

THE EFFICIENT DEPLOYMENT OF POLICE RESOURCES: THEORY AND NEW EVIDENCE FROM A RANDOMIZED DRUNK DRIVING CRACKDOWN IN INDIA

ABHIJIT V. BANERJEE, ESTHER DUFLO, DANIEL KENISTON, AND NINA SINGH

ABSTRACT. A central question in the law and economics literature is whether, with limited resources, police should be deployed randomly over many locations or establish a predictable presence in crime “hot spots”. Critics of static hot spot enforcement have argued that this approach will only lead to crime shifting to other locations. But for short “crackdowns”, this may not apply, as it may take some time for potential law breakers to understand that the police has started a campaign, and the police may take advantage of this period to intervene intensively in the most productive location. We propose a model where criminals progressively learn about policing, and test it using a randomized controlled experiment on an anti-drunk driving campaign that we set up in collaboration with the police department in Rajasthan, where we randomized both whether the crackdown was random across 3 main routes or fixed in the best route, and the intensity of the crackdown. We find clear evidence of people learning over time that a crackdown is occurring, and strategically responding to it. Indeed, learning is quick enough that even for a short campaign the surprise checkpoints still dominate the fixed location approach. We estimate that crackdowns in surprised locations reduced night accidents in the area covered by a particular police station by 17%, and night deaths by 25% over a two month crackdown and 6 weeks following it.

1. INTRODUCTION

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

Date: First Draft August 4, 2013. Please do not cite without permission from authors.
Author affiliations: Banerjee & Duflo: Massachusetts Institute of Technology, Keniston: Yale University, Singh: Rajasthan Police.

However this is only true in the full Nash Equilibrium. The potential criminals need to know the strategy the police is following in order to undo it. In the case of a temporary crackdown by the police this may not be true. At the onset of short, intensive crackdowns, lawbreakers may not be aware that they are being policed to start with and therefore continue to frequent their most favored locations, in which case the police will stop the maximum number of them by concentrating their efforts on the most crime-prone locations. Moreover even after criminals actually find out that those locations are being policed it may be reasonable for the lawbreakers to assume that the intervention was temporary if, in their past experience, no policing was the normal state of being. In that case they will continue to go back to their favored location for a while, until they are persuaded that the checking will continue long enough that it is worth switching locations. Until such a time as this switch happens, it is optimal to continue to police the location that the law breakers have traditionally favored. It is only if checking goes on beyond this point of time that it pays switch away from this location, at least for part of the time.

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

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

This project was undertaken in collaboration with the Rajasthan police, who invested in equipment and resources to fight drunk driving but were uncertain about the right way to organize the campaign. In some areas, chosen at random, only the route judged most conducive for catching drunken drivers was chosen for the crackdown, while in others, checkpoints was randomly assigned across the three most popular routes. In yet other areas no crackdown was implemented, allowing us to evaluate the effectiveness of the crackdown itself. The intensity of the crackdown was also varied in two different ways. First, police stations involved in the intervention were randomly assigned to have 1, 2 or 3 checkpoints per week

in their jurisdictions. Second, incentives for the policemen for doing those checks effectively was also varied at random.

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

We can start to understand why when we look at the effect of the number of police checkpoints on the number of drunks who get caught each night of checking. Fixed checks are conducted at the location best suited for catching drunken drivers while surprise checks are conducted at one of the best three locations, so absent avoidance effects, there should be more drunk drivers caught for fixed checks. In fact we see the opposite: there are marginally fewer (rather than more) vehicles with a drunk driver apprehended at the fixed checkpoints. Furthermore, with fixed checks the number of drunks apprehended per night goes down with the number of past checks, with the number of weeks since the intervention started, and with the assignment to a higher weekly checkpoint frequency. However this effect entirely vanishes with surprise checks.

This finding helps us distinguish between the two reasons why more frequent checks at a fixed location reduces the number of drunks stopped at those locations. One is that people are no longer driving drunk. The other is that they are switching to alternative routes. If people are simply not driving drunk, then an increased frequency of checking should have a similar negative effect on drunks stopped with random checks as it does with fixed checks (albeit with a less sharp slope since with random checks the probability of being checked on the favored route goes up more slowly). The reason we do not see any decrease in drunken drivers caught with random checks is that the drunk drivers have no viable alternative route to switch to and do not decrease their overall propensity to drive drunk— therefore with surprise checkpoints frequency does not matter.

It is of course precisely the fact that the drunk drivers are still savvy enough to switch to other routes that gives random checks their advantage over fixed checks. If they simply stopped driving drunk then the only reason to check routes other than the most popular route would be if the response to additional checks at the favored route is subject to diminishing returns.¹

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

While these predictions are essentially static in nature, the basic trade-off that motivated this paper centers on the fact people may not know that they are being checked to start with and may take time to learn. This too is clearly present in the data. In the first two weeks of the intervention, more drunks are caught with fixed checking when checking is relatively infrequent. This is consistent with the fact people have not yet learnt enough to want to switch and therefore the favored route is the best place to check.

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

The remainder of the paper is organized as follows: section 2 provides some background on drunk driving and law enforcement against drunk driving in India. Section 3 describes our experiment. Section 4 contains the model. Section 5 present results and section 6 concludes.

2. BACKGROUND: ENFORCING DRUNK DRIVING

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

Sobriety checkpoints have been evaluated by a number of studies in a wide variety of contexts, and the general consensus is that these checkpoints significantly reduce

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

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

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

The key to understanding the tradeoff between the increasing returns of concentrating enforcement and the diluting effects of displacement may be the speed at which criminals learn about police crackdowns. Nevertheless, robust evidence on the speed of criminal learning is lacking. Sherman (1990) argues that many of the non-experimental drunken driving crackdowns show post-project decreases

in accidents lasting even longer than the projects themselves. Similarly, the Minneapolis hot spot experiment shows some non-experimental results that hotspot level disorder remained lower a few hours immediately after the police visited and, on an aggregate level, during 6 weeks when the police lowered patrolling intensity. However, other studies find that program effects are short-lived and disappear even while the intervention is ongoing (Braga, 2001). Thus the ability of lawbreakers to adjust to police behavior—both in terms of reducing or shifting criminal behavior during a crackdown and in terms of returning to criminal behavior after a crackdown—remains a central and open empirical question in the field of criminal behavior.

2.1. Drunken Driving Enforcement in India. In India, highway safety laws of all kinds are generally enforced by fixed sobriety checkpoints manned by personnel from the local police station using a “selective breath checkpoint” methodology. Barriers are arranged on the roadway so that passing vehicles are forced to slow down, and the officers on duty signal selected vehicles to pull over as they pass through the barriers. If the roadblock is intended to target drunken driving, police personnel then ask the driver a few questions on his or her identity, destination, etc., while observing the driver’s demeanor and smelling his or her breath. If the police feel the driver may be drunk, then according to the official procedure they will order him or her to blow into a breathalyzer, following the results of which the driver is either charged or released². The printed results of a handheld breathalyzer are considered sufficient proof of drunkenness in court. Once caught, drunken drivers’ vehicles are confiscated by the police, and the driver must appear in court to pay a fine or potentially face jail time, although imprisonment is never observed in our data. The fine amount depends largely on the judgment of the local magistrate, with a maximum fine of Rs. 2000 (roughly \$50) for the first offense. The driver must then return to the police station to recover the vehicle from the police lot.

Even this official procedure leaves many factors undetermined and up to the discretion of the police manning the roadblock. The choice of how many, and which, vehicles to pull over for questioning and potential testing is the most important. Ideally the police should target vehicles with the highest probability of drunkenness, and in conversations the police often noted that if they saw a vehicle with a family, or driven by elderly people they assumed that the driver would be unlikely to be drunk and let it pass. Other considerations may also enter the decision: police may be less likely to stop vehicles with many passengers who would be difficult to deal with at an isolated police station if the vehicle were impounded. Similarly, the police might be hesitant to stop luxury vehicles whose owners might be well

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

connected and cause problems if subjected to breath testing. Unscrupulous police officers are faced with another decision: whether to follow the official ticketing procedure or instead to solicit (or accept) a bribe from the accused drunken driver.

Most importantly, to understand our results, it is important to note that roadblocks of this type occurred extremely rarely prior to the intervention and in the control police stations during the intervention. This rarity is primarily due to the fact that breathalyzers had not been widely distributed to police stations prior to the program, and without a breathalyzer the local police would have needed to take a suspect to the hospital for blood testing in order to be reasonably certain of securing a conviction under the 30mg/100ml threshold for drunken driving established by the Indian Motor Vehicles Act. This procedure is highly inconvenient for the police (who typically have a single vehicle per police station), and seems to have prevented enforcement at all except urban police stations in Jaipur, the state capital. While breathalyzers were distributed in treatment and control prior to the program, control police stations hardly ever used them to carry out roadblocks. In the 925 nights that surveyors visited control police stations, on only 7 (.76%) occasions did they witness the police carrying out a roadblock.

3. THE INTERVENTION

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

First, the location of the roadblocks was randomized. Roadblocks occurred at either the same spot on every occasion, selected by the local chief of police as the location best suited to preventing accidents due to drunken driving or one of a group of three locations, with each night's location chosen at random. The three spots were chosen by the local chief of police as the best three suited to catching drunken drivers. Second, the frequency of the roadblocks was also randomized. Roadblocks were carried out in the police station jurisdiction either 0 (for control stations), 1, 2, or 3 nights per week. Third, the incentives of the personnel carrying out the roadblock. Roadblocks were staffed by either police officers from the local police station (the status quo outside the intervention) or a dedicated team selected from the police reserve force at the district level. These special teams were monitored by GPS devices installed in their vehicles. The design is summarized in Table 1, and in the remainder of this section we explain the design choices in detail.

The program took place in two phases, an initial pilot, from September-early October 2010, and a larger rollout from September to the end of November 2011. The initial pilot covered 2 districts and 40 police stations, and the second covered 10

districts and 183 police stations. Treatment status was assigned randomly, stratified by district, whether a station was located on a national highway, and total accidents between 2008-2010. The assignment of police stations to treatment groups during both rounds of the intervention is reported in Table 1. In addition to the stations reported on this table, 16 and 60 were designated as control stations during each treatment round in which no additional sobriety testing occurred. The 2010 and 2011 interventions were identical in implementation, with the exception that during the pilot all roadblocks occurred twice a week and the program lasted only slightly more than one month. In the analysis and results stated below, we combine data from both intervention periods and control for any time trends using month fixed effects.

3.1. Roadblock locations. To test the central hypothesis that surprise checks lead to more crime reduction than fixed checks, and the implicit mechanism of learning by potential and actual drunken drivers, we randomly assigned police stations to hold their checkpoints at either a single location, or a rotating set of three locations. In the single location group, the police station chief identified the best location in the station's jurisdiction for catching drunken drivers, and the checkpoints were carried out by either local staff or dedicated police lines teams at that location. Fixed checkpoints were carried out at the same place, same time, and to the greatest degree possible, on the same day every week, although scheduling difficulties occasionally made this impossible. In contrast, surprise checkpoints rotated among the three best locations for catching drunken drivers, again as identified by the police station chief. Each police station's rotation was pre-determined in advance by the research team. The differences among locations, in particular that the third best location usually has far fewer passing vehicles than the first, affect our analysis of these two program options. In regressions estimating the overall impact of the different strategies we do not control for checkpoint location fixed effects or characteristics, since these are themselves an outcome of the program choice. When analyzing learning by the public, however, we do include these controls to see how public behavior at a specific checkpoint responds to the enforcement strategy.

3.2. Roadblock frequency and timing. The variation in roadblock frequency was designed to identify the shape of the relationship between police enforcement and driver behavior. Discussions with the police determined that it was not feasible to carry out more than 3 roadblocks per week per police station, thus giving us the final randomization categories of 1, 2, or 3 roadblocks per week (and, of course, 0 in the control group). These frequencies were at the police station level, not the road level; for example in a surprise group police station with a frequency of 2

checkpoints per week each of the three roads would have a roadblock twice every three weeks. Checkpoints were always held at 7:00pm-10:00pm in the evening.

3.3. Roadblock personnel. Previous work with the Indian government (Duflo, Banerjee and Glennerster 2010) and the Rajasthan Police (Banerjee, et al. 2012) suggests that the implementation of government initiatives often decreases dramatically in the medium term if the civil servants implementing the project are not sufficiently motivated. To gain further insight on the role of monitoring and motivation in project implementation, as well as to guard against a failure of the project due poor implementation, the anti-drunken driving campaign was implemented by two different sets of police staff, with different motivation, monitoring, and characteristics. We henceforth refer to these groups as "station teams" and "lines teams". The station teams consisted of the staff of the police stations under whose jurisdiction the checkpoint sites fell—the status quo for special crackdowns of this type in the Rajasthan Police. For these personnel, manning the checkpoints was an additional responsibility on top of their existing duties. They were monitored by the existing police hierarchy, which consisted of a single supervisory officer assigned to a group of roughly 5 police stations, and a district level officer monitoring the implementation of the project in all police stations within the district. These district level officers were limited in their ability to verify that checkpoints were actually being carried out, since many of the police stations were at least an hour's drive from the district headquarters.

The second group of staff implementing the project were drawn from the Police Lines, a reserve force of police kept at the district headquarters who normally perform crowd control and VIP security duties. The police lines are often considered a punishment posting in the Indian police system, since a position in the police lines effectively removes an officer from contact with the public. Because of the undesirable nature of the police lines, these staff were potentially a group for whom positive incentives, specifically a transfer out of the police lines, could be provided relatively easily. However, an objective measure of performance was required to make these incentives more credible. We generated this measure by installing GPS tracking devices in the police vehicles allocated to the police lines staff to travel from their barracks at the district headquarters to the checkpoint locations. These devices provided a real-time display of the vehicle's location to the district-level supervisory officers via a simple internet interface, as well as stored the vehicles' travel routes for future examination. The online interface also recorded the usage of the system by the district supervisory officers, providing a proxy for the degree of monitoring by district. Police lines teams were told of the installation of the GPS devices, and were informed that good performance on this assignment might

improve their chances for transfer out of the police lines. The police lines teams differed from the station teams in their incentives, their monitoring, and perhaps also in their fundamental motivation and ability as police officers.

Disentangling the various effects of these differences is a challenge given the limitations of the program design, but certain features of the setup do allow for some insight. First, the lines teams were monitored only on their attendance at the roadblocks, but not on their conduct once there. Thus any differences once at the roadblock, for instance in terms of vehicles stopped, cannot be directly attributable to the GPS devices. Second, conversations with senior police officials suggest that selection of police staff into the Police Lines is, if anything, negative. Police constables tend to be transferred to the police lines if they have shown themselves to be incompetent in a police station. Hence if we find that lines teams outperform station teams it may be despite, not because of, their inherent aptitude for policing.

4. THEORY

If the drivers knew the strategy adopted by the police and were free to react to it, their problem would be a static choice of whether to cease drunken driving, continue to drink but avoid the checkpoints, or drive drunk on their favored routes in spite of the checkpoints. However, the assumption of complete information is an unlikely one in this context: potential drunken drivers did not know that there will be such an intervention and even after it started there was no official announcement. So people could only learn about it from their own experience or that of others. They also did not know how long the intervention was meant to continue and were not told when it ended. Drivers could only make inferences about the duration of the program and whether it had ended based on their individual and collective experiences.

If we assume that the police pre-committed its strategy and that drivers started from priors about this strategy and then progressively discovered the actual strategy based on their own experience, we have a variant of a “bandit” model, where the novel element is that at some point the payoffs from taking the different routes revert to their pre-intervention levels.

4.1. Model setup. Consider a driver’s problem faced by the possibility of being stopped by the police. There is a population of drivers who drive from the local bar or liquor store to their homes every period. There are two routes back from the bar: A and B. A fraction $\phi > 1/2$ of the drivers prefer route A while the rest prefer route B. The value of making the trip by the preferred route is V_H , while

that by the non-preferred route is $V_L < V_H$. Drivers are forward-looking expected utility maximizers with a discount factor $\delta < 1$.

Drivers are drunk at the time of driving but not so drunk as to be unable to make rational calculations. Drunk drivers get into accidents and the expectation of the private cost from the possibility of accidents to the driver is C_p . There may be in addition a population of sober drivers but we assume that being checked is costless for them and therefore they never change their driving behavior.

The reason for police intervention is that accidents have externalities (other people get hurt, traffic jams) and the expected social cost from the possibility of an accident resulting from a single drunk driver is $C_s > C_p$. Assume that $V_L - C_p > 0$, so that the private cost by itself is not enough to discourage drunk driving.

This is why drunk driving is illegal and drunk drivers who get stopped by the police are prevented from driving till they are sober and in addition are subject to a fine. The disabling effect—drunken drivers' vehicles are confiscated for the night—directly prevents accidents. The deterrent effect of the expectation of the fine and vehicle confiscation has an incentive effect on potential drunk drivers. Assume that the total utility cost to the driver of the fine and the disabling is F , and F is large enough that $V_H - C_p - F < 0$: if the drunk driver was sure of being stopped by the police he would opt to not drive drunk (we call this staying home).

The key elements of this model are the drivers' evolving beliefs about the nature of the police enforcement. These beliefs are formed about three states of the world: checking locations, checking intensity, and the permanence of the enforcement. First, drivers do not know whether the checking is occurring on both routes, or only on route A. Second, they do not know the intensity of the checking, and hence their probability of apprehension. Finally, they are unsure whether the enforcement campaign is permanent or temporary.

However assume that, as is the case in the context of Rajasthan, historically there have been very few occasions when drivers were stopped and as a result, people start from the belief that with probability $\sigma \simeq 1$, they will not be checked. Nevertheless, drivers allow for the possibility that with probability $1 - \sigma$, things will change and there will be some checking.

To limit the number of cases we assume that all drivers have the same beliefs about the nature of the checking and the choice of these beliefs loosely mimics the design of our experiment. Specifically we assume that people know when checking would start if it did start. Designate this to be time zero. Moreover they believe that either there is checking only on route A or checking is equally likely on both routes A and B. Let θ be the probability of the checking being concentrated in location A conditional on their being checking.

Drivers are also uncertain about the frequency of checking. The ex ante probability of being checked in any period given that there is checking is λ_0 with probability π or $\lambda_1 > \lambda_0$ with probability $1 - \pi$. Each driver has the same probability of being checked.³

Even in that case there is a strong possibility that the enforcement campaign may be just short-lived. We capture this by assuming that there are two states of the world which we will label T and P (for temporary and permanent). In one the campaign has a probability α of ending every period. In the other it goes on for ever. The prior probability that we are in the first state is β . Drivers assume that when the campaign ends, it will end at all locations at the same time.

All those drunk drivers who are currently on the road when a check happens are assumed to be disabled and fined. We assume that the fine and its frequency is high enough to get the driver to change routes: Specifically if a driver who prefers A becomes sure that there could be checking in A in a given period (even if it only were at the low rate λ_0) but no checking in B, he would switch to B. In the notation of the model this implies

$$V_H - \lambda_0 F < V_L$$

4.2. Analysis. The structure we imposed on the problem simplifies the analysis considerably. Since there is always checking in A when there is checking on B, no one who prefers B switches to A. And anyone who prefers A, switches to B and discovers that there is checking in B, either stops driving or switches back to A.

4.2.1. The choice between A and B. Let us begin by focusing on the switch between the two routes by assuming that no one stops driving—this will be true when driving always pays off even when there is high frequency checking, i.e. $V_H - \lambda_1 F > 0$. Now consider the decision to switch from A to B for someone who has never switched before. Suppose that he is checked in period S and opts not to switch in the next period. Now consider his choice in some subsequent period S' assuming that the last time he was checked was S . The lack of checks after S increases the likelihood that checking has stopped and reduces the probability that checking is high frequency. Both increase the incentive not to switch. Therefore if the driver does not switch in a period where he is checked he will not switch in a subsequent period without additional checking. Therefore we can focus on periods where there is checking.

Consider any driver who prefers route A and has only driven there, at the point where he has just been checked at time S and has been checked m times in total. To

³This way of modeling checking has the implication that a driver may be checked multiple times during a single trip. However the likelihood of being checked multiple times can be minimized by setting a small enough value for λ .

keep notation manageable, denote the prior probability of high checking intensity at time S as $\Pr[\lambda_1]_m^S = \pi \lambda_1^m (1 - \lambda_1)^{S-m}$ and similarly the prior probability of low intensity checking as $\Pr[\lambda_0]_m^S = (1 - \pi) \lambda_0^m (1 - \lambda_0)^{S-m}$. Then the driver's posterior belief that the checking is temporary is:

$$\begin{aligned} \Pr[T|S, m] &= \frac{\beta(1 - \alpha)^S \left(\Pr[\lambda_1]_m^S + \Pr[\lambda_0]_m^{S,m} \right)}{\beta(1 - \alpha)^S \left(\Pr[\lambda_1]_m^S + \Pr[\lambda_0]_m^S \right) + (1 - \beta) \left(\Pr[\lambda_1]_m^S + \Pr[\lambda_0]_m^S \right)} \\ &= \frac{\beta(1 - \alpha)^S}{\beta(1 - \alpha)^S + (1 - \beta)} \end{aligned}$$

and her belief that the checking is high intensity is:

$$\begin{aligned} \Pr[\lambda_1|S, m] &= \frac{\Pr[\lambda_1]_m^S (\beta(1 - \alpha)^S + (1 - \beta))}{\beta(1 - \alpha)^S \left(\Pr[\lambda_1]_m^S + \Pr[\lambda_0]_m^S \right) + (1 - \beta) \left(\Pr[\lambda_1]_m^S + \Pr[\lambda_0]_m^S \right)} \\ &= \frac{\Pr[\lambda_1]_m^S}{\Pr[\lambda_1]_m^S + \Pr[\lambda_0]_m^{S,m}} = \frac{\pi \lambda_1^m (1 - \lambda_1)^{S-m}}{\pi \lambda_1^m (1 - \lambda_1)^{S-m} + (1 - \pi) \lambda_0^m (1 - \lambda_0)^{S-m}} \end{aligned}$$

It is easy to check that $\Pr[T|S, m]$ only depends on S and that these two events are independent of each other. A higher S increases the probability that the campaign is permanent and an increased m increases the probability that there is high frequency checking ($\lambda = \lambda_1$).

Consider now the decision of the same driver to switch to route B. Every additional time he is checked in route A both state P and state λ_1 become more likely and the value of switching to B goes up. This is because the advantage of B is that with probability $1 - \theta$ there is no checking there. The worse A looks in terms of frequency and permanence of checking, the bigger the option value of finding out whether there is checking in B.

Claim 1. For any $\theta < 1$ there exists a $\underline{\delta}$ close enough to one such that any driver who prefers A and has $\delta > \underline{\delta}$ and has been checked over a long enough period will at least temporarily switch to B and the switch back if he tested.

Proof. A driver who has been checked over a long enough period will put probability close to 1 that the enforcement campaign is permanent and that he knows λ (we make the assumption that is less favorable switching, i.e. $\lambda = \lambda_0$). By our assumption that $V_H - \lambda_0 F < V_L$, it is therefore strictly optimal for him to switch to B if he were certain that B was checking free. The potential for loss from switching comes from the fact that with probability θ there is checking in B as well in which case he is just driving on his less preferred route and gaining nothing. The option of switching to B and then switching back to A if there is checking in B is preferable to remaining in A forever if:

$$(1 - \theta) \frac{V_L}{1 - \delta} + \theta \sum_{s=0}^{\infty} (1 - \lambda_0)^s \lambda_0 \left(\frac{1 - \delta^{s+1}}{1 - \delta} V_L - \delta^s F + \frac{\delta^{s+1}}{1 - \delta} (V_H - \lambda_0 F) \right) > \frac{V_H - \lambda_0 F}{1 - \delta}$$

From this equation, the net gain from switching to road B can be written as:

$$\frac{\theta \lambda_0}{1 - \delta (1 - \lambda_0)} \left(-F + \left(\frac{1 - \delta (1 - \lambda (1 - \theta_0))}{\theta \lambda_0} \right) \frac{V_L - V_H + \lambda_0 F}{1 - \delta} \right)$$

It is evident from this expression that as δ goes to 1, the net gain eventually becomes positive. Hence the result. \square

Of course there is no guarantee that switching only when the driver is more or less certain of permanent checking in A is the optimal strategy: It may be better to switch earlier, before being almost certain that checking in A is for ever. In that case even if there is no checking in B, there would be reason to switch back to A after a while to find out if checking continues.

However it is easy to check that a driver who has only driven on A has a cutoff rule $m^*(S)$, where S is the period of his last check such that if m is the number of checks experienced during period S , he will switch if and only if $m > m^*(S)$. When there is no possibility of switching we set $m^*(S)$ to be S . This follows from the fact that S pins down the probability that checking is permanent and conditional on that, a higher m always means more frequent checking and hence a bigger incentive to switch. From this and the fact that the distribution of m under λ_1 first order stochastically dominates the distribution of m under λ_0 , it follows that:

Claim 2. If the true state is λ_1 the driver who starts in A is more likely to have made his first switch to B before any given point of time after the program starts compared to the case where the true state is λ_0 .

Now consider what happens after the switch to B. Compare two drivers who have the same exact histories in route A but in one case the true state of the world is (λ_1, T) and in other the true state of the world is (λ_0, T) . Given that they face the same histories they will switch from route A at the same time and with the same posteriors.

After they have switched to B we permit their histories to diverge: If they never get checked, they will of course continue to have the same histories and the same beliefs. Therefore they will switch back to A in the same period if they are never checked. Let S_A be that period. However the probability that they will be checked before S_A is different in the two states of the world. Let the number of periods after the switch to B until the driver is first checked in B be denoted by S_c . The expected number of periods until the switch back to A after the switch to B is given by

$$\sum_{S_c=1}^{S_A-1} Prob[S_c < S_A | \lambda_1, T] S_c + (1 - Prob[S_c < S_A | \lambda_1, T]) S_A$$

in state (λ_1, T) and

$$\sum_{S_c=1}^{S_A-1} Prob[S_c < S_A | \lambda_0, T] S_c + (1 - Prob[S_c < S_A | \lambda_0, T]) S_A$$

in state (λ_0, T) . Since there is more frequent checking in state (λ_1, T) , the expected time before the switch is less in that state.

The exact same argument can be made when checking is permanent.

Claim 3. Comparing two drivers who had the exact same experience in terms of being checked on route A in two different locations with differing probabilities of checking, and who have therefore switched to route B at the same time, we expect the first switch-back to A happen earlier where checking is more frequent.

Of course it is not very likely that two drivers in two different locations with differing probabilities of checking will have exact same experience in terms of being checked on route A. The typical driver who switches from A to B when checking is less frequent will do so based on a rather different experience. He is likely to be more persuaded of the permanence of checking than his counterpart in a location where checking is more frequent, and more likely to believe that checking is relatively infrequent. If he is checked he will of course switch back, and this will happen sooner when checking is more frequent. However if he is not checked, either because he is lucky or because checking has stopped, it is not entirely obvious how initially being exposed to λ_1 rather than λ_0 affects the speed of switching back to A. Someone who encountered frequent checking in A when the checking initially started (i.e someone who started in λ_1), is likely to have switched even if he thought that checking would end soon just because of the contemporaneous cost of being checked. By contrast someone who was only checked occasionally (i.e who started in λ_0) may have only switched when he was relatively sure that checking will continue for a long time—as long as he thought that checking has a reasonable chance of being gone by next period he would not have bothered to switch. The absence of any checking in B means that both still retain the priors about the state of the world that they had when they switched. Therefore on one side the person from λ_0 will be more reluctant to switch than the person from λ_1 because he expects checking to continue, but on the other side switching is also less costly for him because he expects infrequent checking. Therefore the effect of a higher λ on the speed with which the driver switches back is potentially ambiguous.

However there are two clear takeaways from this discussion. First:

Claim 4. High frequency checking in both locations, may lead more drunks to be driving on route A compared to route B, at least for some time

This is because the switch back can be (and is probably likely to be) sooner with high frequency checking.

Second, it should also be clear from this discussion that the switch back could be very slow or even never happen: If the driver was sufficiently persuaded of permanence of checking in A before he switched, he may never want to try route A again.⁴

Claim 5. There is no guarantee that the driver who prefers A will eventually switch back to A after checking stops. More generally the switch back can be slow and the effects of the program may continue long after its over.

4.2.2. *The decision to stop.* Now suppose $V_H - \lambda_1 F < 0 < V_H - \lambda_0 F$. A driver who is sufficiently confident that there will be high frequency checking in A in the next period will prefer to stay home rather than drive on A. However he still has the option of switching to B and as long as he is sufficiently patient an argument very similar to that for Claim 1 shows that it is always better for him to try out B before choosing to stay home. It is only after he encounters checking in B as well that he will opt out of driving.

Once he stops driving he gets no new information. However given that checking might stop on its own, he will eventually want to give A another try as long he is not too impatient. If he still finds checking in A he will stop for a while, and then perhaps go back and try again and so on (though eventually he will just stop, because the probability of state P will be close to 1).

By contrast, with low frequency checking, given our assumption that $0 < V_H - \lambda_0 F$, he will never stop driving so the optimal strategy will be as in the previous sub-section: he will eventually switch from A to B and then if there is also checking in B, he will most likely get checked there and switch back to A. Otherwise he will continue to drive on B, and occasionally switch back to A to verify whether checking continues (as long as he is sufficiently patient) .

Hence with a high frequency checking in both locations we will observe a relatively rapid exit from A (the exact argument for this is similar to one made in the previous sub-section) followed by a switch to no driving. Once this happens, drivers will stop learning about the state of checking except to the extent that they take drive from time to time specially for this purpose. By contrast with low frequency checking in both locations the exit from A is slower but the eventual outcome has the drivers back in A driving. Hence the end of checking has no effect with low

⁴This is two-armed bandit property.

frequency checking in both locations but will generate a potentially quite slow increase in drunk driving with high frequency checking in drivers preferred locations (slow because many drivers have temporarily stopped driving and therefore do not learn).

Claim 6. The end of checking might only affect drunk driving where there was high frequency checking on both routes and even there the increase in drunk driving may be very gradual.

4.2.3. *Driving sober.* So far we have not allowed the drunks to remain sober—if they want to not be caught, they have to stay home. This is not unreasonable—perhaps the reason why people drive long distances while drunk is because they went to a party and may be they would not want to attend and not drink. However it is also possible make sure that the driver stays sober so that the rest can enjoy the party.

Let us now allow drivers to drive sober on their favored route and assume that it gives them utility of zero. Also assume that sober drivers can observe whether others are being checked and react accordingly. Allowing drivers to choose to be sober changes one important thing: they are now driving rather than staying home and therefore they will learn faster when checking stops.

Claim 7. The rate of increase in drunk driving after checking stops will be much faster if drivers continue to drive while staying sober when they believe they have a high enough chance of being checked.

4.2.4. *Other extensions.* Nothing essential changes if we assume that $V_H - \lambda_0 F < 0$, so that people want to stay sober (or stay home even with low frequency checking. In this case there is an increase in the number of drunk drivers once the program stops, even with low frequency checking, but, as before, the time trajectory will be flatter with low frequency checking because it is harder to learn that low frequency checking has stopped.

Another possible extension is to relax the assumption that checking is either permanent or follows a Poisson process. Perhaps a more realistic formulation would be to assume that it follows a Pareto process, so that checking is always expected to end even though it may take some time to do so. However the Poisson case is more tractable and the results should be very similar in the two cases.

The introduction of social learning in the sense of people talking to each other about their experiences, will tend to dampen all the effects we have emphasized because learning will be much faster and not so reliant on the individual's own experience.

4.3. **Implications for Policy.** How should a police department with limited checking resources deploy them to best reduce the welfare losses from drunk driving?

Note that there are two effects of checking—a disabling effect (the drunk driver is stopped before he gets into an accident) and an incentive effect (the driver stays home or stays sober).

Suppose the police only checks route A. Then drivers will eventually discover that there is no checking on B, and use that route. We will see some welfare gains because of the disabling effect in the early periods, but these would get dissipated over time as drivers switch away from route A and in the long run we would see no effect.

Alternately they could check both routes but with a frequency that is half of what it was in the previous case. In this case there will either be a long run disabling effect (drivers drive on their preferred routes but get stopped from time to time) or a long run incentive effect (drivers stay sober or stay home) and possibly both. However since most drivers prefer route A, the initial disabling effect would be weaker because checking is less frequent.

The trade off is therefore between a stronger initial effect from doing all the checking in A and the stronger long run effect of checking both A and B. Given that the police knew that checking would not last for ever, the choice between these two options depends on how fast drivers learn about checking. If they learn fast and switch quickly then checking both places is better. How fast they learn, in turn depends on how frequently they drive (relative to the duration of the checking program) and their expectations about how long checking will last. If they expect it to end very soon, they may not bother to switch.

Our empirical analysis will start by looking at the impact of checking just A relative to checking both at random on welfare outcomes such as drunks stopped and deaths avoided. We will then check whether the result from this reduced form comparison accords with what we see in the data about the speed of learning.

5. DATA

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

5.1. Administrative data. This study’s main results on accident and death rates are drawn from accident reports recorded by the police. For each accident on which data has been collected properly we know the police station, date and time of the incident, whether it was “serious” or not (according to the police’s judgement, the criterion used here is subjective), the number of individuals killed or injured, and the types of vehicles involved. We also have some additional information on the

location, including whether it occurred on a highway. Unfortunately we do not know whether drunken driving contributed to the accident.

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

Summary statistics are presented in panel A of Table 2, with statistics presented for control stations. The data, displayed at the police station/month level, shows that control police stations have roughly 0.12 accidents per day and 0.05 deaths. Of these, roughly 1/3rd occur at night. For lack of direct measure of accidents caused by drunkenness the number of night accidents and deaths may provide an outcome that varies more with the level of drunkenness than the overall total. Court records provide a secondary source of administrative data on the project outcomes. For each police station, the court provided a list of all drunken driving cases heard, including the date of the offense, level of the fine, and breathalyzer reading. Finally, the breathalyzers themselves provide a source of data on the actions of the police at the roadblocks. Each time that a breathalyzer is used, it records the date, time, and alcohol level measured. At the end of the project we copied the memory files off the breathalyzers and can use these to examine police conduct. Unfortunately these records are partial, since many of the breathalyzer memories were wiped when their batteries ran out. Since few, if any, of the police personnel knew of the breathalyzer memory recording feature these erasures are likely to be random.

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

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

The summary statistics of the data collected by these monitors is displayed in panels B-D of Table 2. Panel B displays the number of vehicles passing by the check points on the average night, using surveyor counts from the locations identified by the control police stations as where they would have carried out the checkpoints. Fewer vehicles pass the second and third checkpoint locations than the first; this is particularly noticeable in the medians. Overall, police stopped 13.1% of passing vehicles, roughly 105 per checkpoint. The majority of these were motorcycles, of which 11.5% or 40 per checkpoint were stopped. Panels D and E report the number of drunken drivers caught by the police on nights when normal checkpoints were carried out as well as on the final checking night when checks were conducted in a more systematic fashion. On average police caught 1.85 drunk drivers per roadblock, primarily motorcycles. This corresponds to a 2.23% overall drunkenness rate, and a rate of 3.36% for motorcyclists. Car drivers had substantially lower drunkenness rates, perhaps partly due to the fact that many cars in India are driven by professional chauffeurs whose employers would not tolerate drunkenness.

6. RESULTS

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

We begin with a simple summary of the effects of the program on accidents in Table 3. Table 3 regresses the number of accidents and deaths in the area covered by the police station (note that it will include, and likely be dominated by, the 3 main roads on which check points were to happen in surprise stations) during daylight, nighttime, or all on an interaction between a treatment station dummy (any treatment) and a dummy for during or post (60 days post) intervention, including police station and month fixed effects. We find a significant decrease in the number of nightly accidents (1.7%) and deaths (1.7%). There is no impact during the day, which is reassuring since all the check points took place at night.

One of the main hypotheses of interest is the difference between the effectiveness of surprise and fixed check points. Table 4 separates the treatment groups in these two subgroups (keeping the number of check points and the implementation type pooled for the time being), and, as before, interact it with a “during or after intervention dummies”, controlling for month and police station fixed effects. This tables makes clear that the impacts of the intervention on accidents and death are entirely driven by the surprise check points police station. At night time, accidents reduced by 2.6% and death 2.4% in the surprise police station, and both number are significant at the 10% or 5% level. In contrast, there is no significant impacts in the surprise police stations and the point estimates are much smaller (2 to 8 times). The difference is significant for night time accidents, though not for deaths.⁵

Our learning model predicts that the difference between surprise and fixed checking should persist for some time after the intervention ends (since it is not clear it has ended) and should in fact increase over time, and particularly after the end of the intervention (since in the fixed stations it is easy to see that the crackdown has ended). Table 5 provides evidence for both claims. It shows that the difference in the number of accidents between fixed and surprised checks is significantly larger in the few weeks after the intervention ended than before. More generally, table 6 shows that, the difference in the number of night accidents in surprised versus fixed locations increases as the number of weeks since .

This evidence suggests that the basic insight of the crime literature is very powerful: checking at random locations does lead to a greater reduction in accidents and deaths (the ultimate outcome the police wants to affect). This is despite the fact that learning does appear to occur, but it is quick enough that the Nash equilibrium is quickly reached. We now investigate whether basic mechanism we highlight in the introduction is indeed driving the results. Figure 1 shows the number of drunk caught during a given checkpoint night in fixed and surprised location, as a function of the checkpoint frequency to which the station was assigned. We see that more drunks seem to be caught at surprised locations where there is 2 or more checking per week. This is striking because in the fixed location checking are always happening at the best location, while the surprise checks rotate between locations 1, 2 and 3. Moreover, in fixed locations, the number of drunks caught per night of checking declines with the intensity of checking (which was randomly assigned), while it is not the case in the surprise location, suggesting strategic behavior.

If learning is indeed going on, in the first two weeks of the intervention, fixed checks should lead to more drunks caught since it is performed at the best place. We investigate this in figure 2, which focuses only on week 1 and 2. At one check point

⁵Note a surprising increase in day time accidents in surprise police stations.

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

Table 7 present the results from regressions on the number of drunk drivers caught. Column 1 shows that, on the average night, 0.28 more drunk drivers, or 34% are caught on surprise check night. This number is not significant due to large standard errors, but note that we should expect to see fewer of them, rather than more, since the location at the average check point is on average less productive. That this is due to strategic avoidance behavior (rather than by people drinking less) is strongly suggested by the result in column 3. In fixed locations, the more frequent the checks, the fewer drunks are caught per night. The effect entirely disappears in random locations. This is presumably because with frequent checks at fixed location, changing route is the only option for drivers, but there is no route to switch to in the case of surprise checks. Column 4 shows evidence of learning: as the number of weeks passes, the difference between fixed and surprise location increases. Finally column 2 puts the two together, and shows that as the number of past checks (which is a combination of the number of weeks and the random assignment to checking frequency) increases, the difference between surprise and fixed increases. Our final specification is presented in column 5. As we mentioned, in half the police station, we formed teams with better incentives (within the constraints of the rigid rules within the police department). These "police lines" teams were more assiduous in carrying out the checks, and they caught more drunk people. We could thus have expected that it is another way in which the intensity of the program increase, and hence a positive interaction with surprise checks. In fact that coefficient of the interaction is indeed positive, but not significant.

Finally, Table 8 shows similar regressions, but using the number of vehicles passing at check points as the dependent variable. We obtain the same results as for the number of drunks caught, except for the number or weeks: in fixed locations, fewer and fewer vehicle are seen when checking is more frequent, but that effect

disappears in random locations and the difference between fixed and surprised is significantly larger when there have been more checks in the past.

7. CONCLUSION

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

REFERENCES

- [1] Banerjee, Abhijit, Raghavendra Chattopadhyay, Esther Duflo, Daniel Keniston, and Nina Singh. Can Institutions be Reformed from Within? Evidence from a Randomized Experiment with the Rajasthan Police. Working Paper 17912, National Bureau of Economic Research, 2012.
- [2] Becker, Gary. "Crime and Punishment: An Economic Approach." *Journal of Political Economy* 76, no. 2 (1968): 169-217.
- [3] Braga, Anthony A. "The Effects of Hot Spots Policing on Crime." *Annals of the American Academy of Political and Social Science*, 578 (2001): 104-125.
- [4] Clarke, Ronald V. and David Weisburd. "Diffusion of crime control benefits: observations on the reverse of displacement." *Crime Prevention Studies* 2, (1994): 165-184.
- [5] Davis, A., A. Quimby, W. Odero, G. Gururaj, and M. Hajar. Improving Road Safety by Reducing Impaired Driving in. London: DFID Global Road Safety Partnership, 2003.
- [6] Duflo, Esther, Abhijit Banerjee, and Rachel Glennerster. "Putting a band-aid on a corpse: Incentives for nurses in the Indian public health care system." *Journal of the European Economic Association* 6, no. 5 (2010): 487-500.
- [7] Eeckhout, Jan, Nicola Persico, and Petra E. Todd. "A Theory of Optimal Random Crackdowns." *American Economic Review* 100, no. 3 (2010): 1104-1135.
- [8] Elder, Randy W., Ruth A. Shults, David A. Sleet, James L. Nichols, Stephanie Zaza, and Robert S. Thompson. "Effectiveness of Sobriety Checkpoints for Reducing Alcohol-Involved Crashes." *Traffic Injury Prevention* 3 (2002): 266-274.
- [9] Erke, Alena, Charles Goldenbeld, and Truls Vaa. "The effects of drink-driving checkpoints on crashes—A meta-analysis." *Accident Analysis and Prevention* 41 (2009): 914-923.
- [10] Green, Lorraine and David Weisburd. "Policing drug hot spots: the Jersey City drug market analysis experiment." *Justice Quarterly* 12, no. 4 (1995): 711-735.
- [11] Groff, Elizabeth R., Jerry H. Ratcliffe, Travis Taniguchi, and Jennifer D. Wood. "The Philadelphia foot patrol experiment: A randomized controlled trial of police patrol effectiveness in violent crime hotspots." *Criminology* 49, no. 3 (2011): 795-831.
- [12] Habyarimana, James, and William Jack. "Heckle and Chide: Results of a Randomized Road Safety Intervention in Kenya." Working Paper Number 169. Center for Global Development, April 2009.
- [13] Koper, Christopher S. "Just enough police presence: reducing crime and disorderly behavior by optimizing patrol time in crime hot spots." *Justice Quarterly* 12, no. 4 (1995): 649-672.
- [14] Lazear, Edward P. Speeding, Terrorism, and Teaching to the Test. Working Paper 10932, National Bureau of Economic Research, 2004.
- [15] Mookherjee, Dilip and I. P. L. Png. "Marginal Deterrence in Enforcement of Law" *Journal of Political Economy*, Vol. 102, No. 5 (1994): 1039-1066
- [16] Peek-Asa, Corinne. "The Effect of Random Alcohol Screening in Reducing Motor Vehicle Crash Injuries." *American Journal of Preventative Medicine* 18, no. 1S (1999): 57-67.
- [17] Ross, Laurence H. "Punishment as a factor in preventing alcohol-related accidents." *Addiction* 88, no. 7 (1993): 997-1002.
- [18] Sherman, Lawrence W. "Police Crackdowns: Initial and Residual Deterrence." *Crime and Justice* 12, (1990): 1-48.
- [19] Transparency International. India Corruption Study 2005: Corruption in the Indian Police. New Delhi: Center for Media Studies, 2005.

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

TABLES

TABLE 1. Police station treatment assignment

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

TABLE 2. Summary Statistics

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

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

TABLE 3. Pooled Results

	Daylight		Darkness		All
	(1) Accidents	(2) Deaths	(3) Accidents	(4) Deaths	(5) Deaths
Treatment during & post intervention	0.00317 (0.00372)	-0.00258 (0.00285)	-0.00568** (0.00238)	-0.00401* (0.00214)	-0.00391 (0.00311)
Month FE	Yes	Yes	Yes	Yes	Yes
Observations	5083	4717	5083	4717	5083
R^2	0.426	0.228	0.351	0.230	0.326
Mean of control	0.0849	0.0293	0.0329	0.0165	0.0454

Standard errors in parentheses

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

TABLE 4. Fixed vs. Surprise Pooled Results

	Daylight		Darkness	
	(1) Accidents	(2) Deaths	(3) Accidents	(4) Deaths
Fixed checkpoints during & post intervention	-0.00217 (0.00447)	-0.00425 (0.00327)	-0.00162 (0.00290)	-0.00293 (0.00235)
Surprise checkpoints during & post intervention	0.00836* (0.00448)	-0.000939 (0.00357)	-0.00962*** (0.00285)	-0.00507* (0.00289)
Month FE	Yes	Yes	Yes	Yes
Observations	5083	4717	5083	4717
R^2	0.426	0.228	0.352	0.230
Mean of control	0.0849	0.0293	0.0329	0.0165
P-value different coefficients	0.0369	0.380	0.0137	0.486

Standard errors in parentheses

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

TABLE 5. Fixed vs. Surprise, During & Post Intervention

	Daylight		Darkness		All
	(1) Accidents	(2) Deaths	(3) Accidents	(4) Deaths	(5) Deaths
Fixed checkpoints	-0.0000727 (0.00646)	-0.00171 (0.00408)	0.00179 (0.00456)	-0.00214 (0.00284)	-0.00310 (0.00476)
Surprise checkpoints	0.0127* (0.00648)	0.00217 (0.00465)	-0.0108*** (0.00406)	-0.00540 (0.00338)	-0.00299 (0.00550)
Fixed checkpoints post intervention	-0.00332 (0.00506)	-0.00679 (0.00420)	-0.00406 (0.00337)	-0.00387 (0.00295)	-0.00610 (0.00433)
Surprise checkpoints post intervention	0.00544 (0.00545)	-0.00427 (0.00425)	-0.00863** (0.00353)	-0.00460 (0.00382)	-0.00286 (0.00477)
Month FE	Yes	Yes	Yes	Yes	Yes
Observations	5083	4717	5083	4717	5083
R^2	0.427	0.229	0.352	0.230	0.327
Mean of control	0.0849	0.0293	0.0329	0.0165	0.0454

Standard errors in parentheses

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

TABLE 6. Fixed vs. Surprise, Weeks After Project Beginning

	Daylight		Darkness		All
	(1) Accidents	(2) Deaths	(3) Accidents	(4) Deaths	(5) Deaths
Fixed checkpoints × weeks intervention & post	-0.000929*** (0.000349)	-0.000252 (0.000223)	0.0000108 (0.000244)	-0.0000416 (0.000148)	-0.000294 (0.000286)
Surprise checkpoints × weeks intervention & post	-0.000471 (0.000320)	0.000104 (0.000281)	-0.000571** (0.000236)	-0.000205 (0.000208)	-0.000101 (0.000382)
Observations	94276	94276	94276	94276	94276
R^2	0.026	0.011	0.021	0.010	0.016
Mean of control	0.0825	0.0280	0.0356	0.0168	0.0449
Month FE	No	No	No	No	No

Standard errors in parentheses

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

TABLE 7. Drunken Driver Learning

	(1)	(2)	(3)	(4)	(5)
	Drivers and Motorcy- clists				
Surprise checkpoint station	0.280 (0.213)		-1.046* (0.616)		0.112 (0.229)
Number Previous checkpoints		-0.0343*** (0.0129)			
Surprise checkpoint × number previous checkpoints		0.0300* (0.0162)			
Frequency			-0.472** (0.232)		
Surprise checkpoint × frequency			0.682** (0.290)		
Weeks of checking				-0.0829*** (0.0271)	
Surprise checkpoint × weeks of checking				0.0859** (0.0354)	
Police lines team					1.035*** (0.373)
Surprise × police lines team					0.0531 (0.415)
Observations	820	820	820	820	820
Mean of dep. variable	1.172	1.172	1.172	1.172	1.172
District FE	Yes	No	Yes	No	Yes
Police Station FE	No	Yes	No	Yes	No

Standard errors in parentheses

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

TABLE 8. Passing Vehicles Learning

	(1)	(2)	(3)	(4)	(5)
	Passing cars, m.cycles	Passing cars, m.cycles	Passing cars, m.cycles	Passing cars, m.cycles	Passing cars, m.cycles
Treatment	-48.52 (58.45)	-5.645 (56.25)	29.27 (57.84)	22.97* (13.74)	-5.181 (83.85)
Surprise checkpoint station	0.0463 (70.93)	-65.83 (85.17)	-229.4* (132.0)	-19.90 (25.78)	-91.65 (97.25)
Number Previous checkpoints		-12.12*** (2.516)			
Surprise checkpoint × number previous checkpoints		7.700** (3.363)			
Frequency			-80.71*** (27.70)		
Surprise checkpoint × frequency			115.2** (52.70)		
Weeks of checking				-8.922*** (2.440)	
Surprise checkpoint × weeks of checking				-2.521 (3.525)	
Police lines team					-109.8 (112.0)
Surprise × police lines team					208.0 (151.4)
Observations	2792	2792	2792	2792	2792
Mean of dep. variable	445.6	445.6	445.6	445.6	445.6
District FE	Yes	Yes	Yes	Yes	Yes
Police Station FE	No	No	No	No	No

Standard errors in parentheses

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

FIGURES

FIGURE 7.1. Whole intervention learning results

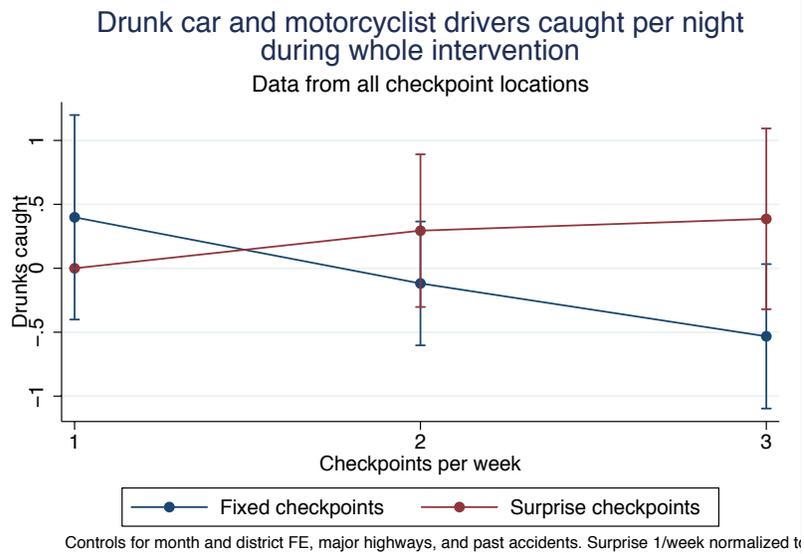
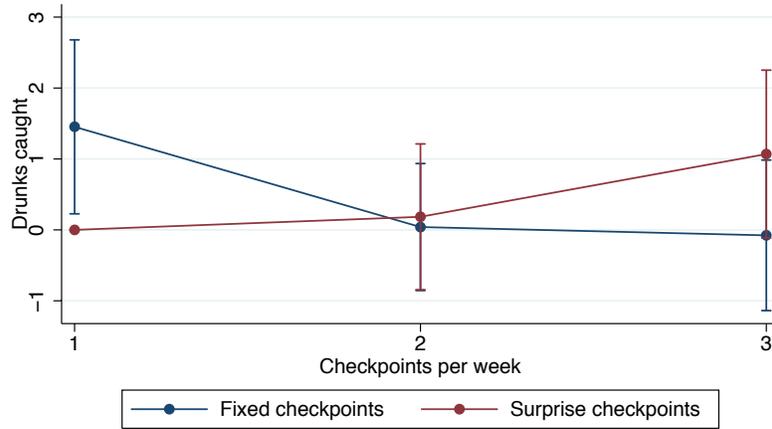


FIGURE 7.2. Pre-learning results

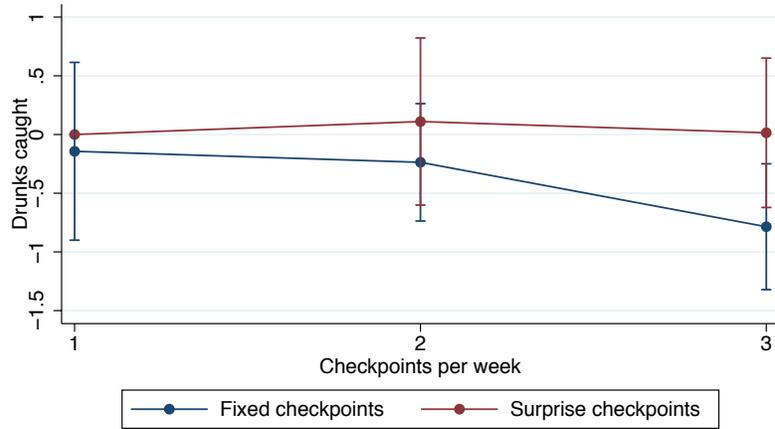
Drunk car and motorcyclists caught per night
in weeks 1, 2 of checkpoints

Data from all checkpoint locations



Controls for month and district FE, major highways, and past accidents. Surprise 1/week normalized to

FIGURE 7.3. All checkpoint locations after learning
 Drunk car and motorcyclist drivers caught per night
 after 6 weeks of checkpoints
 Data from all checkpoint locations



Controls for month and district FE, major highways, and past accidents. Surprise 1/week normalized to