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

Abhijit Banerjee[1], Esther Duflo[1], Louis-Mäel Jean[4], Daniel Keniston[2], and Nina Singh[3]

[1]Massachusetts Institute of Technology
[2]Louisiana State University
[3]Rajasthan Police
[4]J-PAL

August 18, 2024

Abstract

Should police activity be narrowly focused and high force, or widely dispersed but of moderate intensity? Critics of intense "hotspot" policing argue it primarily displaces, not reduces, crime. But if learning about enforcement takes time, the police may take advantage of that 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 three locations or fixed in the best one, 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.

*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 J-PAL Governance Initiative for financial support. Thanks to Steven Berry, Mira Frick, Debraj 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.

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

1 Introduction

A central question in law and economics is how to deploy limited law enforcement resources to maximal effect. One perspective advocates intense, pre-announced crackdowns, to take advantage of potential increasing returns from higher arrest probabilities on criminal activity (Sherman and Weisburd, 1995; Weisburd and Telep, 2014). Another argues that, unless it is possible to police all locations all the time, the deployment should be randomized over time and across potential crime locations, because otherwise offenders would switch to the unpoliced areas and continue their activities with impunity (Clarke and Weisburd, 1994; Mookherjee and Png, 1994).

The experimental evidence on this topic is mixed. While some studies find limited displacement effects (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). In one such study conducted with the police of Bogota, Colombia, Blattman et al. (2017) find suggestive evidence of displacement of property crime away from streets with randomly increased police patrols in, but little displacement of violent crime. Andres and Mobarak (2019) show that Chilean fish vendors learn to disguise illegal sales from regulators, particularly if enforcement is intense.

One factor that may limit displacement is that potential criminals need to know the police strategy in order to undo it. If lawbreakers learn slowly about enforcement activity, then the police will stop the most criminals by concentrating their efforts on the most crime-prone spots (Anwar and Loughran, 2011; Wilson et al., 2017). Moreover, even after criminals discover that those locations are being policed, they may continue to frequent their favored location if, based on their past experience, they expect a brief crackdown. On the other hand, if perpetrators learn about the new initiative quickly enough, and have attractive alternative locations to shift illegal activity, then randomizing enforcement across locations is optimal.

Evidence on the speed of criminal learning about specific police initiatives remains sparse, and mixed. Sherman (1990) argues that many crackdowns show long post-program crime decreases, while other studies find that effects disappear quickly (Weisburd and Telep, 2014). Thus, the ability of lawbreakers to adjust to police efforts—both in terms of reducing or shifting criminal activity during a crackdown, and returning to their practices afterward—remains a central and open empirical question in the field of criminal behavior.

We examine this issue using a randomized crackdown on drunk driving in the Indian state of Rajasthan, combined with a structural model of learning. Some police stations conducted checkpoints only on the route judged most conducive for catching drunk drivers. Other stations randomly rotated checkpoints across the three most promising routes. Still others implemented no crackdown, which allows us to evaluate the effectiveness of the campaign itself. The intensity of the crackdown was also varied in two ways: police stations involved in the intervention were randomly assigned to have one, two, or three checkpoints per week; and the duration of the crackdown was varied from one to three months.

The experiment shows that checkpoints work. Over the duration of the crackdown and the subsequent three months, the number of deaths and accidents at night (when checking took place) decreased by about 25% and 17%, respectively, in all treatment precincts relative to the control areas. Moreover, these effects come mainly from the rotating-check locations, where accidents remained lower even in the 90-day period after enforcement ended. Meanwhile, the number of drunken drivers apprehended per night remained steady with rotating checks, but decreased over time in fixed-check precincts.

Motivated by these data patterns, 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 drivers' initial beliefs and payoffs. These parameters confirm the importance of learning: drivers' estimated priors about the intensity of checking in the most preferred location are nearly uninformative, so learning is very fast. Displacement of drunken driving is heterogeneous: increased checking causes drivers to avoid some locations, but seek out others. To our knowledge, this is the first time a Bandit model of learning has been applied in the policing literature, despite many informal discussions implying the same underlying logic.

With these structural parameters in hand, we can evaluate the effectiveness of alternative enforcement policies. Given the beliefs held by drivers in Rajasthan at the time of the beginning of our intervention, a long (100-day) campaign is optimal to convince drivers that the police are in fact checking on all roads—even those where the population initially has very low expectations of finding a checkpoint. In contrast, if the population had had correct beliefs about the intensity and duration of the police checkpoint strategy, then the optimal campaign duration would have been much shorter—20 days, and highly focused with 90% of checkpoints on the preferred road. In this scenario, the police's main goal is simply to demonstrate that the crackdown has begun through the most visible "show of force".

This paper analyzes a large scale randomized controlled trial conducted in collaboration with a police department as part of a real drunk driving campaign. It informs the optimal deployment of resources in this and future drunk driving campaign, but also contributes to our understanding of the police deployment literature by highlighting the role of the critical parameters in the learning model.

2 Background: Enforcing Drunk Driving Laws

Each year 1.35 million people die in traffic accidents worldwide, with as many as 50 million injured. 90% of these deaths happen in developing countries (World Health Organization, 2018), where death and accident rates are rapidly increasing, even as they fall 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 available data suggests that drinking plays a major role in traffic accidents. A review of studies conducted in low and middle-income countries names alcohol as a factor in 33% to 69% of road fatalities, and between 8% and 29% of accidents (World Health Organization, 2018). Evidence on optimal policing in developing countries, and the enforcement of anti-drunken driving law in particular, is often lacking since police rarely have the staffing or technology needed to measure drivers' alcohol levels. Nevertheless, law enforcement constitutes a major share of government expenditure: in 2018 the State of Rajasthan spent more on policing than on water supply and sanitation, urban development, all types of irrigation, and almost as much (70%) as on all medicine and public health (Comptroller and Auditor General of India, 2019).

Sobriety checkpoints have been evaluated in a wide variety of contexts, with recent meta-analyses (Peek-Asa, 1999; Erke et al., 2009; Elder et al., 2002) suggesting they reduce accidents by about 17% to 20%, and traffic fatalities by roughly the same amount. However, since (to our knowledge) no previous research has used a randomized trial, these results may be biased by the endogenous location and timing of the interventions. Furthermore, the vast majority of research has been conducted in developed countries, and consists of increasing checkpoints above an already high standard of enforcement. Thus, little is known about the impact of sobriety checkpoints in a low-income context, or versus a counterfactual of essentially zero enforcement.

The India police catch drunken drivers at barriers arranged on the roadway, where officers pull over selected vehicles as they pass through. In the Rajasthan intervention, officers were instructed to prevent traffic delays by letting all waiting vehicles pass by if they observed the checkpoint was causing congestion. After stopping a vehicle, police personnel observe the driver's demeanor and smell their breath. If the driver appears drunk, police order them to blow into a breathalyzer, and the driver is either charged or released.[1] Once caught, the drunken driver's vehicle is impounded, and the driver must pay a fine in court (usually Rs. 2000, approximately $50 in 2012) or potentially face jail time, although imprisonment is never observed in our data. Because this process could take as little as one day, the incapacitation effects of arrest or vehicle confiscation are unlikely to be significant in this setting and our model focuses on the deterrent effect of avoiding future apprehension by police.

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. Unscrupulous police officers might also accept (or solicit) a bribe. 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 to face either a demand for cash or a formal citation face a strong incentive to avoid being caught driving drunk.

Sobriety checks and even breathalyzers were extremely rare prior to the intervention and in the control police stations during the intervention. In the 925 nights that surveyors visited control police stations, they only witnessed the police carrying out a checkpoint on seven (0.76%) occasions. Nevertheless, conversations with the police and our personal observations at the checkpoints suggest that all citizens knew that

[1] This "selective breath checkpoint" protocol is commonly used in US sobriety checkpoints as well (Elder et al., 2002).

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

driving under the influence was illegal, though many were surprised to encounter a breathalyzer. Thus, when interpreting the quantitative magnitudes of the reduced-form parameters, one should bear in mind the context of no previous enforcement.

3 The Intervention

The Rajasthan crackdown on drunken driving was implemented as a large-scale randomized controlled trial (RCT), consisting of three overlapping experiments and a control group. The program took place in two phases: an initial pilot, from Sept.–Oct. 2010, and a larger rollout from Sept.–Nov. 2011. The initial pilot covered two districts and 40 police stations, and the second wave covered 11 districts and 183 police stations serving a territory of 125,000 km^2 with a population of approximately 30 million. Treatment status was stratified by district, whether a station was located on a national highway, and total accidents between 2008–2010. 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 below, we combine data from both intervention periods and we control for any time trends using month fixed effects.

The design is summarized in Table 1, and in the remainder of this section we explain the design choices in detail.

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 drivers in their area. In the fixed location group, the checkpoints were (to the greatest degree possible) always carried out at the best location, and on the best day of the week (chosen by the police chief). In contrast, rotating checkpoints moved among the three top locations on random days. Each police station's rotation was randomized in advance by the research team. In all interventions, checkpoints were held from 7:00pm to 10:00pm, which was when peak drunk driving took place, according to the police.

Checkpoint Frequency and Duration

Police stations were randomly assigned to carry out one, two, or three checkpoints per week.[3] These frequencies were at the police station level, not the road level; for example, in a rotating-group police precinct with a frequency of two checkpoints per week, on average each of the three roads would have a checkpoint twice every three weeks. The duration of the crackdown was also randomized at the police station level. In the 2010 intervention, stations conducted checks for between one and 1.5 months; in the 2011 intervention the crackdown lasted between two and three months.

Checkpoint Personnel

Previous work with the Rajasthan Police (Banerjee et al., 2021) suggests that government initiatives may be poorly implemented if civil servants are not motivated. To gain further insight on the role of monitoring and motivation, as well as to guard against a failure of the project due to poor implementation, special teams of officers from the district headquarter's reserve force, the "Police Lines", carried out all checkpoints in a randomly chosen set of stations. The Police Lines were considered by officers an undesirable "punishment" posting. Police Lines teams were monitored by GPS devices in their vehicles and informed that good performance might improve their chances for transfer out of the Police Lines. Analysis of this intervention is complicated by the fact these officers differed from regular staff along many dimensions, including motivation, monitoring, local knowledge, and inherent capabilities.

4 Data

To evaluate the effects of the crackdown on drunken driving we draw on a combination of administrative data on road accidents and deaths, as well as data collected by surveyors on vehicles passing and stopped at checkpoints.

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 have been collected

[3] Discussions with the police determined that it was not feasible to carry out more than three checkpoints per week per police station.

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 are 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.

Summary statistics for control stations are presented in Panel A of Table 2. The data, averaged at the police station/month level, show that control police stations have roughly 0.12 accidents per day, a quarter of which occur at night, and 0.05 deaths (40% occurring at night). The majority of deaths (88.5%) occur in single-fatality accidents.

4.2 Survey Data

We supplement the police administrative information with additional data collected by surveyors sent to monitor randomly selected checkpoint locations. Surveyors visited both on nights when the police were conducting drunk driving checkpoints 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, 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.

The summary statistics of the data collected by these monitors is displayed in Panels B and C of Table 2. Panel B displays the average number of passing vehicles, using surveyor counts from the locations identified by the control police stations as those where they would have carried out the checkpoints. The first and second checkpoint locations have similar amounts of traffic, while the third location seems less busy. The third location was also less likely to be near the police station itself: 11.5% of checkpoint 1 and 6.0% of checkpoint 2 were in front of a police station; this is true for only 2.7% of checkpoint location 3. This may affect driver's prior beliefs about the likelihood of encountering the police in these spots.

The overall magnitude of these numbers reflects the generally low levels of traffic on (mostly) rural Rajasthan roads, with vehicles passing every 10–15 seconds. Surveyor reports and disaggregated vehicle count data suggest that there were few if any cases of the checkpoints causing significant congestion delays. Overall, police stopped 13.1% of passing vehicles, roughly 106 per checkpoint, of which the majority were motorcycles. Panel C shows the effectiveness of the crackdown in catching drunken drivers: on average, police caught 1.89 drunk drivers per checkpoint, primarily motorcyclists.

After the end of the intervention, the surveyors collected data from a special final round of checkpoints. These checks were held once in all precincts, regardless of earlier treatment or control assignment, the week immediately after the intervention had concluded. Final checks were always at the location designated as the second-best place to catch drunken drivers. Hence, for the fixed-checkpoint intervention stations (and control ones) there had previously been no enforcement at these specific locations, whereas for the rotating stations there had previously been checking at these exact spots. On these nights, police were asked to conduct checks either randomly or at a fixed interval of cars (e.g. one in ten gets stopped), which resulted in a 2.26% overall drunkenness rate, and a rate of 3.43% 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 and Results

Driver actions, and hence accident and arrest rates, are complex functions of the utility of drunken driving, the cost of encountering the police, and beliefs about the intensity and duration of checking. However, much of the intuition behind our results is captured by the driver's choice of whether to drink and drive on one road, learn 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 a temporary period, then resume. If they encounter police a second time, they cease permanently. The police crackdown strategy then affects drunken driving via three margins. First, it alerts drivers that enforcement is happening and triggers their temporary sobriety period. Second, it affects the length of the sobriety period. Third, it affects the share of drivers who re-encounter police and thus are convinced to permanently stop offending. The main implications are intuitive.[4]

1. Persistence: The impact of the crackdown will carry on after the last checkpoint is finished because 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.

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 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 regressions of accidents and deaths in the entire area covered by the police station (including, but not restricted to, the three main roads identified for potential checkpoints) during daylight or nighttime. Accident and death data are reported as the daily accident rate, with observations at the police station-month level. When the intervention began or ended mid-month, we split the month into treated and untreated period observations and weight according to the number of days in each. We can thus interpret all outcomes as the change in daily accidents or deaths at the police station level.

We begin with a summary of the effects of the program in Table 3, which displays the impact of any enforcement during and 90 days post intervention, including police station and month fixed effects. We find a significant decrease in the number of nighttime accidents (17%) and deaths (25%).[5] There is no significant impact during

[4]For full details and formalism, please see the working paper version of this study, Banerjee et al. (2019).

[5]These figures refer to the average effect across Police Lines and regular personnel interventions.

the day, which is reassuring since all the checkpoints took place at night, though as we will comment later, there could be a strategic reaction on the daytime accidents. We find no significant evidence that the Police Lines intervention was more or less effective than the status-quo police personnel.

One of the main hypotheses of interest is the difference between the effectiveness of rotating and fixed checkpoints. Columns 1–4 of Table 4 repeats the analysis from Table 3, showing that the positive outcomes are largely driven by the rotating checkpoints police stations. Accidents at night went down by 35% ($p<0.003$) and nighttime deaths dropped 36% ($p<0.080$) in these precincts, while estimates are much smaller in the fixed police stations. The difference between fixed and rotating stations is statistically significant for nighttime accidents, though not for deaths. Interestingly, the disaggregated analysis in Table 4 shows an increase (though non-significant) in daylight accidents in rotating-check stations during the intervention, but no similar result in the fixed-check stations. This may be caused by drivers, either drunk or sober, who shift their travel patterns to earlier times of the day in the rotating checkpoint precincts, but can simply continue driving at the same time on alternate routes in the fixed police stations.

The effect of the intervention should be increasing over time, as more drivers notice the checkpoints, and should persist for some time after the intervention ends. Columns 5 and 6 of Table 4 provide evidence for these claims. We see that the decline in nighttime accidents and deaths in rotating stations is almost equally strong both during and after the intervention. Overall, the evidence suggests that the basic insight of the criminology literature is very powerful: checking at rotating locations causes a greater reduction in accidents and deaths, and this persists over time. In column 7 we examine effects on accidents deemed "serious" by local police. Though the sample size is smaller (this outcome was not measured in certain months), results are qualitatively similar to both accidents and deaths outcomes.

We now turn to the data collected by surveyors on the number of drunken drivers apprehended at the police checkpoints. In Table 5, column 1 shows that about 10% more drunken drivers, are caught at rotating check locations than in fixed locations. This number is not statistically significant, but without strategic behavior we might expect 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 drunk drivers are caught per night. That effect entirely disappears in rotating stations. 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 further 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.[6] The bottom row of Table 5 shows that on nights when a surveyor was present the Police Lines teams performed significantly better than others, catching 1.1–1.4 additional drunken drivers.

Our final measure of the impact of the interventions on drunken driving comes from the last night of checking. Since this final check was always held at checkpoint location 2, we can measure displacement effects even in the fixed-location stations where all other checks were on route 1. The first three columns of Table 6 focus on rotating checkpoints, and confirm drivers' avoidance behavior, with 63% fewer drunken drivers caught in treatment stations than in the control.[7] On average, there are no significant signs of reversion, although in stations with a frequency of three checkpoints per week the deterrent effects are larger and there are some (non-significant) suggestions of drunken driving returning after the final checkpoint. Surprisingly, the results from fixed checkpoints (columns 4–6) also show decreases in final-check drunken drivers relative to the control, albeit smaller. Although one might have expected that fixed checkpoints on road 1 would lead to displacement of drivers onto road 2, in fact the net effect of checks on road 1 reduced them by 32% on roads 2 as well. Here we do find significant evidence of reversion, with the decrease lasting only 39 days. As we will

[6] A potential concern is that many scheduled checkpoints did not occur due to police partial 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, presented in Table A3, are qualitatively the same as those in Table 5. 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.

[7] Since only 109 stations actually conducted the final check (60% of the total), there is a potential concern that compliance may be correlated with the intervention categories. Appendix Table A8 tests for differential compliance across intervention branches. We find no evidence that this occurred.

show below, this result is not inconsistent with the model, since drivers might have had priors that checking on road 2 will also be intense in the event of a crackdown. If they occasionally travel both roads, and notice a checkpoint on road 1, they will then conclude that switching to road 2 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.

Finally, the number of vehicles passing the checkpoints provides further insights into driver learning—unlike drunken driver arrests, passing vehicles could also be measured in control stations. Table 7 divides these outcomes by road to highlight substitution effects. On road 1 (column 1), checkpoints decrease the number of passing vehicles and, as with drunken drivers, the frequency of checking increases avoidance only in the fixed locations. The avoidance effect of rotating checks on roads 2 and 3 is comparable in magnitude (and significance for road 2) to road 1. Consistent with final check results on drunken driving, the displacement effect on road 2 passing vehicles from fixed checking on road 1 is also negative on net, increasingly so with frequency.[8] Consistent with the lack of effect on accidents and deaths, we find no significant effect of the Police Lines intervention on passing vehicles. The large magnitude of the decrease in passing vehicles cannot be explained solely by drunken drivers, who probably make up 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 that come with the checkpoints. Understanding the sources of this reaction is important for assessing the overall welfare implications of the intervention, but beyond the scope of this study.

In Appendix Section A1 we examine additional mechanisms and outcomes beyond our main analyses, with tables in Appendix Section A4; here we briefly summarize their results. The effect of the crackdown was very similar in the 2010 pilot and the 2011 main intervention, though less precisely estimated in the pilot. Stations with more frequent checking exhibited greater decreases in accidents, but our estimates are too noisy to say more about the shape of this relationship. Consistent with

[8]This effect is somewhat obscured by the parameterization of Table 7. Since passing vehicles are decreasing in the number of days after the intervention in fixed stations, and the final checks took place on average about two weeks after the final checkpoint, the net effect of fixed checking is to reduce the number of passing vehicles on route 2 at the time of the final checks, relative to the controls.

relatively isolated and independent nature of police stations in rural Rajasthan, we find no spillover effects to nearby police stations either in terms of accidents, deaths, or drunken drivers apprehended. 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 absence of police misconduct). Neither outcome is significantly different between rotating and fixed interventions, although we find that police implementation in three-checkpoints-per-week stations is lower than in less-intense interventions.

Finally, we consider the possibility that individuals might communicate with each other about police activity or learn while driving sober. 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. We discuss this test further in Appendix Section A1 and present results in Appendix Table A7. In fact, we find no particularly strong reduction in drivers caught in the three days after a prior checkpoint. 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 simple intuition outlined in Section 5 (and formally presented in Banerjee et al., 2019) suggests the mechanisms behind driver and police decisions, estimating the parameters determining drivers' behavior and designing the optimal police enforcement strategy requires a more detailed model.

Suppose agents receive utility d_r ($r \in \{1, 2, 3\}$) from drunken driving on road r, a route potentially checked by police. If they encounter the police, agents experience disutility c.[9] However, drivers are unsure whether they will encounter the police, both because they do not know if a sobriety crackdown is ongoing and because they do not know the intensity and geographic scope of the crackdown if it is in progress. Let

[9]The c parameter encompasses multiple factors affecting driver's disutility of encountering a checkpoint: their subjective probability of engaging with the police, their costs in time and money if they are caught drunk, 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.

driver i's subjective probability that a crackdown is ongoing at time t be π_{it}, and their belief about the probability of encountering a checkpoint on road r, if a crackdown is indeed ongoing, be λ_{irt}. The static deterministic component of utility of driving on road r at time t is then,

$$u_{rt} = d_r + \pi_{it}\lambda_{irt}c \tag{6.1}$$

where we normalize the static utility of the alternative action which we call "staying at home" to 0.[10]

Based on these payoffs, as well as the dynamic value of the information they may learn by driving (elaborated below), agents choose whether to drink and drive, and if so on which road. They perform their highest-utility action, and observe the results: if they drink and drive they may witness a checkpoint, 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.

6.1.1 Evolution of Beliefs

We make the following assumptions to quantify driver learning:

Assumption 1. Drivers know that the anti–drunken driving campaign takes the form of a single discrete interval of enforcement, which ends on each night with probability η. They cannot observe the starting date, ending date, or η with certainty.[11]

Assumption 2. If a crackdown is ongoing, police checkpoints occur randomly at location r with probability λ_r, and are independent across locations. Drivers do not know λ_r with certainty.

Assumption 3. Drunken drivers do not anticipate the crackdown.

Assumption 4. Driver i's time-t belief about the probability of checkpoint on route r, λ_{irt} is distributed Beta $\left(\alpha^{\lambda}_{irt}, \beta^{\lambda}_{irt}\right)$.

Assumption 5. Driver i's time-t belief of the probability that checking ends in period t, η_{it} is distributed Beta $(\alpha^{\eta}_{it}, \beta^{\eta}_{it})$.

Assumption 3 implies that, until they encounter the first checkpoint on the road,

[10] Our 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 discuss this issue in further detail below.

[11] We 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.

drivers' beliefs remain constant. Once a driver observes that the crackdown has begun, each time they travel on a road their belief about the intensity of police checking is updated and becomes more precise. Formally, consider a driver who has chosen to take road r in period t, and let $\mathbb{I}_{irt}$ be an indicator equal to 1 if they saw a checkpoint. By Assumption 4, they update the parameters of their beliefs according to $\alpha^{\lambda}_{ir,t+1} = \alpha^{\lambda}_{irt} + \mathbb{I}_{irt}$ and $\beta^{\lambda}_{ir,t+1} = \beta^{\lambda}_{irt} + (1 - \mathbb{I}_{irt})$. Their subjective probability of being checked on road r the following night is then,

$$\lambda_{ir,t+1} = \frac{\alpha^{\lambda}_{irt} + \mathbb{I}_{irt}}{\alpha^{\lambda}_{irt} + \beta^{\lambda}_{irt} + 1}$$

If the driver had chosen to remain at home or if they had not yet encountered any checkpoints (Assumption 3), their beliefs about checking intensity would not evolve: $\lambda_{ir,t+1} = \lambda_{irt}$. By Assumption 2, beliefs about intensity on other roads are unaffected: observing a crackdown on route r is only informative about the risks of driving on other roads in that it may inform the driver that the sobriety campaign is still in force.

Assumption 5 implies that agents' expectation that the crackdown ends after night $t + 1$ (again, conditional on it being in progress in period $t + 1$), is $\eta_{i,t+1} = \alpha^{\eta}_{it} / \left(\alpha^{\eta}_{it} + \beta^{\eta}_{it} + 1\right)$. Since drivers never definitively observe the end of the crackdown (Assumption 1), η_{it+1} evolves deterministically based solely on the time passed since the driver first encounters a checkpoint. Intuitively, as time passes, drivers infer that if the crackdown is still going on then the chance it will end in each period decreases steadily.[12]

A driver's history of encounters with the police, combined with their beliefs about crackdown intensity and duration, determine their beliefs about whether the crackdown is currently ongoing, π_{it}. If the driver encounters a checkpoint at time t, they know with probability 1 that the crackdown was in force during period t, and thus the probability that it continues in the next period is $\pi_{i,t+1} = 1 - \eta_{it}$. If they drive on road r but see no police checkpoint, then their posterior belief on π_{it} reflects that the crackdown may have ended, or possibly the crackdown is still ongoing but the police did not implement a checkpoint on that night. This posterior belief, adjusted for the fact that the crackdown might end between t and $t + 1$ is

[12]Since the driver never conclusively observes the end of the crackdown, $\alpha^{\eta}_{it} = \alpha^{\eta}_{0}$ for all t implying that $\eta_{i,t+1} = \alpha^{\eta}_{0} / \left(\alpha^{\eta}_{0} + \beta^{\eta}_{it} + 1\right)$ for all t.

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

Since there had never previously been a crackdown on drunk driving in the area where the project was implemented (Assumption 3), we set $\pi_0 = 0$.

These 9 variables, $\Psi_{it} = \left(\left\{\alpha_{irt}^{\lambda}, \beta_{irt}^{\lambda}\right\}_{r=1}^{3}, \alpha_0^{\eta}, \beta_{it}^{\eta}, \pi_{it}\right)$ constitute the state variables in the driver's decision problem.

6.1.2 Dynamic Utility and Choice Probabilities

The evolution of a driver's dynamic utility depends upon the outcome of driving. If the driver witnessed a police checkpoint on road r, we denote their posterior beliefs as $\Psi_{it+1}(check_r)$, if they see no checkpoint, then their beliefs become $\Psi_{i,t+1}(nocheck_r)$. Beliefs of those staying at home are $\Psi_{i,t+1}(home)$.

Drivers' behavior is not entirely deterministic: they experience random shocks ϵ_{irt} 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 drivers' actions, their information about police activity, and hence their behavior in future periods. Dynamic utility, incorporating these preference shocks is then,

$$V(\Psi_{it}) = \max \begin{cases} v_1(\Psi_{it}) + \epsilon_{i1t} & = u_1 + \delta\left(\pi_{it}\lambda_{i1t} V(\Psi_{i,t+1}(check_1)) + (1-\pi_{it}\lambda_{i1t}) V(\Psi_{i,t+1}(nocheck_1))\right) + \epsilon_{i1t} \\ \dots & \dots \\ v_H(\Psi_{it}) + \epsilon_{iHt} & = \delta V\left(\Psi_{i,t+1}^{1}\right) + \epsilon_{iHt} \end{cases} \tag{6.2}$$

Assumption 6: Consistent with much of the empirical learning literature (Ching, Erdem, and Keane, 2013), we assume that the choice-specific shocks in Equation 6.2 are distributed IID extreme value type 1.

Let $x_{s,t-1}$ be the history of checkpoints in police station s until time t, and $\epsilon_{s,t-1}$ be the history of shocks to drivers at that station. Because different drivers experience different shocks, their actions and beliefs differ. Let $h\left(\Psi_{it}; \Psi_0, \{d_r\}_{r=1}^{3}, c, x_{s,t-1}, \epsilon_{s,t-1}\right) \equiv h_{st}(\Psi_{it})$ denote the distribution of these beliefs in police station s, on night t. Then Assumption 6 implies that the fraction of potential drunken drivers who choose to drive on road r on night t is μ_{srt}:

$$\mu_{srt} = \int_i \frac{\exp(v_r(\Psi_{it}))}{\sum_{r'} \exp(v_{r'}(\Psi_{it})) + \exp(v_H(\Psi_{it}))} dh_{st}(\Psi_{it}) \tag{6.3}$$

6.2 Estimation

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

The model contains two sets of parameters. The first relates to the driver's preferences and the initial conditions of their beliefs, $\theta_1 = \left(\{d_r\}_{r=1}^3, c, \left\{\alpha_{r0}^\lambda, \beta_{r0}^\lambda\right\}_{r=1}^3, \alpha_0^\eta, \beta_0^\eta \right)$. These form the 12-element parameter vector θ_1 that determines driver behavior and hence allows the calculation of the optimal police crackdown strategy.[13] The second set of parameters, θ_2, is a vector of auxiliary parameters controlling for other exogenous characteristics of the police station jurisdiction and checkpoint implementation that may affect police effectiveness or the local number of accidents, but not the drivers' learning and decision process. We specify the elements of θ_2 in the estimation section below.

We follow a three-step procedure to generate the predictions of accidents and drunken drivers arrested which we subsequently take to the data. First, we numerically solve for the value function $V(\Psi_{it}, \theta_1)$ as defined in Equation 6.2 by iteration. We make an initial guess for the value function $V(\Psi_{it}, \theta_1)$, solve for the optimal choices based on Equation 6.3, then update $V(\Psi_{it}, \theta_1)$ based upon the implied payoffs. This algorithm converges to a value function which satisfies Equation 6.2 for all possible driver beliefs. The calculation is complicated by the relatively high number (8) of time-varying state variables and the infinite-horizon nature of the problem, making backwards induction impossible.[14]

Second, using this value function, we simulate agents' actions as they encounter the sequence of checkpoints that occurred in each police precinct. We simulate $H = 5,000$

[13] The baseline specification sets the discount rate to $\delta = .95$.

[14] Because 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. 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 grid points, we interpolate.

histories of potential drunken drivers' for each of the 223 police stations, and we follow these sequences through each night of the intervention and afterwards, for 121 days each. Starting from the initial beliefs about checking probabilities and crackdown duration in θ_1, each simulated driver chooses actions with the probabilities expressed in Equation 6.3. Based on the outcomes of their actions, agents' beliefs are updated according to the formulas in Section 6.1.1 and these modify their choice probabilities in the next period. We average these choices across agents to generate an estimate of the share of drunken drivers on each route on every night.[15] We denote the share of simulated agents in station s, drunken driving on road r, at night t as $\mu_{srt}(\theta_1, x_{s,t-1}, \epsilon_{st}) \equiv \tilde{\mu}_{srt}$. We lack data on actual police implementation of 24% of assigned checkpoints; in these cases we use predicted implementation. See Appendix A2 for details.

The third and final step is to map the simulated fraction of agents drunken driving, $\tilde{\mu_{srt}}$ to the number of accidents observed (A_{st}) and drivers arrested (N_{srt}). These outcomes are also a function of the populations of potential drunken drivers, their propensity to get into an accident (to match data on the number of accidents observed), and the effectiveness of local police at catching them (to match data on drivers arrested). Since our model has little to say on these issues, we include two vectors of auxiliary parameters, θ_2^{Acc} and θ_2^{Arr} which are associated with the baseline number of accidents and arrests, respectively.

We model the number of accidents in station-night st as

$$\tilde{A}_{st} = \left(1 + \sum_{r=1}^{3} \tilde{\mu}_{srt}\right) w_{st}^{Acc\prime} \theta_2^{Acc} \tag{6.4}$$

where the exogenous variables affecting the number of accidents are w_{st}^{Acc} and their corresponding coefficients are in the vector θ_2^{Acc}. These controls include police station, month, and day of the week fixed effects. As in the reduced form specifications, the identifying variation thus derives from changes in accident rates within the stations over the course of the intervention. Intuitively, the baseline accident rate, expressed as $w_{st}^{Acc\prime}\theta_2^{Acc}$ in the estimation, is shifted by the share of potential drunken drivers who choose to take any of the three roads modeled: $\sum_{r=1}^{3} \mu_{srt}$. The structural equation

[15]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).

allows for the fact that some accidents are unrelated to drunk driving: even if all drivers remained sober ($\mu_{srt} = 0, \forall s, r, t$), a quantity $w_{st}^{Acc\prime}\theta_2^{Acc}$ road accidents would still occur.

We model the expected number of drunken drivers arrested on road r using a similar functional form,

$$\tilde{N}_{srt} = \tilde{\mu}_{srt} w_{srt}^{Arr\prime}\theta_2^{Arr} \tag{6.5}$$

The vector of controls, w_{srt}^{Arr}, includes all elements of w_{rt}^{Acc} as well as variables related to police effort: whether the checkpoint was conducted by a Police Lines team, the number of vehicles stopped, and the time spent at the checkpoint, the latter two included as third-degree polynomials. The structural equation for arrests does not include an intercept, since if all drivers chose to remain sober there would be no arrests for drunken driving.

The difference between these predicted accident and arrest rates and the observed accidents and arrests in the data, A_{st} and N_{srt}, are,

$$\begin{bmatrix} \zeta_{st}^{Acc} \\ \zeta_{srt}^{Arr} \end{bmatrix} = \begin{bmatrix} A_{st} - \tilde{A}_{st} \\ N_{srt} - \tilde{N}_{srt} \end{bmatrix} \tag{6.6}$$

where ζ_{st}^{Acc} and ζ_{srt}^{Arr} are residual terms capturing factor such as holidays, idiosyncratic weather conditions, shocks to police manpower, etc. Let the stacked vector of residuals be $\zeta(\theta)$. We input these residuals into 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 orthogonal to any unobservable components of checkpoint effectiveness. In practice, we include 21 variables in the instrument matrix, $\boldsymbol{Z}$, divided into one set for the accident moments Z^{Acc}, and another for the arrest moments Z^{Arr}.

The 8 instruments in Z^{Acc} include indicator variables designating whether fixed or rotating strategies were used in a police station, interacted with the number of prior assigned checkpoints (two instruments), and those same variables further interacted with number of weeks since the end of the intervention (two instruments). These capture 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 four more instruments. These capture the very short-term learning effect.

The remaining 13 instruments, Z^{Arr}, create moments from the arrest data. The first two are simply indicators for whether the checkpoint was held on the second- or third-choice roads. The next three instruments are the number of previous checkpoints held at the location of a given night's checkpoint. We allow the own-location effect of checkpoints at location 1 to differ in fixed checking stations (one more instrument). The following three instruments are the opposite: the number of checkpoints previously conducted *at other* locations. The last set of four instruments contain an indicator for data from the final check, interacted with indicators for fixed or rotating checkpoints and both the number of previous checkpoints, and also the days since the previous checkpoint.

Let $\boldsymbol{Z}$ be the block-diagonal matrix of instruments combining exogenous variables $\left[Z^{Arr}\ w^{Arr}\right]$ and $\left[Z^{Acc}\ w^{Acc}\right]$, and $\zeta(\theta)$ be the vector of residuals. The GMM objective function is the standard,

$$\min_{\theta\in\Theta}\left(\boldsymbol{Z}'\zeta(\theta)\right)'\left(\boldsymbol{Z}'\Omega\boldsymbol{Z}\right)^{-1}\left(\boldsymbol{Z}'\zeta(\theta)\right)$$

We employ two-stage GMM, first solving for the values of θ that minimize the GMM criterion with $\Omega = (Z'Z)^{-1}$, 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 (Kennedy and Eberhart, 1995) to identify the neighborhood of the global minimum, then use the Nelder-Mead algorithm to refine the solution.

6.2.1 Identification

The 21 moments constructed from the design of the anti–drunken driving field experiment identify the 12 structural parameters of the drunken driving model. Because the model is overidentified, and due to the complexity of the relationship between the parameters and drivers' forward-looking decision problem, there is not always a simple mapping from the data moments to the parameters that they identify. In this section we first present a heuristic analysis of the ways in which the data from the randomized experiment, and the moments generated from it, are informative about the parameter values. We then discuss the results of a more formal analysis of identification following the procedure suggested by Andrews et al. (2017).

The relative utilities of drunken driving on different routes ($d_1 - d_2$ and $d_1 - d_3$) are identified by the difference in arrest rates on first versus the second and third routes, most clearly on the first night of checking. If fewer drivers are arrested on (say) the second route than the first, we infer the utility from drunken driving there must be lower. The level values of the utilities[16] (d_r) are identified most directly by the normalization of the constant term in Equation 6.4 to 1. This allows us to estimate auxiliary parameters θ_2^{Acc} from baseline data, and calculate the baseline share of drunken drivers from the changes in accidents over time.[17] The arrest data also helps identify utility levels, via the relative effectiveness of the 2nd (and later) checks compared to the first. Intuitively, if the utility of drunken driving on a route is high, then a large share of potential drunken drivers will learn of the police campaign after the first check, and subsequent checkpoints will be much less effective at apprehending drunken drivers. Then, we would expect a large decrease in road accidents the night immediately after a checkpoint if d_r is high, and many drivers witness the checkpoint.

Driver substitution patterns across roads, as well as the differences in accidents across intervention strategies, identify their prior beliefs about the intensity of checking (λ_{r0}). For instance, if intense checkpoints on route 1 lead to fewer arrests on route 2 (as we observe in the data, especially at the time of the final check) then drivers must believe that if there is checking on route 1, there is likely to be checking on route 2. We conclude that $\lambda_{2,0}$ must be large enough to induce this behavior.

The speed of driver learning in stations with different checkpoint intensities, both in the decrease of arrests on different roads as well as the drop in accidents, identifies the precision of drivers' priors about checking intensity and the duration of enforcement. If the number of arrests on a route declines rapidly after it begins being checked, then we conclude that drivers' prior beliefs about the intensity of checking on that route $\left(\alpha_{r0}^{\lambda} + \beta_{r0}^{\lambda}\right)$ are imprecise and their posterior belief about the risk of the route is increased substantially by observing the checkpoints. Conversely, a large difference between the rate of decrease in stations with three checkpoints per week versus one

[16] As in any discrete choice problem, the scale of the parameters and the variance of the idiosyncratic road preference shocks ϵ_{irt} are not separately identified.

[17] More formally, let $\bar{\mu}_{s0}\left(d_1, d_2, d_3\right) = \sum_{r=1}^{3} \tilde{\mu}_{sr0}$ be the baseline share of drunken drivers in station s. If an intervention decreases the fraction of drunken drivers by $\phi\%$ and decreases accidents by ω accidents per night, then we infer the baseline share of drunken drivers to be $\bar{\mu}_{s0} = \omega / \left(\psi w_{st}^{Acc\prime} \theta_2^{Acc}\right)$.

suggests that beliefs are precise and require substantial additional information to shift. The variance of the prior also shifts drivers' dynamic incentives to explore routes to gain knowledge of the police strategy. The extremely rapid learning on road 3 encourages drivers to try this route once they know that the campaign has begun, since they will learn quickly if there are checkpoints in this location.

The speed of agents' return to drunken driving is particularly informative about the beliefs on the duration of the crackdown (η_0 and $\alpha_0^\eta + \beta_0^\eta$). If reversion is quick and the effect of the crackdown on arrests and accidents is (initially) short-lived, this suggests that η_0 is high—drivers expect only a few checkpoints to be conducted. The change in the persistence of the checkpoint effect over time, in particular how it differs according to the police station's crackdown duration, pins down the precision of agents' priors about η. If, for example, the rate of reversion to drunken driving and accidents is very sensitive to the length of the crackdown, then we would infer that agents' priors on η are diffuse and thus $\alpha_0^\eta + \beta_0^\eta$ is low.

6.3 Results: Structural

Table 8 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 drivers' beliefs.

The first three columns 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 absolute values of these coefficients per se cannot be readily interpreted. The relative sizes are informative, however. In the baseline model, at 72,130 the perceived disutility of being caught by the police is vastly greater than the benefit to drinking and driving (which is slightly negative, implying on average people avoid it), showing that sobriety checkpoints are an efficient tool for crime prevention if deployed effectively. Prior to the crackdown, when drivers perceive zero risk of arrest, our parameters imply that 26% would be intoxicated on road 1, 25% on road 2, and 21% on road 3. Although the utility parameters mirror well the relative popularity of roads as was the case in real life (with road 3 significantly less attractive than roads 1 and 2), the implied shares of drunken drivers on each road are not very different suggesting that, as we show below, a rotating strategy will yield the highest reduction in drunkenness.

Columns 4–6 report the model estimates of drivers' beliefs. Drivers initially think that the crackdown will end after just one night with 20.7% probability, yielding an expected duration of the crackdown of only about three days (row 1). This is consistent with what senior police officers told us of past police enforcement of traffic laws, which consisted of only one or two checkpoints per initiative. Perhaps due to long experience with other traffic enforcement crackdowns, drivers' beliefs on this parameter are very slow to evolve—they behave as if they had almost 2000 nights of experience learning the duration of crackdowns (row 2).

Finally, rows 3–8 show the initial priors on the expected intensity of checking on roads 1, 2, and 3. Estimated beliefs are substantially different across all three roads. Driver's initial expectation of encountering a checkpoint on road 1 (14.5%) is 2.6pp larger than on road 2 and over twice as large as on road 3 (6.2%). Conversely, drivers update their beliefs about checking frequency extremely quickly on roads 1 and 3—so fast that the strength of the priors are estimated at virtually zero. What this reflects is that whatever their prior, they converge to the belief there is a crackdown. [18] Updating is slower on road 2, where drivers are as confident as if they had experienced about 8 nights of previous checking.

These parameters help rationalize the apparent paradox 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 significant displacement effect on road 2 in the fixed as well as the rotating stations (Table 6). Even if they observe checking only on road 1, agents' prior beliefs dissuade them from driving on road 2 for the duration of the intervention since the expected cost of driving on road 2 ($\pi_t \times 0.119 \times 72,130 = \pi_t \times 8,583$) is much greater than its utility value (-0.107) even for low beliefs about the probability of police activity (π_t). Because drivers are more confident in their beliefs about road 2, they do not risk experimenting by driving there, and numbers of drunken drivers remain persistently low even after checking is over. In contrast, drivers adapt their beliefs quickly about road 3, so there is a substantial option value to exploratory driving on this road.

[18] The standard error is large and the T statistic is very low. However, inference at the boundary of the parameter space raises substantial econometric complications. When only a single parameter is at the boundary, previous researchers have argued that conventional standard errors are likely to be conservative (Andrews, 1999).

In Figure 9.1 we present graphs of the share of agents drinking and driving in the precincts of each branch of the intervention.[19] The solid blue line shows the share on road 1, the dotted red line on road 2, and the dashed orange line on road 3. The first row of sub-graphs, illustrating the fixed intervention stations, clearly demonstrates the spillover effects resulting from checking on only one route. The share of drunken drivers on all roads drops sharply in the first days of the campaign, particularly in the stations with three checks per week. At first, the number of drivers on road 3 drops as well, since agents infer that if there is checking on road 1, there is also likely to be checking on 3 as well. However, this rebounds quickly as drivers' perceived likelihood of the crackdown continuing drops and they become willing to experiment on road 3 again, where they never personally witnessed a checkpoint. In contrast, the number of drunken drivers on road 2 remain persistently lower. The second row of Figure 9.1 shows the corresponding graphs from the rotating intervention stations. Here there is a decrease in drunken driving on all three roads, leading the program to be more effective overall.

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. 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 less 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 re-running the same reduced-form regressions used in the analysis of the actual data in Section 5 on a simulated dataset. Maintaining the same order of reduced-form regressions as Section 5, we begin with the accidents analysis. Column 1 of Table A10 shows police station-FE analysis of the simulated data predicts that both the fixed and rotating checkpoint locations yield a statistically significant reduction in night accidents, with the rotating

[19] These graphs use the estimated parameters and data on history of checkpoints actually conducted. Drunken driving probabilities are averages across all stations within each arm of the intervention.

strategy also being statistically significantly more effective than the fixed strategy.. This closely matches the actual results in Table 4 which show the rotating stations performing better.

The model's predictions on the dynamics of drunken drivers caught at checkpoints match the reduced-form estimates well. The simulated results in Table A11 show the exact same patterns of decreasing effectiveness in fixed but not rotating police stations, both with and without police station fixed effects, as in the corresponding columns 1–4 of Table 5, . The model is also able to match the results from the final check quite closely. Table A12 (corresponding to Table 6) shows substantially fewer drunken drivers caught on road 2 at the final check, and this difference also holds in the fixed-checkpoint stations where there was no enforcement on road 2, again matching the empirical data.

6.4 Optimal Enforcement Strategy

With the estimated structural parameters in hand, we calculate the optimal crackdown strategy. To simplify the exercise, we set two constraints on the design of this strategy. First, we assume that the police department has only enough budget to carry out 20 checkpoints per police station in every 120-day period.[20] Second, we 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 away from 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,[21] it has the advantage of being more transparent, easier to calculate, and more closely resembling the conditions of the actual Rajasthan campaign. One can interpret this counterfactual strategy as the best possible way to have designed that intervention, subject to the same constraints faced in reality. In Appendix A3 we consider how incomplete program implementation, of the type observed in the

[20]This is in line with what the police could implement in our context: during the second round of the intervention the police stations averaged 22.4 checkpoints over three months, which we were informed by the police was the largest crackdown they could implement.

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

Rajasthan campaign, might influence the optimal strategy.

The estimated parameters imply that the optimal strategy (shown in Table 9) is to spread the 20 checkpoints over 100 days, and to place 8 of them on road 1, dividing the other 12 across the roads 2 and 3. The allocation of checkpoints fairly evenly across roads, similar to the rotating strategy in the actual intervention, derives from two implications of the parameters estimated in the structural model. First, drivers' estimated utilities are not substantially 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 agents' priors concerning 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 67% drop in the rate of drunken driving.

An alternative strategy employed by many police departments is to publicly announce the beginning of a campaign against drunk driving, for example by radio advertisement, with the goal of dissuading drunken drivers immediately. In the second row of Table 9 we examine this alternative strategy, while maintaining the assumption that drivers' beliefs are as estimated at the beginning of the intervention.[22] 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 9 imply that optimal duration should increase slightly—from 100 to 105 days, with the same spatial distribution of checkpoints. In this scenario, the police no longer need to intensify checkpoints to alert drivers that the crackdown has begun; therefore, they can lengthen the program to increase its effectiveness. The prior announcement of the campaign substantially improves its efficacy: the rate of drunken driving now drops by 72%.

As discussed above, the intervention occurred in a setting where 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 has been ongoing for some time and the population is familiar with police strategies. To operationalize this counterfactual we hold constant the estimated utilities $\theta_2 = \left(\{d_r\}_{r=1}^{3}, c\right)$, while imposing that the drivers' beliefs $\theta_1 = \left(\left\{\alpha_{r0}^{\lambda}, \beta_{r0}^{\lambda}\right\}_{r=1}^{3}, \alpha_0^{\eta}, \beta_0^{\eta}\right)$

[22] 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.

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 $\alpha^{\lambda}_{r0} + \beta^{\lambda}_{r0} = \alpha^{\eta}_{0} + \beta^{\eta}_{0} = 10^6$ which captures the fact that beliefs are likely 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; citizens' 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 9 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 20 days, and 90% 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 10% of enforcement on roads 2 and 3, since otherwise, in equilibrium, drivers will know that these routes are unmonitored and immediately shift to them. This strategy is extremely effective: drunken driving decreases by 97%. One reason for the increased effectiveness in equilibrium, relative to the initial campaign against drunk driving in row 1, is that the average duration believed by agents (20 days) is much longer than their estimated crackdown length in the data (about three 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 meaningful chance of apprehension on roads 2 and 3, thus maximizing the negative displacement effect.

The final counterfactual we examine, in the fourth row of Table 9 presents the optimal policy if drivers have equilibrium beliefs about the parameters of the 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 length of greater than 22 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) drunk driving crackdown. 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.[23]

7 Conclusion

This paper presents the results from a randomized experiment in the context of a crackdown on drunk driving in Rajasthan, India designed to test a model of criminal learning and strategic behavior. The central conclusion is that there is clear evidence of learning, such that 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 substantial decrease in road accidents and deaths. However, just as drivers learn about the beginning of police enforcement, they also learn when it ends—after the intervention we see a slow reversion of driver behavior and a return to drinking and driving.

These results provide the data for a structural estimation of the parameters of a 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 in many (but not all) locations change quickly. They also reveal the two main factors underlying the results: the estimated "cost" of encountering a police checkpoint is very high (causing rapid behavioral adjustment), and 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.

The structural 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, would be a relatively long-term effort that spreads 20 checkpoints over 100 days. This ensures that drivers

[23]For 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.

are dissuaded by gradually revising upwards their (initially low) beliefs on duration of the crackdown. But the optimal crackdown is situation-specific: if drivers' 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 a short, intense crackdown.

Our approach might be applied to illuminate the relationship between crime and punishment in other contexts. The most natural application is to hotspot policing in developed countries, where a learning-based model might reconcile heterogeneous results in the existing literature on the persistence of the deterrence effect and the extent of displacement. Our empirical approach might also be applied to individual-level decisions to commit different types of crimes, and criminal learning about the probability of apprehension. Building upon reduced-form work by Wilson et al. (2017) and others, microdata on arrest histories could be analyzed to generate counterfactual law enforcement and incarceration policies designed to minimize crime and recidivism.

With additional development, our methodology might illuminate a greater variety of criminology topics. 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. A relatively new strand of research in the criminology literature (Sloan et al., 2014) applies new knowledge from behavioral economics to criminal behavior. Insights from this behavioral literature could be explicitly tested in structural models, and potentially incorporated into the design of future randomized policing experiments. Finally, the very high predicted efficacy of pre-announcing a crackdown raises the possibility that the police may wish to publicize campaigns 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.

More broadly, our work demonstrates the potential to extend learning models to a range of topics in which agents update beliefs about correlated factors that evolve over time. These might range from traders learning about demand and prices at spatially distinct markets, to farmers choosing a crop under changing climatic conditions, to the study of how police gain information about criminal activity through patrolling—the flip side of the subject of this paper. The main challenge remains computational: the context of checkpoints in Rajasthani villages is consistent with a small set of belief-state variables, but estimating a fully flexible set of beliefs on a larger range of states is likely to remain difficult. For example, in a context such as the city of Bogotá (Blattman et al., 2017), crime deterred from one hotspot might potentially be displaced into dozens of possible locations, substantially increasing the number of parameters. Our model and estimation strategy are best suited to environments in which agents choose between a relatively limited set of options, or in which choices can readily be grouped into a few discrete categories. Nevertheless, in many cases a detailed understanding of the context may suggest a parsimonious model that would open the topic to analysis.

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).

Andrews, D. W. (1999). Estimation when a parameter is on a boundary. *Econometrica 67*(6), 1341–1383.

Andrews, I., M. Gentzkow, and J. M. Shapiro (2017). Measuring the sensitivity of parameter estimates to estimation moments. *The Quarterly Journal of Economics 132*(4), 1553–1592.

Anwar, S. and T. A. Loughran (2011). Testing a bayesian learning theory of deterrence among serious juvenile offenders. *Criminology 49*(3), 667–698.

Banerjee, A., R. Chattopadhyay, E. Duflo, D. Keniston, and N. Singh (2021, February). Improving police performance in Rajasthan, India: Experimental evidence on incentives, managerial autonomy and training. *American Economic Journal: Economic policy 13*, 36–66.

Banerjee, A., E. Duflo, D. Keniston, and N. Singh (2019, September). The efficient deployment of police resources: theory and new evidence from a randomized drunk driving crackdown in India. Technical Report 26224, National Bureau of Economic Research.

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.

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.

Clarke, R. V. and D. Weisburd (1994). Diffusion of crime control benefits: Observations on the reverse of displacement. *Crime prevention studies 2*, 165–184.

Comptroller and Auditor General of India (2019). State Finances, Government of Rajasthan. Technical Report Report 2.

Davis, A., A. Quimby, W. Odero, G. Gururaj, and M. Hijar (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.

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.

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.

Kennedy, J. and R. Eberhart (1995). Particle swarm optimization. In *Proceedings of ICNN'95-International Conference on Neural Networks*, Volume 4, pp. 1942–1948. IEEE.

Mookherjee, D. and I. P. Png (1994). Marginal deterrence in enforcement of law. *Journal of Political Economy 102*(5), 1039–1066.

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.

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

		Randomized Implementation Staffing		
		Police Lines Teams	Police Station Teams	Total
A. Sep.–Oct. 2010 Round: 16 control police stations				
Randomized Checkpoint Strategy	Rotating	5 stations @ 2/week	7 stations @ 2/week	12 stations @ 2/week
	Fixed	6 stations @ 2/week	6 stations @ 2/week	12 stations @ 2/week
	Total	11 stations	13 stations	24 stations
B. Sep.–Nov. 2011 Round: 60 control police stations				
Randomized Checkpoint Strategy	Rotating	8 stations @ 1/week 11 stations @ 2/week 10 stations @ 3/week	10 stations @ 1/week 9 stations @ 2/week 12 stations @ 3/week	18 stations @ 1/week 20 stations @ 2/week 22 stations @ 3/week
	Fixed	9 stations @ 1/week 7 stations @ 2/week 9 stations @ 3/week	14 stations @ 1/week 13 stations @ 2/week 11 stations @ 3/week	23 stations @ 1/week 20 stations @ 2/week 20 stations @ 3/week
	Total	54 stations	69 stations	123 stations
	Grand Total	65 stations	82 stations	41 stations @ 1/week 64 stations @ 2/week 42 stations @ 3/week

Frequency of checking (1, 2, or 3/week) was assigned randomly.

Table 2: Summary Statistics

	Obs.	Mean	SD	Median	Min.	Max.
A. Police station daily accidents and deaths (control stations)						
Accidents	1616	0.12	0.09	0.1	0	0.58
Deaths	1616	0.05	0.05	0.03	0	0.43
Night accidents	1616	0.03	0.04	0.03	0	0.29
Night deaths	1496	0.02	0.03	0	0	0.26
B. Total vehicles passing police checkpoint locations in control stations						
Location 1	238	936.6	720.75	661.50	117	4692
Location 2	243	924.96	911.7	612.00	123	4998
Location 3	274	875.83	869.7	544.0	38	4726
C. Vehicle Checking Outcomes						
Whether checkpoint occurred	1565	62.62%	23.41%			
Total vehicles stopped	866	106.15	108.24	71	1	1169
Drunken drivers caught	866	1.89	2.36	1	0	21
Percent drunk (final check)	44	2.26%	5.24%	5	0	184

Number of observations in Panel A indicates the number of month-station observations in police administrative data. Number of observations in Panels B and rows 2 and 3 of Panel C indicate the number of police checkpoints in the data collected by surveyors. In row 1 of Panel C, the number of observations corresponds to the number of assigned checkpoints, and in row 4 to the number of drunken drivers stopped during the final check. The total vehicle category includes cars, trucks, motorcycles vans, jeeps, buses, auto-rickshaws, and other (mostly tractors). The lower number of night deaths observations is due to the fact that these data are not available for January and February 2012.

Table 3: Pooled Results

	Daylight		Darkness	
	(1) Accidents	(2) Deaths	(3) Accidents	(4) Deaths
Treatment during & post intervention	0.0018 (0.0045)	-0.0037 (0.0035)	-0.0069 (0.0034)	-0.0045 (0.0026)
Police Lines Team during & post intervention	-0.0008 (0.0051)	0.0033 (0.0040)	0.0029 (0.0033)	0.0012 (0.0029)
Month FE	Yes	Yes	Yes	Yes
Police Station FE	Yes	Yes	Yes	Yes
Mean of dep. variable	0.085	0.029	0.033	0.016
N	4724	4724	4724	4724

This table presents the impact of all (pooled) sobriety checkpoint interventions on the number of daily road accidents and deaths per police station, using monthly police administrative reports 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.
Standard errors in parentheses are clustered at the police station level.

Table 4: Fixed vs. Rotating Pooled Results

	Daylight		Darkness				
	(1) Accidents	(2) Deaths	(3) Accidents	(4) Deaths	(5) Accidents	(6) Deaths	(7) Serious Accidents
Fixed checkpoints during & post intervention	-0.0015 (0.0048)	-0.0053 (0.0036)	-0.0027 (0.0039)	-0.0035 (0.0027)	0.0008 (0.0048)	-0.0026 (0.0031)	-0.0010 (0.0032)
Rotating checkpoints during & post intervention	0.0055 (0.0057)	-0.0019 (0.0044)	-0.0116 (0.0038)	-0.0057 (0.0032)	-0.0115 (0.0046)	-0.0058 (0.0037)	-0.0048 (0.0033)
Fixed checkpoints post intervention					-0.0078 (0.0059)	-0.0018 (0.0033)	-0.0041 (0.0035)
Rotating checkpoints post intervention					0.0001 (0.0058)	0.0004 (0.0045)	-0.0025 (0.0040)
Police Lines Team during & post intervention	-0.0014 (0.0052)	0.0030 (0.0041)	0.0036 (0.0032)	0.0014 (0.0029)	0.0035 (0.0032)	0.0013 (0.0029)	0.0010 (0.0026)
Month fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Police Station FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mean of dep. variable	0.085	0.029	0.033	0.016	0.033	0.016	0.018
P-value of test fixed = rotating effect	0.174	0.348	0.013	0.487	0.010	0.391	0.276
N	4724	4724	4724	4724	4724	4724	3324

This table presents the impact of fixed and rotating checkpoint interventions on the number of daily road accidents and deaths per police station, using monthly data from August 2010–October 2012. The during & post intervention variable is positive for the duration of the crackdown and 90 days afterwards, and the post intervention variable is positive for the 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 were taken from police administrative reports, as collected in both treatment and control police stations.
Standard errors in parentheses are clustered at the police station level.

Table 5: Checkpoint Surveys During Intervention

	Drunk drivers and motorcyclists caught			
	(1)	(2)	(3)	(4)
Rotating checkpoint station	0.130	-1.001		
	(0.189)	(0.519)		
Frequency		-0.519		
		(0.199)		
Rotating checkpoint × frequency		0.586		
		(0.235)		
Weeks of checking			-0.102	
			(0.027)	
Rotating checkpoint × weeks of checking			0.093	
			(0.037)	
Number previous checkpoints				-0.042
				(0.012)
Rotating checkpoint × number previous checkpoints				0.033
				(0.017)
Police Lines team	1.118	1.138		
	(0.233)	(0.216)		
District FE	Yes	Yes	Yes	Yes
Police Station FE	No	No	Yes	Yes
Mean of dep. variable	1.257	1.257	1.257	1.257
N	852	852	852	852

This table reports the intensity and dynamic effects of the crackdown on the number of drunken drivers caught. 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, and the pre-intervention accident rate. Standard errors in parentheses are clustered at the police station level.

Table 6: Drunk Drivers Caught on Final Check

	Rotating checkpoints			Fixed Checkpoints		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	-1.229	-1.250	-0.447	-0.628	-1.301	0.560
	(0.473)	(0.799)	(1.269)	(0.480)	(0.574)	(1.463)
Days since last checkpoint		0.001	-0.057		0.033	-0.007
		(0.018)	(0.033)		(0.009)	(0.039)
Frequency			-0.340			-0.954
			(0.368)			(0.541)
Days since last checkpoint×frequency			0.027			0.021
			(0.012)			(0.020)
Police lines teams	0.427	0.429	0.579	0.750	1.014	1.124
	(0.458)	(0.447)	(0.479)	(0.332)	(0.279)	(0.248)
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Control mean	1.953	1.953	1.953	1.953	1.953	1.953
N	77	77	77	74	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 1–3 compare rotating checkpoint stations with controls, and columns 4–6 compare fixed checkpoint police stations with controls. Outcome data were collected by surveyors sent to monitor the final checkpoints. All specifications include controls for whether the police station is located on a major highway and the pre-intervention accident rate.
Robust standard errors are in parentheses.

Table 7: Passing Cars and Motorcycles

	Road 1 (1)	Road 2 (2)	Road 3 (3)
Fixed checkpoint station	13.604 (149.461)	251.425 (186.383)	
Fixed checkpoints × frequency	-89.784 (49.064)	-113.381 (58.054)	
Rotating checkpoint station	-293.102 (136.710)	-356.204 (166.767)	-235.107 (167.547)
Rotating checkpoints × frequency	45.328 (46.596)	38.790 (62.735)	34.811 (70.974)
Fixed checkpoints × days since last checkpoint	2.244 (4.338)	-11.690 (5.125)	
Rotating checkpoints × days since last checkpoint	7.367 (5.774)	1.441 (3.178)	-0.639 (4.648)
Police Lines Team	28.914 (70.814)	164.201 (94.481)	180.596 (133.028)
District FE	Yes	Yes	Yes
Control mean	628.4	610.2	614.2
N	1540	736	568

All columns report outcomes on the number of cars and motorcycles observed passing in different checkpoint locations (including locations where checkpoints never actually occurred, as in the fixed intervention on road 2). Data were 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 are in parentheses.

Table 8: Structural Parameters

Preferences			Beliefs		
(1)	(2)	(3)	(4)	(5)	(6)
Utility of drunken driving on road 1	d_1	-0.0736 (0.0674)	Prior on crackdown ending probability	$\eta_0 = \frac{\alpha_0^\eta}{\alpha_0^\eta+\beta_0^\eta}$	0.207 (0.0197)
Utility of drunken driving on road 2	d_2	-0.0786 (0.1093)	Strength of prior on crackdown ending probability	$\alpha_0^\eta + \beta_0^\eta$	1918.7 (3071)
Utility of drunken driving on road 3	d_3	-.2558 (0.0695)	Prior on road 1 intensity	$\lambda_{1,0} = \frac{\alpha_{1,0}^\lambda}{\alpha_{1,0}^\lambda+\beta_{1,0}^\lambda}$	0.145 (0.0478)
Disutility of encountering checkpoint	c	72,130 (63,014)	Prior on road 2 intensity	$\lambda_{2,0} = \frac{\alpha_{2,0}^\lambda}{\alpha_{2,0}^\lambda+\beta_{2,0}^\lambda}$	0.119 (0.0235)
			Prior on road 3 intensity	$\lambda_{3,0} = \frac{\alpha_{3,0}^\lambda}{\alpha_{3,0}^\lambda+\beta_{3,0}^\lambda}$	0.0618 (0.0251)
			Strength of prior on road 1 intensity	$\alpha_{1,0}^\lambda + \beta_{1,0}^\lambda$	$2.36 * 10^{-7}$ (0.0019)
			Strength of prior on road 2 intensity	$\alpha_{2,0}^\lambda + \beta_{2,0}^\lambda$	8.291 (10.046)
			Strength of prior on road 3 intensity	$\alpha_{3,0}^\lambda + \beta_{3,0}^\lambda$	0.00 (0.003)

This table displays the parameters of the structural model presented in Section 6. Standard errors in parenthesis.

Table 9: Optimal Crackdown Strategy and Duration

	Prior Warning	Duration	Pct. on road 1	Decrease in drunken driving
		(1)	(2)	(3)
Estimated beliefs	No	100 days	40%	67%
	Yes	105 days	40%	72%
Equilibrium beliefs	No	20 days	90%	97%
	Yes	>22 days	5%-95%	100%

This table displays the parameters and effectiveness of an optimal anti–drunken driving campaign that allocates 20 checkpoints over up to 120 days and across 3 locations. Effectiveness is measured by the undiscounted share of agents drinking and driving on any road over the 120 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 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. Column 1 shows the duration of the optimal campaign, and column 2 shows the share of checks allocated to road 1. The remaining checks are equally divided between roads 2 and 3. Column 3 shows 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: Simulated Drunken Driving Probabilities Across Interventions

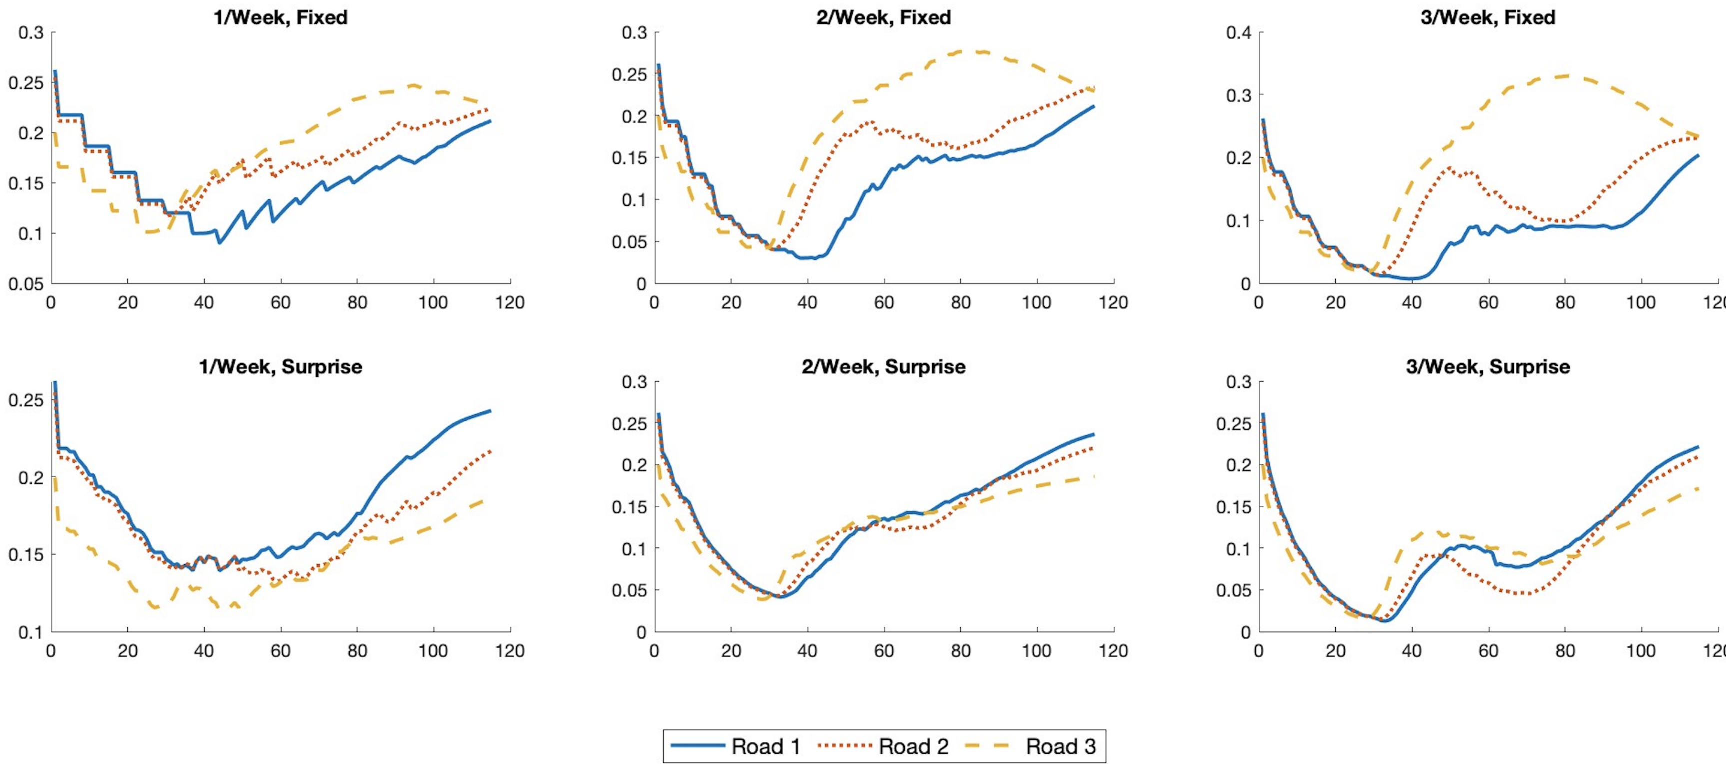

Appendix

A1 Further Reduced-Form Analyses: Robustness and Alternative Mechanisms

Pilot vs. main intervention

Our main dataset contains the pooled data from the 2010 pilot survey in 24 precincts, and the main 2011 intervention in 123 precincts. In Table A1 we estimate the specification from Table 4 separately in each intervention to verify that the results from both of these rounds are consistent with each other and with the pooled results. We find broad agreement between the two sets of estimates. In particular, for both interventions we find that the effect of rotating checkpoints on accidents and deaths in darkness is larger than that of fixed checkpoints, although this is only significant in the main intervention.

Intensity of intervention

Table A2 further decomposes the program effect by the intensity of checking. As the number of rotating checkpoints per week increases, the negative effect on night-time 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 precincts with fixed checkpoints.

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 all serve to attenuate any spillovers.

We quantify the geographical extent of the spillovers in Table A5, which examines accidents and deaths, and Table A6 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 A5 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 A6 are equally non-significant 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 not statistically significant. 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 checkpoint 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.[24]

Table A7 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 statistically 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

[24] In 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.

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 A8 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 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 non-significant) 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 roads 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 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 collecting 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 A9 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 non-significant, 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.

A1.1 Example Sequences of Individual Driver Choices and Outcomes

Drivers' learning behavior can generate complex patterns of travel and staying home. 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 $h(x_t, \epsilon_t; \Psi_0)$. To illustrate, Figure A1 shows the simulated histories of two drivers, the first in a fixed checkpoint station, and the second in a rotating checkpoint station. 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 sequences of checkpoints and parameters used to simulate these histories are those estimated from the data.

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.[25] 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 three 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.

A2 Missing Checkpoint Implementation Data

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 LASSO probit regression of an indicator that the police

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

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 A8), plus variables indicating checkpoints that occurred in the early, middle, or late stage of the intervention, police station fixed effects, the location of the checkpoint, days since the previous checkpoint, and interactions.[26] Using the variables selected by the LASSO algorithm, 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.

A3 Counterfactuals with Incomplete Implementation

The enforcement strategies discussed in the main body of the paper do not take into account potential partial compliance with the crackdown protocol by the police. However, as we show in Table A8, the police are less likely to carry out a crackdown when the intensity of enforcement is high (particularly in the three-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 $\Phi\left(.615 - .139\left(F^*\right)\right)$ where F^* represents the assigned station-level frequency of checking and $\Phi\left(\cdot\right)$ is the Normal CDF.[27] These strategies may be more reasonable than those that assume full compliance, particularly for policies that imply very high-intensity checking.

If we consider partial compliance (A13) the optimal strategy is basically unchanged for the first two scenarios, although somewhat less effective with a decrease in drunken driving of 25%. Again, incomplete police compliance has little effect on the overall strategy but decreases its effectiveness. Looking at row 3, accounting for partial

[26]This regression is analogous to column 1 of Table A8, with the inclusion of a variety of additional control variables as selected by the LASSO probit estimator.

[27]These parameters were estimated from the police implementation data in a probit regression analogous to column 2 of A8, except using a continuous measure of assigned frequency.

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.

A4 Appendix Tables

Table A1: Pilot vs. Main Intervention Results

	Daylight		Darkness	
	(1) Accidents	(2) Deaths	(3) Accidents	(4) Deaths
Panel A: Pilot				
Fixed checkpoints	-0.007	0.018	-0.005	-0.014
during & post intervention	(0.017)	(0.013)	(0.012)	(0.012)
Rotating checkpoints	-0.008	0.00	-0.012	-0.030
during & post intervention	(0.018)	(0.013)	(0.010)	(0.023)
Police Lines team	0.002	-0.005	-0.003	0.019
during & post intervention	(0.020)	(0.013)	(0.012)	(0.017)
Mean of dep. variable	0.082	0.026	0.036	0.016
N	1123	1123	1123	1123
Panel B. Main Intervention				
Fixed checkpoints	0.00	-0.006	-0.002	0.00
during & post intervention	(0.006)	(0.004)	(0.005)	(0.003)
Rotating checkpoints	0.011	0.00	-0.014	-0.004
during & post intervention	(0.007)	(0.005)	(0.005)	(0.003)
Police Lines team	-0.005	0.001	0.005	0.00
during & post intervention	(0.006)	(0.004)	(0.004)	(0.003)
Mean of dep. variable	0.086	0.031	0.032	0.017
N	3601	3601	3601	3601
Month fixed effects	Yes	Yes	Yes	Yes
Police Station FE	Yes	Yes	Yes	Yes

This table presents the impact of fixed and rotating checkpoint interventions on the number of daily road accidents and deaths per police station in the pilot intervention (Aug. 2010-Jan. 2011) and the main intervention (Feb. 2011-Oct. 2012). The during & post intervention variable is positive for the duration of the 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 were taken from police administrative reports, as collected in both treatment and control police stations. Standard errors in parentheses are clustered at the police station level.

Table A2: Fixed vs. Rotating, Intensity of Checking

	Daylight		Darkness	
	(1) Accidents	(2) Deaths	(3) Accidents	(4) Deaths
Rotating checkpoints×1/week during & post intervention	0.0110 (0.0100)	-0.0007 (0.0083)	-0.0099 (0.0064)	0.0035 (0.0055)
Rotating checkpoints×2/week during & post intervention	0.0025 (0.0071)	-0.0035 (0.0044)	-0.0102 (0.0043)	-0.0102 (0.0043)
Rotating checkpoints×3/week during & post intervention	0.0050 (0.0069)	-0.0013 (0.0066)	-0.0152 (0.0048)	-0.0069 (0.0034)
Fixed checkpoints ×1/week during & post intervention	-0.0058 (0.0067)	-0.0036 (0.0057)	-0.0034 (0.0045)	-0.0032 (0.0031)
Fixed checkpoints×2/week during & post intervention	0.0010 (0.0056)	-0.0043 (0.0044)	-0.0022 (0.0057)	-0.0025 (0.0040)
Fixed checkpoints×3/week during & post intervention	-0.0001 (0.0081)	-0.0088 (0.0052)	-0.0031 (0.0049)	-0.0043 (0.0035)
Police Lines team during & post intervention	-0.0013 (0.0052)	0.0033 (0.0040)	0.0035 (0.0032)	0.0017 (0.0028)
Month FE	Yes	Yes	Yes	Yes
Police Station FE	Yes	Yes	Yes	Yes
Mean of dep. variable	0.085	0.029	0.033	0.016
N	4724	4724	4724	4724

This table presents the impact of fixed and rotating checkpoint interventions on the number of daily road accidents and deaths per police station, using monthly data from August 2010–October 2012. The during intervention variable is positive for the duration of the 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 were taken from police administrative reports, as collected in both treatment and control police stations.
Standard errors in parentheses are clustered at the police station level.

Table A3: 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				0.010
× number previous checkpoints				(0.010)
Police Lines team				
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 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 correspond with columns 1–4 of Table 5, except that in Table 5 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 are clustered at the police station level.

Table A4: Substitution Patterns Across Routes

	Drunk drivers and motorcyclists caught		
	First 28 days of intervention	Full intervention without final check	All checkpoint data
	(1)	(2)	(3)
Location 1×# checkpoints at that location	-0.059	-0.043	-0.040
	(0.047)	(0.012)	(0.012)
Location 2×# checkpoints at that location	0.180	-0.092	-0.019
	(0.179)	(0.066)	(0.039)
Location 3×# checkpoints at that location	-0.742	-0.166	-0.185
	(0.518)	(0.075)	(0.072)
Location 1×# checkpoints at other locations	0.050	0.007	-0.005
	(0.094)	(0.017)	(0.015)
Location 2×# checkpoints at other locations	0.023	0.040	-0.029
	(0.133)	(0.033)	(0.014)
Location 3×# checkpoints at other locations	0.206	0.061	0.059
	(0.112)	(0.034)	(0.035)
Police Station FE	Yes	Yes	Yes
Mean of dep. variable	1.390	1.257	1.254
Observations	346	852	960

This table reports the dynamic effects of the crackdown on the number of drunken drivers caught. All outcome variables are based on data collected by surveyors sent to monitor the checkpoints. Rows 1–3 report the effect of the number of prior checkpoints at a location on the number of drunken drivers arrested at that location. Rows 4–6 report the effect on the number of drunken drivers at a given location of the number of prior checkpoints held at other locations. Standard errors in parentheses are clustered at the police station level.

Table A5: 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	0.00331	-0.00216	-0.00567	-0.00394	-0.00347
during & post intervention	(0.00372)	(0.00294)	(0.00235)	(0.00213)	(0.00317)
Treated stations in 10 km	0.00091	0.00079	0.00055	0.00004	0.00077
during & post intervention	(0.00116)	(0.00094)	(0.00100)	(0.00069)	(0.00099)
Treated stations in 20 km	-0.00017	-0.00035	-0.00122	-0.00106	-0.00122
during & post intervention	(0.00197)	(0.00151)	(0.00169)	(0.00104)	(0.00167)
Treated stations in 40 km	-0.00056	-0.00073	0.00072	-0.00074	-0.00135
during & post intervention	(0.00181)	(0.00178)	(0.00104)	(0.00102)	(0.00193)
Panel B. Spillover results by intervention scope					
Fixed checkpoints	-0.002	-0.004	-0.002	-0.003	-0.004
during & post intervention	(0.005)	(0.003)	(0.003)	(0.002)	(0.004)
Rotating checkpoints	0.008	-0.001	-0.009	-0.005	-0.003
during & post intervention	(0.005)	(0.004)	(0.003)	(0.003)	(0.004)
Fixed checkpoints in 10 km	0.006	0.001	0.000	0.001	0.002
during & post intervention	(0.003)	(0.003)	(0.002)	(0.002)	(0.003)
Fixed checkpoints in 20 km	0.000	-0.002	-0.001	-0.003	-0.005
during & post intervention	(0.003)	(0.003)	(0.003)	(0.002)	(0.003)
Fixed checkpoints in 40 km	0.001	0.000	0.000	-0.001	-0.001
during & post intervention	(0.003)	(0.003)	(0.002)	(0.001)	(0.003)
Rotating checkpoints in 10 km	-0.004	0.001	0.001	0.000	0.000
during & post intervention	(0.003)	(0.003)	(0.002)	(0.002)	(0.003)
Rotating checkpoints in 20 km	-0.001	0.001	-0.001	0.001	0.003
during & post intervention	(0.003)	(0.002)	(0.002)	(0.002)	(0.003)
Rotating checkpoints in 40 km	-0.001	-0.002	0.001	-0.001	-0.002
during & post intervention	(0.003)	(0.002)	(0.002)	(0.001)	(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 A6: Drunk Drivers Caught - Spillovers

	Main intervention				Final check	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment station					-1.021	-0.756
					(0.378)	(0.443)
Rotating checkpoints		-0.246				-0.376
		(0.246)				(0.328)
Weeks of checking	-0.063	-0.109	-0.060	-0.108		
	(0.033)	(0.043)	(0.035)	(0.046)		
Rotating checkpoints × weeks of checking		0.068		0.075		
		(0.038)		(0.038)		
Treated stations in 10 km	0.719	0.557			-0.196	-0.025
	(0.258)	(0.516)			(0.157)	(0.443)
Treated stations in 20 km	-0.452	-0.319			-0.137	0.656
	(0.265)	(0.376)			(0.332)	(0.799)
Treated stations in 40 km	-0.413	-0.789			-0.423	-0.057
	(0.199)	(0.344)			(0.284)	(0.390)
Treated stations in 10 km × weeks of checking	-0.045	-0.007	-0.024	0.019		
	(0.019)	(0.042)	(0.023)	(0.054)		
Treated stations in 20 km × weeks of checking	0.050	0.043	0.03	-0.001		
	(0.028)	(0.041)	(0.031)	(0.049)		
Treated stations in 40 km × weeks of checking	0.033	0.061	0.017	0.015		
	(0.026)	(0.040)	(0.025)	(0.041)		
Rotating checkpoints in 10 km		0.243				-0.329
		(0.945)				(0.725)
Rotating checkpoints in 20 km		-0.285				-1.494
		(0.576)				(1.102)
Rotating checkpoints in 40 km		0.742				-0.67
		(0.453)				(0.627)
Rotating checkpoints in 10 km × weeks of checking		-0.051		-0.056		
		(0.088)		(0.095)		
Rotating checkpoints in 20 km × weeks of checking		0.01		0.058		
		(0.069)		(0.074)		
Rotating checkpoints in 40 km × weeks of checking		-0.033		0.018		
		(0.058)		(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 5 and 6 for details on variables and specifications.

Table A7: 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)
Police Lines team during & post intervention				
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 crackdown and up to 90 days afterwards. All data were taken from police administrative reports, as collected in both treatment and control police stations.
Standard errors in parentheses are clustered at the police-station level.

Table A8: Implementation of Intervention

	(1) Checkpoint occurred - OLS	(2) Checkpoint occurred - Probit	(3) Checkpoint occurred - Final Check OLS	(4) Duration of check-point	(5) Number vehicles stopped
Fixed checkpoints			0.006		
			(0.106)		
Rotating checkpoints	-0.001	0.000	0.06	-1.856	-0.235
	(0.046)	(0.153)	(0.067)	(5.443)	(7.893)
Intensity 2/week	0.021	0.047	-0.009	-4.827	-16.280
	(0.049)	(0.154)	(0.065)	(4.982)	(9.455)
Intensity 3/week	-0.105	-0.328	-0.057	-11.701	-15.624
	(0.053)	(0.163)	(0.080)	(5.041)	(9.786)
Locations 2 or 3	-0.024	-0.105		-0.326	-2.786
	(0.038)	(0.132)		(4.779)	(6.090)
Police Lines team	0.284	0.921	0.337	23.653	25.390
	(0.035)	(0.117)	(0.120)	(4.439)	(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.

Table A9: Legal Prosecution of Drunken Drivers

	Cases Processed by Court		
	(1)	(2)	(3)
Rotating checkpoints	4.258	2.871	-0.615
	(4.338)	(3.997)	(3.311)
Drunk drivers caught in station		5.467	2.445
		(1.475)	(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.

Table A10: Simulated Fixed vs. Rotating Pooled Results

	Darkness
	(1) Accidents
Fixed checkpoints during & post intervention	-0.00475 (0.00054)
Rotating checkpoints during & post intervention	-0.00674 (0.00072)
Month FE	Yes
Police Station FE	Yes
Mean of dep. variable	0.03546
N	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. Controls include: fixed or rotating treatment indicators, day of week, urban/rural, national highway, police-station fixed effects and the number of prior accidents in the police station.
Standard errors in parentheses are clustered at the police-station level.

Table A11: Simulated Checkpoint Surveys During Intervention

	Drunk drivers and motorcyclists caught			
	(1)	(2)	(3)	(4)
Rotating checkpoint station	0.138 (0.187)	-1.066 (0.512)		
Frequency		-0.555 (0.197)		
Rotating checkpoint × frequency		0.625 (0.233)		
Weeks of checking			-0.108 (0.024)	
Rotating checkpoint × weeks of checking			0.101 (0.028)	
Number previous checkpoints				-0.044 (0.010)
Rotating checkpoint × number previous checkpoints				0.038 (0.012)
District FE	Yes	Yes	Yes	Yes
Police Station FE	No	No	Yes	Yes
Mean of dep. variable	1.240	1.240	1.240	1.240
N	852	852	852	852

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 5.
Standard errors in parentheses are clustered at the police-station level.

Table A12: Simulated Drunk Drivers Caught on Final Check

	All Stations		Rotating checkpoints		Fixed Checkpoints	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	-1.281	-0.917	-1.555	-1.746	-1.122	-0.512
	(0.389)	(0.747)	(0.512)	(0.553)	(0.285)	(0.861)
Days since last checkpoint	0.014	0.008	0.018	0.063	0.014	-0.012
	(0.010)	(0.021)	(0.013)	(0.024)	(0.011)	(0.029)
Frequency		-0.176		0.048		-0.325
		(0.273)		(0.216)		(0.373)
Days since last checkpoint×frequency		0.003		-0.020		0.014
		(0.009)		(0.009)		(0.015)
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Control mean	2.077	2.077	2.077	2.077	2.077	2.077
Mean treatment effect, @ freq = 2	-1.097	-1.087	-1.318	-1.355	-0.941	-0.964
P-value of mean treatment effect	0.00450	0.00556351	0.00838451	0.00598818	0.00105702	0.0004470978
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 6. 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 are in parentheses.

Table A13: Optimal Crackdown Strategy and Duration

	Prior Warning	Duration	Pct. on road 1	Decrease in drunken driving
		(1)	(2)	(3)
Estimated beliefs	No	100 days	40%	60%
	Yes	99 days	45%	66%
Equilibrium beliefs	No	37 days	90%	90%
	Yes	>22 days	5%-95%	100%

This table displays the parameters and effectiveness of an optimal anti–drunken driving campaign that allocates 20 checkpoints over up to 120 days and across 3 locations. Effectiveness is measured by the undiscounted share of agents drinking and driving on any road over the 120 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 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. Column 1 shows the duration of the optimal campaign, and column 2 shows the share of checks allocated to road 1. The remaining checks are equally divided between roads 2 and 3. Column 3 shows 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.

A5 Appendix Figures

Figure A1: Example of Agents' Simulated Actions and Beliefs