Information and Policing

Shota Ichihashi*

September 1, 2023

Abstract

Agents decide whether to commit crimes based on their heterogeneous returns to crime, or their types. Police have some information about these types and allocate search capacity across the agents to uncover crimes. The police that have full information about types fail to deter any crime, because the ability to predict crimes erodes the deterrent effect of policing. The information structure that minimizes a crime rate is only partially informative and never allows the police to identify who will commit crimes, but it may reveal some of the agents who will not commit crimes.

*Queen's University, Department of Economics. Email: shotaichihashi@gmail.com. I thank Moritz Meyer-ter-Vehn for providing valuable comments on an early draft of this paper. For valuable suggestions and comments, I thank Yeon-Koo Che, Christopher Cotton, Navin Kartik, Qingming Liu, Ruslan Momot, Alex Smolin, and the participants of seminars and talks at Osaka University, Columbia Business School, Toulouse School of Economics, University of Michigan (Ross School of Business), and Queen's University.

1 Introduction

Law enforcement agencies are increasingly relying on data and algorithms to predict crimes (Perry, 2013; Brayne, 2020). They use a variety of data sources, such as criminal records, social media posts, financial records, and local environmental information. Private vendors, such as Palantir and PredPol, also offer predictive algorithms to police departments. This trend, which is driven by the pursuit of more effective law enforcement, has raised a number of concerns. As a case in point, the EU's proposed regulation "Artificial Intelligence Act" classifies a certain use of artificial intelligence to predict crime as a prohibited practice.[1]

Motivated by the recent discussion, I study how the information available to a law enforcement affects its ability to deter crimes. I examine this question from the perspective of information design. The model consists of a unit mass of agents and a law enforcement, which we call the "police." The agents are potential criminals with heterogeneous types that capture their returns to crimes. At the outset, the police observe information about the type of each agent according to an exogenous signal structure. Each agent decides whether to commit a crime, and simultaneously, the police allocate search resources across agents to catch criminals. This simultaneous-move game captures a situation in which agents decide whether to commit crimes—such as illegal parking, tax fraud, or drug trafficking—and then a police officer or tax auditor tries to uncover crimes without directly observing the agents' behavior.

I present three main results. The first result shows that maximal information leads to maximal crime rate: If the police have full information about the agents' types, every agent commits a crime with probability 1 in any equilibrium. The fully informed police will allocate search resources only to the agents who are most likely to commit crimes. This search strategy fails to deter any crime: Each agent is either not searched at all, and thus commits a crime, or is searched with a positive probability but still commits a crime as predicted by the police. The result implies that for policing to have a deterrence role, we may need to restrict the police's information.

The second result characterizes the signal structure that minimizes a crime rate. To

[1] See https://www.europarl.europa.eu/news/en/press-room/20230609IPR96212/meps-ready-to-negotiate-first-ever-rules-for-safe-and-transparent-ai.

do so, I first solve a relaxed problem in which we jointly design the information available to agents and the police. The relaxed problem has a solution where the police receive no information and randomly search any agent with the same probability, whereas each agent learns whether their type exceeds some cutoff. I then show that in the original problem where the agents observe their types, we can construct a signal structure for the police that attains the same outcome as in the relaxed problem. This crime-minimizing signal structure has two properties. First, it enables the police to avoid excessive searches, i.e., in equilibrium, no agent is searched with a probability strictly greater than what is required to deter their crime. Second, the signal structure garbles the agents' types to prevent the police from predicting crimes, i.e., in equilibrium, the likelihood of crime is equalized across all possible signals. The crime-minimizing signal structure is also amenable to comparative statics. For example, it becomes less informative when the police have a lower search capacity or the agents face greater returns to crime.

The above results assume that the police distribute a fixed search capacity across the agents. The third result assumes that the police endogenously choose a total search capacity at a cost: For example, an individual officer might exert different levels of search effort depending on the available information. In such a case, the crime-minimizing signal structure may enable the police to identify some of the agents who will abstain from crimes in equilibrium. Revealing some of these "innocents" could decrease a crime rate, because the police who possess such information can more effectively allocate search resources and thus choose a higher level of search effort. At the same time, the crime-minimizing signal structure never enables the police to identify agents who will commit crimes, because revealing these "criminals" not only distorts the allocation of search resources—as shown in the first two results—but also reduces the police's search effort.

Overall, the paper offers a cautionary tale against the growing use of data and predictive algorithms in law enforcement: Information can improve the allocation of search resources but also erode the deterrence effect of law enforcement. The information structure that strikes this balance may limit or even eliminate the ability of the law enforcer to predict crimes. Moreover, the law enforcer may need to have less information when they have limited resources or operate in an environment prone to crime.

Related work. The paper relates to the literature on Bayesian persuasion and information design (see Kamenica (2019) and Bergemann and Morris (2019) for surveys). Papers such as Lazear (2006), Eeckhout, Persico, and Todd (2010), and Hernández and Neeman (2022) use Bayesian persuasion and related tools to study the problem of disclosing information to players who may take socially undesirable actions. In contrast, I study what information a law enforcer should have about such players. Methodologically, the paper studies information design problems with a continuum of players, states, and actions (for the police), which seems to be less understood than a single-player Bayesian persuasion problem (see Smolin and Yamashita (2022) for a discussion). The information design literature provides conditions under which an optimal signal takes a tractable form, such as monotone partitional signals, censorship policies, and nested intervals (e.g., Guo and Shmaya 2019; Dworczak and Martini 2019; Kolotilin, Mylovanov, and Zapechelnyuk 2022). The crime-minimizing signal structure characterized in this paper does not belong to these classes of signals.

The paper also relates to the economic literature on crime and policing, which starts from Becker (1968). The question of what information about agents should or should not be used for policing is often discussed in the context of racial profiling (Knowles, Persico, and Todd 2001; Persico and Todd 2005; Bjerk 2007; Persico 2009). In terms of the timing and payoffs of the game, my paper builds closely on Persico (2002), who studies whether requiring the police to adopt a fairer search strategy reduces crime. Instead of constraints on police behavior, I consider restrictions on information available to police and study what information renders policing effective. A model of predictive enforcement is also studied by Che, Kim, and Mierendorff (2023), who consider a bandit model that captures the endogenous generation and use of information for law enforcement. To focus on the role of information in policing, the model abstracts away from other important considerations, such as the design of judicial systems, richer responses by potential criminals and victims, as well as the "fairness" of predictive algorithms (e.g., Curry and Klumpp 2009; Cotton and Li 2015; Jung, Kannan, Lee, Pai, Roth, and Vohra 2020; Vasquez 2022; Liang, Lu, and Mu 2022).

2 Model

The model consists of police and a unit mass of agents indexed by $i \in [0,1]$. Each agent i has some underlying returns to crime, or *type*, $x_i \in [0,1]$.[2] Each agent observes their type. We may interpret an agent's type as reflecting the individual characteristics that affect their returns to crime (e.g., legal earning opportunities) or crime opportunities specific to certain locations or times. The agents' types are independently and identically drawn from distribution function $F \in \Delta[0,1]$, which has a positive density f and is commonly known.[3] We use $\mathbb{E}_F[\cdot]$ for the corresponding expectation operator. Also, we use $F(\cdot|\tilde{x} \leq c)$ for the conditional distribution of F on $[0,c]$ and $\mathbb{E}_F[\cdot|\tilde{x} \leq c]$ for the corresponding expectation operator. We denote the uniform distribution over the interval $[a,b]$ by $U[a,b]$.

The police learn information about each agent's type according to a *signal structure* (S,π), which consists of a set S of signals and a collection $\pi = \{\pi(\cdot|x)\}_{x\in[0,1]}$ of conditional distributions $\pi(\cdot|x)$ over signals for each type x. For each agent $i \in [0,1]$, the police observe a signal $s_i \in S$ drawn according to distribution $\pi(\cdot|x_i) \in \Delta S$. Conditional on types, signals are independent across agents. The signal structure is exogenous and commonly known, but only the police observe realized signals.

Given the signal structure, the police and the agents play the following simultaneous-move game: Each agent decides whether or not to commit a crime, and simultaneously, the police allocate search resources across agents. Specifically, the police choose a *search strategy* $p : S \to [0,1]$, where $p(s)$ is the probability of searching agents with signal $s \in S$. The police have a measure $\overline{P} \in (0,1)$ of searches to allocate; thus the police can choose a search strategy $p(\cdot)$ if and only if the total mass of searches does not exceed $\overline{P}$, i.e.,

$$\int_0^1 \int_S p(s)\pi(ds|x)F(dx) \leq \overline{P}. \tag{1}$$

For a given strategy profile, we use *innocents* for agents who commit crimes with probability 0 and *criminals* for agents who commit crimes with probability 1.

[2] For simplicity, I exclude types below 0 or above 1. Under the payoff specification presented below, such agents are nonstrategic and their best responses do not depend on the police's strategy.

[3] Given set X we write ΔX for the set of all probability distributions on X.

An agent's payoff of committing a crime is $x - p$, where x is the agent's type and $p \in [0, 1]$ is the probability of being searched.[4] An agent's payoff of not committing a crime is 0. Thus an agent commits a crime if and only if their type x exceeds the anticipated probability of search. The police's payoff is equal to the *mass of successful searches*, which is defined as the mass of agents who commit crimes and are searched by the police. The literal interpretation of the police's payoff is that whenever the police search an agent who has indeed committed a crime, the police earn a payoff of 1. The solution concept is Bayesian Nash equilibrium, which we refer to as *equilibrium*.

To minimize possible case classifications, we assume that the primitives—i.e., the distribution of returns to crime and the police's search capacity—satisfy the following:

Assumption 1. The primitives, F and $\overline{P}$, satisfy $\overline{P} < \int_0^1 xF(dx)$.

The assumption implies that under any signal structure and any equilibrium, a positive mass of agents commit crimes. Indeed, the inequality $\overline{P} < \int_0^1 xF(dx)$ implies that the police do not have enough search capacity to search each agent i with a probability of at least x_i. As a result, some agents, facing search probabilities below x_i, prefer to commit crimes.

The main analysis focuses on two signal structures: One is a signal structure that provides the police with full information about the agents' types, i.e., a signal structure is such that for every type $x \in [0, 1]$, $\pi(\cdot|x)$ places probability 1 on $s = x$. The other is a signal structure that minimizes a *crime rate*, defined as the mass of agents who commit crimes in equilibrium:

Definition 1. A *crime-minimizing signal structure* is a signal structure that has an equilibrium with the lowest crime rate across all signal structures and equilibria. The corresponding equilibrium is called a *crime-minimizing equilibrium*.

We can view a crime-minimizing signal structure as the choice of a social planner who internalizes the social cost of crime and regulates the information available to the police. Crime minimization is different from the police's objective, which is to maximize the number of successful searches. The police's objective is better interpreted as the preferences of

[4] We obtain qualitatively the same results when an agent's payoff of committing a crime is, in line with Persico (2002), $(1 - p)x - Lp$, i.e., criminals enjoy their returns to crime if they are not searched and incur a loss of L if they are searched. We adopt payoff $x - p$ to simplify exposition.

individual officers or auditors who have career concerns and are rewarded for uncovering crimes, as discussed below.

Remark 1 (Timing and Payoffs). The results of this paper hinge on the assumption that the police and agents move simultaneously. In particular, the police cannot commit to a search strategy in advance. Two reasons support this assumption. First, we view a signal structure as a predictive algorithm used within a law enforcement agency. In such a context, predictions generated by an algorithm (i.e., realized signals) would not be visible to the public, making it difficult for the police to commit to a search strategy as a function of realized signals. Second, in line with the literature on decentralized law enforcement, such as Persico (2002) and Porto et al. (2013), we may view the "police" not as an organization but as individuals such as law enforcement officers and tax auditors. The simultaneous-move assumption then arises from the idea that the action of an individual officer does not directly influence the decisions of potential criminals. Finally, the simultaneous-move assumption offers a partial justification to the assumption that the police care about successful searches but not a crime rate. Indeed, even if the police's payoffs depend both on successful searches and a crime rate, the simultaneous-move assumption implies that the police behave as if they only care about successful searches, taking a crime rate as given.

3 Fully Informed Police

We begin with the analysis of the fully informed police, which we may interpret as reflecting the consequence of unrestricted data collection by a law enforcement agency.

Theorem 1. *If the police have full information, then in any equilibrium, almost every agent commits a crime with probability* 1*.*

Proof. Take any equilibrium. Assumption 1 implies that the police cannot search every type x with a probability weakly above x. Hence a positive mass of types who expect search probabilities strictly below their types commit crimes with probability 1. If another positive mass of types commit crimes with probability strictly below 1, the police could increase successful searches by shifting search probabilities away from these types to the types who surely commit crimes. This is a contradiction, so almost every agent commits a crime. $\square$

The result captures the idea that the police's ability to predict crimes—combined with their incentive to uncover crimes and a lack of commitment in search strategy—eliminates the deterrence effect of policing. For example, suppose that $F = U[0, 1]$ and $\overline{P} = \frac{1}{4}$. If the police have no information, every agent faces search probability $\overline{P}$, so the equilibrium crime rate is $1 - F(\overline{P}) = \frac{3}{4}$. In contrast, when the police have full information, there are multiple equilibria, all of which have crime rate 1: In one equilibrium, the police adopt search strategy $p(x) = \frac{x}{2}$ for every $x \in [0, 1]$. Given this search strategy, all agents strictly prefer to commit crimes. Moreover, because every agent is equally likely to be committing a crime, the police are indifferent across all search strategies that exhaust search capacity $\overline{P}$, which includes search strategy p. In another equilibrium, the police search type $x \leq \frac{1}{\sqrt{2}}$ with probability x and do not search any type $x > \frac{1}{\sqrt{2}}$. In this equilibrium, each type $x \leq \frac{1}{\sqrt{2}}$ is indifferent yet commits a crime with probability 1. It cannot be an equilibrium if some agents break ties and do not commit a crime, because the police would then redirect their search toward types above $\frac{1}{\sqrt{2}}$.

4 Crime-Minimizing Signal Structure

Having established that restricting the police's information is necessary for crime deterrence, we turn to the question of what signal structure minimizes a crime rate. The analysis consists of two steps. First, we study a "relaxed problem," in which we minimize a crime rate by designing a joint information structure for the police and the agents. We show that the solution to this problem is to disclose no information to the police and reveal the agents whether their types exceed a cutoff. The resulting crime rate becomes a lower bound of the possible crime rates in the original problem, in which the agents observe their types. Second, we turn to the original setup and construct a signal structure that attains this lower bound.

4.1 Relaxed Problem

To define the relaxed problem, we modify the model as follows: The agents do not directly observe their types. Instead, the information of the agents and the police is determined by an *extended signal structure*, (S_P, S_A, π). Here, S_P and S_A are the sets of signals for the police and agents, respectively, and $\pi = \{\pi(\cdot|x)\}_{x\in[0,1]}$ is the collection of conditional probability distributions on $S_P \times S_A$ for each type. If agent i has type x_i, the police observe s_i^P and agent i observes s_i^A, where $(s_i^P, s_i^A) \sim \pi(\cdot|x_i)$. The rest of the game remains the same: Each agent observes s_i^A and decides whether to commit a crime, and simultaneously, the police choose a search strategy to maximize successful searches. The following result characterizes a crime-minimizing extended signal structure.

Lemma 1. *In the relaxed problem, the following extended signal structure minimizes a crime rate: The police learn no information, e.g., $S_P = \{\phi\}$, and each agent learns whether their type exceeds cutoff $\hat{c} \in (0,1)$ that uniquely solves*

$$\mathbb{E}_F[\tilde{x}|\tilde{x} \leq \hat{c}] = \overline{P}. \tag{2}$$

In equilibrium, the police search every agent with probability $\overline{P}$, and each agent commits a crime if and only if their type exceeds $\hat{c}$.

Proof. Take any extended signal structure and any equilibrium. The proof consists of three steps. First, by the "revelation principle" of information design (e.g., Bergemann and Morris 2019), we can replace the agents' signals with action recommendations.[5] Specifically, we set $S_A = \{crime, not\}$ and assume that in equilibrium, each agent commits a crime after observing signal *crime* and not after observing signal *not*.

Second, we replace the police's signal with an uninformative signal, e.g., $S_P = \{\phi\}$. This weakly decreases search probabilities allocated to signal *crime*, because the police aim to

[5] Formally, take any extended signal structure and equilibrium. Let $a_i(y)$ be the equilibrium probability that agent i takes action $a \in \{crime, not\}$ after observing signal $y \in S_A$; $\pi_{AP}(\cdot|x, S)$ be the conditional distribution of an agent's signal when their type is x and the police's signal is in S; and $\pi_P(\cdot|x)$ be the conditional distribution of the police's signal given type x. We define the new signal structure as $\hat{\pi}(S \times \{a\}|x) = \pi_P(S|x) \int_0^1 \int_{S_A} a_i(y)\pi_{AP}(dy|x, S)di$ for each $S \subset S_P$ and action recommendation $a \in \{crime, not\}$.

maximize successful searches. As a result, the obedience constraints get relaxed, ensuring that the agents continue to follow action recommendations.[6]

The relaxed problem now reduces to Bayesian persuasion: An agent receives a payoff of $x - \overline{P}$ from committing a crime, and we disclose information about type $x \sim F$ (in the form of action recommendations) to the agent in order to minimize the probability of committing a crime. As shown in Kamenica and Gentzkow (2011), the solution is to disclose whether type x exceeds a cutoff $\hat{c}$ defined by $\mathbb{E}_F[\tilde{x}|\tilde{x} \leq \hat{c}] = \overline{P}$. Type distribution has a density and satisfies Assumption 1, so the cutoff $\hat{c} \in (0,1)$ exists and is unique. $\square$

Lemma 1 implies that in our original setup, the minimum crime rate is attained if all types below $\hat{c}$ do not commit crimes. Such an outcome cannot arise if the agents observe their types but the police have no information, because the resulting random search induces types between $\overline{P}$ and $\hat{c}$ to commit crimes. However, the next section shows that we can provide the police with partial information to attain the same crime rate.

4.2 Characterizing the Crime-Minimizing Signal Structure

We now turn to the original setup, in which the agents observe their types. First, we define a class of signal structures:

Definition 2. For each $c \in (0,1)$, the *truth-or-noise signal structure with cutoff* c, denoted by (S_c, π_c), is the following signal structure: The signal space S_c is $[0, c]$; for each $x \leq c$, distribution $\pi_c(\cdot|x)$ draws $s = x$ with probability 1; and for any $x > c$, distribution $\pi_c(\cdot|x)$ is independent of x and equals $F(\cdot|\tilde{x} \leq c)$.

To better understand the truth-or-noise signal structures, consider the police's posterior belief on types induced by a signal generated by (S_c, π_c) (see Figure 1). The signal can be the "truth" (i.e., an agent's type is below c and the signal is equal to their true type) or a "noise" (i.e., an agent's type is above c and the signal is a realization of a random draw from $F(\cdot|\tilde{x} \leq c)$). As a result, the posterior belief places positive probabilities both on types

[6] The obedience constraint for signal *crime* (resp. signal *not*) is that an agent's expected payoff from committing a crime conditional on signal *crime* is non-negative (resp. non-positive). When we simply say the obedience constraints, they refer to the constraints for both signals.

below and above c. Specifically, the posterior induced by signal $s \in [0, c]$ contains a point mass $F(c)$ of type s and a mass $1 - F(c)$ of types distributed according to $F(\cdot|\tilde{x} > c)$. The posterior beliefs (indexed by their point masses) are distributed according to $F(\cdot|\tilde{x} \leq c)$, so they average to prior distribution F.

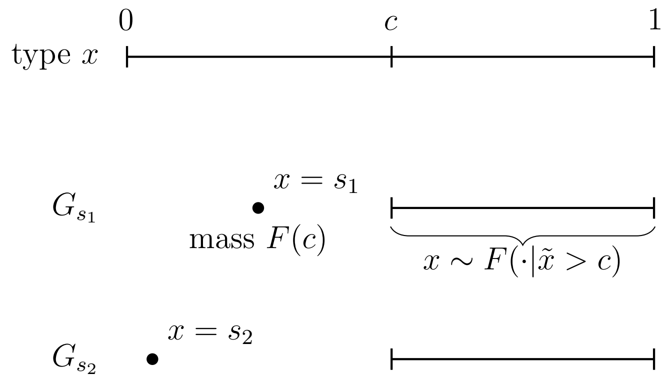

Figure 1: Posterior distribution G_s of types conditional on signal $s = s_1, s_2$ under (S_c, π_c).

Theorem 2. *Let $\hat{c}$ denote the cutoff defined by equation (2). The truth-or-noise signal structure with cutoff $\hat{c}$ is a crime-minimizing signal structure.*

Proof. It suffices to show that signal structure $(S_{\hat{c}}, \pi_{\hat{c}})$ has an equilibrium in which types above $\hat{c}$ commit crimes and types below $\hat{c}$ do not. Consider the following strategy profile: The police adopt search strategy $p^*(s) = s$ for every $s \in [0, \hat{c}]$, and each agent commits a crime if and only if their type exceeds $\hat{c}$. We show that this strategy profile is an equilibrium. First, from the police's perspective, each agent is committing a crime with probability $1-F(\hat{c})$ conditional on any signal. Thus the police are indifferent across all search strategies that exhaust search capacity $\overline{P}$. Search strategy p^* indeed exhausts the search capacity because of equation (2), i.e., $\overline{P} = \mathbb{E}_F[\tilde{x}|\tilde{x} \leq \hat{c}]$. The strategy of each agent is also optimal: Any agent with type $x \leq \hat{c}$ knows that the police will observe signal $s = x$ and search them with probability x, so the agent is indifferent and be willing to abstain from a crime. Agents with types above $\hat{c}$ will be searched with probability at most $\hat{c}$, so they commit crimes. Hence the strategy profile described above is an equilibrium. $\square$

The crime-minimizing signal structure has two properties: First, it prevents the police from predicting crimes by equalizing the likelihood of crime across all signals. This property minimizes the distortion highlighted by Theorem 1, where the police focus their search resources on agents most likely to commit crimes. Various signal structures have this first property, including the one that discloses no information. The second property of the crime-minimizing signal structure is that it reveals partial information and enables the police to reduce wasteful searches: Under the crime-minimizing signal structure, the signals are differentiated according to the lowest possible types, and the equilibrium search rates are tailored to these types. Hence the police will never search agents with a probability greater than what is minimally necessary to deter crimes.

The crime-minimizing signal structure admits the following comparative statics:

Corollary 1. *The crime-minimizing signal structure in Theorem 2 becomes less Blackwell informative if (i) search capacity $\overline{P}$ is smaller or (ii) type distribution F is greater in the likelihood ratio order.*[7]

Proof. For Part (i), suppose that search capacity decreases from $\overline{P}_2$ to $\overline{P}_1$. Let (S_1^*, π_1^*) and (S_2^*, π_2^*) denote the respective crime-minimizing signal structures. Equation (2) implies that cutoff $\hat{c}_2$ under $\overline{P}_2$ is greater than cutoff $\hat{c}_1$ under $\overline{P}_1$. We can use this property to show that (S_2^*, π_2^*) is more informative than (S_1^*, π_1^*). Indeed, we can replicate (S_1^*, π_1^*) by garbling signals drawn from (S_2^*, π_2^*) so that whenever a signal drawn from (S_2^*, π_2^*) falls in the interval $(\hat{c}_1, \hat{c}_2]$, we redraw a new signal randomly from distribution $F(\cdot|\tilde{x} \leq \hat{c}_1)$. For Part (ii), note that as F increases in the likelihood ratio order, the right-hand side of equation (2) (i.e., $\overline{P} = \mathbb{E}_F[\tilde{x}|\tilde{x} \leq \hat{c}]$) increases.[8] To maintain the equality, cutoff $\hat{c}$ must decrease. By the same argument as the proof of Part (i), the decrease in the cutoff implies a less informative crime-minimizing signal structure. □

Corollary 1 implies that when the environment is more prone to crime, the police should have less information for the maximum crime deterrence. For example, if the police have a lower search capacity, then everything else equal, more agents will commit crimes. The

[7]Take two type distributions, G and F, that have positive densities on $[0, 1]$. We say that distribution *G is greater than F in the likelihood ratio order* if $\frac{g(x)}{f(x)}$ increases in x.

[8]If G is greater than F in the likelihood ratio order, then for any measurable event $A \subset [0, 1]$, $G(\cdot|\tilde{x} \in A)$ first-order stochastically dominates $F(\cdot|\tilde{x} \in A)$ (see, e.g., Theorem 1.C.6 of Shaked and Shanthikumar (2007)).

police will then use information to identify agents who are more likely to commit crimes and allocate a large fraction of scarce search resource to these agents. To mitigate this distortion, the crime-minimizing information structure reveals less information to the police and makes it harder to identify who face high returns to crime. The result counters the idea that better prediction technology could partly compensate for the lack of physical resources in law enforcement.[9]

5 Endogenous Search Capacity

We now study the case in which the police can increase the total number of searches at cost. Formally, we assume that the police can choose any search strategy p at cost $C(P)$, where P is the total mass of searches induced by p, i.e.,

$$P \triangleq \int_0^1 \int_S p(s)\pi(ds|x)F(dx).$$

The cost function, $C(\cdot)$, is strictly increasing, strictly convex, differentiable, and satisfies $C'(0) < 1 < C'\left(\int_0^1 xF(dx)\right)$.[10] The police's payoff is the mass of successful searches minus cost $C(P)$. The rest of the model, such as the agents' payoffs and the timing, remains the same.

5.1 Relaxed Problem with Endogenous Search Capacity

First, we study the relaxed problem and solve for an extended signal structure (S_P, S_A, π) that minimizes the equilibrium crime rate. Figure 2 describes the solution: Similar to the case of the exogenous search capacity (Lemma 1), each agent receives signal *crime* or *not* depending on whether their type exceeds some cutoff c^*, and in equilibrium, they follow action recommendations. The police's signal is a garbling of the agents' signals. In particular,

[9] A police chief quoted in Pearsall (2010) says that a predictive policing algorithm "is the perfect tool to help departments become more efficient as budgets continue to be reduced."

[10] Under any signal structure and any equilibrium, the first inequality implies that the police search a positive mass of agents, whereas the second inequality, which plays the same role as Assumption 1, ensures a positive crime rate.

the police privately identify and choose not to search a fraction α^* of agents who receive signal *not*. For the remaining population, the police apply the same search rate $\rho^* > 0$.

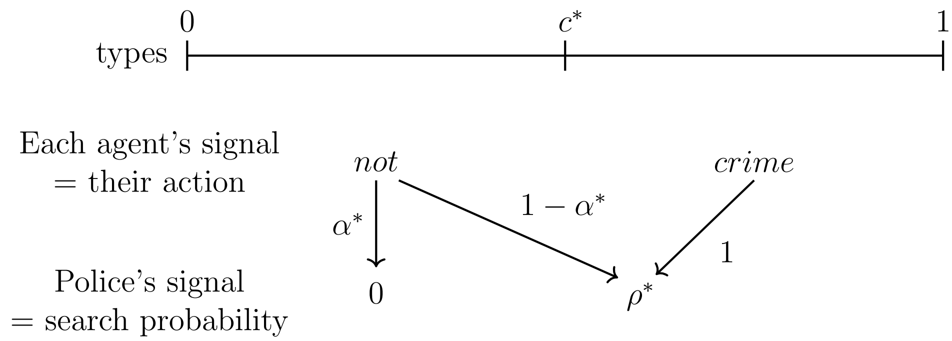

Figure 2: A solution to the relaxed problem with endogenous search capacity. The police identify a fraction α^* of innocents.

To simplify exposition, we prepare some terminologies (in what follows, we refer an extended signal structure simply as a signal structure). Take any signal structure (S_P, S_A, π) and the strategies of agents. For each $s \in S_P$, its *posterior crime rate*, $r(s) \in [0, 1]$ refers to the probability with which an agent commits a crime conditional on the police's signal s.[11] The following lemma restricts the class of signals we need to consider:

Lemma 2. *In the relaxed problem, take any signal structure and any equilibrium, denoting its crime rate by r. There is some $c \in [0, 1]$ such that the same crime rate arises under a signal structure and an equilibrium with the following properties:*

1. *Each agent receives signal "crime" and signal "not" if $x > c$ and $x < c$, respectively. In equilibrium, the agents follow action recommendations.*

2. *The police's signal is a garbling of an agent's signal.*[12] *In equilibrium, each signal of the police leads to a distinct posterior crime rate.*

---

[11]For every signal $t \in S_A$ of an agent, let $a(t)$ be the average probability with which an agent commits a crime after observing signal t. Then the posterior crime rate is equal to $r(s) = \mathbb{E}[a(\tilde{t})|s]$, where the expectation is with respect to an agent's signal $\tilde{t}$ conditional on the police's signal s.

[12]By "garbling," we mean that there exist conditional distributions of the police's signal given an agent's signal, i.e., $\pi_P(\cdot|crime), \pi_P(\cdot|not) \in \Delta S_P$, such that for any $S \subset S_P$, $a \in \{crime, not\}$, and $x \in [0, 1]$, we have $\pi(S \times \{a\}|x) = \pi_P(S|a)\pi_A(a|x)$. Here, $\pi_A(s|x)$ is the probability of an agent's signal being a conditional on type x.

Proof. Take any signal structure (S'_P, S'_A, π') and any equilibrium. Let p' denote the police's equilibrium search strategy. First, as in Lemma 1, we can replace S'_A with $S_A = \{crime, not\}$ and assume that each agent follows the action recommendation in equilibrium.

Second, we replace each signal $s' \in S'_P$ of the police with its posterior crime rate $r(s')$, resulting in a new signal space, $S_P \subset [0, 1]$. This replacement reduces the police's information because different signals may have the same posterior crime rate. We then assume that the police adopt search strategy $p(y) \triangleq \mathbb{E}[p'(s')|r(s') = y]$ for each $y \in S_P$, where the expectation is with respect to the police's original signal $s' \in S'_P$ conditional on posterior crime rate y. The police find it optimal to adopt p because it ensures the same payoff as p' despite having less information. The agents' incentives remain the same, because strategies p' and p result in the same expected search rates conditional on each signal in S_A.

Finally, let $\hat{\pi} \in \Delta(S_P \times S_A)$ denote the joint distribution of the police's signal (i.e., posterior crime rate) and an agent's signal (i.e., action recommendation). Let $\hat{\pi}(\cdot|s) \in \Delta S_P$ denote the associated conditional distribution of the police's signal given an agent's signal $s \in \{crime, not\}$. We modify the signal structure as follows. First, given equilibrium crime rate $r \in [0, 1]$, we assume that types above and below cutoff $c \triangleq F^{-1}(1 - r)$ receive signals *crime* and *not*, respectively. This modification relaxes the obedience constraints for the agents, so they are willing to follow action recommendations. Second, we assume that conditional on each agent i's signal $s_i \in \{crime, not\}$, the police observe signal $y_i \sim \hat{\pi}(\cdot|s_i) \in \Delta S_P$ (regardless of i's type). This modification preserves the joint distribution of posterior crime rates and action recommendations across the population. As a result, the agents continue to follow action recommendations, and the police optimally choose search straetgy p. The resulting equilibrium has the desired properties and attains crime rate r. □

The second part of Lemma 2 means that in equilibrium, the police receive a noisy signal of an agent's behavior. We can impose this restriction upon solving the relaxed problem, because the police care about the likelihood of crimes but do not directly care about the agents' types. Using the lemma, we show that the crime-minimizing signal structure of the relaxed problem takes the form described in Figure 2.

Proposition 1 (Relaxed Problem with Endogenous Search Capacity). *In the re-*

laxed problem, the following signal structure (S_P^*, S_A^*, π^*)*, characterized by tuple* $(\rho^*, c^*, \alpha^*) \in (0,1)^2 \times [0,1)$*, minimizes a crime rate:*

1. *The signal space is* $S_A^* = \{crime, not\}$ *for the agents and* $S_P^* = \{0, \rho^*\}$ *for the police. In equilibrium, the agents follow action recommendations, and the police choose search strategy* $p(0) = 0$ *and* $p(\rho^*) = \rho^*$*.*

2. *For every type* $x > c^*$*, the realized signals are* $(\rho^*, crime)$ *with probability* 1*. For every type* $x < c^*$*, the realized signals are* $(0, not)$ *and* (ρ^*, not) *with probabilities* α^* *and* $1 - \alpha^*$*, respectively.*

3. *Tuple* (ρ^*, c^*, α^*) *satisfies* $\mathbb{E}_F[\tilde{x}|\tilde{x} \leq c^*] = (1-\alpha^*)\rho^*$*, i.e., the agents who observe signal "not" are indifferent between committing a crime and not.*

Proof. Take any signal structure (S_P, S_A, π) and any equilibrium with the properties described in Lemma 2. Part 2 of the lemma ensures that the police's signal is a garbling of an agent's signal, generated by conditional distributions $\hat{\pi}(\cdot|crime), \hat{\pi}(\cdot|not) \in \Delta S_P$. It is without loss to assume that the police's signals are recommended search probabilities the police follow in equilibrium.

First, we show $S_P \subset \{0, \rho, 1\}$ for some $\rho \in (0,1)$. If S_P contains multiple interior search rates $\rho, \rho' \in (0,1)$, they must have the same posterior crime rate, i.e., $r(\rho) = r(\rho')$. For example, if $r(\rho) < r(\rho')$, the police would profitably deviate by shifting search probabilities from signal ρ to ρ' without changing the total search capacity. However, $r(\rho) = r(\rho')$ contradicts Part 2 of Lemma 2 that each signal leads to a distinct posterior crime rate.[13] As a result, S_P contains at most one interior search rate, i.e., $S_P \subset \{0, \rho, 1\}$ for some ρ.

The unique interior search rate ρ (if exists) satisfies two properties. First, $\{\rho, 1\} \subset S_P$ implies $r(\rho) < r(1)$, i.e., the agents whom the police search with a higher probability (here, 1) have a higher posterior crime rate. Second, the police equate the marginal cost of search with the marginal probability of detecting a crime. Hence total search P solves $C'(P) = r(\rho)$.

[13]To be precise, this argument only implies that there exists some $\rho \in (0,1)$ such that the ex ante probability of a signal belonging to $S_P \cap (0,1) \setminus \{\rho\}$ is 0 (instead of this set being empty). However, we can replace all signals in $S_P \cap (0,1) \setminus \{\rho\}$ with signal ρ without affecting the equilibrium crime rate.

In the second step, we show that if $\{\rho, 1\} \subset S_P$, we can replace signals ρ and 1 with the same signal σ to increase the search probability allocated to signal *not*. Indeed, after pooling signals ρ and 1, the posterior crime rate $r(\sigma)$ for signal σ satisfies $r(\sigma) > r(\rho)$. Thus if we let the police choose an optimal search strategy, the police will choose total search capacity $\tilde{P} > P$, because their marginal return to search at P is now $r(\sigma) - C'(P) > r(\rho) - C'(P) = 0$. This pooling also reduces the police's information and hence the mass of successful searches. As a result, pooling signals ρ and 1 increases the expected search probability conditional on signal *not* and relaxes its obedience constraint.

We now have $S_P = \{0, \sigma\}$ or $\{\sigma\}$ for some $\sigma > 0$, and the obedience constraint for signal *not* holds.[14] We then gradually increase cutoff type c defined in Lemma 2 while maintaining conditional distributions $\hat{\pi}(\cdot|crime)$ and $\hat{\pi}(\cdot|not)$ that garble an agent's signals for the police. Increasing cutoff c decreases the mass of agents with signal *crime* and the posterior crime rate for every signal. Thus as c increases, the police will choose lower search rates for every signal. At some cutoff $c^* \geq c$, the obedience constraint for signal *not* binds.[15] The obedience constraint for signal *crime* also holds, because the assumption $C'\left(\int_0^1 xF(dx)\right) > 1$ implies that the police's optimal searches never make all agents weakly prefer to abstain from crimes.

In the last step, we consider two cases. If $S_P = \{\sigma\}$, we obtain the desired result where $\alpha^* = 0$ (signal 0 is irrelevant if $\alpha^* = 0$). Otherwise, the signal space is $\{0, \sigma\}$. In this case, we replace the police's signal with the corresponding posterior crime rate, resulting in a signal space $\{r_0, r_\sigma\}$. We then split signal r_0 into signals t_0 and t_σ that have posterior crime rates 0 and $r_\sigma > r_0$, respectively. The police's optimal strategy remains the same, i.e., the police search signal r_σ with a positive probability and never search signals t_0 and t_σ. But signals t_σ and r_σ have the same posterior crime rate, so we can pool them into one signal, say u_σ. The police's signal is now binary, i.e., signal r_0 has posterior crime rate 0 and search rate 0, and signal u_σ has a positive crime rate and search probability. In terms of action recommendations, the police's signal space becomes $S_P^* = \{0, \rho^*\}$ for some $\rho^* \in (0, 1]$.

[14] We cannot have $S_P = \{0\}$ because if the police do not search, all agents commit crimes, which incentivizes the police to choose a positive search rate because of $C'(0) < 1$.

[15] The reason is as follows. If $c^* = 1$, then all agents receive signal *not*, the police choose search probability 0, and the obedience constraint for signal *not* is violated. Thus at some $c^* \in (0, 1)$, the agents become indifferent between the two actions after receiving signal *not*.

We now have signal structure (S_P^*, S_A^*, π^*) such that: $S_P^* = \{0, \rho^*\}$ and $S_A^* = \{crime, not\}$; the agents with types above and below cutoff c^* receive signals *crime* and *not*, respectively; and the police's signal divides the population into two groups, one with zero posterior crime rate and the other with a positive posterior crime rate. This signal structure and equilibrium satisfy Parts 1 and 2 of the proposition. Part 3 follows from the second step. □

To see why partial revelation to the police can be optimal when the search capacity is endogenous, suppose that a mass $r \in (0, 1)$ of agents are committing crimes. Compare the following two cases: In Case 0, the police have no information. In Case α, the police can privately identify a fraction α of innocents (or equivalently, a mass $\alpha(1-r)$ of innocents). Moving from Case 0 to Case α affects the police's strategy in two ways. First, the innocents are on average less likely to be exposed to search, because the police do not search a fraction α of them. Second, the police choose a higher total search capacity in Case α than Case 0, because the police can detect a crime with probability $\frac{r}{1-\alpha(1-r)} > r$ by searching a mass $1-\alpha(1-r)$ of unidentified agents in Case α, whereas the probability of detecting a crime is r in Case 0. The crime-minimizing signal structure involves partial revelation (i.e., $\alpha^* > 0$ in Proposition 1) when the second effect dominates, so that the overall costs for the agents of committing crimes increase as we move from Case 0 to Case α.

At the same time, consistent with the case of exogenous search capacity, the crime-minimizing signal structure does not enable the police to identify criminals. To see why, consider Case β in which the police can identify a fraction β of criminals. In contrast to Case α, moving from Case 0 to Case β could reduce the police's effort, because the probability of detecting a crime in the unidentified population is $\frac{(1-\beta)r}{1-\beta r} < r$. Moreover, as in the baseline model, the information about criminals distorts the allocation of searches and reduces the deterrent effect of search. Thus, allowing the police to predict crimes could increase a crime rate by both reducing search effort and distorting the allocation of it.

5.2 Uniform-Quadratic Example

We now turn to the original setup in which agents observe their types. One candidate signal structure that may implement the same outcome as in the relaxed problem (Proposition 1) is the following: First, the police privately identify a fraction α^* of agents who have types below c^*. Second, as to the remaining population, the police observe the truth-or-noise signal with cutoff c^*, i.e., a signal coincides with an agent's true type if the type is below c^*, and otherwise the signal is a noise drawn from $F(\cdot|\tilde{x} \le c^*)$. In equilibrium, the police search agents with signal s with probability $\frac{s}{1-\alpha^*}$. Facing such a strategy, types below c^* are indeed willing to not commit crimes, because any agent with type $x < c^*$ believes that they will be searched with probability $(1-\alpha^*)\frac{x}{1-\alpha^*} = x$.

However, in general, the above strategy profile may not be an equilibrium. For example, the police's search strategy may be infeasible because it may specify a search rate above 1. Also, some types above c^*, who anticipate search probability $\frac{\mathbb{E}[\tilde{x}|\tilde{x} \le c^*]}{1-\alpha^*}$, may abstain from crimes, which leads to a different outcome from the relaxed problem. At the same time, the signal structure constructed above indeed minimizes a crime rate in some settings (see the appendix for the proof):

Claim 1. *In the original setup where the agents observe their types, suppose that $F = U[0,1]$ and $C(P) = \frac{L}{2}P^2$ with $L \ge \frac{3+\sqrt{5}}{4} \approx 1.31$. There is a signal structure that attains the same crime rate as in the relaxed problem (Proposition 1). This signal structure reveals a fraction $\alpha^* \in [0,1)$ of the agents with types below c^* and discloses the truth-or-noise signal with cutoff c^* regarding the rest of the agents. If $L < 2$, we have $\alpha^* = 2 - \sqrt{2L} > 0$. If $L \ge 2$, we have $\alpha^* = 0$, i.e., the police observe the truth-or-noise signal with cutoff c^*.*

As an example, we set $L = \frac{72}{49}$ and describe the crime-minimizing outcome. The signal structure becomes as follows: If an agent has type $x \le \frac{7}{12}$, the police observe signal 0 or x with probability $\frac{2}{7}$ or $\frac{5}{7}$, respectively. If an agent has type $x > \frac{7}{12}$, the police observe signal s that is drawn from $U[0, \frac{7}{12}]$ independent of the true type. Thus signal $s = 0$ reveals that an agent's type is below $\frac{7}{12}$, whereas any signal $s \in (0, \frac{7}{12}]$ indicates that the agent's type is s or uniformly distributed on $[\frac{7}{12}, 1]$ with equal probability.

In equilibrium, the police search any agent with signal $s \in (0, \frac{7}{12}]$ with probability $\frac{7}{5}s$ and never search agents with signal 0. If an agent has type $x < \frac{7}{12}$, the expected probability of search is $\frac{5}{7} \cdot \frac{7}{5}x = x$, so they weakly prefer to abstain from a crime. If an agent has type $x \ge \frac{7}{12}$, the expected probability of search is $\mathbb{E}_{x \sim U[0,\frac{7}{12}]}[\frac{7}{5}x] < x$, so they prefer to commit a crime. Finally, any signal indicates that the agent is committing a crime with probability $\frac{1}{2}$, so the police are indifferent about how to distribute search capacity across agents with positive signals.[16] The above search strategy leads to total search capacity $P^* = \frac{49}{144}$. The police's marginal search cost is then $C'(P^*) = LP^* = \frac{1}{2}$, which indeed equals to the marginal probability of detecting a crime.

References

Becker, Gary S (1968), "Crime and punishment: an economic approach." *Journal of Political Economy*, 76, 169–217.

Bergemann, Dirk and Stephen Morris (2019), "Information design: A unified perspective." *Journal of Economic Literature*, 57, 44–95.

Bjerk, David (2007), "Racial profiling, statistical discrimination, and the effect of a colorblind policy on the crime rate." *Journal of Public Economic Theory*, 9, 521–545.

Blackwell, David (1951), "Comparison of experiments." *Proceedings of the second Berkeley symposium on mathematical statistics and probability.*

Blackwell, David (1953), "Equivalent comparisons of experiments." *The annals of mathematical statistics*, 24, 265–272.

Brayne, Sarah (2020), *Predict and Surveil: Data, discretion, and the future of policing.* Oxford University Press, USA.

Che, Yeon-Koo, Jinwoo Kim, and Konrad Mierendorff (2023), "Predictive enforcement." *mimeo.*

Cotton, Christopher and Cheng Li (2015), "Profiling, screening, and criminal recruitment." *Journal of Public Economic Theory*, 17, 964–985.

Curry, Philip A and Tilman Klumpp (2009), "Crime, punishment, and prejudice." *Journal of Public economics*, 93, 73–84.

[16]The posterior crime rate for signal $s > 0$ is $\frac{1-F(c^*)}{1-F(c^*)+(1-\alpha)F(c^*)} = \frac{\frac{5}{12}}{\frac{5}{12}+\frac{5}{7}\cdot\frac{7}{12}} = \frac{1}{2}$.

Dworczak, Piotr and Giorgio Martini (2019), "The simple economics of optimal persuasion." *Journal of Political Economy*, 127, 1993–2048.

Eeckhout, Jan, Nicola Persico, and Petra E Todd (2010), "A theory of optimal random crackdowns." *American Economic Review*, 100, 1104–35.

Guo, Yingni and Eran Shmaya (2019), "The interval structure of optimal disclosure." *Econometrica*, 87, 653–675.

Hernández, Penélope and Zvika Neeman (2022), "How bayesian persuasion can help reduce illegal parking and other socially undesirable behavior." *American Economic Journal: Microeconomics*, 14, 186–215.

Jung, Christopher, Sampath Kannan, Changhwa Lee, Mallesh Pai, Aaron Roth, and Rakesh Vohra (2020), "Fair prediction with endogenous behavior." In *Proceedings of the 21st ACM Conference on Economics and Computation*, 677–678.

Kamenica, Emir (2019), "Bayesian persuasion and information design." *Annual Review of Economics*, 11, 249–272.

Kamenica, Emir and Matthew Gentzkow (2011), "Bayesian persuasion." *American Economic Review*, 101, 2590–2615.

Knowles, John, Nicola Persico, and Petra Todd (2001), "Racial bias in motor vehicle searches: Theory and evidence." *Journal of Political Economy*, 109, 203–229.

Kolotilin, Anton, Timofiy Mylovanov, and Andriy Zapechelnyuk (2022), "Censorship as optimal persuasion." *Theoretical Economics*, 17, 561–585.

Lazear, Edward P (2006), "Speeding, terrorism, and teaching to the test." *The Quarterly Journal of Economics*, 121, 1029–1061.

Liang, Annie, Jay Lu, and Xiaosheng Mu (2022), "Algorithmic design: Fairness versus accuracy." In *Proceedings of the 23rd ACM Conference on Economics and Computation*, 58–59.

Pearsall, Beth (2010), "Predictive policing: The future of law enforcement." *National Institute of Justice Journal*, 266, 16–19.

Perry, Walt L (2013), *Predictive policing: The role of crime forecasting in law enforcement operations*. Rand Corporation.

Persico, Nicola (2002), "Racial profiling, fairness, and effectiveness of policing." *American Economic Review*, 92, 1472–1497.

Persico, Nicola (2009), "Racial profiling? detecting bias using statistical evidence." *Annu. Rev. Econ.*, 1, 229–254.

Persico, Nicola and Petra E Todd (2005), "Passenger profiling, imperfect screening, and airport security." *American Economic Review*, 95, 127–131.

Porto, Edoardo Di, Nicola Persico, and Nicolas Sahuguet (2013), "Decentralized deterrence, with an application to labor tax auditing." *American Economic Journal: Microeconomics*, 5, 35–62.

Shaked, Moshe and J George Shanthikumar (2007), *Stochastic orders*. Springer.

Smolin, Alex and Takuro Yamashita (2022), "Information design in concave games." *arXiv preprint arXiv:2202.10883*.

Vasquez, Jorge (2022), "A theory of crime and vigilance." *American Economic Journal: Microeconomics*, 14, 255–303.

Appendix: Proof of Claim 1

The proof consists of two steps: We first use Proposition 1 to solve the relaxed problem. We then turn to the original setup and construct a signal structure for the police that attains the same crime rate as in the relaxed problem.

To solve the relaxed problem, we focus on signal structures that take the form described in Figure 2. Instead of parameters (c^*, α^*, p^*), we use (c, α, p) to indicate that they may not be the crime-minimizing signal structure. Recall that α is the probability with which the police observe signal 0 conditional on that an agent observes signal *not*. Because we focus on an equilibrium in which types above some cutoff commit crimes, a crime rate r pins down the cutoff type through $c = F^{-1}(1-r)$.

We fix $\alpha \in [0,1]$ arbitrarily and then determine cutoff type c and the unique positive search probability ρ from the mutual best responses of the agents and the police. By Part 3 of Proposition 1, the equilibrium crime rate $r(\alpha)$ is determined by the condition that the police's optimal search strategy given crime rate $r(\alpha)$ makes the agents who observe signal *not* indifferent between committing a crime and not.

We derive the expected search probability given signal *not*. As in Figure 2, if the crime rate is r, the posterior crime rate for signal ρ is $\frac{r}{r+(1-r)(1-\alpha)}$. The police's mass of searches P then solves the first-order condition $C'(P) = LP = \frac{r}{r+(1-r)(1-\alpha)}$, or $P = \frac{1}{L} \cdot \frac{r}{r+(1-r)(1-\alpha)}$. Thus the expected search probability conditional on signal *not* is $(1-\alpha)\frac{P}{r+(1-r)(1-\alpha)}$, or

$$I(r,\alpha) \triangleq \frac{(1-\alpha)r}{L[r+(1-r)(1-\alpha)]^2}.$$

The binding obedience constraint for signal *not* is written as

$$\begin{aligned} & I(r,\alpha) = \mathbb{E}_{x\sim U[0,1]}[x|x \leq F^{-1}(1-r)] \\ \iff & I(r,\alpha) = \frac{1-r}{2}. \end{aligned} \tag{3}$$

At $r=0$, we have $I(0,\alpha) = 0 < \frac{1}{2}$. At $r=1$, we have $I(1,\alpha) = \frac{1-\alpha}{L} > 0$. Moreover, $I(r,\alpha)$ is concave in r.[17] Thus as a function of r, $I(r,\alpha)$ crosses $\frac{1-r}{2}$ exactly once and from below, i.e., equation (3) has a unique solution in terms of r. Let $r(\alpha)$ be the solution. Then the minimal crime rate in the relaxed problem is given by $r^* \triangleq \min_{\alpha\in[0,1]} r(\alpha)$. However, instead of solving this minimization problem, we first derive $\alpha(r) \triangleq \arg\max_{\alpha\in[0,1]} I(r,\alpha)$ and then determine the minimal crime rate r^* through $I(r,\alpha(r)) = \frac{1-r}{2}$. Because $I(r,\alpha)$ is concave in

[17]We have $\frac{\partial I}{\partial r} = L^{-1}\frac{(1-\alpha)(1-\alpha(1+r))}{1-\alpha(1-r)}$ and $\frac{\partial^2 I}{\partial r^2} = L^{-1}\frac{-2\alpha(1-\alpha)^2}{[1-\alpha(1-r)]^2} \leq 0$.

α,[18] we can use the first-order condition to solve $\max_{\alpha\in[0,1]} I(r,\alpha)$. The solution is as follows:

$$\alpha(r) = \begin{cases} \frac{1-2r}{1-r} & \text{if} \quad r \le \frac{1}{2}, \\ 0 & \text{if} \quad r \ge \frac{1}{2}. \end{cases}$$

and

$$I(r,\alpha(r)) = \begin{cases} \frac{1}{4L(1-r)} & \text{if} \quad r \le \frac{1}{2}, \\ \frac{r}{L} & \text{if} \quad r \ge \frac{1}{2}. \end{cases}$$

Solving $I(r,\alpha(r)) = \frac{1-r}{2}$, we obtain the minimized crime rate under the relaxed problem:

$$r^* = \begin{cases} 1 - \frac{1}{\sqrt{2L}} & \text{if} \quad L \le 2, \\ \frac{L}{2+L} & \text{if} \quad L \ge 2. \end{cases}$$

Because the type distribution is uniform, a crime rate of r^* means that an agent commits a crime if and only if their type exceeds $c^* = 1 - r^*$.

We now show that if $L \ge L^* = \frac{3+\sqrt{5}}{4}$, we can implement crime rate r^* in the original problem. To do so, we modify the signal structure in Theorem 2 as follows: If $x < 1 - r^*$, with probability $\alpha^* = \alpha(r^*)$, the police observe signal 0. Other parts of the signal structure follow the truth-or-noise signal structure with cutoff c^*: The police observe signal x with probability $1 - \alpha^*$ if $x \le c^*$ and observe signal $s \sim F(\cdot|\tilde{x} \ge c^*)$ whenever $x \ge c^*$.

If $L \ge L^*$, this signal structure has an equilibrium in which each agent commits a crime if and only if $x > c^*$, and the police search agents with signal s with probability $\frac{s}{1-\alpha^*}$. The agents' strategies are optimal: Any agent with a type below c^* is indifferent between committing a crime and not because they anticipate search probability $(1-\alpha^*)\frac{x}{1-\alpha^*} = x$ in expectation. Any type $x \ge c^*$ will face a search probability of $\frac{1-r^*}{2(1-\alpha(r^*))} < 1 - r^* \le x$, where the first inequality uses $\alpha^* < 1/2$, which follows from $L \ge L^*$. The police's strategy is also optimal: The police never search signal 0 and are indifferent regarding how to allocate a given mass of searches across the positive signals. The choice of a total search capacity is

[18]We have $\frac{\partial I}{\partial \alpha} = L^{-1}\frac{-r[2r-1+\alpha(1-r)]}{1-\alpha(1-r)}$ and $\frac{\partial^2 I}{\partial \alpha^2} = L^{-1}\frac{-2r^2(1-r)}{[1-\alpha(1-r)]^2} \le 0$.

also optimal: Indeed, the posterior crime rate for any positive signal is $\frac{r^*}{r^*+(1-r^*)(1-\alpha^*)}$, so the total search capacity induced by the above search strategy equates the marginal cost with the marginal crime rate because of the police's first-order condition in the relaxed problem. Also, $L \geq L^*$ ensures that the highest search probability $\frac{c^*}{1-\alpha^*}$ is below 1, so the police's strategy is well-defined. Finally, if $L \in [L^*, 2)$, then we have $r^* = 1 - \frac{1}{\sqrt{2L}}$ and $\alpha(r^*) = \frac{1-2r^*}{1-r^*} = 2 - \sqrt{2L}$. If $L \geq 2$, we have $\alpha^* = 0$, so the police's signal reduces to the truth-or-noise signal structure.

$\square$