step function and computes the average under-reported amount for different levels of q^* for treatment (red dots) and control (black dots). The dotted lines on the bottom represent the MSE from each regression.

Figure 15: Under-report

Control Mean: 614.09
Last 3 months pooled: -88.40* (53.00)

Notes: Each point on the graph is the difference between the average amount of revenue treatment and control drivers under-report in a particular month. The data was collected from owner and driver surveys - see Figure 5 for a description of how under-reported revenue was computed. Note that because installations were rolled out over time, Month 1 (through 6) represent the first month (through the sixth month) since *installation*.

Figure 16: Company Outcomes

Control Mean: 7126.94
Last 3 months pooled: -6.99 (177.28)

Control Mean: 3275.61
Last 3 months pooled: 361.70* (196.33)

Notes: Each point on the graphs is the average difference between treatment and control in a particular month for a particular outcome. Note that because installations were rolled out over time, Month 1 (through 6) represent the first month (through the sixth month) since *installation*. The outcome in the first panel is the amount of revenue collected by the driver (captured through daily driver surveys). The outcome in the second panel is amount of profits the business generates, where $profits = revenue - costs - driver\,salary$ - which are captured from owner and driver daily reports.

Figure 17: Salary per hour

Control Mean: 59.91
Last 3 months pooled: -0.09 (2.61)

Notes: Each point on the graph is the difference in driver salary per hour between treatment and control drivers. The data was collected from daily driver surveys. Note that because installations were rolled out over time, Month 1 (through 6) represent the first month (through the sixth month) since *installation*.

Figure 18: Speeding

Notes: Each point on the graphs is the average difference between treatment and control in a particular month for a particular outcome. Note that because installations were rolled out over time, Month 1 (through 6) represent the first month (through the sixth month) since *installation*. The outcome in the first panel is maxspeed, while the outcome in the second panel is average speed. Both are measured directly by the tracker.

Figure 19: Over-acceleration and Over-speeding

Control Mean: 0.08
Last 3 months pooled: 0.00 (0.02)

Control Mean: 0.70
Last 3 months pooled: 0.00 (0.17)

Notes: Each point on the graphs is the average difference between treatment and control in a particular month for a particular outcome. Note that because installations were rolled out over time, Month 1 (through 6) represent the first month (through the sixth month) since *installation*. The outcome in the first panel is over-acceleration alerts, while the outcome in the second panel is over-speeding alerts. Both are measured directly by the tracker

61

Figure 20: Sharp-turning and Sharp-braking

Control Mean: 0.38
Last 3 months pooled: 0.08 (0.07)

Control Mean: 0.79
Last 3 months pooled: 0.17 (0.15)

Notes: Each point on the graphs is the average difference between treatment and control in a particular month for a particular outcome. Note that because installations were rolled out over time, Month 1 (through 6) represent the first month (through the sixth month) since *installation*. The outcome in the first panel is sharp-turning alerts, while the outcome in the second panel is sharp braking alerts. Both are measured directly by the tracker.

Figure 21: Accidents

Control Mean: 0.02
Last 3 months pooled: 0.02 (0.01)

Notes: Each point on the graph is the difference in the number of accidents treatment and control drivers get into. The data was collected from daily owner/driver daily surveys, and then validated by an enumerator who called the owner directly. Note that because installations were rolled out over time, Month 1 (through 6) represent the first month (through the sixth month) since *installation*.

Appendix 1: Tables

Table 10: Submitting daily survey report

	(1) Owner report submitted	(2) Driver report submitted
Info treatment group	0.042	0.029
	(0.040)	(0.037)
Control Mean of DV	0.45	0.55
Day FE	Y	Y
Route FE	Y	Y
Matatu N	255	255
Matatu-Day N	46,920	46,920

The data are from the owner and drivers that submitted data throughout the study period. The dependent variable is a binary indicator for whether the owner (Column 1) or driver (Column 2) submitted a report that day. Standard errors are clustered at the village level. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels.

Table 11: Differential Revenue Reporting

	(1) Revenue	(2) Revenue
Mileage	8.848***	
	(1.089)	
Info treatment group	-259.461	
	(267.894)	
Mileage X Treat	1.241	
	(1.491)	
Mileage quartile 1		-1489.213***
		(209.834)
Mileage quartile 2		-1248.806***
		(203.465)
Mileage quartile 3		-552.104***
		(180.865)
Mileage quartile 1 X Treat		-331.924
		(266.860)
Mileage quartile 2 X Treat		30.419
		(181.920)
Mileage quartile 3 X Treat		-119.329
		(204.766)
Mileage quartile 4 X Treat		38.374
		(235.454)
Control Mean of DV	7126.94	7126.94
Day FE	Y	Y
Route FE	Y	Y
Matatu N	241	241
Matatu-Day N	22,436	23,514

The data are from days where we have both the driver reports (reported revenue), and the tracking data (number of miles). The dependent variable is the amount of revenue the driver reports. In the first column we regress revenue on the number of miles the vehicle travelled, and indicator for treatment, and an interaction between the two terms. The interaction captures the differential relationship between mileage and reported revenue between treatment and control. We might be concerned that drivers are reporting differently in the treatment than the control group. The coefficient on the interaction term is not significant however. Column 2 further investigates whether there is differential reporting across different quartiles of the mileage distribution. Again, the interaction terms are all insignificant. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels. Note there are 14 drivers that did not report revenue at all during the 6 months of the study.

Table 12: Contract

	(1) Daily Target	(2) Daily Target	(3) Daily Target
Treat X Month1	-35.141		
	(68.842)		
Treat X Month2	-46.502		
	(91.193)		
Treat X Month3	-63.407		
	(86.703)		
Treat X Month4	-94.940		
	(83.842)		
Treat X Month5	-120.861		
	(92.118)		
Treat X Month6	-134.693		
	(97.085)		
Treat X First 3 months		-51.322	
		(79.284)	
Treat X Last 3 months		-112.564	
		(86.747)	
Treat X Trend			-9.332
			(17.225)
Control Mean of DV	3057.38	3057.38	3057.38
Controls	Y	Y	Y
Day FE	Y	Y	N
Route FE	Y	Y	Y
Matatu N	237	237	237
Matatu-Day N	15,884	15,884	15,542

This table shows the results of three different regression specifications. The first column shows the treatment effects broken out by treatment month, as specified in the main body of the paper (through figures). The second column pools the first three months and the last three months - to show the effects in the "early" versus "later" months of the experiment. The final column demonstrates the results of a regression of the outcome of interest on a linear time trend, treatment, and the interaction of the two. The interaction term reveals how treatment trended differently over time relative to the control group. The outcome of interest in this table is the target. Note there are 237 matatus included in these regressions because 18 owners failed to report the target throughout the study (balanced between treatment and control). Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels. Standard errors are clustered at the matatu level.

Table 13: Deviceon

	(1) Hours tracking device on	(2) Hours tracking device on	(3) Hours tracking device on
Treat X Month1	-0.842		
	(0.658)		
Treat X Month2	0.623		
	(0.590)		
Treat X Month3	0.986*		
	(0.554)		
Treat X Month4	0.691		
	(0.640)		
Treat X Month5	1.468**		
	(0.715)		
Treat X Month6	1.474*		
	(0.757)		
Treat X First 3 months		0.388	
		(0.503)	
Treat X Last 3 months		1.119*	
		(0.637)	
Treat X Trend			0.177
			(0.207)
Control Mean of DV	14.79	14.79	14.79
Controls	Y	Y	Y
Day FE	Y	Y	N
Route FE	Y	Y	Y
Matatu N	254	254	254
Matatu-Day N	45,654	45,654	44,444

This table shows the results of three different regression specifications. The first column shows the treatment effects broken out by treatment month, as specified in the main body of the paper (through figures). The second column pools the first three months and the last three months - to show the effects in the "early" versus "later" months of the experiment. The final column demonstrates the results of a regression of the outcome of interest on a linear time trend, treatment, and the interaction of the two. The interaction term reveals how treatment trended differently over time relative to the control group. The outcome of interest in this table is the number of hours the device was on (as a proxy for effort). Note there are 254 matatus included in these regressions because 1 device was faulty (the matatu was in an accident). Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels. Standard errors are clustered at the matatu level.

Table 14: Mileage

	(1)	(2)	(3)
	Kilometers	Kilometers	Kilometers
Treat X Month1	-4.302		
	(5.966)		
Treat X Month2	2.053		
	(5.601)		
Treat X Month3	7.802		
	(5.318)		
Treat X Month4	4.411		
	(5.700)		
Treat X Month5	9.640		
	(6.386)		
Treat X Month6	13.132*		
	(6.917)		
Treat X First 3 months		2.911	
		(5.095)	
Treat X Last 3 months		8.365	
		(5.793)	
Treat X Trend			1.533
			(1.650)
Control Mean of DV	96.64	96.64	96.64
Controls	Y	Y	Y
Day FE	Y	Y	N
Route FE	Y	Y	Y
Matatu N	254	254	254
Matatu-Day N	45,654	45,654	44,444

This table shows the results of three different regression specifications. The first column shows the treatment effects broken out by treatment month, as specified in the main body of the paper (through figures). The second column pools the first three months and the last three months - to show the effects in the "early" versus "later" months of the experiment. The final column demonstrates the results of a regression of the outcome of interest on a linear time trend, treatment, and the interaction of the two. The interaction term reveals how treatment trended differently over time relative to the control group. The outcome of interest in this table is the number of hours the device was on (as a proxy for effort). Note there are 254 matatus included in these regressions because 1 device was faulty (the matatu was in an accident). Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels. Standard errors are clustered at the matatu level.

Table 15: Repair Cost

	(1) Repair Cost	(2) Repair Cost	(3) Repair Cost
Treat X Month1	67.838		
	(70.046)		
Treat X Month2	-50.802		
	(73.520)		
Treat X Month3	-125.119		
	(79.387)		
Treat X Month4	-186.449**		
	(89.291)		
Treat X Month5	-187.876**		
	(94.119)		
Treat X Month6	-226.720**		
	(104.752)		
Treat X First 3 months		-46.459	
		(70.294)	
Treat X Last 3 months		-187.616**	
		(89.675)	
Treat X Trend			-40.116*
			(20.406)
Control Mean of DV	483.48	483.48	483.48
Controls	Y	Y	Y
Day FE	Y	Y	N
Route FE	Y	Y	Y
Matatu N	238	238	238
Matatu-Day N	15,881	15,881	15,539

This table shows the results of three different regression specifications. The first column shows the treatment effects broken out by treatment month, as specified in the main body of the paper (through figures). The second column pools the first three months and the last three months - to show the effects in the "early" versus "later" months of the experiment. The final column demonstrates the results of a regression of the outcome of interest on a linear time trend, treatment, and the interaction of the two. The interaction term reveals how treatment trended differently over time relative to the control group. The outcome of interest in this table is the amount of repairs the owner incurred. Note there are 238 matatus included in these regressions because 17 owners failed to report the target throughout the study (balanced between treatment and control).Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels. Standard errors are clustered at the matatu level.

Table 16: Repair Cost (Binary)

	(1) Large Repairs	(2) Large Repairs	(3) Large Repairs
Treat X Month1	0.033		
	(0.030)		
Treat X Month2	-0.019		
	(0.030)		
Treat X Month3	-0.030		
	(0.032)		
Treat X Month4	-0.048		
	(0.033)		
Treat X Month5	-0.066**		
	(0.033)		
Treat X Month6	-0.081**		
	(0.038)		
Treat X First 3 months		-0.010	
		(0.029)	
Treat X Last 3 months		-0.060*	
		(0.032)	
Treat X Trend			-0.009
			(0.006)
Control Mean of DV	0.17	0.17	0.17
Controls	Y	Y	Y
Day FE	Y	Y	N
Route FE	Y	Y	Y
Matatu N	238	238	238
Matatu-Day N	15,881	15,881	15,539

This table shows the results of three different regression specifications. The first column shows the treatment effects broken out by treatment month, as specified in the main body of the paper (through figures). The second column pools the first three months and the last three months - to show the effects in the "early" versus "later" months of the experiment. The final column demonstrates the results of a regression of the outcome of interest on a linear time trend, treatment, and the interaction of the two. The interaction term reveals how treatment trended differently over time relative to the control group. The outcome of interest in this table is the number of large repairs the owner incurred. Note there are 238 matatus included in these regressions because 18 owners failed to report the target throughout the study (balanced between treatment and control).Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels. Standard errors are clustered at the matatu level.

Table 17: Made Target

	(1) Made Target	(2) Made Target	(3) Made Target
Treat X Month1	-0.099**		
	(0.045)		
Treat X Month2	-0.013		
	(0.046)		
Treat X Month3	0.068		
	(0.048)		
Treat X Month4	0.125**		
	(0.051)		
Treat X Month5	0.083		
	(0.053)		
Treat X Month6	0.059		
	(0.053)		
Treat X First 3 months		-0.010	
		(0.043)	
Treat X Last 3 months		0.084*	
		(0.048)	
Treat X Trend			0.028**
			(0.011)
Control Mean of DV	0.43	0.43	0.43
Controls	Y	Y	Y
Day FE	Y	Y	N
Route FE	Y	Y	Y
Matatu N	237	237	237
Matatu-Day N	15,888	15,888	15,546

This table shows the results of three different regression specifications. The first column shows the treatment effects broken out by treatment month, as specified in the main body of the paper (through figures). The second column pools the first three months and the last three months - to show the effects in the "early" versus "later" months of the experiment. The final column demonstrates the results of a regression of the outcome of interest on a linear time trend, treatment, and the interaction of the two. The interaction term reveals how treatment trended differently over time relative to the control group. The outcome of interest in this table is whether or not the driver made the target. Note there are 237 matatus included in these regressions because 17 owners failed to report the target throughout the study (balanced between treatment and control).Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels. Standard errors are clustered at the matatu level.

Table 18: Under-report

	(1) shade	(2) shade	(3) shade
Treat X Month1	-6.782		
	(51.685)		
Treat X Month2	-89.596*		
	(48.896)		
Treat X Month3	-103.442**		
	(52.358)		
Treat X Month4	-105.062**		
	(53.112)		
Treat X Month5	-66.293		
	(59.854)		
Treat X Month6	-106.336*		
	(61.400)		
Treat X First 3 months		-69.041	
		(47.008)	
Treat X Last 3 months		-88.403*	
		(52.955)	
Treat X Trend			-15.213
			(12.967)
Control Mean of DV	521.20	521.20	521.20
Controls	Y	Y	Y
Day FE	Y	Y	N
Route FE	Y	Y	Y
Matatu N	215	215	215
Matatu-Day N	7,426	7,426	7,320

This table shows the results of three different regression specifications. The first column shows the treatment effects broken out by treatment month, as specified in the main body of the paper (through figures). The second column pools the first three months and the last three months - to show the effects in the "early" versus "later" months of the experiment. The final column demonstrates the results of a regression of the outcome of interest on a linear time trend, treatment, and the interaction of the two. The interaction term reveals how treatment trended differently over time relative to the control group. The outcome of interest in this table is the amount revenue drivers under-report. Note there are 215 matatus included in these regressions because 17 owners failed to report the target throughout the study (balanced between treatment and control); and 14 drivers didn't report their revenue/salary (balanced between treatment and control) - both of which are required to compute this measure - as detailed in the main figures. We also needed both owners and drivers to report in a particular day to compute this measure. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels. Standard errors are clustered at the matatu level.

Table 19: Revenue

	(1) Revenue	(2) Revenue	(3) Revenue
Treat X Month1	-338.338		
	(229.414)		
Treat X Month2	-61.938		
	(187.910)		
Treat X Month3	-6.492		
	(180.413)		
Treat X Month4	124.687		
	(184.769)		
Treat X Month5	54.217		
	(193.679)		
Treat X Month6	-207.192		
	(205.937)		
Treat X First 3 months		-131.349	
		(176.217)	
Treat X Last 3 months		-6.992	
		(177.277)	
Treat X Trend			-26.140
			(58.137)
Control Mean of DV	7126.94	7126.94	7126.94
Controls	Y	Y	Y
Day FE	Y	Y	N
Route FE	Y	Y	Y
Matatu N	241	241	241
Matatu-Day N	22,436	22,436	22,107

This table shows the results of three different regression specifications. The first column shows the treatment effects broken out by treatment month, as specified in the main body of the paper (through figures). The second column pools the first three months and the last three months - to show the effects in the "early" versus "later" months of the experiment. The final column demonstrates the results of a regression of the outcome of interest on a linear time trend, treatment, and the interaction of the two. The interaction term reveals how treatment trended differently over time relative to the control group. The outcome of interest in this table is reported revenue. Note there are 241 matatus included in these regressions because 14 owners failed to report the revenue throughout the study (balanced between treatment and control).Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels. Standard errors are clustered at the matatu level.

Table 20: Profits

	(1) Gross Profits	(2) Gross Profits	(3) Gross Profits
Treat X Month1	-178.812		
	(233.078)		
Treat X Month2	55.929		
	(206.280)		
Treat X Month3	83.617		
	(217.394)		
Treat X Month4	441.369**		
	(222.362)		
Treat X Month5	451.597**		
	(212.750)		
Treat X Month6	172.690		
	(227.400)		
Treat X First 3 months		-13.173	
		(191.659)	
Treat X Last 3 months		361.696*	
		(196.326)	
Treat X Trend			51.548
			(75.291)
Control Mean of DV	3275.61	3275.61	3275.61
Controls	Y	Y	Y
Day FE	Y	Y	N
Route FE	Y	Y	Y
Matatu N	216	216	216
Matatu-Day N	10,406	10,406	10,277

This table shows the results of three different regression specifications. The first column shows the treatment effects broken out by treatment month, as specified in the main body of the paper (through figures). The second column pools the first three months and the last three months - to show the effects in the "early" versus "later" months of the experiment. The final column demonstrates the results of a regression of the outcome of interest on a linear time trend, treatment, and the interaction of the two. The interaction term reveals how treatment trended differently over time relative to the control group. The outcome of interest in this table is profits. Note there are 216 matatus included in these regressions because 18 owners failed to report the target throughout the study (balanced between treatment and control); and 14 drivers didn't report their revenue/salary (balanced between treatment and control) - both of which are required to compute this measure - as detailed in the main figures. We also needed both owners and drivers to report in a particular day to compute this measure. Asterisks indicate statistical significance at the 1% ***, 5% **, and 10% * levels. Standard errors are clustered at the matatu level.

Appendix 2: Model derivation

Step 1: Owner's punishment

We assume the owner's signal is noisy but unbiased, so that $\hat{q}$ is defined as follows:

$$\hat{q} = q - \sigma$$

$$\sigma \sim U\left(-\frac{1}{\alpha}, \frac{1}{\alpha}\right) \qquad f(\hat{q}) = \frac{1}{q + \frac{1}{\alpha} - \left(q - \frac{1}{\alpha}\right)} = \frac{\alpha}{2}$$

where α is the precision of the owners' signal about true revenue. Any monitoring technology we introduce will provide more information to the owner about driver behavior. This will increase the precision of the owner's signal about what revenue should be, which gives the driver less leeway to significantly under-report revenue on a particular day. In the case of our monitoring technology specifically, the owner can observe the number of kilometers driven, and where the driver was operating at any point in time. Owners can use this information to estimate the number of trips to and from the city center, which provides a more accurate measure of total daily revenue.[21] With $\hat{q}$ defined in this way, the owners can be sure that real revenue q falls in the interval:

$$\left[\hat{q} - \frac{1}{\alpha} \ , \ \hat{q} + \frac{1}{\alpha}\right]$$

On days when the driver reports making the target, the owner receives the target amount and does not punish the driver. On days when the driver does not report making the target, the owner will punish them if they can be sure the driver is lying to them. In other words, they will punish if the reported revenue comes in below the possible range for q. This assumes the owner is really unwilling to punish a driver incorrectly, which makes sense because firing costs are high in this setting. The actual punishment applied is some function of the difference between this lower bound, $\hat{q} - \frac{1}{\alpha}$, and the reported amount $\tilde{q}$ (owners are less upset on days where the driver reports below the target and they know for a fact that conditions were difficult). We assume for simplicity that this function is

---

[21]While they do not know the exact number of passengers that board, they know that the vehicle generally waits at the terminal until it fills up.

linear.

$$E[punishment] = E\left[\left(\hat{q} - \frac{1}{\alpha}\right) - \tilde{q} \mid \hat{q} - \frac{1}{\alpha} > \tilde{q}\right] \cdot \Pr\left(\hat{q} - \frac{1}{\alpha} > \tilde{q}\right)$$

$$= E\left[\hat{q} - \tilde{q} - \frac{1}{\alpha} \mid \hat{q} > \tilde{q} + \frac{1}{\alpha}\right] \cdot \Pr\left(\hat{q} > \tilde{q} + \frac{1}{\alpha}\right)$$

$$= \int_{\tilde{q}+\frac{1}{\alpha}}^{q+\frac{1}{\alpha}} \left(\hat{q} - \tilde{q} - \frac{1}{\alpha}\right) \cdot f(\hat{q}) d\hat{q}$$

$$= \frac{\alpha}{2} \int_{\tilde{q}+\frac{1}{\alpha}}^{q+\frac{1}{\alpha}} \left(\hat{q} - \tilde{q} - \frac{1}{\alpha}\right) d\hat{q}$$

$$= \frac{\alpha}{2} \left[\frac{\hat{q}^2}{2} - \left(\tilde{q} + \frac{1}{\alpha}\right) \hat{q} \Big|_{\tilde{q}+\frac{1}{\alpha}}^{q+\frac{1}{\alpha}}\right]$$

$$= \frac{\alpha}{2} \left[\frac{1}{2}\left(q + \frac{1}{\alpha}\right)^2 - \left(\tilde{q} + \frac{1}{\alpha}\right)\left(q + \frac{1}{\alpha}\right) - \frac{1}{2}\left(\tilde{q} + \frac{1}{\alpha}\right)^2 + \left(\tilde{q} + \frac{1}{\alpha}\right)^2\right]$$

$$= \frac{\alpha}{2} \left[\frac{1}{2}\left(q + \frac{1}{\alpha}\right)^2 - \left(\tilde{q} + \frac{1}{\alpha}\right)\left(q + \frac{1}{\alpha}\right) + \frac{1}{2}\left(\tilde{q} + \frac{1}{\alpha}\right)^2\right]$$

$$= \frac{\alpha}{4} \left[\left(q + \frac{1}{\alpha}\right)^2 - 2\left(\tilde{q} + \frac{1}{\alpha}\right)\left(q + \frac{1}{\alpha}\right) + \left(\tilde{q} + \frac{1}{\alpha}\right)^2\right]$$

$$= \frac{\alpha}{4} \left[\left(q + \frac{1}{\alpha}\right)^2 - \left(\tilde{q} + \frac{1}{\alpha}\right)^2\right]$$

$$= \frac{\alpha}{4} (q - \tilde{q})^2$$

Step 2: Solve the agent's optimal shading amount

Below we provide the full derivation of the optimal shading amount

$$\frac{\partial U^D}{\partial \tilde{q}} = -1 + \frac{\alpha}{2}(q - \tilde{q}) = 0$$

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

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

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

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

Step 3: Switch point

Below we provide the full derivation of the switch point

$$q - T - \beta r = (q - \tilde{q}) - \frac{\alpha}{4}(q - \tilde{q})^2 - \beta r$$

$$q - T = \left(q - q + \frac{2}{\alpha}\right) - \frac{\alpha}{4}\left(q - q + \frac{2}{\alpha}\right)^2$$

$$q - T = \frac{2}{\alpha} - \frac{\alpha}{4}\left(\frac{4}{\alpha^2}\right)$$

$$q - T = \frac{1}{\alpha}$$

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

Step 4: Driver's optimal choice of effort and risk

The driver chooses effort to maximize his utility

$$\max_{e,r} \quad E[(q - T - \beta r) \mid q \geq q^*] \cdot Pr(q \geq q^*) + E[(q - \tilde{q}) - \frac{\alpha}{4}(q - \tilde{q})^2 - \beta r \mid q < q^*] \cdot Pr(q < q^*) - h(e,r)$$

Simplifying

$$\max_{e,r} \quad E\left[(q - T - \beta r) \mid q \geq q^*\right](1 - F(q^*)) + E\left[\frac{1}{\alpha} - \beta r \mid q < q^*\right] \cdot F(q^*) - h(e,r)$$

Expressing in terms of exogenous variable

$$\max_{e,r} \quad E\left[(e + r\varepsilon - T - \beta r) \mid \varepsilon \geq \frac{q^* - e}{r}\right] \cdot \left(1 - F_\varepsilon\left(\frac{q^* - e}{r}\right)\right) + E\left[\frac{1}{\alpha} - \beta r \mid \varepsilon < \frac{q^* - e}{r}\right] \cdot F_\varepsilon\left(\frac{q^* - e}{r}\right) - h(e,r)$$

Using integral notation:

$$L = \int_{\frac{q^* - e}{r}}^{\infty} (e + r\varepsilon - T - \beta r) f_\varepsilon(\varepsilon) d\varepsilon + \int_0^{\frac{q^* - e}{r}} \left(\frac{1}{\alpha} - \beta r\right) f_\varepsilon(\varepsilon) d\varepsilon - h(e,r)$$

Taking the derivative with respect to e

$$\frac{\partial L}{\partial e} = \int_{\frac{q^* - e}{r}}^{\infty} 1 \cdot f_\varepsilon(\varepsilon) d\varepsilon + \frac{1}{r}(q^* - T - \beta r) f_\varepsilon\left(\frac{q^* - e}{r}\right) + 0 - \frac{1}{r}\left(\frac{1}{\alpha} - \beta r\right) f_\varepsilon\left(\frac{q^* - e}{r}\right) - h'_e$$

$$= \int_{\frac{q^* - e}{r}}^{\infty} f_\varepsilon(\varepsilon) d\varepsilon - h'_e$$

$$\rightarrow \underbrace{1 - F_\varepsilon\left(\frac{q^* - e}{r}\right) - h'_e = 0}_{F.O.C}$$

Taking the derivative with respect to r

$$\frac{\partial L}{\partial r} = \int_{\frac{q^*-e}{r}}^{\infty} (\varepsilon - \beta) \cdot f_\varepsilon(\varepsilon) d\varepsilon + \left(\frac{q^*-e}{r^2}\right) \cancel{(q^* - T - \beta r) f_\varepsilon \left(\frac{q^*-e}{r}\right)} +$$

$$\int_0^{\frac{q^*-e}{r}} (-\beta) \cdot f_\varepsilon(\varepsilon) d\varepsilon - \cancel{\left(\frac{q^*-e}{r^2}\right)\left(\frac{1}{\alpha} - \beta r\right) f_\varepsilon \left(\frac{q^*-e}{r}\right)} - h'_r$$

$$= \int_{\frac{q^*-e}{r}}^{\infty} \varepsilon f_\varepsilon(\varepsilon) d\varepsilon - h'_r - \beta \left(\int_{\frac{q^*-e}{r}}^{\infty} f_\varepsilon(\varepsilon) d\varepsilon + \int_0^{\frac{q^*-e}{r}} f_\varepsilon(\varepsilon) d\varepsilon \right)$$

$$\to \underbrace{\int_{\frac{q^*-e}{r}}^{\infty} \varepsilon f_\varepsilon(\varepsilon) d\varepsilon - h'_r - \beta = 0}_{F.O.C}$$

Next we investigate how a change in T affects effort and risk:

$$\begin{bmatrix} \frac{\partial e}{\partial T} \\ \frac{\partial r}{\partial T} \end{bmatrix} = - \begin{bmatrix} \frac{\partial^2 L}{\partial e^2} & \frac{\partial L}{\partial r \partial e} \\ \frac{\partial L}{\partial e \partial r} & \frac{\partial^2 L}{\partial r^2} \end{bmatrix}^{-1} \begin{bmatrix} \frac{\partial L}{\partial T \partial e} \\ \frac{\partial L}{\partial T \partial r} \end{bmatrix}$$

$$= - \frac{1}{\underbrace{\text{Determinant}}_{S.O.C \text{ for Hessian} > 0}} \begin{bmatrix} \frac{\partial^2 L}{\partial r^2} & -\frac{\partial L}{\partial r \partial e} \\ -\frac{\partial L}{\partial e \partial r} & \frac{\partial^2 L}{\partial e^2} \end{bmatrix} \begin{bmatrix} \frac{\partial L}{\partial T \partial e} \\ \frac{\partial L}{\partial T \partial r} \end{bmatrix}$$

Taking each term in turn:

$$\frac{\partial^2 L}{\partial r^2} = 0 - \frac{\partial}{\partial r}\left(\frac{q^* - e}{r}\right)\left(\frac{q^* - e}{r}\right)f_\varepsilon\left(\frac{q^* - e}{r}\right) - h''_{rr} - 2\beta$$

$$= \left(\frac{q^* - e}{r^2}\right)\left(\frac{q^* - e}{r}\right)f_\varepsilon\left(\frac{q^* - e}{r}\right) - h''_{rr} - 2\beta$$

$$= \underbrace{\frac{1}{r}\left(\frac{q^* - e}{r}\right)^2 f_\varepsilon(\cdot) - h''_{rr}}_{S.O.C \ < \ 0}$$

$$\frac{\partial^2 L}{\partial e^2} = f_\varepsilon(\cdot)\left(\frac{1}{r}\right) - h''_{ee}$$

$$= \underbrace{\frac{1}{r}f_\varepsilon(\cdot) - h''_{ee}}_{S.O.C \ < \ 0}$$

$$\frac{\partial L}{\partial e \partial r} = 0 - \left[-\frac{1}{r}\left(\frac{q^* - e}{r}\right)f_\varepsilon\left(\frac{q^* - e}{r}\right)\right] - h''_{er}$$

$$= \underbrace{\left(\frac{q^* - e}{r^2}\right)f_\varepsilon(\cdot) - h''_{er}}_{< \ 0}$$

$$\frac{\partial L}{\partial r \partial e} = \underbrace{-f_\varepsilon(\cdot)\left(\frac{q^* - e}{r^2}\right) - h''_{er}}_{< \ 0}$$

$$\frac{\partial L}{\partial T \partial e} = -\frac{1}{r}f_\varepsilon\left(\frac{q^* - e}{r}\right)$$

$$= \underbrace{-\frac{1}{r}f_\varepsilon(\cdot)}_{< \ 0}$$

$$\frac{\partial L}{\partial T \partial r} = 0 - \frac{1}{r}\left(\frac{q^* - e}{r}\right)f_\varepsilon\left(\frac{q^* - e}{r}\right)$$

$$= \underbrace{-\left(\frac{q^* - e}{r^2}\right)f_\varepsilon(\cdot)}_{> \ 0}$$

We can sign most of these terms because of 1) second order conditions and 2) the shape of the distribution of revenue (q), which is skewed to the left, and the fact that drivers make the target 44% of the time. This means $(q^* - e < 0)$. Note, we would expect the cross partial to be negative because as the driver increases risk (fatter tails), they are less likely to make the target, which

means the returns to making the target decrease and effort will be reduced. Putting it altogether:

$$\frac{\partial e}{\partial T} = -\underbrace{\frac{1}{\text{Determinant}}}_{+}\left[\overbrace{\frac{\partial^2 L}{\partial r^2}}^{-}\cdot\overbrace{\frac{\partial L}{\partial T\partial e}}^{-} - \overbrace{\frac{\partial L}{\partial r\partial e}}^{-}\cdot\overbrace{\frac{\partial L}{\partial T\partial r}}^{+}\right]$$

$$< 0$$

$$\frac{\partial r}{\partial T} = -\underbrace{\frac{1}{\text{Determinant}}}_{+}\left[-\overbrace{\frac{\partial L}{\partial e\partial r}}^{-}\cdot\overbrace{\frac{\partial L}{\partial T\partial e}}^{-} + \overbrace{\frac{\partial^2 L}{\partial e^2}}^{-}\cdot\overbrace{\frac{\partial L}{\partial T\partial r}}^{+}\right]$$

$$> 0$$

Next we investigate how a change in α affects effort and risk:

$$\begin{bmatrix}\frac{\partial e}{\partial \alpha}\\\frac{\partial r}{\partial \alpha}\end{bmatrix} = -\begin{bmatrix}\frac{\partial^2 L}{\partial e^2} & \frac{\partial L}{\partial r\partial e}\\\frac{\partial L}{\partial e\partial r} & \frac{\partial^2 L}{\partial r^2}\end{bmatrix}^{-1}\begin{bmatrix}\frac{\partial L}{\partial \alpha\partial e}\\\frac{\partial L}{\partial \alpha\partial r}\end{bmatrix}$$

$$= -\underbrace{\frac{1}{\text{Determinant}}}_{S.O.C\text{ for Hessian} > 0}\begin{bmatrix}\frac{\partial^2 L}{\partial r^2} & -\frac{\partial L}{\partial r\partial e}\\-\frac{\partial L}{\partial e\partial r} & \frac{\partial^2 L}{\partial e^2}\end{bmatrix}\begin{bmatrix}\frac{\partial L}{\partial \alpha\partial e}\\\frac{\partial L}{\partial \alpha\partial r}\end{bmatrix}$$

Computing the additional terms

$$\frac{\partial L}{\partial \alpha\partial e} = f_\varepsilon\left(\frac{q^* - e}{r}\right)\left(\frac{1}{r\cdot\alpha^2}\right)$$

$$= \frac{1}{\alpha^2\cdot r}f_\varepsilon(\cdot)$$

$$\frac{\partial L}{\partial \alpha\partial r} = 0 + \frac{1}{r\cdot 4\alpha^2}\left(\frac{q^* - e}{r}\right)f_\varepsilon\left(\frac{q^* - e}{r}\right)$$

$$= \frac{1}{\alpha^2\cdot r}\left(\frac{q^* - e}{r}\right)f_\varepsilon(\cdot)$$

Putting it altogether:

$$\frac{\partial e}{\partial \alpha} = -\underbrace{\frac{1}{\text{Determinant}}}_{+} \left[\overbrace{\frac{\partial^2 L}{\partial r^2}}^{-} \cdot \overbrace{\frac{\partial L}{\partial \alpha \partial e}}^{-} - \overbrace{\frac{\partial L}{\partial r \partial e}}^{-} \cdot \overbrace{\frac{\partial L}{\partial \alpha \partial r}}^{+} \right]$$

$$> 0$$

$$\frac{\partial r}{\partial \alpha} = -\underbrace{\frac{1}{\text{Determinant}}}_{+} \left[-\overbrace{\frac{\partial L}{\partial e \partial r}}^{-} \cdot \overbrace{\frac{\partial L}{\partial \alpha \partial e}}^{-} + \overbrace{\frac{\partial^2 L}{\partial e^2}}^{-} \cdot \overbrace{\frac{\partial L}{\partial \alpha \partial r}}^{+} \right]$$

$$< 0$$

Finally we investigate how a change in β affects effort and risk:

$$\begin{bmatrix} \frac{\partial e}{\partial \beta} \\ \frac{\partial r}{\partial \beta} \end{bmatrix} = - \begin{bmatrix} \frac{\partial^2 L}{\partial e^2} & \frac{\partial L}{\partial r \partial e} \\ \frac{\partial L}{\partial e \partial r} & \frac{\partial^2 L}{\partial r^2} \end{bmatrix}^{-1} \begin{bmatrix} \frac{\partial L}{\partial \beta \partial e} \\ \frac{\partial L}{\partial \beta \partial r} \end{bmatrix}$$

$$= -\underbrace{\frac{1}{\text{Determinant}}}_{S.O.C \text{ for Hessian} > 0} \begin{bmatrix} \frac{\partial^2 L}{\partial r^2} & -\frac{\partial L}{\partial r \partial e} \\ -\frac{\partial L}{\partial e \partial r} & \frac{\partial^2 L}{\partial e^2} \end{bmatrix} \begin{bmatrix} 0 \\ -1 \end{bmatrix}$$

Putting it altogether:

$$\frac{\partial e}{\partial \beta} = -\underbrace{\frac{1}{\text{Determinant}}}_{+} \left[\overbrace{\frac{\partial L}{\partial r \partial e}}^{-} \cdot \overbrace{1}^{+} \right]$$

$$> 0$$

$$\frac{\partial r}{\partial \beta} = -\underbrace{\frac{1}{\text{Determinant}}}_{+} \left[\overbrace{\frac{\partial^2 L}{\partial e^2}}^{-} \cdot \overbrace{(-1)}^{-} \right]$$

$$< 0$$

Step 5: Owner's optimal reporting choice

Constrained case

The owner chooses T to maximize his utility:

$$\max_{T} \quad T \cdot Pr(q \geq q^*) + E[\tilde{q} \mid q < q^*] \cdot Pr(q < q^*) - \gamma(r) \qquad \text{s.t}$$

$$E\left[q - T - \beta r \mid q \geq q^*\right] \cdot Pr(q \geq q^*) + E\left[q - \tilde{q} - \frac{\alpha}{4}(q - \tilde{q})^2 - \beta r \mid q < q^*\right] \cdot Pr(q < q^*) - h(e^*, r^*) > R$$

Expressing in terms of exogenous variables:

$$\max_{T} \quad T \cdot Pr\left(\varepsilon \geq \frac{q^* - e^*}{r}\right) + E\left[e + r\varepsilon - \frac{1}{2\alpha} \mid \varepsilon < \frac{q^* - e^*}{r}\right] \cdot Pr\left(\varepsilon < \frac{q^* - e^*}{r}\right) - \gamma(r) \qquad \text{s.t}$$

$$E\left[(e + r\varepsilon - T - \beta r) \mid \varepsilon \geq \frac{q^* - e}{r}\right] \cdot Pr\left(\varepsilon \geq \left(\frac{q^* - e}{r}\right)\right) + E\left[\frac{1}{\alpha} - \beta r \mid \varepsilon < \frac{q^* - e}{r}\right] \cdot Pr\left(\varepsilon < \frac{q^* - e}{r}\right) - h(e, r) \geq 0$$

Translating into integral notation:

$$L = T \underbrace{\int_{\frac{q^* - e^*}{r*}}^{\infty} f_{\varepsilon}(\varepsilon) d\varepsilon}_{A} + \underbrace{\int_{0}^{\frac{q^* - e^*}{r*}} \left(e + r\varepsilon - \frac{1}{2\alpha}\right) f_{\varepsilon}(\varepsilon) d\varepsilon}_{B} - \gamma(r) +$$

$$\lambda \left[\underbrace{\int_{\frac{q^* - e^*}{r*}}^{\infty} \left(e + r\varepsilon - T - \beta r\right) f_{\varepsilon}(\varepsilon) d\varepsilon}_{C} + \underbrace{\int_{0}^{\frac{q^* - e^*}{r*}} \left(\frac{1}{\alpha} - \beta r\right) f_{\varepsilon}(\varepsilon) d\varepsilon - h(e, r)}_{D} \right]$$

82

Taking the derivative with respect to T

$$= \underbrace{\int_{\frac{q^*-e^*}{r*}}^{\infty} f_\varepsilon(\varepsilon)d\varepsilon + T\left[0 - \frac{\partial}{\partial T}\left(\frac{q^*-e^*}{r^*}\right)f_\varepsilon(\cdot)\right]}_{A'} + \underbrace{\int_0^{\frac{q^*-e^*}{r*}}\left(\frac{\partial e}{\partial T} + \frac{\partial r}{\partial T}\varepsilon\right)f_\varepsilon(\varepsilon)d\varepsilon + \frac{\partial}{\partial T}\left(\frac{q^*-e^*}{r^*}\right)\left(T - \frac{1}{\alpha}\right)f_\varepsilon(\cdot)}_{B'} -\gamma'(r)\frac{\partial r}{\partial T}$$

$$+ \lambda\left[\underbrace{\int_{\frac{q^*-e^*}{r*}}^{\infty}\left(\frac{\partial e}{\partial T} + \frac{\partial r}{\partial T}\varepsilon - 1 - \left(\beta\cdot\frac{\partial r}{\partial T}\right)\right)f_\varepsilon(\varepsilon)d\varepsilon - \frac{\partial}{\partial T}\left(\frac{q^*-e^*}{r^*}\right)\frac{1}{\alpha}f_\varepsilon(\cdot)}_{C'}\right.$$

$$\left.+ \underbrace{\int_0^{\frac{q^*-e}{r}} -\left(\beta\cdot\frac{\partial r}{\partial T}\right)f_\varepsilon(\varepsilon)d\varepsilon + \frac{\partial}{\partial T}\left(\frac{q^*-e^*}{r^*}\right)\frac{1}{\alpha}f_\varepsilon(\cdot) - h'\left(\frac{\partial e}{\partial T} + \frac{\partial r}{\partial T}\right)}_{D'}\right]$$

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

$$+ \lambda\left[\underbrace{-(1 - F_\varepsilon(\cdot)) + \frac{\partial e}{\partial T}(1 - F_\varepsilon(\cdot)) + \frac{\partial r}{\partial T}\int_{\frac{q^*-e^*}{r*}}^{\infty}\varepsilon f_\varepsilon(\varepsilon)d\varepsilon - \frac{\partial}{\partial T}\cancel{\left(\frac{q^*-e^*}{r^*}\right)}\frac{1}{\alpha}f_\varepsilon(\cdot)}_{C'}\right.$$

$$\left.+ \underbrace{\frac{\partial}{\partial T}\cancel{\left(\frac{q^*-e^*}{r^*}\right)}\frac{1}{\alpha}f_\varepsilon(\cdot) - h'\left(\frac{\partial e}{\partial T} + \frac{\partial r}{\partial T}\right)}_{D'} - \beta\frac{\partial r}{\partial T}\right]$$

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

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

Taking the derivative with respect to λ

$$\int_{\frac{q^*-e^*}{r*}}^{\infty}(e + r\varepsilon - T - \beta r)f_\varepsilon(\varepsilon)d\varepsilon + \int_0^{\frac{q^*-e^*}{r*}}\left(\frac{1}{\alpha} - \beta r\right)f_\varepsilon(\varepsilon)d\varepsilon - h(e, r) = 0$$

83

Next we apply the IFT to understand how T changes with α.

$$
\begin{bmatrix} \frac{\partial T}{\partial \alpha} \\ \frac{\partial \lambda}{\partial \alpha} \end{bmatrix} = - \begin{bmatrix} \frac{\partial^2 L}{\partial T^2} & \frac{\partial L}{\partial \lambda \partial T} \\ \frac{\partial L}{\partial T \partial \lambda} & \frac{\partial^2 L}{\partial \lambda^2} \end{bmatrix}^{-1} \begin{bmatrix} \frac{\partial L}{\partial \alpha \partial T} \\ \frac{\partial L}{\partial \alpha \partial \lambda} \end{bmatrix}
$$

$$
= -\frac{1}{\text{Determinant}} \begin{bmatrix} \frac{\partial^2 L}{\partial \lambda^2} & -\frac{\partial L}{\partial \lambda \partial T} \\ -\frac{\partial L}{\partial T \partial \lambda} & \frac{\partial^2 L}{\partial T^2} \end{bmatrix}^{-1} \begin{bmatrix} \frac{\partial L}{\partial \alpha \partial T} \\ \frac{\partial L}{\partial \alpha \partial \lambda} \end{bmatrix}
$$

$$
= -\frac{1}{\text{Determinant}} \begin{bmatrix} 0 & -\frac{\partial L}{\partial \lambda \partial T} \\ -\frac{\partial L}{\partial T \partial \lambda} & \frac{\partial^2 L}{\partial T^2} \end{bmatrix}^{-1} \begin{bmatrix} \frac{\partial L}{\partial \alpha \partial T} \\ \frac{\partial L}{\partial \alpha \partial \lambda} \end{bmatrix}
$$

Taking each term in turn:

$$
\frac{\partial L}{\partial \lambda \partial T} = -(1 - F_\varepsilon(\cdot)) + \frac{\partial e}{\partial T}(1 - F_\varepsilon(\cdot)) + \frac{\partial r}{\partial T} \int_{\frac{q^* - e^*}{r*}}^{\infty} \varepsilon f_\varepsilon(\varepsilon) d\varepsilon + \frac{\partial r}{\partial T}(-\beta) - h'\left(\frac{\partial e}{\partial T} + \frac{\partial r}{\partial T}\right)
$$

$$
= -(1 - F_\varepsilon(\cdot)) + \underbrace{\frac{\partial e}{\partial T}\left(1 - F_\varepsilon(\cdot) - h'_e\right)}_{F.O.C = 0} + \underbrace{\frac{\partial r}{\partial T}\left(\int_{\frac{q^* - e^*}{r*}}^{\infty} \varepsilon f_\varepsilon(\varepsilon) d\varepsilon - h'_r - \beta\right)}_{F.O.C = 0}
$$

$$
= \underbrace{-(1 - F_\varepsilon(\cdot))}_{< 0}
$$

$$
\frac{\partial L}{\partial \alpha \partial \lambda} = \int_{\frac{q^* - e^*}{r*}}^{\infty} \left(\frac{\partial e}{\partial \alpha} + \frac{\partial r}{\partial \alpha}\varepsilon - \beta\frac{\partial r}{\partial \alpha}\right) f_\varepsilon(\varepsilon) d\varepsilon - \cancel{\frac{\partial}{\partial \alpha}\left(\frac{q^* - e^*}{r*}\right)\left(\frac{1}{\alpha} - \beta r\right) f_\varepsilon(\cdot)}
$$

$$
- \int_0^{\frac{q^* - e^*}{r*}} \left(\frac{1}{\alpha^2} + \beta\frac{\partial r}{\partial \alpha}\right) f_\varepsilon(\varepsilon) d\varepsilon + \cancel{\frac{\partial}{\partial \alpha}\left(\frac{q^* - e^*}{r*}\right)\left(\frac{1}{\alpha} - \beta r\right) f_\varepsilon(\cdot)} - h'\left(\frac{\partial e}{\partial \alpha} + \frac{\partial r}{\partial \alpha}\right)
$$

$$
= \int_{\frac{q^* - e^*}{r*}}^{\infty} \left(\frac{\partial e}{\partial \alpha} + \frac{\partial r}{\partial \alpha}\varepsilon - \beta\frac{\partial r}{\partial \alpha}\right) f_\varepsilon(\varepsilon) d\varepsilon - \int_0^{\frac{q^* - e^*}{r*}} \left(\frac{1}{\alpha^2} + \beta\frac{\partial r}{\partial \alpha}\right) f_\varepsilon(\varepsilon) d\varepsilon - h'\left(\frac{\partial e}{\partial \alpha} + \frac{\partial r}{\partial \alpha}\right)
$$

$$
= \frac{\partial e}{\partial \alpha}(1 - F_\varepsilon(\cdot)) + \frac{\partial r}{\partial \alpha}\int_{\frac{q^* - e^*}{r*}}^{\infty} \varepsilon f_\varepsilon(\varepsilon) d\varepsilon - \int_0^{\frac{q^* - e^*}{r*}} \left(\frac{1}{\alpha^2}\right) f_\varepsilon(\varepsilon) d\varepsilon - h'\left(\frac{\partial e}{\partial \alpha} + \frac{\partial r}{\partial \alpha}\right) - \beta\frac{\partial r}{\partial \alpha}
$$

$$
= \underbrace{\frac{\partial e}{\partial \alpha}(1 - F_\varepsilon(\cdot) - h'_e)}_{F.O.C = 0} + \underbrace{\frac{\partial r}{\partial \alpha}\left(\int_{\frac{q^* - e^*}{r*}}^{\infty} \varepsilon f_\varepsilon(\varepsilon) d\varepsilon - h'_r - \beta\right)}_{F.O.C = 0} - \int_0^{\frac{q^* - e^*}{r*}} \left(\frac{1}{\alpha^2}\right) f_\varepsilon(\varepsilon) d\varepsilon
$$

$$
= \underbrace{-\frac{1}{\alpha^2}F_\varepsilon(\cdot)}_{<0}
$$

Putting it all together:

$$\frac{\partial T}{\partial \alpha} = -\frac{1}{0 - \left(\frac{\partial L}{\partial T \partial \lambda}\right)^2}\left[-\frac{\partial L}{\partial \lambda \partial T} \cdot \frac{\partial L}{\partial \alpha \partial \lambda}\right]$$

$$= \underbrace{\frac{1}{\left(\frac{\partial L}{\partial T \partial \lambda}\right)^2}}_{+}\left[-\underbrace{\frac{\partial L}{\partial \lambda \partial T}}_{-} \cdot \underbrace{\frac{\partial L}{\partial \alpha \partial \lambda}}_{-}\right]$$

$$< 0$$

We can also apply the IFT to understand how T changes with β.

$$\begin{bmatrix}\frac{\partial T}{\partial \beta} \\ \frac{\partial \lambda}{\partial \beta}\end{bmatrix} = -\begin{bmatrix}\frac{\partial^2 L}{\partial T^2} & \frac{\partial L}{\partial \lambda \partial T} \\ \frac{\partial L}{\partial T \partial \lambda} & \frac{\partial^2 L}{\partial \lambda^2}\end{bmatrix}^{-1}\begin{bmatrix}\frac{\partial L}{\partial \beta \partial T} \\ \frac{\partial L}{\partial \beta \partial \lambda}\end{bmatrix}$$

$$= -\frac{1}{\text{Determinant}}\begin{bmatrix}\frac{\partial^2 L}{\partial \lambda^2} & -\frac{\partial L}{\partial \lambda \partial T} \\ -\frac{\partial L}{\partial T \partial \lambda} & \frac{\partial^2 L}{\partial T^2}\end{bmatrix}^{-1}\begin{bmatrix}\frac{\partial L}{\partial \beta \partial T} \\ \frac{\partial L}{\partial \beta \partial \lambda}\end{bmatrix}$$

$$= -\frac{1}{\text{Determinant}}\begin{bmatrix}0 & -\frac{\partial L}{\partial \lambda \partial T} \\ -\frac{\partial L}{\partial T \partial \lambda} & \frac{\partial^2 L}{\partial T^2}\end{bmatrix}^{-1}\begin{bmatrix}\frac{\partial L}{\partial \beta \partial T} \\ \frac{\partial L}{\partial \beta \partial \lambda}\end{bmatrix}$$

We compute the only new term required:

$$\frac{\partial L}{\partial \beta \partial \lambda} = \int_{\frac{q^* - e^*}{r^*}}^{\infty} -r f_\varepsilon(\varepsilon) d\varepsilon - \frac{\partial}{\beta}\left(\frac{q^* - e^*}{r^*}\right)\cancel{(q^* - T - \beta r)} f_\varepsilon(\cdot)$$

$$+ \int_0^{\frac{q^* - e^*}{r^*}} -r f_\varepsilon(\varepsilon) d\varepsilon + \frac{\partial}{\beta}\left(\frac{q^* - e^*}{r^*}\right)\left(\frac{1}{\alpha}\cancel{- \beta r}\right) f_\varepsilon(\cdot)$$

$$= -r$$

Putting it all together:

$$\frac{\partial T}{\partial \beta} = \underbrace{\frac{1}{\left(\frac{\partial L}{\partial T \partial \lambda}\right)^2}}_{+}\left[-\underbrace{\frac{\partial L}{\partial \lambda \partial T}}_{-} \cdot \underbrace{\frac{\partial L}{\partial \beta \partial \lambda}}_{-}\right]$$

$$< 0$$

Unconstrained case

The owner chooses T to maximize his utility:

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

Expressing in terms of exogenous variables:

$$\max_{T} \quad T \cdot Pr\left(\varepsilon \geq \frac{q^* - e^*}{r}\right) + E\left[e + r\varepsilon - \frac{1}{2\alpha} \mid \varepsilon < \frac{q^* - e^*}{r}\right] \cdot Pr\left(\varepsilon < \frac{q^* - e^*}{r}\right) - \gamma(r)$$

Translating into integral notation:

$$L = T\underbrace{\int_{\frac{q^*-e^*}{r*}}^{\infty} f_\varepsilon(\varepsilon)d\varepsilon}_{A} + \underbrace{\int_{0}^{\frac{q^*-e^*}{r*}}\left(e + r\varepsilon - \frac{1}{2\alpha}\right)f_\varepsilon(\varepsilon)d\varepsilon}_{B} - \gamma(r)$$

Taking the derivative with respect to T

$$= \underbrace{\int_{\frac{q^*-e^*}{r*}}^{\infty} f_\varepsilon(\varepsilon)d\varepsilon + T\left[0 - \frac{\partial}{\partial T}\cancel{\left(\frac{q^*-e^*}{r*}\right)f_\varepsilon(\cdot)}\right]}_{A'} + \underbrace{\int_{0}^{\frac{q^*-e^*}{r*}}\left(\frac{\partial e}{\partial T} + \frac{\partial r}{\partial T}\varepsilon\right)f_\varepsilon(\varepsilon)d\varepsilon + \frac{\partial}{\partial T}\left(\frac{q^*-e^*}{r*}\right)\cancel{\left(T - \frac{1}{\alpha}\right)}f_\varepsilon(\cdot)}_{B'} - \gamma'(r)\frac{\partial r}{\partial T}$$

$$= \underbrace{\int_{\frac{q^*-e^*}{r*}}^{\infty} f_\varepsilon(\varepsilon)d\varepsilon}_{A'} + \underbrace{\int_{0}^{\frac{q^*-e^*}{r*}}\left(\frac{\partial e}{\partial T} + \frac{\partial r}{\partial T}\varepsilon\right)f_\varepsilon(\varepsilon)d\varepsilon - \frac{\partial}{\partial T}\left(\frac{q^*-e^*}{r*}\right)\left(\frac{1}{\alpha}\right)f_\varepsilon(\cdot)}_{B'} - \gamma'(r)\frac{\partial r}{\partial T}$$

$$= \underbrace{(1 - F_\varepsilon(\cdot))}_{A'} + \underbrace{\frac{\partial e}{\partial T}F_\varepsilon(\cdot) + \frac{\partial r}{\partial T}\left(\int_{0}^{\frac{q^*-e^*}{r*}}\varepsilon f_\varepsilon(\varepsilon)d\varepsilon\right) - \frac{1}{\alpha}f_\varepsilon(\cdot)\frac{\partial}{\partial T}\left(\frac{q^*-e^*}{r*}\right)}_{B'} - \gamma'(r)\frac{\partial r}{\partial T}$$

Next we apply the IFT to understand how T changes with α:

$$\frac{\partial T}{\partial \alpha} = -\frac{\frac{\partial L}{\partial \alpha \partial T}}{\frac{\partial^2 L}{\partial T^2}}$$

We know that $\frac{\partial^2 L}{\partial T^2} < 0$ because it's a second order condition. Now let's see if we can express $\frac{\partial L}{\partial \alpha \partial T}$ as a function of the S.O.C

$$\frac{\partial e}{\partial T} = -\frac{1}{D}\left[\frac{\partial^2 L}{\partial r^2} \cdot \frac{\partial L}{\partial T \partial e} - \frac{\partial L}{\partial r \partial e} \cdot \frac{\partial L}{\partial T \partial r}\right]$$

$$\frac{\partial e}{\partial \alpha} = -\frac{1}{D}\left[\frac{\partial^2 L}{\partial r^2} \cdot \frac{\partial L}{\partial \alpha \partial e} - \frac{\partial L}{\partial r \partial e} \cdot \frac{\partial L}{\partial \alpha \partial r}\right]$$

$$\frac{\partial r}{\partial T} = -\frac{1}{D}\left[-\frac{\partial L}{\partial e \partial r} \cdot \frac{\partial L}{\partial T \partial e} + \frac{\partial^2 L}{\partial e^2} \cdot \frac{\partial L}{\partial T \partial r}\right]$$

$$\frac{\partial r}{\partial \alpha} = -\frac{1}{D}\left[-\frac{\partial L}{\partial e \partial r} \cdot \frac{\partial L}{\partial \alpha \partial e} + \frac{\partial^2 L}{\partial e^2} \cdot \frac{\partial L}{\partial \alpha \partial r}\right]$$

Where

$$\frac{\partial L}{\partial T \partial e} = -\frac{1}{r} f_\varepsilon(\cdot)$$

$$\frac{\partial L}{\partial \alpha \partial e} = \frac{1}{\alpha^2 r} f_\varepsilon(\cdot)$$

$$\frac{\partial L}{\partial T \partial r} = -\frac{1}{r} f_\varepsilon(\cdot) \left(\frac{q^* - e}{r} \right)$$

$$\frac{\partial L}{\partial \alpha \partial r} = \frac{1}{\alpha^2 r} f_\varepsilon(\cdot) \left(\frac{q^* - e}{r} \right)$$

Therefore we can write:

$$\frac{\partial e}{\partial \alpha} = -\frac{1}{\alpha^2} \frac{\partial e}{\partial T}$$

$$\frac{\partial}{\partial \alpha} \left[\frac{\partial e}{\partial T} \right] = -\frac{1}{\alpha^2} \frac{\partial}{\partial T} \left[\frac{\partial e}{\partial T} \right]$$

$$\frac{\partial r}{\partial \alpha} = -\frac{1}{\alpha^2} \frac{\partial r}{\partial T}$$

$$\frac{\partial}{\partial \alpha} \left[\frac{\partial r}{\partial T} \right] = -\frac{1}{\alpha^2} \frac{\partial}{\partial T} \left[\frac{\partial e}{\partial T} \right]$$

Also as a sanity check:

$$\frac{\partial}{\partial T} \left(\frac{q^* - e}{r} \right) = \frac{\left(1 - \frac{\partial e}{\partial T} \right) r - (q^* - e) \frac{\partial r}{\partial T}}{r^2}$$

$$\frac{\partial}{\partial \alpha} \left(\frac{q^* - e}{r} \right) = \frac{- \left(\frac{1}{\alpha^2} - \frac{\partial e}{\partial \alpha} \right) r - (q^* - e) \frac{\partial r}{\partial \alpha}}{r^2} = \frac{- \left(\frac{1}{\alpha^2} + \frac{1}{\alpha^2} \frac{\partial e}{\partial T} \right) r - (q^* - e) \left(-\frac{1}{\alpha^2} \right) \frac{\partial r}{\partial T}}{r^2} = -\frac{1}{\alpha^2} \frac{\partial}{\partial T} \left(\frac{q^* - e}{r} \right)$$

Finally:

$$\frac{\partial L}{\partial T^2} = -f_\varepsilon(\cdot) \frac{\partial}{\partial T} \left(\frac{q^* - e}{r} \right) + \frac{\partial^2 e}{\partial T} F_\varepsilon(\cdot) + \frac{\partial e}{\partial T} f_\varepsilon(\cdot) \frac{\partial}{\partial T} \left(\frac{q^* - e}{r} \right) + \frac{\partial^2 r}{\partial T} \left(\int_0^{\frac{q^* - e^*}{r^*}} \varepsilon f_\varepsilon(\varepsilon) d\varepsilon \right)$$

$$+ \frac{\partial r}{\partial T} \left[0 + \frac{\partial}{\partial T} \left(\frac{q^* - e}{r} \right) \left(\frac{q^* - e^*}{r^*} \right) f_\varepsilon(\cdot) \right] - \frac{1}{\alpha} f_\varepsilon'(\cdot) \left[\frac{\partial}{\partial T} \left(\frac{q^* - e}{r} \right) \right]^2 - \frac{1}{\alpha} f_\varepsilon(\cdot) \frac{\partial}{\partial T} \left[\frac{\partial}{\partial T} \left(\frac{q^* - e}{r} \right) \right] - \frac{\partial}{\partial T} \left[\gamma'(r) \frac{\partial r}{\partial T} \right]$$

$$= - \left(1 - \frac{\partial e}{\partial T} \right) f_\varepsilon(\cdot) \frac{\partial}{\partial T} \left(\frac{q^* - e}{r} \right) + \frac{\partial^2 e}{\partial T} F_\varepsilon(\cdot) + \frac{\partial^2 r}{\partial T^2} \left(\int_0^{\frac{q^* - e^*}{r^*}} \varepsilon f_\varepsilon(\varepsilon) d\varepsilon \right) + \frac{\partial r}{\partial T} f_\varepsilon(\cdot) \frac{\partial}{\partial T} \left(\frac{q^* - e}{r} \right) \left(\frac{q^* - e^*}{r^*} \right)$$

$$- \frac{1}{\alpha} f_\varepsilon'(\cdot) \left[\frac{\partial}{\partial T} \left(\frac{q^* - e}{r} \right) \right]^2 - \frac{1}{\alpha} f_\varepsilon(\cdot) \frac{\partial}{\partial T} \left[\frac{\partial}{\partial T} \left(\frac{q^* - e}{r} \right) \right] - \left[\gamma''(r) \left(\frac{\partial r}{\partial T} \right)^2 + \gamma'(r) \frac{\partial^2 r}{\partial T^2} \right]$$

And:

$$\frac{\partial L}{\partial \alpha \partial T} = -f_\varepsilon(\cdot)\frac{\partial}{\partial \alpha}\left(\frac{q^*-e}{r}\right) + \frac{\partial}{\partial \alpha}\left[\frac{\partial e}{\partial T}\right]F_\varepsilon(\cdot) + \frac{\partial e}{\partial T}f_\varepsilon(\cdot)\frac{\partial}{\partial \alpha}\left(\frac{q^*-e}{r}\right) + \frac{\partial}{\partial \alpha}\left[\frac{\partial r}{\partial T}\right]\left(\int_0^{\frac{q^*-e^*}{r*}}\varepsilon f_\varepsilon(\varepsilon)d\varepsilon\right)$$

$$+ \frac{\partial r}{\partial T}\left[0 + \frac{\partial}{\partial \alpha}\left(\frac{q^*-e}{r}\right)\left(\frac{q^*-e^*}{r*}\right)f_\varepsilon(\cdot)\right] - \frac{1}{\alpha}f_\varepsilon'(\cdot)\frac{\partial}{\partial \alpha}\left(\frac{q^*-e}{r}\right)\frac{\partial}{\partial T}\left(\frac{q^*-e}{r}\right) - \frac{1}{\alpha}f_\varepsilon(\cdot)\frac{\partial}{\partial \alpha}\left[\frac{\partial}{\partial T}\left(\frac{q^*-e}{r}\right)\right]$$

$$+ \frac{1}{\alpha^2}f_\varepsilon\left[\frac{\partial}{\partial T}\left(\frac{q^*-e}{r}\right)\right] - \frac{\partial}{\partial \alpha}\left[\gamma'(r)\frac{\partial r}{\partial T}\right]$$

$$= -\left(1 - \frac{\partial e}{\partial T}\right)f_\varepsilon(\cdot)\frac{\partial}{\partial \alpha}\left(\frac{q^*-e}{r}\right) + \frac{\partial}{\partial \alpha}\left[\frac{\partial e}{\partial T}\right]F_\varepsilon(\cdot) + \frac{\partial}{\partial \alpha}\left[\frac{\partial r}{\partial T}\right]\left(\int_0^{\frac{q^*-e^*}{r*}}\varepsilon f_\varepsilon(\varepsilon)d\varepsilon\right) + \frac{\partial r}{\partial T}f_\varepsilon(\cdot)\frac{\partial}{\partial \alpha}\left(\frac{q^*-e}{r}\right)\left(\frac{q^*-e^*}{r*}\right)$$

$$- \frac{1}{\alpha}f_\varepsilon'(\cdot)\frac{\partial}{\partial \alpha}\left(\frac{q^*-e}{r}\right)\frac{\partial}{\partial T}\left(\frac{q^*-e}{r}\right) - \frac{1}{\alpha}f_\varepsilon(\cdot)\frac{\partial}{\partial \alpha}\left[\frac{\partial}{\partial T}\left(\frac{q^*-e}{r}\right)\right]$$

$$+ \frac{1}{\alpha^2}f_\varepsilon\left[\frac{\partial}{\partial T}\left(\frac{q^*-e}{r}\right)\right] - \left[\gamma''(r)\left(\frac{\partial r}{\partial \alpha}\right)\left(\frac{\partial r}{\partial T}\right) + \gamma'(r)\frac{\partial}{\partial \alpha}\frac{\partial r}{\partial T}\right]$$

$$= -\frac{1}{\alpha^2}\left[-\left(1 - \frac{\partial e}{\partial T}\right)f_\varepsilon(\cdot)\frac{\partial}{\partial T}\left(\frac{q^*-e}{r}\right)\right] - \frac{1}{\alpha^2}\left[\frac{\partial^2 e}{\partial T^2}\right]F_\varepsilon(\cdot) - \frac{1}{\alpha^2}\left[\frac{\partial^2 r}{\partial T^2}\right]\left(\int_0^{\frac{q^*-e^*}{r*}}\varepsilon f_\varepsilon(\varepsilon)d\varepsilon\right)$$

$$- \frac{1}{\alpha^2}\left[\frac{\partial r}{\partial T}f_\varepsilon(\cdot)\frac{\partial}{\partial T}\left(\frac{q^*-e}{r}\right)\left(\frac{q^*-e^*}{r*}\right)\right] - \frac{1}{\alpha^2}\left[-\frac{1}{\alpha}f_\varepsilon'(\cdot)\left[\frac{\partial}{\partial T}\left(\frac{q^*-e}{r}\right)\right]^2\right] - \frac{1}{\alpha^2}\left[-\frac{1}{\alpha}f_\varepsilon(\cdot)\frac{\partial}{\partial T}\left[\frac{\partial}{\partial T}\left(\frac{q^*-e}{r}\right)\right]\right]$$

$$+ \frac{1}{\alpha^2}f_\varepsilon\left[\frac{\partial}{\partial T}\left(\frac{q^*-e}{r}\right)\right] + \frac{1}{\alpha^2}\left[\gamma''(r)\left(\frac{\partial r}{\partial T}\right)\left(\frac{\partial r}{\partial T}\right) + \gamma'(r)\frac{\partial}{\partial T}\frac{\partial r}{\partial T}\right]$$

$$= \underbrace{-\frac{1}{\alpha^2}\left(\frac{\partial^2 L}{\partial T^2}\right)}_{-} + \underbrace{\frac{1}{\alpha^2}f_\varepsilon\left[\frac{\partial}{\partial T}\left(\frac{q^*-e}{r}\right)\right]}_{+}$$

Putting it altogether

$$\frac{\partial T}{\partial \alpha} = -\frac{\overbrace{\frac{\partial L}{\partial \alpha \partial T}}^{+}}{\underbrace{\frac{\partial^2 L}{\partial T^2}}_{-}}$$

$$> 0$$

Applying the same analysis to β, is much simpler because $\frac{\partial}{\partial\beta}\frac{\partial r}{\partial T}$ and $\frac{\partial}{\partial\beta}\frac{\partial e}{\partial T}$ equal zero. Moreover:

$$\frac{\partial}{\partial\beta}\left(\frac{q^*-e^*}{r^*}\right)=\frac{-\frac{\partial e}{\partial\beta}r-(q^*-e^*)\frac{\partial r}{\partial\beta}}{(r^*)^2}<0$$

$$\frac{\partial}{\partial T}\left(\frac{q^*-e^*}{r^*}\right)=\frac{\left(1-\frac{\partial e}{\partial T}\right)r-(q^*-e)\frac{\partial r}{\partial T}}{r^2}>0$$

$$\frac{\partial}{\partial\beta}\left[\frac{\partial}{\partial T}\left(\frac{q^*-e^*}{r^*}\right)\right]=\frac{r^2\left[\left(1-\frac{\partial e}{\partial T}\right)\frac{\partial r}{\partial\beta}+\frac{\partial e}{\partial\beta}\frac{\partial r}{\partial T}\right]-2r\frac{\partial r}{\partial\beta}\left[\left(1-\frac{\partial e}{\partial T}\right)r-(q^*-e)\frac{\partial r}{\partial T}\right]}{(r^*)^4}$$

$$=\frac{r^2\left(1-\frac{\partial e}{\partial T}\right)\frac{\partial r}{\partial\beta}+r^2\frac{\partial e}{\partial\beta}\frac{\partial r}{\partial T}-2r^2\left(1-\frac{\partial e}{\partial T}\right)\frac{\partial r}{\partial\beta}+2r(q^*-e)\frac{\partial r}{\partial\beta}\frac{\partial r}{\partial T}}{(r^*)^4}$$

$$=\frac{-r^2\left(1-\frac{\partial e}{\partial T}\right)\frac{\partial r}{\partial\beta}+r^2\frac{\partial e}{\partial\beta}\frac{\partial r}{\partial T}+2r(q^*-e)\frac{\partial r}{\partial\beta}\frac{\partial r}{\partial T}}{(r^*)^4}$$

$$>0$$

Then

$$\frac{\partial L}{\partial\beta\partial T}=-f_\varepsilon(\cdot)\frac{\partial}{\partial\beta}\left(\frac{q^*-e^*}{r^*}\right)+\frac{\partial}{\partial\beta}\cancel{\left[\frac{\partial e}{\partial T}\right]F_\varepsilon(\cdot)}+\frac{\partial e}{\partial T}f_\varepsilon(\cdot)\frac{\partial}{\partial\beta}\left(\frac{q^*-e}{r}\right)+\frac{\partial}{\partial\beta}\cancel{\left[\frac{\partial r}{\partial T}\right]\left(\int_0^{\frac{q^*-e^*}{r^*}}\varepsilon f_\varepsilon(\varepsilon)d\varepsilon\right)}$$

$$+\frac{\partial r}{\partial T}\left[0+\frac{\partial}{\partial\beta}\left(\frac{q^*-e^*}{r^*}\right)\left(\frac{q^*-e^*}{r^*}\right)f_\varepsilon(\cdot)\right]-\frac{1}{\alpha}f'_\varepsilon(\cdot)\frac{\partial}{\partial\beta}\left(\frac{q^*-e^*}{r^*}\right)\frac{\partial}{\partial T}\left(\frac{q^*-e^*}{r^*}\right)-\frac{1}{\alpha}f_\varepsilon(\cdot)\frac{\partial}{\partial\beta}\left[\frac{\partial}{\partial T}\left(\frac{q^*-e^*}{r^*}\right)\right]$$

$$-\frac{\partial}{\partial\beta}\left[\gamma'(r)\frac{\partial r}{\partial T}\right]$$

$$=-f_\varepsilon(\cdot)\frac{\partial}{\partial\beta}\left(\frac{q^*-e^*}{r^*}\right)+\frac{\partial e}{\partial T}f_\varepsilon(\cdot)\frac{\partial}{\partial\beta}\left(\frac{q^*-e}{r}\right)$$

$$+\frac{\partial r}{\partial T}\left[\frac{\partial}{\partial\beta}\left(\frac{q^*-e^*}{r^*}\right)\left(\frac{q^*-e^*}{r^*}\right)f_\varepsilon(\cdot)\right]-\frac{1}{\alpha}f'_\varepsilon(\cdot)\frac{\partial}{\partial\beta}\left(\frac{q^*-e^*}{r^*}\right)\frac{\partial}{\partial T}\left(\frac{q^*-e^*}{r^*}\right)-\frac{1}{\alpha}f_\varepsilon(\cdot)\frac{\partial}{\partial\beta}\left[\frac{\partial}{\partial T}\left(\frac{q^*-e^*}{r^*}\right)\right]$$

$$-\left[\gamma''(r)\frac{\partial r}{\partial\beta}\cdot\frac{\partial r}{\partial T}+\gamma'(r)\cancel{\frac{\partial}{\partial\beta}\frac{\partial r}{\partial T}}\right]$$

$$=\overbrace{-f_\varepsilon(\cdot)\frac{\partial}{\partial\beta}\left(\frac{q^*-e^*}{r^*}\right)}^{+}+\overbrace{\frac{\partial e}{\partial T}f_\varepsilon(\cdot)\frac{\partial}{\partial\beta}\left(\frac{q^*-e}{r}\right)}^{+}+\overbrace{\frac{\partial r}{\partial T}\left[\frac{\partial}{\partial\beta}\left(\frac{q^*-e^*}{r^*}\right)\left(\frac{q^*-e^*}{r^*}\right)f_\varepsilon(\cdot)\right]}^{+}$$

$$\underbrace{-\frac{1}{\alpha}f'_\varepsilon(\cdot)\frac{\partial}{\partial\beta}\left(\frac{q^*-e^*}{r^*}\right)\frac{\partial}{\partial T}\left(\frac{q^*-e^*}{r^*}\right)}_{+}-\underbrace{\left[\gamma''(r)\frac{\partial r}{\partial\beta}\cdot\frac{\partial r}{\partial T}\right]}_{+}-\underbrace{\frac{1}{\alpha}f_\varepsilon(\cdot)\frac{\partial}{\partial\beta}\left[\frac{\partial}{\partial T}\left(\frac{q^*-e^*}{r^*}\right)\right]}_{-}$$

$$>0$$

Putting it altogether

$$\frac{\partial T}{\partial \beta} = -\frac{\overbrace{\dfrac{\partial L}{\partial \beta \partial T}}^{+}}{\underbrace{\dfrac{\partial^2 L}{\partial T^2}}_{-}}$$

$$> 0$$