The Efficient Deployment of Police Resources: Theory and New Evidence from a Randomized Drunk Driving Crackdown in India

Abhijit Banerjee[1], Raghabendra Chattopadhyay[2], Esther Duflo[1], Daniel Keniston[3], and Nina Singh[4]

[1]Massachusetts Institute of Technology
[2]Indian Instutite of Managment, Calcutta
[3]Yale University
[4]Rajasthan Police

Current version May 1, 2017. Please do not cite without permission from authors.

Abstract

A central question in law and economics is whether, with limited resources, police activity should be narrowly focused and high force, or widely-dispersed but of moderate intensity. Critics of intense "hot spot" enforcement argue that this approach will only lead to crime shifting to other locations or other times. But if law breakers take time to learn that enforcement has begun, the police may take advantage of this period to intervene intensively in the most productive location. We propose a model where criminals progressively learn about policing, and structurally estimate its parameters using a randomized controlled experiment on an anti-drunk driving campaign with the police department in Rajasthan, India. In each station, sobriety checkpoints were either randomly rotated on 3 main routes or fixed in the best route, and the intensity of the crackdown was cross-randomized. We find clear evidence of driver learning about the crackdown and strategically responding to it, causing rotating checkpoints to dominate the fixed location approach. Our structural estimates allow us to design the optimal policy: we find that 50% of checks should occur on roads with, ex-ante, less criminal activity. We estimate that crackdowns in rotating locations reduced night accidents in the area covered by a particular police station by 17%, and night deaths by 25% over a two month crackdown and 6 weeks following it.

1 Introduction

How to deploy limited law enforcement resources to maximal effect is one of the central questions of law and economics. One perspective advocates intense, pre-announced crackdowns, to take advantage of assumed increasing returns of arrest probabilities on criminal activity (Eeckhout, Persico and Todd, 2010; Ross, 1993) . Others argue that unless it is possible to police all locations all the time, the deployment should be randomized over time and across the potential locations where the wrongdoing could take place, because otherwise the wrongdoers would simply switch to the unpoliced locations and continue their activities with impunity (Clarke and Weisburd, 1994; Mookherjee and Png, 1994)

However this is only true in the full Nash Equilibrium. Potential criminals need to know the strategy the police is following in order to undo it. At the onset of short, intensive crackdowns, lawbreakers may not be aware that they are being policed and therefore continue to frequent their most favored locations, in which case the police will stop the maximum number of them by concentrating their efforts on the most crime-prone locations. Moreover even after criminals actually find out that those locations are being policed, they may assume that the intervention was temporary if, in their past experience, policing strategies involved temporary crackdowns. In that case they will continue to go back to their favored location for a while, until they are persuaded that the checking will continue long enough that it is worth permanently switching locations. Until such a time as this switch happens, it is optimal to continue to police the location that the law breakers have traditionally favored.

Therefore if the police are planning a new initiative involving policing hitherto unpoliced locations for a fixed period of time it is not clear whether it should simply focus on the most popular location or randomize across many potential locations. The answer will depend on, among other things, the speed of learning and potential criminals' beliefs about the duration of enforcement. If people learn about the new initiative fast enough, then randomizing enforcement across locations is optimal. On the other hand, if people learn slowly the police is better off focusing on the areas that most lawbreakers prefer.

In this paper we use a randomized field experiment combined with a structural model of learning to examine this issue. We study the effects of an anti-drunk driving crackdown in the Indian state of Rajasthan which was implemented in a randomized fashion. Road safety is emerging as one of the top public health problems in developing countries, and although the precise share of accidents due to alcohol is unknown, it is likely to be important. Until recently, the issue did not receive as much attention as other policing and health issues, and there is thus considerable uncertainty about the best way to organize police action against drunk driving.

This project was undertaken in collaboration with the Rajasthan police, who invested in equipment and resources to fight drunk driving but were uncertain about the right way to organize the campaign. In some areas, chosen at random, only the route judged most conducive for catching drunken drivers was chosen for the crackdown–we refer to these as "fixed" checkpoints. In others, checkpoints were randomly assigned across the three most popular routes–we refer to these as the "rotating" checkpoints. In yet other areas no crackdown was implemented, allowing us to evaluate the effectiveness of the crackdown itself. The intensity of the crackdown was also varied in two different ways. First, police stations involved in the intervention were randomly assigned to have 1, 2 or 3 checkpoints per week in their jurisdictions. Second, the duration of the crackdown was varied, with some stations randomly ending enforcement up to one month earlier than others. Finally, incentives for the policemen for doing those checks effectively was also varied at random.

The experimental results clearly show that checking works. Over the two to three months duration of the crackdown and the subsequent three months (starting with the last check), the number of accidents at night (when the checking took place) went down by about 17% and the number of deaths from accidents at night decreased by about 25%. This salutary effect comes mainly from the rotating check locations. Only those effects are significant and they are much larger, and in the case of accidents the difference is statistically significant.

We can start to understand the reason for the greater effectiveness of the rotating checks when we look at the number of drunk drivers caught each night of checking. Since fixed checks are conducted at the location best suited to catch drunken drivers, so absent avoidance effects, there should be more drunk drivers caught for fixed checks. In fact we see the opposite: after the first 2 weeks there are marginally fewer (rather than more) vehicles with a drunk driver apprehended at the fixed checkpoints. Furthermore, with fixed checks the number of drunks apprehended per night goes down with the number of past checks, with the number of weeks since the intervention started, and with the assignment to a higher weekly checkpoint frequency. All of these effects vanish with rotating checks. It is of course precisely the fact that the drunk drivers are still savvy enough to switch to other routes that gives random checks their advantage over fixed checks. If they simply stopped driving drunk then the only reason to check routes other than the most popular route would be if the response to additional checks at the favored route is subject to diminishing returns.[1]

Some corroborating evidence that learning matters comes from looking at the post-intervention period. First the effect on accidents of the randomized intervention continues over the period ninety days after the intervention was concluded. The effect on deaths is also of the same magnitude as that in the intervention period, though not significant. Second, the difference between fixed and rotating checks is even larger after the intervention period. The difference between random checks and the fixed checks is consistent with the idea that the end of checking is easier to detect with a fixed intervention. Overall however learning is sufficiently fast that a campaign that rotates across locations dominates.

The data from the experiment are informative as to which of the tested strategies was most effective, but also suggestive of a more general model of drunken driver behavior that could be applied towards calculating the optimal crackdown strategy. To inform this counterfactual policy choice, we develop a multi-armed Bandit model of driver behavior in which drivers choose actions both to maximize static payoffs as well as to learn about police strategies. We fit this model to the experimental data and structurally estimate the parameters of driver's initial beliefs and payoffs. These parameters, in turn, allow us to evaluate counterfactual enforcement strategies and select the one which would be most effective in preventing drunken driving. This contribution is, to our knowledge, the first time a Bandit model of learning has been applied in the policing literature.

The remainder of the paper is organized as follows: Section 2 provides some background on drunk driving and law enforcement against drunk driving in India. Section 3 describes our experiment. Section 4 presents a tractable model illustrating the factors affecting driver behavior and police effectiveness. Sections 5 and 6 present the data and reduced form results

[1] This was pointed out by Lazear (2005) and Eeckhout, Persico and Todd (2010). The latter suggest that there is in fact increasing returns, at based on data from Belgian speeding enforcement.

of the experiment. The structural estimation technique and results are presented in Section 7, along with the implications for the optimal policy. Finall, Section 8 concludes.

2 Background: Enforcing Drunk driving

Each year 1.2 million people die in traffic accidents worldwide, with as many as 50 million injured. A staggering 85% of these deaths happen in developing countries (Davis, et al. 2003). Moreover, death and accident rates are rapidly increasing in developing countries even though these are rates decreasing in the developed world (Davis, et al. 2003, WHO 2004). By 2030, traffic accidents will be the third or fourth most important contributor to the global disease burden, and will account for 3.7 percent of deaths worldwide, twice the projected share for malaria (Habyarimana and Jack 2009). Estimates of drunk-driving frequency vary widely across countries and across studies. The role of alcohol in road accidents is also difficult to measure, especially in developing countries where police often lack the manpower and technology to measure drivers' alcohol levels. The available evidence suggests, however, that alcohol does play a major role in traffic accidents. According to a review of studies conducted in low-income countries, alcohol is present in between 33% and 69% of fatally injured drivers, and between 8% and 29% of drivers who were involved in crashes but not fatally injured (WHO 2004).

Sobriety checkpoints have been evaluated by a number of studies in a wide variety of contexts, and the general consensus is that these checkpoints significantly reduce traffic accidents and fatalities. Several recent meta-analyses (Peek-Asa 1999, Erke, Goldenbeld and Vaa 2009, Elder, et al. 2002) suggest that sobriety checkpoint programs reduce accidents by about 17% to 20%, and traffic fatalities by roughly the same amount. These results are not entirely conclusive, however, since most of the existing literature has struggled with a variety of challenges and limitations. First, no previous research has been conducted in the framework of a randomized trial, leaving even the few studies that employ multivariate statistical analysis open to concerns of endogeneity based on the location and timing of the interventions. Second, the vast majority of research has been conducted in developed countries, and consists of increases in checkpoints over and above what is already a relatively high standard of enforcement. Thus little is known about the impact of carryout out sobriety checkpoints versus a counterfactual of essentially zero enforcement.

The law and economics and crime prevention literatures have investigated the impact of crackdowns more broadly, often under the rubric of "hot-spot" policing. Much of the literature argues for increasing returns in enforcement effort: that low levels of policing have essentially no effect, but intense inforcement can be effective in reducing crime. Perhaps the strongest experimental evidence comes from Sherman and Weisburd's (1995) evaluation of the the Minneapolis hot spot policing experiment, a randomized field trial that led to a 25%-50% decrease in disorder events in hot spots that received roughly twice as many police patrols as control areas. Unfortunately this experiment, like most others, provides no information on crime in areas from which the extra police resources were taken, thereby making a clear test of increasing returns impossible.

A natural concern (Sherman, 1990) lies in the possibility of crime displacement to other locations-if criminals are relatively mobile and do not have strong preferences over locations, then crackdowns may simply shift crime from one place to another and not decrease overall lawbreaking. Here, experimental evidence is mixed. While some studies find limited displacement effects of foot patrols in crime hotspots into surrounding areas, (Groff et al., 2011) others find that crime in control areas near hotspots actually decreased relative to more distant control areas (Green and Weisburd, 1995, Dell 2015). This heterogeneity in results is perhaps not surprising, since displacement effects depend heavily on the type of crime, for instance drug sales vs. drunken driving, and the police enforcement approach, in particular the length of the crackdown.

The key to understanding the trade-off between the increasing returns of concentrating enforcement and the diluting effects of displacement may be the speed at which criminals learn about police crackdowns and when they ended–obviously if they continue to believe that the crack-down continues when it has ended the effect would be more durable. Nevertheless, robust evidence on the speed of criminal learning is lacking. Sherman (1990) argues that many of the non-experimental drunken driving crackdowns show post-project decreases in accidents lasting even longer than the projects themselves. Similarly, the Minneapolis hot spot experiment shows some non-experimental results that hotspot level disorder remained lower a few hours immediately after the police visited and, on an aggregate level, during 6 weeks when the police lowered patrolling intensity. However, other studies find that program effects are short-lived and disappear even while the intervention is ongoing (Braga, 2001). Thus the ability of lawbreakers to adjust to police behavior-both in terms of reducing or shifting criminal behavior during a crackdown and in terms of returning to criminal behavior after a crackdown-remains a central and open empirical question in the field of criminal behavior.

2.1 Drunken Driving Enforcement in India

In India, highway safety laws of all kinds are generally enforced by fixed sobriety checkpoints manned by personnel from the local police station using a "selective breath checkpoint" methodology. Barriers are arranged on the roadway so that passing vehicles are forced to slow down, and the officers on duty signal selected vehicles to pull over as they pass through the barriers. If the checkpoint is intended to target drunken driving, police personnel then ask the driver a few questions on his or her identity, destination, etc., while observing the driver's demeanor and smelling his or her breath. If the police feel the driver may be drunk, then according to the official procedure they will order him or her to blow into a breathalyzer, following the results of which the driver is either charged or released[2]. The printed results of a handheld breathalyzer are considered sufficient proof of drunkenness in court. Once caught, drunken drivers' vehicles are confiscated by the police, and the driver must appear in court to pay a fine or potentially face jail time, although imprisonment is never observed in our data. The fine amount depends largely on the judgment of the local magistrate, with a maximum fine of Rs. 2000 (roughly $50) for the first offense . The driver must then return to the police station to recover the vehicle from the police lot.

[2]This protocol is commonly used in US sobriety checkpoints as well (Elders, Zaza 2002)

Even this official procedure leaves many factors undetermined and up to the discretion of the police manning the checkpoint. The choice of how many, and which, vehicles to pull over for questioning and potential testing is the most important. Ideally the police should target vehicles with the highest probability of drunkenness, and in conversations the police often noted that if they saw a vehicle with a family, or driven by elderly people they assumed that the driver would be unlikely to be drunk and let it pass. Unscrupulous police officers are faced with another decision: whether to follow the official ticketing procedure or instead to solicit (or accept) a bribe from the accused drunken driver. However, as long as our crackdown strategies do not differentially affect the integrity of the police, the legality of the punishment received by arrested drunken drivers is not key to understanding which enforcement strategy would be most effective. Drivers expecting either a bribe or a formal citation face a strong incentive to avoid being arrested for inebriated driving.

Sobriety checks of this kind were extremely rary prior to the intervention and in the control police stations during the intervention. The main reason for the absence of checks the fact that breathalyzers had not been widely distributed to police stations prior to the program, and without a breathalyzer the local police would have needed to take a suspect to the hospital for blood testing in order to be reasonably certain of securing a conviction under the 30mg/100ml threshold for drunken driving established by the Indian Motor Vehicles Act. The police who typically have a single vehicle per police station see that as too much trouble and as a result, most police stations other than some police stations in Jaipur, the state capital, essentially do not try to implement check points. Some time before our expeirment started, breathalyzers were distributed in treatment and control stations, but the intended change in police practices had not yet taken place when the experiment started. We observe that during the experiment control police stations hardly ever used them to carry out checkpoints. In the 925 nights that surveyors visited control police stations, on only 7 (.76%) occasions did they witness the police carrying out a checkpoint.

3 The Intervention

The anti-drunken driving program was implemented as a randomized control trial (RCT), consisting of a control group (with no program) and an intervention group with three overlapping experiments, each varying a different aspect of how the campaign was implemented:

First, the location of the checkpoints was randomized into one of two strategies: either the spot best suited to preventing accidents due to drunken driving, or rotated among three locations, with each night's site chosen at random. All locations were chosen by the local chief of police as the best suited to catching drunken drivers. Second, the frequency of the checkpoints was randomized to be 0 (for control stations), 1, 2, or 3 nights per week. Third, the duration of the program was randomized: the crackdown phased out in certain stations up to a month prior to the end in other stations.

The program took place in two phases, an initial pilot, from September-early October 2010, and a larger rollout from September to the end of November 2011. The initial pilot covered 2 districts and 40 police stations, and the second covered 10 districts and 183 police stations. Treatment status was assigned randomly, stratified by district, whether a station was located on a national highway, and total accidents between 2008-2010. The assignment of police stations to treatment groups during both rounds of the intervention is reported in Table 1. The 2010 and 2011 interventions were identical in implementation, with the exception that during the pilot all checkpoints occurred twice a week and the program lasted only slightly more than one month. In the analysis and results stated below, we combine data from both intervention periods and control for any time trends using month fixed effects.

Together with the enforcement strategy RCT, we also randomly varied the incentives of the personnel carrying out the checkpoint. Checkpoints were staffed by either police officers from the local police station (the status quo outside the intervention) or a dedicated team selected from the police reserve force at the district level. These special teams were monitored by GPS devices installed in their vehicles. The design is summarized in Table 1, and in the remainder of this section we explain the design choices in detail.

3.1 Checkpoint locations

To test the central hypothesis that rotating checks lead to more crime reduction than fixed checks, and the implicit mechanism of learning by potential and actual drunken drivers, we randomly assigned police stations to hold their checkpoints at either a single location, or a rotating set of three locations. In the single location group, the police station chief identified the best location in the station's jurisdiction for catching drunken drivers, and the checkpoints were carried out by either local staff or dedicated police lines teams at that location. Fixed checkpoints were carried out at the same place, same time, and to the greatest degree possible, on the same day every week, although scheduling difficulties occasionally made this impossible. In contrast, rotating checkpoints moved among the three best locations for catching drunken drivers, again as identified by the police station chief. Each police station's rotation was pre-determined in advance by the research team. In all cases, the three checkpoint locations were selected by police chiefs prior to assignment of the police station to a treatment arm.

The differences among locations, in particular that the third best location usually has far fewer passing vehicles than the first, affect our analysis of these two program options. In regressions estimating the overall impact of the different strategies we do not control for checkpoint location fixed effects or characteristics, since these are themselves an outcome of the program choice. When analyzing learning by the public, however, we seperate the 1st road from the 2nd and 3rd routes to see how public behavior at a specific checkpoint responds to the enforcement strategy.

3.2 Checkpoint frequency and duration

The variation in checkpoint frequency was designed to identify the shape of the relationship between the intensity of police enforcement and driver behavior. Discussions with the police determined that it was not feasible to carry out more than 3 checkpoints per week per police station, thus giving us the final randomization categories of 1, 2, or 3 checkpoints per week (and, of course, 0 in the control group). These frequencies were at the police station level, not the road level; for example in a rotating group police station with a frequency of 2 checkpoints per week each of the three roads would have a checkpoint twice every three weeks. Checkpoints were always held at 7:00pm-10:00pm in the evening which was when pea drunk drving took place accordung to the police.

The date of the last checkpoint, and hence the duration of the crackdown, was also randomized at the police station level. In the shorter, 2010 intervention stations phased out randomly over the first 2 weeks of October. In the longer, 2011 intervention, some police stations stopped checks after two months (Sept.-Oct.), while others continued for a full three months (Sept.-Nov.). This variation assists in seperately identifying the post-intervention effects of the crackdown from any seasonal trends in drunken driving, since at the same date different stations had either ceased to hold checkpoints or were continuing them. It also gives us more insights into potential drunken drivers' return to driving under the influence: in early December 2011 we conducted one final round of checks. At this point some stations had not carried out a prior checkpoint for a whole month, whereas others had performed checks only a few days earlier.

3.3 Checkpoint personnel

Previous work with the Indian government (Banerjee, Duflo and Glennerster 2010) and the Rajasthan Police (Banerjee, et al. 2012) suggests that the implementation of government initiatives often decreases dramatically in the medium term if the civil servants implementing the project are not sufficiently motivated. To gain further insight on the role of monitoring and motivation in project implementation, as well as to guard against a failure of the project due poor implementation, the anti-drunken driving campaign was carried out by two sets of police staff, with different motivation, monitoring, and characteristics. In one group, checkpoints were manned by the staff of the police stations under whose jurisdiction the checkpoint sites fell—the status quo in the Rajasthan Police. The second group were drawn from the Police Lines, a reserve force of police often considered to be a punishment posting. Police lines teams were monitored by GPS devices in their vehicles, and were informed that good performance on this assignment might improve their chances for transfer out of the police lines. The police lines teams thus differed from the station teams in their incentives, their monitoring, and perhaps also in their fundamental motivation and ability as police officers. We reserve a complete analysis of this portion of the intervention for future work, and simply control for it when relevant in the analysis of the other branches of the intervention.

4 Theory

4.1 Setup

If police enforcement strategy were fixed in perpetuity, and drivers knew this strategy, their problem would be a static choice of whether to cease drunken driving, continue to drink but avoid the checkpoints, or drive drunk on their favored routes in spite of the checkpoints. However, the assumption of complete information is an unlikely one in this context: potential drunken drivers did not know that there will be such an intervention and even after it started there was no official announcement. So potential criminals could only learn about it from their own experience or that of others. Furthermore, drivers are likely to believe that policing strategies are not fixed, but rather take the form of discrete "crackdown" periods of intense enforcement. They also did not know how long the intervention was meant to continue and were not told when it ended. Drivers could only make inferences about the duration of the program and whether it had ended based on their individual and collective experiences.

To illustrate the trade-offs inherent to designing policy in an environment in which agents can learn to avoid enforcement, we analyze a simplified version of the structural model that we take to the data in Section 7.1. The model falls into the class of "restless two-armed Bandit" models, in which agents learn about the payoffs of choices while the payoffs themselves are changing over time. A homogeneous set of agents of mass 1 choose between drunken driving, from which they get utility flow d, or an alternate activity from which they get utility s. We refer to this activity as "staying home" or "sobriety", although we later consider the possibility that their alternative may be drunken driving in another location. For the population relevant to this model, $d > s$, so that in the absence of any enforcement they would prefer to drink and drive. Checkpoints are held with instantaneous hazard rate λ, (which we refer to as "intensity") and the cost to the agent of getting caught drunk driving is $c < 0$. We assume that

$$d + \lambda c < s \tag{4.1}$$

so that if agents were sure that a crackdown is ongoing, this would be sufficient to (temporarily) end drunken driving.

Drunken drivers only learn about the existence of the crackdown when they first encounter a checkpoint. We assume that at this moment they know with certainty that a crackdown of intensity λ is happening, although they do not know for how long it has gone on. However, all agents believe that with the same probability $1 - p_0$ the enforcement is only temporary, and if this is the case the police have a hazard rate η of ending enforcement. Conversely, agents believe that with probability p_0 the police have adopted a strategy of permanently enforcing drunken driving laws.[3] Thus agents' decision to cease drinking and driving after the announcement of anti-drunken driving campaign may be only temporary. After some time, it may be optimal for them to return to drunken driving if they believe that the crackdown is likely enough to have ended.

The goal of the police is to choose λ and η to minimize the number of drunken drivers subject to a budget constraint. We assume, in this simplified example, that police commit to a $\{\lambda, \eta\}$ policy and that these values are public knowledge, but the start date and permanence of the crackdown is not. If checking is permanent then we show the police choice is trivial: they should implement the highest intensity of checking that their budget permits. However, if the police have a policy of temporary crackdowns, or if the budget is too low to support a permanent level of λ that satisfies equation 4.1, then there is clearly a

[3]While the case of a permanent "crackdown" is extreme, it allows for a simpler exposition of the main factors affecting agents' and policing strategy. Our structural model in Section 7.1 relaxes this assumption and allows drivers to learn about the permanence of the campaign from a continuous distribution of expected crackdown durations.

trade-off between a shorter, more intense crackdown versus a more prolonged, lower intensity crackdown. In this case, under restrictions detailed below, the optimal policy may be to implement the longest, lowest intensity crackdown that satisfies 4.1.

The analysis is complicated by the fact that drivers may make several switches between periods of sobriety and exploratory bouts of drinking and driving. To enable a clearer exposition, we assume that the parameters are such that after first learning of the enforcement campaign agents immediately commence a single temporary interval of sobriety before returning to drunken driving. If they subsequently encounter a checkpoint, they then choose to permanently cease drinking and driving. This behavior will always be optimal for sufficiently high values of p_0. Further details are in Appendix 11.1.

Consider a driver at time 0, immediately after he has first learned that a crackdown is in force. His valuation of not drinking and driving for τ periods before resuming is:

$$V(\tau;p_0) = \left(\frac{1-e^{-r\tau}}{r}\right)s + e^{-r\tau}\left(p_0\left[W_{perm}\right] + (1-p_0)\left[\left(1-e^{-\eta\tau}\right)\frac{d}{r} + e^{-\eta\tau}W_{temp}\right]\right) \quad (4.2)$$

including a subjective discount factor of r. This valuation incorporates the three possible states of the world upon his return to drunken driving. The campaign may be permanent, in which case he receives (the net present value of) utility W_{perm}. Alternatively, it may have ended during his interval of sobriety, in which case he enjoys utility d in perpetuity. Finally, it may be temporary, but still continuing at the time he returns to drunken driving; in this case he receives W_{temp}. Both W_{temp} and W_{perm} incorporate the driver's strategy of forever ending drunken driving if checked again, and can be written:

$$W_{perm} = \frac{d+\lambda\left(c+\frac{s}{r}\right)}{(\lambda+r)} \quad W_{temp} = \frac{\frac{d}{r}(\eta+r)+\lambda\left(c+\frac{s}{r}\right)}{(\eta+\lambda+r)}$$

Not surprisingly, holding all else constant, higher intensity of checking (λ) decreases both W_{perm} and W_{temp}, and longer lasting crackdowns (low η) decrease the value of W_{temp}. Solving equation 4.2 for the utility-maximizing sobriety interval τ yields

$$\tau^* = \frac{1}{\eta}\ln\left(\frac{(1-p_0)(\eta+r)\left(\frac{d}{r}-W_{temp}\right)}{r\left((1-p_0)\frac{d}{r}+p_0W_{perm}-\frac{s}{r}\right)}\right)$$

Intuitively, in deciding how long to temporarily stop drinking and driving, agents trade off two opposing forces: a longer delay postpones the time until they can start drinking and driving again, but it improves their expected utility of returning to illegal drinking if checking is temporary. A temporary halt to drunken driving will only be optimal if

$$(1-p_0)d + p_0 rW_{perm} > s$$

that is, if the long-run perceived value of attempted drunken driving is greater than the alternative activity at the driver's initial beliefs about the permanence of the crackdown. The more that agents believe that checking is permanent (p_0 high) the longer their interval of temporary sobriety will be.

The attractiveness of the outside option s also influences the duration of time during which agents temporarily cease drinking and driving. It can be shown that $\partial\tau^*/\partial(d-s) < 0$, meaning that the greater the utility that drivers get from drinking and driving on their preferred road where the police might catch them (relative to their outside option), the shorter will be their temporary period of sobriety. This has implications for our interpretation of s: if the alternative is to drink and drive on another (unpatrolled) road in which case the outside option is good ($d-s$ small) we would expect a long discouragement period. In contrast, if the outside option is to cease drinking entirely ($d-s$ large), we would expect a shorter discouragement period.

4.2 Government's objective function

The police want to minimize the number of people drinking and driving in all future periods. To accomplish this, they have a budget which allows for B checkpoints in net present value terms, discounted by the government discount factor ρ. For simplicity we assume the budget is flexible in the sense that the budget constraint must only bind in expectation, but not ex-post:[4]

$$B = \int_0^\infty e^{-\eta t}\eta\left(\int_0^t e^{-\rho x}\lambda dx\right)dt = \frac{\lambda}{\rho+\eta} \tag{4.3}$$

Because of the stationarity of the crackdown strategy, every driver's expected amount of future drunken driving after being apprehended is the same: we denote this as Δ. The police objective function can then be written as the sum of two terms,

$$\begin{aligned}\min_{\eta,\lambda} D &= \int_0^\infty e^{-\rho x}\left(\int_0^x \left(e^{-(\eta+\lambda)z}\eta\right)dz + e^{-(\eta+\lambda)x} + e^{-(\eta+\lambda)x}\lambda\Delta\right)dx \\ &= \frac{\eta+\rho}{\rho(\eta+\lambda+\rho)} + \frac{\lambda\Delta}{(\eta+\lambda+\rho)}\end{aligned} \tag{4.4}$$

The first term captures drunken driving by agents who have not yet encountered a checkpoint; if the crackdown is not permanent ($\eta > 0$) some will continue to drink and drive, oblivious to the crackdown. The second term is the discounted total of all lawbreaking by drivers who will, at some point, encounter a checkpoint.

If checking is permanent, then the government's objective function is simply:

$$D_{perm} = \frac{1}{\lambda+\rho} + \frac{\lambda}{(\lambda+\rho)}\left(e^{-\rho\tau}\left(\frac{1}{\lambda+\rho}\right)\right)$$

and since $\partial\tau^*/\partial\lambda > 0$, the optimal policy is to simply set $\lambda = B\rho$ as long as equation 4.1 is satisfied.

If the police choose a temporary crackdown, their policy choices are more complex. Once a drunken driver has been caught, his expected future amount of drunken driving is,

[4] In our later counterfactuals we impose that the budget constraint must also be satisfied ex-post.

$$\begin{aligned}\Delta_{temp} &= \left(1-e^{-\eta\tau}\right)\left[\frac{e^{-\rho\tau}}{\rho}\right]+e^{-\eta\tau}e^{-\rho\tau}\int_{0}^{\infty}e^{-(\lambda+\eta+\rho)t}\left(1+\eta\frac{1}{\rho}+\lambda 0\right)dt \\ &= e^{-\rho\tau}\left(\left(1-e^{-\eta\tau}\right)\frac{1}{\rho}+e^{-\eta\tau}\left(\frac{\eta+\rho}{\rho\left(\eta+\lambda+\rho\right)}\right)\right) \qquad (4.5)\end{aligned}$$

Again, there are two terms, the first accounting for the possibility that checking ends while the agent is abstaining from drunken driving, and the second for the chance he returns during the crackdown and may be caught. Holding τ constant, the comparative statics are straightforward: a more intense crackdown (high λ) and/or a longer lasting one (low η) both decrease the number of drunken drivers. With unlimited resources the police would naturally want to carry out constant checkpoints over a very long period.

However, a change in police enforcement policy will also affect drivers' decisions of how long to temporarily cease drinking and driving. This induces four countervailing forces on driver behavior: First, higher intensity checking clears offenders from the road faster, reducing the first term in equation 4.4. Second, a higher intensity λ increases τ which makes the return to drunken driving occur further in the future, which is desirable for a police department that seeks to end drunken driving in the present ($\rho > 0$). Third, if drivers recommence drinking before checking has ended, they face a higher chance of being caught, and then ending drunk driving permanently. On the other hand, a higher probability of encountering a checkpoint induces drivers to lengthen their interval of non-drunk driving τ, thereby increasing the probability that checking ends ends while they are staying home and that they continuing drinking and driving permanently after their return. Overall, the first three factors dominate, and an unconstrained police department would want to make checking as intense as possible.

The calculus changes dramatically when budget constraint is binding. Higher intensity crackdowns now necessitate shorter (expected) duration $1/\eta$, and induce *shorter* τ intervals from the agents. Conversely, the constrained choice of λ has no effect on drunken driving by un-caught drivers (the first term in equation 4.4) so the total number of drunken drivers collapses to:

$$D = \frac{1}{\rho\left(1+B\right)}+\frac{B}{1+B}\left[\frac{e^{-\rho\tau}}{\rho}\left(1-e^{-\eta\tau}\frac{B}{1+B}\right)\right]$$

and some algebraic manipulations show that $\partial D/\partial\lambda > 0$. Thus our previous result for an unconstrained police department is reversed: subject to the budget constraint in 4.3 the optimal police strategy is to have a prolonged crackdown of the minimal intensity necessary to dissuade drunken drivers.

We consider several extensions that move the simple theory developed above closer to the structural model estimated in Section 7.1. First, suppose that drivers are unsure about the intensity of the crackdown. They believe that with probability q_t it is λ_h, and with probability $1-q_t$ it is $\lambda_l < \lambda_h$. If $q_0\lambda_h + (1-q_0)\,\lambda_l$ is low enough, drivers will not immediately cease drinking upon being stopped, but will instead continue to drink and drive until their beliefs about q_t and p_t become pessimistic enough that they either temporarily or permanently stay home. This is likely to to occur sooner if checking is more intense, and thus drivers will begin their interval of sobriety when they are still relatively confident that the program is temporary (p_t low). However, applying the finding that $\partial\tau/\partial p_0 > 0$ from the simpler model, a higher λ intensity will cause, on average, lower p_t and shorter τ, and hence a more rapid return to the road. This reinforces the trade-off in the choice of λ: police can either clear the roads of drunken drivers quickly, but face a relatively rapid reversion, or tolerate more drunkenness in the short run in exchange for more permanent dissuasion.

We next consider the case in which drunken drivers can switch to less preferred roads where they hope that police enforcement will be lower. Let road A denote the best road for drunken driving (perhaps it is near the local bar), and road B designate an alternate less convenient choice, with respective utilities $d_A > d_B$. If drivers are unsure whether the police know their avoidance routes they may attempt to drink and drive on the less preferred road for a temporary period before either returning to the preferred road or, if stopped on the alternate road, stay home temporarily or permanently. Adding this option to the previous model causes the number of cases to become cumbersome to analyze algebraically[5] (though our numerical model below does allow for it). To provide clearer insights, we consider only the case when police are beginning a permanent checking routine, but drivers believe that enforcement might be temporary. We assume that priors are the same on each road, and that parameters are such that drunken drivers, once caught on road A, will cease offending for a period of exactly $\bar{\tau}$, before returning to drink and drive on road B. If d_A is sufficiently close to d_B, drivers will remain on road B for indefinitely long, since posteriors on checking risk on road A will always be weakly higher. If subsequently apprehended on road B, they stop permanently. The total number of drunken drivers on both roads are

$$\begin{aligned} D &= \int_0^\infty e^{-(\lambda_A+\rho)t}dt + e^{-\rho\bar{\tau}}\int_0^\infty e^{-\rho t}\left(\int_0^t \left(e^{-\lambda_A x}\lambda_A\right)e^{-\lambda_B(t-x)}dx\right)dt \qquad (4.6) \\ &= \frac{1}{\lambda_A+\rho} + e^{-\rho\bar{\tau}}\frac{\lambda_A}{(\rho+\lambda_B)(\lambda_A+\rho)} \end{aligned}$$

where we can once again divide the total number of law-breakers into two groups: the first term in equation 4.6, which enumerates drivers who have not been stopped on road A. The second term accounts for those who have moved to drink and drive on road B.

This simplified example captures many of the key trade-offs between fixed and rotating checkpoints. It clear that a fixed crackdown on road A with a high intensity λ_A will push drunken driving onto road B, while high intensity in B will allow the large number of drunken drivers already on A to continue undisturbed. The exact trade-off between λ_A and λ_B depends on the police discount factor and drivers' beliefs (through $\bar{\tau}$); a few examples illustrate the extreme cases. First, consider a highly patient police force with $\rho = 0$: in this case $D = 1/\lambda_A + 1/\lambda_B$ and the optimal choice is clearly to have equal enforcement

[5] If drivers' priors on checking on road B are sufficiently low, strategies will be as follows: drivers stopped on road A move to road B: if they encounter a checkpoint quickly, they cease driving temporarily, and then return to A and the analysis proceeds as above. If they are stopped again on B after a longer duration, they cease driving permanently. If not stopped on road B, they return to road A after some time. If subsequently stopped on A, they return to B where they continue indefinitely unless they encounter a checkpoint in which case they cease driving permanently.

However, if priors on checking on B are high enough, then different checking experiences generate another set of equally complex potential histories.

on all roads. On the other hand, if police are less patient and $\bar{\tau}$ is large, perhaps because drivers suspect that checking may indeed be permanent, then the optimal policy is to focus on clearing road A quickly, and thus to have high λ_A and low λ_B. This rationale would be strengthened further if $\bar{\tau}$ could be exogenously increased, for example through jail time.

We have so far neglected the possibilities of social learning about police enforcement or agents safely gaining information by exploratory driving while sober. Both of these would serve to speed up the impact of the crackdown on reducing drunken driving, and reduce its persistence after the checkpoints have ended. In the extreme case in which all potential law-breakers know exactly the beginning and end of the crackdown, their optimal strategy would clearly be to remain sober for exactly the duration of the enforcement, and the optimal crackdown would spread the minimal effective intensity over as long a period as possible. In particular, social learning eliminates the advantage that an intense crackdown has in rapidly informing criminals that enforcement is occurring, while sober driver eliminates the need for long-term precautionary sobriety. Thus to the extent we see slow learning, or long persistence in the data, we can infer that non-risky learning is not the dominant source of information for potential drunken drivers.

In summary, the model implies the following predictions on the reduced form impacts of the intervention:

1. Persistence: The impact of the crackdown will carry on after the last checkpoint is finished: some drivers will be temporarily sober, and others will have permanently ceased drinking and driving.

2. Scope: A crackdown implemented on just one road will be less effective than one implemented on all (or many) roads, even at a lower intensity.

   (a) Furthermore, the intensity of the crackdown will not matter for overall drunken driving if it is implemented on just one road, since no drivers are pushed into sobriety. Intensity does matter for multi-road checking, since higher intensity will clear the road faster and (potentially) induce some drivers to permanently cease inebriated driving.

3. Reversion and intensity: A more intense crackdown causes a shorter period of temporary sobriety, and thus a more rapid reversion after the crackdown has ended.

4. Reversion and scope: A crackdown implemented at many locations, forcing drivers into sobriety, will exhibit faster post-crackdown reversion to drinking and driving on the most preferred road than one in which drivers can avoid the police by drunken driving on an alternate route.

5 Data

To evaluate the effects of the anti-drunken driving campaign we draw on a combination of administrative data on road accidents and deaths, as well as data on vehicles passing and stopped at checkpoints collected by surveyors hired for this program.

5.1 Administrative data

This study's main results on accident and death rates are drawn from accident reports recorded by the police. For each accident on which data has been collected properly we know the police station, date and time of the incident, whether it was "serious" or not (according to the police's judgment, the criterion used here is subjective), the number of individuals killed or injured, and the types of vehicles involved. We also have some additional information on the location, including whether it occurred on a highway. Unfortunately we do not know whether drunken driving contributed to the accident.

We collected monthly accident data from August 2010 through October 2012. For January and February 2012 the data is not disaggregated by day and night–which is problem since the intervention was always in the evening. Our main results exclude these two months, but we also show pooled results for day and night together and then use the data from these two months.

Summary statistics are presented in panel A of Table 2, with statistics presented for control stations. The data, displayed at the police station/month level, shows that control police stations have roughly 0.12 accidents per day and 0.05 deaths. Of these, roughly 1/3rd occur at night. For lack of a direct measure of accidents caused by drunkenness the number of night accidents and deaths may provide an outcome that varies more with the level of drunkenness than the overall total.

5.2 Survey data

We supplemented the police administrative data with additional data on the implementation of the checkpoints collected by surveyors sent to monitor a set of randomly selected checkpoint locations both on nights when the police were conducting anti-drunken driving checking and on nights when they were not, as well as at locations near the control police stations identified as the best checkpoint sites prior to those stations being assigned to the control group. After arriving at the designated stretch of road, the surveyor counted the number of passing vehicles, categorizing them by type into motorcycles, cars, luxury cars, trucks, autorickshaws, buses, and other. If the police were conducting a checkpoint, the surveyor also counted the number of vehicles stopped and the number that proved to be drunk. Finally, the surveyor recorded the arrival and departure dates of the police from the checkpoint location.

At the end of the usual monitoring of checkpoint and non-checkpoint locations during the intervention, the surveyors also collected data from a special final round of checkpoints held in the week immediately after the main portion of the program had concluded. These checks were held once in all stations, regardless of earlier treatment or control assignment. On these nights police were asked to set aside their normal practice of stopping only vehicles with a higher probability of containing drunken drivers and conduct checks either randomly or at a fixed interval of cars (e.g. one in ten get stopped). Surveyors were present for all of these final checks, where they recorded both the rate of drunken drivers as well as the extent to which the police carried out checkpoints randomly.

The summary statistics of the data collected by these monitors is displayed in panels

B-D of Table 2. Panel B displays the number of vehicles passing by the check points on the average night, using surveyor counts from the locations identified by the control police stations as where they would have carried out the checkpoints. Fewer vehicles pass the second and third checkpoint locations than the first; this is particularly noticeable in the medians. Overall, police stopped 13.1% of passing vehicles, roughly 105 per checkpoint. The majority of these were motorcycles, of which 11.5% or 40 per checkpoint were stopped. Panels D shows the effectiveness of the crackdown in catching drunken drivers in treatment stations: on average police caught 1.85 drunk drivers per checkpoint, primarily motorcycles. Our best estimates of the true fraction of drunken drivers come from the final checking night when checks were conducted in a more systematic fashion in controls stations which had previously not had any enforcement activity. Panel E shows the results from these checks: a 2.23% overall drunkenness rate, and a rate of 3.36% for motorcyclists. Car drivers had substantially lower drunkenness rates, perhaps partly due to the fact that many cars in India are driven by professional chauffeurs whose employers would not tolerate drunkenness.

6 Reduced Form Results

We report program results as coefficients from OLS or fixed effect (FE) regressions of accidents and deaths on all program categories.

We begin with a simple summary of the effects of the program on accidents in Table 3. The primary outcome is the number of accidents and deaths in the area covered by the police station (including, and likely dominated by, the 3 main roads on which check points were to happen in rotating stations) during daylight, nighttime, or both. Table 3 displays coefficients on the interaction between a treatment station dummy (any treatment) and a dummy for during or 90 days post intervention, including police station and month fixed effects. We find a significant decrease in the number of nightly accidents (17%) and deaths (23%). There is no significant impact during the day, which is reassuring since all the check points took place at night, though as we will comment later, there could be a strategic reaction on the daily accidents.

One of the main hypotheses of interest is the difference between the effectiveness of rotating and fixed check points. Table 4 separates the treatment groups in these two subgroups (keeping the number of check points pooled for the time being), and, as before, interacting treatment indicators with variables denoting dates during or after the intervention, controlling for month and police station fixed effects. This tables makes clear that the impacts of the intervention on accidents and death are entirely driven by the rotating check points police station. At night time, accidents reduced by 29% and death 30% in the rotating police station, and both number are significant at the 10% or 5% level. In contrast, there is no significant impacts in the rotating police stations and the point estimates are much smaller (2 to 8 times). The difference is significant for night time accidents, though not for deaths. Note a surprising increase in day time accidents in rotating police stations.

Our learning model predicts that the effect of the intervention should be increasing over time, as more drivers are notice the checkpoints and begin avoidance, and should persist for some time after the intervention ends, since it is not clear it has ended and many drivers will have been temporarily or permanently pushed into sobriety. Table 5 provides evidence for these claims. We see that the decrease in nighttime accidents in rotating stations is equally strong both during and after the intervention. Coefficients on nighttime deaths show a similar pattern, though they are not significant. Interestingly, the disaggregated analysis in Table 5 shows an increase in daylight accidents in rotating check stations during the intervention, but no similar result in the fixed check stations. This may be due to dedicated drunken drivers who shift their alcohol consumption earlier in the day in the rotating stations, but can simply continue drunken driving at the same time on alternate routes in the fixed stations.

This evidence suggests that the basic insight of the crime literature is very powerful: checking at rotating locations do lead to a greater reduction in accidents and deaths (the ultimate outcome the police wants to affect). Figure 10.2a shows the number of drunk caught during a given checkpoint night in fixed and rotating location, as a function of the checkpoint frequency to which the station was assigned. We see that more drunks seem to be caught at rotating locations where there is 2 or more checking per week. This is striking because in the fixed location checking are always happening at the best location, while the rotating checks rotate between locations 1, 2 and 3. Moreover, in fixed locations, the number of drunks caught per night of checking declines with the intensity of checking (which was randomly assigned), while it is not the case in the rotating location, suggesting that in the fixed interventions drivers are shifting to alternative routes and doing so more when the frequency of checks is higher but in the rotating intervention there was no alternative route to switch to.

If learning is indeed going on, in the first two weeks of the intervention, fixed checks should lead to more drunks caught since it is performed at the best place. We investigate this in Figure 10.2b, which focuses only on week 1 and 2. At one check point per week, it is true: the number of drunk drivers caught is significantly greater. However this effect disappears when checking becomes more frequent. Even in the first two weeks, random checking does as well as the fixed location in stations with 2 or more 3 checks per week. This is probably because there is learning even within the first two weeks and it is easier to learn about the fixed location than when the location jumps around. More frequent checks speeds up learning and boosts this effect which counteracts the advantage of the fixed location. In other words, people appear to learn extremely fast about the police's strategy, especially when it is fixed and as a result, random checks tend to dominate even when the intervention is new. Figure 10.2c shows the data after 6 weeks, once learning is presumably complete: there is a clear pattern where fewer and fewer drunks are caught per night when the checking is more frequent at fixed locations, but not at random locations.

Table 7 present the results from regressions on the number of drunk drivers caught. Column 1 shows that, on the average night, 0.123 more drunk drivers, or 10% are caught at rotating check locations. This number is not significant due to large standard errors, but note that we should expect to see fewer, rather than more drivers apprehended, since the rotating checkpoints are held at less productive locations. That this is due to strategic avoidance behavior (rather than by people drinking less) is strongly suggested by the result in column 2. In fixed locations, the more frequent the checks, the fewer drunks are caught per night. The effect entirely disappears in random locations. This is presumably because with frequent checks at a fixed location, changing routes is an attractive option for drivers, but there is no desirable route to switch to in the case of rotating checks. Column 3 shows evidence of learning: as the number of weeks passes, the difference between fixed and rotating location increases. Finally column 4 puts the two together, and shows that as the number of past checks (which is a combination of the number of weeks and the random assignment to checking frequency) increases, the difference between rotating and fixed increases. Both columns 3 and 4 include police station fixed effects, so the impacts of different checking strategies over time are estimated from entirely within-police station variation.[6]

Columns 5-8 of Table 7, presents the analogous regressions using the number of passing cars and motorcycles as the outcome variable. Since these could be counted by surveyors posted at potential checkpoint locations in the control stations, we can estimate, in row 1, the overall impact of the intervention relative to no checking. Column 5 shows that this impact is substantial: compared to control stations the intervention caused a 28% decrease in vehicles passing the checkpoint locations. In contrast to the impact on drunken drivers, these effects are not significantly different in rotating versus fixed checkpoint stations. Estimates without station fixed effects (Column 6) are too noisy to be informative, but once station FEs are introduced we find very significant decreases in passing vehicles over time (columns 7, 8) in both the rotating and fixed stations. The large magnitude of the decrease in passing vehicles cannot be explained solely by drunken drivers, who comprise only 2-3% of all traffic. However we are limited in our ability to reconcile the arrest and passing vehicle data because we do not know how non-drunken drivers perceived the checkpoints: for example, if other violators (say drivers without a license) were also attempting to avoid checkpoints they may generate different patterns of traffic and arrests. Or they might just be reacting to the delays and harassment cased by the checkposts. Understanding the sources of this reaction is important from assessing the overall welfare implications of the intervention, but is beyond the scope of this study.

The checkpoint survey data provides more direct insights into the drivers' learning that the intervention had ended. Table 8 displays the results of the intervention on the number of vehicles passing the former checkpoint sites on the nights after the end of the intervention. Columns 1 and 2 show counts at the former checkpoint on road 1; consistent with the theory we find 43%-59% fewer vehicles from both rotating and fixed checks, approximately the same as during the intervention. Though evidence of reversion is noisy the coefficient in column 2 is much larger and close to significant only in police stations with rotating checkpoints; this is consistent with the theoretical prediction that driver's avoidance interval is shorter when the outside option to risking arrest on road 1 is relatively worse. The theory also suggests differential effects of rotating and fixed checks on roads 2 and 3, since drivers in fixed stations will move to these routes while drivers in rotating checkpoint stations will avoid them. Results in columns 3 and 4 confirm these predictions, with an increase

[6] A potential concern is that many scheduled checkpoints did not occur due to police non-compliance with the intervention. If this 'attrition' is correlated with the potential number of drunken drivers caught, it may bias the estimates of program impacts. We test this by recoding all nights when police did not perform a scheduled checkpoint as 0 drunk drivers caught. The results, in Table A1 are qualitatively the same as those in Table 7. While the absolute magnitude of coefficients in columns 3 and 4 are smaller, likely due to the introduced measurement error, the signs and relative magnitudes are the same as in the main results.

in passing drivers in fixed stations (though not significant, p = .166), and a decrease in rotating stations. Strikingly, as time passes after the intervention, the number of vehicles on roads 2 and 3 decreases in fixed stations, presumably as drivers revert to their preferred route. In rotating stations we see no significant signs of reversion.

Our final direct measure of the impact of the interventions occurred at the last night of checking. As describe above, this final check was always held at checkpoint location #2, regardless of the police station's treatment group assignment during the main intervention. Thus even fixed checkpoint stations and control stations held these final checks on route 2, despite having previously conducted them on route 1 (in the fixed stations) or never (in the controls). Since all checks in control police stations were conducted by teams from the police lines, we control for the type of police team conducting the checkpoint to avoid conflating the lines team and control station effects. Table 9 shows these results from the 109 police stations which conducted this check.[7]

Consistent with earlier results, we find strong evidence of drunken drivers avoiding checkpoints. Overall, the number of drunken drivers caught in a station with any intervention is 67% lower than in the control stations. We do observe significant signs of drivers returning to illegal drunkenness, but this is slow: at the estimated coefficient values it would take 91 days for the impact of the crackdown to fully dissipate. Column 2 shows that the avoidance effect increases with the frequency of checking, and that overall reversion may be increasing with the frequency of checking, although the reversion results are not significant. Results from the rotating checkpoint stations (columns 3 and 4) are not substantially different, with the exception that reversion may be slower but is now significantly increasing in the intensity of the crackdown. This confirms the model prediction that drivers temporary intervals of sobriety are shorter with more intense checking (if they expect the police to be budget constrained).

Perhaps surprisingly, the results from fixed checkpoints also show very significant decreases in final check drunken drivers relative to the control (columns 5, 6). Recall that these final checkpoints were at location #2, but that in fixed stations there was only checking at location #1. Thus the net effect of checks on one roads seems to have been to decrease them on others as well. As we will show below, this result is not fundamentally inconsistent with the model, since drivers might have had priors that checking on road 2 would be intense in the event of a crackdown. If they occasionally travel both roads, and notice a checkpoint on road 1, they would they conclude that any drunken driving might be too risky.

7 Structural Estimation

7.1 Empirical Model

While the simplified model presented in Section 4 gives insights into the mechanisms behind drivers and police decisions, mapping the results of the RCT into the optimal police enforcement strategy requires a more detailed model of driver behavior. We thus augment

[7]Since only 109 stations actually conducted the final check (60% of the total), there is a potential concern that compliance may be correlated with the intervention categories. Appendix Table A2 tests for differential compliance across intervention branches. We find no evidence to suggest this occurred.

the theory along several dimensions. First, we extend the set of road options to two: the most preferred road (A, or road 1 in the data) and all other roads (collectively designated B, or roads 2 and 3 in the data). We allow drivers to be uncertain about police checking intensity on each road, with driver i's beliefs about the probability of encountering a checkpoint distributed Beta $\left(\tilde{\alpha}_{irt}^{\lambda}, \tilde{\beta}_{irt}^{\lambda}\right)$, $r \in \{A, B\}$, where $\tilde{\alpha}_{irt}^{\lambda}$ and $\tilde{\beta}_{irt}^{\lambda}$ are the parameters of the Beta distribution of enforcement intensity on road r. Consistent with the notation above, we define the driver's expectation of the intensity of checking on road r as $\lambda_{irt} = \tilde{\alpha}_{irt}^{\lambda} / \left(\tilde{\alpha}_{irt}^{\lambda} + \tilde{\beta}_{irt}^{\lambda}\right)$. Instead of modeling driver's beliefs about the probability of a crackdown ending as discretely either η or 0, we now allow drivers to have continuous beliefs about the probability that the crackdown ends each period, distributed Beta $\left(\tilde{\alpha}_{it}^{\eta}, \tilde{\beta}_{it}^{\eta}\right)$ with expectation $\eta_{it} = \tilde{\alpha}_{it}^{\eta} / \left(\tilde{\alpha}_{it}^{\eta} + \tilde{\beta}_{it}^{\eta}\right)$. Both λ_{irt} and η_{it} are defined as conditional on the checking being ongoing at the start of the current period. Finally, drivers have beliefs about the probability that a crackdown is currently ongoing, π_{it}. These 7 variables, $\Psi_t = \left\{\tilde{\alpha}_{At}^{\lambda}, \tilde{\beta}_{At}^{\lambda}, \tilde{\alpha}_{Bt}^{\lambda}, \tilde{\beta}_{Bt}^{\lambda}, \tilde{\alpha}_{t}^{\eta}, \tilde{\beta}_{t}^{\eta}, \pi_t\right\}$ constitute the state variables in the driver's decision problem.

At the beginning of each period, agents choose whether to drink and drive on their preferred road A, a less preferred road(s) B, or to abstain from drinking and driving by staying at home.[8] They perform their highest-utility action, and observe the results: if they did drink and drive they may or may not witness a check point, while if they choose to remain home they learn nothing about police activity on that night. Based upon these observations, they update their beliefs on the state of checking during period t. Their posterior beliefs are then adjusted for the possibility η_t that the crackdown ended between periods, and become the priors in the next period, Ψ_{it+1}.

We assume drivers' behavior is not entirely deterministic: on certain nights they may experience random shocks that cause them to return to drinking and driving when it would otherwise be too risky, or to stay at home when they might otherwise go out for a drink. These shocks then affect driver's actions, and their information about police activity, and hence their beliefs in future periods. Thus even if all drivers have the same priors at the beginning of the crackdown (as we assume they do), by period t they will have a distribution of state variables $h\left(\Psi_0, x_t, \epsilon_t\right)$ where x_t denotes the history of checkpoints the station jurisdiction from periods 1 to t, and ϵ_t denotes the past history of shocks experienced by all drivers.

If the utility of driving drunk on road r is d_r and the expected cost when encountering a police checkpoint is c,[9] the static expected utility of driving on road r at time t is

$$u_r\left(\Psi_{it}\right) = u_r = d_r + \pi_{it}\lambda_{irt}c \tag{7.1}$$

where we normalize the static utility of staying at home to 0. The evolution of driver's dynamic utility depends upon the outcome of driving: For example, if the driver witnessed

[8] We continue to assume that the cost of exploratory driving while sober is always above the return it would generate, and thus do not include this option.

[9] The c parameter encompasses multiple factors affecting driver's disutility of encountering a checkpoint: their subjective probability of being detected by the police, their costs in time and money if they are detected, the time and anxiety of waiting to be inspected at a checkpoint even if not detected, etc. Since our experiment is not well-suited to disentangle these factors, we simply model the aggregate expected disutility.

a police checkpoint on road A, we denote his posterior beliefs as Ψ^a_{it+1}, if he drives on A but sees no checkpoint, then his beliefs become $\Psi^{na}_{i,t+1}$. Beliefs following travel on road B are defined analogously as $\Psi^b_{i,t+1}, \Psi^{nb}_{i,t+1}$, and those staying at home are $\Psi^h_{i,t+1}$. Dynamic utility is,

$$V(\Psi_{it}) = \max \begin{cases} v_A(\Psi_{it}) & = u_A + \delta\left(\pi_{it}\lambda_{iAt}V\left(\Psi^a_{i,t+1}\right) + (1-\pi_{it}\lambda_{iAt})V\left(\Psi^{na}_{i,t+1}\right)\right) + \epsilon_{iAt} \\ v_B(\Psi_{it}) & = u_B + \delta\left(\pi_{it}\lambda_{iBt}V\left(\Psi^b_{i,t+1}\right) + (1-\pi_{it}\lambda_{iBt})V\left(\Psi^{nb}_{i,t+1}\right)\right) + \epsilon_{iBt} \\ v_H(\Psi_{it}) & = \delta V\left(\Psi^h_{i,t+1}\right) + \epsilon_{iHt} \end{cases} \tag{7.2}$$

The evolution of beliefs is a simple application of Bayes rule, thanks to the Beta distribution and our specification of beliefs regarding the intensity of checking (λ_{irt}) and the probability of a crackdown ending (η_{it}) as conditional on the crackdown being ongoing at the beginning of the period. Consider a driver who has chosen to take road r in period t, and saw no checkpoints. Posteriors beliefs on the expected checking intensity are

$$\lambda_{ir,t+1} = \frac{\tilde{\alpha}^\lambda_{irt}}{\tilde{\alpha}^\lambda_{irt} + \tilde{\beta}^\lambda_{irt} + 1}$$

where the $\tilde{\alpha}^\lambda_{irt}$ term in the numerator is replaced with $\tilde{\alpha}^\lambda_{irt} + 1$ if the driver instead passes a checkpoint. If the driver had chosen to remain at home, her beliefs about checking intensity would not evolve: $\lambda_{irt} = \lambda_{ir,t+1}$

Similarly, the posterior distribution of beliefs on the probability of the crackdown ending, η_t is

$$\eta_{t+1} = \frac{\tilde{\alpha}^\eta_0}{\tilde{\alpha}^\eta_0 + \tilde{\beta}^\eta_t + 1}$$

In this case η_{t+1} evolves deterministically based solely on the time passed since the driver first encounters a checkpoint, and thus does not depend upon the choices made by the driver or the outcomes that she witnesses: conditional on the crackdown being ongoing in period t, $\Pr[\eta = \tilde{\eta}|X] = (1-\tilde{\eta})^t$. Intuitively, as time passes drivers infer that if the crackdown is still ongoing then the chance that it ends in each period is likely to be increasingly low, and update their beliefs about η and the future duration of the crackdown accordingly.[10]

The final element about which the driver learns is the probability that the crackdown is indeed in progress. If the driver encounters a checkpoint on either road, she knows with certainty that the crackdown was in force during period t, and thus that the probability of checking in the next period is $\pi_{i,t+1} = 1 - \eta_{it}$. If she drives on road r but sees no police checkpoint, then her posterior reflects the fact that this observation may either be because the crackdown has ended, or because the police did not implement a checkpoint on that night although the crackdown is still ongoing. Thus her posterior belief, adjusted for the fact that the crackdown might end between t and $t+1$ is

[10] Since the driver never conclusively observes the end of the crackdown, $\tilde{\alpha}^\eta_t = \tilde{\alpha}^\eta_0$ for all t. This reduces the number of endogenously evolving state variables to 6.

$$\pi_{it+1} = (1-\eta_{it})\frac{\pi_{it}(1-\lambda_{irt})}{\pi_{it}(1-\lambda_{irt})+(1-\pi_{it})}$$

This learning behavior can generate complex patterns of travel and staying home by the simulated drivers. To illustrate these strategies, Figure 10.1 shows the simulated histories of two drivers, the first in a station with only checking on road A, and the second in a station that also has checking on alternative roads B. The parameters used for the simulation are those estimated from the data, but for illustrative purposes we set $\epsilon_{irt} = 0 \ \forall r$ so that all changes of road are due to endogenous driver choices rather than exogenous shocks. The sequence of checkpoints are also chosen from two real stations in the data. The background of the graphs denotes the road on which the driver is traveling: green for road A, yellow for road B, and red if the potential driver is staying home. The solid black line (and the left-hand axis) shows his belief regarding the probability that the checking is ongoing at the beginning of the period, π_t. The dashed blue line and dotted red line show, respectively, λ_{At} and λ_{Bt}, and their scale is shown on the right-hand axis.

Both drivers encounter a checkpoint on the first night of the crackdown, causing them to shift to road B. Encountering the checkpoint causes their beliefs about λ_A and π to spike upwards, but their initial priors are such that encountering a single checkpoint results in only a perceived ∼20% probability that the crackdown will be in force the next night. Nevertheless, the driver in Panel 1 remains remains on road B for two more nights as π_t continues to fall. Note that during this period his belief on λ_A remains constant, since traveling road B gives no new information on the intensity on A. He then returns to A for 5 more nights, during which time both π and λ_A decrease. However, he is stopped again on night 8, which leads to a larger spike in λ_A and π, and a much longer stint on road B. The process repeats once more, and the driver finally remains on road B until period 57, at which point he correctly infers that checking has ended on road A and returns to his preferred route.

The driver in panel 2 encounters a checkpoint on road B shortly after her initial switch. However, because estimated priors on crackdown intensity on road B are less malleable than on road A, this does not cause her to cease drinking and driving. Instead, she later returns to A, and it is only after she is stopped 3 more times that she briefly gives up driving altogether in period 17. However by this point her posterior on λ_B and π has risen considerably, so that each subsequent checkpoint she encounters on B causes her to remain home for longer and longer intervals. After encountering a final checkpoint in period 56, she forgoes driving for 13 nights until she returns to road B, and after 11 more nights, to road A. Note that height of the spikes in π_t become greater and the (negative) subsequent slope of the π_t graph becomes less steep after each successive checkpoint. This is caused by drivers' posteriors about η_t becoming lower and lower. Since driving reveals more information about the crackdown than staying at home, the slope of π_t become steeper during periods when the agent is actively driving versus when she is at home–this is particularly evident in the switch around period 69 in panel 2.

Estimation

We estimate the structural model using an simulated method of moments approach. In short, we first estimate a set of moments from the data, measuring the within-station changes over time in the numbers of drunken drivers apprehended. We then simulate a large number of potential drunken drivers in each station, using the model to predict their behavior as they encounter the same sequence of checkpoints as the real-life drivers. We re-estimate the same set of moments using this simulated data, and search the parameter space until we find the vector that best matches the empirical to the simulated moments.

The model contains two sets of parameters. The first are the initial conditions of drivers' beliefs, $\theta_1 = \left\{\tilde{\alpha}^{\lambda}_{A0}, \tilde{\beta}^{\lambda}_{A0}, \tilde{\alpha}^{\lambda}_{B0}, \tilde{\beta}^{\lambda}_{B0}, \tilde{\alpha}^{\eta}_{0}, \tilde{\beta}^{\eta}_{0}\right\}$. Since there had never previously been an anti-drunken driving crackdown in the area where the project was implemented, we set $\pi_0 = 0$. In addition to θ_1, we estimate three preference parameters: the driver's utility of drunken driving on roads A and B, and their disutility of encountering a checkpoint: $\theta_2 = \{d_A, d_B, c\}$. Combined, these form the parameter vector that determine driver behavior and hence allow the calculation of the optimal police crackdown strategy: $\theta = \left\{\tilde{\alpha}^{\lambda}_{A0}, \tilde{\beta}^{\lambda}_{A0}, \tilde{\alpha}^{\lambda}_{B0}, \tilde{\beta}^{\lambda}_{B0}, \tilde{\alpha}^{\eta}_{0}, \tilde{\beta}^{\eta}_{0}, d_A, d_B, c\right\}$.

Numerically solving the model is complicated by the relatively high number (6) of continuous state variables and the infinite-horizon nature of the problem, making backwards induction impossible. Because this case falls into the category of restless correlated Bandit models, the standard Gittins Index results that simplify calculating the optimal solution in more standard Bandit models do not apply. Many researchers have devised heuristic methods to solve bandit problems which asymptotically approach the optimal solution in maximizing payoffs as the number of periods increases. However, agents' exploration behavior in these heuristic approaches may be very different than under the optimal strategy. We thus calculate the optimal strategy using value function iteration, but do so on a finite grid of the state space. When future states fall between the grid points, we interpolate.

To estimate the model, we make two additional assumptions. First, we assume that the choice-specific shocks in 7.2 are distributed IID extreme value type 1. These additional shocks allow us to rationalize observing a few drivers continuing to be drunk on routes that the model would otherwise tell us that all drunken drivers should have abandoned for fear of checking. They also smooth the objective function, which would otherwise consist of a step function in which large masses of drivers might change behavior with a small change in parameters Incorporating these shocks, the fraction of potential drunken drivers who choose to drive on road r on night t is μ^r_t:

$$\mu_{srt} \quad = \quad \int_i \frac{\exp\left(v_r\left(\Psi_{ist}\right)\right)}{\sum_{r'}\exp\left(v_{r'}\left(\Psi_{ist}\right)\right) + \exp\left(v_H\left(\Psi_{ist}\right)\right)} dh\left(\Psi_{ist}; \Psi_0, x_{s,t-1}, \epsilon_{s,t-1}\right) \qquad (7.3)$$

Second, we assume that each police station has a constant mass of *potential* drunken drivers M_s, and that the local police's efficiency, that is their probability of catching a drunken driver passing the checkpoint, is a constant Ω_s. Other time-variant factors that might affect the number of violators apprehended are captured in the vector w_{st} , with their effects parameterized by the vector ϕ. Thus conditional on an agent choosing to drink and drive on a road where the police have a checkpoint, their probability of getting caught is

$\Omega_s \exp(\phi' w_{st})$. The total number of drunken drivers caught on road r, night t is then

$$N_{srt} \quad = \quad \mu_{srt} * M_s * \Omega_s * \exp(w'_{st}\phi) \tag{7.4}$$

The choice of moments is closely linked to the design of the experiment. The first set of moments are estimated from the "fixed" branch of the experiment: we calculate the percentage change in drunken drivers caught from the first 3 weeks of the program to the 3rd-7th weeks of the program, as well as the percentage change from the first 3 weeks to weeks 8+ of the program. We calculate these 2 ratios separately for each of the 3 intensity branches of the treatment: the 1, 2, and 3 checkpoints per week sets of stations. Since the fixed checkpoints were only held in the best location for catching drunken drivers ("road A"), this generates a total of $2 \times 3 = 6$ moments. Similarly, for the "dispersed" wing of the treatment, we calculate the same 6 moments for checkpoints on road A, as well as the analogous moments for the checkpoints on the other roads (roads B). This generates an additional $2 \times 2 \times 3 = 12$ moments from data on the dispersed stations. Finally, we take the ratio of the number of drivers caught on roads A and road B, for one additional moment. This gives a total of $6 + 12 + 1 = 19$ moments.

Police stations differ both in their baseline numbers of potential drunken drivers M_s, as well as the effectiveness of the local police, Ω_s. Since this variation is not informative for estimating our model, we control for these factors by estimating the model including police station fixed effects. In addition to these fixed effects, the number of drivers arrested also depends upon the number of vehicles stopped by the police, the selectiveness of the police in stopping vehicles, and whether that particular night was desirable for drunken driving. To control for these time-variant factors, we estimate the empirical moments controlling for 3rd order polynomials of the number of passing vehicles, number of vehicles stopped by police, and fraction of vehicles stopped by police in the w_{st} vector. We also include controls for the day of the week.

Let z_{srt} be a vector of indicator variables denoting the particular treatment$\times$time period$\times$road category into which the data from the checkpoint in station s on night t falls. The omitted category, for which $z_{srt} = 0$ indicates checkpoints on road A in the first 3 weeks of checking. We estimate the empirical moments, which also correspond to the reduced form ITT impacts of the intervention, with the following specification.

$$N_{srt} \quad = \quad \exp(z'_{srt}\beta)\, M_s \Omega_s \exp(w'_{srt}\phi) + \zeta_{st}$$

Assuming ζ_{st} is distributed Poisson , we can estimate the reduced form impacts of the intervention β and the effects of other time-variant checkpoint conditions ϕ with a fixed effects Poisson specification where $D_s\Omega_s$ denote police station-level fixed effects. Let T_s be the number of nights that surveyors observed checkpoints in station s. After concentrating out the fixed effects, the log likelihood is (see, eg. Wooldridge, 1999)

$$\ln L(\beta, \phi) \propto \sum_s \sum_t N_{rst} \ln \left(\frac{\exp(z'_{srt}\beta + z'_{srt}\phi)}{\sum_{\tau=1}^{T_s} \exp(z'_{sr\tau}\beta + z'_{sr\tau}\phi)} \right) \tag{7.5}$$

which is the criterion that we maximize.

The estimated $\hat{\beta}$ reduced form coefficients have a straightforward interpretation as the differential changes over time in the number of drivers caught under different checkpoint strategies. Let $\beta_{rp\lambda}$ be the reduced form coefficient on the element of m_{st} that indicates data from road r in period p, in a station that had c checkpoints per week. Holding all else constant, the ratio of drunken drivers caught on road r in period p versus period 0, in stations with intensity c is then

$$\frac{\mathbb{E}\left[N_{rpc}\right]}{\mathbb{E}\left[N_{r0c}\right]} = \frac{D_s \exp\left(z_{srt}\hat{\beta}\right)\Omega_s \exp\left(w'_{st}\phi\right)}{D_s\Omega_s \exp\left(w'_{st}\phi\right)} = \exp\left(\hat{\beta}_{rpc}\right)$$

To estimate the parameters of the model, we maximize the analogous fixed-effects Poisson likelihood using data simulated from the model. For each station in the data we generate $H = 2000$ simulated histories of potential drunken drivers' choices as they potentially encounter the sequence of checkpoints that was randomly assigned to that station in the intervention. The simulated probability of a potential drunken driver being on road r in station s on night t is then the analogue to equation 7.4,

$$\tilde{N}_{srt} = \tilde{D}_s\tilde{\mu}_{srt}\tilde{\Omega}_s \exp\left(\tilde{z}'_{srt}\phi\right)$$

where $\tilde{\mu}_{srt}$ is defined as in Equation 7.3, with the integral over the distribution of state variables approximated by the mean of the 2000 simulated agents.

Since all empirical moments are estimated with station-level fixed effects, in the simulations we can normalize the fixed station-level mass of drunken drivers and police efficiency terms to one: $\tilde{D}_s = \tilde{\Omega}_s = 1$. Similarly, we hold the intensity of police checking, number of passing vehicles, and day of the week constant by setting $\exp\left(\tilde{w}'_{srt}\phi\right) = 1$. Thus the simulated number of drunken drivers on road r, station s, night t is just $\tilde{N}_{rst} = \tilde{\mu}_{rst}$, which we the use as the "data" to estimate the analogue of equation 7.5 on the simulated data:

$$\ln L\left(\tilde{\beta}\right) \propto \sum_s \sum_t \tilde{\mu}_{rst}\left(\ln\left(\frac{\exp\left(m'_{srt}\tilde{\beta}\right)}{\sum_{\tau=1}^{T_s}\exp\left(m'_{sr\tau}\tilde{\beta}\right)}\right)\right)$$

Finally, we choose θ to minimize the criterion $Q\left(\theta\right) = \left(\hat{\beta} - \hat{\tilde{\beta}}\left(\theta\right)\right)^2$ with respect to the vector of structural parameters θ.

7.2 Results: Structural

The estimated moments are shown in Table 10. The coefficients from the more flexibly estimated FE Poisson model qualitatively match the results of the more constrained OLS regressions discussed earlier: We see a rapid decrease in drunken drivers caught on road A in the fixed stations, and smaller and slower decreases on roads A and B in the rotating stations. Perhaps surprisingly, there are signs of a non-monotonic relationship between effectiveness and frequency on road A, with stations having 2 checks per week showing smaller declines over time in drunken drivers than both 1 and 3 check per week stations. Closer examination reveals that this is driven by low numbers of drunken drivers caught in the first 2 weeks in these stations, a phenomenon that, as we show below, our model struggles to explain.

Table 13 presents the results estimation of the structural parameters. To simplify the interpretation of the results, we report functions of the parameter that are more easily interpretable than the fundamental parameters of driver's beliefs.

Rows 1-3 display the parameters of the drivers' utility functions. Since the location and scale of these coefficients are determined by the value of not drinking and driving (normalized to 0) and the standard deviation of the shocks ($\pi/6$) the values of these coefficients per-se cannot be readily interpreted. The relative sizes are informative, however: the value of driving on road A is somewhat larger than road B, as expected by the police staff who ordered the routes. Both are much smaller than the implied disutility of encountering a checkpoint, showing that sobriety checkpoints are an efficient tool for crime prevention of deployed effectively. The small utility of drinking and driving (relative to the shock variance) implies that, on the average night with no risk of police enforcement, 42% of potential drunk drivers choose road A, 39% road B, and 19% remain sober. While these results, taken alone, are not surprising, they provide in inputs to designing a more effective strategy for preventing drunken driving.

Rows 4-9 show the initial priors on the expected intensity of checking on roads A and B, and on the probability that the crackdown ends in each night. Drivers seems to expect that the checkpoints are one-off events: their beliefs on the probability of a checkpoint on either road are close to zero (rows 4, 6) and they believe that the crackdown has a 91% chance of ending after 1 night. Rows 5, 7, and 9 report a measure of the precision of driver's priors about checking intensity and duration. These can be interpreted as analogous to the number of trials in the standard Beta model of the probability of success, or the number of trips the driver has made on each road prior to the beginning of the program. Drivers learn most quickly about the duration of the crackdown, with an estimated $\tilde{\alpha}_0^\eta + \tilde{\beta}_0^\eta = 4.3$. Learning about the intensity on road A is also reasonably fast, with an estimated 22.5 "prior trips". In contrast, learning about crackdown intensity on road B is very slow, with estimated 93.3 "prior trips". This parameter estimate is necessary to rationalize the empirical fact that drunken driving is very slow to respond to the crackdown on road B. It may also be due to the fact that "route" B is actually a 2 roads, and thus harder to learn about than A.

An advantage of our method of moments approach is that we can directly examine which empirical moments are best fit by the model. Since there are more moments (19) than parameters (9) not all results of the experiment will be perfectly explained. The simulated moments at the parameter vector that best fits the data are displayed in column 3 of Table 10. Overall, the data fits the model quite well. Clearly the main descrepancy lies in the explaining the greater numbers of drunken drivers caught in 2-night/week stations versus 1-night/week stations. While the model is capable of generating this non-monotonicity at other parameter values, the vector of parameters that best fits the rest of the data implies that the number of violators arrested per night will be decreasing in enforcement intensity. However, we note that the numbers of drunken drivers caught at intensity 1 versus 2 are usually not statistically different.

The final checks offer an additional opportunity for an out of sample test of the model. Table 12 shows moments estimated from the final check data, including the decrease in drunken drivers caught at each branch of the intervention (relative to control) and the rate at which this decrease dissipates over time. As in the more parametric results in Table 9, we see a roughly equal decrease in the number of criminals apprehended at both the rotating and fixed stations. Column 3 displays the results implied from the model estimated from the data during the intervention. The out-of-sample predictions for the rotating stations match the data in sign and relative magnitude, although their predicted absolute decrease is much smaller than the actual decrease. In contrast, at these parameter values the model predicts that more, rather than less, drivers should be using road #2 in the fixed police stations. The inability to match these moments is not a fundamental limitation of the model–column 4 shows that, if driver's priors on checking on roads B are sufficiently high (here set at .5 for illustrative purposes), there would be decreased drunkeness on both routes in fixed police stations. These beliefs would also better fit the magnitude of the decrease in rotating stations. However, as we see in column 4 of 10, negative priors on route B intensity would imply a large rapid decrease in the number of drunken drivers caught on road B during the intervention, whereas in reality the number of criminals apprehended on these routes changed only gradually due to the enforcement program.

While the model does not perfectly fit the out-of-sample data we do note that the final check results are both quite noisy, and collected using a different police checking methodology than in the main portion of the intervention. Furthermore, since they were almost all collected on one of two nights, they are subject to potential shocks that do not wash out over the course of the intervention, for instance of weather or traffic makes certain routes less desirable. Finally, it is possible that in some police station jurisdictions, roads A and B are not geographically independent as assumed. If so, this might cause a mechanical decrease in drunken driving on road B if some drivers must cross point B in order to get to A, and are dissuaded from driving on A due to the program.

7.3 Optimal Enforcement Strategy

With the estimated structural parameters in hand, we proceed to calculating the optimal crackdown strategy. To simplify the exercise, we set two constraints on the design of this strategy. The first to impose an ex-post budget constraint: we assume that the police department has enough budget to carry out 20 checkpoints per police station in every 90 days.[11] The second is to limit the police optimization problem to two parameters: the duration of the crackdown, and the fraction of checkpoints to be conducted off of the primary road (that is, on road B). Once these parameters are set, the specific timing of the 20 checkpoints is determined randomly with uniform probability on each nights, subject to the constraint that any station cannot carry out two checkpoint on the same night.

While this constrained optimal strategy would, by definition, be less effective to than the unconstrained optimal strategy, it has the advantages of being more transparent and easy to calculate, as well as adhering more closely to the theoretical problem analyzed in Section 4.2.[12] Due to the discreteness of the number of checkpoints, the strategy can be

[11] This is close to the average number of checkpoints per station conducted in the second round of the intervention (22.4 over 3 months) which we were informed by the police was the largest crackdown they would implement.

[12] In particular, the police might want to have crackdown intensities change in more complex patterns over

defined as the (integer) number of checks to be conducted on road B and the number of nights to conduct checks, and then solved by a simple grid search over the possible values.

The estimated parameters imply that the optimal strategy is to spread the 20 checkpoints over 76 days, and to divide them evenly between roads A and B (10 on each). These results are consistent with the implications of the simple theoretical model developed in Section 4: the police are most effective when they spread the crackdown diffusely over a longer period. Of the 90 day cycle considered, only two weeks would not be included in the crackdown. The equal division of checkpoints between roads is a natural implication of two results: first, that road A is only weakly preferred to road B, with 7.6% more drunken drivers on the first night of checking. Second, since learning is slower on road B it takes more intense police enforcement to alter behavior on that road.

8 Conclusion

This paper presents the results from a randomized experiment on the enforcement of drunk driving campaign in Rajasthan India, which was set up to test a model of strategic behavior which augments the classic model in the literature with learning. The key conclusion is that, although there is clear evidence of learning, the central insight of the literature that random checking of potential locations is a better use of scarce policing resources remain valid, even in this case. There is indeed learning, but it is extremely quick: we see some evidence of it even in the first two weeks of a crackdown, when the intensity of checking is high enough. This has implications for drunk driving, a public health issue which is gaining in salience and importance, and for other situations where crackdown solutions may be appropriate.

time.

References

[1] Banerjee, Abhijit, Raghabendra Chattopadhyay, Esther Duflo, Daniel Keniston, and Nina Singh. Can Institutions be Reformed from Within? Evidence from a Randomized Experiment with the Rajasthan Police. Working Paper 17912, National Bureau of Economic Research, 2012.

[2] Becker, Gary. "Crime and Punishment: An Economic Approach." Journal of Political Economy 76, no. 2 (1968): 169-217.

[3] Braga, Anthony A. "The Effects of Hot Spots Policing on Crime." Annals of the American Academy of Political and Social Science, 578 (2001): 104-125.

[4] Clarke, Ronald V. and David Weisburd. "Diffusion of crime control benefits: observations on the reverse of displacement." Crime Prevention Studies 2, (1994): 165-184.

[5] Davis, A., A. Quimby, W. Odero, G. Gururaj, and M. Hijar. Improving Road Safety by Reducing Impaired Driving in. London: DFID Global Road Safety Partnership, 2003.

[6] Dell, Melissa, "Trafficking Networks and the Mexican Drug War" American Economic Review, 105 (6), 2015, 1738-1779.

[7] Duflo, Esther, Abhijit Banerjee, and Rachel Glennerster. "Putting a band-aid on a corpse: Incentives for nurses in the Indian public health care system." Journal of the European Economic Association 6, no. 5 (2010): 487-500.

[8] Eeckhout, Jan, Nicola Persico, and Petra E. Todd. "A Theory of Optimal Random Crackdowns." American Economic Review 100, no. 3 (2010): 1104–1135.

[9] Elder, Randy W., Ruth A. Shults, David A. Sleet, James L. Nichols, Stephanie Zaza, and Robert S. Thompson. "Effectiveness of Sobriety Checkpoints for Reducing Alcohol-Involved Crashes." Traffic Injury Prevention 3 (2002): 266-274.

[10] Erke, Alena, Charles Goldenbeld, and Truls Vaa. "The effects of drink-driving checkpoints on crashes–A meta-analysis." Accident Analysis and Prevention 41 (2009): 914-923.

[11] Green, Lorraine and David Weisburd. "Policing drug hot spots: the Jersey City drug market analysis experiment." Justice Quarterly 12, no. 4 (1995): 711-735.

[12] Groff, Elizabeth R., Jerry H. Ratcliffe, Travis Taniguchi, and Jennifer D. Wood. "The Philadelphia foot patrol experiment: A randomized controlled trial of police patrol effectiveness in violent crime hotpots." Criminology 49, no. 3 (2011): 795-831.

[13] Habyarimana, James, and William Jack. "Heckle and Chide: Results of a Randomized Road Safety Intervention in Kenya." Working Paper Number 169. Center for Global Development, April 2009.

[14] Koper, Christopher S. "Just enough police presence: reducing crime and disorderly behavior by optimizing patrol time in crime hot spots." Justice Quarterly 12, no. 4 (1995): 649-672.

[15] Lazear, Edward P. Speeding, Terrorism, and Teaching to the Test. Working Paper 10932, National Bureau of Economic Research, 2004.

[16] Mookherjee, Dilip and I. P. L. Png. "Marginal Deterrence in Enforcement of Law" Journal of Political Economy, Vol. 102, No. 5 (1994): 1039-1066

[17] Peek-Asa, Corinne. "The Effect of Random Alcohol Screening in Reducing Motor Vehicle Crash Injuries." American Journal of Preventative Medicine 18, no. 1S (1999): 57-67.

[18] Ross, Laurence H. "Punishment as a factor in preventing alcohol-related accidents." Addiction 88, no. 7 (1993): 997-1002.

[19] Sherman, Lawrence W. "Police Crackdowns: Initial and Residual Deterrence." Crime and Justice 12, (1990): 1-48.

[20] Transparency International. India Corruption Study 2005: Corruption in the Indian Police. New Delhi: Center for Media Studies, 2005.

[21] WHO. World Report on Road Traffic Injury Prevention. Geneva: World Health Organization, 2004.

9 Tables

Table 1: Police station treatment assignment

		Implementation Staff	
A. Sep. 2010 Round: 16 control police stations			
		Police Lines Teams	Police Station Teams
Checkpoint Strategy	Rotating	5 stations @ 2/week	7 stations @ 2/week
	Fixed	6 stations @ 2/week	6 stations @ 2/week
B. Sep.-Nov. 2011 Round: 60 control police stations			
		Police Lines Teams	Police Station Teams
Checkpoint Strategy	Rotating	8 stations @ 1/week 11 stations @ 2/week 10 stations @ 3/week	10 stations @ 1/week 9 stations @ 2/week 12 stations @ 3/week
	Fixed	9 stations @ 1/week 7 stations @ 2/week 9 stations @ 3/week	14 stations @ 1/week 13 stations @ 2/week 11 stations @ 3/week

Table 2: Summary Statistics

	Obs.	Mean	SD	Median	Min.	Max.
A. Police station daily mean accidents and deaths (Control stations)						
Accidents	77357	0.12	0.36	0	0	4
Deaths	77357	0.05	0.27	0	0	13
Night Accidents	77357	0.04	0.19	0	0	3
Night Deaths	77357	0.02	0.16	0	0	13
B. Total vehicles passing police checkpoint locations in control stations						
Location 1	238	941.02	726.48	672.5	117	4862
Location 2	244	932.66	914.84	612	123	4998
Location 3	256	895.33	888.9	571	38	4743
C. Vehicles stopped by police at checkpoints						
Total	837	105.28	108.26	69	1	1180
Motorcycles	837	39.9	47.04	25	0	357
Cars	837	22.16	35.24	10	0	435
Trucks	837	19.52	35.25	9	0	580
D. Drunk drivers caught by police at checkpoints						
Total	837	1.85	2.36	1	0	21
Motorcycles	837	1.03	1.63	0	0	14
Cars	837	0.2	0.59	0	0	7
Trucks	837	0.23	0.61	0	0	5
E. Percentage found drunk in control police stations at final check						
Total	4988	2.23%	2.18%			
Motorcycles	2202	3.36%	3.25%			
Cars	1383	0.72%	0.72%			
Trucks	571	1.93%	1.89%			
F. Police checkpoint attendance						
Checkpoint occurred	1580	62.50%	23.45%			
Arrived on time	980	54.54%	24.79%			
Stayed until 10:00pm	980	72.23%	20.06%			

Omitted vehicle categories are vans, jeeps, buses, autorickshaws, and other (mostly tractors). The lower number of night deaths observations is due to the fact that this data is not available for January and February 2012. These months are omitted from the rest of the analysis.

Table 3: Pooled Results

	Daylight		Darkness		
	(1) Accidents	(2) Deaths	(3) Accidents	(4) Deaths	(5) Deaths
Treatment during & post intervention	0.00326	-0.00219	-0.00561**	-0.00393*	-0.00355
	(0.00372)	(0.00289)	(0.00237)	(0.00213)	(0.00313)
Month FE	Yes	Yes	Yes	Yes	Yes
Observations	5080	4714	5080	4714	5080
R^2	0.426	0.229	0.350	0.229	0.327
Mean of control	0.0849	0.0293	0.0329	0.0165	0.0454

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Fixed vs. Rotating Pooled Results

	Daylight		Darkness		
	(1) Accidents	(2) Deaths	(3) Accidents	(4) Deaths	(5) Deaths
Fixed checkpoints during & post intervention	-0.00199 (0.00448)	-0.00389 (0.00330)	-0.00153 (0.00288)	-0.00282 (0.00234)	-0.00451 (0.00346)
Rotating checkpoints during & post intervention	0.00838* (0.00447)	-0.000525 (0.00363)	-0.00959*** (0.00284)	-0.00502* (0.00288)	-0.00261 (0.00411)
Month FE	Yes	Yes	Yes	Yes	Yes
Observations	5080	4714	5080	4714	5080
R^2	0.427	0.229	0.351	0.229	0.327
Mean of control	0.0849	0.0293	0.0329	0.0165	0.0454
P-value of test fixed = Rotating effect	0.0398	0.378	0.0130	0.473	0.660

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5: Fixed vs. Rotating, During & Post Intervention

	Daylight		Darkness		
	(1) Accidents	(2) Deaths	(3) Accidents	(4) Deaths	(5) Deaths
Fixed stations during intervention	-0.00089	-0.00081	0.00234	-0.00199	-0.00199
	(0.00644)	(0.00411)	(0.00452)	(0.00274)	(0.00476)
Fixed stations post intervention	-0.00265	-0.00732*	-0.00438	-0.00383	-0.00632
	(0.00495)	(0.00416)	(0.00344)	(0.00301)	(0.00440)
Rotating stations during intervention	0.01111*	0.00286	-0.00987**	-0.00515	-0.00202
	(0.00648)	(0.00478)	(0.00410)	(0.00340)	(0.00556)
Rotating stations post intervention	0.0064	-0.00451	-0.00921**	-0.00478	-0.00294
	(0.00551)	(0.00415)	(0.00354)	(0.00394)	(0.00481)
Month FE	Yes	Yes	Yes	Yes	Yes
Observations	5080	4714	5080	4714	5080
R^2	0.426	0.230	0.350	0.229	0.327
Mean of control	0.0849	0.0293	0.0329	0.0165	0.0454

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Fixed vs. Rotating, Intensity of Checking

	Daylight		Darkness		
	(1) Accidents	(2) Deaths	(3) Accidents	(4) Deaths	(5) Deaths
Rotating checkpoints×1/week during & post intervention	0.00892 (0.00827)	0.0007 (0.00769)	-0.00739 (0.00534)	0.00432 (0.00538)	0.00248 (0.00821)
Rotating checkpoints×2/week during & post intervention	0.00583 (0.00637)	-0.00179 (0.00374)	-0.00941** (0.00380)	-0.00932** (0.00399)	-0.00737 (0.00502)
Rotating checkpoints×3/week during & post intervention	0.01061* (0.00619)	0.00016 (0.00628)	-0.01156*** (0.00413)	-0.00607* (0.00339)	-0.00062 (0.00641)
Fixed checkpoints ×1/week during & post intervention	-0.00694 (0.00631)	-0.0016 (0.00546)	-0.00247 (0.00387)	-0.00227 (0.00319)	-0.00295 (0.00490)
Fixed checkpoints×2/week during & post intervention	0.00195 (0.00656)	-0.00314 (0.00419)	-0.00136 (0.00433)	-0.00184 (0.00359)	-0.00164 (0.00526)
Fixed checkpoints×3/week during & post intervention	-0.00228 (0.00626)	-0.00726 (0.00479)	-0.0007 (0.00439)	-0.00342 (0.00322)	-0.00950* (0.00510)
Month FE	Yes	Yes	Yes	Yes	Yes
Observations	5080	4714	5080	4714	5080
R^2	0.426	0.230	0.350	0.229	0.327
Mean of control	0.0849	0.0293	0.0329	0.0165	0.0454

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7: Checkpoint surveys during intervention

	Drunk drivers and motorcyclists caught				Cars and motorcycles passing			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment					-145.143*	-56.589		
					(86.469)	(133.879)		
Rotating checkpoint station	0.123	-0.989*			-8.269	-34.895		
	(0.187)	(0.512)			(72.071)	(81.652)		
Frequency		-0.505**				-44.792		
		(0.195)				(43.426)		
Rotating checkpoint × frequency		0.576**				21.751		
		(0.232)				(17.482)		
Weeks of checking			-0.092***				-17.365***	
			(0.026)				(3.878)	
Rotating checkpoint × weeks of checking			0.085**				-5.962	
			(0.035)				(8.909)	
Number Previous checkpoints				-0.039***				-8.201***
				(0.012)				(1.977)
Rotating checkpoint × number previous checkpoints				0.031*				-1.315
				(0.016)				(4.053)
Observations	866	866	866	866	2843	2843	2843	2843
Mean of dep. variable	1.237	1.237	1.237	1.237	528.6	528.6	528.6	528.6
District FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Police Station FE	No	No	Yes	Yes	No	No	Yes	Yes

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ All specifications include controls for whether the police station is located on a major highway, the pre-intervention accident rate, and assignment to the police lines intervention.

Table 8: Passing vehicles post-crackdown

	Road 1		Roads 2 & 3	
	(1)	(2)	(3)	(4)
Treatment	-214.1**		-13.89	
	(71.45)		(43.40)	
Days since last checkpoint	1.761		-3.776	
	(4.003)		(3.217)	
Fixed checkpoints		-224.4**		191.4
		(79.33)		(128.0)
Fixed checkpoint		0.155		-11.77**
×days since last checkpoint		(3.556)		(5.261)
Rotating checkpoints		-289.0**		-107.0**
		(115.6)		(40.37)
Rotating checkpoint		17.13		1.593
×days since last checkpoint		(15.91)		(4.637)
Observations	145	145	347	347
R-squared	0.170	0.192	0.107	0.116
District FE	Yes	Yes	Yes	Yes
Control mean	488.0	488.0	501.5	501.5

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ All specifications include controls for whether the police station is located on a major highway, the pre-intervention accident rate, and assignment to the police lines intervention.

Table 9: Drunk drivers caught on final check

	All Stations		Rotating checkpoints		Fixed Checkpoints	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	-1.274**	-0.223	-1.195	-0.440	-1.254*	0.711
	(0.452)	(0.615)	(0.807)	(1.250)	(0.560)	(1.366)
Days since last checkpoint	0.017*	-0.013	0.001	-0.057	0.033***	-0.011
	(0.008)	(0.025)	(0.017)	(0.032)	(0.008)	(0.037)
Frequency		-0.510*		-0.317		-1.012*
		(0.268)		(0.368)		(0.511)
Days since last last checkpoint		0.015		0.026*		0.023
×frequency		(0.012)		(0.012)		(0.019)
Observations	109	109	78	78	75	75
R-squared	0.179	0.188	0.182	0.194	0.128	0.152
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Control mean	1.909	1.909	1.909	1.909	1.909	1.909
Mean treatment effect, @ freq = 2	-1.045	-1.020	-1.182	-1.125	-0.819	-0.847
P-value of mean treatment effect	0.0153	0.0141	0.0747	0.0951	0.126	0.100

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ All specifications include controls for whether the police station is located on a major highway, the pre-intervention accident rate, and assignment to the police lines intervention.

Table 10: Main Moments - Drunk drivers caught during intervention

	Empirical		Simulated	Simulated $\lambda_{B0} = .5$	Empirical		Simulated	Simulated $\lambda_{B0} = .5$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Fixed Stations - Road A				**Rotating Stations - Road A**			
	Mid Intervention							
1 check/week	-0.131	-0.235	-0.1968	0.0301	-0.027	-0.226	0.0086	-0.0683
	(0.212)	(0.193)			(0.319)	(0.291)		
2 checks/week	-0.171	-0.139	-0.4153	-0.2731	0.297	0.139	-0.1399	0.0644
	(0.228)	(0.216)			(0.238)	(0.294)		
3 checks/week	-0.884	-0.734	-0.5987	-0.5122	-0.452	-0.491	-0.3450	0.0901
	(0.281)	(0.264)			(0.329)	(0.363)		
	Late Intervention							
1 check/week	-0.546	-0.533	-0.4068	-0.0928	0.005	-0.349	-0.0456	-0.0048
	(0.194)	(0.208)			(0.358)	(0.314)		
2 checks/week	-0.298	-0.360	-0.7803	-0.6436	-0.255	-0.279	-0.2816	0.0615
	(0.328)	(0.303)			(0.229)	(0.198)		
3 checks/week	-1.578	-1.332	-0.9250	-0.9154	-0.510	-0.397	-0.4659	-0.2125
	(0.295)	(0.290)			(0.324)	(0.375)		
	Rotating Stations - Road B							
	Mid Intervention							
1 check/week	0.376	0.232	-0.0243	-0.3141				
	(0.373)	(0.323)						
2 checks/week	0.293	0.111	-0.0198	-0.7664				
	(0.223)	(0.214)						
3 checks/week	0.216	0.233	-0.0621	-1.2056				
	(0.325)	(0.289)						
	Late Intervention							
1 check/week	0.257	0.024	-0.0509	-0.8292				
	(0.374)	(0.360)						
2 checks/week	-0.047	-0.248	-0.1081	-1.5945				
	(0.290)	(0.304)						
3 checks/week	-0.494	-0.418	-0.2081	-1.6913				
	(0.298)	(0.272)						
Road B	-0.105	-0.006	-0.2656	-1.8697				
	(0.250)	(0.215)						
Observations	761	761			761	761		
Controls for police stops, passing vehicles, day of week	No	Yes	No	No	No	Yes	No	No

Table 12: Final Check Moments - Drunk drivers caught

	Empirical		Simulated	Simulated - $\lambda_{B0} = .5$
	(1)	(2)	(3)	(4)
Fixed Stations - Road B				
1 check/week	-0.117	-1.583	0.3094	-1.5525
	(0.644)	(0.867)		
2 checks/week	-1.614	-2.995	0.4426	-1.9750
	(0.833)	(0.967)		
3 checks/week	-6.735	-8.742	0.4854	-2.2036
	(3.959)	(3.267)		
1 check/week ×days since last checkpoint	-0.011	-0.001	-0.0076	0.0416
	(0.024)	(0.029)		
2 checks/week ×days since last checkpoint	0.026	0.036	-0.0112	0.0571
	(0.032)	(0.036)		
3 checks/week ×days since last checkpoint	0.194	0.223	-0.0124	0.0672
	(0.122)	(0.102)		
Rotating Stations - Road B				
1 check/week	-0.367	-1.116	-0.1262	-1.6044
	(0.792)	(0.885)		
2 checks/week	-0.868	-1.291	-0.2175	-2.1512
	(0.974)	(1.068)		
3 checks/week	-1.378	-2.645	-0.3183	-1.9173
	(0.672)	(0.800)		
1 check/week ×days since last checkpoint	-0.079	-0.076	0.0036	0.0440
	(0.070)	(0.079)		
2 checks/week ×days since last checkpoint	-0.008	-0.018	0.0067	0.0631
	(0.036)	(0.037)		
3 checks/week ×days since last checkpoint	0.024	0.045	0.0091	0.0510
	(0.026)	(0.030)		
Observations	104	104	104	104
Controls for police stops, passing vehicles, day of week	No	Yes	No	No

Table 13: Structural Parameters

Interpretation	Parameters	Estimated value: $\delta = 0.90$	Estimated value: $\delta = 0.00$
Utility of drunken driving on road A	d_A	0.7597	
Utility of drunken driving on road B	d_B	0.6879	
Utility cost of encountering checkpoint	c	30.4592	
Prior on road A intensity	$\lambda_{A0} = \frac{\tilde{\alpha}^{\lambda}_{A0}}{\tilde{\alpha}^{\lambda}_{A0}+\tilde{\beta}^{\lambda}_{A0}}$	0.0024	
Strength of prior on road A intensity	$\tilde{\alpha}^{\lambda}_{A0} + \tilde{\beta}^{\lambda}_{A0}$	22.4680	
Prior on road B intensity	$\lambda_{B0} = \frac{\tilde{\alpha}^{\lambda}_{B0}}{\tilde{\alpha}^{\lambda}_{B0}+\tilde{\beta}^{\lambda}_{B0}}$	0.00026	
Strength of prior on road B intensity	$\tilde{\alpha}^{\lambda}_{B0} + \tilde{\beta}^{\lambda}_{B0}$	93.2608	
Prior on crackdown duration	$\eta_0 = \frac{\tilde{\alpha}^{\eta}_{0}}{\tilde{\alpha}^{\eta}_{0}+\tilde{\beta}^{\eta}_{0}}$	0.9144	
Strength of prior on road crackdown duration	$\tilde{\alpha}^{\eta}_{0} + \tilde{\beta}^{\eta}_{0}$	4.2987	

10 Figures

Figure 10.1: Simulated Histories

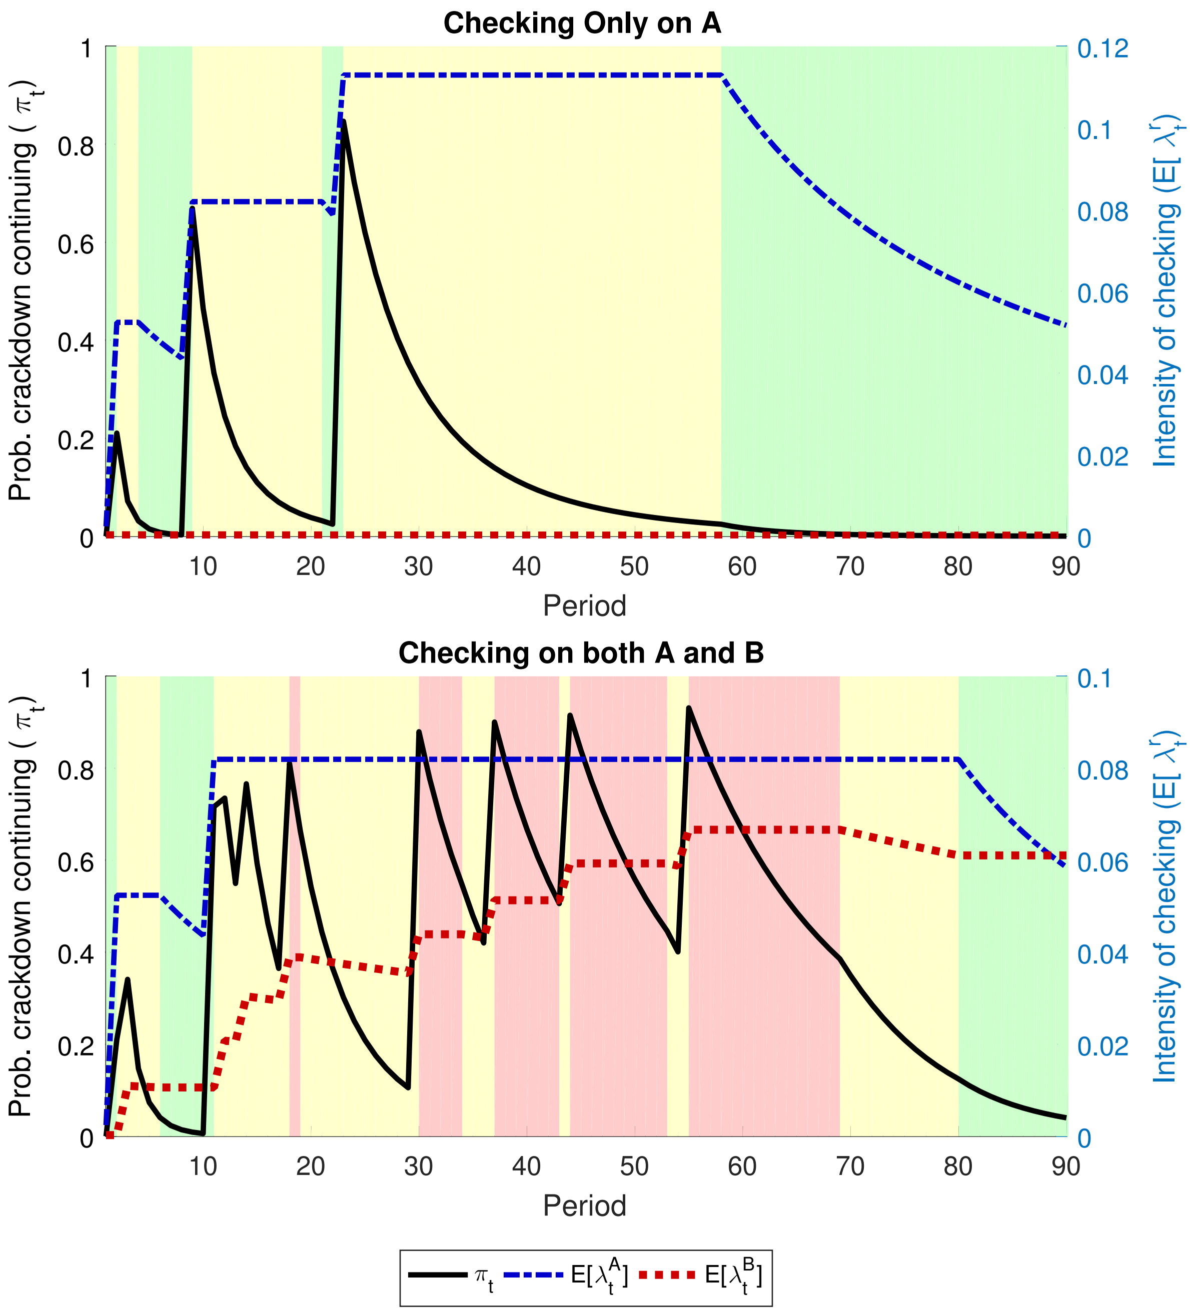

Figure 10.2: Drunken drivers caught over time

(a) Whole intervention learning results

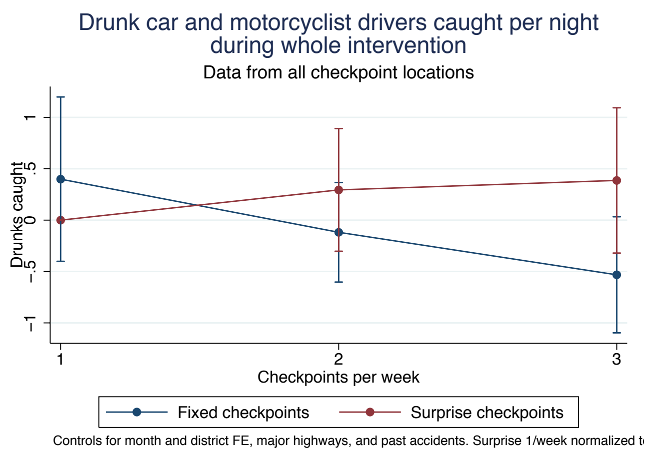

(b) Pre-learning results

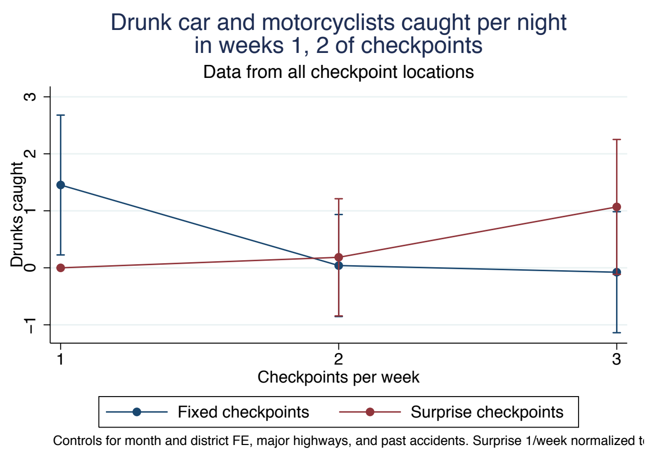

(c) Post-learning results

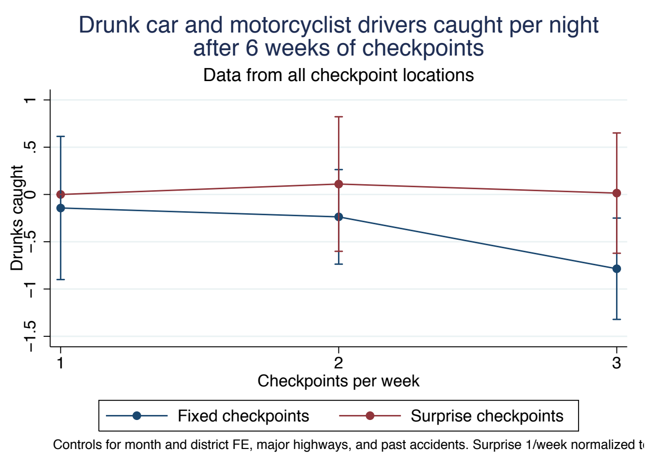

11 Appendix

11.1

Here we establish the conditions necessary for drivers' strategies to be the following: Upon learning of the change in enforcement they cease drinking and driving for a finite interval τ. After τ has elapsed, they return to drinking and driving. If they never subsequently encounter a checkpoint, they continue to drink and drive. If they do encounter a checkpoint, they give up drinking and driving permanently.

Define $\bar{p}$ as the belief about the permanence of checking such that if a driver is checked and his belief is that checking is permanent with probability $\bar{p}$, then he will be indifferent between a policy of remaining on the safe road forever, or remaining sober for a finite period τ before returning to the risky road. $\bar{p}$ is implicitely defined by:

$$\frac{s}{r} = \left(\frac{1-e^{-r\bar{\tau}}}{r}\right)s + \frac{e^{-r\bar{\tau}}}{r}\left(\bar{p}\left[W_{perm}\right] + (1-\bar{p})\left[\left(1-e^{-\eta\bar{\tau}}\right)d + e^{-\eta\bar{\tau}}W_{temp}\right]\right)$$

$$\bar{\tau} = \frac{1}{\eta}\ln\left(\frac{(1-\bar{p})(\eta+r)(d-W_{temp})}{r\left((1-\bar{p})d+\bar{p}W_{perm}-s\right)}\right)$$

Since at any time t the driver's belief that checking is permanent, p_t, is defined as

$$p_t = \frac{p_0}{p_0 + (1-p_0)e^{-\eta t}}$$

as time passes p_t reaches the point such that even if any temporary crackdown was certain to have ended, staying home is still preferable to drunk driving. At this point the denominator in the expression for $\bar{\tau}$ becomes negative, and the driver will cease driving permanently if stopped. Thus $\bar{p}$ is defined,

$$\begin{aligned} d - \bar{p}d + \bar{p}W_{perm} &= s \\ \bar{p} &= \frac{d-s}{d-W_{perm}} \end{aligned}$$

and the time that $\bar{p}$ is reached is implicitly defined as

$$e^{-\eta t} = \frac{p_0}{1-p_0}\left(\frac{s-W_{perm}}{d-s}\right)$$

In order for the proposed strategy to be optimal, it must be the case that $\tau \geq \bar{t}$: after the driver's initial period at home, if he returns and is checked permantely not drinking and driving will be preferable to another temporary interval of sobriety. Thus

$$\begin{aligned}
\tau &\geq \bar{t} \\
\frac{p_0}{1-p_0}\left(\frac{s-W_{perm}}{d-s}\right) &\geq \frac{r\left(\left(1-p_0\right)d+p_0W_{perm}-s\right)}{\left(1-p_0\right)\left(\eta+r\right)\left(d-W_{temp}\right)} \\
p_0\left(\frac{s-W_{perm}}{d-s}\right) &\geq \frac{r\left(d-p_0d+p_0W_{perm}-s\right)}{\left(\eta+r\right)\left(d-W_{temp}\right)} \\
p_0\left(\frac{s-W_{perm}}{d-s}\right) &\geq \frac{r\left(d-s-p_0\left(d-W_{perm}\right)\right)}{\left(\eta+r\right)\left(d-W_{temp}\right)}
\end{aligned}$$

which is bound to hold for p_0 high enough.

Details on updating

Updating on checking intensity:

If beliefs about checking intensity are modeled as being distributed Beta(α, β), then we require two state variables for the conditional expectation of the intensity under each potential scope of checking.

$$\begin{aligned}
\text{Stopped on } A\text{: } \lambda_t^A &= \frac{\alpha_t^{\lambda,A}+1}{\alpha_t^{\lambda,A}+\beta_t^{\lambda,A}+1} \\
\text{Not stopped on } A\ \lambda_t^A &= \frac{\alpha_t^{\lambda,A}}{\alpha_t^{\lambda,A}+\beta_t^{\lambda,A}+1} \\
\text{Stopped on } B\text{: } \lambda_t^A &= \frac{\alpha_t^{\lambda,A}}{\alpha_t^{\lambda,A}+\beta_t^{\lambda,A}} \\
\text{Not stopped on } B\text{: } \lambda_t^A &= \frac{\alpha_t^{\lambda,A}}{\alpha_t^{\lambda,A}+\beta_t^{\lambda,A}} \\
\text{Stay home: } \lambda_t^A &= \frac{\alpha_t^{\lambda,A}}{\alpha_t^{\lambda,A}+\beta_t^{\lambda,A}}
\end{aligned}$$

with the transitions for λ_t^B defined analogously.

Updating on checking permanence:

Let $g_t(\rho)$ be the distribution of the driver's beliefs on this probability in period t. These beliefs evolve in a simple way:

$$g_{t+1}(\pi|X) = \frac{\alpha_t^\pi}{\alpha_t^\pi+\beta_t^\pi+1}$$

Specification of prior beliefs:

For computational convenience, we define the state variables as the mean of the intensity distributions and the number of trips on each road:

$$\begin{aligned} \mu_t &= \frac{\alpha}{\alpha+\beta} \\ Trips_t &= \alpha+\beta \end{aligned}$$

Based upon these, we can recover the

$$\begin{aligned} \beta &= \frac{\alpha}{\mu} - \alpha \\ \beta &= \frac{Trips-\beta}{\mu} - Trips + \beta \\ 0 &= Trips - \beta - \mu Trips \\ \beta &= (1-\mu)\, Trips \\ \alpha &= Trips - \beta \\ \alpha &= Trips - (1-\mu)\, Trips \\ \alpha &= \mu Trips \end{aligned}$$

Based upon these state variables, the beta distribution can be rewritten as:

$$\text{Beta}\,(\alpha, \beta) = \text{Beta}\,(\mu Trips, (1-\mu)\, Trips)$$

We specify the parameters to be estimated as the initial prior on the intensity, μ_0 and the initial number of "trips" that the driver has taken, $Trips_0$. The prior number of trips can be interpreted as a shifter of the precision of the driver's initial beliefs about checking. Thus after τ trips made after the beginning of the crackdown, the driver's total trips are $Trips_t = Trips_\tau + Trips_0$.

For the π parameter, the driver's expectation that the checking ends in period t can be written as:

$$\begin{aligned} \frac{\alpha_t^\pi}{\alpha_t^\pi + \beta_t^\pi + t} &= \frac{\alpha_t^\pi}{Trips_0^\pi + t} \\ &= \frac{\alpha_t^\pi}{Trips_0^\pi} \frac{Trips_0^\pi}{Trips_0^\pi + t} \\ &= \frac{\mu_0^\pi Trips_0^\pi}{Trips_0^\pi + t} \end{aligned}$$

Table A1: Checkpoint surveys during intervention

	Drunk drivers and motorcyclists caught			
	(1)	(2)	(3)	(4)
Rotating checkpoint station	0.096 (0.136)	-0.813* (0.432)		
Frequency		-0.362*** (0.135)		
Rotating checkpoint × frequency		0.458** (0.188)		
Weeks of checking			-0.044** (0.019)	
Rotating checkpoint × weeks of checking			0.036 (0.025)	
Number Previous checkpoints				-0.016** (0.007)
Rotating checkpoint × number previous checkpoints				0.010 (0.010)
Observations	1352	1352	1352	1352
Mean of dep. variable	0.806	0.806	0.806	0.806
District FE	Yes	Yes	Yes	Yes
Police Station FE	No	No	Yes	Yes

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ All specifications include controls for whether the police station is located on a major highway, the pre-intervention accident rate, and assignment to the police lines intervention.

11.2 Tables

Table A2: Final check differential compliance

	(1) Checkpoint occurred	(2) Checkpoint occurred	(3) Checkpoint occurred
Fixed checkpoints	-0.0165		
	(0.104)		
Rotating checkpoints	0.0347		0.0543
	(0.0555)		(0.108)
Intensity 1/week		0.0324	0.00558
		(0.0713)	(0.106)
Intensity 2/week		0.0259	-0.00355
		(0.0624)	(0.103)
Intensity 3/week		-0.0202	-0.0514
		(0.0930)	(0.134)
Police lines teams	0.336**	0.342**	0.337**
	(0.119)	(0.122)	(0.120)
Observations	182	182	182
R-squared	0.166	0.166	0.168
District FE	Yes	Yes	Yes

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ All specifications include controls for whether the police station is located on a major highway and the pre-intervention accident rate.