The Supply of Bribes:
Evidence from Roadway Tolls in the D.R. Congo

Otis Reid and Jonathan Weigel

April 7, 2019

1 Abstract

Why is corruption so resilient? While many recent papers examine interventions seeking to discipline bureaucrats, less is known about the elasticity of citizens' supply of bribes. This randomized controlled trial explores citizen bribe payment at roadway tolls in Kananga, D.R. Congo. We offer financial and social incentives to motorcycle taxi drivers to bring receipts proving that they paid the legal toll. Observing a 7 to 10 percentage point increase in legal transactions due to financial incentives, we estimate an elasticity of citizen supply of bribes ranging from -0.45 to -0.95. Social incentives have no effect. We argue that limited responsiveness to large financial incentives reflects the fact that bribe payment reduces driver time at tolls by nearly 70%. Toll officers know their payer population is highly time constrained and so endogenously create an elaborate payment procedure to increase the time costs of paying the tax instead of a bribe. This result suggests that citizen-side anticorruption interventions may have limited effectiveness in settings in which (i) citizens are time constrained and bribes expedite the transaction, or (ii) bureaucrats can endogenously increase the hassle costs of citizen compliance.

2 Introduction

Corruption is often deemed detrimental to economic outcomes (Ferraz et al., 2012; Kaufmann and Wei, 1999; Méon and Sekkat, 2005). It also limits the ability of governments to raise revenue. Although some interventions appear to reduce corruption, the magnitudes of such reductions are often small. For example, in Olken (2007), the threat of government audits of a road-building project in Indonesia decreased estimated leakage by 8 percentage points; but even with a 100% probability of an audit, the level of missing expenditures remained at 19%. Indeed, perhaps the more striking result from this and other studies is just how persistent corrupt transactions appear to be despite interventions targeted against them.

The persistence of corruption could be explained by a range of factors. In many corrupt interactions, there is both a bribe-taking bureaucrat and a bribe-paying citizen. Recent research has largely focused on the bureaucrat side. For example, monitoring technologies have proven effective in disciplining absenteeism (Dhaliwal and Hanna, 2013; Duflo et al., 2012) and monetary incentives have increased tax collector effort, albeit also increasing bribe levels (Khan et al., 2016). The implicit view motivating these papers is that corruption is perpetuated by low-quality institutions that create weak or perverse incentives for bureaucrats. Less is known about the citizen side of corruption: in a literature review on corruption, Olken and Pande (2012) mention "bureaucrats" or "bureaucracy" 46 times; they mention 'citizen' 5 times. However, it is possible that citizens play a role in perpetuating corruption. In particular, an older theoretical literature notes how corruption can "grease the wheels" in settings of low institutional quality: paying a bribe might enable citizens to access public services faster than navigating the red tape in a bloated bureaucracy (Leys, 1965; Lui, 1985). In short, citizens might supply bribes because they increase the efficiency of their transactions with the state.

This paper explores citizen supply of corruption in the context of roadway tolls in Kananga, Democratic Republic of Congo (DRC). Motorcycle taxi drivers were offered financial and social incentives to bring receipts proving that they paid the legal toll. The goal of the financial incentive treatments was to estimate the elasticity of citizen supply of bribes with respect to the price of complying with the toll. The financial incentives were either 1000 Congolese Francs (FC) — about \$1 (half the price of the toll) — or 2000 FC (the full price of the toll) for drivers with proper receipts. The goal of the social incentives was to estimate the extent to

which corruption is affected by the perceived social value of tax compliance. Social incentives included (1) a pledge by the researchers to contribute 2000 FC to a widows' shelter in Kananga, and (2) a pledge to contribute 2000 FC to the provincial government to subsequently transfer to a widows' shelter in Kananga.[1] The goal of these treatments is to test if citizens are willing to pay bribes because they perceive it to have low social cost in light of high-level government corruption. If citizens trust our pledge to directly give widows 2000 FC for each valid receipt more than they trust the government to follow through on transferring this money, they should respond differentially to the former incentive relative to the latter. Finally, to net out experimenter demand effects, drivers in the control group were also asked to bring receipts but without any additional reward.

Participants who completed the baseline survey were randomized into different experimental conditions for their next trip outside of Kananga. After each of two three-week periods, they returned for a follow-up visit and were subsequently re-assigned to a different treatment group, such that each participant was assigned to three different treatments over the course of the experiment. Non-attriting participants were always assigned once to a financial treatment, once to a social incentive, and once to the control group.

Figure 1 summarizes intent-to-treat (ITT) results. The financial incentive treatments caused a 4 percentage-point increase in the probability of legal toll payments. IV estimates — in which driver treatment recall of their treatment status is instrumented with true treatment status — are larger, increasing the effect size to 11 percentage points. The implied elasticity of citizen supply of bribes with respect to the effective price of the toll is negative, as expected, but is relatively small in absolute value, ranging from -0.45 to -0.95 depending on estimates of drivers' discount factors. Our estimated effects on the magnitude of the equilibrium bribe are mixed, though a Tobit model suggests negative effects consistent with the simple bargaining theory we introduce in Section 5.1. The social incentive groups are not statistically distinguishable from the control group.

This paper contributes to the growing field experimental literature on corruption. Prior work has demonstrated the effectiveness of monitoring technologies in decreasing absenteeism in schools (Duflo et al., 2012) and public-sector health facilities (Dhaliwal and Hanna, 2013). Khan et al. (2016) find that incentive pay for tax

---

[1]We chose a widows' shelter for this treatment because pilot surveys we conducted in February 2016 revealed that such shelters were the modal charity given to by this subject population.

collectors in Pakistan reduces the frequency of bribes and boosts government revenues. Most similar is Bertrand et al. (2007) who find that individuals in India who are promised financial rewards if they obtain a driver's license quickly are indeed more likely to obtain a license, but they are also more likely to pay bribes and they are worse drivers on average. Past work predominantly focuses on the bureaucrat side, and ours is the first field experiment to incentivize citizens to forego paying bribes. The closest non-experimental study is Naritomi (2015), a study of a government policy in Brazil incentivizing consumers to obtain receipts for final products. In that context, she finds sizable decreases in evasion of the value-added tax in sectors with customers (not firms) as final consumers, consistent with other studies documenting the effectiveness of third-party information and VAT compliance (e.g. Pomeranz (2015). A likely explanation for the different results we observe is that consumers are not trading off time costs and financial costs to the same degree as they are in the toll setting we analyze.[2] Our study suggests that efforts to intervene on the citizen side of a given bribe-taking transaction may be hampered when bribes function as an "efficient grease" to an otherwise slow bureaucratic process. This may be particularly true when toll officers can strategically increase the cost of compliance to elicit more bribes.

Our results therefore offer experimental evidence in support of the "greasing the wheels" hypothesis on corruption. This hypothesis comes from an older theoretical literature noting several channels by which corruption could improve efficiency in a non-Coasian, second-best world. Huntington (1968) memorably summarizes this view: "The only thing worse than a society with a rigid, over-centralized, dishonest bureaucracy is one with a rigid, over-centralized, honest bureaucracy." We focus on the "speed money" mechanism from this literature: corruption can speed up bureaucratic processes that are otherwise beset by red tape (Leys, 1965; Lui, 1985).[3] If paying bribes enables individuals to obtain a needed government service faster

---

[2]An alternative explanation is that in Naritomi's setting, citizens are not publicly breaking a corrupt bargain, but instead are privately defecting by turning in their receipts.

[3]The literature offers two other reasons why corruption could increase efficiency. First, corruption could enable individuals to dodge bad public policy (Leff, 1964). If a regulation does more harm than good, and firms can bypass it by paying bribes, again corruption could enhance the economy's efficiency. As Méon and Sekkat (2005) put it: "Graft may simply be a hedge against bad public policies." Second, firms making corrupt bids for government contracts could approximate a competitive auction and outperform other allocation rules — such as government favoritism — if these bids reflect the underlying efficiency of the bidding firms (Beck and Maher, 1986; Lien, 1986). These two mechanisms are not relevant in our toll setting, so we focus on the "speed money" view of corruption.

than by navigating the bureaucracy, corruption can improve the efficiency of the transaction. Note that this argument takes institutional inefficiencies as exogenous. But, as Bardhan (1997) points out, "The distortions are not exogenous to the system and are instead often part of the built-in corrupt practices of a patron-client political system." Given these relationships, we expect endogenous growth of distortionary red tape. Our study speaks to the effects of corruption in those situations, not to the overall welfare effects of corruption in general.

The observational literature on the effects of corruption on growth finds mixed results. Several analyses of the average effects of corruption on growth reject the "grease the wheels" hypothesis (Fisman and Svensson, 2007; Mauro, 1995; Méon and Sekkat, 2005), but other papers have found support for it in settings of low institutional quality. Indeed, Méon and Sekkat (2005) observed that average negative effects of corruption on growth could belie heterogeneity by institutional quality. Corruption could promote efficiency in settings with bad institutions — consistent with the "grease the wheels" hypothesis — even as it causes misallocation in settings with good institutions. Along these lines, Méon and Weill (2010) find cross-country evidence that corruption is associated with lower efficiency costs in settings of low institutional quality. Similarly, Dreher and Gassebner (2011) find that corruption is positively correlated with firm entry in heavily regulated countries. On a more micro level, Vial and Hanoteau (2010) find that Indonesian manufacturing plants that paid more bribes also grew more during the 1975-1995 period.[4] Kato and Sato (2015) reach a similar conclusion in a study of manufacturing in India. Our paper seeks to complement this literature by providing experimental, well-identified evidence on this topic in a relevant low-capacity context.

We begin by discussing the context of our study in Section 3. We then present the experimental design in Section 4 and a simple theory in Section 5. We discuss our data in Section 6 and we then present our results in Section 8. Finally, we discuss the interpretation of our results in Section ?? and offer concluding remarks.

---

[4]This finding contradicts evidence from Uganda that follows an identical firm-level estimation strategy (Fisman and Svensson, 2007). Vial and Hanoteau (2010) argue that this reflects the forward-looking nature of Indonesian officials during the Suharto period and the long-term deals they struck with firms to maximize rents overtime; they characterize firms and officials in Uganda as dealing with greater uncertainty and so more likely to extract more today at the expense of tomorrow.

3 Context

3.1 Setting

The Democratic Republic of Congo ranks 147th out of 168 countries in the Corruption Perceptions Index of 2015 according to Transparency International. Our experiment takes place in Kananga, a city of roughly 1 million (the fourth largest in the country) and the capital of the Kasai Central province. In a quasi-random sample of households, Lowes et al. (2016) find that self-reported median monthly household income is approximately $70, or $111 at PPP. The local currency is the Congolese Franc (FC) and during the period of our study, 1,000 FC was worth $1.00-$1.03.

Motorcycle taxis (known as "motos") are the most common form of transportation in Kananga and the surrounding areas. Moto trips out of town take an average of one full day to complete. All routes out of the city pass by a toll station, where motorcycles must stop and pay the tax of 2,000 FC, show a receipt that they have already paid at that toll within the last 48 hours, or bribe the toll officer. Toll officers are occasionally rotated to new posts, but this occurs only rarely — during our sample period, we do not observe any re-assignment of toll officers. Due to the extremely poor quality of roads in the area, it is very difficult to avoid passing by one of these tolls stations when leaving Kananga. Paying the full toll amounts to about 13% of the median pre-toll estimated trip profit (15,000 FC). On about 85% of trips, motorcycle taxis transport passengers; on the remainder of such trips, they transport goods.

Our experiment occurred from May to September, 2016, though 94% of data collection was completed by the end of August. May to August is the dry season, when travel is most common.[5]

3.2 Study Population

Motorcycle taxi drivers (known as "motards") are 100% male and generally in their mid twenties. Almost all of them have reached secondary school, speaking at least

---

[5]In mid August, there was a rare outbreak of violence near the city, when a prominent local sub-chief of the predominant tribe launched a rebellion against the government, culminating in an attack on the local airport on September 23-24. He is believed to have been killed in the ensuing fighting, but the violence and deployment of soldiers reduced travel in the direction of his home territory substantially. Fighting related to this rebellion is ongoing in the area.

some French in addition to the local language of Tshiluba. The median driver has been driving for 4-5 years.

As we discuss further in Section 4.1, to participate, a driver first needed to complete an interview with an enumerator somewhere in the city of Kananga. The sampling process was not explicitly random, but we believe it approximated a (partial) census of motorcycle taxi drivers. In the initial interview, drivers were invited to come to a baseline visit if they had taken at least 1 trip outside the city in the prior 2 months. Importantly, however, we did not randomize drivers into treatment or control unless they completed the baseline interview. This reduced attrition during the study considerably.

In Table 1 we show observables gathered during the initial interview between drivers who were invited to the baseline survey and showed up compared to those who did not show up. Those who showed up do not substantially differ in observables from those who never came to the office. This suggests that our experimental estimates are generalizable to the broader population of drivers in Kananga.

4 Experimental Design

4.1 Timing

Our experiment was conducted over three phases, as follows:

1. **Recruitment:** we sent enumerators to intersections throughout Kananga where motorcycle taxis were known to gather or pass frequently. Our enumerators stopped individual motards and administered a brief survey to gather a small number of covariates and to determine eligibility for participation in our experiment. Drivers were paid 500 FC (about $0.50) for participation in this short survey and were eligible to earn an additional 200 FC in an accompanying game. Motards were eligible to participate if they had taken a trip outside the city of Kananga in the prior 2 months. Individuals who were eligible were invited to a baseline visit on a randomly selected day in the following three weeks.[6] 1,616 drivers were interviewed, of whom 1,219 were invited to a baseline visit.

---

[6]We used stickers, phone numbers, and a screening question to prevent the same individuals from joining the sample twice, as well as comparing photos of individuals reporting similar names to remove people from the sample who successfully entered twice.

2. **Baseline:** any driver who was invited for and attended a baseline visit was assigned to treatment and entered the experimental sample. Drivers were paid 3,000 FC for attending the baseline visit and were eligible to earn up to an additional 1,200 FC in an accompanying game. Of the 1,219 eligible drivers, 912 came for a baseline visit and became part of the experimental sample. Table 1 shows that motards who showed up for the baseline survey do not differ substantially from those who did not.

3. **Follow-up visits (3 total):** at the baseline, the driver was asked to return in 22 days (or 23 days, if the 22nd day was a Sunday) for a follow-up visit to collect outcomes from the treatment round. During this follow-up visit, the driver received an invitation to return to the office at the completion of the next treatment round in 3 weeks (21 days). During the second follow-up visit, the driver was again invited back in another 3 weeks at the conclusion of the last treatment round. In other words, drivers completed a maximum of 4 total visits: baseline, plus 3 follow-up visits conducted at the end of each treatment round). Drivers were paid 3,000 FC for attending each follow-up visit as a show-up fee to incentivize attendance.

An illustrative timeline is shown in Figure 2.[7]

4.2 Treatments

At each of the three follow-up visits, drivers were assigned to one of four treatments, or the control group. Specifically, each driver received exactly one assignment from each of the following treatment categories, given in a random order. Within each category, drivers were randomized across treatments. For example, within the financial incentive category, every driver received either the FC1000 or FC2000 treatment, but not both.

1. **Control:**

   - *Control*: the driver was asked to bring a receipt from his next trip through a toll, but no reward was offered.

2. **Financial Incentive:**

---

[7]Note that for a specific driver, since the date of baseline was random, this process could be as long as 13 weeks, if the baseline visit occurred at the end of the baseline period.

- *FC1000*: the driver was asked to bring a receipt from his next trip through a toll and told he would be paid 1000 FC (50% of the toll price) for each receipt that he brought up to a maximum of 2 receipts.

- *FC2000*: the driver was asked to bring a receipt from his next trip through a toll and told he would be paid 2000 FC (100% of the toll price) for each receipt that he brought up to a maximum of 2 receipts.

3. **Social Incentive:**

- *Charity*: the driver was asked to bring a receipt from his next trip through a toll and was told that for each receipt he brought, we (the research group) would donate 2000 FC to a home for widows in the city of Kananga, up to a maximum of 2 receipts.

- *Government*: the driver was asked to bring a receipt from his next trip through a toll and was told that for each receipt he brought, we (the research group) would give 2000 FC to the government of Kananga to, in turn, transmit to a home for widows in the city of Kananga, up to a maximum of 2 receipts.

We discuss the power advantages of this panel design, as well as some of the assumptions required to evaluate it in Section 11.3.1.

In addition to these main treatments, we cross-randomized a social norms intervention at the individual level. For this intervention, individuals' treatment status was constant across rounds. Selected participants were told by enumerators the proportion of drivers who reported paying the full amount at the toll during the baseline survey. The text read as follows: "Now, I'm going to give you an update about information we've learned speaking to motorcycle taxi drivers in Kananga over the past 3 weeks. In particular, did you know that 62 percent of motorcycle taxi drivers paid 2000 FC to the DGRKOC at the toll in their last trip? Is that 62 percent higher, lower, or the same as you would have expected?" Selected participants were given a chance to respond to these questions. This information was repeated at the first two follow-up visits.

The hypothesis we wish to test with this intervention is that individuals participate in petty corruption partly because they believe it is the status quo. Because 62% is likely to be construed as a high rate of tax compliance in this setting — where we estimate true tax compliance at 13% — this information should surprise

8

individuals and potentially move their priors about the prevalence of bribe payment at the toll.[8] Indeed, 80.5% of individuals in the first follow-up survey admitted that this rate of full payment was equal to or higher than what they would have expected. We thus hypothesized that this treatment will increase drivers' efforts to obtain a receipt.

Simultaneously, there were two interventions at the toll-day level. Driver recall of their date of travel was poor, which means that we have trouble evaluating these treatments (since we cannot match a specific trip to whether or not treatment was active on that toll-day), but for completeness, we describe both of them here.

1. *Enumerator presence*: Beginning one month into the study, we assigned an enumerator to spend 4 days per week at the tolls around the city (as there are 4 main tolls and 7 days in a week, the probability of a given toll being monitored on a given day was $\frac{4}{7} * \frac{1}{4} = 14.3\%$).[9] The purpose of the enumerator was to validate that drivers were taking the trips that they claimed to have taken (and were reporting all trips) by doing very short surveys with study participants as they passed the tolls. The enumerator was *not* told to monitor the toll officer, but his presence at the toll was clear and, anecdotally, officers were suspicious that he might report on them.

2. *Financial incentive for toll officer*: Beginning at the same time as the enumerator monitoring, we began offering financial incentives to toll officers to issue more receipts. Specifically, on randomly selected days, we offered the toll officers 500 FC for each receipt they issued over a threshold set at the 25th percentile of *reported* receipts in the pre-period (this excludes a small number of inframarginal receipts). The randomization was set so that an average of every 5th day was treated, but the period between treatments was randomized.[10]

---

[8]This information was presented in terms of drivers reported payments, as reported to us and thus was truthful.

[9]The enumerator did not work on Sundays, so the probability of monitoring varied between Sunday and the other days, but this provides an approximation.

[10]The initial day of treatment was randomized. Then, the pause between each episode was randomized as follows: 3 days (65 percent of the time), 4 days (17 percent), 5 days (12 percent), 6 days (6 percent).

4.3   Experimental Measures

We use experimental measures of honesty, tax morale, and patience of drivers, which are described briefly in this section.

4.3.1   Honesty

We measured honesty using an experimental game administered in the field by the enumerator using a tablet during the recruitment stage. The participant received the following instructions (in French or Tshiluba):

"Now, I will ask you to play a short game to evaluate your chance to win a small amount of money. In one minute, a number will randomly appear on the screen. It will be a number between 1 and 10. Before the number appears on the screen, I will ask you to think of one of those numbers in your head, but not say it out loud. If the number you are thinking about appears on the screen, you will tell me so and I will give you 200 FC. If a different number appears, then tell me that it was different and you will not get a reward."

The enumerators then demonstrated two examples, one in which the participant won and one in which he lost. Participants were asked whether they understood the rules of the game or needed further explanation. If the participant understood the game, then he was asked to imagine a number. Once he had imagined a number, the enumerator would then pass to the next screen, where a randomly selected number would appear. The enumerator would then ask:

"Did you guess the number that appeared on the screen?"

If all participants were honest, we would expect an average of 10% of participants to guess correctly. However, participants could also lie about having guessed the number. In practice, 38.4% of participants claimed to have won the game, suggesting a significant rate of dishonesty.

4.3.2   Tax Morale

We measured tax morale, defined by Luttmer and Singhal (2014) as "nonpecuniary motivations for tax compliance" using an experimental game administered at our office during the baseline visit. This game is similar to the Random Allocation Game as described in Lowes et al. (2017). In the game, the enumerator presented the participant with a six-sided die and told the participant:

"Now, I would like to play a game using this dice. Each side has a different number of dots from 1 to 6. I am going to face away from you. While I am facing away, please roll the dice 4 times in the box, so that only you can see it. When you are done, I will ask you the total number of dots that were on the dice during those 4 rolls. At the end of the survey, I will give you 50 FC for each dot that you rolled during those four rolls and we will give the government 1200 FC (the maximum) minus the amount that we give you."

The enumerators then demonstrated several examples and administered test questions to make sure that the participant had understood the rules. The enumerator then faced away or left the room while the participant completed his rolls. If all participants were completely honest, we would expect the average amount received by a participant to be 700 FC (14 being the average over 4 rolls). In practice, the average was 779 FC, with excess mass in the right tail, as seen in Figure 3.

This game measures the experimentally relevant parameter of "willingness to deprive the government of revenue for personal gain," since participants were informed repeatedly that any money that they did not win from the game would go to the provincial government (also the recipient of toll revenue).[11]

4.3.3   Anticipated Discount Factor

We measured each participant's "anticipated discount factor" using choices over money today versus at the participant's next visit. At the baseline visit and the first two follow-up visits, the participant was given two choices between an amount of money at the current moment and a larger amount of money in the future. Those decisions are summarized in Table 2.

Drivers were told that one of the two decisions would be selected by the tablet to occur in real life, so they were incentivized to think seriously about their preferred option. If they received money at a future visit, that amount was clearly delineated so that they understood that we were, in fact, making good on our promise to deliver the payment.

We use the term "anticipated discount factor" to emphasize that this measure

---

[11]We find much less theft in this sample than in Lowes et al. (2017), working in the same context, but with a different sub-population. There are several reasons why this could be true, but we suspect that the most important difference was that the study population came to our office for this activity, where they may have felt more uncomfortable cheating than at home as in Lowes et al. (2017). This also explains why cheating in the "honesty game" described in Section 4.3.1 was significantly higher, as that game was played outside of our compound, in the field.

11

combines three elements: (1) the participant's true underlying time discount factor (2) the participant's perceived likelihood of returning for his follow-up visit and (3) the participant's trust that we would deliver on our promise in the future (relative to now). The anticipated discount factor is the product of these three factors. We want to emphasize that this measure is precisely the experimentally relevant parameter. Conceptually, there is no difference between this decision and the decision of the driver to demand a receipt at the toll in exchange for a promised reward. In both cases, he is trading off a short-term benefit against a future reward, which he may discount if (1) he values the future little, (2) he anticipates *not* coming to his next visit, or (3) he expects the experimenters not to honor their word. Thus, the decision taken in this game is informative about the driver's valuation of our promises of a reward — even if it is true that this game does not measure patience alone.[12]

Many drivers selected the low return options. Specifically, if we combine the 3 times that drivers made these decisions (at Baseline, at Follow-Up 1, and at Follow-Up 2), we see the matrix of realized choices in Table 3, where the vertical axis is for choice 1 and the horizontal axis is for choice 2.

Many drivers change their choices across rounds, so only 21.4% of drivers always choose the immediate amount (indicating a stable weekly β, bounded above by 0.63) and only 15.5% of drivers always choose to wait (indicating a stable weekly β bounded below by 0.74). Overall, these results suggest that drivers are, on average, highly impatient or present biased. We will return to this fact when we interpret the results in Section **??** below.

5   Theory

5.1   Basic Bargaining

This section sketches a simple bargaining model applicable to this setting to elucidate the mechanism behind the financial incentive treatments. Section 11.1 in the Appendix explores a possible mechanism behind the social incentive treatments.

Consider a driver and a toll agent. The driver (subscript i) receives an individual return from completing the trip π distributed $F(\cdot)$. This value is the driver's return

---

[12]Importantly, we do not seek to distinguish between hyperbolic and non-hyperbolic preferences (unlike most recent work on time preferences, e.g. Andreoni et al. (2015)). This is because in our setting, all rewards (both in this experimental game and decisions about seeking a receipt) are in the future relative to an immediate pay-off or cost, so any hyperbolic factor would always be active. Thus, it is not conceptually important to separately estimate it.

net of all costs (e.g. fuel and motorcycle rental) besides paying at the toll.[13] The driver also has tax morale, χ distributed $H(\cdot)$, which captures his intrinsic valuation of completing a legitimate transaction and obtaining a receipt (not including other incentives), and a trip-specific time cost, ξ distributed $G(\cdot)$ with mean 0. This time cost will only be paid if the driver decides to wait for a receipt, so it represents the driver's trip-specific shock to the value of getting a receipt.

The driver has two potential actions: $\{t, b\}$, where t indicates demanding a receipt and paying the full tax, and b indicates paying a bribe. If the driver pays t, the toll agent (subscript a) receives w, a piece-rate value of reporting the traffic. Otherwise, the driver pays b as a transfer to the toll agent. We assume that the toll agent can always refuse a bribe offer and instead issue a receipt, but cannot refuse to issue a receipt. However, in practice, the driver will never refuse a bribe because, conditional on the driver's action, he is always at least as well off accepting a bribe in place of the tax.[14]

The payoff of the driver if he demands a receipt is given below:

$$V_i(\pi, \xi, t) = \pi - t - \xi + \chi \tag{1}$$

If the driver does not demand a receipt, his payoff is instead:

$$W_i(\pi, \xi, b) = \pi - b \tag{2}$$

The agent has a corresponding set of payoffs. If he issues a receipt, his payoff is:

$$V_a(\pi, \xi, t) = w \tag{3}$$

If he takes a bribe, his payoff is:

$$V_a(\pi, \xi, t) = b \tag{4}$$

The amount b is determined by Nash bargaining with the toll agent, who has bargaining weight δ. This yields the following simple equation for the bribe, where the agent receives his outside option w plus a δ share of the surplus generated by collusion.

---

[13] We assume that the value of $\pi \geq t$ so trip completion is not affected.

[14] Subscripts are largely suppressed in what follows for visual simplicity.

$$b^* = w + \delta \underbrace{(t - \chi - w + \xi)}_{\text{surplus from collusion}} \qquad (5)$$

If the surplus from collusion is negative, the driver will demand a receipt. if it is positive, the driver will instead pay a bribe.

5.1.1 Treatment Effects

We can think of any of the treatments outlined in Section 4.2 as being a shock $k \geq 0$ to the driver's return to getting a receipt. Then the solution to the Nash bargaining problem is a slight adjustment to (5):

$$b^* = w + \delta \underbrace{(t - \chi - w - k + \xi)}_{\text{surplus from collusion}} \qquad (6)$$

Again, any time the surplus from collusion is positive, the driver will still pay a bribe. However, an increase to k will cause the share of drivers who pay bribes to fall (weakly). The share of drivers getting a receipt, $\bar{t}$, will change as follows in response to a change in k:

$$\frac{\partial \bar{t}}{\partial k} = \frac{\partial \bar{t}}{\partial w} = m(t - k - w) \text{ where } m(\cdot) \text{ is the PDF of } \chi_i - \xi_i \qquad (7)$$

To evaluate the effect on the equilibrium average bribe, $\bar{b}$, we can do a similar exercise:

$$\frac{\partial \bar{b}}{\partial k} = -\delta(1 - \underbrace{(t - k - w)m(t - k - w)}_{\text{selection effect}}) \qquad (8)$$

This equation makes clear that there are two off-setting forces: (1) the direct, "causal" effect of increasing k on bribe levels, which *ceteris paribus* decreases the bribe level by δk, and (2) the selection effect of removing the marginal individuals from engaging in bribery. In this model, the marginal individuals are those for whom the value of collusion is already very low — i.e. those individuals already paying low bribes — so removing those individuals causes the average bribe to rise. In Section 7.3 below, we explain our strategy for separating these two effects.

6 Data

Our data are reported from drivers at follow-up visits to our office. At these visits, we ask them about trips they have taken in the period between their last visit and the current visit — the period during which a given treatment applies.

6.1 Outcomes

We focus on two main outcomes. The first outcome is presence of a valid receipt corresponding to a trip taken by the driver. In our context, presence of a valid receipt is evidence of non-participation in corruption. In all but one of the toll stations around Kananga, the toll officers are equipped with electronic receipt issuing machines.[15] These machines record all receipts issued, and the toll officers are responsible for depositing an amount of cash equal to the receipts issued at the conclusion of each reporting period (generally, each week). As a result, if an officer issues a receipt, he cannot steal money associated with that receipt. If no receipt is issued, then there is no tracking mechanism for the money and, according to the office workers who took the reports, toll officers never turned in any money that was not backed by receipts. Only 8.4% of drivers in the control group produced any receipt proving payment at the toll and the average number of receipts (in the control group), conditional on producing any receipt, was 1.14.

Our second outcome is the amount paid as a bribe. There are several ways to measure this:

- *Self-reported amount paid*: we asked drivers how much they paid at the toll. We count this amount as a "bribe amount" if they do not have a receipt — this is important to avoid conflating effects on receipts issued with effects on equilibrium bribe paid, as those effects have different policy implications. Roughly 37% of drivers report paying less than the official rate. Conditional on underpaying, the average reported discount is 62%; the median discount is 50%.

- *Box amount*: for each of the same transactions, we asked drivers to privately record the amount that they had paid on a slip of paper and to put that paper in a sealed cardboard box. The paper slips were clearly marked with an

---

[15]The final toll station uses specialized receipt pads, which we also accepted.

individual-specific ID code, so the drivers knew that the information was iden-
tifiable, but it removes any direct embarrassment/social desirability associated
with admitting to a bribe in the presence of the enumerator.[16]

- *Self-reported "arrangement"*: for each trip past a toll, we also asked drivers
  to report whether they had an "arrangement" with the toll officer. The term
  "arrangement" is used locally to refer to an extralegal agreement with the
  toll officer, including underpayment of the toll or, potentially, an agreement
  to allow the driver to avoid paying on his return trip. This is thus a coarse
  measure of bribery that may be less sensitive than stating an amount paid.

All of the measures are highly correlated. Since the driver's self-report has
the lowest measurement error (since unusual amounts could be discussed with the
enumerator to check that they reflected our preferred definition of cost), it is our
preferred measure. However, we recognize the possibility of social desirability bias
in reporting potentially illegal activity. In Figure 4, we provide a chart of bribes
paid in each round. Note that this intentionally excludes all payments associated
with a receipt, since paying the full amount for a receipt would not be a bribe.

The figure suggests that there may have been some under-reporting of bribes. In
particular, there is a clear shift in mass from people reporting paying 2000 FC (the
legal amount) to lower amounts, starting in round 3, after they have encountered
the research team several times and a stock of trust has been built. However, the
amount of mass that shifts is small and there is no further shift in round 4. This
provides some suggestive evidence that the effect of social desirability may be small:
(a) many drivers are willing to report payments that are clearly illegal (any amount
under 2000 cannot be explained by claiming a lost receipt or the like), (b) the shift
in reporting after repeated contact with our research team is relatively small, and
(c) there is no further shift during the 4th visit, suggesting that few individuals are
marginal with respect to their willingness to reveal illegal behavior. This suggests
(though does not prove) that most people are honestly revealing their amount paid.

It is worth noting the mass of drivers who pay a bribe of 2000 FC. Are these
drivers who wanted to pay the legal toll but were denied a receipt from the officer?
This appears unlikely, since only 9% of drivers report having asked for a receipt but

---

[16]One problem with this measure is that there is a very large mass point at 200 FC, which we
attribute to a lack of attention/innumeracy (we believe that the vast majority of these people meant
to indicate 2000 FC). As a result, we recode all responses of 200 FC as 2000 FC if the individual
stated in his self-report that he paid 2000 FC.

being denied one. Rather, our interpretation is paying a bribe of 2000 FC ensures rapid crossing of the toll, which is valuable to drivers who have, by definition, a higher opportunity cost of time. We provide more evidence consistent with this interpretation in Section **??**.

6.2   Covariates

As we pre-specified in our Pre-Analysis Plan, we use the following individual covariates:

- $HONESTY_i$: a dummy that is 1 if the driver reported winning the "honesty game" (as described in 4.3.1 above) in our recruitment survey and 0 otherwise

- $INITRECEIPT_i$: a dummy that is 1 if the driver brings a receipt to the baseline visit and 0 otherwise

- $OWNBIKE_i$: a dummy that is 1 if the driver reports owning his own bike and 0 otherwise

- $TAXMORALE_i$: value from 4 to 24, as described in section 4.3.2 above

- $EDUCATION_i$: a set of dummies for different education levels (no schooling, primary completed, secondary completed, tertiary completed)

- $INCOME_i$: a measurement of income at baseline[17]

- $EXPERIENCE_i$: years of experience as moto driver

- AGE_i: age in years

We can see in Tables 4 and 5 that, by design, these covariates are balanced across treatments. A combined test of joint significance across all variables and treatments (shown at the bottom of all columns) confirms that there is no systematic difference across treatment and control.

In addition, as pre-specified, we use the following trip-specific covariates:

- $EXEMPT_{ir}$: a dummy that is 1 if the driver reports carrying a document or passenger (e.g. senior government official) who exempted him from paying the toll and is 0 otherwise

---

[17]We use a measure of consumption, amount spent on cellphone airtime, in the past week, as we think it is measured with more accuracy.

- $AVOID_{ir}$: a dummy that is 1 if the driver reports avoiding the toll and is 0 otherwise

- $BOSSMONEY_{ir}$: a dummy that is 1 if the driver reports being given money from his boss to pay the toll and is 0 otherwise

- $ENUMPRESENCE_{ir}$: a dummy that is 1 if the driver passes a toll on a day during which we are validating traffic and 0 otherwise

- $TRIPVALUE_{ir}$: a continuous variable that is equal to the driver's self-reported profitability of the trip

- γ_{toll}: toll fixed effect

- γ_e: enumerator fixed effect

The trip-specific covariates may be problematic because they are post-treatment. As a result, if treatment induces drivers to change their behavior, such as whether or not they evade the toll, then including these covariates will bias our estimates of the effects of treatment. However, a large majority of trips (81%) are passenger directed, which means drivers are unlikely to manipulate these trip features. Moreover, in Table 6, we analyze whether any of the trip-specific covariates are unbalanced across treatments and find no systematic differences.

7 Empirical Specification

7.1 ITT

Our pre-specified specification is a standard intent-to-treat regression run at the individual-by-trip-level. For an individual i on round r taking trip t, we run:

$$Y_{irt} = \beta_0 + \beta_1 FC1000_{ir} + \beta_2 FC2000_{ir} + \beta_3 CHARITY_{ir} \qquad (9)$$
$$+\beta_4 GOVERNMENT_{ir} + \beta_5 TOLLPAY_{irt} + \beta_6 ENUM_{irt} + \gamma_r + \varepsilon_{irt}$$

The six coefficients measure the causal effects of each treatment, as outlined in Section 4.2. We cluster all standard errors at the level of the driver to allow for arbitrary serial correlation within individuals. In some specifications, we add a

18

vector of driver controls, $X_i'\theta$ or a driver fixed effect, γ_i. As noted, we also sometimes include a vector of trip-specific controls, $X_{irt}'\phi$.

As a robustness check, we also report results using an ITT specification at the individual-round level. This specification is robust to the possibility of endogenous trip misreporting, which we discuss in detail in Section 11.5.[18] This alternative specification is similar to equation 9, but it is not possible to look at the two toll-level treatments, $TOLLPAY_{irt}$ and $ENUM_{irt}$, without looking at individual trips.[19] For completeness, we also show this individual-by-round regression below.

7.2 IV

We also show results in an IV framework. Unlike many RCTs, we do not have any compliance problems in terms of receiving assigned treatment: conditional on showing up to a visit, every participant received the correct, randomly assigned treatment. However, since our treatment is a promise of a reward for a given action (bringing a receipt), we have the problem that participants may not be able to *recall* their treatment assignment correctly. Participants who forget their reward assignment (or forget that any reward is possible) are unlikely to respond to treatment. This motivates the following first stage, where $RECALL(T)$ is a dummy for recalling a given treatment, T. Note that since there are 4 treatments, there are four endogenous regressors and thus four first-stage equations.[20]

---

[18]There appears to be imbalance between trips reported between the control and the four treatment groups. We believe that this is a false positive for three reasons. First, there are no systematic differences in misreporting when we compare across the treatment groups that *drivers themselves thought they were in*. Second, social desirability bias is an unconvincing explanation of the pattern of results we observe because the control was an "active" control: drivers were asked to bring receipts, even if there was no financial or social incentive. Finally, this increase in trip-taking relative to the control implies implausibly large back-of-the-envelope trip-cost elasticities of -2.6 to -5.2. Each of these points is examined in detail in Section 11.5.

[19]Since the timing of those treatments was randomized, causal identification of the main treatments is unaffected.

[20]The four endogenous regressors are $RECALL(FC1000)$, $RECALL(FC2000)$, $RECALL(CHARITY)$, and $RECALL(GOVERNMENT)$. Note that there are two conditions under which all four of these endogenous regressors will have a value of 0 (i.e. the respondent does not recall any of them). One is if the respondent explicitly recalls that he is in the control. The other is if the respondent does not remember his treatment status at all. This means that we implicitly assume that the answer "don't know" is equivalent to recalling that one is part of the control, rather than representing some probability distribution over the treatments. Given the experimental design, this is the most natural assumption.

$$RECALL(T)_{irt} = \beta_0 + \beta_1 FC1000_{ir} + \beta_2 FC2000_{ir} + \beta_3 CHARITY_{ir} \quad (10)$$
$$+\beta_4 GOVERNMENT_{ir} + \gamma_r + \varepsilon_{irt}$$

The exclusion restriction for this regression is that treatment assignment only affects outcomes through its effect on what participants believe their treatment to be. *A priori* this is likely to hold: it is difficult to imagine that assignment to a certain treatment would have effects unless it was remembered by the participant. The main source of violations to this assumption would be if treatment assignment was remembered by the participant initially, thereby affecting his actions, but then forgotten prior to the follow-up interview. We believe that this is unlikely, as we think that participants who change their actions in response to treatment are likely to recall that treatment during the interview. We paid participants an additional 100 FC if they correctly recalled their treatment status to attenuate cheap talk problems.

In Table 7 we show the matrix of recalled treatment against treatment assignment. Correct recollections lie in the main diagonal. From this matrix, it is easy to see that treatment recall was highly imperfect: only 40.7% of participants correctly recall treatment. Nonetheless, treatment assignment is clearly predictive of recalled treatment. In Table 8, we show this more formally by running the first stage separately for each outcome and we see that there is a very strong first stage across all treatments.

One clear pattern of note is that many people in the control group "recall" having been assigned to a treatment category other than control. There are two potential reasons for this. First, some participants may have been attempting to "game" the system and fool the enumerator into believing that they are part of a different group. Based on enumerator reports, we suspect drivers did attempt this strategy occasionally in spite of the financial incentive to answer honestly. Second, we learned that there was confusion about the reassignment of people in a treatment condition in round 1 to the control in round 2. In particular, some participants did not understand that they were losing their reward offer and switching to a new condition in which they were asked to bring a receipt, but would not receive any financial compensation, nor would we donate to any social cause on their behalf.[21]

---

[21]When such confusion became evident after starting the second treatment round of the study,

In this case, we would have a set of individuals who are genuinely confused about their treatment status. Note that both of the problems should be fixed by the IV strategy: the IV identifies the LATE *precisely for the population who is induced into believing they are treated by the instrument*, that is to say, by treatment assignment.

7.3 Estimating Causal Effects on Bribes Levels

Estimating the causal effect of treatment on bribes paid is challenging for two reasons. First, as discussed in Section 6.1, bribe levels may be misreported; however, as noted earlier, we do not believe that this is a major issue in our setting.

Second, since treatment affects whether or not a receipt is issued, then as outlined in Section 5.1.1, treatment will affect both the pool of bribe payers (selection effect) as well as having a direct effect on the amount that they pay (causal treatment effect). We can estimate the total effect with a naive regression using amount paid (conditional on not having a receipt) as the outcome. To separate the effects, we use the following approaches.

- Fixed effects regression

    - If we refer back to Section 5.1.1, there are two components that produce selection, individual tax morale χ_i and trip-specific time-cost ξ_{it}. Since tax morale is fixed within individuals, then *if there were no trip-specific shocks*, the use of individual fixed effects would eliminate the selection effect. Assume, for the moment, that there are no trip-specific shocks. In this case, if we observe two bribery incidents for the same driver (under different treatments), we can interpret the difference between the bribes paid under each condition as the causal effect of treatment. Any individual factor, such as tax morale, is netted out by the within-individual comparison.

    - In the presence of trip-specific shocks, this result will be biased upwards in proportion to the magnitude of $Var(\xi_{it})$, as discussed in Section 5.1.1.

- Tobit

    - Again, in Section 5.1.1, if we take the structure of the model seriously, then a driver chooses to demand a receipt if and only if the latent surplus

---

we noted this in an addendum to our PAP uploaded during the study.

from bribery is negative.[22] Since the equilibrium bribe is w (the value of issuing a receipt to the toll officer) plus a share of the surplus, then whenever we observe a receipt it reflects a latent bribe that is left-censored at w. To see this, note that the minimum bribe is $w + \delta(0)$. We can use a Tobit model to account for this censoring. We set the "bribe" paid by the driver to be w for all cases in which the driver received a receipt.

– One issue is that w is *ex ante* unknown. However, we can estimate w by looking at the lowest bribes reported in the data, which, in the model, must be equal to w. In Section 6.1, we can see that roughly 15% of drivers report a bribe of 0 (but no receipt), so we set $w = 0$. That toll officers receive no utility from issuing receipts aligns with the fact that they are paid by monthly salary, irrespective of how many receipts they issue, i.e. how much money they deposit in the state coffers.

8   Results

8.1   Main Results

We now turn to the main results. First, we estimate equation 9, which is on the individual-by-round level. The regressions in Table 9 use as the dependent variable a dummy for bringing a valid receipt, while the regressions in Table 10 use a count variable for the total number of receipts reported by the driver in a given round.[23] Both tables suggest that the $FC1000$ and $FC2000$ treatments had a small but significant effects on the probability that drivers brought receipts. In other words, the financial incentives seemed to have induced a small subset of drivers to abstain from corruption at the toll and demand that a legitimate transaction take place. According to OLS estimates, drivers in the financial incentive treatments were about 4 percentage points more likely to have brought receipts to an interview. Specifications including controls (column 2) and individual fixed effects (column 3) look quite similar. According to IV estimates, which isolate the effect on drivers who remembered their treatments, the effect is larger: 10 percentage points for the $FC1000$ treatment and $FC2000$ treatment.

---

[22]If we wish to take the theory less seriously, the ideal solution would be to use a Heckman selection correction. However, finding an instrument that predicts selection into bribery, but does not (separately) affect bribe levels, is very difficult in a bargaining set up. We nonetheless explore available options in Appendix Section 11.6.

[23]As noted above, the data are pooled across rounds of the experiment.

In contrast, the social incentive treatments have no effect on average. Promising donations directly or via the government does not motivate the average driver to negotiate for a receipt at Kananga's tolls. This null average result suggests that perceptions about the low social value of payments to the state do not explain citizen supply of bribes. Even when their money would go toward spending on public goods with higher probability, drivers were not more likely to request that a legitimate transaction take place. However, when we examine participants who are children of widows, a pre-specified sub-group for which we expect a larger effect of the donation treatments, these results change. Treatment effects are larger and marginally significant for this subgroup. As a result, when we estimate ϕ (the perception of government corruption) as in Section 11.1, we estimate a confidence interval that covers all values between 0 and 1. In addition, the social norms treatment had no effect. Drivers do not appear to participate in corruption simply because they think everyone else is doing the same. The null effect for the social norms treatment remains true in all estimations that follow, so we drop that coefficient for simplicity.

Tables 11 and 12 consider results on the individual-by-trip level. The number of observations decreases relative to the individual-by-round analysis because many drivers reported taking no trips in the roughly three week period between office visits. The dependent variable is whether the driver brought a receipt corresponding to his (reported) trip. As noted in Section 11.5, these results should be interpreted with caution, due to the issues with misreporting.

In Table 11, we can effectively sign the bias and so we should think of the results shown there as upper-bounds on the true effects. Having noted that caveat, the financial incentive treatments caused a 7 percentage point increase in receipt-bringing relative to the control. As in the individual-by-round analysis, the estimated effect is essentially the same using individual controls, trip controls, and individual fixed effects. The social incentive treatments do not have a consistent effect across specifications. As indicated in equation 9, these regressions also include (i) a dummy ($ENUM_{irt}$) that equals 1 if on a given trip there was an enumerator validating traffic at the toll, and (ii) a dummy ($TOLLPAY_{irt}$) that equals 1 if on a given trip the toll officer incentive was available at the relevant toll. The coefficient on $ENUM_{irt}$ is for the most part positive, but also never significant.[24]

---

[24]Drivers often had difficulty recalling dates of travel. This makes it complicated to match the toll-level treatments to specific trips and thus biases the treatment effects on both of these treatments to zero.

Next we turn to the results on equilibrium bribes. In these regressions, the dependent variable is the reported amount paid at the toll. Results are essentially identical if we use the box measure.[25] The first three columns of Table 12 show estimations of equation 9 with no covariates, with individual covariates, and with both individual and trip-level covariates. In the first two columns, the coefficients on the main treatments are negative, but never statistically significant.

Then, in Table 13, we attempt to establish the pure causal effect of treatment on equilibrium bribes, as described in Section 7.3. The results vary across methods. The fixed effect regression produces the peculiar result that the two different rewards have divergent effects on bribe amounts. One rationalization of this result is an extensive form bargaining game where toll officers can use a costly hassling technology to separate high- and low-time-cost drivers.[26]

The Tobit methodology suggests large negative effects on bribe payment. This largely results from the fact that, in the Tobit, we assume all cases with receipts have a bribe of 0 or less (left-censored at 0) and, as seen in Table 11, we know that receipts-per-reported-trip rise significantly. As a result, the model estimates a negative causal effect on bribes. This is consistent with the model sketched above.[27]

8.2 First-degree price discrimination?

This section probes the bribe results further by considering to what extent driver and trip covariates predict the magnitude of bribes paid. We find weak evidence that toll officers engage in first-degree price discrimination.

In Table 14, we show how the amount that a driver pays at the toll changes with respect to a proxy for driver consumption/wealth (amount spent on phone credit), the driver's estimate of the revenue he earned from the trip, and the two honesty/tax

---

[25]Note that these regressions are on the individual-by-trip level; it is not possible to recreate the bribery analysis on the individual-by-round analysis.

[26]The flavor of this model would be as follows. Toll officers are not able to observe the reward promised to the driver and must pay a cost if they wish impose a time hassle on the driver. In equilibrium, officers mix over hassling or not. Drivers have either low or high time cost, implying that the penalty of being hassled is low or high. Conditional on being hassled, drivers can either pay their time cost or "surrender" and pay the maximum bribe (2000 FC). All high time cost drivers will surrender, conditional on being hassled. The divergent result for the low and high reward arises as follows. For a small reward, only low cost drivers demand a receipt. Whereas for a high reward, there is a pooling equilibrium where many drivers gamble on asking for a receipt, but the high cost drivers will surrender and pay a high bribe if they are hassled. This generates higher bribes for the 2000 FC reward group, but lower bribes in the 1000 FC reward group.

[27]In the Heckman model considered in Appendix Section 11.6, these effects disappear, and the data appear quite noisy.

morale measures from Section 4.3. Column (1) shows the results for the sample of transactions where the driver does not have a receipt, showing that reported bribe levels do respond to these measures. This offers some evidence that toll officers engage in first-degree price discrimination. However, we should not exaggerate this claim. The coefficient on driver revenue is small and only marginally significant. Moreover, column (2) limits the sample to transactions in which the driver gets a receipt, showing that for drivers who "exit" bribery, they face a flat cost, even when they are wealthier or their trip is more valuable. Thus, although toll officers may have some ability to set different prices of bribes for different drivers, the results are far from perfect price discrimination by a monopolist.

In addition, as predicted by the theory in Section 5.1, tax morale weakly predicts bribe magnitude among drivers who do not have receipts. In particular, drivers who won more in the tax morale game (discussed in Section 4.3.2) and thus have *lower* tax morale pay higher bribes. Intuitively, individuals who intrinsically value paying official taxes to the government must be compensated with a lower bribe, or they will select out of the bargaining process, pay the full tax, and demand a receipt. Conversely, those with low tax morale pay relatively higher bribes. As expected, this effect disappears for those who get receipts.

Overall, our results indicate that the elasticity of corruption with respect to monetary incentives is negative, but relatively small. We first show that this result is unchanged when we calculate the full elasticities, but that the elasticity is sensitive to the driver discount rate, which might have important implications for anti-corruption policy. Second, we argue that drivers are fairly inelastic in their supply of corruption in this setting because bribes increase the efficiency of toll transactions. In other words, bribes appear to function as an "efficient grease."

8.3 Elasticity of Corruption

We can estimate the precise elasticity of supply for bribes per trip with respect to an incentive in the following way:

$$\epsilon = \frac{\Delta\% \text{ receipts-per-trip}}{\Delta\% \text{ monetary cost of compliance}} \tag{11}$$

The numerator can be written as follows. Let c index the result in the control and m index the result under monetary incentive M.

$$\Delta\% \text{ receipts per trip} = \frac{\frac{\text{Receipts}_m}{\text{Trips}_m} - \frac{\text{Receipts}_c}{\text{Trips}_c}}{\frac{\text{Receipts}_c}{\text{Trips}_c}} \quad (12)$$

Receipts can be calculated directly from the estimates of the total increase in receipts as seen in Figure 10, over the control average.[28] We use the estimates from column (1), the baseline ITT specification. However, to calculate the true number of trips, we need to account for the potential misreporting discussed in Section 11.5, which we show in equation (13). Let T_x represent *reported* total trips under condition x (either treatment or control) and let U_x be the under-reporting rate (i.e. the percent of trips recorded by the auditor, but not reported in interviews) under condition x.

$$\text{Trips}_x = \frac{T_x}{1 - U_x} \quad (13)$$

We can use estimates of T_c and T_M from Figure 27 (column (2)). For the estimates immediately below, we assume that U_c and U_M are the same, per our logic in Section 11.5, but we use estimates of U_c and U_M from Figure 28 (column (1)) in our additional analysis in Section 11.5.1 in the Appendix.

Finally, the denominator is calculated by discounting the potential reward faced by a given driver by the amount of time between his trip and his appointment.[29] We can see this formally in equation (14) below:

$$\Delta\% \text{ monetary cost of compliance} = -\frac{M * \beta^w}{2000} \quad (14)$$

To estimate this value, we need to know β, the driver discount rate, and w, the time between when the driver passes the toll (i.e. when he makes the decision

---

[28]Technically, we will use the constant in the regression, which is the control average, net of round fixed effects. We will generally use the constant instead of the control mean throughout for this reason.

[29]Note that here we are using the "list price" of compliance (i.e. imagining changing the price of 2000 FC) as our baseline cost measure, not the "marginal financial cost" of compliance. The "marginal financial cost" of compliance is [List price - Bribe], where the bribe that a driver would have faced is unobserved in any instance where he takes up a receipt. We focus on the "list price" both for ease of exposition and because it is a policy-relevant: our results speak to the effect of lowering the list price, which is directly in the government's control. Since the marginal financial cost of compliance is bounded above by the list price (and can be as low as 0 for people currently paying a bribe of 2000 FC), looking at the marginal financial cost of compliance would cause us to estimate that drivers are *even more* inelastic than we already estimate them to be.

to participate in corruption or not) and his interview.[30] w is known to the driver because at the time of his decision to seek a receipt, he knows the timing of his next appointment. w represents the number of weeks between a driver's trip and his next scheduled appointment, which we can calculate by subtracting the date of the driver's next appointment from his reported date of travel.[31] β is the driver's weekly discount factor, which we estimated using the methods described in Section 4.3.3. This discount factor includes both literal time preference, as well as any other factors that affect the driver's belief that he will receive a future payoff (e.g. trust that the experimenters will honor payments in the future), which is the empirically relevant discount factor.

One issue here is that we do not have a precise estimate of driver β, due to the small number of time preference decisions that we offered. As a result, if a driver was always impatient, then all we know is that his $\beta \leq 0.63$, while a driver who was always patient has $\beta \geq 0.74$. Meanwhile, drivers that sometimes were patient and sometimes were impatient have intermediate discount factors. The median driver took a patient decision 1 in 3 times (twice out of 6 possible opportunities). This suggests a median β of around 0.63 or slightly below, but we will show sensitivity to changing this parameter for our estimates in Table 29.[32] We then calculate the average β^w over the entire sample (w and β are uncorrelated with treatment) for the different values of β that we consider.

Table 15 shows our estimated elasticities for the two treatment dummies and separately for a specification where the $FC1000$ and $FC2000$ are combined linearly.[33] We bootstrap the standard errors using 5,000 draws. Our preferred estimates are in the last column, which combines the two monetary rewards into a single estimate. These estimates suggest that on average, the elasticity is negative, but close to 0. Our 95% confidence interval allows us to reject elasticities larger (more negative) than -0.95 for a low β or smaller (less negative) than -0.45 with no discounting. Citizen supply of bribes is thus relatively inelastic in this setting.

---

[30]M is the (known) value of the incentive (either 1000 FC or 2000 FC).

[31]Some drivers missed their appointments by large margins, creating a negative w. However, this is more plausibly interpreted as a very large w, since most of those drivers did not come to the office until they were sought out by our enumerators and thus may not have been considering collecting the reward at all. We drop any w past -2 days (and replace those in $[-2, 0]$ with the empirical time until appointment), though we believe this still underestimates the "true" w.

[32]β elicitation was always done prior to treatment assignment, so, by design, there is no relationship between period-wise β and treatment assignment in that period.

[33]We construct a variable that equals 1000 in the $FC1000$ group, 2000 in the $FC2000$ group, and 0 for the control group, charity treatment, and government treatment.

The table also shows the importance of time preferences in the elasticity of corruption with respect to financial incentives. Since driver discount rates are very high on average, the present value of the reward they face for compliance with the toll is much lower than the face value of the reward. This finding is policy relevant. Even if the value of abstaining from corruption is high — but that value arrives in the future — high rates of discounting may perpetuate citizen supply of bribes. Liquidity constraints and other factors contributing to discounting in developing countries may thus indirectly fuel citizens' participation in corruption.

8.4 Bribes as an "efficient grease"

The high discount rates we observe among drivers suggest the most compelling explanation for the limited responsiveness of drivers to monetary incentives: bribe payment may increase the efficiency of transactions in the toll setting we study. We consider two types of evidence supporting this interpretation: (1) reported duration of toll transactions when drivers demand receipts versus when they do not demand receipts, and (2) differential treatment effects by driver time costs.

The clearest evidence that bribes increase efficiency in this toll setting comes from drivers' reports about typical transactions when they do and do not demand a receipt. Obtaining a receipt is slow due to bureaucratic procedures and the extent to which toll officers deliberately delay the process to try to extract bribes. Although the toll officers wait on the road in front of the toll, they leave their handheld receipt printers in an office located a short distance from the road. To get a receipt, drivers must park their motorbikes and walk down to this office with the toll officer. The officer manually enters on the receipt printer the name of the motard, the type of vehicle, the chassis number, and the name of the agent. Then, he handwrites these same pieces of information in a large ledger. He might ask to see their driver's license and motorbike registration documents, too. In total, this process takes an estimated 15 minutes on average. It can take considerably longer if there are many other vehicles at the toll — whom the officer will likely see first if they do not require receipts — or if there is a problem with any of a driver's documents.

On the other hand, drivers can speed up the interaction considerably by paying the toll money and not asking for the receipt. They do not even half to dismount their motorbike in this case. "To pass through the toll rapidly," one enumerator recalled, "motards prepare their money in advance, and ... then they don't ask

for the receipt."[34] Drivers often do not ask for a receipt even after paying the full 2000 FC to speed up the interaction. To avoid confusion, drivers and officers have a code system that communicates their intent to pay a quick bribe without explicitly saying so. "Brother, here is your coffee," is how motards frequently communicate their intent to pay a bribe.[35] According to enumerators' estimates, paying a bribe cuts the time of the toll transaction by nearly 70% (an average of 11.25 minutes). This large difference in the time cost of passing through the toll reinforces the notion that bribes may increase efficiency of toll transactions. Moreover, given that drivers have large observed discount rates, this time savings is likely very consequential in their decisions to ask for a receipt.

Table 16 shows results of the trip-level regression (equation 9) of treatment on whether or not the driver has an associated receipt in different samples.[36] In column (1) we show our baseline results as reported in Table 11. In columns (2) and (3), we limit our sample to trips during which the driver reported having a passenger and not having a passenger, respectively.

Responsiveness to the financial incentives is more than twice as large when drivers do not have a passenger. This finding is consistent with drivers' reports that they are more time constrained when carrying passengers and thus less willing to wait for a receipt.

In sum, the evidence suggests that citizen supply of corruption may be inelastic in this setting because bribes substantially reduce the time costs associated with a toll transaction. Bribes are quite literally "speed money" Bardhan (1997). As noted in the Introduction, speeding up bureaucratic procedures is one of the principal reasons identified in the literature why corruption can in certain circumstances "grease the wheels" and improve efficiency (Leys, 1965; Lui, 1985). Instead of absorbing the time cost associated with bureaucratic red tape at the toll, a bribe enables drivers to pass quickly, thereby facilitating the flow of goods and people in and out of Kananga.

---

[34]Enumerator Interview, August 3, 2017, Kananga.

[35]Enumerator Interview, August 3, 2017, Kananga. Bardhan (1997) notes that it is common to have specific words for "speed money," noting examples from the Philippines in particular. "Coffee" and "beer" are two such examples from Kananga.

[36]Due to the trip non-reporting issue discussed earlier, the results here are caveated by the aforementioned problems related to trip non-reporting. Nevertheless, since all of the columns use the same specification, they contain valuable information in reference to one another.

9 Conclusion

This experiment provides three key contributions to the corruption literature. First, it is one of the first studies to experimentally vary the returns to corruption in the field for citizens (as opposed to bureaucrats) and provides a template for future research in this area. Second, it credibly estimates the elasticity of tax compliance with respect to price. Third, it offers evidence that citizens can be inelastic in their supply of corruption in settings in which bribes increase the speed of official transactions. It thus provides experimental evidence in support of the "grease the wheels" hypothesis about corruption and economic efficiency.

References

Andreoni, James, Michael A Kuhn, and Charles Sprenger, "Measuring time preferences: A comparison of experimental methods ," *Journal of Economic Behavior and Organization*, 2015, *116*, 451–464.

Banerjee, Abhijit V, Shawn Cole, Esther Duflo, and Leigh Linden, "Remedying Education: Evidence From Two Randomized Experiments in India," *Quarterly Journal of Economics*, 2007, *122* (3), 1235–1264.

Bardhan, Pranab, "Corruption and Development : A Review of Issues," *Journal of Economic Literature*, 1997, *35* (3), 1320–1346.

Beck, Paul J. and Michael W. Maher, "A comparison of bribery and bidding in thin markets," *Economics Letters*, 1986, *20* (1), 1–5.

Bertrand, Marianne, Simeon Djankov, Rema Hanna, and Sendhil Mullainathan, "Obtaining a Driver's License in India: An Experimental Approach to Studying Corruption," *Quarterly Journal of Economics*, 2007, (November), 1639–1676.

Dhaliwal, Iqbal and Rema Hanna, "Deal with the Devil: The Successes and Limitations of Bureaucratic Reform in India," 2013.

Dreher, Axel and Martin Gassebner, "Greasing the wheels? The impact of regulations and corruption on firm entry," *Public Choice*, 2011, *155* (3-4), 413–432.

Duflo, Esther, Rema Hanna, and Stephen P. Ryan, "Incentives work: Getting teachers to come to school," *American Economic Review*, 2012, *102* (4), 1241–1278.

Ferraz, Claudio, Frederico Finan, and Diana B Moreira, "Corrupting learning Evidence from missing federal education funds in Brazil," *Journal of Public Economics*, 2012, *96* (9-10), 712–726.

Fisman, Raymond and Jakob Svensson, "Are corruption and taxation really harmful to growth? Firm level evidence," *Journal of Development Economics*, 2007, *83* (1), 63–75.

Huntington, Samuel, *Political Order in Changing Societies*, New Haven, CT: Yale University Press, 1968.

Kato, Atsushi and Takahiro Sato, "Greasing the wheels? The effect of corruption in regulated manufacturing sectors of India," *Canadian Journal of Development Studies/Revue canadienne d'études du développement*, 2015, *36* (4), 459–483.

Kaufmann, D and SJ Wei, "Does "grease money" speed up the wheels of commerce?," 1999, (8209).

Khan, Adnan, Asim Khwaja, and Benjamin Olken, "Tax Farming Redux: Experimental Evidence on Performance Pay for Tax Collectors," *The Quarterly Journal of Economics*, 2016, (August).

Leff, Nathaniel H, "Economic development through bureaucratic corruption," *American behavioral scientist*, 1964, *8* (3), 8–14.

Leys, Colin, "What is the Problem about Corruption?," *The Journal of Modern African Studies*, 1965, *3* (2), 215–230.

Lien, Da-Hsiang Donald, "A note on competitive bribery games," *Economics Letters*, 1986, *22* (4), 337–341.

Lowes, Sara, Nathan Nunn, James A Robinson, and Jonathan L Weigel, "The evolution of culture and institutions: Evidence from the Kuba Kingdom," *Econometrica*, 2017, *85* (4), 1065–1091.

Lui, Francis T, "An equilibrium queuing model of bribery," *Journal of political economy*, 1985, *93* (4), 760–781.

Luttmer, Erzo F P and Monica Singhal, "Tax Morale," *Journal of Economic Perspectives*, 2014, *28* (4), 149–168.

Mauro, Paolo, "Corruption and Growth," *Quarterly Journal of Economics*, 1995, *110* (3), 681–712.

Méon, Pierre Guillaume and Khalid Sekkat, "Does corruption grease or sand the wheels of growth?," *Public Choice*, 2005, *122* (1-2), 69–97.

Méon, Pierre-Guillaume and Laurent Weill, "Is Corruption an Efficient Grease?," *World Development*, 2010, *38* (3), 244–259.

Naritomi, Joana, "Consumers as Tax Auditors," 2015.

Olken, Benjamin A., "Monitoring Corruption: Evidence from a Field Experiment in Indonesia," *Journal of Political Economy*, 2007, *115* (2), 200–249.

_ **and Rohini Pande**, "Corruption in Developing Countries," *Annual Review of Economics*, 2012, *4* (1), 479–509.

Pomeranz, Dina, "No Taxation without Information: Deterrence and Self-Enforcement in the Value Added Tax," *American Economic Review*, 2015, *105* (8), 2539–2569.

Vial, Virginie and Julien Hanoteau, "Corruption, manufacturing plant growth, and the Asian paradox: Indonesian evidence," *World Development*, 2010, *38* (5), 693–705.

10 Tables and Figures

10.1 Figures

Notes: Treatment effects are estimated in a single regression using equation 9, so results include round fixed effects. The dependent variable is whether the driver received any good receipts fort the trip. The constant was excluded, so all results can be compared relative to the control group. Standard errors are clustered at the individual level. The confidence interval shown is a 95% confidence interval.

Figure 1: Effects of treatment on trips with receipts (ITT)

Notes: Illustrative timeline for a participant who was recruited on the first day of recruitment and requested to arrive for his baseline interview on the first day of the baseline interview period. Assigned timelines could be as much as three weeks longer than this timeline by assigning a baseline interview in week 4 instead of week 1. Gaps between follow-up interviews were always three weeks.

Figure 2: Illustrative timeline of the experiment
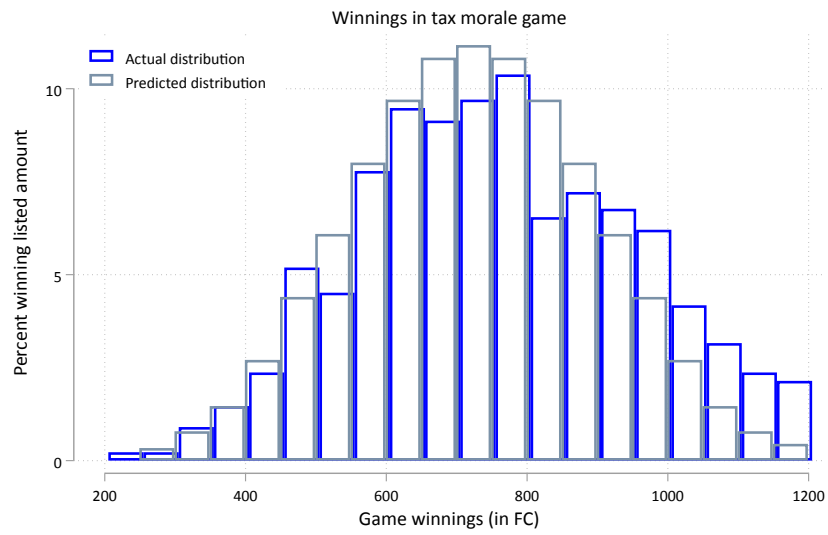
Notes: Dark blue indicates the empirical histogram of game winnings in the tax morale game described in Section 4.3.2. The pale blue indicates the predicted distribution that should arise by chance on average if there was no cheating.

Figure 3: Distribution of realized tax morale game winnings against the predicted distribution

Distribution of reported bribes by visit number

Notes: Empirical histogram of self-reported bribe payments (people with valid receipts are excluded from this graph). The histogram for Visit 3 is off center by 10 FC relative to Visit 2 and the histogram for Visit 4 is off center by 20 FC relative to Visit 2. Each bin is 100 FC wide.

Figure 4: Bribes by round (each round is jittered by 10 FC for visibility)

10.2 Tables

	(1)	(2)
VARIABLES	In Study	Out of Study
How old were you at your last birthday?	27.97	26.92
	(5.855)	(6.776)
Secondary or tertiary education	0.975	0.983
	(0.157)	(0.128)
Owns bike	0.420	0.380
	(0.494)	(0.486)
Weekly phone credit spend (in FC)	3,247	2,816
	(3,566)	(2,743)
Years as motard	5.097	4.713
	(3.029)	(2.933)
Trust in foreign researchers (1-4)	3.508	3.463
	(0.855)	(0.893)
Won dishonesty game	0.372	0.430
	(0.484)	(0.496)
Observations	866	300

Notes: This table shows summary statistics for our sample. Column (1) shows summary statistics for the population of drivers who did join the study. Column (2) shows summary statistics for drivers who were invited to join the study (based on the recruitment survey), but did not join.

Table 1: Summary statistics

Choice	Immediate Amount	Amount at Next Visit	Weekly Discount Factor (β) for Indifference
1	200 FC	500 FC	0.74
2	100 FC	400 FC	0.63

Notes: This table shows the trade-offs faced faced by motards in the two different versions of the discount factor game. The first line shows one of the choices and the second line shows the other choice. The columns show, respectively, the amount received if demanded immediately, the amount received if he waited until the following visit, and minimum weekly discount factor that would justify choosing to wait for the higher amount.

Table 2: Choices over amounts

	Chose to Wait	Chose Immediate Amount
Chose to Wait	31.2%	10.5% FC
Chose Immediate Amount	13.5% FC	44.8% FC

Notes: This table shows the decisions made by the motards in the game. The unit of analysis is an interview round, so the results in each cell show the percentage of all visits that fell into that cell. The first row indicates that the motard chose to wait for 500 FC instead of taking 200 FC immediately. The opposite is true in the second row. The first column indicates that the motard chose to wait for 400 FC instead of taking 100 FC immediately instead of taking 100 FC immediately. The opposite is true in the second column. Since all participants faced both decisions at each visit, each visit can be allocated to one of the four resulting cells.

Table 3: Empirical distribution of choices over all rounds

VARIABLES	(1) Dishonesty game	(2) Receipt at baseline	(3) Owns bike	(4) Dice game score
FC1000	0.02	-0.02**	0.04**	0.16
	(0.02)	(0.01)	(0.02)	(0.14)
FC2000	-0.01	0.02**	-0.02	-0.11
	(0.02)	(0.01)	(0.02)	(0.14)
Charity	-0.00	0.00	0.03*	0.15
	(0.02)	(0.01)	(0.02)	(0.14)
Government	0.00	-0.01	-0.03*	-0.22
	(0.02)	(0.01)	(0.02)	(0.14)
Constant	0.37***	0.04***	0.40***	15.78***
	(0.02)	(0.01)	(0.02)	(0.19)
Observations	2,414	2,484	2,481	2,475
R-squared	0.00	0.00	0.01	0.00
Ind. FE	N	N	N	N
Ind. Controls	N	N	N	N
Control avg.	0.374	0.0487	0.430	15.64
Joint p-value	0.205	0.205	0.205	0.205

Notes: This is a table testing balance across treatment categories. It is estimated using using equation (**??**), so all results include round fixed effects (not reported). The dependent variable is noted at the top of each column. The p-value from a test of joint significance of all regression coefficients across all variables tested is listed at the bottom of each column. Standard errors clustered by individual. $^*p < 0.1, ^{**}p < 0.05, ^{***}p < 0.01$.

Table 4: Test of balance across individual-level controls.

VARIABLES	(1) Education	(2) Weekly airtime spend (FC)	(3) Age	(4) Experience
FC1000	0.01	82.23	0.07	0.00
	(0.02)	(218.38)	(0.36)	(0.19)
FC2000	-0.01	-51.07	0.01	-0.03
	(0.02)	(216.27)	(0.36)	(0.19)
Charity	0.02	-88.38	0.08	0.04
	(0.02)	(215.74)	(0.36)	(0.19)
Government	-0.03	135.98	0.02	0.04
	(0.02)	(218.68)	(0.36)	(0.19)
Constant	3.04***	3,265.73***	28.01***	5.19***
	(0.02)	(176.44)	(0.29)	(0.15)
Observations	2,481	2,481	2,481	2,481
R-squared	0.00	0.00	0.00	0.00
Ind. FE	N	N	N	N
Ind. Controls	N	N	N	N
Control avg.	3.061	3284	28.08	5.255
Joint p-value	0.205	0.205	0.205	0.205

Notes: This is a table testing balance across treatment categories. It is estimated using using equation (**??**), so all results include round fixed effects (not reported). The dependent variable is noted at the top of each column. The p-value from a test of joint significance of all regression coefficients across all variables tested is listed at the bottom of each column. Standard errors clustered by individual. $^*p < 0.1,^{**}p < 0.05,^{***}p < 0.01$.

Table 5: Test of balance across individual-level controls (continued).

Variable category	Test of joint significance p-value
Toll chosen	.55
Enumerator at interview	.78
Exempt from toll	.67
Trip revenue	.44
Toll allowance from boss	.18

Notes: This table tests joint signficance of all treatments within a variable or variable family (for categoric variables). Results are estimated using equation (9). Standard errors are clustered by individual.

Table 6: Test of balance across trip-level controls.

Assigned treatment	Remembered treatment							
	Control	1000FC	2000FC	Charity	Gov.	Other	DK	Total
Control	17	13	16	12	5	8	30	100
1000FC	3	58	7	6	2	8	17	100
2000FC	3	2	68	5	2	9	11	100
Charity	4	5	9	57	0	8	17	100
Gov.	6	7	14	9	30	10	24	100
Total	8	16	22	17	7	8	22	100

Notes: This table shows the treatments recalled by participants (columns) for each treatment assigned to participants (rows). All values are in percent terms, such that the total for each treatment assignment adds up to 100% when adding across all of the columns (i.e. each cell shows the percent of people who recall the treatment listed in that column, among those that are assigned the treatment listed in that row).

Table 7: Treatment remembered versus treatment assigned

VARIABLES	(1) Recall FC1000	(2) Recall FC2000	(3) Recall Charity	(4) Recall Government
FC1000	0.46***	-0.09***	-0.06***	-0.03***
	(0.03)	(0.02)	(0.02)	(0.01)
FC2000	-0.10***	0.52***	-0.07***	-0.03***
	(0.01)	(0.03)	(0.02)	(0.01)
Charity	-0.07***	-0.07***	0.44***	-0.04***
	(0.02)	(0.02)	(0.03)	(0.01)
Government	-0.05***	-0.02	-0.03*	0.25***
	(0.02)	(0.02)	(0.02)	(0.02)
Observations	2,487	2,487	2,487	2,487
R-squared	0.28	0.27	0.25	0.17
Ind. FE	N	N	N	N
Ind. Controls	N	N	N	N
Control avg.	0.125	0.159	0.122	0.0462

Table 8: IV first stage (each column is a regression with the outcome of recalling the treatment named at the top of the column)

	(1)	(2)	(3)	(4)	(5)
VARIABLES	OLS	OLS	FE	IV	IV FE
FC1000	0.03*	0.04**	0.05**	0.10*	0.14**
	(0.02)	(0.02)	(0.02)	(0.05)	(0.06)
FC2000	0.05***	0.05***	0.04**	0.12**	0.10**
	(0.02)	(0.02)	(0.02)	(0.05)	(0.05)
Charity	-0.00	-0.00	0.01	0.03	0.06
	(0.02)	(0.02)	(0.02)	(0.05)	(0.05)
Government	0.00	-0.00	-0.00	0.04	0.02
	(0.02)	(0.02)	(0.02)	(0.08)	(0.09)
Observations	2,487	2,402	2,467	2,487	2,467
R-squared	0.01	0.03	0.42	0.00	0.42
Ind. FE	N	N	Y	N	Y
Ind. Controls	N	Y	N	N	N
Control avg.	0.0839	0.0818	0.0833	0.0833	0.0833

Table 9: Main effects on bringing a receipt to the follow-up visit (unconditional on reporting haven taken a trip in the intervening period)

	(1)	(2)	(3)	(4)	(5)
VARIABLES	OLS	OLS	FE	IV	IV FE
FC1000	0.04*	0.04*	0.05**	0.11*	0.16**
	(0.02)	(0.02)	(0.02)	(0.06)	(0.07)
FC2000	0.04**	0.04**	0.03	0.11**	0.09
	(0.02)	(0.02)	(0.02)	(0.05)	(0.06)
Charity	-0.01	-0.00	0.00	0.02	0.05
	(0.02)	(0.02)	(0.02)	(0.06)	(0.06)
Government	0.00	-0.00	-0.00	0.04	0.03
	(0.02)	(0.02)	(0.02)	(0.09)	(0.11)
Observations	2,487	2,402	2,467	2,487	2,467
R-squared	0.00	0.02	0.43	0.00	0.43
Ind. FE	N	N	Y	N	Y
Ind. Controls	N	Y	N	N	N
Control avg.	0.0961	0.0931	0.0956	0.0956	0.0956

Table 10: Main effects on number of valid receipts brought to the follow-up visit (unconditional on reporting a trip in the intervening period)

VARIABLES	(1) OLS	(2) OLS	(3) OLS	(4) FE	(5) FE	(6) IV	(7) IV FE
FC1000	0.07**	0.08***	0.07**	0.10***	0.08***	0.22***	0.25**
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.08)	(0.10)
FC2000	0.07***	0.07***	0.08***	0.07**	0.07*	0.17***	0.19**
	(0.03)	(0.03)	(0.03)	(0.03)	(0.04)	(0.06)	(0.09)
Charity	0.00	0.01	0.01	0.05*	0.06*	0.08	0.18**
	(0.02)	(0.03)	(0.03)	(0.03)	(0.03)	(0.07)	(0.08)
Government	0.03	0.02	0.02	0.05	0.02	0.18	0.12
	(0.03)	(0.03)	(0.03)	(0.03)	(0.04)	(0.14)	(0.20)
Officer	-0.02	-0.03	-0.03	-0.03	-0.03	-0.04	-0.04
	(0.03)	(0.03)	(0.04)	(0.04)	(0.05)	(0.03)	(0.05)
Enumerator	0.06	0.05	0.04	0.02	0.00	0.05	0.01
	(0.04)	(0.04)	(0.04)	(0.05)	(0.05)	(0.04)	(0.05)
Observations	1,723	1,659	1,418	1,510	1,262	1,723	1,262
R-squared	0.01	0.03	0.08	0.44	0.47	0.01	0.46
Ind. FE	N	N	N	Y	Y	N	Y
Ind. Controls	N	Y	Y	N	N	N	N
Trip Controls	N	N	Y	N	Y	N	Y
Control avg.	0.124	0.120	0.137	0.115	0.131	0.124	0.131

Table 11: Main effects on receipts per trip.

	(1)	(2)	(3)
VARIABLES	OLS	OLS	OLS
FC1000	-46.53	-29.99	-25.56
	(84.77)	(84.60)	(80.84)
FC2000	-141.96**	-160.85**	-149.13**
	(69.55)	(72.00)	(67.76)
Charity	-63.60	-45.12	-62.89
	(74.67)	(75.04)	(75.98)
Government	-191.61*	-167.37	-92.01
	(111.23)	(113.86)	(99.52)
Officer	17.01	25.73	-52.04
	(75.29)	(76.41)	(76.76)
Enumerator	-63.80	-62.15	-43.01
	(80.97)	(82.87)	(80.99)
Observations	1,355	1,305	1,168
R-squared	0.02	0.04	0.15
Ind. FE	N	N	N
Ind. Controls	N	Y	Y
Trip Controls	N	N	Y
Control avg.	1528	1534	1537

Table 12: Total (selection and causal) estimates on equilibrium bribes.

VARIABLES	(1) Tobit	(2) Fixed Effects
Reward: 1000 FC	-191.13**	-224.82***
	(95.29)	(61.52)
Reward: 2000 FC	-300.21***	-68.67
	(96.82)	(75.94)
Reward: 2000 FC donation to charity	-52.46	-89.26
	(95.09)	(69.90)
Reward: Gov. donates 2000 FC to charity	-126.03	-135.76*
	(97.86)	(74.70)
Toll officer treatment in effect	97.06	85.91
	(115.62)	(86.71)
Enumerator validating traffic at toll	-195.84	-53.98
	(124.83)	(97.53)
Observations	1,611	1,397
Ind. FE	N	Y
Ind. Controls	N	N
Trip Controls	N	N
Control avg.	1326	1357

Notes: The dependent variable is the bribe (amount paid not backed by a receipt) paid by the driver. In column (1), we show the results of a Tobit regression based on equation 9 where bribes are considered to be censored below at 0 and all drivers who receive a receipt (and do not overpay for said receipt) have a bribe of 0. In column (2), we show the result of a regression where we restrict the sample to drivers who are observed to not get receipts (but who report traveling) in two or more different rounds and the outcome is the amount paid as a bribe in each round — this regression includes individual fixed effects, but is otherwise the same as equation 9. Standard errors clustered by individual. $^*p < 0.1,^{**}p < 0.05,^{***}p < 0.01$

Table 13: Tobit results on bribes.

	(1)	(2)
	Amount paid to officer	
VARIABLES	No receipt	Got receipt
Weekly phone credit expenditure (1000s of FC)	13.00**	0.78
	(6.16)	(1.73)
Driver estimated trip earnings (1000s of FC)	0.92*	-0.87
	(0.52)	(0.57)
Tax morale game winnings (in FC)	0.28*	0.06
	(0.16)	(0.11)
Won dishonesty game	8.16	14.25
	(61.87)	(23.40)
Observations	1,178	236
R-squared	0.03	0.06
Ind. FE	N	N
Ind. Controls	N	N
Trip Controls	N	N
Control avg.	1534	1956

Notes: All columns are estimated using equation **??**, so all results include round fixed effects (not reported). The dependent variable is amount of money paid to the toll officer, as reported by the driver at his follow-up interview. Column (1) restricts the sample to trips where the driver did not have a valid receipt. Column (2) restricts the sample to trips where the driver did have a valid receipt. Standard errors clustered by individual. $^*p < 0.1,$ $^{**}p < 0.05,$ $^{***}p < 0.01$.

Table 14: Covariates of amount paid when getting a receipt and not.

Beta	FC1000	FC2000	Per FC
0.5	-1.71	-0.96	-0.95
	[-3.55, -0.12]	[-1.97, -0.18]	[-1.84, -0.25]
0.63	-1.37	-0.77	-0.77
	[-2.86, -0.10]	[-1.58, -0.15]	[-1.48, -0.20]
0.74	-1.16	-0.65	-0.65
	[-2.41, -0.08]	[-1.34, -0.12]	[-1.25, -0.17]
1	-0.81	-0.45	-0.45
	[-1.67, -0.06]	[-0.93, -0.09]	[-0.87, -0.12]

Notes: This table shows estimates of the elasticity of corruption (not getting a receipt) with respect to instantaneous monetary cost of compliance (i.e. cost of getting a receipt), using the calculations described in equation (11). Standard errors reflect the empirical 95% confidence interval from 5,000 bootstrap iterations. Bootstrapping was done using a block bootstrap, blocked by individual participant. The rows show the results using different assumed values of beta. Column (1) shows the results using only the effects of the 1000 FC treatment. Column (2) shows the results using only the effects of the 2000 FC treatment. Column (3) uses both treatments simultaneously (treating them linearly).

Table 15: Estimated elasticities of receipt-getting with respect to instantaneous monetary reward for various values of β. Bootstrapped 95% confidence interval shown in square brackets below each estimate.

| | (1) | (2) | (3) |
VARIABLES	Baseline	Has passenger	No passenger
FC1000	0.07**	0.05*	0.12*
	(0.03)	(0.03)	(0.07)
FC2000	0.07***	0.05*	0.14**
	(0.03)	(0.03)	(0.06)
Charity	-0.00	-0.01	0.07
	(0.02)	(0.03)	(0.06)
Government	0.03	0.03	0.03
	(0.03)	(0.03)	(0.06)
Officer	0.00	0.00	0.06
	(0.03)	(0.03)	(0.07)
Observations	1,723	1,357	270
R-squared	0.01	0.01	0.04
Ind. FE	N	N	N
Ind. Controls	N	N	N
Trip Controls	N	N	N
Control avg.	0.124	0.138	0.0833

Notes: All columns are estimated using equation (9), so all results include round fixed effects (not reported). The dependent variable is the number of valid (incentivized) receipts per trip reported by the driver at his follow-up interview. Column (2) restricts the sample to trips where the driver reported having a passenger when he passed the toll. Column (3) restricts the sample to trips where the driver reported not having a passenger when he passed the toll. Standard errors clustered by individual. $^*p < 0.1,^{**}p < 0.05,^{***}p < 0.01$.

Table 16: Receipts per reported trip under different conditions.

11 Appendix

11.1 A Simple Theory of Social Incentives

The social incentives will not form the focus of this paper. Nonetheless, here we provide a simple theory of their intended impact. Consider an individual who places social value $\gamma \in [0, +\infty)$ on each dollar spent on a given social cause (here, donations to a home for widows). Assume that they place probability p on the research team actually making the promised payment to the widow's home. Then, the charity treatment, a promised payment of m dollars, is valued in expectation at $mp\gamma$.

Now, consider the government treatment, which has the exact same social cause, but involves transmission through the government. Now, there is an additional term, $\phi \in [0, 1]$, which represents the individual's perceived percentage of the money that actually will reach the social cause. In general, ϕ could be less than 1 either due to fraud or generic waste in government. In our setting, we believe that the parameter is best interpreted as perceived fraud, since the money is being transmitted directly as cash to the widow's home, not spent on goods or services (where waste might be a concern). This means that the government treatment, a promised payment of m dollars, is valued in expectation at $mp\phi\gamma$. Notice that the m, p, and γ parameters are all held constant, which means that comparing responses to the government and charity treatments allows us to recover perceived ϕ.

Note that this simple theory will break down if the framing of the government treatment changes the individual's utility function. For example, if the framing of the government treatment causes people to increase their belief that the experimenters are, in fact, working for the government, then that could affect the utility from the government treatment. If individuals increase their subjective view that the experimenters are working for the government by ε and place utility δ on obeying the wishes of the government, then the utility from the government treatment will be $mp\phi\gamma + \varepsilon\delta$ and comparing the two treatments will yield a biased estimate of ϕ. We believe that this is unlikely to have occurred — only 6.3% of people reported that they believed the research office was run by the government and this self-report was uncorrelated with treatment status — but we cannot explicitly rule it out.

11.2 Additional derivations

As noted, a driver gets a receipt if the surplus from collusion is negative, that is if

$$t - \chi - w - k + \xi \leq 0 \tag{15}$$

The share of drivers getting the receipt, $\bar{t}$, is thus given by:

$$\bar{t} = P(\xi - \chi \leq k + w - t) = \int m(k + w - t) \tag{16}$$

where m is the PDF of $\xi - \chi$.

Then it is easy to see that

$$\frac{\partial \bar{t}}{\partial k} = \frac{\partial \bar{t}}{\partial w} = m(k + w - t). \tag{17}$$

The average equilibrium bribe, $\bar{b}$, is

$$\bar{b} = w + \delta(t - w - k + E(\xi - \chi)) = w + \delta(t - w - k + \int (k + w - t)m(k + w - t)) \tag{18}$$

So the change in the average bribe in response to a change in the driver's financial outside option is:

$$\frac{\partial \bar{b}}{\partial k} = -\delta(1 - (k + w - t)m(k + w - t)) \tag{19}$$

11.3 Panel Structure

11.3.1 Discussion of the Panel Structure

As we laid out in Section 4.1 and Section 4.2, each participant in our study received different treatments in succession over the three rounds of the experiment. This design is novel: to our knowledge, ours is the first to randomize within individuals in a field setting in economics.[37] Obviously, this design is not possible for many experiments, which require a longer run period (e.g. an agricultural study that covers an entire growing period) or are built around a single event (e.g. an anti-vote buying intervention around a single election). However, a similar design may be

---

[37]Similar designs have been used in some lab experiments and in medical experiments (see for a review). The most similar field experiment design to our knowledge is Banerjee et al. (2007) on the Balsakhi program in India, but individuals in their study never experience multiple different treatments.

possible for many shorter term interventions and we believe it is instructive to lay out and discuss some of the advantages and disadvantages of the panel design here.

The advantage of using a panel design is clear: increased power. Since each participant appears multiple times, total observations $T * N$ increase. Since individuals may have correlated errors over time, we cluster at the level of the participant. This means that *effective* observations do not rise at the same rate as total observations. However, the fact that we change the treatment status for each participant in each round mitigates this issue substantially. Even when clustering standard errors at the level of the individual, the fact that treatment changes across rounds for each participant creates within-participant (i.e. within cluster) variation. This means that the loss of power from clustering is small. If we were to have left participants in the same treatment repeatedly, then the panel structure would provide a smaller power advantage.[38] Due to the randomized order of treatments, the experiment is equally valid with and without individual fixed effects. Fixed effects will improve power if the individual error component is large relative to the individual-by-round error and we observe multiple observations for a sufficient portion of the sample. They will worsen power if the opposite is true.

The disadvantages of using a panel design are less clear. The main concern is inter-temporal spillovers. If treatment effects persist across periods, then this could create bias in our estimates. These spillovers could take several forms. First, particularly concerning for our design, participants could be confused about changes in their treatment status Because our treatment is a promise of a certain reward (or lack thereof), if participants do not realize that their treatment has changed, they will continue to respond as though in their prior treatment group, thereby biasing estimates. One sub-group in our experiment exhibited nontrivial confusion over treatment status, which we discuss in Appendix 11.3.2. That said, our results are robust to removing this group. Second, participant behavior in one round (induced by treatment) could have persistent effects in future rounds, even if participants fully understand that their treatment status has changed. For example, if relationships

---

[38]In a simple three-period, two-treatment simulation where 30% of the variation is constant within person and 70% is randomly assigned each person-period, power is 8 percentage points higher when changing treatment each round (i.e. C-T1-T2) versus keeping each individual in the same treatment twice (i.e. C-T1-T1 or C-T2-T2) and 34 percentage points higher than a design in which individuals are in the same treatment/control category across all three rounds (i.e. C-C-C vs. T1-T1-T1 vs. T2-T2-T2). The power advantage of changing treatment each round is even larger (about 19 percentage points) when we add individual fixed effects (and individual fixed effects are not possible in the design where treatment status never changes).

between drivers and toll officers are important, and demanding a receipt affects those relationships, then a prior treatment that changes driver behavior could have persistent effects through its effect on the driver-toll officer relationship. Third, there could be a differential effect of treatment after a driver has previously been exposed to a different treatment (e.g. a disappointment or surprise effect). We explore these two effects in detail in Appendix 11.3.3. They do not appear to create bias in our estimates of contemporaneous treatment effects.

11.3.2 Follow-Up Visit 2 Controls

As noted in Section 7.2, there was an issue in Follow-Up Visit 2 where participants who were assigned to the control did not understand that they were being removed from their prior treatment. The issue was that participants who were in the control in that round were asked to bring receipts (with no mention of a reward), but we did not *explicitly* emphasize that they would not receive any other reward (regardless what past treatment they had). As a result, many participants in the control group likely believed that they were still in their old treatment category. We corrected this for Follow-Up Visit 3, but we will use this section to explore the implications of that implementation issue.[39]

Table 18 shows the difference in treatment recall for the affected round versus the other two rounds (pooled here, though they look similar when examined separately). There are two important points to notice from this table. The first is that the issue of people in the control incorrectly identifying themselves as being part of treatment is much bigger in Follow-Up Visit 2 than in the other rounds. In total, in Follow-Up Visit 2, a full *72 percent* of respondents in the control group identified themselves as being part of one of the four actual treatments versus only 32 percent in the other visits. The second is that in the treatment groups, this problem essentially disappears. Averaging across the 4 other treatments, in Follow-Up Visit 2, an average of 50 percent of people correctly identify their actual treatment. In the other two visits, an average of 55 percent of people correctly identify their actual treatment, an economically and statistically identical number.

In Table 19, we examine this issue in more detail by focusing on the control groups, identified earlier as the most likely place for confusion. In this table, we restrict the sample to people who were assigned to the control in Follow-Up Visit 2

---

[39]We pre-registered this issue once it was discovered part of the way through the experiment.

(top panel) or in Follow-Up Visit 3 (bottom panel). Thus, if there were no confusion issues (or misremembered treatments), we should see all of the mass in the first column (i.e. everyone should remember that they are in the control). We then compare their recalled treatment to the treatment group they were assigned in the prior (not current) round. For example, for Follow-Up Visit 2, we compare their recalled treatment at Follow-Up Visit 2 to the treatment they were assigned at Follow-Up Visit 1.[40] If it is in fact true that in Follow-Up Visit 2, people did not realize that they were no longer assigned to their prior treatment (and instead were in the control), then we should expect to see many more people match their remembered treatment to their former (but not current) treatment status. This is exactly what we see. For example, at Follow-Up Visit 2, 74% of people who were assigned to the FC1000 treatment at Follow-Up Visit 1 (but are in fact in the control) "remember" that they are in the FC1000 treatment — but that same quantity is only 38% at Follow-Up Visit 3. Averaging across treatments, 59 percent of the control group in Follow-Up Visit 2 incorrectly believe they are keeping their old treatment versus only 34 percent in Follow-Up Visit 3.

These facts suggest that our anecdotal information about the confusion with the Follow-Up Visit 2 control group is correct. We would expect this confusion to bias our results downwards, since if much of the control group believed that it was treated then they would behave as if they were treated, pushing the difference between control and treatment towards zero.

In Table 20, we show the results from estimating equation 9. In column (1), we show our baseline results, as seen in Table 11 column (1). In columns (2) and (3), we remove the entirety of Follow-Up Visit 2, without and with individual fixed effects, respectively. In columns (4) and (5), we just remove the control group for Follow-Up Visit 2, which we believe is the most reasonable way to deal with this issue. As we anticipated, the results are larger, though not statistically significantly different from the baseline results.

11.3.3 Path Dependence and Reputation

One area of additional substantive interest is the effect of the panel structure, with repeated exposure to different treatments over the course of the experiment. The

---

[40]Since no one was assigned the same treatment multiple times, and we are restricting to people who were actually assigned to the control in the current round, there is no one assigned to the control in the prior round.

treatment history of a given participant could have effects through three channels. First, the panel structure could result in respondent confusion. If respondents do not understand that they are changing treatment categories, then there will be a mechanical persistence of treatment effects.[41] Second, prior treatment exposure could affect participant behavior through an information or learning channel. If prior exposure leads participants to change their behavior at the tolls, then they may learn new information about the ease of getting receipts or the possibility of changing bribe levels that could be persistent. Third, and relatedly, prior treatment exposure may have persistent effects through a reputation channel. If treatment induces a change in driver behavior, even for a single period or interaction, then it may persistently affect their relationship with a given toll officer. The direction of this effect is not *a priori* clear. A toll officer could hassle them more in the future as punishment or become more acquiescent to their demands for a receipt (as he learns that they are a more honest type than he believed).

In Table 21, we estimate a modified version of equation 9, in which we regress the number of valid receipts brought in a given round on treatment in that round. In column (1), we replicate our baseline results, pooling across all rounds, as in column (1) of Table 10. In column (2), we use only the data from the first follow-up visit (where there cannot be any history effects, since it includes only the original treatment assignment). In column (3), we add a set of controls for the treatment history of the participant. In particular, we code up a set of dummy variables that have a value of 1 if the participant has *previously* been exposed to a given treatment and are 0 otherwise.[42] Since the order of treatments was randomly assigned for each individual, these dummies are valid, exogenous regressors.

In column (4), we add a control for whether the driver has ever brought a valid receipt to any previous interview. This control combines the driver's type (honest or not) with any "treatment effect" of previously being induced into getting a receipt by treatment. In columns (5) and (6), we attempt to disentangle these two margins by using the same treatment history variables from column (3) as instruments for having previously presented a valid receipt. This approach is only valid under the strong assumption that prior treatment affects current period behavior *only* through its effect on inducing drivers to get receipts in the past. This would require us to rule out the possibility of respondent confusion or any reputation effects that do

---

[41]We explore a particular case of this in Section 11.3.2.

[42]In follow-up visit 1, all of these variables are 0.

not operate through getting a receipt.[43] We do not believe that this assumption is necessarily likely to hold, but we include it as an additional data point.

Overall, we believe that this exercise provides a number of interesting insights for the reader. First, the main (current period) treatment effects are relatively consistent across periods, no matter the specification. This suggests that our main effects are relatively unaffected by the panel structure of the experiment, which may be of interest to other experiments designing future interventions. One exception to this is that the results using only the data from the first follow-up are somewhat stronger than the results when we pool all rounds. Second, there is some evidence that prior treatments have a persistent effect on behavior, though the point estimates are generally insignificant.[44] Finally, there is some suggestive evidence for a persistence channel through past receipt-bringing, but this evaporates when individual fixed effects are added in column (6). This suggests to us that persistence is unlikely to be first-order in our context.

Next, we show several results designed to dig deeper into the issue of reputation. The earlier analysis explored the potential long-run effects of changing a participant's reputation through prior treatment status. Here, we explore in more detail the potential effects that *forward-looking* reputation concerns might have on treatment effects. In particular, we might think that drivers with more exposure to a given set of toll officers would be less inclined to take up treatment if treatment affected their reputation. Since the treatment is temporary, but relationships presumably last for many periods, drivers might be unwilling to risk their reputations if they believe that they will continue to have significant exposure to a toll officer after the treatment period.[45] We do not have data on driver expectations about which tolls they are likely to frequent in the future, but we use data on self-reported past history of toll usage as a proxy.[46] We show the results using this proxy in Table 22.

In column (1), we repeat our analysis from column (1) of Table 11, using spec-

---

[43]For example, if a participant demanded a receipt but failed to get one, this could affect future interactions with the toll officer and would thus constitute an exclusion restriction violation.

[44]Note that some of this effect is likely due to the confusion among members of the follow-up visit 2 control group as discussed in Section 11.3.2.

[45]As discussed in Section 3.1, toll officer identities are constant during the period we study and tend to remain so for long periods.

[46]We validate this proxy by regressing the toll passed by the participant on his self-reported history. For each of the main tolls, prior history is a statistically significant and economically meaningful predictor of toll chosen, suggesting that prior history is a valuable proxy for the driver's expectations about his future interactions.

ification (9).[47] In column (2), we restrict our sample to the set of *trips* where the driver reported having visited that toll at least once in the two months before our study period began. In column (3), we use the same sample, but now we repeat the specification from column (5) of Table 21, in which we use prior treatment as an instrument for past receipt bringing. These results indicate that forward-looking reputation concerns are unlikely to be important for explaining the relatively low responsiveness of participants to treatment in our experiment. In particular, comparing columns (2) and (3), we see that, if anything, drivers visiting tolls that they have frequented in the past (and are likely to continue to frequent in the future) are *more* responsive to treatment than the population as a whole, though the difference is not statistically significant. Likewise, when we repeat the analysis looking at the effect of having been previously induced into bringing a receipt in column (3), we see that the point estimate on prior receipt is smaller and insignificant compared to its analogue in column (5) of Table 21. Again, this suggests that forward-looking receipt effects are small — if permanent relationships were heavily affected by participants demanding receipts, we would expect this effect to be larger.

11.3.4 Further Discussion of Charity/Government Treatments

In this section, we elaborate on the charity and government donation treatments, which do not receive much attention in the main paper. As discussed in Section 4.2, in these treatment conditions, we promised to donate 2000 FC to a community of poor widows called "Bikolo wa moyo", or "Well of life." This is a place where widows who lack other family to support them are able to live. They regularly beg in the street as a source of income and so their presence in the city is salient to most residents. The distinction between the charity and government treatments was that in the charity treatment, we promised to deliver the money ourselves, while in the government treatment, we said we would give the money to the government, which would be responsible for transmitting the funds to the widows. As noted in 11.1, comparing behavioral responses to these treatments provides us with an estimate of the degree to which perceived government corruption makes drivers undervalue the social benefits of their tax payments (in this case, as proxied by donations to widows).[48]

---

[47] We do not show the coefficients on enumerator presence and the toll officer treatment for visual simplicity.

[48] We can compare these treatments to the control or to the 2000 FC treatment as well. Such comparisons would allow us to estimate bounds on the relative utility value of donations to widows.

To repeat the basic theory, the "utility value" of m FC of donations as delivered by the researchers is $mp\gamma$, where p is the probability that the research team delivers on their promise and γ is the participant's subjective valuation of 1 FC of donations to the widows. We can compare this to the utility value of m FC of donations as delivered by the government, which is $\phi mp\gamma$. Note that we now have an additional term $\phi \in [0, 1]$, which is the individual's perceived percentage of the money that will actually reach the widow home (i.e. will not be stolen or redirected). However, the other terms in the equation, m, p, and γ, should not be affected by the government treatment. Furthermore, due to the random assignment of treatment, we should expect that p, and γ should be equal (in expectation) in the two treatment groups (m is held constant at 2000 FC by the experiment).[49] Thus, in theory, identifying ϕ is straightforward. Under the assumptions of this simple theory, we can identify the average ϕ in the population from the following equation:

$$\hat{\phi} = \frac{\hat{\beta}_{GOV}}{\hat{\beta}_{CHARITY}} = \frac{\phi m \hat{p} \gamma}{m \hat{p} \gamma}$$

However, one concern is that this estimation may be low powered. Under the assumption that it is roughly costless to donate one's own money to the widow's home, then $\beta_{CHARITY} \leq \beta_{FC2000}$, because any individual could always take the 2000 FC that we offer and donate it. If at least a portion of the population has a $\gamma < 1$ (as seems very likely), then even abstracting away from p, then it will be more difficult to detect effects for either the charity or government treatments, relative to those for the cash treatments. In the extreme, if all people had $\gamma = 0$ (no valuation of donations to widows), then it would be impossible to estimate ϕ in that population, since both treatments would yield a zero result. One attempt to limit this issue is to isolate a sub-population where we expect γ to be high, namely individuals whose own mothers are widowed or who have lost both of their parents (in which case, it is plausible that their mother was a widow for a period of time). These sub-populations may, due to their personal experience, have a higher social weight on donations to widows.

In Table 17, we estimate equation (9) using the number of receipts brought

However, this question is not of fundamental economic interest and we do not attempt to answer it here.

[49]Note that we cannot separately identify p and γ — however, separate identification of those parameters is not important to this exercise, as by dividing the coefficients, we should cancel them out.

	(1)	(2)	(3)	(4)	(5)
VARIABLES	Original	Widow	Widow (no R2 control)	Widow/orphan	Widow/orphan (no R2 control)
Charity	-0.01	0.03	0.12*	0.03	0.10*
	(0.02)	(0.05)	(0.06)	(0.04)	(0.05)
Govt.	0.00	-0.03	0.02	-0.03	0.01
	(0.02)	(0.04)	(0.05)	(0.03)	(0.04)
Observations	2,487	651	507	874	682
R-squared	0.00	0.01	0.02	0.01	0.01
Ind. FE	N	N	N	N	N
Control avg.	0.0961	0.111	0.111	0.111	0.111

Standard errors clustered by individual. $^*p<0.1,^{**}p<0.05,^{***}p<0.01$

Table 17: Results by closeness to widows.

per round (unconditional on having travelled) as the outcome. For simplicity of exposition, we only show β_{GOV} and $\beta_{CHARITY}$, though all other treatments are also in the regression. In column (1), we show our original specification (the same as column (1) of Table 10). In column (2) we limit our sample to only those individuals whose fathers (but not mothers) have died. In column (3) we keep the same sample, but remove the Follow-Up Visit 2 controls, for reasons discussed in Section 11.3.2. Column (4) is the same as column (2), but also adds individuals who have lost both parents. Finally, column (5) again removes the Follow-Up Visit 2 controls. The results in the "high γ" samples suggest that ϕ is low. In columns (3) and (5), a comparison of the coefficients on the charity and government treatments yields p-values of 0.17 and 0.13, respectively.[50] Notwithstanding our inability to statistically distinguish between the coefficients with precision, the ratios of the coefficients are suggestive of a low value of ϕ, estimated at 0.20 in column (3) and 0.11 in column (5). If these point values are taken seriously, they would suggest that one reason for tax non-compliance is the strong perception of government corruption, which reduces the social benefits of a dollar of tax revenue by 80-90%.[51]

---

[50] A strong case could be made *a priori* that we should do a one-sided test instead of a two-sided test, since we pre-specified that the effect of the charity treatment should be larger than that of the government treatment. In this case, our p-values would be 0.09 and 0.06, respectively.

[51] As a note of caution, when we bootstrap a 90 confidence interval for the value of ϕ, our estimated range is very large, $[-0.77, 1.19]$, using the specification in column (3), and $[-0.96, 1.15]$, using the specification in column (5).

11.4 Robustness and Additional Analysis

11.4.1 Attrition

In this section, we consider the issue of attrition. Table 23 contains the results of regressions in which the outcome was a dummy that equaled 1 if a driver failed to show up for the following visit. For example, in the first column, drivers receive a 1 if they attrited between rounds 1 and 2 (i.e. between the baseline visit and the first follow up visit). The results suggest that the charity and government treatments *decreased* the probability of attrition between baseline and follow-up visit 1 relative to the control. Likewise, the 1000 FC reward treatment, the 2000 FC reward treatment, and the government treatment *decreased* the probability of attrition between visits 2 and 3; and the government treatment *decreased* the probability of attrition between visits 3 and 4. Note that the sample changes slightly across columns, as, for example, in order to be observed in column (2), one cannot have attrited between rounds 1 and 2. Finally, in column (4), for the population of drivers who do not attrit, we regress the number of days between appointments on treatment status. If compliance with our protocol was perfect, we would have an average number of days between appointments of exactly 21. Instead, drivers take an average of 23 days, but there is no statistically or economically meaningful relationship between treatment and days between appointments, which suggests that we do not need to be worried about any effect of having more time to take trips on outcomes.

Overall, the most concerning of these findings is that in each of the three rounds, the government treatment appears to have had an effect. The direction of this finding is perhaps surprising, as it suggests that drivers in the government treatment were least likely to attrit (indeed, the coefficient is roughly 100% of the control mean, suggesting almost no attrition in this group). Ex ante, we believed that individuals in the government treatment would be most likely to attrit, over concerns that we (the experimenters) were involved with the government and might arrest drivers who lacked certain documents. However, the opposite appears to have been the case. One theory is that drivers in the government treatment may have believed that they would be found by the police or otherwise sought by the authorities if they did not return, but we cannot know for sure.

However, despite the slight evidence of differential attrition across treatment groups, overall attrition across rounds was not large: 3% between baseline and follow-up 1, 2% between follow-ups 1 and 2, and about 5% between follow-ups

2 and 3. Overall, 87.6% of all baseline attendees completed all 4 visits. In our analysis, whether we add individual fixed effects, which means that all comparisons are within individual, or we can restrict to the set of individuals who attended all visits, there is neither a statistically or economically meaningful effect on the coefficients of interest.

11.4.2   Attempts to Get Receipt

In this section, we elaborate on the possible coercive power of the toll agents and the role of our treatments in encouraging "receipt-seeking effort" in addition to actual receipts gotten. In particular, we are interested in whether the elasticity of this effort may be greater than that of the actual success-rate, which would provide evidence that toll officers have some degree of coercive power.

We do not have a perfect measure of receipt-seeking effort. However, for any trips in which people did not have a receipt, we asked people why.[52]   Table 25 shows that 9 percent of respondents reported asking for a receipt, but having the agent refuse to grant them one.[53] Since this data is self-reported, we treat this as only suggestive evidence of driver effort. Drivers may feel social desirability bias to provide an excuse as to why they lack a receipt, even if the truth is that they did not pay the toll. However, we believe that there is likely still some signal in this measure, as we show in more detail later.

In Table 24, we use our trip-level regression specification, equation 9. Note that this means that we are conditioning on an individual reporting the trip to us. In column (1), we estimate our main effects, but restricting to a sample that includes only (a) trips with receipts and (b) the first trip without a receipt. These estimates are very similar to those in Table 11, even though the sample is somewhat different. In column (2), we look only at the distribution of excuses (ignoring trips with receipts) and observe a positive, but insignificant effect of treatment on reporting trying (but failing) to get a receipt. Note, however, that this specification will treat as missing data any drivers who report only a sole trip, for which they have a receipt, and it is contaminated by selection bias. If treatment induces drivers to

---

[52]Importantly for the estimation that follows, we only asked people once for their excuse, even if they had multiple trips without a receipt. Thus, in the estimation that follows, we restrict our sample to trips where the driver has a receipt *or* to the first trip for which he lacks one.

[53]We can also include the 2.5 percent of respondents who said that the machine was broken or out of power, both of which may be excuses used by agents who refuse to issue receipts, but the results are similar.

demand receipts more intensely, it could plausibly both induce drivers to go from not demanding receipts to demanding them (but failing) *and* to go from demanding receipts (but failing) to demanding them more stridently and succeeding. This latter effect would cause us to underestimate the effect of treatment on effort in column (2). Thus, in column (3), we define a dummy that is 1 if either the participant brought a receipt or reported requesting one, but being denied.

Overall, these results suggest that toll officers may have some coercive power over drives, but it appears to be quite limited. The point estimates in column (2) are positive for the financial incentives, but, notably, we cannot reject that the coefficients in columns (1) and (3) are the same. If there was a substantial margin on which toll officers were able to resist demands from drivers for receipts, we would expect that incentivized drivers would be much more likely to report trying and failing to obtain receipts (due to their desire for the reward). These results are, however, only suggestive. It could also be the case that, in equilibrium, there are few unsuccessful demands in part because drivers know that toll officers are likely to resist.

11.4.3 Receipt Loss

One concern with interpreting our estimates as an elasticity of corruption with respect to incentives would be that we systematically mismeasured our outcome. If we measured our outcome equally accurately in both treatment and control (even with noise), this would not bias our estimator. But if our mismeasurement was not centered at zero and was more severe in treatment, this would be a serious concern. One obvious way this could occur would be if drivers commonly lost their receipts. Loss of receipts is not zero centered (no one "accidentally finds" receipts, so there are no positive shocks) and, importantly, one can only lose a receipt if it was requested in the first place. This suggests that lost receipts could be significantly more severe for our treatment groups, which are more likely to demand receipts (as we know from Section 8.1). Thus, if drivers commonly lose receipts, this could substantially bias our results towards zero.

Anecdotally, we do not believe that losing receipts is common. Drivers generally store the receipts with their money. Since there is no evidence that drivers are careless in potentially losing bills (even low quality Congolese francs close to disintegration), we believe that they would treat their receipts with the same care. Indeed, the receipts we received had often been kept for extended periods, as mea-

sured not only by the date of the receipt, but also by the wear-and-tear on the receipt itself. As a result, we believe that the vast majority of reported "loss" of receipts by participants is not truthful.

However, we also test directly for evidence that drivers are losing receipts. In Table 25, we show the reasons that drivers report for not having a receipt, *conditional on reporting having taken a trip*.[54] First, by far the most common reason for missing a receipt is self-reporting having underpaid the toll. The next most common reason is having left one's receipt at home, followed by reporting losing the receipt. Any driver who left their receipt at home could have fetched the receipt and brought it for whichever reward to which they were entitled. The fact that none did suggests most of these reports are untruthful.

Nevertheless, we also test more directly for evidence of lost receipts in Table 26. In column (1), we estimate our equation 9, restricting to the sample of (a) trips with receipts and (b) the first trip without a receipt, which shows that the financial incentives induced participants to bring more receipts. In column (2) we regress our treatments on a dummy variable for reporting having lost one's receipt. If losing receipts were an issue, we would expect that the treatments that show strong treatment effects in our main tables would also predict more people reporting having lost their receipts. In column (3), we construct a second dummy variable for either reporting losing one's receipt or bringing a valid receipt. As expected, the treatment effects in column (3) cannot be distinguished from the effects in column (1).[55] This does not prove that no receipts were lost. However, it tells us that the magnitude of any "lost receipt" effect must be small, since otherwise we would expect to see an identical pattern to treatment (since only receipts gotten in the first place could be subsequently lost, something that only treated individuals are likely to do).

11.5   Trip Misreporting

As discussed briefly, drivers do not report their trips with perfect accuracy. We explore possible consequences of trip reporting errors in this section. Table 27 shows the results with a dummy variable indicating the driver took at least one trip (column 1) and a count variable of the total trips taken (column 2) on the

---

[54]Note that since this specification requires having reported a trip, then even if there is some under-reporting of trips, this likely overestimates the share of people reporting having lost their receipts.

[55]If anything, the effects are somewhat smaller for the financially incentivized groups, which is not surprising, if the receipts were more valuable to them and thus were better cared for.

left-hand side, regressed on the treatment dummies. Negative coefficients on the government treatment suggests that either individuals in this treatment group went on slightly fewer trips or, perhaps more likely, they reported taking fewer trips. Although there is no reason that drivers in the government donation treatment group should have feared admitting to trips, it is possible that any mention of the government induced concerns among drivers, leading them to underreport trips. This measurement error would likely bias the estimated effect of the government donation treatment on receipts-per-trip upward, if we assume that unreported trips lacked receipts, as seems plausible. This does not meaningfully affect our results, since we did not detect a treatment effect for this treatment arm.

Another way to test for underreporting of trips is to consider the set of trips validated by enumerators at tolls. As discussed in Section 4.2, the validation was done by an enumerator stationed at the toll who flagged down participants (who had been warned to expect this possibility) and completed a very short survey with them. Because we started validating trips after a month and we only validated four toll-days per week, the set of validated trips is small: only 169 observations.[56] However, it is still possible to see whether the treatments appear to have affected misreporting among this subset. More specifically, Table 28 shows the results from a regression of a dummy variable for unreported trips on the treatments (restricting the sample to the 169 validated trips).[57]

The first column suggest that under-reporting is lower in the control group relative to *all four* treatment conditions. A joint F-test that the coefficients on the four treatment indicators are different from zero has a p-value of 0.096 using standard inference or of 0.13 using randomization inference.[58] We believe that the difference between the control and the treatment groups is likely a false positive for several reasons.

First, as shown in column 2 of Table 28, there are no systematic differences in misreporting when we compare across the treatment groups that *drivers themselves thought they were in.* Driver recall is obviously endogenous. However, if drivers were

---

[56]203 driver-trips were observed, but only 169 could be matched to the sample on driver ID number, driver name, or phone number.

[57]Trips are considered unreported if, at his follow-up interview, the driver reported no trip within 7 days of the date of the validated trip (when the enumerator completed a short survey with a given motard at the toll). The 7 day window is used to account for the relatively poor recall of drivers about their exact date of travel.

[58]With a small sample (169 trips), the assumptions behind asymptotic normality may not hold, so randomization inference provides robust p-values in this case.

strategically under-reporting, then using the endogenously recalled treatment should *strengthen* the "effect" of treatment on misreporting: precisely those drivers who believed themselves to be treated should be the ones under-reporting trips. However, we do not see that effect here. Instead, there is no effect using recalled treatment instead of assigned treatment.[59] Related to this, in column 3, we restrict to the sample of drivers who specifically reported that they did not remember (at all) their treatment status. This group is small (only 31 trips), but the results are suggestive. Despite being the group of drivers who should be *least* likely to respond to treatment, by virtue of their total uncertainty about their treatment, the "treatment effects" in this population are very large and indeed are larger (significantly so, for some treatments) than those for the population as a whole.

Second, one might imagine that drivers were embarrassed about admitting to trips for which they did not get receipts. However, such embarrassment would exist in the control group, too, given that it was an "active" control: enumerators asked drivers in this group to bring receipts, using the same wording they used with drivers in other treatment groups (absent language about receiving a reward). Holding constant experimenter demand effects was precisely the objective of this active control. Moreover, it makes little sense why they would try to hide trips that had already been recorded by the research team.[60]

Finally, further evidence comes from economic theory. If the effects estimated in column 1 of Table 28 are not false positives, then this implies that each of the four treatment groups increased the total number of trips by around 20% (since reported trips in Table 27 are essentially the same in treatment and control). This is implausible for two reasons. First, if the treatments induced more trip taking, it must have been because the promised inducements were decreasing the effective cost of trips for drivers. But this can only be true if they anticipated bringing receipts and thus obtaining the incentives we offered.[61] Given the low rate of receipt-bringing, it is difficult to imagine that such a cost calculation could explain a 20% increase in trip taking. For example, drivers brought about .07 additional receipts per reported trip as seen in Table 11 when in either the FC1000 or FC2000 treatment group,

---

[59]This stands in stark contrast to the results of a regression of bringing a valid receipt on endogenously recalled treatment — recalled treatment strongly predicts receipt-bringing for all treatments — suggesting that this measure has valuable information content.

[60]Another interpretation is that drivers were confused about the validation survey, thinking that if they completed a survey at the toll they were not supposed to report it again at the office. However, it is difficult to explain why this would not have also applied for the control group.

[61]Or if there were large effects on average bribe paid, for which we do not find strong evidence.

on a base of .124 receipts. Even without any discounting of the reward (which we explore below), an expected return of 1000*.194 = 194 FC (or even 2000*.194 = 388 FC) is small. For reference, the cost of a trip from the center of town to one of the tolls usually costs about 5000 FC. Thus, for even a 388 FC cost difference to generate 20% more travel would imply an implausibly large back-of-the-envelope trip-cost elasticity of -2.6 (or -5.2 for the FC1000 treatment group).

Additional evidence for this theory comes from the social treatments (Charity and Government). The estimated misreporting rate is roughly the same across all four treatments. However, we find no effect from the Charity and Government treatments on receipt-bringing, so it is difficult to understand how these treatments would be causing drivers to take more trips, even if drivers did value the donations to the widow's group. That all four treatments groups have the same rate of estimated misreporting suggests that the lower rate of misreporting in the control group is more likely than not an artifact of noisy data. However, we wanted to flag the issue in the interest of maximal research disclosure. See Section 11.5.1 for more details and estimates of the elasticities (see Section 8.3 for calculations) when we assume that the misreporting is real.

11.5.1 Corruption Estimates Under Trip Misreporting

Even though we have reason to believe that the different rates of estimated trip misreporting in the treatment groups relative to the control is likely a false positive (for reasons explained above), in this section we consider its implications if true. If we assume that all unreported trips lacked receipts, then trip-level estimations will bias the effect of treatment upwards. This is because our results suggest that treated individuals were less likely to report their trips, which lowers the denominator of any regression with an outcome of receipts-per-trip. Thus, ignoring the unreported trips would lead us to conclude that the receipt-bringing rate in the treatment groups was higher than it actually was.

The best strategy to deal with this bias is to conduct our analysis on the individual-by-round level instead of the individual-by-trip level, as it is still well-defined to look at the effect of incentives on receipts-per-interview-period without controlling for or normalizing by the number of trips reported, as reported in Table 10 in the main text. We also then inflate the number of trips reported by our estimates of the rate at which the treatment groups tended to underreport trips, as reported in Table 28 in the main text. Together, these strategies allow us to

account for trip non-reporting in the elasticity calculations from Section 8.3, which we summarize in Table 29.

As we might have expected from our discussion earlier, these results are much closer to zero (and indeed can neither statistically nor economically be distinguished from zero). This is because, by taking the under-reporting results at face value, we end up with a rate of receipts per "true" trip that is indistinguishable between the control and treatment groups. Since the rate of receipt bringing is then unchanged, it is unsurprising that we observe an elasticity of zero. These results suggest that the responsiveness of corruption to financial incentives is even more inelastic than we estimated earlier, which only further deepens the puzzle. Indeed, these results suggest that driver's high value of time, combined with the partial coercive power of toll officers means that corruption may be even more important as a way to "grease the wheels" of this bureaucracy.

11.6 Heckman selection correction

As noted in Section 7.3, the Tobit methodology we employ in the body of the paper takes the model very seriously. A more general solution would be to use a Heckman selection correction with an instrument that predicts selection into bribery, but does not (separately) affect bribe levels. However, it is difficult to imagine a predictor of bribery that would not also affect bribe levels because the decision to bribe rather than get a receipt is directly a function of expected returns to each action under any bargaining framework.

Notwithstanding this concern, in this section, we consider a Heckman selection correction using the presence of the enumerator at the toll on the day of travel as our instrument. The logic for this instrument is that bribery happens outside the view of the enumerator, in a small hut by the toll. The presence of the enumerator at the toll may provide some impetus on the part of the driver or toll officer to issue a receipt; but, conditional on agreeing to seek a corrupt bargain, the negotiation happens outside the view of the enumerator, who thus does not have a direct impact on the outcome. We present this result not as an ironclad solution to the selection issue, but rather another attempt –along with the fixed effects and tobit strategies discussed in section 7.3 –with the two other methods to provide a broader picture of the possible effects.

Assigned treatment	Remembered treatment							
	Control	1000FC	2000FC	Charity	Gov.	Other	DK	Total
Control	4	20	26	18	8	8	16	100
1000FC	3	55	5	11	2	9	15	100
2000FC	3	2	68	5	4	8	10	100
Charity	5	7	14	52	0	8	14	100
Gov.	4	7	15	12	25	12	24	100
Total	4	19	25	20	8	9	16	100

Assigned treatment	Remembered treatment							
	Control	1000FC	2000FC	Charity	Gov.	Other	DK	Total
Control	23	9	11	9	3	7	37	100
1000FC	3	60	7	3	2	7	17	100
2000FC	2	3	68	6	1	9	12	100
Charity	4	4	6	59	0	8	19	100
Gov.	7	7	14	7	32	9	24	100
Total	10	15	20	16	7	8	24	100

Notes: This table shows the treatments recalled by participants (columns) for each treatment assigned to participants (rows). All values are in percent terms, such that the total for each treatment assignment adds up to 100% when adding across all of the columns (i.e. each cell shows the percent of people who recall the treatment listed in that column, among those that are assigned the treatment listed in that row). The top panel shows the results only for Follow-up Visit 2, while the bottom panel excludes Follow-Up Visit 2.

Table 18: Treatment remembered versus treatment assigned, as a percent of respondents assigned to that treatment group, in Follow-Up Visit 2 (top panel) and in the two other rounds (bottom panel)

12 Appendix Tables and Figures

12.1 Appendix Tables

	Remembered treatment							
Treatment from previous round	Control	1000FC	2000FC	Charity	Gov.	Other	DK	Total
1000FC	3	73	3	0	2	5	15	100
2000FC	7	3	77	0	0	10	4	100
Charity	4	4	10	53	1	9	19	100
Gov.	3	2	9	17	29	11	29	100
Total	4	20	26	18	8	8	16	100

	Remembered treatment							
Treatment from previous round	Control	1000FC	2000FC	Charity	Gov.	Other	DK	Total
1000FC	29	40	6	4	3	4	14	100
2000FC	37	3	40	8	0	2	10	100
Charity	20	6	12	40	0	8	14	100
Gov.	22	4	4	25	22	6	15	100
Total	27	14	15	19	6	5	13	100

Notes: This table shows the treatments recalled by participants (columns) for each treatment assigned to participants in the prior round (rows), restricted to the set of participants who were assigned to the control in the current round. All values are in percent terms, such that the total for each treatment assignment adds up to 100% when adding across all of the columns (i.e. each cell shows the percent of people who recall the treatment listed in that column, among those that were assigned the treatment listed in that row in the previous round). The top panel shows the results only for Follow-up Visit 2, while the bottom panel shows the results only for Follow-Up Visit 3.

Table 19: Treatment remembered versus treatment assigned in the previous round (not the current round), looking at people who were actually assigned the control in the current round (as a percent of those respondents), in Follow-Up Visit 2 (top panel) and in Follow-Up Visit 3 (bottom panel)

| | (1) | (2) | (3) | (4) | (5) |
VARIABLES	Original	No R2	No R2	No R2 Control	No R2 Control
FC1000	0.08**	0.11***	0.09**	0.11***	0.11***
	(0.03)	(0.04)	(0.05)	(0.04)	(0.04)
FC2000	0.08***	0.12***	0.14***	0.12***	0.09*
	(0.03)	(0.03)	(0.05)	(0.03)	(0.05)
Charity	-0.01	0.03	0.02	0.03	0.03
	(0.02)	(0.03)	(0.04)	(0.03)	(0.04)
Government	0.03	0.06*	0.03	0.06*	0.04
	(0.03)	(0.03)	(0.05)	(0.03)	(0.05)
Observations	1,830	1,246	989	1,444	1,203
R-squared	0.02	0.02	0.51	0.02	0.48
Ind. FE	N	N	Y	N	Y
Control avg.	0.121	0.0868	0.0714	0.121	0.112

Standard errors clustered by individual. $^{*}p<0.1,^{**}p<0.05,^{***}p<0.01$

Table 20: Effects of treatment removing Follow-Up Visit 2 in whole or in part

VARIABLES	(1) Original	(2) R1 only	(3) History	(4) Prior Receipt	(5) IV	(6) IV FE
FC1000	0.04*	0.07*	0.06**	0.04**	0.06**	0.05*
	(0.02)	(0.04)	(0.03)	(0.02)	(0.03)	(0.03)
FC2000	0.04**	0.13***	0.06***	0.04**	0.05**	0.02
	(0.02)	(0.04)	(0.02)	(0.02)	(0.02)	(0.03)
Charity	-0.01	0.07*	0.01	-0.00	0.00	0.00
	(0.02)	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)
Government	0.00	0.05	0.02	0.00	0.01	-0.01
	(0.02)	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)
Ever FC1000 in past			0.02			
			(0.03)			
Ever FC2000 in past			0.07**			
			(0.03)			
Ever charity in past			0.03			
			(0.03)			
Ever govt in past			0.05			
			(0.03)			
Ever past receipt				0.14***	0.44**	-0.14
				(0.03)	(0.21)	(0.62)
Observations	2,487	857	2,487	2,484	2,484	2,464
R-squared	0.00	0.02	0.01	0.02	-0.05	0.49
Ind. FE	N	N	N	N	N	Y
Control avg.	0.0961	0.0638	0.0961	0.0962	0.0962	0.0957

Table 21: Effects of treatment when controlling for treatment history.

VARIABLES	(1) Original	(2) Toll repeaters	(3) Toll repeaters - IV
FC1000	0.07**	0.10**	0.12***
	(0.03)	(0.04)	(0.05)
FC2000	0.07***	0.09**	0.10***
	(0.03)	(0.03)	(0.04)
Charity	0.00	-0.02	0.00
	(0.02)	(0.03)	(0.04)
Government	0.03	0.03	0.04
	(0.03)	(0.04)	(0.04)
Ever brought receipt previously			0.19
			(0.18)
Observations	1,723	1,063	1,063
R-squared	0.01	0.03	0.05
Ind. FE	N	N	N
Control avg.	0.124	0.142	0.142

Standard errors clustered by individual. $^*p<0.1,^{**}p<0.05,^{***}p<0.01$

Notes: All columns are estimated using equation (9), so all results include round fixed effects (not reported). The dependent variable is the fraction of reported trips for which the driver brought a valid (incentivized) receipt. Column (1) includes all trips. Column (2) restricts to drivers who report traveling past a toll that they had reported visiting at least once in the two months prior to our study. Column (3) instruments for having ever brought a receipt to a prior round using prior treatment history as instruments, while still restricting to trips where drivers passed tolls they had reported visiting in the two months prior to the start of our study. Standard errors clustered by individual. $^*p < 0.1,^{**}p < 0.05,^{***}p < 0.01$.

Table 22: Effects of treatment for toll regulars.

VARIABLES	(1) R1-R2	(2) R2-R3	(3) R3-R4	(4) Days between appointments
FC1000	-0.0254	-0.0251**	-0.0305	0.0102
	(0.0187)	(0.0124)	(0.0231)	(0.399)
FC2000	-0.0275	-0.0242*	-0.0232	0.141
	(0.0180)	(0.0129)	(0.0231)	(0.382)
Charity	-0.0396**	0.0143	-0.0202	-0.0407
	(0.0159)	(0.0200)	(0.0239)	(0.362)
Government	-0.0323*	-0.0246*	-0.0480**	-0.119
	(0.0175)	(0.0127)	(0.0198)	(0.358)
Constant	0.0528***	0.0318***	0.0702***	23.17***
	(0.0129)	(0.0105)	(0.0152)	(0.221)
Observations	905	856	835	2,486
R-squared	0.008	0.011	0.006	0.000
Avg. Dep. Variable	0.0319	0.0222	0.0503	23.17

Notes: This regression is estimated using equation (**??**), so all results include round fixed effects (not reported). For columns (1)-(3), the dependent variable is a dummy that is 1 if the driver stopped participating between the two survey rounds listed at the top of the column and 0 if the driver continued participating. The sample is restricted to the set of drivers who participated in the prior round of the survey. In column (4), the dependent variable is the number of days taken between survey visits, conditional on participating in the follow-up visit (averaged across all rounds). Standard errors clustered by individual. $^*p < 0.1,^{**} p < 0.05,^{***} p < 0.01$.

Table 23: Attrition by round.

VARIABLES	(1) Got a receipt	(2) Tried to get receipt	(3) Tried or got receipt
FC1000	0.08**	0.04	0.10***
	(0.04)	(0.03)	(0.04)
FC2000	0.09***	0.02	0.09**
	(0.03)	(0.03)	(0.03)
Reward: Charity	0.01	0.01	0.02
	(0.03)	(0.02)	(0.04)
Reward: Govt.	0.03	-0.00	0.03
	(0.03)	(0.02)	(0.04)
Observations	1,287	1,025	1,287
R-squared	0.01	0.01	0.01
Ind. FE	N	N	N
Control avg.	0.169	0.0842	0.239

Notes: All columns are estimated using equation (9), so all results include round fixed effects (not reported). In column (1), the dependent variable is the fraction of reported trips for which the driver brought a valid (incentivized) receipt. In column (2), the dependent variable is a dummy that is 1 if the driver reported demanding a receipt, but not receiving one, and 0 otherwise. In column (3), the dependent variable is 1 if the trip resulted in either a receipt or a demand for a receipt that was refused, and 0 otherwise. Columns (1) and (3) restrict the sample to the set of reported trips for which the driver reports receiving a receipt or the first trip for which he does not have a receipt. Column (2) restricts the sample to only the first trip for which a driver does not have a receipt. Standard errors clustered by individual. $^*p < 0.1,$ $^{**}p < 0.05,$ $^{***}p < 0.01.$

Table 24: Treatment effects on reporting attempting to get a receipt.

Underpaid toll	35.4
Left receipt at home	18.0
Lost receipt	16.5
Asked, but agent refused	9.3
Didn't ask for receipt	8.7
No toll on route	3.9
Exempt from toll	3.3
Machine was broken	1.2
Machine out of power	1.2
Other	0.9
No agents due to rebel group	0.7
Out of money	0.6
Boss or client took receipt	0.4
Total	100.0

Table 25: Self-reported reasons for missing a receipt from a trip.

	(1)	(2)	(3)
VARIABLES	Got a receipt	Reported losing receipt	Lost or got a receipt
FC1000	0.08**	-0.02	0.05
	(0.04)	(0.03)	(0.04)
FC2000	0.09***	-0.00	0.07*
	(0.03)	(0.03)	(0.04)
Charity	0.01	0.01	0.02
	(0.03)	(0.03)	(0.04)
Government	0.03	0.03	0.05
	(0.03)	(0.04)	(0.04)
Observations	1,287	1,025	1,287
R-squared	0.01	0.00	0.01
Ind. FE	N	N	N
Control avg.	0.169	0.163	0.305

Notes: All columns are estimated using equation (9), so all results include round fixed effects (not reported). In column (1), the dependent variable is the fraction of reported trips for which the driver brought a valid (incentivized) receipt. In column (2), the dependent variable is a dummy that is 1 if the driver reported losing a receipt, and 0 otherwise. In column (3), the dependent variable is 1 if the trip resulted in either a receipt or a (reported) lost receipt, and 0 otherwise. Columns (1) and (3) restrict the sample to the set of reported trips for which the driver reports receiving a receipt or the first trip for which he does not have a receipt. Column (2) restricts the sample to only the first trip for which a driver does not have a receipt. Standard errors clustered by individual. $^{*}p < 0.1,^{**}p < 0.05,^{***}p < 0.01$.

Table 26: Treatment effects on likelihood of reporting a lost receipt.

| | (1) | (2) |
VARIABLES	Any trip	Total trips taken
FC1000	-0.0134	-0.0517
	(0.0266)	(0.0500)
FC2000	-0.0188	-0.0480
	(0.0273)	(0.0722)
Charity	-0.0382	-0.0434
	(0.0270)	(0.0550)
Government	-0.0533*	-0.148***
	(0.0276)	(0.0544)
Constant	0.647***	1.084***
	(0.0241)	(0.0626)
Observations	2,487	2,487
R-squared	0.031	0.034
Control avg.	0.526	0.792

Notes: This regression is estimated using equation (**??**), so all results include round fixed effects (not reported). The dependent variable in column (1) is a dummy that is 1 if the driver reported one or more trips and 0 otherwise. The dependent variable in column (2) is a count of the number of trips reported by the driver. Standard errors clustered by individual. $^*p < 0.1,^{**}p < 0.05,^{***}p < 0.01$.

Table 27: Overall trip reporting.

| | (1) | (2) | (3) |
| | OLS | Recalled | If DK Treatment |
VARIABLES			
FC1000	0.198*	-0.0475	0.703***
	(0.116)	(0.115)	(0.226)
FC2000	0.200*	-0.115	0.794***
	(0.104)	(0.0982)	(0.115)
Charity	0.205*	-0.0140	0.499*
	(0.121)	(0.111)	(0.283)
Government	0.233**	0.0293	0.390
	(0.117)	(0.161)	(0.265)
Constant	0.282***	0.459***	
	(0.0961)	(0.107)	
Observations	169	169	31
R-squared	0.078	0.048	0.335
Trip non-reporting in control	0.404	0.404	0.404

Notes: This regression is estimated using equation (**??**), so all results include round fixed effects (not reported). The dependent variable is a dummy that is 1 if the driver did not report a given trip in his follow-up interview and 0 if he did report the trip. The sample is restricted to the set of trips that were validated by our enumerator at the toll. Column (1) shows the results using assigned treatment as the treatment. Column (2) shows the results using recalled treatment as the treatment. Column (3) restricts the sample to drivers who correctly recalled their treatment at their follow-up visit. Standard errors clustered by individual. $^{*}p < 0.1,^{**}p < 0.05,^{***}p < 0.01$.

Table 28: Trips recorded by validator, but not reported by drivers.

β	**FC1000**	**FC2000**	**Per FC**
0.5	-0.06	-0.08	-0.28
	[-2.41, 2.04]	[-1.23, 0.92]	[-1.29, 0.68]
0.63	-0.04	-0.06	-0.23
	[-1.94, 1.64]	[-0.99, 0.74]	[-1.04, 0.55]
0.74	-0.04	-0.05	-0.19
	[-1.64, 1.39]	[-0.84, 0.63]	[-0.88, 0.47]
1	-0.03	-0.04	-0.13
	[-1.14, 0.97]	[-0.58, 0.44]	[-0.61, 0.32]

Notes: This table shows estimates of the elasticity of corruption (not getting a receipt) with respect to instantaneous monetary cost of compliance (i.e. cost of getting a receipt), using the calculations described in equation (11). Standard errors reflect the empirical 95% confidence interval from 5,000 bootstrap iterations. Bootstrapping was done using a block bootstrap, blocked by individual participant. The rows show the results using different assumed values of beta. Column (1) shows the results using only the effects of the 1000 FC treatment. Column (2) shows the results using only the effects of the 2000 FC treatment. Column (3) uses both treatments simultaneously (treating them linearly).

Table 29: Estimated elasticities of receipt-getting with respect to instantaneous monetary reward for various values of β. Bootstrapped 95% confidence interval shown in square brackets below each estimate.

VARIABLES	(1) Heckman
Reward: 1000 FC	99.19
	(70.33)
Reward: 2000 FC	-4.48
	(70.24)
Reward: 2000 FC donation to charity	-3.98
	(75.88)
Reward: Gov. donates 2000 FC to charity	0.84
	(72.65)
Toll officer treatment in effect	-40.76
	(78.49)
Reward: 1000 FC	-0.21**
	(0.10)
Reward: 2000 FC	-0.23**
	(0.11)
Reward: 2000 FC donation to charity	0.10
	(0.12)
Reward: Gov. donates 2000 FC to charity	-0.11
	(0.11)
Toll officer treatment in effect	0.13
	(0.14)
Enumerator validating traffic at toll	-0.20*
	(0.11)
Observations	1,519
Ind. FE	N
Ind. Controls	N
Trip Controls	N
Control avg.	1520

Notes: The dependent variable is the bribe (amount paid not backed by a receipt) paid by the driver. In column (1), we show the results of a Heckman selection correction based on equation 9 where bribes are considered to be missing for all drivers who receive a receipt. The top panel shows the second stage results on bribes paid and the bottom panel shows the selection equation. The presence of an enumerator at the toll on the date of travel is the excluded instrument. Standard errors clustered by individual. $^*p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table 30: Heckman results on bribes.