

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

Abhijit Banerjee¹, Esther Duflo¹, Daniel Keniston², and Nina Singh³

¹Massachusetts Institute of Technology

²Louisiana State University

³Rajasthan Police

August 18, 2019

Abstract

Should police activity should be narrowly focused and high force, or widely-dispersed but of moderate intensity? Critics of intense “hot spot” policing argue it primarily displaces, not reduces, crime. But if learning about enforcement takes time, the police may take advantage of this period to intervene intensively in the most productive location. We propose a multi-armed bandit model of criminal learning and structurally estimate its parameters using data from a randomized controlled experiment on an anti-drunken driving campaign in Rajasthan, India. In each police station, sobriety checkpoints were either rotated among 3 locations or fixed in the best location, and the intensity of the crackdown was cross-randomized. Rotating checkpoints reduced night accidents by 17%, and night deaths by 25%, while fixed checkpoints had no significant effects. In structural estimation, we show clear evidence of driver learning and strategic responses. We use these parameters to simulate environment-specific optimal enforcement policies.

Key Words: Learning Models, Choice Modeling, Information Acquisition, Illegal Behavior, Law Enforcement, Crime Prevention

*We thank John Firth and Pankaj Verma for their outstanding contributions as research assistants on this project, and the Rajasthan Research Institute for coordinating data collection. We are grateful to the JPAL Governance Initiative for financial support. Thanks to Steven Berry, Mira Frick, Debray Ray, Nicholas Ryan and Christopher Udry for helpful suggestions, and seminar participants at Auburn Univ., Boston Univ., the Univ. of Chicago, Oxford Univ., Stanford Univ., Harvard Univ., Univ. of Pennsylvania, Northwestern Univ., LSE, LSU, and Yale Univ. for comments. The findings and opinions expressed in this paper are those of the authors and do not in any way engage the responsibility of the Rajasthan Police.

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 potential increasing returns from higher arrest probabilities on criminal activity (Eeckhout et al., 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 are 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 most criminals 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 frequent their favored location for a while, until they are persuaded that the checking will continue long enough that it is worth changing behavior. 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 monitoring hitherto unpoliced locations for a fixed period of time, it is not clear whether they 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 are 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, (Habyarimana and Jack, 2011) road safety 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 promising 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 were 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, the number of deaths from accidents at night (when the checking took place) decreased by about 25% in all treatment stations relative to the control stations. There was essentially no change in the number of deaths during the day time. The number of accidents at night also went down by about 17%, though there was some marginally significant increase in non-fatal accidents in day time. This salutary effect comes mainly from the rotating check locations. Those effects are much larger, only those are significant, and in the case of accidents the difference with fixed checkpoints 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, 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 drunken drivers apprehended per night goes down with the number and intensity of past check, while these effects are absent with rotating checks. This is consistent with the view that drunk drivers are still savvy enough to switch to other routes, which is what gives random checks their advantage over fixed checks.

This pattern of choices suggests that prospective drunken drivers are learning about police behavior, but allow for the possibility that the behavior changes over time, much as in a Restless Bandit model of learning. Some corroborating evidence that learning matters comes from looking at the post-intervention period. First, the effect on accidents of the rotating intervention continues over 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, consistent with the idea that the end of checking is easier to detect with a fixed intervention. Overall, the reduced form evidence is consistent that learning is sufficiently fast that a campaign that rotates across locations dominates one that focuses on the most crime-prone areas.

The data from the experiment are 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, despite many informal discussions implying the same underlying logic.

The structural model specifies agents' preferences for drunken driving as a function of their static utility of drunken driving in various locations, and their beliefs about the intensity and duration of police activity. Agents' strategies involve maximizing their static utility, and gaining information about police activity that will decrease the future danger of getting caught; we solve for these optimal policies using value function iteration. Using these policies, we can simulate the behavior of a large

number of agents (in each police station) as they encounter the set of checkpoints implemented during the anti-drunken driving campaign and adjust their beliefs and behaviors accordingly. We then match these simulated behaviors with data on road accidents and drunken drivers apprehended. Since the actual implementation of the checkpoints was imperfect and potentially endogenous, we construct moments using the randomly assigned checkpoint schedule as an instrument and find parameters of drivers' preferences and beliefs that minimize the resulting GMM criterion function.

Overall, our structural results confirm the qualitative findings implied by the reduced form results: drivers' estimated priors about the intensity of checking in the most preferred location are nearly uninformative, implying that they rapidly update their beliefs about the intensity of police enforcement in that location after just a single encounter. While drivers initially believe that, on average, the crackdown will last about 3 days, their high disutility of getting stopped by police ensures that a strong deterrence effect persists long after the checkpoints have ended and agents' beliefs about the probability of apprehension have decreased substantially. The parameters suggest that displacement of drunken driving is heterogeneous: increased checking in the preferred location decreases drunken driving in the 2nd location, but displaces criminal activity towards the less popular 3rd location.

With these structural parameters in hand, we can evaluate the effectiveness of potential counterfactual enforcement policies. We find that, given the beliefs held by drivers in Rajasthan at the time of the beginning of our intervention, the optimal crackdown would have been longer than the sobriety campaign actually implemented (82 days) and would have put 35% of the checkpoints on the most preferred road, similar to the actual rotating intervention. This long duration is necessary to convince drivers that the police are in fact checking on all roads—even those in which the population initially has very low expectations of finding a checkpoint. However, this is because we start from the premise that we have a one-off intervention. If such interventions happened regularly, in the long run equilibrium the population would probably have close to correct beliefs about the intensity and duration of the police checkpoint strategy. In this case, the optimal campaign duration is much shorter—21 days, and highly focused with 75% of checkpoints on the preferred road. In this scenario, the police's main goal is simply to demonstrate that the crackdown has begun through a "show of force" on the most commonly traveled road. Finally, we consider how pre-announcing the beginning of the campaign would affect the

design and impact of the campaign. An initial announcement has a moderate effect when drivers' beliefs are as measured in the actual intervention, but if drivers have equilibrium beliefs the announcement of the campaign is fully effective in ending drunken driving for the duration of the campaign and afterwards.

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 A1 presents a tractable model illustrating the factors affecting driver behavior and police effectiveness. Sections 4 and 5 present the data and reduced form results of the experiment. The structural estimation technique and results are presented in Section 6, along with the implications for the optimal policy. Finally, Section 7 concludes.

2 Background: Enforcing Drunk Driving

Each year 1.35 million people die in traffic accidents worldwide, with as many as 50 million injured. A staggering 90% of these deaths happen in developing countries (World Health Organization, 2018). Moreover, death and accident rates are rapidly increasing in developing countries even though these rates are falling in the developed world (Davis et al., 2003; Peden et al., 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% of deaths worldwide, twice the projected share for malaria (Habyarimana and Jack, 2011). The role of alcohol in road accidents is 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 a factor 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 (World Health Organization, 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 et al., 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, because most of the existing

literature has struggled with a variety of challenges and limitations. First, to our knowledge no previous research has used a randomized trial, raising 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 sobriety checkpoints in a low income context, or versus a counterfactual of essentially zero enforcement.

The law and economics and crime prevention literatures, recently reviewed by Chalfin and McCrary (2017), 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: switching from broad but low-level policing to more focused and intense patrols can be effective in reducing crime (Weisburd and Telep, 2014). Perhaps the strongest experimental evidence comes from the Sherman and Weisburd (1995) evaluation of 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. A natural concern (Sherman, 1990) lies in the possibility of crime being displaced 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, (Ratcliffe et al., 2011) others find that crime in control areas near hotspots actually decreased relative to more distant control areas (Weisburd and Green, 1995; Dell, 2015). The only such study conducted with police outside a high income country, Blattman et al. (2017) finds suggestive evidence of displacement of property crime away from streets with randomly increased police patrols in Bogota, but little displacement of violent crime. More evidence on displacement comes from recent work in Chile by Andres and Mobarak (2019), who show that vendors learn to disguise illegal fish sales from regulators and argue that more intense enforcement facilitates this learning. This heterogeneity in results is perhaps not surprising, since displacement effects depend heavily the location-specificity of the crime and the police enforcement approach, in particular the length of the crackdown.

One of the factors potentially driving the success of hot-spot policing is the

“experiential” effect: that (potential) criminals learn from their encounters with the police. Several recent studies (Anwar and Loughran, 2011; Wilson et al., 2017) model this learning in a Bayesian framework, showing evidence that encounters with the police cause individuals to update their beliefs on the probability of detection. While this literature studies learning, it does not model the “exploration” incentive—that individuals are interested in “seeing what they can get away with”. Moreover robust evidence on the speed of criminal learning about specific police initiatives remains 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. However, other studies find that program effects are short-lived and disappear quickly after the intervention (Weisburd and Telep, 2014). 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.

Finally, this study relates to a large theoretical and empirical literature on estimating models of learning (reviewed by Ching et al. (2017) and Foster and Rosenzweig (2010) in the context of developing countries). This paper differs from much of the existing criminology and development economics literature in that we estimate a structural learning model and apply it to the determination of optimal policy (for an exception in a very different context, see Bai, 2018). In contrast to the marketing and industrial organization literature, we use variation induced by a randomized experiment, rather than cross-market conditions. The structure of the learning model differs from the previous literature as well: instead of learning about a fixed parameter, such as the quality of a product, agents in our model learn about a police crackdown in which the intensity of enforcement varies over time in a correlated manner across different checkpoint locations.

2.1 Drunken Driving Enforcement in India

In India, highway safety laws of all kinds are generally enforced by fixed checkpoints manned by personnel from the local police station. Police in our intervention used the “selective breath checkpoint” methodology: barriers were arranged on the roadway, and officers signaled selected vehicles to pull over as they passed through the barriers.

Police personnel then questioned the driver on their identity, destination, etc., while observing the driver’s demeanor and smelling their breath. If the police suspected the driver to be drunk, then according to the official procedure they would order them to blow into a breathalyzer, following the results of which the driver was either charged or released.¹ The printed results of a handheld breathalyzer are considered sufficient proof of drunkenness in court. Once caught, drunken drivers’ vehicles were confiscated by the police, and the driver was required to 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 parking lot.

Even this official procedure leaves many factors to the discretion of the police, in particular the choice of how many, and which, vehicles to pull over for questioning and potential testing. 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,² 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 rare prior to the intervention and in the control police stations during the intervention. Outside the intervention, breathalyzers had not been widely distributed and without a breathalyzer it would have been extremely difficult for the police to prosecute a suspected drunken driver. As a result, with the exception of some police stations in Jaipur, the state capital, the police essentially did not implement sobriety check points. We observe that during the experiment control police stations hardly ever used them to carry out checkpoints.

¹This protocol is commonly used in US sobriety checkpoints as well.(Elder et al., 2002)

²We examine this issue in Section 5 using data from courts and find no relationship between checkpoint strategy and legal prosecution of drunken drivers.

In the 925 nights that surveyors visited control police stations, on only 7 (0.76%) occasions did they witness the police carrying out a checkpoint.

Conversations with the police and our personal observations at the checkpoints suggest that all suspects knew that driving under the influence was illegal, though many seemed quite surprised to encounter a checkpoint with a breathalyzer. Thus when interpreting the quantitative magnitudes of the reduced form parameters, we should bear in mind the context of no previous enforcement. Specific results might differ if the experiment had been done in an area with a long history of sobriety checkpoints, although we argue that the qualitative evidence of learning and displacement reveal more general facts about law enforcement. The structural estimation allows us address this question more directly by estimating optimal enforcement policies both with drivers' beliefs as estimated from the context of the intervention, and also with beliefs that are in expectation correct, as one would imagine in areas where citizens are very familiar with police strategies.

3 The Intervention

The anti-drunken driving program was implemented as a large-scale 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 according to the police, 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. In all interventions, checkpoints were held from 7:00pm to 10:00pm which was when peak drunk driving took place according to the police.

The program took place in two phases, an initial pilot, from September-early October 2010, and a larger roll-out from September to the end of November 2011. The initial pilot covered 2 districts and 40 police stations, and the second covered 11 districts and 183 police stations serving a territory of 125,000 km² with a population

of approximately 30 million. 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 in 2010 all checkpoints occurred twice a week and the program lasted at most 1.5 months. 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 and 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 theory of learning by potential and actual drunken drivers, we randomly assigned police stations to hold checkpoints at either a single location, or a rotating set of three locations. Before randomization, the police chief selected and rank-ordered the three best spots to catch drunken driver in their area. In the fixed location group, the checkpoints were always carried out at the best location, at the 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. Each police station's rotation was predetermined in advance by the research team.

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 consider the entire area under the police's jurisdiction and we do not control for checkpoint location fixed effects or characteristics, since these are themselves an outcome of the program choice. In the structural estimation of the learning model, however, we allow the public's preferences and beliefs about checkpoint probability to differ across the 3 locations in order 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.

The date of the last checkpoint, and hence the duration of the crackdown, was also randomized at the police station level. In the shorter, September 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 separately 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 over a month, whereas others had performed checks only a few days earlier³.

3.3 Checkpoint personnel

Previous work with the Indian government (Banerjee et al. (2008)) 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

³At the final check, the duration of time since the last checkpoint ranged from 1 to 44 days in different police stations.

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. Generally we found that the police lines teams were more likely to carry out intended check point. A complete analysis of this portion of the intervention is not within the scope of this paper, but we control for it when relevant in the analysis of the other branches of the intervention.

4 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.

4.1 Administrative Data

This study’s main results on accident and death rates are drawn from administrative accident reports by the police. For each accident on which data has been collected properly we know the police station, date and time of the incident, the number of individuals killed or injured, and the types of vehicles involved. Unfortunately we do not know whether drunken driving contributed to the accident, nor can we reliably link accidents with the sobriety checkpoints.

We obtained monthly accident data from August 2010 through October 2012 and daily data from August 2010-December 2011. For January and February 2012 the data is not disaggregated by day and night—which is unfortunate since the intervention was always in the evening when most drunken driving presumably occurs. 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.

4.2 Survey Data

We supplemented the police administrative data with additional data on police activity at the checkpoints collected by surveyors sent to monitor randomly selected checkpoint locations. Surveyors visited both on nights when the police were conducting anti-drunken driving checking and on nights when they were not, as well as at locations that the control police stations had 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.

After 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, and they were always held at the location designated as the second best location to catch drunken drivers. Hence in the fixed checkpoint intervention stations (and the control stations) there had previously been no enforcement at these specific locations, whereas in the rotating stations there had previously been checking at these exact spots. 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 the fraction of drunken drivers.

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 average fraction of drunken drivers in the underlying population come from the final checking night when checks were conducted in a more systematic fashion in control 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.

5 Reduced Form Predictions & Results

Driver actions, and hence accident and arrest rates, are a complex function of the utility of drunken driving, the cost of encountering the police, and beliefs about the intensity of checking, its status, and duration. However, much of the intuition behind our results can be shown in a simplified model in which drivers choose only whether to drink and drive on one road, learn only about whether the crackdown is permanent or temporary, and follow very simple strategies. Suppose that all drivers are initially unaware of the crackdown, and upon encountering the police for the first time choose to cease drinking and driving for temporary period then resume. If they encounter police a second time, they cease permanently. The police crackdown then affects drunken driving via three margins. First, it alerts drivers that enforcement is ongoing and triggers their temporary sobriety period. Second, it affects the length of the sobriety period. Third, it affects the share that are re-encountered and thus convinced to permanently stop offending. The police choice of intensity versus duration of the crackdown, and their overall budget, affects each of these margins differently. For brevity, we relegate the formal details of the model to Appendix Section A1.8, and present its main results to guide interpretation of the reduced form findings.

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, as long as drivers are strategic and learning is fast enough.
 - (a) If drivers initially believe that police would (almost) never put checkpoints on alternate routes, the intensity of a fixed one-road crackdown will not matter for overall drunken driving, since no drivers are pushed into sobriety. Intensity does matter for multi-road checking, since higher intensity will clear the roads faster and (potentially) induce some drivers to permanently cease inebriated driving.
3. Reversion and intensity: In equilibrium, drivers understand that the police face a binding budget constraint in which more intense crackdowns are necessarily shorter. So when potential criminals observe an intense crackdown they cease offending for a shorter interval and thus a revert more rapidly to crime 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.

We test these predictions using reduced form results 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 entire area covered by the police station (including, but not restricted to, the 3 main roads on which check points were to happen in rotating stations) during daylight, nighttime, or both. All regressions are at the police station-month level, except in cases where the intervention began or ended mid-month in which case we split the month into treated and untreated observations and weight according to the number of days in each. 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 (25%). There is no significant impact during the day or in the merged day and night data (column 5), 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. These coefficients imply that over the 17,616 station-nights of the crackdown and post-intervention period, approximately 70 deaths were averted because of the campaign.

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 (pooling observations with different numbers of check points 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 table makes clear that the impacts of the intervention on accidents and death are largely driven by the rotating check points police stations. Accidents at night went down by 29% and night deaths dropped 30% in these police stations, both significant at the 10% or 5% level. In contrast, there are no significant impacts in the fixed police stations and the point estimates are much smaller (2 to 8 times). The difference between fixed and rotating stations is significant for night time accidents, though not for deaths. Note that there is an 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 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 decline in nighttime accidents in rotating stations is almost 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. Overall, the evidence suggests that the basic insight of the crime and punishment literature is very powerful: checking at rotating locations does lead to a

greater reduction in accidents and deaths (the ultimate outcome the police wants to affect).

Table 6 further decomposes the program effect by the intensity of checking. As the number of rotating checkpoints per week increases, the negative effect on night accidents increases (column 3), further demonstrating the effectiveness of this intervention strategy. The effect of rotating monitoring on road deaths also increases from 1 to 2 checkpoints per week, though the coefficient on 3 checkpoints per week is (insignificantly) smaller than 2 per week. The coefficients values do not suggest increasing returns to intensity in this range of police enforcement although the standard errors are large at this level of treatment disaggregation. As predicted by the model, increased intensity has no effect on night accidents or deaths in stations in police stations with fixed checkpoints.

We now turn to the data collected by surveyors on the number of drunken drivers apprehended at the police checkpoints. Figure 9.1a shows the drunken drivers 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 as long there are 2 or more checks per week, more inebriated drivers seem to be caught at rotating stations despite the fact that rotating checks are held in ex-ante worse locations. Moreover, in fixed locations, the number of drivers caught per night of checking declines with the intensity of checking (which was randomly assigned), while this 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.

Even if learning is going on, fixed checks should apprehend more drunken drivers than random checks at the start of checking since they are performed at the best place. We investigate this in Figure 9.1b, which focuses only on week 1 and 2. At one check point per week, this is what we find: the number of drunk drivers caught is significantly greater in fixed. 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 checks per week. This is probably because there is learning already within the first two weeks and it is easier to learn about the fixed location than when the location jumps around. 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 early in the intervention period. Figure

9.1c 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 analogous regressions on the number of drunk drivers caught. Column 1 shows that, on the average night, 0.13 more drunk drivers, or 10% more, are caught at rotating check locations than in fixed locations. This number is not significant, but note that without strategic behavior 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. That effect entirely disappears in random locations. Since we see no significant effects of the fixed checkpoint intervention on road accidents and deaths, these results suggest 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 increases, the difference between rotating and fixed grows. 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.⁴ Strikingly, these results are robust to limiting the data to the 194 checkpoints held on road 1 in either rotating stations with 3 checks per week, or fixed stations with 1 checkpoint per week. In this sample, checkpoint 1 is manned once per week in both types of intervention, but we still find a significant negative effect of the number of previous checkpoints in fixed stations, and no such effect in rotating stations.

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.

⁴A 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.

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 probably comprise only 2-3% of all traffic. However we do not know how non-drunken drivers perceived the checkpoints. It could be that other violators (say drivers without a license) were also attempting to avoid checkpoints. Or they might just be reacting to the delays and harassment caused by the checkpoints. Understanding the sources of this reaction is important for 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 how the drivers learn about the end of the intervention. 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%-50% 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 on the number of days since the final checkpoint 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 for location 1 is shorter when the checkpoints rotate because there is no easy outside option that allows drunk drivers to avoid all checks. 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 in passing drivers in fixed stations (though not significant, $p = .166$), and a decline in rotating stations. As time passes after the intervention, the number of vehicles on roads 2 and 3 significantly decreases in fixed stations, presumably as drivers revert to their preferred route. Curiously however, we see no corresponding increase in passing vehicles on road 1 as days go by after the intervention. In rotating stations we see no signs of reversion.

Our final direct measure of the impact of the interventions occurred at the last night of checking. As described above, this final check was always held at checkpoint location #2, including in the fixed checkpoint stations and control stations which had previously conducted them exclusively on route 1 (in the fixed stations) or never (in the controls). This allows us to measure the displacement effect even in the fixed-location stations which, by definition, is only possible to do a single time before affecting the subsequent beliefs of drivers in that area. 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.⁵

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 as time passes since the last crackdown, 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 and is 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).

Surprisingly, the results from fixed checkpoints also show very significant decreases in final check drunken drivers relative to the control (columns 5, 6), even though final checkpoints were at location #2, but in fixed stations there was only checking at location #1. Thus one might have expected that checking on road 1 would lead to displacement of drivers onto road 2. However, in fact the net effect of checks on road 1 seems to have been to reduce them on roads 2 as well. As we will show below, this result is not inconsistent with the model, since drivers might have had priors

⁵Since 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.

that checking on road 2 would also be intense in the event of a crackdown. If they occasionally travel both roads, and notice a checkpoint on road 1, they would then conclude that any drunken driving might be too risky. These results highlight the value of estimating crime displacement patterns, and developing a structural model to interpret the results and design policy accordingly.

In addition to our analyses of the main impacts of the anti-drunken driving campaign and its reduced form effects on driver behavior, for robustness we examine several additional mechanisms and outcomes. These analyses are discussed at length in Appendix Section A2 with tables in Appendix Section A3; here we briefly summarize their results. We examine spillover effects to nearby police stations, both in terms of accidents and deaths, as well as the number of drunken drivers apprehended. Consistent with relatively isolated and independent nature of police stations in rural Rajasthan, we find no significant spillovers. If different treatment arms (particularly the rotating vs. fixed interventions) were correlated with different police behavior on other dimensions, this might affect our interpretation of the results. We examine the relationship between the intervention and police implementation of checkpoints, as well as the share of drunken drivers who paid a fine in court (a proxy for lack of police misconduct). Neither outcome is significantly different between rotating and fixed interventions, although we find that police implementation in 3 checkpoints per week stations is lower than in less intense interventions. Finally, we examine the possibility that individuals might also learn about police activity while driving sober or through social learning from friends. If this were the case, we might expect a very rapid decrease in drivers caught the day or two after a checkpoint as news spreads rapidly, even if only a few individuals personally witnessed the checkpoint. In fact, we find no particularly strong effects of checkpoints in the very short-term. Thus while we cannot fully rule out forms of learning outside the model, we find no evidence of their implications in the data.

6 Structural Estimation

6.1 Empirical Model

While the simplified model introduced in Section 5 (and formally presented in Section A1.8) suggests the mechanisms behind drivers' and police decisions, estimating the

parameters of drivers' preferences and beliefs and designing the optimal police enforcement strategy requires a more detailed model of driver behavior. First, let the utility that agents receive from drunken driving be d_r for $r \in \{1, 2, 3\}$, the roads potentially checked by police. If they encounter the police, agents experience disutility c .⁶ Drivers are uncertain about police checking intensity on each road; driver i 's beliefs about the probability of encountering a checkpoint on road r at night t is distributed Beta $(\alpha_{irt}^\lambda, \beta_{irt}^\lambda)$, where α_{irt}^λ and β_{irt}^λ are the parameters of the Beta distribution with expected intensity $\lambda_{irt} = \alpha_{irt}^\lambda / (\alpha_{irt}^\lambda + \beta_{irt}^\lambda)$. Similarly, we assume drivers believe that the crackdown has some probability of ending each night, with this probability distributed Beta $(\alpha_{it}^\eta, \beta_{it}^\eta)$ with expectation $\eta_{it} = \alpha_{it}^\eta / (\alpha_{it}^\eta + \beta_{it}^\eta)$. 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 12 variables, $\Psi_{it} = \left(\left\{ \alpha_{irt}^\lambda, \beta_{irt}^\lambda \right\}_{r=1}^3, \alpha_{it}^\eta, \beta_{it}^\eta, \pi_{it} \right)$ constitute the state variables in the driver's decision problem.

6.1.1 Choice Probabilities

At the beginning of each period, agents choose whether to drink and drive, and if so on which road, or take an alternative action which we refer to as abstaining from drinking and driving by staying at home.⁷ They perform their highest-utility action, taking into account possible value of information, and observe the results: if they drink and drive they may witness a check point, while if they 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_{it} that the crackdown ended between periods, and these beliefs become the priors in the next period, Ψ_{it+1} .

Consistent with the implementation of anti-drunken driving campaign by the Rajasthan police, we assume that drivers are not aware of the start date of the

⁶The 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.

⁷Our primary approach thus assumes that the cost of exploratory driving while sober is always above the return it would generate, and thus does not include this option. We examine this issue in further detail below.

enforcement.⁸ This implies that, until they encounter the first checkpoint on the road, drivers' beliefs remain constant. Period 0 beliefs about the duration of the crackdown (η_{i0}) should therefore be interpreted as pertaining to the remaining period of the enforcement campaign. In control police stations, where drivers never encounter a checkpoint, we assume drivers' beliefs remain constant throughout the intervention period.

We also assume that drivers believe that the crackdown consists of a single discrete interval of enforcement and does not start and stop several times in short succession. These beliefs correspond with the actual reality of Rajasthan Police crackdowns, and make the specification of learning behavior much simpler. In particular, if a driver is checked for the second (or more) time on a given night she is sure that the crackdown has persisted for all nights since her initial encounter with the police. Therefore each of her previous driving experiences, including those in which she did not encounter the police, is equally informative of the intensity of the current crackdown.

Drivers' behavior is not entirely deterministic: they experience random shocks that encourage them to drink and drive when it would otherwise be too risky, or to stay at home when they might otherwise go out for a drink. These shocks affect driver's actions, 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 $h(x_t, \epsilon_t; \Psi_0)$ where x_t denotes the history of checkpoints in the police station from periods 1 to t , and ϵ_t denotes the past history of shocks for all drivers in the jurisdiction. An individual driver's beliefs in period t , $\Psi_{it}(x_t, \epsilon_{it}; \Psi_0)$, are a function of his own shocks ϵ_{it} as well as the the history of checkpoints in his station; for conciseness we suppress this dependence in subsequent notation.

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

$$u_r(\Psi_{it}) = u_r = d_r + \pi_{it}\lambda_{irt}c \tag{6.1}$$

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. If

⁸We considered randomizing an information campaign announcing the crackdown across police stations. However, since we could not guarantee that all drivers would be informed in the time available for this campaign, we defer the analysis of this counterfactual policy to Section 6.4 based upon the estimated structural parameters.

the driver witnessed a police checkpoint on road 1, we denote his posterior beliefs as Ψ_{it+1} (*check*₁), if sees no checkpoint, then his beliefs become $\Psi_{i,t+1}$ (*nocheck*₁). Beliefs following travel on roads 2 and 3 are defined analogously, and those staying at home are $\Psi_{i,t+1}$ (*home*). Dynamic utility, incorporating the preference shocks is then,

$$V(\Psi_{it}) = \max \begin{cases} v_1(\Psi_{it}) + \epsilon_{i1t} & = u_1 + \delta(\pi_{it}\lambda_{i1t}V(\Psi_{i,t+1}(\textit{check}_1)) + (1 - \pi_{it}\lambda_{i1t})V(\Psi_{i,t+1}(\textit{nocheck}_1))) + \epsilon_{i1t} \\ \dots & \dots \\ v_H(\Psi_{it}) + \epsilon_{iHt} & = \delta V(\Psi_{i,t+1}^1) + \epsilon_{iHt} \end{cases} \quad (6.2)$$

Consistent with much of the empirical learning literature (Ching, Erdem, and Keane, 2013), we assume that the choice-specific shocks in 6.2 are distributed IID extreme value type 1. Incorporating these shocks, the fraction of potential drunken drivers who choose to drive on road r on night t is μ_{rt} :

$$\mu_{rt} = \int_i \frac{\exp(v_r(\Psi_{it}))}{\sum_{r'} \exp(v_{r'}(\Psi_{it})) + \exp(v_H(\Psi_{it}))} dh(\Psi_{it}; \Psi_0, \{d_r\}_{r=1}^3, c, x_{t-1}, \epsilon_{t-1}) \quad (6.3)$$

6.1.2 Evolution of beliefs

Once drivers have made their choice of road and observed any police activity, the evolution of their beliefs is a simple application of Bayes rule to the Beta distribution. Let \mathbb{I}_{it} be an indicator equal to 1 if the agent has witnessed a checkpoint on any road prior to night t . Agents' expectation that the crackdown ends after night t (again, conditional on it being in progress in period t), is $\eta_{it} = \tilde{\alpha}_0^\eta / (\tilde{\alpha}_0^\eta + \beta_{it}^\eta)$. Because drivers never conclusively observe that the crackdown has ended, their beliefs about its duration evolve only as the β_{it}^η term in this Beta distribution grows over time:

$$\beta_{i,t+1}^\eta = \beta_{it}^\eta + \mathbb{I}_{it}$$

In this case η_{it+1} evolves deterministically based solely on the time passed since the driver first encounters a checkpoint: 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.⁹

In contrast, beliefs about the intensity of the crackdown do depend upon drivers' actions and observations after they are aware of the police campaign. Consider a

⁹Since the driver never conclusively observes the end of the crackdown, $\tilde{\alpha}_t^\eta = \tilde{\alpha}_0^\eta$ for all t . This reduces the number of endogenously evolving state variables to 11.

driver who has chosen to take road r in period t , and saw a checkpoint. Posteriors beliefs on the expected checking intensity on road r , are

$$\lambda_{ir,t+1} = \frac{\alpha_{irt}^\lambda + 1}{\alpha_{irt}^\lambda + \beta_{irt}^\lambda + 1}$$

If the driver does not pass a checkpoint in period t , but has encountered one in the past (and thus knows the crackdown is underway) the $\alpha_{irt}^\lambda + 1$ term in the numerator is replaced with α_{irt}^λ but the denominator remains the same. If the driver had chosen to remain at home, her beliefs about checking intensity would not evolve: $\lambda_{irt} = \lambda_{ir,t+1}$. As mentioned above, if the driver has never encountered a checkpoint, his beliefs remain static: $\lambda_{irt} = \lambda_{ir0}$.

Once a driver observes that the crackdown has begun, each time he travels on a road his belief about the intensity of police checking becomes more precise. This is modeled through the β_{irt}^λ term, which evolves according to $\beta_{ir,t+1}^\lambda = \beta_{irt}^\lambda + 1$. As above, β_{irt}^λ remains unchanged in the period before the driver is aware that the crackdown has begun, and if the driver chooses to remain at home or drive on any road other than r .

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 probability 1 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

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

This learning behavior can generate complex patterns of travel and staying home by the drivers. To illustrate these strategies, Figure 9.2 shows the simulated histories of two drivers, the first in a station with only checking on road 1, and the second in a station that also has checking on alternative roads 2 and 3. The parameters used

to simulate these histories are those estimated from the data, and 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: blue for road 1, red for 2, green for 3, and gray 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 blue, red, and green lines show, respectively, λ_{1t} , λ_{2t} and λ_{3t} , and their scale is shown on the right-hand axis.

The driver in Panel A encounters a checkpoint on road 1 on night 23 of the crackdown. His beliefs about λ_1 spike upwards since (as we discuss below) drivers have very diffuse priors about the initial probability of checking on road 1. His belief about the likelihood that the crackdown is ongoing (π_t) also shoots up, reaching about 75%. For several weeks afterwards the driver has a positive belief that a crackdown is ongoing but, due to low prior beliefs about the probability of a crackdown on road 2 and (especially) 3, his drunken driving activity on these alternate roads is only slightly decreased. Finally, at night 74 the driver returns to road 1 and, not encountering a checkpoint, revises his posterior belief about the intensity of checking on road 1 downward to around 55%. Further trips on road 1 in periods 84 and 89 cause further reductions in λ_1 and eventually driver activity returns to normal.

Panel B depicts the history of a (particularly unlucky) driver in a police station with checking on all roads.¹⁰ This driver initially encounters a checkpoint on road 2 at night 20. Again, π_t rises to about 75%, but since his priors on λ_2 are very precise, there is not a visible increase in $\lambda_{2,21}$. The driver's awareness that the crackdown might be ongoing is sufficient to dissuade her from driving on roads 1 and 2, although she continues to drive on road 3 where λ_3 is almost 0. Unfortunately (for her) she encounters a checkpoint on road 3 at night 39, which increases her posterior on λ_3 to about .34 and sends π_t over .8. After about 3 weeks of sobriety, she re-attempts to drive drunk on road 2 at night 65, only to encounter a police checkpoint immediately. This scares her away from drunken driving for the remainder of the intervention. 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.

¹⁰This history was chosen for illustrative purposes and contains more interaction with the police than typical.

6.1.3 Alternative Empirical Models

An important element of the theory presented above is that agents change their beliefs about the parameters of police strategy over time. To evaluate the role of this learning in explaining the patterns observed in the data, we estimate an alternative “No Learning” model in which agent’s beliefs on $\{\lambda_r\}_{r=1}^3$ and η remain constant. Equivalently, one might also say that prior beliefs about these parameters are infinitely precise. This alternative model does still include learning about whether the crackdown is ongoing—thus when drivers encounter a checkpoint their $\pi_t = 1$, but the value of η (the rate at which π_t decays) is not affected by the number or timing of encounters with the police.

A fundamental result of the theoretical literature on learning in economics is that forward looking agents value information and base their actions partly on the incentive to increase their quantity of information. Yet empirical work on learning has often found that this “exploration” motive has a minimal effect on improving model fit (Ching et al., 2017). Specifying this alternative, “myopic” model can be done by simply setting $\delta = 0$ in Equation 6.2. The assumption of myopia is particularly interesting to explore in this setting, given the substantial literature on alcohol induced myopia and its effect on the choice to drive drunk (Giancola et al., 2010).

The model presented above does not include either social learning or sober driving for the purpose of learning about the presence of police enforcement. Both of these extensions would allow drivers to learn about police activity without paying the cost of encountering a checkpoint, thereby speeding learning about the beginning and end of the crackdown. Unfortunately our intervention is limited in the information it can reveal about such behaviors, since none of the treatment wings specifically manipulated either of these mechanisms. We therefore focus our empirical investigation of these factors on a single, admittedly extreme, alternative model in which all drivers in a police station are instantaneously aware of any checkpoint that occurs on any road in that area. This captures the case in which social learning is perfect and/or drivers are able to costlessly check all roads for police activity every night. Learning is exogenous, driver beliefs are homogeneous, and the agent’s problem is equivalent to the static choice of the highest expected payoff road. Details of this alternative model are presented in Appendix Section A2.1.

6.2 Estimation

We estimate the structural model using simulated method of moments. In short, we first match the number of drunken drivers caught and road accidents on each night with the corresponding numbers predicted by simulating the model given the history of checkpoints and checkpoint locations in each station. This generates a residual which we use to create a GMM criterion which minimizes the correlation of the residuals and instruments generated from the random assignment of police stations to the different treatment arms.

The model contains three sets of parameters. The first are the initial conditions of drivers' beliefs, $\theta_1 = \left(\left\{ \alpha_{r0}^\lambda, \beta_{r0}^\lambda \right\}_{r=1}^3, \alpha_0^\eta, \beta_0^\eta \right)$. 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 4 preference parameters: the driver's utility of drunken driving on roads 1, 2, and 3, and their disutility of encountering a checkpoint: $\theta_2 = \left(\left\{ d_r \right\}_{r=1}^3, c \right)$.¹¹ Combined, these form the 12-element parameter vector that determines driver behavior and hence allows the calculation of the optimal police crackdown strategy. The third set of parameters, θ_3 , is a vector of nuisance parameters controlling for other exogenous characteristics of the police station jurisdiction and checkpoint implementation that may affect police effectiveness or the local number of potential drunken drivers, but not the drivers' learning and decision process. We discuss the elements of θ_3 further in the estimation section below.

We follow a four step procedure to generate the model predictions of the number of accidents and drunken drivers caught which we subsequently take to the data. First, we numerically solve for the value function $V(\Psi_{it})$. This is complicated by the relatively high number (8) of continuous state variables and the infinite-horizon nature of the problem, making backwards induction impossible.¹² Although heuristic algorithms (e.g. Scott (2010)) can achieve near-optimal payoffs in bandit problems as the number of periods increases, agents' exploration behavior in these heuristic approaches may be very different than under the optimal strategy. We therefore prefer to directly compute the optimal driver strategy using value function iteration, but do so on a finite grid of the state space. When the future state falls between the

¹¹The baseline specification sets the discount rate to $\delta = .95$, but we present results with alternative values.

¹²Because this case falls into the category of restless correlated Bandit models, the Whittle Index results that simplify calculating the optimal solution in most restless Bandit models do not apply.

grid points, we interpolate.

Second, we simulate agents' actions as they encounter the sequence of checkpoints that occurred in each police station. For each of the 223 stations in the data we generate $H = 5,000$ simulated histories of potential drunken drivers' choices on each night of the crackdown and for up to 90 days afterwards. However, as discussed above not all of the assigned checkpoints were actually conducted by the police and we only have data on the implementation status of 76% of the checkpoints (either through surveyor observation or GPS tracking of the police vehicle). Thus in 24% of cases we do not know with certainty whether a checkpoint was actually implemented. We incorporate this fact into the simulations by allowing the fraction of agents who encounter an assigned checkpoint to be less than 1 whenever the actual implementation status is unknown, and specifically to depend upon the treatment group of the police station. We first run a probit regression of an indicator that the police implement the checkpoint on the fully interacted set of all intervention categories (including the police line intervention, which affects implementation, as we saw in Table A2), plus variables indicating checkpoints that occurred in the early, middle, or late stage of the intervention, and other police station characteristics.¹³ From this regression we predict each unobserved checkpoint's compliance probability and use this to determine how many of the simulated agents choosing to drive on that road would encounter the checkpoint on that night. For example, suppose 100 simulated agents in station i find it optimal to drive drunk on road 2 on night t , when there was a checkpoint assigned on road 2, and station i would be predicted to implement the checkpoint with 75% probability. Then $100 \times .75 = 75$ agents would experience the checkpoint, and hence update their beliefs negatively about checking on road 2, while the remaining 25 would see no police presence and update positively.

The third step is to calculate each driver's probability of travel on each road using the choice probability in Equation 6.3, and average these over all simulated drivers to generate an estimate of the share of drunken drivers on each route on every night. To smooth the objective function, we calculate the mass of drivers on each road using summed choice probabilities, not the realized choices. This is an application of the smoothing technique discussed in Bruins et al. (2018).

¹³This regression is analogous to column 1 of Table A2, with the inclusion of the interactions of the intervention categories as a variety of additional control variables. Specifically, we include dummies for the day of the week, whether the station is on a National Highway, the total past accidents in the station, urban/rural status, 2010 or 2011 intervention, and district fixed effects.

The fourth and final step is to incorporate the simulated probability of drunken driving on each road r (μ_{srt}) into model of the number of accidents observed (A_{st}) and drivers arrested (N_{srt}). This requires modeling three additional factors: the total number of potential drunken drivers, their propensity to get into an accident (to match them with the number of accidents observed), and the local police effectiveness at catching them (to match data on drivers arrested). Since our model has little to say on these issues, we model them as two vectors of nuisance parameters, θ_3^{Acc} and θ_3^{Arr} which are associated with the number of accidents and arrests, respectively.

We model the baseline accident rate as a reduced form function of district fixed effects and indicators for urban or rural location, proximity to a National Highway, day of the week, and the number of past accidents. We also include dummy variables for assignment to the fixed and rotating intervention treatments to control for any random baseline differences in assignment; as described below our identifying variation derives from changes in accident rates between the stations over the course of the intervention. The vector of these controls for a given station-night is w_{st}^{Acc} and the corresponding coefficients on these variables are the vector θ_3^{Acc} . This baseline accident rate, expressed as $\exp(w_{st}^{Acc'}\theta_3^{Acc})$ in the estimation, is shifted by the intervention since it is multiplied by the share of potential drunken drivers who choose to take any of the 3 roads modeled: $\sum_{r=1}^3 \mu_{srt}$. The structural equation for the number of accidents also includes an intercept to account for the presence of non-drunken driving related accidents: even if all drivers remained sober ($\mu_{srt} = 0 \forall s, r, t$), road accidents would still occur.

We include the same vector of controls when modeling baseline police effectiveness in catching drunken drivers, but also include variables related to the checkpoint itself: the number of vehicles stopped, and the time spent at the checkpoint, each included as a 3rd degree polynomial. We also include controls for whether the checkpoint occurred on the final check or was conducted by a Police Lines team, since these checkpoints may have been differentially better or worse at identifying law-breakers. We combine all these controls into a checkpoint-level vector w_{srt}^{Arr} , with their impact on the number of drunks caught parameterized by θ_3^{Arr} , which enters the model as $\exp(w_{srt}^{Arr'}\theta_3^{Arr})$. This baseline is then shifted by the model's prediction of the share of drivers drunk on that road-night, μ_{srt} . We do not include an intercept in the structural equation for arrests, since if all drivers chose to remain sober there would be no arrests for drunken driving.

The total number of accidents in police station s on night t , A_{st} , and arrests of

drunken drivers, N_{srt} , caught on road r is,

$$A_{st} = \left(1 + \sum_{r=1}^3 \mu_{srt}\right) \exp\left(w_{st}^{Acc'} \theta_3^{Acc}\right) + \zeta_{st}^{Acc} \quad (6.4)$$

$$N_{srt} = \mu_{srt} \exp\left(w_{srt}^{Arr'} \theta_3^{Acc}\right) + \zeta_{srt}^{Arr} \quad (6.5)$$

where ζ_{st}^{Acc} and ζ_{srt}^{Arr} are IID mean-0 error terms capturing factor such as holidays, idiosyncratic weather conditions, shocks to police manpower, etc. Let the stacked vector of structural errors be ζ . If all scheduled checkpoints were conducted as planned, then each station's checking history would be random, and $E[\zeta|w_{srt}, h_{srt}] = 0$. In this case we could estimate the structural θ parameters by simply searching for the θ vector that minimizes ζ subject to some loss function. However, since only about 63% of the scheduled checkpoints occurred, we must take into account the possibility of correlation between the checkpoint histories and the residuals. This might occur, for example, if police are more active in implementing the intervention when or where there is a larger problem with drunken driving.

To address this source of bias, we employ a non-linear GMM estimator that uses functions of the *assigned* checkpoints as instruments. Since this assignment was randomized by the intervention team, as discussed in Section 3, it is by definition orthogonal to any unobservable components of checkpoint effectiveness. In practice we include 18 variables in the instrument vector, Z , divided into one set for the accident moments Z^{Acc} , and another for the arrest moments Z^{Arr} .

The 8 instruments in Z^{Acc} contain interactions of indicator variables designating whether fixed and rotating strategies were used in a police station with the number of prior assigned checkpoints (2 instruments), and those same interactions further interacted with number of weeks since the end of the intervention (2 instruments). These identify the differential decrease in accidents in fixed and surprise police stations during the intervention, and the eventual reversion of accident rates in both types of stations. Z^{Acc} also includes indicators for the night immediately after a checkpoint, and the interaction of these indicators with the number of previous checkpoints. Both of these variables are separated by fixed and rotating checkpoints, for a total of 4 more instruments. These allow identification from the very short-term learning effect, which helps to identify the share of drivers who witness a checkpoint and are thus deterred from drunken driving on the following night.

The vector of instruments based upon survey and arrest data, Z^{Arr} contains 10 elements. The first 2 are simply indicators for whether the checkpoint was held on the 2nd or 3rd choice roads—these identify the relative utility of the different roads. The next 4 instruments are the interactions of the number of previous checkpoints held in a police station with a indicator for whether the police station had fixed or rotating checkpoints and, if the checkpoints were rotating, separate instruments for each type of road on which the checkpoint was held (1,2, or 3). The last 4 instruments come from the final check, and contain the interaction between indicators for fixed or rotating checkpoints with a indicator variable for data from the final checks, as well as these same variables interacted with the days since the previous checkpoint.

Letting \mathbf{Z} be the matrix of instruments and $\boldsymbol{\varepsilon}$ be the vector of residuals, the GMM objective function is the the standard,

$$\min_{\theta \in \Theta} (\mathbf{Z}'\zeta(\theta))' (\mathbf{Z}'W\mathbf{Z})^{-1} (\mathbf{Z}'\zeta(\theta))$$

We employ 2-stage GMM, first solving for the values of θ that minimize the GMM criterion with $W = (Z'Z)^{-1}$, then re-estimating the parameters using the optimal weighting matrix. The estimation is complicated by the non-smooth criterion function and the presence of local minima. To ensure that we have identified the global minimum we first use the Particle Swarm global optimization algorithm to identify the neighborhood of the global minimum, then use the Nelder-Mead algorithm to refine the solution.

6.3 Results: Structural

6.3.1 Primary Specification

Table 10 presents the results estimation of the structural parameters. To simplify the interpretation of the results, we report transformations of the structural parameters that are more easily interpretable than the fundamental beta distribution of driver's beliefs.

Rows 1-4 display the parameters of the drivers' utility functions. Since the location and scale of these coefficients are determined by the value of the outside option to drinking and driving (normalized to 0) and the standard deviation of the extreme value shocks ($\pi/6$) the values of these coefficients per se cannot be readily interpreted.

The relative sizes are informative, however. At 5,842 the perceived disutility of being caught by the police is hugely greater than the benefit to drinking and driving, showing that sobriety checkpoints are an efficient tool for crime prevention if deployed effectively. Consider a pre-crackdown scenario in which drivers perceive 0 risk of being caught drunk: this corresponds to the first night of checking in each police station. The estimated parameters suggest that the majority of potential drunken drivers (63.3%) will not be driving drunk on any of the roads under consideration, while 14.4% would be inebriated on road 1, 11.0% on road 2, and 11.3% on road 3. While the standard errors on this parameter are large, this is due to the fact that the GMM criterion is not very sensitive to changes in the cost of encountering a checkpoint once it becomes so large that drivers are dissuaded for long periods once they learn of the campaign. For example, the GMM analogue to the LR test (Newey and West, 1987) of the hypothesis that $c = 0$ yields a χ^2 statistic of 36.2, significantly rejecting the null.¹⁴ While these results, taken alone, are not surprising, they provide inputs to designing a more effective strategy for preventing drunken driving.

Rows 5 and 6 contain the driver's priors on the duration of the crackdown. Drivers initially believe that the crackdown will end with 20.7% probability after just one night and that the expected duration of the crackdown is only about 3 days. This is consistent with what senior police officers told us of past police enforcement of traffic laws, which consisted of only 1 or 2 checkpoints per initiative. However, the drivers do adjust their beliefs about enforcement duration if they witness a prolonged police crackdown. The initial precision of priors on η implies that if a driver witnesses a second checkpoint 30 days after they first become aware that the crackdown has begun, their belief about the probability of the campaign ending the next night drops to 13.70%. Standard errors on these parameters are quite large, although once again if we test parameter restrictions implying larger deviations from the estimated parameters the LR tests strongly reject the restrictions. For example, testing $\eta = 0$ (permanent crackdowns) yields an χ^2 statistic of 112.4, while a test of the restriction $\eta = 0.95$ (very short crackdowns) yields a χ^2 statistic of 23.0, both significantly rejecting the restrictions.

Finally, rows 7-12 show the initial priors on the expected intensity of checking on

¹⁴Since the restriction $c = 0$ implies that driver behavior is totally unaffected by the police crackdown and thus the μ_{srt} shifters are constant, this restriction can also be seen as a broader specification test of the whole model.

roads 1,2, and 3. Estimated beliefs are substantially different across all three roads. On road 1, we estimate a prior on the checking intensity of $\lambda_{1,0} = .090$, but extremely rapid learning of $\alpha_{1,0}^\lambda + \beta_{1,0}^\lambda = 6.1 * 10^{-7}$. These extremely diffuse priors imply that the prior mean is only relevant for agents who know that a crackdown is ongoing, but have not yet been stopped on road 1, that is after they have been stopped on road 2 or 3 in the rotating station. Hence standard errors on $\lambda_{1,0}$ are very large. Once drivers encounter a checkpoint on road 1 their posterior rises to almost 1 (as in Panel A of Figure 9.2). These very diffuse priors rationalize the empirical observation that the number of drunken drivers on the road 1 decreases very soon after the intervention has begun.

Conversely, on road 2 agents have a much lower but non-zero prior on checking intensity of $\lambda_{0B} = 0.006$, though their priors are much more precise with $\tilde{\alpha}_{0B}^\lambda + \tilde{\beta}_{0B}^\lambda = 8,807$. These parameter help rationalize the empirical result that the intervention was most effective at reducing accidents in the rotating stations (Table 4), but at the time of the final check there was a roughly equal displacement effect on the road 2 in the fixed and rotating stations (Table 9). This implies that agents are effectively dissuaded from driving on road 2 for the duration of the intervention since the expected disutility of driving on road 2 ($\pi_t \times .006 \times 5,942 = \pi_t \times 35.6$) is sufficient to prevent travel on road 2 even for quite low beliefs about the probability of police activity (π_t). The deterrence effect occurs even when drivers have not personally witnessed a checkpoint on road 2 (as in the fixed checkpoint intervention). However, after the intervention, once π_t falls very low, drivers' average beliefs about λ_2 remain low even in the rotating stations, since they are not substantially revised upwards by experience. Thus the share of drunken drivers in road 2 in the post-intervention period is not very different in the rotating or fixed police stations.

Finally, the estimated parameters suggest that drivers' prior beliefs about the probability of police patrolling the 3rd road are extremely low: $\lambda_{3,0} = 7.567 * 10^{-7}$. Encountering checkpoints on road 1 (as in the fixed intervention) has little impact on individuals' expected utility of drunken driving on the 3rd road, since agents do not expect to encounter the police there even if a crackdown is ongoing. However, if they do observe a checkpoint on road 3 (as in the rotating intervention) they update their posteriors very quickly because $\alpha_{3,0}^\lambda + \beta_{3,0}^\lambda = 0.009$. These parameter help rationalize the finding in the data that the fixed checkpoint location intervention is not effective at reducing drunken driving (because agents move to road 3), but the

rotating intervention can be effective because after moving to road 3 agents learn very rapidly about police presence there.

In Figure 9.3 we present graphs of the share of agents drinking and driving in the stations of each branch of the intervention. The blue-solid line shows the share on road 1, the red-dotted line on road 2, and the orange-dashed line on road 3. The first row of sub-graphs, illustrating the fixed intervention stations, demonstrates clearly the spillover effects resulting from checking on only 1 route. The share of drunken drivers on road 1 drops sharply in the first days of the campaign, particularly in the stations with 3 checks per week. At first, the number of drivers on road 2 drops as well, since agents infer that if there is checking on route 1, there is also likely to be checking on 2 as well. However, this rebounds quickly as drivers' perceived likelihood of the crackdown continuing drops and they become willing to drive on road 2 again, where they never personally witnessed a checkpoint. In contrast, the number of drunken drivers on route 3 increases over the first month of the campaign as agents are displaced from routes 1 and 2. This causes the overall impact of the intervention to be limited in the fixed stations. The second row of Figure 9.3 shows the corresponding graphs from the rotating intervention stations. Here we see a decrease in drunken driving on all 3 roads, leading the program to be more effective overall. The rate of decrease is slowest on road 3, since drivers' perceptions of risk on this road are not strongly affected by checking on roads 1 and 2.

Can this model of driver behavior provide a reasonable approximation to the observed results of the crackdown implemented by the Rajasthan Police? While we cannot conduct an out-of-sample test (there was no variation in enforcement outside this program), it is not a forgone conclusion that the model results would closely match the reduced form results in Section 5. This is partly mechanical—there are substantially more moments (or reduced form results) than estimated parameters (12), and the moments estimated do not exactly correspond to the reduced form results. More substantively, the functional forms of Bayesian updating and the assumptions of common Beta-distributed priors place considerable structure on the impact of crackdowns on driver behavior. For example on a given road, increased enforcement must lead to fewer drunken drivers, and the magnitude of this impact must be decreasing in the long-run.

Given these degrees of freedom, we evaluate model fit by simulating a data set of accidents and drunken drivers caught based upon the model, estimated parameters,

and actual police checkpoint implementation patterns, then re-running the same reduced form regressions used in the analysis of the actual data in Section 5 on this simulated data set. Maintaining the same order of reduced form regressions as Section 5, we begin with the accidents analysis. Here the moments matched by the model are quite different from the reduced form regression specifications which include police station (rather than district) fixed effects, additive rather than multiplicative controls, and are at the monthly (instead of daily) level. Column 1 Table A7 shows police station-FE analysis of the simulated data predict that fixed and rotating checkpoint locations perform similarly in reducing night accidents, with a slight edge for the fixed strategy. Since this contrasts with the (admittedly noisy) actual results in Table 4 which show the rotating stations performing better, what explains this deviation? In column 2 we show that switching to a district-FE specification restores the result that the rotating intervention reduces accidents more than the fixed-checkpoint strategy, albeit at a different level of variation (daily vs. monthly). Since the only police-station level data that enters the simulation is the performance of the staff in implementing checkpoints, this suggests that the discrepancy is due to correlations between police compliance carrying out checkpoints and the fixed treatment status (within district). These are captured by the police station fixed effects, but not in the structural estimation.

The model's predictions on the dynamics of drunken drivers caught at checkpoints match the reduced form estimates well. The simulated results in Table A8 show the exact same patterns of decreasing effectiveness in fixed but not rotating police stations as in the corresponding columns 1-4 of Table 7, including those with police station fixed effects (columns 3-4). The model is also able to match the results from the final check quite closely. Table A9 (corresponding to Table 9) shows substantially fewer drunken drivers caught at the #2 road at the final check, and that this difference is are greater in the fixed checkpoint stations where there was no enforcement on road 2. Simulated reversion is faster in the fixed-checkpoint stations (columns 5 vs. 3)—also matching the empirical data. Thus, with the exception of the police-station FE accident regressions, the model can match the main empirical findings well. In particular, the apparent paradox that the fixed checkpoint intervention causes greater displacement of drunken drivers on the unpatrolled road 2 but is less effective than the rotating intervention at catching drunken drivers on road 1 or stopping accidents (with district FEs) can be rationalized by the beliefs and structure of the model.

6.3.2 Alternative Specifications

In addition to our primary model, we present three alternative specifications to illustrate the ways in which the different parameters and assumptions affect the estimated results. In the first of these, we eliminate all learning by drivers about checkpoint intensity or the duration of the crackdown. This corresponds to setting the parameters $\{\beta_{0r}^\lambda\}_{r=1}^3 = \infty, \beta_0^\eta = \infty$, so that drivers initial priors about λ_r and η remain constant over time. We do allow drivers to learn about whether the crackdown is ongoing (π_t), so that they correctly conclude that the crackdown is still continuing if they observe a checkpoint. Since this specification is nested within our primary model, the difference in estimated parameters and policies demonstrates the potential bias in not taking learning into account. Given the importance of learning in the baseline model, it is perhaps not surprising that the value of the GMM criterion is almost twice as high in the “no learning” model. The values of d_r increase slightly, causing an increase in the initial rate of drunken driving and thus faster learning about the beginning of the crackdown. Perhaps the main difference is that $\lambda_{3,0}$ is now substantially higher at 0.022. This value is sufficient to cause substantial deterrence on road 3 even in the fixed checkpoint intervention—thus the no-learning model predicts that the fixed checkpoint campaign should be more effective than it actually was.

Our second alternative specification assumes that agents are capable of learning about all parameters of the crackdown, but are myopic when selecting their choice of road. This corresponds to setting $\delta = 0$ in the driver’s value function, as opposed the value of $\delta = .95$ used in the main specification. Re-estimating the model under this alternative specification column 3 of Table 10 shows very little change in the estimated parameter values or model fit. This lack of sensitivity to the discount rate can be attributed to the very high cost of being apprehended by the police relative to the low value of active exploration. For example, consider an agent who is stopped by the police on road 1 on the first night of the crackdown, when beliefs are at their most malleable. The estimated parameters imply that the dynamic value of driving on road 1, which captures the value of the information gained through experimentation, is only 1.06% higher than the dynamic value of remaining at home where no new information will be gathered. Yet the expected cost of driving on road 1 is $-4,122.23$ versus a utility value of 8.98 for not driving at all. Thus while a bandit model of the type estimated here might, in principle, have a substantial role for exploration and

an incentive for information gathering, the estimated parameters in this context are such that this effect is essentially absent.

The fourth and final alternative specification we examine is the case of costless learning about the state of the crackdown, as discussed in Section 6.1.3. This could correspond either to a case in which agents can costlessly “explore” all roads while sober to monitor police actions, or in which drivers perfectly inform each other whenever anyone passes a checkpoint (as in the app “waze”, which was not available to our drivers at the time but may be available in other contexts). Perhaps not surprisingly, this alternative model generates hugely different parameter estimates than the baseline model. Results in column 4 of Table 10 show values of d_r around 42—implying that agents drink and drive essentially every night in the absence of the intervention. This high utility of drunken driving is combined with an equally high cost of getting caught (47,568), and the belief that police crackdowns will end immediately with probability of almost 1 ($\eta = 0.999$). Similarly, priors on police enforcement on each road are estimated at almost 1, although standard errors on these parameters are very high. The implied pattern of driver actions is a dramatic fluctuations between all agents drinking and driving, then all agents remaining sober for a few days, then all agents driving again, etc., as shown in Figure A1 which displays the simulated actions of drivers in this model. Although this pattern of behavior appears counterintuitive, the value of the GMM criterion is only moderately higher than the baseline model (0.0097 vs. 0.0087). However, this similarity is due to the alternately large and small values of the residuals in the costless learning model canceling out in the summation of the GMM criterion. If, instead, we compare 1 minus the ratio of the sum of squared residuals to the total sum of squared variation (an informal analogue of the R^2 statistic) for the data on drunken drivers caught, we find that it is 0.35 for the baseline model, but only 0.10 for the model with costless learning. Thus, on a night-by-night basis, the model with individual learning seems to deliver a much better fit than if we were to assume perfect social (or costless) learning.

6.4 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 is 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.¹⁵ The second is to limit the police optimization problem to two parameters: the duration of the crackdown in terms of nights, and the fraction of checkpoints to be conducted off of the primary road (that is, on roads 2 and 3). Once these parameters are set, the specific timing of the 20 checkpoints is determined randomly with uniform probability on each night of the campaign, subject to the constraint that any station cannot carry out two checkpoints on the same night.

While this strategy would, by definition, be less effective than an unconstrained optimal strategy, it has the advantages of being more transparent, easier to calculate, and adhering more closely to the theoretical problem analyzed in Section A1.2.¹⁶ Since the actual Rajasthan anti-drunken driving intervention also complies with these constraints, one can interpret this counterfactual strategy as the best possible way to have designed the intervention that was carried out in Rajasthan, subject to the same constraints faced in reality.

The enforcement strategies suggested above do not take into account potential partial compliance with the crackdown protocol by the police. However, as we show in Table A2, the police are less likely to carry out a crackdown when the intensity of enforcement is high (particularly in the 3-checkpoint per week intervention). To incorporate this factor into the design of the counterfactual policy, we adjust the probability of a crackdown occurring to be $\Phi(.615 - .139(F^*))$ where F^* represents the assigned station-level frequency of checking and $\Phi(\cdot)$ is the Normal CDF.¹⁷ These strategies may be more reasonable than those that assume full compliance, particularly for policies that imply very high intensity checking.

The estimated parameters imply that the optimal strategy incorporating partial compliance (shown in Table 11) is to spread the 20 checkpoints over 82 days, and to place 7 of them on road *A*, dividing the other 13 across the roads 1 and 2. These results are consistent with the implications of the simple theoretical model developed in Section A1: the police are most effective when they spread the crackdown diffusely

¹⁵This 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.

¹⁶In particular, the police might want to have crackdown intensities change in more complex patterns over time.

¹⁷These parameters were estimated from the police implementation data in a probit regression analogous to column 2 of A2, except using a continuous measure of assigned frequency.

over a longer period. Of the 90 day cycle considered, almost all would be included in the crackdown. The allocation of checkpoints fairly evenly across roads derives from two implications of the parameters estimated in the structural model. First, drivers' estimated utilities are not hugely different across routes, thus there is no need to increase the risk of apprehension in any particular location in order to effectively dissuade criminals there. Second, because agent's priors on the risk of driving on roads 2 and (especially) 3 are low, the police must allocate forces to all areas in order to prevent diversion of criminal activity away from road 1. This campaign causes a 36% drop in the rate of drunken driving. If we consider partial compliance (Columns 4-6) the optimal strategy is basically unchanged, although somewhat less effective with a decrease in drunken driving of 25%.

An alternative strategy employed by many police departments is to announce the beginning of a sobriety campaign publicly, for example by radio advertisement, with the goal of dissuading drunken drivers immediately. In the second row of Table 11 we examine this alternative strategy, while maintaining the assumption that drivers' beliefs are as estimated at the beginning of the intervention. In fact, the Rajasthan police did consider a pre-intervention awareness campaign, so this counterfactual reveals the outcomes if that awareness campaign had been implemented. Computationally, to solve for the optimal policy we simply set $\pi_0 = 1$, and evaluate the grid of policy options as before.

The results, in row 2 of Table 11 imply that optimal duration should fall slightly—from 82 to 77 days, while the fraction of checkpoints on road 1 drops from 35% to 20%. The reason for this shift is straightforward: if the police announce a crackdown, most citizens will assume that road 1 will be targeted and will cease driving there following the announcement even without actually witnessing a checkpoint. However, since citizens (essentially) do not believe that there will be checking on road 3, the only way for the police to dissuade them from driving there is to conduct checkpoints at that location. The prior announcement of the campaign substantially improves its efficacy: the rate of drunken driving now drops by 46%. Again, incomplete police compliance has little effect on the overall strategy but decreases its effectiveness.

As discussed above, the intervention occurred in a setting in which there had previously been no enforcement of drunken driving laws. Thus the parameters estimated, and the optimal crackdown policies estimated from those parameters, could only extend to analogous contexts. However, the full structural model allows us to investigate

how these policy implications would differ in other settings, for example in cases where checking had been ongoing for some time and the population was familiar with police strategies. To operationalize this counterfactual we hold constant the estimated utilities $\theta_2 = (\{d_r\}_{r=1}^3, c)$, while imposing that the driver's beliefs $\theta_1 = (\{\alpha_{r0}^\lambda, \beta_{r0}^\lambda\}_{r=1}^3, \alpha_0^\eta, \beta_0^\eta)$ be equal, in expectation, to the police strategies and have very low variance.

To specify these equilibrium beliefs, first we set the number of prior trips, a measure of the precision of the prior beliefs, at $\alpha_{r0}^\lambda + \beta_{r0}^\lambda = \alpha_0^\eta + \beta_0^\eta = 10^6$ which captures the fact that beliefs are likely to be slow to change once citizens are very experienced. Second, we define the beliefs about police intensity to be correct in expectation. Let Q_r^* be the number of checkpoints allocated to road r for a crackdown lasting T days; citizen's long-run equilibrium beliefs about checkpoint intensity are $\lambda_{r,0} = Q_r^*/T$. Likewise citizens beliefs about the ending probability of the campaign are, $\eta_0 = 1/T$.

The third row of Table 11 shows that the “long-run” optimal policy estimates are substantially different from the optimal crackdown given the estimated beliefs. In this case a short and highly focused campaign is best: a duration of 21 days, and 75% of checkpoints on road 1. Because there is minimal learning, the goal of the crackdown is simply to inform as many drivers as soon as possible that checking has started and thus clear the roads rapidly of drunken drivers. The reasoning for the emphasis on road 1 is that most drunken drivers are located on road 1, so focusing the crackdown on this location is the fastest way to show them that the police are enforcing the law. It is still optimal to put 25% of enforcement on roads 2 and 3, since otherwise in equilibrium drivers will know that these locations are safe and will immediately shift to these routes. This strategy is extremely effective: drunken driving decreases by 90%. One reason for the increased effectiveness in equilibrium, relative to the initial sobriety campaign in row 1, is that beliefs about the average duration (21 days) are much longer than agents' estimated beliefs about crackdown duration in the data (about 3 days). Thus the initial discovery that the police are enforcing the law creates a much longer period of sober driving. Furthermore, once drivers encounter police on road 1, they know there is a significant chance of apprehension on roads 2 and 3, thus maximizing the negative displacement effect. In this situation accounting for partial compliance does make a difference in the strategy, extending the optimal duration to 43 days and shifting the location of 95% of checkpoints to road 1. However, the

impact on drunken driving is not very sensitive to changes in checkpoint allocation in this scenario—putting to 50% of checks on road 1 increases drunken driving by only 3%.

The final counterfactual we examine, in the fourth row of Table 11 presents the optimal policy if drivers have equilibrium beliefs about the parameters of the sobriety campaign and the police can give full prior warning of the beginning of the campaign. In the case the campaign is extremely effective under virtually all parameters of the crackdown. The effectiveness of the campaign is increasing in the duration of the crackdown—with a duration of greater the 31 days the numerical computation rounds the share of drunken drivers down to exactly 0 and the program is 100% effective! Effectiveness is not highly sensitive to the allocation across routes, as long as there is some non-zero probability of a checkpoint at all locations. While this counterfactual admittedly pushes the model quite far from its original setting, it does highlight the effectiveness of prior warning about checkpoints in a known (in equilibrium) anti-sobriety campaign. This corresponds quite closely with the strategies of many US police departments, whose announcements often specify not only the beginning of the campaign, but also the exact locations and times of the checkpoints.¹⁸

7 Conclusion

This paper presents the results from a randomized experiment on the enforcement of drunk driving campaign in Rajasthan India, which was designed to test a model of criminal learning and strategic behavior. The central conclusion is that there is clear evidence of learning, hence police interventions focused on the single location with the highest prior concentration of criminal activity are rapidly undone by the diversion of criminal activity to other areas. In contrast, an intervention spread across multiple, initially less promising, locations causes a significant decrease in road accidents and deaths. However, just as drivers learn about the beginning of police enforcement, they also learn that it has come to an end—we see a slow reversion of driver behavior and a return to drink and driving after the intervention.

These results provide the data for a structural estimation of the parameters of a

¹⁸For example, the Connecticut state police announce a checkpoint at the “Intersection of Route 67/Mountain Road in Oxford — 7 p.m. Dec. 30 to 3 a.m. Dec. 31. [2018]” <https://www.ctpost.com/local/article/Connecticut-State-Police-announce-DUI-checkpoints-13486655.php>.

model of learning by potential drunken drivers. The structural parameters confirm many of the qualitative implications of the reduced form data: drivers' priors on the intensity of checking on many (but not all) locations change quickly. They also show the two main factors underlying the main results: the estimated "cost" of encountering a police checkpoint is very high (causing rapid behavioral adjustment), but there are certain locations where citizens do not expect the police to be checking. This combination is central in reproducing the results found in the data.

These parameters provide the basis for the evaluation of a range of counterfactual policies. We find that the optimal crackdown, given the conditions and constraints present in Rajasthan at the time of the project, implies a relatively long-term effort that spreads 20 checkpoints over 82 days. This ensures that drivers are dissuaded for a long time by gradually revising their (initially low) beliefs on duration of the crackdown upwards. This emphasizes the situation-specificity of the optimal crackdown: if driver's initial beliefs were consistent with a longer campaign, it would not be necessary to convince them by conducting such a long intervention. Indeed, when we alter our assumptions to require citizens' beliefs about police activity to be correct in expectation, we find that the optimal police strategy becomes one of a short, intense crackdown mostly done in one location.

Our results suggest several areas where future research could contribute greatly to illuminating the relationship between crime and punishment. There is little evidence on the role of social learning in the effectiveness of police enforcement strategies, although in certain contexts this may be quite important. This would be more challenging to study, perhaps requiring individual level data on (potential) law-breakers. A relatively new strand of research in the criminology literature (Sloan et al., 2014) investigates the insights that behavioral economics might bring to understanding criminal behavior. Insights from the behavioral literature could be incorporated into the design of future randomized policing experiments, and explicitly tested in structural modeling of the outcomes. Finally, the very high predicted efficacy pre-announcing a crackdown raises the possibility that the police may wish to announce crackdowns that are subsequently never actually implemented. Because citizens would come to expect this, the equilibrium effectiveness of such a policy is another important question that we leave to future research.

References

- Andres, G.-L. and A. M. Mobarak (2019). Slippery fish: Enforcing regulation under subversive adaptation. Technical Report 12179, Institute for the Study of Labor (IZA).
- Anwar, S. and T. A. Loughran (2011). Testing a bayesian learning theory of deterrence among serious juvenile offenders. *Criminology* 49(3), 667–698.
- Bai, J. (2018). Melons as lemons: Asymmetric information, consumer learning and quality provision. Mimeo, Harvard Kennedy School.
- Banerjee, A., R. Chattopadhyay, E. Duflo, D. Keniston, and N. Singh (2012, March). Improving police performance in Rajasthan, India: Experimental evidence on incentives, managerial autonomy and training. Working Paper 17912, National Bureau of Economic Research.
- Banerjee, A. V., E. Duflo, and R. Glennerster (2008). Putting a band-aid on a corpse: incentives for nurses in the indian public health care system. *Journal of the European Economic Association* 6(2-3), 487–500.
- Blattman, C., D. Green, D. Ortega, and S. Tobón (2017). Place-based interventions at scale: The direct and spillover effects of policing and city services on crime. Technical Report 23941, National Bureau of Economic Research.
- Bruins, M., J. A. Duffy, M. P. Keane, and A. A. Smith Jr (2018). Generalized indirect inference for discrete choice models. *Journal of Econometrics* 205(1), 177–203.
- Chalfin, A. and J. McCrary (2017). Criminal deterrence: A review of the literature. *Journal of Economic Literature* 55(1), 5–48.
- Ching, A. T., T. Erdem, and M. P. Keane (2013). Learning models: An assessment of progress, challenges, and new developments. *Marketing Science* 32(6), 913–938.
- Ching, A. T., T. Erdem, and M. P. Keane (2017). *Empirical Models of Learning Dynamics: A Survey of Recent Developments*, pp. 223–257. Cham: Springer International Publishing.
- Clarke, R. V. and D. Weisburd (1994). Diffusion of crime control benefits: Observations on the reverse of displacement. *Crime prevention studies* 2, 165–184.

- Davis, A., A. Quimby, W. Odero, G. Gururaj, and M. Hajar (2003). Improving road safety by reducing impaired driving in. *London: DFID Global Road Safety Partnership*.
- Dell, M. (2015). Trafficking networks and the mexican drug war. *The American Economic Review* 105(6), 1738–1779.
- Eeckhout, J., N. Persico, and P. E. Todd (2010). A theory of optimal random crackdowns. *American Economic Review* 100(3), 1104–1135.
- Elder, R. W., R. A. Shults, D. A. Sleet, J. L. Nichols, S. Zaza, and R. S. Thompson (2002). Effectiveness of sobriety checkpoints for reducing alcohol-involved crashes. *Traffic Injury Prevention* 3(4), 266–274.
- Erke, A., C. Goldenbeld, and T. Vaa (2009). The effects of drink-driving checkpoints on crashes: a meta-analysis. *Accident Analysis & Prevention* 41(5), 914–923.
- Foster, A. D. and M. R. Rosenzweig (2010). Microeconomics of technology adoption. *Annu. Rev. Econ.* 2(1), 395–424.
- Giancola, P. R., R. A. Josephs, D. J. Parrott, and A. A. Duke (2010). Alcohol myopia revisited: Clarifying aggression and other acts of disinhibition through a distorted lens. *Perspectives on Psychological Science* 5(3), 265–278. PMID: 26162159.
- Habyarimana, J. and W. Jack (2011). Heckle and chide: Results of a randomized road safety intervention in kenya. *Journal of Public Economics* 95(11), 1438–1446.
- Hörner, J. and A. Skrzypacz (2017). *Learning, Experimentation, and Information Design*, Volume 1 of *Econometric Society Monographs*, pp. 63–98. Cambridge University Press.
- Mookherjee, D. and I. P. Png (1994). Marginal deterrence in enforcement of law. *Journal of Political Economy* 102(5), 1039–1066.
- Newey, W. K. and K. D. West (1987). Hypothesis testing with efficient method of moments estimation. *International Economic Review*, 777–787.
- Peden, M., R. Scurfield, D. Sleet, D. Mohan, A. A. Hyder, E. Jarawan, C. D. Mathers, et al. (2004). World report on road traffic injury prevention.

- Peek-Asa, C. (1999). The effect of random alcohol screening in reducing motor vehicle crash injuries. *American Journal of Preventive Medicine* 16(1), 57–67.
- Ratcliffe, J. H., T. Taniguchi, E. R. Groff, and J. D. Wood (2011). The Philadelphia foot patrol experiment: A randomized controlled trial of police patrol effectiveness in violent crime hotspots. *Criminology* 49(3), 795–831.
- Ross, H. L. (1993). Punishment as a factor in preventing alcohol-related accidents. *Addiction* 88(7), 997–1002.
- Scott, S. L. (2010). A modern Bayesian look at the multi-armed bandit. *Applied Stochastic Models in Business and Industry* 26(6), 639–658.
- Sherman, L. W. (1990). Police crackdowns: Initial and residual deterrence. *Crime and justice* 12, 1–48.
- Sherman, L. W. and D. Weisburd (1995). General deterrent effects of police patrol in crime "hot spots": A randomized, controlled trial. *Justice quarterly* 12(4), 625–648.
- Sloan, F. A., L. M. Eldred, and Y. Xu (2014). The behavioral economics of drunk driving. *Journal of health economics* 35, 64–81.
- Weisburd, D. and L. Green (1995). Policing drug hot spots: The Jersey City drug market analysis experiment. *Justice Quarterly* 12(4), 711–735.
- Weisburd, D. and C. W. Telep (2014). Hot spots policing: What we know and what we need to know. *Journal of Contemporary Criminal Justice* 30(2), 200–220.
- Wilson, T., R. Paternoster, and T. Loughran (2017). Direct and indirect experiential effects in an updating model of deterrence: A research note. *Journal of Research in Crime and Delinquency* 54(1), 63–77.
- World Health Organization (2018). Global status report on road safety 2018. Technical report, Geneva, Geneva.

8 Tables

Table 1: Police station treatment assignment

		Implementation Staff	
A. Sep.-Oct. 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	10 stations @ 1/week
		11 stations @ 2/week	9 stations @ 2/week
		10 stations @ 3/week	12 stations @ 3/week
	Fixed	9 stations @ 1/week	14 stations @ 1/week
		7 stations @ 2/week	13 stations @ 2/week
		9 stations @ 3/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%			

The total vehicle category also includes 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.

Table 3: Pooled Results

	Daylight		Darkness		Day & Night
	(1) Accidents	(2) Deaths	(3) Accidents	(4) Deaths	(5) Deaths
Treatment during & post intervention	0.0032 (0.0037)	-0.0023 (0.0029)	-0.0056** (0.0024)	-0.0040* (0.0021)	-0.0036 (0.0031)
Month FE	Yes	Yes	Yes	Yes	Yes
Police Station FE	Yes	Yes	Yes	Yes	Yes
Mean of dep. variable	0.085	0.029	0.033	0.016	0.045
N	5090	4724	5090	4724	5090

This table presents the impact of all (pooled) sobriety checkpoint interventions on the number of monthly road accidents and deaths from August 2010-October 2012. The during & post intervention variable is positive for the duration of the sobriety crackdown and 90 days afterwards. Accident/death counts have been re-normalized to the per-day level. Each observation corresponds to a police-station month, with months that span the beginning or end of the intervention divided into 2 observations using daily accident/death data and weighted accordingly. All data was taken from police administrative reports, as collected in both treatment and control police stations. Hourly death data was not available for some months, hence columns 2 and 4 have fewer observations. Column 5 displays total deaths (daylight and nighttime) during the entire study period, including for dates when the hour of the accident was not recorded. Standard errors in parentheses clustered at the police station level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Fixed vs. Rotating Pooled Results

	Daylight		Darkness		Day & Night
	(1) Accidents	(2) Deaths	(3) Accidents	(4) Deaths	(5) Deaths
Fixed checkpoints during & post intervention	-0.0021 (0.0045)	-0.0041 (0.0033)	-0.0016 (0.0029)	-0.0029 (0.0023)	-0.0046 (0.0035)
Rotating checkpoints during & post intervention	0.0084* (0.0045)	-0.0005 (0.0036)	-0.0096*** (0.0028)	-0.0050* (0.0029)	-0.0026 (0.0041)
Month fixed effects	Yes	Yes	Yes	Yes	Yes
Police Station FE	Yes	Yes	Yes	Yes	Yes
Mean of dep. variable	0.085	0.029	0.033	0.016	0.045
P-value of test fixed = rotating effect	0.0398	0.378	0.0130	0.473	0.660
N	5090	4724	5090	4724	5090

This table presents the impact of fixed and rotating checkpoint interventions on the number of monthly road accidents and deaths from August 2010-October 2012. The during & post intervention variable is positive for the duration of the sobriety crackdown and 90 days afterwards. Accident/death counts have been re-normalized to the per-day level. Each observation corresponds to a police-station month, with months that span the beginning or end of the intervention divided into 2 observations using daily accident/death data and weighted accordingly. All data was taken from police administrative reports, as collected in both treatment and control police stations. Hourly death data was not available for some months, hence columns 2 and 4 have fewer observations.

Standard errors in parentheses clustered at the police station level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

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

	Daylight		Darkness		Day & Night
	(1) Accidents	(2) Deaths	(3) Accidents	(4) Deaths	(5) Deaths
Fixed stations during intervention	-0.0014 (0.0064)	-0.0011 (0.0041)	0.0021 (0.0045)	-0.0021 (0.0027)	-0.0024 (0.0047)
Fixed stations post intervention	-0.0025 (0.0049)	-0.0073* (0.0042)	-0.0043 (0.0034)	-0.0039 (0.0030)	-0.0062 (0.0044)
Rotating stations during intervention	0.0112* (0.0065)	0.0029 (0.0048)	-0.0099** (0.0041)	-0.0052 (0.0034)	-0.0020 (0.0056)
Rotating stations post intervention	0.0065 (0.0055)	-0.0044 (0.0041)	-0.0092** (0.0035)	-0.0048 (0.0039)	-0.0029 (0.0048)
Month fixed effects	Yes	Yes	Yes	Yes	Yes
Police Station FE	Yes	Yes	Yes	Yes	Yes
Mean of dep. variable	0.085	0.029	0.033	0.016	0.045
N	5090	4724	5090	4724	5090

This table presents the impact of fixed and rotating checkpoint interventions on the number of monthly road accidents and deaths from August 2010-October 2012. The during intervention variable is positive for the duration of the sobriety crackdown, and the post intervention variable is positive for 90 days afterwards. Accident/death counts have been re-normalized to the per-day level. Each observation corresponds to a police-station month, with months that span the beginning or end of the intervention divided into 2 observations using daily accident/death data and weighted accordingly. All data was taken from police administrative reports, as collected in both treatment and control police stations. Hourly death data was not available for some months, hence columns 2 and 4 have fewer observations.

Standard errors in parentheses clustered at the police station level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Fixed vs. Rotating, Intensity of Checking

	Daylight		Darkness		Day & Night
	(1) Accidents	(2) Deaths	(3) Accidents	(4) Deaths	(5) Deaths
Rotating checkpoints×1/week during & post intervention	0.0090 (0.0083)	0.0007 (0.0077)	-0.0074 (0.0053)	0.0043 (0.0054)	0.0025 (0.0082)
Rotating checkpoints×2/week during & post intervention	0.0059 (0.0064)	-0.0017 (0.0037)	-0.0094** (0.0038)	-0.0093** (0.0040)	-0.0073 (0.0050)
Rotating checkpoints×3/week during & post intervention	0.0106* (0.0062)	0.0002 (0.0063)	-0.0116*** (0.0041)	-0.0061* (0.0034)	-0.0006 (0.0064)
Fixed checkpoints ×1/week during & post intervention	-0.0073 (0.0063)	-0.0023 (0.0054)	-0.0027 (0.0038)	-0.0026 (0.0031)	-0.0035 (0.0049)
Fixed checkpoints×2/week during & post intervention	0.0020 (0.0066)	-0.0031 (0.0042)	-0.0013 (0.0043)	-0.0019 (0.0036)	-0.0016 (0.0053)
Fixed checkpoints×3/week during & post intervention	-0.0022 (0.0063)	-0.0072 (0.0048)	-0.0007 (0.0044)	-0.0035 (0.0032)	-0.0095* (0.0051)
Month FE	Yes	Yes	Yes	Yes	Yes
Police Station FE	Yes	Yes	Yes	Yes	Yes
Mean of dep. variable	0.085	0.029	0.033	0.016	0.045
N	5090	4724	5090	4724	5090

This table presents the impact of fixed and rotating checkpoint interventions on the number of monthly road accidents and deaths from August 2010-October 2012. The during intervention variable is positive for the duration of the sobriety crackdown, and the post intervention variable is positive for 90 days afterwards. Accident/death counts have been re-normalized to the per-day level. Each observation corresponds to a police-station month, with months that span the beginning or end of the intervention divided into 2 observations using daily accident/death data and weighted accordingly. All data was taken from police administrative reports, as collected in both treatment and control police stations. Hourly death data was not available for some months, hence columns 2 and 4 have fewer observations.

Standard errors in parentheses clustered at the police station level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7: Checkpoint surveys during intervention

	Drunk drivers and motorcyclists caught				Cars and motorcycles passing			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment					-150.267*	-65.989		
					(86.260)	(130.658)		
Rotating checkpoint station	0.130 (0.189)	-1.001* (0.518)			0.208 (74.392)	-32.860 (83.159)		
Frequency		-0.519** (0.199)				-42.536 (43.020)		
Rotating checkpoint × frequency		0.586** (0.235)				25.996 (18.340)		
Weeks of checking			-0.099*** (0.026)				-17.528*** (4.201)	
Rotating checkpoint × weeks of checking			0.091** (0.035)				1.506 (5.375)	
Number previous checkpoints				-0.041*** (0.012)				-8.507*** (2.182)
Rotating checkpoint × number previous checkpoints				0.032* (0.016)				2.024 (2.604)
District FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Police Station FE	No	No	Yes	Yes	No	No	Yes	Yes
Mean of dep. variable	1.257	1.257	1.257	1.257	527.0	527.0	527.0	527.0
N	852	852	852	852	2645	2645	2645	2645

This table reports the intensity and dynamic effects of the sobriety crackdown on the number of drunken drivers caught (columns 1-4) and vehicles passing the checkpoint location (columns 5-8). All outcome variables are based on data collected by surveyors sent to monitor the checkpoints. The frequency of checking variable is the number of checkpoints per week: 1, 2, or 3. The weeks of checking variable is the number of weeks that have elapsed since the first checkpoint. The number of previous checkpoints is the number of checkpoints assigned prior to the given night, after the start of the intervention. 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. Standard errors in parentheses clustered at the police station level. * $p < 0.10$. ** $p < 0.05$. *** $p < 0.01$

Table 8: Passing cars and motorcycles post-crackdown

	Road 1		Roads 2 & 3	
	(1)	(2)	(3)	(4)
Treatment	-214.1** (71.45)		-13.87 (43.41)	
Days since last checkpoint	1.761 (4.003)		-3.774 (3.216)	
Fixed checkpoints		-224.4** (79.33)		191.5 (128.0)
Fixed checkpoint ×days since last checkpoint		0.155 (3.556)		-11.77** (5.261)
Rotating checkpoints		-289.0** (115.6)		-107.0** (40.37)
Rotating checkpoint ×days since last checkpoint		17.13 (15.91)		1.597 (4.636)
District FE	Yes	Yes	Yes	Yes
Control mean	488.0	488.0	501.5	501.5
N	145	145	347	347

All columns report outcomes on the number of cars and motorcycles observed passing the former checkpoint locations in the weeks after the sobriety checkpoint intervention was over. Data collected by surveyors monitoring the locations. All specifications include controls for whether the police station is located on a major highway, the pre-intervention accident rate, and whether the surveyor was counting 1-way or 2-way traffic. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 9: Drunk drivers caught on final check

	All Stations		Rotating checkpoints		Fixed Checkpoints	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	-1.322** (0.447)	-0.286 (0.663)	-1.250 (0.799)	-0.447 (1.269)	-1.301** (0.574)	0.560 (1.463)
Days since last checkpoint	0.017* (0.008)	-0.012 (0.026)	0.001 (0.018)	-0.057 (0.033)	0.033*** (0.009)	-0.007 (0.039)
Frequency		-0.503 (0.280)		-0.340 (0.368)		-0.954 (0.541)
Days since last last checkpoint \times frequency		0.015 (0.012)		0.027* (0.012)		0.021 (0.020)
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Control mean	1.953	1.953	1.953	1.953	1.953	1.953
Mean treatment effect, @ freq = 2	-1.089	-1.064	-1.236	-1.177	-0.856	-0.875
P-value of mean treatment effect	0.0114	0.0105	0.0594	0.0778	0.118	0.102
N	108	108	77	77	74	74

This table reports the impact of the interventions on the number of drunken car and motorcycle drivers caught at the final check conducted after the end of the intervention in all police stations, including control stations. Columns 1 and 2 compare pooled treatment police stations with control stations, columns 3 and 4 compare rotating checkpoint stations with controls, and columns 5 and 6 compare fixed checkpoint police stations with controls. Outcome data collected by surveyors sent to monitor the final checkpoints. All specifications include controls for whether the police station is located on a major highway, the pre-intervention accident rate, and whether the final check was conducted by a team from the local station or district Police Lines.

Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 10: Structural Parameters

Interpretation	Parameters	Baseline Model (1)	No Learning (2)	Myopic (3)	Riskless Learning (4)
Utility of drunken driving on road 1	d_1	-1.541 (0.766)	-1.355 (0.219)	-1.379 (0.749)	42.787 (48.607)
Utility of drunken driving on road 2	d_2	-1.809 (0.873)	-1.436 (0.212)	-1.626 (0.786)	42.667 (49.575)
Utility of drunken driving on road 3	d_3	-1.780 (0.734)	-1.415 (0.215)	-1.656 (0.468)	42.838 (48.485)
Disutility of encountering checkpoint	c	5,942.684 (20,111.04)	7,771.940 (28,512.67)	5,594.131 (20,5972.21)	47,567.846 (1.003 * 10 ⁶)
Prior on crackdown ending probability	$\eta_0 = \frac{\alpha_0^\eta}{\alpha_0^\eta + \beta_0^\eta}$	0.315 (0.643)	0.148 (0.033)	0.287 (0.681)	0.999 (0.087)
Strength of prior on crackdown ending probability	$\alpha_0^\eta + \beta_0^\eta$	20.670 (101.230)	-	26.264 (165.883)	70.539 (805.219)
Prior on road 1 intensity	$\lambda_{1,0} = \frac{\alpha_{1,0}^\lambda}{\alpha_{1,0}^\lambda + \beta_{1,0}^\lambda}$	0.090 (4.603)	0.226 (0.756)	0.094 (2.098)	0.999 (0.351)
Strength of prior on road 1 intensity	$\alpha_{1,0}^\lambda + \beta_{1,0}^\lambda$	6.050*10 ⁻⁷ (0.685)	-	3.654 * 10 ⁻⁷ (0.455)	258.516 (3388.131)
Prior on road 2 intensity	$\lambda_{2,0} = \frac{\alpha_{2,0}^\lambda}{\alpha_{2,0}^\lambda + \beta_{2,0}^\lambda}$	0.006 (0.030)	0.031 (0.098)	0.007 (0.124)	0.999 (0.757)
Strength of prior on road 2 intensity	$\alpha_{2,0}^\lambda + \beta_{2,0}^\lambda$	8,807.176 (2.555 * 10 ⁶)	-	11,992.953 (5.926 * 10 ⁶)	157.103 (1330.756)
Prior on road 3 intensity	$\lambda_{3,0} = \frac{\alpha_{3,0}^\lambda}{\alpha_{3,0}^\lambda + \beta_{3,0}^\lambda}$	3.060*10 ⁻⁶ (0.008)	0.022 (0.069)	2.035 * 10 ⁻⁷ (0.007)	0.999 (0.345)
Strength of prior on road 3 intensity	$\alpha_{3,0}^\lambda + \beta_{3,0}^\lambda$	0.009 (12.360)	-	0.004 (5.102)	1405.658 (11014.146)
Value of GMM Criterion		.0087	.0175	.0090	0.0097

Structural parameters estimated using techniques and specifications described in Section 6. Standard errors in parentheses calculated using delta method.

Table 11: Optimal Crackdown Strategy and Duration


		Implementation Compliance					
		Full			Partial		
Prior Warning		Duration	Pct. on road 1	Decrease in drunken driving	Duration	Pct. on road 1	Decrease in drunken driving
		(1)	(2)	(3)	(4)	(5)	(6)
Estimated beliefs	No	82 days	35%	36.1%	83 days	35%	25.0%
	Yes	77 days	20%	46.6%	81 days	35%	36.5%
Equilibrium beliefs	No	21 days	75%	89.9%	43 days	95%	71.7%
	Yes	>31 days	35%	100%	>31 days	35%	100%

This table displays the parameters and effectiveness of the optimal anti-drunken driving campaign that allocates 20 checkpoints over up to 90 days and across 3 locations. Effectiveness is measured by the undiscounted share of agents drinking and driving on any road over the 90 days after the first checkpoint. Parameters used in rows 1 and 2 (“Estimated beliefs”) are all as estimated in the baseline model. Parameters used in rows 3 and 4 are as estimated from the data for drivers’ utilities, but with beliefs constrained to equal police strategies in expectation; see details in Section 6.4. Counterfactual strategies estimated in rows 1 and 3 assume (as was the case in the empirical intervention) that drivers have no prior warning of the beginning of drunken driving enforcement. Rows 2 and 4 assume that all drivers are informed of the beginning of the campaign. Columns 1 and 4 show the duration of the optimal campaign, and columns 2 and 5 show the share of checks allocated to road 1. The remaining checks are equally divided between roads 2 and 3. Columns 3 and 6 show the decrease in drunken driving induced by the campaign relative to the share of drivers drinking and driving on the night just prior to the first checkpoint.

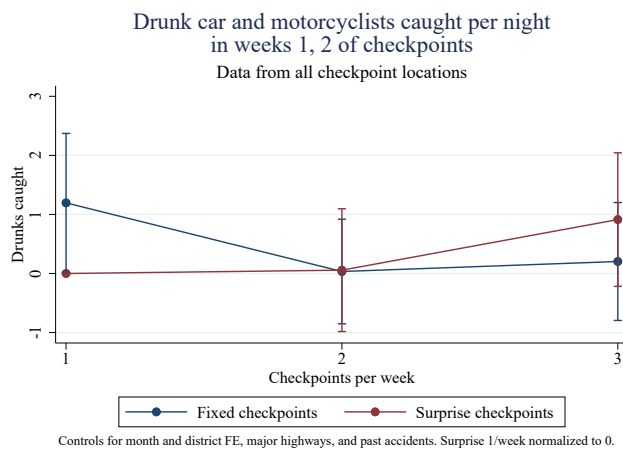
9 Figures

Figure 9.1: Drunken drivers caught over time

(a) Whole intervention learning results



(b) Pre-learning results



(c) Post-learning results

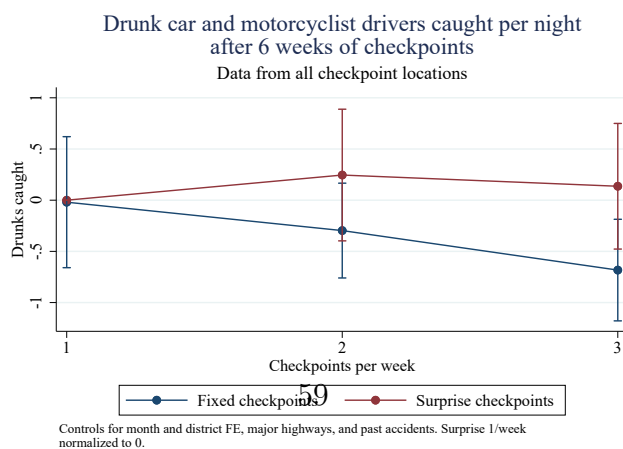


Figure 9.2: Example of Agents' Simulated Actions and Beliefs

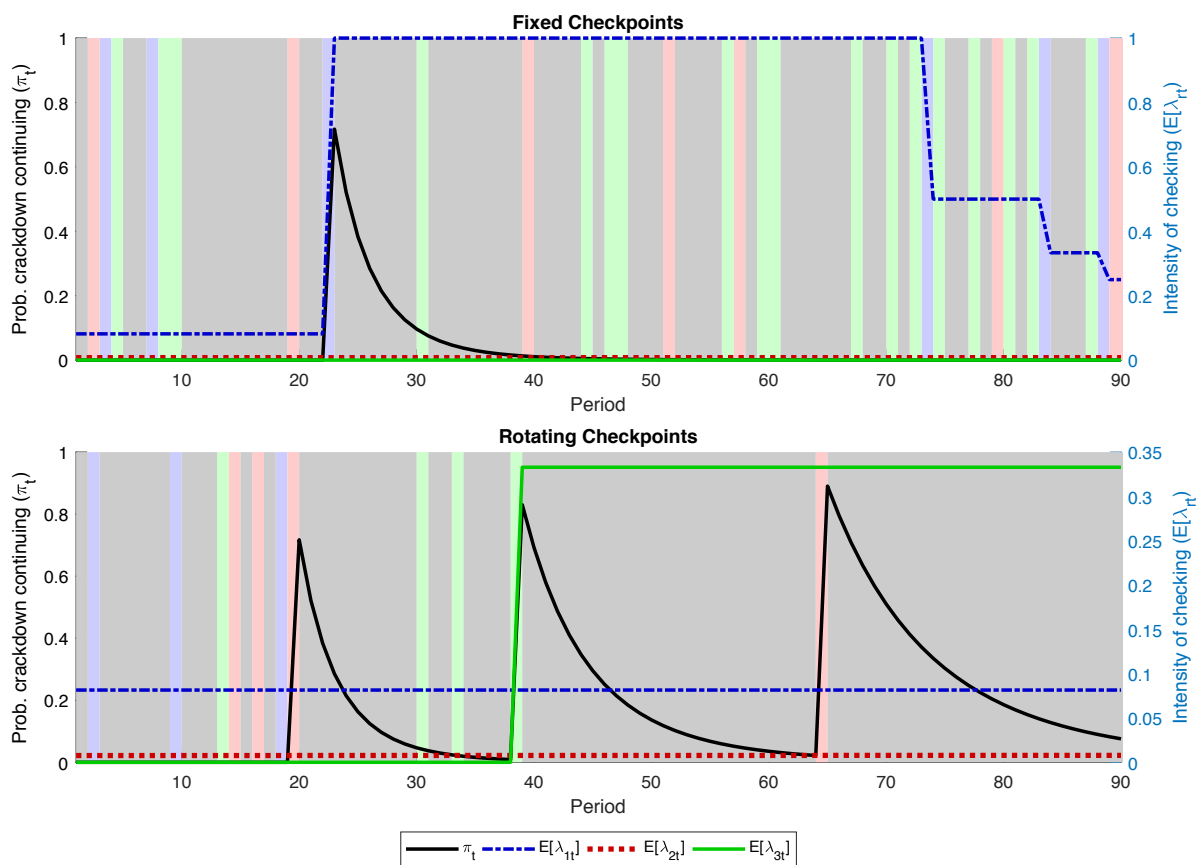
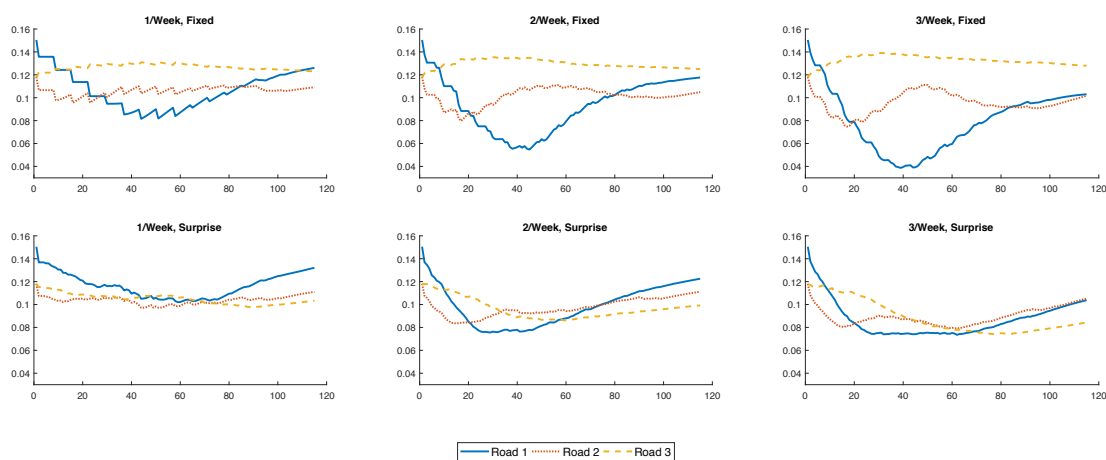


Figure 9.3: Simulated Drunken Driving Probabilities Across Interventions



Appendix

A1 Theory

A1.1 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 would be such an intervention and even after it started there was no official announcement. So potential law-breakers could only learn about it from their own experience or that of others. Furthermore, based on their experience with enforcement of other laws, drivers are likely to believe that anti-drunken driving 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 6.1. 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

$$d + \lambda c < s \tag{A1.1}$$

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.¹⁹ However, all agents believe that with the same prior 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, while with probability p_0 they believe that the police have adopted a strategy of permanently enforcing drunken driving laws.²⁰ 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 A1.1, then there is clearly a 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 A1.1.

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 make use of the fact that because A1.1 holds, after first learning of the enforcement campaign agents immediately commence a single temporary interval of sobriety before returning to drunken driving. The following proposition shows conditions under which drivers attempt a single (rather than multiple)

¹⁹In particular they have no priors about when the checks started and therefore do not update on the probability of their being permanent based on when they first observe a check.

²⁰While 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 6.1 relaxes this assumption and allows drivers to learn about the permanence of the campaign from a continuous distribution of expected crackdown durations.

period of temporary sobriety before giving up on drunken driving.

Proposition. *There exists a value p^* such that for $p_0 \geq p^*$ the agent's optimal strategy is to cease drunken driving for a finite interval following their first encounter with police, and stop permanently following the second encounter with police.*

Proof. Let \bar{p} be 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 giving up drunken driving forever, or an alternative of remaining sober for an additional a finite period τ , then stopping permanently if caught again. Since the driver's belief that checking is permanent, p_t ,

$$p_t = \frac{p_0}{p_0 + (1 - p_0) e^{-\eta t}} \quad (\text{A1.2})$$

is increasing in t , any driver stopped at a time such that $p_t > \bar{p}$ will cease drinking and driving forever. \bar{p} is implicitly defined by:

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

Substituting the optimal stopping time τ and solving for \bar{p} yields:

$$\bar{p} = \frac{d - s}{d - W_{perm}} \quad (\text{A1.3})$$

We now show conditions which ensure that, for p_0 high enough, the optimal temporary stopping time τ after the first check is such that $p_\tau > \bar{p}$. Combining the expressions for \bar{p} and p_t yields an implicit expression for \bar{t} , the duration of time after the agent's first check when \bar{p} is reached:

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

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 of sobriety, if he returns and is checked then permanently not drinking and driving will be preferable to another temporary interval of sobriety.

This is true if,

$$\tau \geq \bar{t}$$

$$p_0 \frac{(\eta + r)}{r} \left(\frac{s - W_{perm}}{d - s} \right) \geq 1 - p_0 \left(\frac{d - W_{perm}}{d - W_{temp}} \right)$$

which is bound to hold for p_0 high enough. \square

Consider a driver at time zero, the moment after the first encounter with a checkpoint, and hence immediately after he has first learned that a crackdown is in force. Let W_{perm} be the net expected present value of future utility discounted to time zero, in the state of the world where checking is permanent and W_{temp} is the corresponding number in the state of the world where checking is temporary. His valuation of not drinking and driving for τ periods before resuming is:

$$V(\tau; p_0) = \left(\frac{1 - e^{-r\tau}}{r} \right) s + e^{-r\tau} \left(p_0 [W_{perm}] + (1 - p_0) \left[(1 - e^{-\eta\tau}) \frac{d}{r} + e^{-\eta\tau} W_{temp} \right] \right) \quad (\text{A1.4})$$

if his subjective discount factor of r . This valuation incorporates the three possible states of the world upon his return to drunken driving. With probability p_0 the campaign is 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 A1.4 for the utility-maximizing sobriety interval τ yields

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

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 A1.1 holds and

$$(1 - p_0) \frac{d}{r} + p_0 W_{perm} > \frac{s}{r}$$

which just says, that being sober for ever does worse than always drinking and driving without ever caught as long as checking is temporary (the best case scenario) and getting the average expected value of the chosen strategy if it is permanent. If this condition fails, they should just stop permanently. 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.

A1.2 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 Q 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:²¹

$$Q = \int_0^\infty e^{-\eta t} \eta \left(\int_0^t e^{-\rho x} \lambda dx \right) dt = \frac{\lambda}{\rho + \eta} \quad (\text{A1.5})$$

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 .

²¹In our later counterfactuals we impose that the budget constraint must also be satisfied ex-post.

The police objective function can then be written as the sum of two terms,

$$\begin{aligned} \min_{\eta, \lambda} D &= \int_0^\infty e^{-\rho x} \left(\int_0^x \left(e^{-(\eta+\lambda)z} \eta \right) dz + e^{-(\eta+\lambda)x} + e^{-(\eta+\lambda)x} \lambda \Delta \right) dx \quad (\text{A1.6}) \\ &= \frac{\eta + \rho}{\rho(\eta + \lambda + \rho)} + \frac{\lambda \Delta}{(\eta + \lambda + \rho)} \end{aligned}$$

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 $\partial\tau^*/\partial\lambda > 0$, $\frac{\partial D_{perm}}{\partial\lambda} < 0$, the optimal policy is to simply set $\lambda = Q\rho$ as long as equation A1.1 is satisfied (otherwise just set $\lambda = 0$, since the drivers just ignore checking).

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,

$$\begin{aligned} \Delta_{temp} &= (1 - e^{-\eta\tau}) \left[\frac{e^{-\rho\tau}}{\rho} \right] + e^{-\eta\tau} e^{-\rho\tau} \int_0^\infty e^{-(\lambda+\eta+\rho)t} \left(1 + \eta \frac{1}{\rho} + \lambda 0 \right) dt \\ &= e^{-\rho\tau} \left((1 - e^{-\eta\tau}) \frac{1}{\rho} + e^{-\eta\tau} \left(\frac{\eta + \rho}{\rho(\eta + \lambda + \rho)} \right) \right) \quad (\text{A1.7}) \end{aligned}$$

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 reduce the drunk driving. 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. 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 ($\rho > 0$). On the other hand, a higher τ makes it less likely that drivers recommence drinking before checking has ended; this has the cost that they are less likely to be caught a second time, and therefore end drunk driving permanently. It can be shown however that on net an unconstrained police department would want to make checking as intense as possible.

The calculus can change dramatically when budget constraint is binding. Higher intensity crackdowns now necessitate shorter (expected) duration $1/\eta$, and in all our numerical examples,²² induces *shorter* τ intervals from the agents. Because of the countervailing effects of λ and η increasing λ while respecting the budget constraint has no effect on drunken driving by uninformed drivers (the first term in equation A1.6). The expression for the amount of drunk driving collapses to:

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

and some algebraic manipulations show that $\partial D/\partial\lambda > 0$, except if $\partial\tau^*/\partial\lambda \gg 0$, which, as we mentioned, is never the case in all our simulations. Thus our previous result for an unconstrained police department is reversed: subject to the budget constraint in A1.5 the optimal police strategy is to have a prolonged crackdown of the minimal intensity necessary to satisfy A1.1 and thus dissuade drunken drivers.

A1.3 Extensions

We consider several extensions that move the simple theory developed above closer to the structural model estimated in Section 6.1. First, suppose that drivers are unsure about the intensity of the crackdown. At night t 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

²²We have not however been able to show this analytically.

home. This is likely 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). A higher λ intensity will cause, on average, lower p_t and, based on the finding $\partial\tau/\partial p_0 > 0$ from the simpler model, 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 1 denote the best road for drunken driving (perhaps it is near the local bar), and road 2 designate an alternate less convenient choice, with respective utilities $d_1 > d_2$. If drivers are unsure whether the police know their avoidance routes they may attempt to drink and drive on the less preferred road temporarily 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 (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 on the duration/permanence of the crackdown are the same on each road, and that parameters are such that drunken drivers, once caught on road A , will cease offending for a time interval of length $\bar{\tau}$, 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

$$\begin{aligned} D &= \int_0^\infty e^{-(\lambda_A + \rho)t} dt + e^{-\rho\bar{\tau}} \int_0^\infty e^{-\rho t} \left(\int_0^t (e^{-\lambda_A x} \lambda_A) e^{-\lambda_B(t-x)} dx \right) dt \quad (\text{A1.8}) \\ &= \frac{1}{\lambda_A + \rho} + e^{-\rho\bar{\tau}} \frac{\lambda_A}{(\rho + \lambda_B)(\lambda_A + \rho)}. \end{aligned}$$

This expression once again divides the total number of law-breakers into two groups: the first term in equation A1.8, which represents 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 is 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 mean less enforcement on A and therefore a higher fraction of drunken drivers initially on A will 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 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. Incorporating either of these would substantially complicate the model; social learning in a similar environment is itself the subject of an active theoretical literature recently summarized in Hörner and Skrzypacz (2017). In reduced form, risk-free exploration would speed up the impact of the crackdown on reducing drunken driving, and reduce its persistence after the checkpoints have ended. While a similar generalization is not possible in the case of social learning without further assumptions, in the extreme case in which all potential law-breakers learn the exact beginning and end of the crackdown from the experience of others, 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. It is reasonable to assume that there is some amount of social learning in the world, but as we will see, there is a good deal of persistence in driver behavior suggesting that learning is slow and social learning is imperfect at best.

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, as long as drivers are strategic and learning is fast enough.
 - (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.

A2 Further Reduced Form Analyses: Robustness & Alternative Mechanisms

Between-Station Spillovers

While our focus has been on the within-station main effects and spillovers of the sobriety checkpoints, it is possible there may also have been cross-station spillovers. In principle these might be positive or negative: drivers avoiding checkpoints near their own stations might move into different police jurisdictions, or citizens might hear that a crackdown has begun in a nearby town and infer that it has also begun in their own area. On the other hand, the isolated nature of many police stations in rural Rajasthan, the fact that they cover a very large catchment area, and the history of relatively independent enforcement would both serve to attenuate any spillovers.

We quantify the geographical extent of the spillovers in Table A3, which examines accidents and deaths, and Table A4 which examines the number of drunken drivers caught at the checkpoints. In short, neither table shows robust evidence of between

station spillovers. Panel A of Table A3 displays the pooled effects of any intervention on accidents and deaths, analogous to Table 3, now including controls for the number of stations with any checkpoints in 10, 20, and 40 kilometer radii. None are significant, and size of the main effects are very similar to those in Table 3. Panel B disaggregates the main effects and spillovers by fixed and rotating police stations, analogous to Table 4. Of the 30 spillover effects estimated, only 1 is significant at the 5% level and it relates to the outcome of daytime accidents, suggesting it may be the result of chance.

Results from the checkpoint surveys in Table A4 are equally insignificant once we control for police station fixed effects. Columns 1 and 2, which do not include station effects (only district fixed effects), show that police stations with many treated close neighbors (stations in 10 km) have more drunken drivers caught, whereas stations with many treated areas far away (20 and 40 kms) have fewer. However, as in Blattman et al. (2017), the distribution of treated neighbors is not random, so this likely reflects the fact that areas with dense police stations are inherently different (more densely populated, higher crime, etc). Controlling for PS fixed effects in columns 3 and 4 removes this bias, and accordingly the coefficients on the spillover terms in columns 3 and 4 are smaller and insignificant. Similarly, column 5 shows that proximity to other treated police stations has no effect on the number of drunken drivers caught in the final check. Thus we find very limited evidence that inter-station spillovers are important in the prosecution of drunken drivers, confirming our focus on the intra-station dynamics of criminal behavior.

Social Learning and Sober Exploration

Our primary model focuses on driver learning from personal experience encountering police checkpoints while engaged in risky drunken driving. However, it is possible that individuals might also learn about police activity while driving sober or through social learning from friends who have encountered the police checkpoints. While these mechanisms are difficult to identify solely with data on aggregate accidents and arrests, we can gain some insights from the immediate response of road accidents to police checkpoints. Consider a model in which agents engage in drunken driving only occasionally, say with 10% probability, and thus only 10% of individuals personally witness a checkpoint. If drivers learn only through personal experience, the immediate impact of the check point will be muted, since only 1% ($= .1 \times .1$) of the informed agents would have engaged in drunken driving the following day anyway. Any effects

of the crackdown will appear gradually and smoothly. In contrast, if there is costless exploration or social learning, one might expect a large immediate effect of a checkpoint on drunken driving and hence on road accidents the following night, since a much larger fraction of the potential drunken driving population will be informed. Thus the empirical drop in road accidents on the night immediately following a checkpoint may be informative about other mechanisms of individual learning.²³

Table A5 presents regression specifications examining the immediate impact of checkpoints on night accidents and deaths in the following days (unlike our main results, these regressions are estimated using daily data). Columns 1 and 2 focus on the impact after exactly 1 night, with separate coefficients for the fixed and rotating stations. The results show no significant 1-night post checkpoint effect on accidents or deaths, for either intervention. The coefficient on main effect of the surprise treatment remains negative, statistically significant, and of a similar magnitude as in the main specification. However, this specification may be overly restrictive if there is social learning but the news of the intervention takes some time to disseminate. Therefore in columns 3 and 4 we include 2 and 3 day lags of the checkpoint indicator. Once again, we find no significant immediate effects of the checkpoints. Thus while we cannot fully rule out forms of learning outside the model, we find no evidence of their implications in the data.

Implementation of the Crackdown Intervention

As we anticipated when designing the Police Lines intervention, implementation of the assigned checkpoints was far from perfect. Of the 1,565 checkpoints that the surveyors visited, only 980 (63%) were actually implemented by the police. This partial compliance would substantially modify the interpretation of our results if it were correlated with the treatment assignment in the intervention, for instance if fixed checkpoints were less likely to be carried out. In Table A2 we examine the impact of the treatments on police implementation of the intervention. Columns 1-3 focus on whether the checkpoint occurred (using OLS, Probit, and focusing on the final checks only), while columns 4 and 5 examine the duration of the checkpoint and the number of vehicles stopped, respectively. The most striking finding is that the Police

²³In principle, the number of arrests on the following night would also be informative when police implemented another checkpoint the next night. However, since consecutive checkpoints may have affected arrest numbers through multiple channels (for example, police effort and morale) we focus on the accident data.

Lines teams substantially outperformed the station-based teams on all outcomes—for example column 1 shows they were 44% more likely to carry out a checkpoint. We also find that the police perform worse in stations with 3 checkpoints assigned per week, perhaps due to fatigue or burn-out. This affects the cardinal interpretation of our reduced form results: areas with 3 checkpoints per week did not have exactly 3 times greater intensity than those with 1 checkpoint per week. However, realized enforcement intensity is still increasing in assigned intensity. Reassuringly, we find no evidence that implementation is correlated with treatment assignment to fixed or rotating checkpoint locations. There is a fairly large (though insignificant) negative coefficient on the variable indicating that the checkpoint was held off the main road, which is likely due to the fact that some of the routes 2 and 3 had low volumes of traffic. In conclusion, the partial implementation of the intervention does not change the main qualitative conclusions of the reduced form results. It does, however, require that we use instrumental variables in the structural estimation, an approach we that we outline below.

Legal or Informal Sanctions for Drunken Drivers

After stopping a drunken driver, police may fail to follow the legal ticketing procedure, either because they choose to release the driver with a warning or because they prefer to take a bribe rather than register an official ticket. If this behavior were uniformly prevalent in the Rajasthan Police, it might weaken the strength of the intervention and reduce the effects. In fact, as our results show, drivers seemed very concerned about the consequences of being stopped while drunken driving, so punishments seem to have been quite strong whether legal or extra-judicial. A more serious issue is the potential correlation between the strength or legality of the punishment and the intervention. We examined this issue by collected the court records on drunken driving cases legally prosecuted over the course of the intervention, since all drunken drivers were required to report to court to pay their fine. These documents contained the name of the police station issuing the ticket, so we were able to generate a measure of the number of legal tickets issued by each police station over the course of the intervention. Table A6 contains the regression of this measure on dummy variables for various intervention categories. In column 1 we see that regressing legal cases on an indicator variable for rotating checkpoint police stations yields a large, though insignificant, coefficient. This is not surprising since this intervention was

relatively more effective. In columns 2 and 3 we control for the number of drunken drivers that surveyors observed being caught in the police station (a noisy measure of the total) and indicator variables for the other interventions. Introducing these controls substantially reduces the size and significance of the rotating checkpoints variable, as expected. Thus while we cannot observe the exact share of offenders legally prosecuted, we find no evidence that the legal treatment of drunken drivers is correlated with the rotating or fixed checkpoint intervention.

A2.1 Driver beliefs under costless social learning or riskless exploration

As a robustness check to our baseline model of independent individual learning, we estimate an alternate model in which we assume that all drivers are informed after a checkpoint has occurred. Driver decisions, conditional on beliefs, are therefore not forward looking because the driver’s evolution of beliefs in the next period does not depend on his contemporaneous decision to drink and drive. This implies that, within a police station, all agents have the same beliefs about the intensity and duration of the police crackdown. The set of possible outcomes is reduced to four: no checkpoint occurs, or a checkpoint on road $r \in \{1, 2, 3\}$. Let \mathbb{I}_t be an indicator variable equal to one if any checkpoints have been held on any roads prior to t , that is if the crackdown has begun. Evolution of beliefs are then,

No checkpoint:

$$\begin{aligned}\pi_{t+1} &= (1 - \eta_t) \pi_t \\ \tilde{\beta}_{t+1}^\eta &= \tilde{\beta}_t^\eta + \mathbb{I}_t \\ \lambda_{r,t+1} &= \frac{\tilde{\alpha}_{rt}^\lambda}{\tilde{\alpha}_{rt}^\lambda + \tilde{\beta}_{rt}^\lambda + \mathbb{I}_t} \\ \tilde{\beta}_{r,t+1}^\lambda &= \tilde{\beta}_{r,t}^\lambda + \mathbb{I}_t\end{aligned}$$

Checkpoint on r :

$$\begin{aligned}\pi_{t+1} &= (1 - \eta_t) \\ \tilde{\beta}_{t+1}^\eta &= \tilde{\beta}_t^\eta + 1 \\ \lambda_{r,t+1} &= \frac{\tilde{\alpha}_{rt}^\lambda + 1}{\tilde{\alpha}_{rt}^\lambda + \tilde{\beta}_{rt}^\lambda + 1} \\ \tilde{\beta}_{r,t+1}^\lambda &= \tilde{\beta}_{r,t}^\lambda + 1 \\ \lambda_{r',t+1} &= \frac{\tilde{\alpha}_{r't}^\lambda}{\tilde{\alpha}_{r't}^\lambda + \tilde{\beta}_{r't}^\lambda + 1} \\ \tilde{\beta}_{r',t+1}^\lambda &= \tilde{\beta}_{r',t}^\lambda + 1\end{aligned}$$

A3 Appendix Tables

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)
Police Station FE	No	No	Yes	Yes
District FE	Yes	Yes	Yes	Yes
Mean of dep. variable	0.806	0.806	0.806	0.806
Observations	1352	1352	1352	1352

This table reports the intensity and dynamic effects of the sobriety crackdown on the number of drunken drivers caught, with the number of drunken drivers caught set to 0 if the police did not implement the checkpoint. Results corresponds with columns 1-4 Table 7, except that in Table 7 checkpoints that were not implemented are dropped from the data.

All outcome variables are based on data collected by surveyors sent to monitor the checkpoints. The frequency of checking variable is the number of checkpoints per week: 1, 2, or 3. The weeks of checking variable is the number of weeks that have elapsed since the first checkpoint. The number of previous checkpoints is the number of checkpoints assigned prior to the given night, after the start of the intervention. 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.

Standard errors in parentheses clustered at the police station level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A2: Implementation of intervention

	(1) Checkpoint occurred - OLS	(2) Checkpoint occurred - Probit	(3) Checkpoint occurred - Final Check OLS	(4) Duration of checkpoint	(5) Number vehicles stopped
Fixed checkpoints			0.006 (0.106)		
Rotating checkpoints	-0.001 (0.046)	0.000 (0.153)	0.06 (0.067)	-1.856 (5.443)	-0.235 (7.893)
Intensity 2/week	0.021 (0.049)	0.047 (0.154)	-0.009 (0.065)	-4.827 (4.982)	-16.280* (9.455)
Intensity 3/week	-0.105* (0.053)	-0.328** (0.163)	-0.057 (0.080)	-11.701** (5.041)	-15.624 (9.786)
Locations 2 or 3	-0.024 (0.038)	-0.105 (0.132)		-0.326 (4.779)	-2.786 (6.090)
Police lines teams	0.284*** (0.035)	0.921*** (0.117)	0.337** (0.120)	23.653*** (4.439)	25.390*** (6.886)
District FE	Yes	Yes	Yes	Yes	Yes
Mean of dep. variable	0.641	0.641	0.599	159.293	62.537
N	1353	1353	182	867	867

The outcome variable in columns 1-3 is an indicator equal to 1 if the assigned checkpoint was implemented by the police during the main intervention. The outcome variable in column 4 is the analogous indicator for the final checks. The outcome variable in column 5 is the duration of the checkpoint in minutes (from the time the police arrived, to when they stopped checking), and the outcome in column 6 is the number of cars and motorcycles checked. Columns 5 and 6 use only data conditional on the checkpoint being conducted. All data on outcomes were collected by surveyors monitoring the police checkpoints. All specifications include controls for whether the police station is located on a major highway, the pre-intervention accident rate.

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

Table A3: Accidents and Deaths - Spillovers

	Daylight		Darkness		Day & Night
	(1) Accidents	(2) Deaths	(3) Accidents	(4) Deaths	(5) Deaths
Panel A. Spillovers main effects					
Treated station during & post intervention	0.00331 (0.00372)	-0.00216 (0.00294)	-0.00567** (0.00235)	-0.00394* (0.00213)	-0.00347 (0.00317)
Treated stations in 10 km during & post intervention	0.00091 (0.00116)	0.00079 (0.00094)	0.00055 (0.00100)	0.00004 (0.00069)	0.00077 (0.00099)
Treated stations in 20 km during & post intervention	-0.00017 (0.00197)	-0.00035 (0.00151)	-0.00122 (0.00169)	-0.00106 (0.00104)	-0.00122 (0.00167)
Treated stations in 40 km during & post intervention	-0.00056 (0.00181)	-0.00073 (0.00178)	0.00072 (0.00104)	-0.00074 (0.00102)	-0.00135 (0.00193)
Panel B. Spillover results by intervention scope					
Fixed checkpoints during & post intervention	-0.002 (0.005)	-0.004 (0.003)	-0.002 (0.003)	-0.003 (0.002)	-0.004 (0.004)
Rotating checkpoints during & post intervention	0.008* (0.005)	-0.001 (0.004)	-0.009*** (0.003)	-0.005* (0.003)	-0.003 (0.004)
Fixed checkpoints in 10 km during & post intervention	0.006** (0.003)	0.001 (0.003)	0.000 (0.002)	0.001 (0.002)	0.002 (0.003)
Fixed checkpoints in 20 km during & post intervention	0.000 (0.003)	-0.002 (0.003)	-0.001 (0.003)	-0.003* (0.002)	-0.005 (0.003)
Fixed checkpoints in 40 km during & post intervention	0.001 (0.003)	0.000 (0.003)	0.000 (0.002)	-0.001 (0.001)	-0.001 (0.003)
Rotating checkpoints in 10 km during & post intervention	-0.004 (0.003)	0.001 (0.003)	0.001 (0.002)	0.000 (0.002)	0.000 (0.003)
Rotating checkpoints in 20 km during & post intervention	-0.001 (0.003)	0.001 (0.002)	-0.001 (0.002)	0.001 (0.002)	0.003 (0.003)
Rotating checkpoints in 40 km during & post intervention	-0.001 (0.003)	-0.002 (0.002)	0.001 (0.002)	-0.001 (0.001)	-0.002 (0.003)
Month FE	Yes	Yes	Yes	Yes	Yes
Police Station FE	Yes	Yes	Yes	Yes	Yes
Mean of control	0.0849	0.0293	0.0329	0.0165	0.0454
N	5090	4724	5090	4724	5090

See note on the analogous regressions in Tables 3 and 4 for details on variables and specifications.

Table A4: Drunk Drivers Caught - Spillovers

	Main intervention				Final check	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment station					-1.021** (0.378)	-0.756* (0.443)
Rotating checkpoints		-0.246 (0.246)				-0.376 (0.328)
Weeks of checking	-0.063* (0.033)	-0.109** (0.043)	-0.060* (0.035)	-0.108** (0.046)		
Rotating checkpoints × weeks of checking		0.068* (0.038)		0.075* (0.038)		
Treated stations in 10 km	0.719*** (0.258)	0.557 (0.516)			-0.196 (0.157)	-0.025 (0.443)
Treated stations in 20 km	-0.452* (0.265)	-0.319 (0.376)			-0.137 (0.332)	0.656 (0.799)
Treated stations in 40 km	-0.413** (0.199)	-0.789** (0.344)			-0.423 (0.284)	-0.057 (0.390)
Treated stations in 10 km × weeks of checking	-0.045** (0.019)	-0.007 (0.042)	-0.024 (0.023)	0.019 (0.054)		
Treated stations in 20 km × weeks of checking	0.050* (0.028)	0.043 (0.041)	0.03 (0.031)	-0.001 (0.049)		
Treated stations in 40 km × weeks of checking	0.033 (0.026)	0.061 (0.040)	0.017 (0.025)	0.015 (0.041)		
Rotating checkpoints in 10 km		0.243 (0.945)				-0.329 (0.725)
Rotating checkpoints in 20 km		-0.285 (0.576)				-1.494 (1.102)
Rotating checkpoints in 40 km		0.742 (0.453)				-0.67 (0.627)
Rotating checkpoints in 10 km × weeks of checking		-0.051 (0.088)		-0.056 (0.095)		
Rotating checkpoints in 20 km × weeks of checking		0.01 (0.069)		0.058 (0.074)		
Rotating checkpoints in 40 km × weeks of checking		-0.033 (0.058)		0.018 (0.061)		
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Police Station FE	No	No	Yes	Yes	No	No
Mean of dep. variable	1.237	1.237	1.237	1.237	1.22	1.22
N	866	866	866	866	109	109

See note on the analogous regressions in Tables 7 and 9 for details on variables and specifications.

Table A5: Short-term Effect of Checkpoints on Accidents

	Darkness			
	(1) Accidents	(2) Deaths	(3) Accidents	(4) Deaths
Fixed checkpoints during & post intervention	-0.00145 (0.00357)	-0.00285 (0.00235)	-0.00089 (0.00385)	-0.00284 (0.00257)
Rotating checkpoints during & post intervention	-0.00878*** (0.00326)	-0.00393 (0.00323)	-0.00883** (0.00356)	-0.00352 (0.00359)
Fixed checkpoint 1 night before	0.00374 (0.00554)	0.00374 (0.00437)	0.00502 (0.00522)	0.00422 (0.00425)
Rotating checkpoint 1 night before	-0.0014 (0.00544)	-0.00342 (0.00403)	-0.00145 (0.00542)	-0.00347 (0.00402)
Fixed checkpoint 2 nights before			-0.00614 (0.00498)	-0.00293 (0.00383)
Rotating checkpoint before 2 nights			-0.00074 (0.00555)	0.00122 (0.00406)
Fixed checkpoint 3 nights before			0.00133 (0.00541)	0.00227 (0.00424)
Rotating checkpoint 3 nights before			0.00083 (0.00575)	-0.00366 (0.00396)
Month FE	Yes	Yes	Yes	Yes
Police Station FE	Yes	Yes	Yes	Yes
Mean of control	0.03558	0.01682	0.03558	0.01682
N	94,276	94,276	94,276	94,276

This table presents the impact of fixed and rotating checkpoint interventions on the number of daily road accidents and deaths from August 2010 through December 2011. The variables indicating a rotating/variable checkpoint 1, 2, or 3 nights before are station-level indicators equal to 1 if a checkpoint was carried out on any road in the past 1, 2, or 3 nights. The during & post intervention variable is positive for the duration of the sobriety crackdown and up to 90 days afterwards. All data was taken from police administrative reports, as collected in both treatment and control police stations.

Standard errors in parentheses clustered at the police station level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A6: Legal Prosecution of Drunken Drivers

	Cases Processed by Court		
	(1)	(2)	(3)
Rotating checkpoints	4.258 (4.338)	2.871 (3.997)	-0.615 (3.311)
Drunk drivers caught in station		5.467*** (1.475)	2.445 (1.602)
2 checkpoints per week			10.625*** (3.236)
3 checkpoints per week			17.663*** (4.072)
Police lines teams			17.685*** (4.915)
District FE	Yes	Yes	Yes
Mean of dep. variable	25.55	25.55	25.55
N	112	112	112

The outcome variable is the number of drunken driving tickets that were processed by the criminal court. Criminal cases in the court records were linked to the police station from which the ticket originated. Several local courts were not willing to share data. All specifications include controls for whether the police station is located on a major highway, the pre-intervention accident rate, and district fixed effects.

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

Table A7: Simulated Fixed vs. Rotating Pooled Results

	Darkness	
	(1)	(2)
	Accidents	Accidents
Fixed checkpoints during & post intervention	-0.00181 (0.00023)	-0.00046 (0.00020)
Rotating checkpoints during & post intervention	-0.00137 (0.00067)	-0.00132 (0.00023)
Month FE	Yes	Yes
Police Station FE	Yes	No
District FE	Yes	Yes
Mean of dep. variable	0.03578	0.03578
N	55692	55692

Results shown on nighttime accident data simulated for the Aug.-Jan. 2010 and Aug.-Jan. 2011 period using the baseline model and parameters as described in Section 6. Specifications are analogous to those estimated from real data in Table 4, with the exception of a different set of controls in Column 2. Column 2 include district fixed effects and all control variables that enter the structural estimation: fixed or rotating treatment indicators, day of week, urban/rural, national highway, and the number of prior accidents in the police station. Standard errors in parentheses clustered at the police station level.

Table A8: Simulated checkpoint surveys during intervention

	Drunk drivers and motorcyclists caught			
	(1)	(2)	(3)	(4)
Rotating checkpoint station	0.128 (0.116)	-0.413 (0.292)		
Frequency		-0.472 (0.121)		
Rotating checkpoint × frequency		0.287 (0.124)		
Weeks of checking			-0.090 (0.017)	
Rotating checkpoint × weeks of checking			0.071 (0.022)	
Number previous checkpoints				-0.040 (0.007)
Rotating checkpoint × number previous checkpoints				0.025 (0.010)
District FE	Yes	Yes	Yes	Yes
Police Station FE	No	No	Yes	Yes
Mean of dep. variable	1.237	1.237	1.237	1.237
N	866	866	866	866

Results shown on number of simulated drunken drivers apprehended during the 2010 and 2011 interventions, using the baseline model and parameters as described in Section 6. The specifications are analogous to those in columns 1-4 of Table 7.

Standard errors in parentheses clustered at the police station level.

Table A9: Simulated drunk drivers caught on final check

	All Stations		Rotating checkpoints		Fixed Checkpoints	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	-0.794 (0.140)	-0.412 (0.383)	-0.760 (0.188)	-0.804 (0.358)	-0.902 (0.170)	-0.136 (0.429)
Days since last checkpoint	0.009 (0.005)	0.001 (0.015)	0.004 (0.009)	0.013 (0.023)	0.017 (0.008)	-0.001 (0.015)
Frequency		-0.185 (0.157)		0.013 (0.160)		-0.394 (0.207)
Days since last checkpoint \times frequency		0.004 (0.005)		-0.004 (0.008)		0.009 (0.007)
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Control mean	1.729	1.729	1.729	1.729	1.729	1.729
Mean treatment effect, @ freq = 2	-0.676	-0.666	-0.705	-0.712	-0.676	-0.685
P-value of mean treatment effect	0.000524	0.000740	0.00278	0.00209	0.000257	0.000849
N	108	108	77	77	74	74

This table reports the reduced form correlations between the interventions and the number of simulated drunken car and motorcycle drivers caught at the final check conducted after the end of the intervention in all police stations, including control stations. The specifications are analogous to those estimated with real data in Table 9. Columns 1 and 2 compare pooled treatment police stations with control stations, columns 3 and 4 compare rotating checkpoint stations with controls, and columns 5 and 6 compare fixed checkpoint police stations with controls.

Robust standard errors in parentheses.

A4 Appendix Figures

Figure A1: Simulated Drunken Driving Probabilities Across Interventions - Social Learning

