Information and Policing

Shota Ichihashi[*]

March 3, 2025

Abstract

Agents decide whether to commit a crime based on their private types, which capture their heterogeneous returns from a crime. The police have information about these types. The police search agents, without commitment, to detect crime subject to a search capacity constraint. The deterrent effect of policing is lost when the police have full information about agents' types. The crime-minimizing information structure prevents the police from identifying agents who face high returns from a crime, while still allowing them to adjust their search intensities based on the types of agents who face low returns from a crime. The result extends to the case in which the police endogenously choose search capacity at a cost.

Keywords: information design, crime, policing, privacy

[*]Queen's University, Department of Economics. Email: shotaichihashi@gmail.com. I thank the editor, Marco Ottaviani, the associate editor, and anonymous referees for their comments and suggestions. I am especially grateful to Christopher Cotton and Moritz Meyer-ter-Vehn for providing valuable feedback on the paper at different stages. For valuable suggestions and comments, I also thank Yeon-Koo Che, Navin Kartik, Qingmin Liu, Ruslan Momot, Alex Smolin, and participants in seminars and lectures at Osaka University, Columbia Business School, Toulouse School of Economics, University of Michigan (Ross School of Business), Queen's University, Michigan State University, University of Bristol, University of Warwick, and University of Oxford.

1 Introduction

Law enforcement agencies increasingly rely on data and algorithms to predict crime (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. Although this trend aims to improve law enforcement effectiveness, it also raises concerns. For example, the EU's "Artificial Intelligence Act" prohibits certain uses of AI for crime prediction.[1]

Motivated by recent discussions, I study how information available to law enforcement affects its ability to deter crime, when the law enforcer cannot commit to how the information will be used. I examine this question from the perspective of information design.

The model consists of a unit mass of agents and a law enforcement decision maker, called the "police." Agents have private types, which capture their heterogeneous returns from a crime. At the outset, the police observe a signal about each agent's type, according to an exogenous signal structure. Each agent decides whether to commit a crime (e.g., tax evasion, fraud, trafficking), and then the police, without observing individual actions, choose how to distribute their limited search capacity across different types of agents.

The situation is modeled as a simultaneous-move game because neither the police nor agents observe the others' actions prior to choosing their own strategy. Although the police might ideally want to reduce the overall crime rate, the simultaneous-move assumption means they cannot directly influence it through their search strategy. Therefore, their focus shifts to detecting crime instead of deterring them.

I first show that information restriction is necessary for crime deterrence: If the police have full information about agents' types and face a limited search capacity, then in any equilibrium every agent commits a crime with probability 1. Intuitively, the fully informed police will focus their limited search capacity on agents who commit a crime with positive probability and never search for those who abstain from crime. However, if agents expect not to be searched, they will commit a crime for sure. The only outcome consistent with these conditions is that all agents commit a crime: The police allocate search capacity so

[1] See https://www.europarl.europa.eu/news/en/press-room/20230609IPR96212/meps-ready-to-negotiate-first-ever-rules-for-safe-and-transparent-ai, accessed March 3, 2025.

that 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 because this probability is too low to deter their crime.

The main result characterizes the signal structure that minimizes the crime rate. As shown in the first result, providing the police with detailed information may erode the deterrent effect of searches because the police will prioritize searching agents who are most actively engaged in crime in equilibrium. However, if the police have no information, they may waste search resources by searching some agents with unnecessarily high probability, even when these agents have low incentives to commit a crime. The crime-minimizing signal structure provides the police with information to reduce wasteful searches without eroding the deterrent effect.

Specifically, to derive the crime-minimizing signal structure, I first solve a relaxed problem in which I choose a joint signal structure for the police and agents to minimize the crime rate under the assumption that agents do not observe their types. The relaxed problem has a solution in which the police receive no information and randomly search agents with a constant probability, and each agent learns only whether their type exceeds some cutoff. The relaxed problem identifies a lower bound of possible crime rates that can arise in the original problem. I then turn to the original problem—in which agents observe their types—and construct a signal structure for the police that attains the same outcome as the solution to the relaxed problem.

The crime-minimizing signal structure shares a limited amount of information with the police: It does not allow them to identify agents with the highest propensity for criminal activity. Instead, it pools such high-crime types with lower-crime types, ensuring equal screening across these groups, which effectively deters crime among the lower types. At the same time, the crime-minimizing signal structure assigns different signals to different low types, which allows the police to avoid searching agents with excessively high probability when these agents can easily be deterred.

The model shows how increased data or technology that improves law enforcement's ability to predict crime could backfire, leading to higher criminal activity. At the same time, the crime-minimizing information policy is not necessarily to prohibit the use of any information, but to structure the information shared with police to ensure that policing efforts are directed toward those who can be deterred from crime.

In Section 4, I demonstrate the relevance of the crime-minimizing signal structure in several extensions. First, I show that a version of the crime-minimizing signal structure remains optimal when the objective is to minimize a combination of the crime rate and the search costs imposed on innocents (Section 4.1). Second, I solve the case in which the police have full commitment power (Section 4.2). Third, I show that the crime-minimizing signal structure applies to a model in which a designer partially controls both a search strategy and the police's information (Section 4.3). Finally, I provide a sufficient condition under which the main result holds when the police can endogenously choose total search capacity at a cost (Section 4.4).

Contribution to the Literature. This paper is related to the literature on the application of Bayesian persuasion and information design to law enforcement.[2] While previous research (e.g., Lazear 2006; Eeckhout, Persico, and Todd 2010; Hernández and Neeman 2022) has explored how disclosing information *to* potential criminals can deter crime, this paper examines how providing information *about* potential criminals to police affects crime. The model provides insights into the potential unintended consequences of the increasing use of data and technology to predict and prevent crime, including the use of crime prediction software and mass surveillance systems (see Berk (2021) for a survey).

Methodologically, this paper addresses an information design problem with many players, states, and actions (for the police)—a setting that is less understood than Bayesian persuasion with one sender and one receiver (see Smolin and Yamashita (2022) for a discussion). I contribute to the literature by introducing a relaxed-problem approach that simplifies analysis of this potentially complex information design problem.

A key feature of my model is that the player receiving information (the police) is distinct from those agents whose behavior we ultimately want to influence. This leads to a novel type of information manipulation, embodied in the "truth-or-noise" signal structures that

[2]See Kamenica (2019) and Bergemann and Morris (2019) for surveys of Bayesian persuasion and information design.

solve my problem. These structures satisfy the constraints of both police and agent behavior and differ substantially from previously studied optimal signal structures, such as monotone partitional signals, censorship policies, and nested intervals (e.g., Dworczak and Martini 2019; Guo and Shmaya 2019; Kolotilin, Mylovanov, and Zapechelnyuk 2022).[3]

This paper also contributes to the economic literature on crime and policing, which starts from Becker (1968). A question regarding what potentially observable information about agents should or should not be used for policing is discussed in the context of racial profiling (Knowles, Persico, and Todd 2001; Persico and Todd 2005; Bjerk 2007; Persico 2009).

My model's structure—particularly the timing of actions and the way payoffs are determined—is similar to that of Persico (2002), who examines whether mandating police to use a more equitable search strategy reduces crime. His model can be interpreted as a model of information provision to police; a "fair" search strategy is equivalent to providing police no information about agent type. While Persico (2002) examines the impact of a given level of information on crime reduction, I take a broad approach by considering all possible information structures and identifying the one that minimizes crime. This allows me to uncover a general principle about optimal information provision to law enforcement, independent of specific parameter values.[4]

The rest of the paper is organized as follows. Section 2 presents the baseline model and the benchmark in which the police have full information. Section 3 characterizes the crime-minimizing signal structure. In Section 4, I consider several extensions that relax various assumptions in the baseline model.

[3]The relaxed problems examined in this paper are also related to Bayesian persuasion with moral hazard, in which the state distribution is endogenously determined by a player's action (Rodina, 2017; Boleslavsky and Kim, 2018; Zapechelnyuk, 2020; Hörner and Lambert, 2021).

[4]My model abstracts away from other important considerations, such as the design of judicial systems, richer responses by potential criminals and victims, the fairness of predictive algorithms, and the endogenous generation of data (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; Che, Kim, and Mierendorff 2024; Li and Cotton 2023).

2 Model

The model consists of police and a unit mass of agents. Each agent $i \in [0,1]$ has some underlying returns on crime, called their *type*, $x_i \in [0,1]$. Each agent observes their type, which captures the individual characteristics that affect their propensity for crime, including their opportunity cost of crime (e.g., legal earning opportunities) or crime opportunities specific to certain locations or times. Types are independently and identically drawn from distribution function $F \in \Delta[0,1]$, which has a positive density f and is commonly known.[5] $\mathbb{E}_F[\cdot]$ is the expectation operator under F. Also, $F(\cdot|\tilde{x} \leq c)$ is the conditional distribution of F on $[0,c]$ and $\mathbb{E}_F[\cdot|\tilde{x} \leq c]$ is the corresponding expectation operator.

The police learn 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 $\pi(\cdot|x) \in \Delta S$ over signals for each type x. For each agent $i \in [0,1]$, the police observe a signal $s_i \in S$ drawn from distribution $\pi(\cdot|x_i)$. Conditional on types, signals are independent across agents. The signal structure is exogenous and commonly known, but only the police observe realized signals.[6]

Given the signal structure, the police and agents play the following simultaneous move game: Each agent decides whether to commit a crime and, simultaneously, 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.,

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

Agent's Payoff The payoff for an agent for committing a crime is $x - \rho$, where x is the agent's type and ρ is the search probability the agent faces. The payoff for not committing

[5] I write ΔX for the set of all probability distributions on a set X. See, e.g., Sun (2006) for a formal treatment of a continuum of independent random variables and the corresponding law of large numbers.

[6] Because the signal structure is exogenous, we should view the results as comparative statics that examine how information affects the deterrent effect of policing. A related but separate question is which signal structure would be chosen if information were endogenous; the answer would depend on the objective of the player who chooses the signal structure and whether this choice is observable to the agents.

a crime is 0.[7] For a given signal structure and a search strategy, an agent's expected payoff for committing a crime is written as

$$x - \int_S p(s)\,\mathrm{d}\pi(s|x).$$

Successful Search and Crime Rate Fix any strategy profile, and let $a(x)$ denote the probability with which agents of type x commit a crime. Define the *mass of successful searches* as the mass of agents who commit a crime and are searched by the police, i.e.,

$$\sigma \triangleq \int_0^1 \int_S p(s)\,\mathrm{d}\pi(s|x)a(x)\,\mathrm{d}F(x). \tag{2}$$

Define the *crime rate* as the mass of agents who commit a crime, i.e.,

$$r \triangleq \int_0^1 a(x)\,\mathrm{d}F(x). \tag{3}$$

Police's Payoff The police's payoff is given by

$$\sigma - \lambda r, \tag{4}$$

where parameter $\lambda \in \mathbb{R}_+$ captures the degree to which the police prioritize deterring crime over uncovering crime. λ is unconstrained, so the police may place a large weight on crime deterrence. However, under the simultaneous-move assumption, the police take the crime rate r as exogenous because it does not directly depend on their search strategy (see (3)). Therefore, for any given λ, in equilibrium, the police act as if their payoffs depend only on

[7] This payoff specification is equivalent to the following richer setup: The payoff for committing a crime is $U(y, \rho)$, which strictly increases in an agent's type y; strictly decreases in search probability ρ; and has a threshold search probability $\hat{p}(y)$ that solves $U(y, \hat{p}(y)) = 0$ for each type y (without loss, the payoff for not committing a crime is normalized to 0). An agent will commit a crime if $\hat{p}(y) > \rho$. But once I redefine the agent's type as $x = \hat{p}(y)$, this richer setup and the original setup lead to the same set of best responses by agents under any search strategy. This setup subsumes $U(y, \rho) = (1 - \rho)y - L\rho$—i.e., an agent enjoys the returns on crime if they are not searched, but incurs a loss of L if they are.

the mass of successful searches, i.e., $\lambda = 0$.[8]

The solution concept is Bayesian Nash equilibrium (BNE). For expositional simplicity, I focus on equilibria in which the agents of the same type commit a crime with the same probability. Hereafter, an *equilibrium* refers to a BNE that satisfies this condition.

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

Assumption 1. The primitives, F and $\overline{P}$, satisfy

$$\overline{P} < \int_0^1 x \, \mathrm{d}F(x). \tag{5}$$

The assumption implies a fundamental limitation on police resources: They cannot completely eliminate crime. To achieve zero crime, each agent of type x would need to be searched with a probability of at least x. However, Assumption 1 states that the police do not have sufficient capacity for this. As a result, a positive mass of agents, anticipating search probabilities below their types, commit a crime.

Consider a benchmark scenario in which the police have no information—i.e., the distribution $\pi(\cdot|x)$ over signals is independent of type x. In this case, every agent faces the same search probability, equal to the police's total search capacity $\overline{P}$.[9] Because every type $x > \overline{P}$ faces a search probability less than their return from crime, they commit a crime. This results in a total crime rate of $1 - F(\overline{P})$.

The main focus is on a signal structure that minimizes the crime rate (3) 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*.

---

[8]This observation does not imply that the police are better off when the crime rate is high: If λ is high but the crime rate is also high, the police in this model could have been better off—in the sense of obtaining a greater payoff—had they publicly committed to a different search strategy or faced a different signal structure. The model is not suitable for analyzing police welfare, because we can derive arbitrary welfare implications depending on λ.

[9]The reason is as follows. When the underlying signal structure is uninformative, every agent faces search probability $\int_S p(s) \, \mathrm{d}\pi(s|x)$, which is independent of x. Then, $\int_S p(s) \, \mathrm{d}\pi(s|x) > \overline{P}$ violates the search capacity constraint. The other case $\int_S p(s) \, \mathrm{d}\pi(s|x) < \overline{P}$ also leads to a contradiction, because the police could increase search rates and thus the mass of successful searches. Hence $\int_S p(s) \, \mathrm{d}\pi(s|x) = \overline{P}$ must hold.

The equilibrium search strategy differs from the crime-minimizing strategy. Even though the police might care about the crime rate, their actions do not deter crime but only affect the amount of detected crime. Therefore, they maximize the mass of successful searches. As a result, simply giving the police more information does not necessarily reduce crime. Instead, the ideal information structure for minimizing crime strategically limits what the police know and thus encourages a search strategy that effectively deters criminal activity.

2.1 Discussion of Timing and Payoffs

The paper's results hinge on two assumptions: First, the police and agents move without observing the actions of the other group; in particular, the police cannot publicly pre-commit to a search strategy. Second, the police are rewarded for uncovering crime. Together, these assumptions imply that the police act to maximize the number of successful searches. This section motivates these assumptions (see Section 4.2 for consequences of relaxing them).

First, the literature has considered police pre-committing to their strategy as well as police moving simultaneously with potential criminals (see, e.g., Eeckhout et al. (2010) for the former). The validity of each assumption depends on the context. For example, the literature on decentralized law enforcement, such as Persico (2002) and Porto et al. (2013), adopts the simultaneous-move assumption by viewing the "police" as a collection of individual officers or auditors. This assumption arises from the idea that the action of each individual officer does not directly influence the decisions of potential criminals (Online Appendix C formalizes this idea). The police's lack of commitment might also stem from the notion that a signal structure captures a predictive algorithm used by a law enforcement agency. Predictions generated by an algorithm (i.e., signals) would be invisible to the public and too complex to describe in advance, making it difficult for the police to commit to a predetermined search strategy.

The validity of the second assumption—that the police are (partly) rewarded for uncovering crime—also depends on the context. For example, papers such as Eeckhout et al. (2010) present evidence favoring pure crime minimization. In contrast, papers such as Baicker and Jacobson (2007), Makowsky and Stratmann (2011), Nagin et al. (2015), and Owens and Ba

(2021) argue that the incentives of police are tied to detecting and resolving crime.[10] Stashko (2023) analyzes police stop-and-search data in the U.S. and concludes that "empirical evidence is consistent with a model of arrest maximization and inconsistent with a model of crime minimization."

2.2 Preliminary Analysis: The Fully Informed Police

To demonstrate that restricting the police's information is necessary for crime deterrence, I begin with an analysis of the fully informed police. Specifically, suppose that the signal structure is such that for every type $x \in [0,1]$, $\pi(\cdot|x)$ places probability 1 on $s = x$.

Theorem 0. *Suppose that the police have full information. An equilibrium exists, and in any equilibrium, almost every agent commits a crime with probability* 1*.*

Proof. First, I construct an equilibrium. Suppose that the police adopt a search strategy $p^*(x) = \frac{x\overline{P}}{\mathbb{E}_F[\tilde{x}]}$ for all $x \in [0,1]$. This search strategy induces the total search capacity of $\int_0^1 p^*(x)\,\mathrm{d}F(x) = \overline{P}$ and thus is feasible. Assumption 1 implies that $\overline{P} < \mathbb{E}_F[\tilde{x}]$, which means that $p^*(x) < x$ for all $x \in (0,1]$. As a result, every agent (except possibly $x = 0$) commits a crime with probability 1. The police then find it optimal to choose any search strategy that exhausts search capacity $\overline{P}$, including p^*. Thus we obtain an equilibrium.

Second, take any equilibrium. The police cannot choose $p(x) \geq x$ for (almost) every x, because by Assumption 1 it violates the search capacity constraint. Thus, the set $X \triangleq \{x \in [0,1] : x > p(x)\}$ has a positive mass, and any type in X commits a crime with probability 1. Let Y be the set of types that commit a crime with probability strictly below 1. If Y has a positive mass, the police must be allocating a positive mass of searches to types in Y, because otherwise types in Y would commit a crime for sure. We then obtain a contradiction, because the police could increase the mass of successful searches by shifting

[10] Nagin et al. (2015) note that rewards from making arrests "are institutionalized in organizational performance metrics that emphasize arrest and clearance," while "generally no tangible reward is given for preventing crime in the first place, precisely because it is a nonevent and therefore is difficult to measure." Similarly, Owens and Ba (2021) observe that CompStat-style reviews "create strong incentives for officers to provide a high level of engagement—in particular, to make arrests, issue citations, and lower crime."

search probabilities from types in Y to X.[11] Thus, the set Y must have measure zero—i.e., almost every agent commits a crime with probability 1. $\square$

The logic of Theorem 0 is simple: The police aim to maximize the mass of successful searches, and with full information, they can identify which agents will commit a crime in equilibrium. Thus, the police never search the agents who abstain from a crime. However, agents also commit a crime if they expect not to be searched. The only outcome consistent with these conditions is that all agents commit a crime. The result implies that the police's ability to predict crime—combined with their incentive to uncover crime and the lack of commitment in search strategy—eliminates the deterrent effect of policing.

In the literature, Goldman and Pearl (1976) and Persico (2002) provide sufficient conditions on feasible search strategies or signal structures under which providing police with information increases crime. Theorem 0 adds to this literature by showing that full information leads to maximal crime under a mild condition (Assumption 1).[12]

In contrast to the crime rate, the police's equilibrium strategy is not unique. For example, suppose that $F = U[0,1]$ (i.e., the uniform distribution on $[0,1]$) and $\overline{P} = \frac{1}{4}$. In the proof of Theorem 0, I constructed an equilibrium in which the police set $p(x) = \frac{x}{2}$ for every $x \in [0,1]$. The following is another equilibrium: The police search any type $x \leq \frac{1}{\sqrt{2}}$ with probability x and never search any type $x > \frac{1}{\sqrt{2}}$. In this equilibrium, each type $x \leq \frac{1}{\sqrt{2}}$ is indifferent between committing a crime and not, yet breaks ties for committing a crime. Indeed, if types below $\frac{1}{\sqrt{2}}$ did not commit a crime, the police would profitably deviate and search types above $\frac{1}{\sqrt{2}}$ instead. In general, under full information, a strategy profile is an equilibrium if and only if (i) the search capacity $\overline{P}$ is allocated across agents in such a way that almost every agent weakly prefers to commit a crime, and (ii) almost every agent commits a crime with probability 1.

[11] The police's deviation would be profitable even if the police's payoffs depend on a crime rate, because the deviation does not change the agents' actions (and thus the crime rate) under the simultaneous-move assumption.

[12] In the terminology of my paper, Proposition 4 of Persico (2002) shows that for a sufficiently small $\overline{P}$, providing police with a binary signal structure (i.e., $|S| = 2$) that satisfies a certain condition results in a higher crime rate than no information. Goldman and Pearl (1976) study a related complete-information zero-sum game between an inspector and an inspectee and construct an equilibrium with crime rate 1 (see p. 194). However, their proof does not apply to my model, which is not a zero-sum game.

Remark 1 (Type-dependent Payoffs). Theorem 0 relies on the assumption that the police's payoff from successful searches is independent of the underlying types of agents. To see this, suppose that, in the previous example, the police's payoff from searching an agent who committed a crime is equal to their type, x (in the baseline model, this payoff is 1 for any type x). Then, in equilibrium, the police search any type $x \geq \frac{1}{\sqrt{2}}$ with probability x and never search any type $x < \frac{1}{\sqrt{2}}$. Any type $x \geq \frac{1}{\sqrt{2}}$ commits a crime with probability $\frac{1}{\sqrt{2}x}$, and any type $x < \frac{1}{\sqrt{2}}$ commits a crime with probability 1.[13] The crime rate is strictly below 1, because the police with type-dependent payoffs can credibly search agents even when they commit a crime with low probability. In Appendix A, I construct an equilibrium when the police have full information and earn arbitrary type-dependent payoffs from searching criminals.

3 Crime-minimizing Signal Structure

I now turn to the main analysis: characterization of the crime-minimizing signal structure. The analysis consists of two steps. First, I study a relaxed problem in which agents do not directly observe their own types. In this relaxed setting, I design an information structure for the police and agents to minimize crime. Because agents have less information in this relaxed version, the lowest crime rate we achieve here serves as a theoretical limit for how low the crime rate could possibly be in the original problem. In the second stage, I return to the original setting, in which agents know their types, and construct a specific information structure that achieves this theoretical minimum crime rate.

3.1 Relaxed Problem

The relaxed problem is defined as follows: Agents know the type distribution F but do not directly observe their type. Instead, information is provided by a *joint signal structure*,

[13] The police's search strategy is optimal. Indeed, searching any type $x < \frac{1}{\sqrt{2}}$ leads to a payoff of at most $\frac{1}{\sqrt{2}}$, whereas searching any type $x \geq \frac{1}{\sqrt{2}}$ leads to a payoff equal to $x \cdot \frac{1}{\sqrt{2}x} = \frac{1}{\sqrt{2}}$. Thus, a search strategy is optimal if and only if it searches only types above $\frac{1}{\sqrt{2}}$ and exhausts the search capacity $\overline{P}$. This condition holds for the search strategy presented here.

(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 i observes s_i^A and decides whether to commit a crime, and simultaneously, the police choose a search strategy $p : S_P \to [0,1]$ to maximize the mass of successful searches. The following result characterizes a crime-minimizing joint signal structure.

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

$$\mathbb{E}_F[\tilde{x}|\tilde{x} \leq \hat{c}] = \overline{P}. \tag{6}$$

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 joint signal structure (S_P, S_A, π) and equilibrium. Let $p : S_P \to [0,1]$ be the equilibrium search strategy and $r \in (0,1)$ the crime rate.[14] The proof consists of three steps. First, following the revelation principle of information design (e.g., Bergemann and Morris 2019), replace the agents' signals with action recommendations that replicate the equilibrium behavior of each type. The signal space for agents is now $S_A^* = \{crime, not\}$ and the obedience constraints hold:

$$\mathbb{E}[\tilde{x} - p(\tilde{s}^P)|crime] \geq 0 \quad \text{and} \quad \mathbb{E}[\tilde{x} - p(\tilde{s}^P)|not] \leq 0, \tag{7}$$

where the expectations in the first and second inequalities are with respect to the agent's type $\tilde{x}$ and the police's signal $\tilde{s}^P$ conditional on action recommendations *crime* and *not*, respectively. The obedience constraints ensure that each agent commits a crime after observing

[14] Assuming $r \in (0,1)$ is without loss of generality. By Assumption 1, there is no equilibrium with $r = 0$. Also, to solve the relaxed problem, we do not need to consider an equilibrium with $r = 1$, because we can ensure $r < 1$ by providing full information to agents and no information to the police.

signal *crime* and not after signal *not*. It holds that

$$\mathbb{E}[p(\tilde{s}^P)|crime] \geq \overline{P} \geq \mathbb{E}[p(\tilde{s}^P)|not]. \tag{8}$$

Indeed, if one inequality is violated, the other inequality must also be violated because of the binding search capacity constraint. However, if $\mathbb{E}[p(\tilde{s}^P)|crime] < \overline{P} < \mathbb{E}[p(\tilde{s}^P)|not]$, the mass of successful searches is $r\mathbb{E}[p(\tilde{s}^P)|crime] < r\overline{P}$. The police would then deviate and randomly search every agent with probability $\overline{P}$ and secure a higher mass of successful searches, $r\overline{P}$. This is a contradiction.

Second, replace the police's signal with an uninformative signal, $S_P^* = \{\phi\}$. Under the resulting signal structure (S_P^*, S_A^*, π^*), the police can only search every agent with probability $\overline{P}$. Inequalities (7) and (8) then imply

$$\mathbb{E}[\tilde{x}|crime] - \overline{P} \geq 0 \quad \text{and} \quad \mathbb{E}[\tilde{x}|not] - \overline{P} \leq 0,$$

i.e., the obedience constraints continue to hold.

Finally, the above two steps imply that, to derive the crime-minimizing joint signal structure, I can without loss of generality assume that the police receive no information, each agent receives an action recommendation, and the obedience constraints hold. The problem can be framed as a Bayesian persuasion. Here, an agent (i.e., the receiver) obtains payoff $x - \overline{P}$ for committing a crime and 0 for not. The sender discloses information about type $x \sim F$ to minimize the probability of the receiver's committing a crime. As shown in Gentzkow and Kamenica (2016) (Section IV.A), the solution is to disclose whether type x exceeds a cutoff $\hat{c}$ that solves (6), which deters types below $\hat{c}$ from committing a crime. The type distribution has a density and satisfies Assumption 1, so the cutoff $\hat{c} \in (0, 1)$ exists and is unique. □

Lemma 1 implies that, in the original setup, the minimum crime rate is attained if all types below $\hat{c}$ abstain from committing a crime. However, this outcome cannot arise if agents observe their type and the police have no information, because the resulting random search induces the types between $\overline{P}$ and $\hat{c}$ to commit a crime. Nevertheless, the next section shows that providing the police with partial information achieves the same crime rate.

3.2 Characterizing the Crime-minimizing Signal Structure

I now turn to the original problem, in which agents observe their type. First, I 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 (i.e., the truth); and for any $x > c$, distribution $\pi_c(\cdot|x)$ is independent of x and equals $F(\cdot|\tilde{x} \leq c)$ (i.e., the noise).

To understand the truth-or-noise signal structures, consider the police's beliefs about agent types after observing a signal realization from (S_c, π_c) (see Figure 1). For agent types below c, the signal coincides with their true type, whereas for types above c, the signal is a random draw from $F(\cdot|\tilde{x} \leq c)$, independent of the true type. Importantly, the police cannot determine whether a given signal realization reflects the true type or is noise. As a result, the posterior belief places positive probabilities on both types below and above c. Specifically, the posterior induced by signal $s \in [0,c]$ contains a point mass $F(c)$ on type s and a mass $1 - F(c)$ of types distributed according to $F(\cdot|\tilde{x} > c)$.[15] The posterior beliefs (indexed by their point mass) are distributed according to $F(\cdot|\tilde{x} \leq c)$ and average to prior distribution F.

[15]To see this, let $f(s|\tilde{x} \leq c)$ denote the conditional density associated with $F(s|\tilde{x} \leq c)$ evaluated at $x = s$. Roughly, the "probability" with which signal s is realized is $f(s|\tilde{x} \leq c)$. The "probability" of the joint event—in which the realized signal and the true type are both equal to s—is equal to the probability of $x \leq c$, which is $F(c)$, multiplied by the "probability" of the true type s conditional on $x \leq c$, which is $f(s|\tilde{x} \leq c)$. Therefore, the posterior probability of type s conditional on signal s is $\frac{F(c)f(s|\tilde{x} \leq c)}{f(s|\tilde{x} \leq c)} = F(c)$.

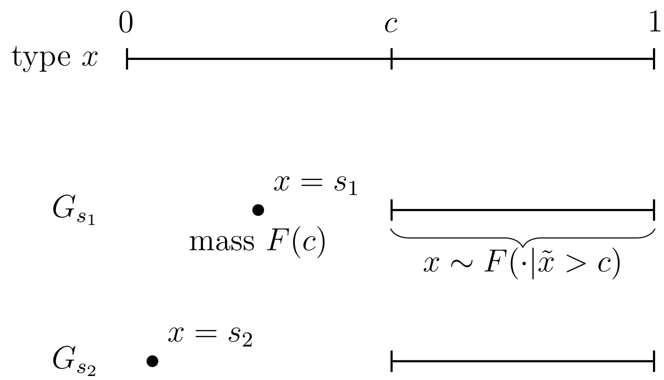

Figure 1: Posterior distribution G_s of types conditional on signal $s = s_1, s_2$ under (S_c, π_c).

Theorem 1. *Let $\hat{c} \in (0, 1)$ denote the cutoff defined by equation (6), i.e., $\overline{P} = \mathbb{E}_F[\tilde{x}|\tilde{x} \leq \hat{c}]$. 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 below $\hat{c}$ do not commit a crime. 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 $\hat{c}$. This 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 any search strategy that exhausts search capacity $\overline{P}$ is optimal for the police. Search strategy p^* indeed exhausts the search capacity because of equation (6). 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 willing to abstain from committing a crime. Types above $\hat{c}$ will be searched with probability at most $\hat{c}$, so they commit a crime. Hence the strategy profile described above is an equilibrium and attains the same crime rate as the solution to the relaxed problem described in Lemma 1. □

The optimality of a truth-or-noise signal structure stems from two forces. On the one hand, providing the police with detailed information may erode the deterrent effect of searches: The police use the information to target those who are not deterred from crime. On the other hand, providing the police with too little information leads to wasteful searches, because the police may search some agents with unnecessarily high probability when those agents have a low incentive to commit a crime. A truth-or-noise signal structure provides the police with partial information, which allows the police to reduce wasteful searches without eroding the deterrent effect.

Specifically, the crime-minimizing signal structure has two key features. First, it prevents the police from identifying or focusing search resources on only high-type agents, because every signal contains the equal fraction $F(\hat{c})$ of types below $\hat{c}$.[16] This prevents the police from focusing too much search capacity on the highest-type individuals, who are less likely to be deterred by the threat of being caught. Second, the crime-minimizing signal structure enables the police to search each type $x < \hat{c}$ with the minimum necessary probability needed to deter crime, because no two types below $\hat{c}$ are pooled together. This reduces wasteful searches and increases the number of agents who can be deterred from crime. These two features work together to minimize the overall crime rate.

3.3 Crime-reducing Signal Structure with Limited Information

The crime-minimizing signal structure prevents the police from targeting high-type agents while allowing them to search low-type agents according to their type. To achieve both goals, the signal structure pools high types with a type below $\hat{c}$ while never pooling multiple types below $\hat{c}$ into a single signal. However, implementing such a signal structure may require detailed information about types below $\hat{c}$, which may not be feasible in practice. Nevertheless, the idea behind the crime-minimizing signal structure can be useful even when we cannot use signal structures that allow for full separation of the low types.

For example, consider a monotone-partitional signal structure that partitions $[0, 1]$ into subintervals $0 = x_0 < x_1 < \cdots < x_{N-1} < x_N = 1$. This could represent a city divided into N areas, each assigned a "risk level" for crime. Instead of sharing this information directly with the police, it can be garbled similarly to the crime-minimizing signal structure: We create $K < N$ segments, where each segment $k = 1, \dots, K$ contains all types in (x_{k-1}, x_k) and the fraction $\frac{F(x_k)-F(x_{k-1})}{F(x_K)}$ of types in $(x_K, 1]$. Every segment contains the same fraction $1 - F(x_K)$ of types above x_K, and no two subintervals below x_K are pooled together.

[16] A consequence of this property is that the crime rate is equal to $1 - F(\hat{c})$ across all signals, and thus the police are indifferent between any search strategies that exhaust $\overline{P}$. However, the police's indifference is not unique to this signal structure and holds for, e.g., any deterministic signal structure (see Online Appendix A).

This garbling serves two purposes: First, the police cannot search types above x_K exclusively, because they are pooled with types below x_K. Second, the police can still search types below x_K with different probabilities, depending on which subinterval these types belong to. As shown below, this garbling—which is a discrete version of the truth-or-noise signal structure—may achieve a lower crime rate than the original partitional signal structure.

Example 1. Suppose $F = U[0,1]$ and $\overline{P} = 1/2$.[17] Consider a signal structure (S, π) that sends signal s_L if $x \in [0, 1/3]$, s_M if $x \in (1/3, 2/3]$, and s_H if $x \in (2/3, 1]$. In equilibrium under (S, π), the police adopt search strategy $p(s_L) = 1/6$, $p(s_M) = 1/2$, and $p(s_H) = 5/6$.[18] Within each signal, half of the agents abstain from committing a crime, so the equilibrium crime rate is $1/2$, the same as under no information.

Now, garble (S, π) by pooling all types in $[0, 1/3]$ with half of those in $(2/3, 1]$ to form signal s_L^*, and all types in $(1/3, 2/3]$ with the remaining half of $(2/3, 1]$ to form signal s_M^*. Let (S^*, π^*) denote the resulting signal structure. Under (S^*, π^*) and $F = U[0,1]$, a fraction $2/3$ of agents within each signal have types below $2/3$. In equilibrium, the police search s_L^* and s_M^* with probabilities $1/3$ and $2/3$, respectively, which deters all types below $2/3$ from committing a crime. The crime rate is thus $1/3 < 1/2$.

Intuitively, the signal structure (S^*, π^*) prevents the police from identifying types above $2/3$. This reduces searches directed at high types and increases those directed at low types, which are more responsive to searches. At the same time, it allows the police to search types in $[0, 1/3]$ and $(1/3, 2/3]$ with distinct probabilities, leading to more effective allocation of search resources across low types than under no information.

Remark 2 (Multiplicity of Crime-minimizing Signals). The crime-minimizing signal structure is not unique in either the relaxed or the original problem. The relaxed problem has multiple solutions, such as those in Lemma 1 and Theorem 1. The original problem also has multiple solutions: For example, let $F = U[0,1]$ and $\overline{P} = \frac{1}{4}$. Solving equation (6)

[17] This violates Assumption 1 as $\int_0^1 x \,\mathrm{d}F(x) = \overline{P}$, but the same observation holds for any $\overline{P} < 1/2$.

[18] This search strategy deters all types in $[0, 1/6] \cup [1/3, 1/2] \cup [2/3, 5/6]$ from committing a crime, equalizing the crime rate of every signal to $1/2$. It also exhausts the search capacity, as $\frac{1}{3}\left(\frac{1}{6} + \frac{1}{2} + \frac{5}{6}\right) = \frac{1}{2} = \overline{P}$. Thus, the search strategy is optimal.

gives $\hat{c} = \frac{1}{2}$. Another crime-minimizing signal structure, different from the one in Theorem 1, is as follows: Signal $s \in [0, \frac{1}{2}]$ is realized with probability 1 if and only if $x = s$ or $x = 1 - s$. Conditional on signal s, the posterior belief assigns equal probability to $x = s$ and $x = 1 - s \geq \hat{c}$. Since $F(\hat{c}) = \frac{1}{2}$, the proof of Theorem 1 applies directly. Thus, this signal structure achieves the same crime rate as that in Theorem 1.

Remark 3 (Fragility of the Crime-minimizing Signal Structure). The signal structure $(S_{\hat{c}}, \pi_{\hat{c}})$ in Theorem 1 exhibits two kinds of fragility. First, it achieves the minimal crime rate only if the agents who are indifferent break ties by not committing a crime. Indeed, under $(S_{\hat{c}}, \pi_{\hat{c}})$, there is another equilibrium with crime rate 1. In this equilibrium, the police use the same search strategy p^* as in the crime-minimizing equilibrium, and all types below $\hat{c}$ break ties and commit a crime. Second, the signal structure fails to deter crime at all if the police's search capacity is slightly misspecified: If the crime-minimizing signal structure $(S_{\overline{P}}, \pi_{\overline{P}})$ is designed and employed for a certain value of $\overline{P}$ but the actual search capacity is $\overline{P} - \epsilon$ with $\epsilon > 0$, then, according to logic similar to that in Theorem 0, all agents commit a crime in a unique equilibrium. At the same time, these kinds of fragility are not necessary for a signal structure to reduce crime (relative to no information). For example, the binary signal structure (S^*, π^*) in Example 1 has a unique equilibrium—which is the one described in the example—and the equilibrium is robust to the slight misspecification of search capacity $\overline{P}$. Appendix B formalizes these ideas.

4 Discussion

I examine the consequences of modifying various assumptions in the baseline model.

4.1 Balancing the Crime Rate and Costs Imposed on Innocents

The cost of policing for innocent people has been an important consideration in the discussion of policing (see, e.g., Section VII of Persico (2002)). In this section, I characterize a signal structure that balances such costs and the crime rate. The timing and payoffs are the same as in the baseline model (Section 2).

Given any signal structure (S, π) and any strategy profile, the *search mass on innocents* is defined as the mass of agents who did not commit a crime but are searched, i.e.,

$$\int_0^1 \int_S p(s)\,\mathrm{d}\pi(s|x)[1 - a(x)]\,\mathrm{d}F(x), \tag{9}$$

where $a(x)$ is the probability with which type x commits a crime. I take (9) as the cost of policing for innocent agents. The minimal search mass on innocents consistent with crime rate r is

$$\int_0^{F^{-1}(1-r)} x\,\mathrm{d}F(x), \tag{10}$$

which arises if all types below $F^{-1}(1-r)$ are searched with a probability equal to their type and abstain from a crime. When $r = 1 - F(\hat{c})$, where $\hat{c}$ is the cutoff type in Theorem 1, the minimal search mass on innocents is achieved by the crime-minimizing signal structure.

The equilibrium crime rate is at least $1 - F(\hat{c})$ across all signal structures. Thus, I fix an arbitrary crime rate $r \in [1 - F(\hat{c}), 1]$ and characterize a signal structure and an equilibrium that minimize the search mass on innocents subject to achieving the crime rate r. If we explicitly formulated an information designer who cares about the crime rate and costs imposed on innocents, the optimal choice of the designer would belong to the class of signal structures characterized below.

Proposition 1. *Let $\hat{r} = 1 - F(\hat{c})$ be the minimal crime rate of the baseline model. For any $r \in [\hat{r}, 1]$, there exists a signal structure and an equilibrium that attain the minimal search mass on innocents given by (10) and crime rate r.*

In Appendix C, I construct a signal structure under which any type $x > F^{-1}(1-r)$ commits a crime and any type $x < F^{-1}(1-r)$ faces search probability x and chooses not to commit a crime. This outcome attains a crime rate of r and the minimal search mass on innocents (10). The signal structure resembles a truth-or-noise signal structure with one modification: For any type $x > F^{-1}(1-r)$, the signal reveals the type with some type-dependent probability. Whenever that happens, the police, knowing that the agent will commit a crime, set search probability equal to 1. The probability of revelation is such that in expectation, any type above $F^{-1}(1-r)$ still commits a crime. As the probability of revelation increases, the police more easily find criminals. Consequently, the police allocate more search resources to criminals, which results in a lower search mass on innocents and a higher crime rate.

4.2 Police with Full Commitment Power

In this section, I assume that the police have full information and can commit to any search strategy upfront. Formally, the police first choose a search strategy, and after observing it, each agent chooses whether to commit a crime. The police's payoff is given by

$$\hat{\lambda}\sigma - (1-\hat{\lambda})r, \tag{11}$$

where σ is the mass of successful searches (2), r is the crime rate (3), and $\hat{\lambda} \in [0,1]$. I focus on an equilibrium that maximizes the police's payoff.

Given that the police have full information, the minimum search capacity needed for attaining crime rate r is

$$\int_0^{F^{-1}(1-r)} t \,\mathrm{d}F(t), \tag{12}$$

which arises when the police search every type $t \le F^{-1}(1-r)$ with probability t and deter them from a crime. The mass of successful searches consistent with the crime rate r is then at most $\overline{P} - \int_0^{F^{-1}(1-r)} t \,\mathrm{d}F(t)$. The police can indeed attain this upper bound and the crime rate r simultaneously; to do so, the police search every type $t \le F^{-1}(1-r)$ with probability t while searching every type $t > F^{-1}(1-r)$ with a probability of at most t, so that (i) the police exhaust their search capacity $\overline{P}$, (ii) all types above $F^{-1}(1-r)$ commit a crime, and (iii) no types below $F^{-1}(1-r)$ commit a crime.

The police's problem thus reduces to finding crime rate r that maximizes their payoff,

$$\hat{\lambda}\left(\overline{P} - \int_0^{F^{-1}(1-r)} t\mathrm{d}F(t)\right) - (1-\hat{\lambda})r,$$

subject to $r \ge r^*$. Here, $r^* \in (0,1)$ is the minimum crime rate under commitment, which solves

$$\int_0^{F^{-1}(1-r^*)} t \, \mathrm{d}F(t) = \overline{P}. \tag{13}$$

Equivalently, the police's problem can be stated as the choice of a cutoff type $x = F^{-1}(1-r)$, i.e.,

$$\max_{x\in[0,\overline{x}]} \hat{\lambda}\left(\overline{P} - \int_0^x t\mathrm{d}F(t)\right) - (1-\hat{\lambda})[1-F(x)],$$

where $\overline{x} \triangleq F^{-1}(1-r^*)$. The derivative of the police's objective with respect to x is

$$-\hat{\lambda}x + 1 - \hat{\lambda},$$

which crosses 0 at most once and always from above as x increases. Solving the first-order condition and taking into account the constraint $x \leq \overline{x}$, we can derive the optimal cutoff:

$$x^* = \min\left(\frac{1-\hat{\lambda}}{\hat{\lambda}}, \overline{x}\right).$$

Proposition 2. *Suppose that the police can commit to their search strategy upfront, have full information, and face payoffs (11). In the police-preferred equilibrium, the crime rate is $1 - F\left(\min\left(\frac{1-\hat{\lambda}}{\hat{\lambda}}, \overline{x}\right)\right)$, which increases from $1 - F(\overline{x})$ to 1 as the weight $\hat{\lambda}$ on successful searches increases from 0 to 1.*

If $\hat{\lambda} = 1$, the police maximize the mass of successful searches. In this case, Proposition 2 implies that the police choose a search strategy under which all agents commit a crime, leading to the maximal crime rate and successful searches. This resembles Theorem 0 but is different: Theorem 0 holds for any $\hat{\lambda} \in (0, 1]$, whereas in the current case, a crime rate of 1 arises only when $\hat{\lambda} = 1$.

If $\hat{\lambda} = 0$, the police aim to minimize the crime rate. In this case, the police search each type $x < \overline{x}$ with probability x. In equilibrium, all such agents are indifferent and thus abstain from committing a crime.

The police with commitment power can attain a lower crime rate than the police with no commitment power. Indeed, the cutoff type for the crime-minimizing outcome in Theorem 1 solves $\int_0^{\hat{c}} \frac{t}{F(\hat{c})} \, \mathrm{d}F(t) = \overline{P}$ (which is (6)), whereas the cutoff type in this commitment model with $\hat{\lambda} = 0$ solves $\int_0^{\overline{x}} t \, \mathrm{d}F(t) = \overline{P}$. Since $F(\hat{c}) < 1$, we must have $\hat{c} < \overline{x}$. Therefore, compared with the baseline model, the police with commitment power can also deter the types in $[\hat{c}, \overline{x}]$ from committing a crime.

Finally, the crime-minimizing outcome under full commitment (i.e., crime rate $1 - F(\overline{x})$) could arise even when the police and agents move simultaneously, provided that the police face different payoffs from those in the baseline model. For example, suppose the police's payoff equals the mass of agents who did *not* commit a crime but are searched, i.e., (9). This captures a case in which the police prefer to search innocents: They want to confirm that citizens followed the law. In equilibrium, the police search each type $x \leq \overline{x}$ with probability x and never search other types. Correspondingly, agents commit a crime if and only if $x > \overline{x}$. The police have no incentive to deviate, because they prefer to search agents who did not commit a crime rather than those who did.[19]

4.3 Designer with Partial Commitment Power

A version of Theorem 1 holds even when a search strategy is partially determined by an entity, different from the police, that has some commitment power. I clarify this point using the following variation of the model.

First, I introduce a player called the *designer*, whose goal is to minimize the crime rate. The designer could represent a law enforcement organization, with the police as individual enforcers. The designer faces an exogenous signal structure, $(\overline{S}, \overline{\pi})$. For simplicity, assume $\overline{S}$ is finite. Each $t \in \overline{S}$ is called a *pre-signal*, where $m(t) \triangleq \int_0^1 \overline{\pi}(t|x) \, \mathrm{d}F(x)$ denotes the mass of agents who receive t, and F_t denotes the posterior type distribution conditional on t. Assume that each F_t has a positive density on $[0, 1]$. The designer has search capacity $\overline{P}$ that satisfies Assumption 1.

The designer commits to the allocation of search capacity and the information available to the police. First, the designer chooses an *allocation policy* $\overline{a} : \overline{S} \to \mathbb{R}_+$, subject to the

[19] This is not the only specification that attains the crime rate $1 - F(\overline{x})$ under the simultaneous-move assumption. For example, suppose the police's objective is to minimize the crime rate (3). In this case, the police's payoffs do not directly depend on their search strategy. Thus, any search strategy constitutes an equilibrium, including the crime-minimizing strategy under full commitment.

capacity constraint:

$$\sum_{t \in \overline{S}} \overline{a}(t) = \overline{P}, \tag{14}$$

where $\overline{a}(t)$ is the mass of searches allocated to pre-signal t.

Second, for each $t \in \overline{S}$, the designer chooses a signal structure (S_t, π_t), which captures the information that the police can use to search agents with each pre-signal t. As in the baseline model, the police cannot commit to how to use this information.

After the designer commits to an allocation policy $\overline{a}$ and signal structures $\{(S_t, \pi_t)\}_{t \in \overline{S}}$, the game unfolds as follows: The pre-signal of each agent is realized and observed by the agent and the police. In addition, each agent privately observes their type. Then, for each $t \in \overline{S}$, the agents with pre-signal t and the police play the game described in Section 2, where the mass of agents is $m(t)$ (instead of 1); the type distribution is F_t; and the search capacity is $\overline{a}(t)$.[20]

The police's payoff equals the total mass of successful searches aggregated across all pre-signals. The designer's objective is to minimize the total mass of agents who commit a crime.[21] I focus on the designer's preferred equilibrium, which minimizes the crime rate across all equilibria. If $|\overline{S}| = 1$, the model reduces to the baseline model.

As an example, imagine that a police organization (i.e., the designer) has some information about how different times of the day or locations in a city are associated with crime opportunities. This information is captured by pre-signals, $(\overline{S}, \overline{\pi})$, where each pre-signal represents a particular time or location.[22] The organization commits to the allocation of officers across different times or locations, which corresponds to an allocation policy $\overline{a}$. The

[20]As in the baseline model, the police have no commitment power. However, if the underlying signal $(\overline{S}, \overline{\pi})$ available to the designer is fully informative, the designer can specify the search capacity allocated to each type and thus directly control the police's search strategy. In this case, the model is equivalent to the one in Section 4.2 with $\hat{\lambda} = 0$.

[21]Formally, let $r(t)$ denote the fraction of agents who commit a crime conditional on pre-signal t, and $d(t)$ denote the fraction of agents who commit a crime and are searched conditional on pre-signal t. The police's payoff is $\sum_{t \in \overline{S}} d(t)m(t)$, and the designer's payoff is $-\sum_{t \in \overline{S}} r(t)m(t)$.

[22]For example, people may face higher returns from illegal parking during certain times of the day. Another example is tax auditing for businesses: Some types of businesses, such as large corporations, may face higher returns from tax evasion. In such scenarios, the assumption that agents' types are exogenous implies that agents do not change the timing of their commute or their business type based on opportunities for illegal parking or tax evasion.

organization can also provide the officers assigned to each pre-signal t with additional information (S_t, π_t), generated by a predictive algorithm.[23] However, neither the organization nor officers can pre-commit to how the information will be used.

The designer can replicate the same outcome as when the police have no information in the baseline model. To do so, the designer allocates search capacity $\overline{a}(t) = m(t)\overline{P}$ and chooses an uninformative (S_t, π_t), such as $S_t = \{\phi\}$, for each t. The police can then only search each pre-signal t with probability $\overline{P} = \frac{m(t)\overline{P}}{m(t)}$, because they use search capacity $m(t)\overline{P}$ to randomly search a mass $m(t)$ of agents. Alternatively, the designer might provide the police with more information, but the resulting outcome could involve more crime. How should the designer balance control over the allocation of search capacity with the use of more information?

Theorem 1 helps us solve the designer's problem. To see this, suppose that the designer has chosen an allocation policy $\overline{a}$. The game between the police and the agents with pre-signal t is equivalent to our baseline model in which the prior type distribution is $F = F_t$ and the search capacity is $\overline{P} = \frac{\overline{a}(t)}{m(t)}$. Here, I scale the mass of each pre-signal to 1 (instead of $m(t)$) and the search capacity to $\frac{\overline{a}(t)}{m(t)}$, so that Theorem 1 applies verbatim.

Theorem 1 implies that for each pre-signal t, the designer should provide the police with the truth-or-noise signal structure with cutoff $c_t = c_t\left(\frac{\overline{a}(t)}{m(t)}\right)$, which solves

$$\mathbb{E}_{F_t}\left[\tilde{x} | \tilde{x} \leq c_t\right] = \frac{\overline{a}(t)}{m(t)}. \tag{15}$$

The cutoff c_t does not exist if $\frac{\overline{a}(t)}{m(t)} > \mathbb{E}_{F_t}[\tilde{x}]$, in which case I set $c_t = 1$. The minimized crime rate conditional on pre-signal t is thus $1 - F_t\left(c_t\left(\frac{\overline{a}(t)}{m(t)}\right)\right)$. The designer's problem reduces to choosing an allocation policy to minimize the overall crime rate:

$$\min_{\overline{a}} \sum_{t \in \overline{S}} m(t)\left[1 - F_t\left(c_t\left(\frac{\overline{a}(t)}{m(t)}\right)\right)\right] \tag{16}$$

subject to the capacity constraint, (14). The following result formalizes the above observation.

[23] As shown in Appendix D, the police can be viewed as a population of individual officers.

Proposition 3. *The designer chooses an allocation policy $\overline{a}$ that solves* (16) *and sets each* (S_t, π_t) *as the truth-or-noise signal structure with cutoff* $c_t\left(\frac{\overline{a}(t)}{m(t)}\right)$*. Under the designer's optimal allocation policy,* $c_t\left(\frac{\overline{a}(t)}{m(t)}\right)$ *solves* (15) *for each* $t \in \overline{S}$.

Appendix D proves this result and solves the designer's problem in a simple example. Furthermore, Online Appendix B extends Theorem 0 to the current setup by assuming that the designer can control the allocation of search capacity based on pre-signals but must disclose full information to the police.

4.4 Endogenous Search Capacity

In this section, I extend the baseline model to allow the police to strategically choose total search capacity. Here, the police can choose any search strategy p at cost $C(P)$, where P represents the mass of searches induced by p:

$$P \triangleq \int_0^1 \int_S p(s)\, \mathrm{d}\pi(s|x)\, \mathrm{d}F(x).$$

The police's objective is to maximize the mass of successful searches (2) minus the cost $C(P)$. The cost function $C(\cdot)$ is strictly increasing, strictly convex, differentiable, and satisfies

$$C'(0) < 1 < C'\left(\int_0^1 x\, \mathrm{d}F(x)\right).$$

The first inequality ensures that the police search a positive mass of agents under any signal structure and any equilibrium. The second inequality ensures that the equilibrium crime rate is positive, which plays the same role as Assumption 1 played in earlier sections. The rest of the model, including the agents' payoffs and the timing, remains the same.

First, Theorem 0 still holds in this extended setup: Full information leads to the highest possible crime rate. Thus, for the remainder of this section, I focus on a crime-minimizing signal structure with endogenous search capacity (technical details are provided in Appendix E).

To analyze this, I begin with the relaxed problem, in which agents do not directly observe their own type and the joint signal structure, (S_P, S_A, π), is chosen to minimize the equilibrium crime rate. Figure 2 illustrates a solution to the relaxed problem. Similar to the case with exogenous search capacity (Lemma 1), agents receive a signal (either *crime* or *not*) based on whether their type exceeds a cutoff, c^*. In equilibrium, they follow an action recommendation. The police's signal is a garbling of the agents' signals and may differ from the uninformative signal. Notably, the police may now privately identify and, therefore, choose not to search a fraction α^* of agents who receive signal *not*. For the remaining population, the police apply the uniform search rate, $\rho^* > 0$.

Intuition and Key Results The partial revelation of agents' types (i.e., $\alpha^* > 0$) incentivizes the police to choose a higher search capacity, since they can more effectively target criminals and increase the mass of successful searches. However, the partial revelation also erodes the deterrence effect of searches, as demonstrated in the baseline model. When the targeting effect dominates the deterrence loss, we obtain $\alpha^* > 0$.

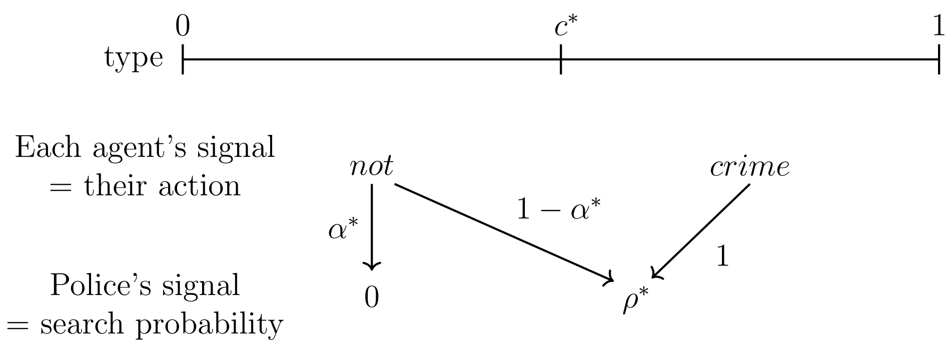

Figure 2: A solution to the relaxed problem with endogenous search capacity. The police identify a fraction α^* of the agents who do not commit a crime.

When $\alpha^* = 0$, the crime-minimizing joint signal structure takes qualitatively the same form as Lemma 1: The police remain uninformed while agents receive a cutoff-based signal. In this case, truth-or-noise signal structures continue to solve the original problem. The following result provides a condition for $\alpha^* = 0$ and relies on the following restriction on the cost function.

Assumption 2. The police's search cost function takes the form of

$$C(P) = \frac{L}{1+\beta} P^{1+\beta}$$

for some $\beta > 0$ and $L > 1$.

Proposition 4. *Consider the police with endogenous search capacity. Suppose that the cost function $C(\cdot)$ satisfies Assumption 2, where (F, β, L) satisfies*

$$\mathbb{E}_F\left[\tilde{x} \,\middle|\, \tilde{x} < F^{-1}\left(\frac{\beta}{1+\beta}\right)\right] > ((1+\beta)L)^{-\frac{1}{\beta}}. \tag{17}$$

Then, $\alpha^* = 0$ holds in the relaxed problem, and the crime-minimizing signal structure is a truth-or-noise signal structure.

For a fixed β, inequality (17) is more likely to hold when (i) the police face higher search costs (i.e., L is high) or (ii) agents face greater returns from a crime (i.e., $F^{-1}\left(\frac{\beta}{1+\beta}\right)$ is high). Under these conditions, the marginal cost of search in equilibrium is high. In such a case, there is little advantage to revealing some agents with signal *not*, as doing so will have little effect on the police's search effort and only erodes the deterrent effect of searches. In this case, the crime-minimizing signal structure in the relaxed problem keeps the police uninformed and, correspondingly, the solution to the original problem becomes a truth-or-noise signal structure.

A full, general analysis of the original problem under the extension is beyond the paper's scope. However, Proposition 8 in Appendix E provides an example in which a generalization of a truth-or-noise signal structure solves the original problem when we have $\alpha^* > 0$ in the relaxed problem.

5 Conclusion

This paper provides insights into the consequences of crime prediction technologies and algorithmic policing. As law enforcement agencies increasingly rely on data-driven decision-making, this paper shows that better information could undermine deterrence and lead to higher overall crime rates. This suggests that simply increasing the predictive power of policing tools, whether through AI models, data aggregation, or machine learning, may not necessarily improve public safety. At the same time, characterization of the crime-minimizing signal structure suggests that the optimal policy may involve regulating how information is structured and used, rather than banning these technologies outright.

The methodology developed in this paper has broader applications beyond crime and policing. Information design is relevant to fields such as tax enforcement, counterterrorism, and regulatory compliance, in which authorities must allocate enforcement efforts based on limited information. The approach developed here, using relaxed-problem formulations to determine optimal information structures, can be applied to those domains to design policies that balance enforcement and deterrence.

Finally, future research may explore crime-minimizing information structures under institutional and legal constraints. In practice, law enforcement agencies may face political and legal limitations on the type of data they can collect or use. Understanding how to design crime-minimizing policies under such constraints would further enhance the applicability of the results.

References

Baicker, Katherine and Mireille Jacobson (2007), "Finders keepers: Forfeiture laws, policing incentives, and local budgets." *Journal of Public Economics*, 91, 2113–2136.

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.

Berk, Richard A (2021), "Artificial intelligence, predictive policing, and risk assessment for law enforcement." *Annual Review of Criminology*, 4, 209–237.

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.

Boleslavsky, Raphael and Kyungmin Kim (2018), "Bayesian persuasion and moral hazard." *Available at SSRN 2913669.*

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 (2024), "Predictive enforcement." *arXiv preprint arXiv:2405.04764.*

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.

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.

Gentzkow, Matthew and Emir Kamenica (2016), "A Rothschild-Stiglitz approach to Bayesian persuasion." *American Economic Review*, 106, 597–601.

Goldman, A. J. and M. H. Pearl (1976), "The dependence of inspection-system performance on levels of penalties and inspection resources." *Journal of Research of the National Bureau of Standards, B. Mathematical Sciences*, 80B, 189–236.

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.

Hörner, Johannes and Nicolas S Lambert (2021), "Motivational ratings." *Review of Economic Studies*, 88, 1892–1935.

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.

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." *Quarterly Journal of Economics*, 121, 1029–1061.

Li, Cheng and Christopher Cotton (2023), "Profiling restrictions in a model of law enforcement and strategic crime." *European Journal of Law and Economics*, 55, 511–532.

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.

Makowsky, Michael D and Thomas Stratmann (2011), "More tickets, fewer accidents: How cash-strapped towns make for safer roads." *Journal of Law and Economics*, 54, 863–888.

Nagin, Daniel S, Robert M Solow, and Cynthia Lum (2015), "Deterrence, criminal opportunities, and police." *Criminology*, 53, 74–100.

Owens, Emily and Bocar Ba (2021), "The economics of policing and public safety." *Journal of Economic Perspectives*, 35, 3–28.

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." *Annual Review of Economics*, 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.

Rodina, David (2017), "Information design and career concerns." *mimeo*.

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

Stashko, Allison (2023), "Do police maximize arrests or minimize crime? Evidence from racial profiling in US cities." *Journal of the European Economic Association*, 21, 167–214.

Sun, Yeneng (2006), "The exact law of large numbers via Fubini extension and characterization of insurable risks." *Journal of Economic Theory*, 126, 31–69.

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

Zapechelnyuk, Andriy (2020), "Optimal quality certification." *American Economic Review: Insights*, 2, 161–176.

Appendix

A Type-dependent Payoffs of the Police

I modify the baseline model so that the police's payoff of searching an agent who commits a crime is $1 + \eta g(x)$, where x is the agent's type, $\eta \geq 0$ is an exogenous parameter that captures the degree of type-dependency, and $g : [0,1] \to \mathbb{R}_+$ is an arbitrary nonnegative continuous function such that $g(x) = 0$ for some $x \in [0,1]$. Thus, if the police face signal structure (S, π) and adopt search strategy p and type x commits a crime with probability $a(x)$, the police's payoff—the mass of successful searches that reflects the type-dependent priority to search—is

$$\int_0^1 \int_S p(s)\,\mathrm{d}\pi(s|x)a(x)[1 + \eta g(x)]\,\mathrm{d}F(x).$$

To extend Theorem 0, define $\gamma > 0$ as follows:

$$\int_{\{x:g(x)\geq\gamma\}} x\,\mathrm{d}F(x) = \overline{P}. \tag{A.1}$$

For simplicity, we assume that γ exists.[24] Then, we obtain the following result.

Proposition 5. *Suppose the police have type-dependent payoffs and full information. There exists an equilibrium with the following property: The police search the agents whose type satisfies $g(x) \geq \gamma$ with probability x. These agents are indifferent and commit a crime with probability*

$$a^*(x) = \frac{1 + \eta\gamma}{1 + \eta g(x)}. \tag{A.2}$$

The police search any type x with $g(x) < \gamma$ with probability 0, and these agents commit a crime with probability 1. As $\eta \to 0$, the equilibrium crime rate converges to 1.

Proof. Consider the strategy profile described in the statement. The agents' strategies are optimal. The police's payoff of searching any type x with $g(x) \geq \gamma$ is $a^*(x)(1 + \eta g(x)) = 1 + \eta\gamma$, whereas the payoff of searching any type $g(x) < \gamma$ is at most $1 + \eta\gamma$. As a result, the police's search strategy is optimal if it only searches types x with $g(x) \geq \gamma$ and exhausts search capacity $\overline{P}$. The candidate strategy satisfies the first property by construction and the second property because of (A.1). Finally, in this equilibrium, the probability of any agent committing a crime is at least $\frac{1+\eta\gamma}{1+\eta \max_{x\in[0,1]} g(x)}$. Because γ is independent of η, as $\eta \to 0$, the crime rate converges to 1. □

[24]Under Assumption 1, γ exists if for any $y \in g([0,1])$, its inverse image $g^{-1}(\{y\}) \subseteq [0,1]$ has measure zero.

B Omitted Formalism for Remark 3

As I discussed in Remark 3, the crime-minimizing signal structure in Theorem 1 is fragile in two ways. In this appendix, I elaborate on those points and then show that, as claimed in the remark, signal structure (S^*, π^*) in Example 1 does not have the same kinds of fragility.

First, consider the crime-minimizing signal structure of Theorem 1. For any crime rate $r \in [1 - F(\hat{c}), 1]$ greater than the minimized level, there exists an equilibrium in which the crime rate is r and the police continue to adopt the same search strategy as in the crime-minimizing equilibrium. In this equilibrium, any type above $\hat{c}$ commits a crime with probability 1, and any type below $\hat{c}$, who is indifferent, commits a crime with probability $1 - \frac{1-r}{F(\hat{c})}$. The mass of agents who abstain from crime is $F(\hat{c}) \cdot \left[1 - \left(1 - \frac{1-r}{F(\hat{c})}\right)\right] = 1 - r$, so the crime rate is r.

Second, the crime-minimizing signal structure fails to deter a crime at all if the police's search capacity is misspecified. To see this, suppose that we construct the crime-minimizing signal structure, denoted by $(S_{\overline{P}}, \pi_{\overline{P}})$, for a certain value of $\overline{P}$, but the actual search capacity turns out to be $\overline{P} - \epsilon$ with $\epsilon > 0$. The signal structure $(S_{\overline{P}}, \pi_{\overline{P}})$ generates signals between 0 and $\hat{c}$, and the search capacity $\overline{P}$ is just enough for the police to make all types below $\hat{c}$ indifferent. Thus, if the search capacity is $\overline{P} - \epsilon$, for any search strategy, there is always a positive measure of signals associated with crime rate 1. The same argument as in the proof of Theorem 0 implies that every signal generated by the signal structure $(S_{\overline{P}}, \pi_{\overline{P}})$ must have crime rate 1. Therefore, if the search capacity is slightly overestimated, the equilibrium crime rate jumps from the minimized level to 1.

The binary signal structure (S^*, π^*) in Example 1 reduces the crime rate relative to no information, but it does not exhibit the same kind of fragility. First, I show that the equilibrium described in Example 1, where $p(s_L^*) = 1/3$ and $p(s_M^*) = 2/3$, is unique. For instance, if $p(s_L^*) > 1/3$ in equilibrium, the binding search capacity constraint implies $p(s_M^*) < 2/3$. Then, the crime rate for signal L becomes strictly lower than that for signal M. The police would then profitably deviate by shifting search probabilities from signal L to signal M, which leads to a contradiction. Symmetrically, $p(s_L^*) < 1/3$ leads to a contradiction. Therefore, a search strategy can be part of an equilibrium only if $p(s_L^*) = 1/3$ and $p(s_M^*) = 2/3$, which implies that the equilibrium described in Example 1 is unique.

Moreover, the equilibrium crime rate is continuous with respect to $\overline{P}$ at $\overline{P} = 1/2$. To see this, suppose we construct the signal structure (S^*, π^*) for $\overline{P} = 1/2$, but the actual search capacity is $\overline{P}_\epsilon = \overline{P} - \epsilon$ for some ϵ. I allow $\epsilon < 0$, in which case the actual search capacity is underestimated. Let $\hat{p}(s_L^*)$ and $\hat{p}(s_M^*)$ denote the equilibrium search probabilities for signals L and M, respectively. Here, the binding search capacity constraint $0.5\hat{p}(s_L^*) + 0.5\hat{p}(s_M^*) = \overline{P}_\epsilon$ implies $\hat{p}(s_L^*) = 2(\overline{P}_\epsilon - 0.5\hat{p}(s_M^*))$. The value $\hat{p}(s_M^*)$ is uniquely determined by the condition that the crime rates for signals L and M are equalized, i.e.,

$$F_L\left(2(\overline{P}_\epsilon - 0.5\hat{p}(s_M^*))\right) = F_M(\hat{p}(s_M^*)), \tag{B.3}$$

where F_L and F_M are the type distributions conditional on signals L and M, respectively. At $\hat{p}(s_M^*) = 0$, the LHS is strictly greater than the RHS. If $\hat{p}(s_M^*)$ is large enough, the RHS is strictly greater than the LHS. Also, F_L and F_H are continuous. By the intermediate value theorem, there exists a unique value of $\hat{p}(s_M^*)$ that solves equation (B.3). The resulting $(\hat{p}(s_L^*), \hat{p}(s_M^*))$ and the agents' (unique) best responses constitute a unique equilibrium. Note also that the solution $\hat{p}(s_M^*)$ to equation (B.3) is continuous with respect to ϵ at $\epsilon = 0$. Therefore, even if the search capacity is slightly misspecified, the resulting game has a unique equilibrium, and the crime rate remains close to $1/3$, which is the level associated with (S^*, π^*) when $\overline{P} = 1/2$.

C Proof of Proposition 1

I prove Proposition 1. Note that if $r = 1$, any equilibrium under full information has crime rate 1 and the search mass on innocents is equal to zero (see Theorem 0), which satisfies the condition in the proposition. For $r \in (\hat{r}, 1)$, the following result presents a signal structure and an equilibrium that have the properties described in Proposition 1.

Proposition 6. *Let $\hat{r}$ be the crime rate under the crime-minimizing outcome in Theorem 1. Take any $r \in (\hat{r}, 1)$. Define $x(r) \triangleq F^{-1}(1-r)$. There exists some $\hat{x} \in (x(r), 1)$ such that the following signal structure and equilibrium attain crime rate r and the minimal search* *mass on innocents in* (10) *i.e.,* $\int_0^{F^{-1}(1-r)} x \, \mathrm{d}F(x)$*. If* $x \leq x(r)$*, the signal realization equals* x *with probability* 1*. If* $x \in (x(r), \hat{x})$*, the signal is independent of the agent's type and drawn from* $F(\cdot|\tilde{x} \leq x(r))$*. If* $x > \hat{x}$*, the signal equals* x *with probability* $\alpha(x) = \frac{x-\mathbb{E}[\tilde{x}|\tilde{x}<x(r)]}{1-\mathbb{E}[\tilde{x}|\tilde{x}<x(r)]}$ *or independent of* x *and drawn from* $F(\cdot|\tilde{x} \leq x(r))$ *with probability* $1-\alpha(x)$*. The police's equilibrium search strategy is as follows:*

$$p(s) = \begin{cases} s & \text{if} \quad s \leq x(r) \\ 1 & \text{if} \quad s > \hat{x}. \end{cases} \tag{C.4}$$

In this equilibrium, the agents commit a crime if and only if $x > x(r)$*.*

Proof. For each $x > x(r)$, define $\alpha(x)$ as the solution to the equation

$$x = \alpha(x) + (1-\alpha(x))\mathbb{E}[\tilde{x}|\tilde{x} < x(r)], \tag{C.5}$$

or

$$\alpha(x) = \frac{x - \mathbb{E}[\tilde{x}|\tilde{x} < x(r)]}{1 - \mathbb{E}[\tilde{x}|\tilde{x} < x(r)]} \in [0,1]. \tag{C.6}$$

Equation (C.5) means that type x is indifferent between committing and not committing a crime if they face search rate 1 with probability $\alpha(x)$ and search rate $\mathbb{E}[\tilde{x}|\tilde{x} < x(r)]$ with probability $1-\alpha(x)$.

Fix $\hat{x} > x(r)$ and consider the signal structure described in the statement. The set of possible signals is $[0, x(r)] \cup [\hat{x}, 1]$. Suppose that the police adopt the search strategy (C.4). Then, any type $x \in [0, x(r)]$ faces search probability x, so they are indifferent between committing and not committing a crime; hence, it is optimal for types below $x(r)$ to abstain from a crime. Any type $x \in (x(r), \hat{x})$ faces a search rate of at most $x(r)$ and strictly prefers to commit a crime. Finally, any type $x \geq \hat{x}$ faces search rate 1 with probability $\alpha(x)$ and the expected search rate of $\mathbb{E}[\tilde{x}|\tilde{x} < x(r)]$ with probability $1-\alpha(x)$. By (C.5), such types are indifferent between committing and not committing a crime, so it is optimal to commit a crime.

To sum up, given the signal structure and the police's strategy (C.4), there exist the agents' best responses such that all agents below type $x(r)$ abstain from committing a crime and thus the resulting crime rate is r.

It remains to show that the police's strategy is feasible and optimal. To begin with, note that the police's search strategy described above satisfies their search capacity constraint with equality if $\hat{x}$ satisfies

$$\mathbb{E}[\tilde{x}|\tilde{x} < x(r)] \left[F(\hat{x}) + \int_{\hat{x}}^{1} (1 - \alpha(t))f(t)\,\mathrm{d}t\right] + \int_{\hat{x}}^{1} \alpha(t)f(t)\,\mathrm{d}t = \overline{P}. \tag{C.7}$$

I show that there is an $\hat{x} \in (x(r), 1)$ that solves (C.7). If $\hat{x} = 1$, the right-hand side (RHS) becomes weakly greater, because the left-hand side (LHS) is

$$\mathbb{E}[\tilde{x}|\tilde{x} < x(r)] \leq \mathbb{E}[\tilde{x}|\tilde{x} < \hat{c}] = \overline{P},$$

where the equality comes from equation (6). If $\hat{x} = x(r)$, then the LHS exceeds the RHS in (C.7) for the following reason. In this case, there is no type in $(\hat{x}, x(r))$, so all types are indifferent between committing and not committing a crime under the search strategy described above. Such a search strategy must violate Assumption 1, which means that the LHS exceeds the RHS.

The intermediate value theorem implies that there is some $\hat{x} \in (x(r), 1]$ that solves (C.7). Therefore, for such an $\hat{x}$, the police's search strategy satisfies their search capacity constraint.

It remains to show that the police have no profitable deviations. The posterior crime rate for signal $s \in [\hat{x}, 1]$ is 1. For signals in $[0, x(r)]$, the posterior crime rate is less than 1 and equalized across all signals. Thus, the police's strategy is optimal so long as it satisfies $p(s) = 1$ for all $s \in [\hat{x}, 1]$ and exhausts search capacity $\overline{P}$. By construction, the police's strategy described above satisfies these conditions.

Finally, in this equilibrium, every type below $x(r)$ chooses not to commit a crime and faces a search rate x. The corresponding search mass on innocents is $\int_0^{x(r)} x\,\mathrm{d}F(x)$. $\square$

D Proof of Proposition 3

I solve the designer's problem in two steps. First, I fix an allocation policy $\overline{a}$ arbitrarily and then solve for the optimal signal structures. As discussed in the main text, if $\mathbb{E}_{F_t}[\tilde{x}] \geq \frac{\overline{a}(t)}{m(t)}$, the equation $\mathbb{E}_{F_t}[\tilde{x}|\tilde{x} \leq c] = \frac{\overline{a}(t)}{m(t)}$ has a unique solution $c_t\left(\frac{\overline{a}(t)}{m(t)}\right)$, and the designer can attain the minimal crime rate by providing the police with the truth-or-noise signal structure with cutoff $c_t\left(\frac{\overline{a}(t)}{m(t)}\right)$. If $\mathbb{E}_{F_t}[\tilde{x}] < \frac{\overline{a}(t)}{m(t)}$, define $c_t\left(\frac{\overline{a}(t)}{m(t)}\right) = 1$. In this case, the minimal crime rate is indeed $1 - F_t\left(c_t\left(\frac{\overline{a}(t)}{m(t)}\right)\right) = 0$, because the designer can provide the police with full information, and the police can deter all crimea crime (by the agents with pre-signal t) with search strategy $p(x) = x$ for all $x \in [0, 1]$. Overall, the crime-minimizing (S_t, π_t) leads to the equilibrium crime rate of $1 - F_t\left(c_t\left(\frac{\overline{a}(t)}{m(t)}\right)\right)$. The designer's optimal allocation policy then solves (16).

I now show that $\mathbb{E}_{F_t}[\tilde{x}] < \frac{\overline{a}(t)}{m(t)}$ never arises at the optimal allocation policy chosen by the designer. To see why, note that even in this extended model, the equilibrium crime rate is strictly positive for at least one pre-signal, say $t' \in \overline{S}$, because the designer faces a search capacity $\overline{P}$ that satisfies Assumption 1. Suppose to the contrary that $\mathbb{E}_{F_t}[\tilde{x}] < \frac{\overline{a}(t)}{m(t)}$ holds. Then the designer can improve its payoff in the following way. First, the designer slightly decreases the search capacity $\overline{a}(t)$ allocated to pre-signal t by $\epsilon > 0$ so that $\mathbb{E}_{F_t}[\tilde{x}] < \frac{\overline{a}(t)-\epsilon}{m(t)}$ holds. The resulting crime rate for pre-signal t remains 0. Then, the designer increases the search capacity allocated to a pre-signal (say t') that has a positive crime rate by ϵ. Originally, the signal structure for pre-signal t' was the truth-or-noise signal with cutoff c where $\mathbb{E}_{F_{t'}}[\tilde{x}|\tilde{x} \leq c] = \frac{\overline{a}(t)}{m(t)}$. For a small $\epsilon > 0$, the cutoff now increases to $c_\epsilon > c$ that solves $\mathbb{E}_{F_{t'}}[\tilde{x}|\tilde{x} \leq c_\epsilon] = \frac{\overline{a}(t)+\epsilon}{m(t)}$. The designer can then replace $(S_{t'}, \pi_{t'})$ with the truth-or-noise signal with cutoff c_ϵ. Because $c_\epsilon > c$, this will reduce the equilibrium crime rate for pre-signal t' as well as the overall crime rate. But then we obtain a contradiction. Therefore, under the designer's optimal search strategy, $\mathbb{E}_{F_t}[\tilde{x}] \geq \frac{\overline{a}(t)}{m(t)}$ holds, so $\mathbb{E}_{F_t}[\tilde{x}|\tilde{x} \leq c] = \frac{\overline{a}(t)}{m(t)}$ has a unique solution c for every $t \in \overline{S}$. This proves the last part of Proposition 3. $\square$

E Omitted Materials for Section 4.4

E.1 Extending Theorem 0

Proposition 7. *Consider the case of endogenous search capacity. Suppose that the police have full information. An equilibrium exists, and in any equilibrium, almost every agent* *commits a crime with probability* 1.

Proof. First, I construct an equilibrium. Pick a unique $\bar{x}$ that solves $C'\left(\int_0^{\bar{x}} t \, \mathrm{d}F(t)\right) = 1$. Consider the strategy profile such that the police adopt a search strategy $p^*(x) = x$ for $x \leq \bar{x}$ and $p^*(x) = 0$ for $x > \bar{x}$, and all agents commit a crime with probability 1. The agents' strategies are optimal. The police's search strategy is also optimal, because the police's payoff from searching mass P of agents is $P - C(P)$, so the strategy is optimal if the police search mass P^* of agents where $C'(P^*) = 1$. Search strategy p^* has this property by construction.

Second, take any equilibrium. Let p be the equilibrium search strategy and $P = \int_0^1 p(x) \, \mathrm{d}F(x)$ be the total mass of searches under p. We have $C'(P) \leq 1$, because otherwise the police could profitably deviate by slightly lowering the mass of searches. Combining $C'(P) \leq 1$ with $1 < C'\left(\int_0^1 x \, \mathrm{d}F(x)\right)$, we obtain $P < \int_0^1 x \, \mathrm{d}F(x)$. We can then apply the same argument as the proof of Theorem 0: The set $X \triangleq \{x \in [0,1] : x > p(x)\}$ has a positive mass, and any type in X commits a crime with probability 1. Let Y be the set of types that commit a crime with probability strictly below 1. If Y has a positive mass, the police must be allocating a positive mass of searches to types in Y, because otherwise they would commit a crime. We then obtain a contradiction, because the police could increase the mass of successful searches without changing the search cost $C(P)$ by maintaining the total search capacity at P and shifting search probabilities from types in Y to X. Thus, the set Y has measure zero, i.e., almost every agent commits a crime with probability 1. □

E.2 Solving the Relaxed Problem with Endogenous Search Capacity

The following lemma shows that the joint signal structure illustrated by Figure 2 solves the relaxed problem under endogenous search capacity.

Lemma 2. *Consider endogenous search capacity. In the relaxed problem, the following joint 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. *For every type $x > c^*$, the realized signal is $(s_i^P, s_i^A) = (\rho^*, crime)$ with probability* 1*. For every type $x < c^*$, the realized signal is $(0, not)$ or (ρ^*, not) with probability α^* or $1-\alpha^*$, respectively. In equilibrium, the agents and the police follow the action recommendations.*

2. *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.*

In what follows, a joint signal structure is referred to as a signal structure. Take any signal structure (S_P, S_A, π) and any strategies chosen by agents. For each $s \in S_P$, the *posterior crime rate*, $r(s) \in [0,1]$, is defined as the probability that an agent commits a crime conditional on the police's signal s (but not conditional on the agent's signal).[25]

To prove Lemma 2, I first prove the following lemma to restrict the class of signal structures that need to be considered.

Lemma 3. *Consider the relaxed problem with endogenous search capacity, and take any signal structure (S'_P, S'_A, π') and any equilibrium with crime rate $r \in (0,1)$. There is some $c \in (0,1)$ such that the same crime rate arises under a signal structure (S_P, S_A, π) and an equilibrium with the following properties:*

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

2. *The police's signal is a garbling of an agent's signal, i.e., there exist conditional distributions of the police's signal given an agent's signal, denoted by $\hat{\pi}_P(\cdot|crime), \hat{\pi}_P(\cdot|not) \in \Delta S_P$, such that for any $S \subset S_P$, $a \in \{crime, not\}$, and $x \in [0,1]$, we have $\pi(S \times \{a\}|x) = \hat{\pi}_P(S|a)\pi_A(a|x)$, where $\pi_A(a|x) \triangleq \pi(S_P \times \{a\}|x)$. In equilibrium, each signal of the police leads to a distinct posterior crime rate.*

Proof. Take any signal structure (S'_P, S'_A, π') and any equilibrium with crime rate $r \in (0,1)$. Let p' denote the police's equilibrium search strategy. First, as in Lemma 1, replace the

[25] Let $a(t)$ denote the probability that an agent commits a crime after observing signal $t \in S_A$. The posterior crime rate is given by $r(s) = \mathbb{E}[a(\tilde{t})|s]$, where the expectation is with respect to the agent's signal $\tilde{t}$ conditional on the police's signal $s \in S_P$.

signal space S'_A of agents with $S_A = \{crime, not\}$ and assume that each agent follows the action recommendation in equilibrium.

Second, replace each signal $s' \in S'_P$ of the police with its posterior crime rate $r(s')$ induced by the agents' strategies. Let $S_P \subset [0,1]$ be the resulting signal space for the police. This modification reduces the police's information, because different signals in S'_P may have the same posterior crime rate. I then assume that the police adopt search strategy

$$p(y) \triangleq \mathbb{E}[p'(s')|r(s') = y], \forall y \in S_P, \tag{E.8}$$

where the expectation is with respect to the police's original signal $s' \in S'_P$ conditional on that the posterior crime rate associated with the signal equals y. The police find it optimal to adopt p because they can ensure the same mass of successful searches as p' despite having less information under S_P than under S'_P. The agents' incentives remain the same (i.e., the obedience constraints continue to hold), because strategies p' and p induce the same expected search probability conditional on each action recommendation. Indeed, as explained below, it holds that for each $a \in \{crime, not\}$, $\mathbb{E}_{s'}[p'(s')|a] = \mathbb{E}_y[p(y)|a]$, or more specifically,

$$\mathbb{E}_{s'}[p'(s')|a] = \mathbb{E}_y[\mathbb{E}_{s'}[p'(s')|r(s') = y, a]|a] = \mathbb{E}_y[\mathbb{E}_{s'}[p'(s')|r(s') = y]|a] = \mathbb{E}_y[p(y)|a]. \tag{E.9}$$

Here, $\mathbb{E}_{s'}[\cdot|a]$ is the expectation with respect to the police's signal $s' \in S'_P$ conditional on action recommendation a; $\mathbb{E}_y[\cdot|a]$ is the expectation with respect to the police's signal $y \in S_P$ as a posterior crime rate conditional on a; $\mathbb{E}_{s'}[\cdot|r(s') = y, a]$ is the expectation with respect to the police's signal $s' \in S'_P$ conditional on posterior crime rate y and action recommendation a; and $\mathbb{E}_{s'}[\cdot|r(s') = y]$ is the expectation with respect to the police's signal $s' \in S'_P$ conditional only on posterior crime rate y.

I now explain why the equalities in (E.9) hold. The first equality is from the law of iterated expectation. The last equality is from the definition of search strategy p in (E.8). To show the second equality, suppose to the contrary that $\mathbb{E}_{s'}[p'(s')|r(s') = y, a] \neq \mathbb{E}_{s'}[p'(s')|r(s') = y]$ for some $y \in S_P$ and $a \in \{crime, not\}$, which is equivalent to

$$\mathbb{E}_{s'}[p'(s')|r(s') = y, crime] \neq \mathbb{E}_{s'}[p'(s')|r(s') = y, not]. \tag{E.10}$$

Hereafter, we fix such a y. Let $r^{-1}(y) \subseteq S'_P$ be the set of all signals $s' \in S'_P$ such that $r(s') = y$. The condition (E.10) means that there is some set $T \subsetneq r^{-1}(y)$ such that

$$\Pr(s' \in T | r(s') = y, crime) > \Pr(s' \in T | r(s') = y, not) \tag{E.11}$$

and thus

$$\Pr(s' \in T^c | r(s') = y, crime) < \Pr(s' \in T^c | r(s') = y, not) \tag{E.12}$$

where $T^c \triangleq r^{-1}(y) \setminus T$. For each $U \in \{T, T^c\}$, we obtain

$$\frac{\Pr(crime|s' \in U)}{\Pr(not|s' \in U)} = \frac{\Pr(crime|s' \in U, r(s') = y)}{\Pr(not|s' \in U, r(s') = y)} = \frac{\Pr(crime|r(s') = y)}{\Pr(not|r(s') = y)} \cdot \frac{\Pr(s' \in U|r(s') = y, crime)}{\Pr(s' \in U|r(s') = y, not)}, \tag{E.13}$$

where the first equality holds because $s' \in U$ implies $r(s') = y$, and the second equality follows from Bayes' rule. Combining (E.11), (E.12), and (E.13), we obtain

$$\frac{\Pr(crime|s' \in T)}{\Pr(not|s' \in T)} > \frac{\Pr(crime|s' \in T^c)}{\Pr(not|s' \in T^c)},$$

or equivalently, $\Pr(crime|s' \in T) > \Pr(crime|s' \in T^c)$. This contradicts the fact that signals in T and T^c have the same posterior crime rate y. Therefore, the second equality in (E.9) is valid.

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 $\hat{\pi}_P(\cdot|a) \in \Delta S_P$ denote the associated conditional distribution of the police's signal given an agent's signal $a \in \{crime, not\}$. I then modify the signal structure as follows. First, given the equilibrium crime rate $r \in (0, 1)$, assume that types above and below cutoff $c \triangleq F^{-1}(1 - r)$ receive signals *crime* and *not*, respectively. Second, assume that conditional on each agent i's signal $a_i \in \{crime, not\}$, the police observe signal $y_i \sim \hat{\pi}_P(\cdot|a_i) \in \Delta S_P$ (regardless of i's type). These modifications (i) preserve the joint distribution of posterior crime rates and action recommendations across the population and (ii) relax the obedience constraints for the agents. As a result, the police optimally choose search strategy p defined in (E.8), and the agents continue to follow action recommendations. The resulting equilibrium has the properties stated in the lemma. □

I now prove Lemma 2.

Proof of Lemma 2. Take any signal structure (S_P, S_A, π) and any equilibrium that satisfy the properties described in Lemma 3 and have crime rate $r \in (0,1)$. Part 2 of the lemma ensures that the police's signal is a garbling of an agent's signal, generated by conditional distributions $\hat{\pi}_P(\cdot|crime), \hat{\pi}_P(\cdot|not) \in \Delta S_P$. Without loss, assume that the police's signals are recommended search probabilities that the police follow in equilibrium.

First, I 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 3 that each signal leads to a distinct posterior crime rate.[26]

The unique interior search rate ρ (if exists) satisfies two properties. First, $\{\rho, 1\} \subset S_P$ implies $r(\rho) < r(1)$, because if $r(\rho) > r(1)$, the police would profitably deviate by shifting search masses from signal 1 to signal ρ (recall $r(\rho) \neq r(1)$ from Lemma 3). Second, the police equate the marginal cost of search with the marginal probability of detecting a crime. Hence, the equilibrium total search capacity P solves $C'(P) = r(\rho)$. Otherwise, the police would profitably deviate by slightly changing $p(\rho)$.

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, with the agents' strategies fixed, after we pool 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 the marginal return on searches at P is now $r(\sigma) - C'(P) > r(\rho) - C'(P) = 0$. This pooling also reduces the police's information about the agents' signals (and their behavior) and thus increases the fraction of searches that go to signal *not*. As a result, if we fix the agents' strategies but let the police adopt

[26] 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.

an optimal search strategy, pooling signals ρ and 1 increases the expected search probability conditional on signal *not* and relaxes its obedience constraint.

The pooling procedure in the previous paragraph changes the police's signal space and the distribution of the police's signal conditional on the agent's action recommendation. However, to reduce notational burden, I continue using S_P for the police's signal space and $(\hat{\pi}_P(\cdot|crime), \hat{\pi}_P(\cdot|not))$ for the conditional distributions of the police's signal.

We now have a signal structure and a strategy profile such that $S_P = \{0, \sigma\}$ or $\{\sigma\}$ for some $\sigma > 0$, and the obedience constraint for signal *not* holds (possibly strictly).[27] If $S_P = \{0, \sigma\}$, we have $r(0) < r(\sigma)$, because otherwise the police would deviate by shifting search masses from signal σ to signal 0. ★ We then continuously increase cutoff type c defined in Lemma 3 while (i) letting the agents follow action recommendations (i.e., they receive signal *crime* and thus commit a crime if and only if $x > c$); (ii) maintaining conditional distributions $(\hat{\pi}_P(\cdot|crime), \hat{\pi}_P(\cdot|not))$ that garble the agents' signals to create the police's signals, and (iii) letting the police adopt the optimal search strategy at any given c. Increasing cutoff c continuously changes the expected search probabilities the agents face conditional on signals *crime* and *not*. At $c = 1$, all agents receive and follow 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 committing crime and not after receiving signal *not*. The obedience constraint for signal *crime* also holds at the same c^*, because the assumption $C'\left(\int_0^1 x \, \mathrm{d}F(x)\right) > 1$ implies that the police's optimal searches never make all agents weakly prefer to abstain from committing a crime.

At this point, we have a signal structure and an equilibrium such that: each agent receives signal *crime* if $x > c^*$ or *not* if $x < c^*$, they follow action recommendations, and those who receive signal *not* are indifferent between the two actions; and the police's signal space S_P is $\{0, \sigma\}$ or $\{\sigma\}$. If $S_P = \{0, \sigma\}$, the inequality $r(0) < r(\sigma)$ continues to hold even though we increased cutoff c, because whether $r(0) < r(\sigma)$ holds depends only on $\frac{\hat{\pi}_P(0|crime)}{\hat{\pi}_P(0|not)}$ and $\frac{\hat{\pi}_P(\sigma|crime)}{\hat{\pi}_P(\sigma|not)}$, not on c.

In the last step, we consider two cases. If $S_P = \{\sigma\}$, we obtain the desired result where

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

$\alpha^* = 0$ (note that $(s_i^P, s_i^A) = (0, not)$ is irrelevant if $\alpha^* = 0$).

Otherwise, the signal space is $\{0, \sigma\}$. In this case, we first change the police's signals back to the corresponding posterior crime rates, which results in a signal space $\{r_0, r_\sigma\}$ with $r_0 < r_\sigma$. 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.

We now need to consider two cases. One is when $p(\sigma) \in (0, 1)$, i.e., before splitting signal r_0 into t_0 and t_σ, the search rate for signal r_σ is interior. In this case, we must have $C'(P) = r_\sigma$ and the police do not search signal r_0 at all before splitting. Thus, after splitting, it continues to be optimal for the police to allocate search mass P to signal r_σ and not search signal t_σ or t_0. In particular, searching agents with signal t_σ (in addition to those with r_σ) will entail a higher marginal cost than $C'(P)$. Now, signals t_σ and r_σ have the same posterior crime rate, so we can pool them into the same signal, say u_σ. After this pooling, the police still find it optimal to apply search capacity P to signal u_σ and search capacity 0 to signal t_0, and the expected search probability conditional on each action recommendation remains the same.

The other case is when $p(\sigma) = 1$, i.e., before splitting signal r_0 into t_0 and t_σ, the police search signal r_σ with probability 1. In this case, we may have $C'(P) < r_\sigma$ before splitting. Thus, after splitting, the police continue to search agents with signal r_σ with probability 1, and may also prefer to search agents with signal t_σ with a positive probability. As a result, the obedience constraint for signal *not* may hold strictly—the agents used to face search mass P after signal r_σ and no search after signal r_0, but they now face search mass P after signal r_σ and possibly a positive search mass after signal t_σ. We then proceed as follows: First, we pool signals t_σ and r_σ into signal u_σ. As in the previous paragraph, signals t_σ and r_σ have the same posterior crime rate, and this pooling does not change the police's search capacity or the agents' incentives. If the obedience constraint for signal *not* holds strictly, we increase cutoff c from c^* as we did in the second step (the procedure ★ above) until the cutoff hits the point at which the obedience constraint for signal *not* binds and the one for signal *crime* holds. Redefine c^* as this new cutoff.

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 a positive search probability. In terms of action recommendations, the police's signal space becomes $S_P^* = \{0, \rho^*\}$ for some $\rho^* \in (0, 1]$.

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; the agents who receive signal *not* are indifferent; 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 Part 1. Part 2 holds because I have constructed the outcome so that the agents with signal *not* are indifferent.

Finally, note that at the beginning of this proof, I chose a signal structure (S_P, S_A, π) whose corresponding equilibrium crime rate r lies in $(0, 1)$. This choice is without loss of generality in the following sense. We do not need to consider $r = 0$ because it never arises: If no agent commits a crime, the police with a strictly increasing cost function would never search, which, in turn, implies that everyone would commit a crime. Similarly, we do not need to consider $r = 1$ to derive the crime-minimizing joint signal structure, because we can provide the agents with full information and the police with no information. The resulting equilibrium crime rate is strictly less than 1, because $C'(0) < 1$. □

E.3 Solving the Original Problem: The Case of $\alpha^* = 0$

The following lemma shows that if $\alpha^* = 0$ holds in the solution to the relaxed problem characterized by Lemma 2, a truth-or-noise signal structure solves the original problem with endogenous search capacity.

Lemma 4. *Consider the case of endogenous search capacity. Suppose that the solution to the relaxed problem described in Lemma 2 has $\alpha^* = 0$, i.e., the police receive no information. Then, the crime-minimizing signal structure of the original problem is the truth-or-noise signal structure with cutoff c^*, where c^* and the equilibrium search capacity P^* jointly solve*

$$\mathbb{E}_F[\tilde{x} | \tilde{x} \leq c^*] = P^*, \text{ and} \tag{E.14}$$

$$1 - F(c^*) = C'(P^*). \tag{E.15}$$

Proof. Equation (E.14) and the same argument as Theorem 1 imply that under the truth-or-noise signal structure with cutoff c^*, there exists a strategy profile such that: the agents optimally commit a crime if $x > c^*$ and not if $x < c^*$; the police adopt search strategy $p(s) = s$ for every $s \in [0, c^*]$; and the police cannot increase their payoffs by changing the search strategy while keeping the total search capacity P^* fixed. Thus, it remains to show that the police have no profitable deviation in terms of changing the total search capacity. The police's payoff from search capacity P is $(1 - F(c^*))P - C(P)$, because the posterior crime rate is equalized to be $1 - F(c^*)$ across all signals. Due to the convexity of $C(\cdot)$, equation (E.15) is sufficient for the optimality of P^*. □

We are now ready to prove Proposition 4.

Proof of Proposition 4. Thanks to Lemma 4, it suffices to show that $\alpha^* = 0$ holds in the relaxed problem. Thus, I use Lemma 2 to solve the relaxed problem and show that $\alpha^* = 0$ holds under inequality (17).

I focus on signal structures that take the form described in Figure 2. Instead of parameters (ρ^*, c^*, α^*), we use (ρ, c, α) 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*. When types above some cutoff commit a crime, a crime rate r pins down the cutoff type through $c = F^{-1}(1 - r)$.

I fix $\alpha \in [0, 1]$ arbitrarily and then determine the cutoff type c and the unique positive search probability ρ from the mutual best responses of the agents and the police. By Part 2 of Lemma 2, 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.

I 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 $C'(P) = LP^\beta = \frac{r}{r+(1-r)(1-\alpha)}$, or

$$P = \frac{1}{L^{\frac{1}{\beta}}} \cdot \left(\frac{r}{r + (1 - r)(1 - \alpha)} \right)^{\frac{1}{\beta}}.$$

For the moment, I ignore the constraint $\rho \leq 1$ and verify it later. The expected search probability conditional on signal *not* is $(1-\alpha)\frac{P}{r+(1-r)(1-\alpha)}$, or

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

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

$$I(r,\alpha) = \mathbb{E}_F[\tilde{x}|\tilde{x} \le F^{-1}(1-r)]. \tag{E.16}$$

At $r=0$, we have $I(0,\alpha) = 0 < \mathbb{E}_{x\sim F}[x]$. At $r=1$, we have $I(1,\alpha) = \frac{1-\alpha}{L^{\frac{1}{\beta}}} \ge 0$. Also, $I(r,\alpha)$ and $\mathbb{E}_{\tilde{x}\sim F}[\tilde{x}|\tilde{x} \le F^{-1}(1-r)]$ are continuous in r. Thus equation (E.16) has a solution. Let $r(\alpha)$ denote the smallest solution. 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, I first derive $\alpha(r) \triangleq \arg\max_{\alpha\in[0,1]} I(r,\alpha)$ and then determine the minimal crime rate r^* through $I(r,\alpha(r)) = \mathbb{E}_F[\tilde{x}|\tilde{x} \le F^{-1}(1-r)]$. Hereafter, I restrict attention to $\alpha \in [0,1)$ and $r \in (0,1)$, because $\alpha = 1$ or $r \in \{0,1\}$ cannot be a part of a crime-minimizing equilibrium.

We have

$$\frac{\partial}{\partial\alpha}\log I(r,\alpha) = -\frac{1}{1-\alpha} + \frac{1+\beta}{\beta}\cdot\frac{1-r}{1-\alpha(1-r)}.$$

Note that $(1-\alpha)(1-\alpha(1-r)) > 0$, but the expression

$$(1-\alpha)(1-\alpha(1-r))\frac{\partial}{\partial\alpha}\log I(r,\alpha) = -r + \frac{1}{\beta}(1-r)(1-\alpha)$$

changes its sign at most once from positive to negative as α increases, and $\frac{\partial}{\partial\alpha}\log I(r,\alpha) < 0$ for α close to 1. As a result, $\frac{\partial}{\partial\alpha}\log I(r,\alpha)$ is either always negative or changes its sign exactly once from positive to negative. Examining the first-order condition $\frac{\partial}{\partial\alpha}\log I(r,\alpha) = 0$, we obtain the following solution:

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

and

$$I(r,\alpha(r)) = \begin{cases} \frac{1}{L^{\frac{1}{\beta}}} \cdot \frac{\beta}{(1+\beta)^{\frac{1+\beta}{\beta}}} \frac{1}{1-r} & \text{if} \quad r \leq \frac{1}{1+\beta}, \\ \left(\frac{r}{L}\right)^{\frac{1}{\beta}} & \text{if} \quad r \geq \frac{1}{1+\beta}. \end{cases}$$

Now, go back to the equation

$$I(r,\alpha(r)) = \mathbb{E}_{\tilde{x}\sim F}[\tilde{x}|\tilde{x} \leq F^{-1}(1-r)].$$

The left-hand side is strictly increasing, and the right-hand side is strictly decreasing in r. Hence the equilibrium crime rate r^* is unique. Moreover, we have $r^* \geq \frac{1}{1+\beta}$ and thus $\alpha(r^*) = 0$ when

$$I\left(\frac{1}{1+\beta}, 0\right) < \mathbb{E}_{\tilde{x}\sim F}\left[\tilde{x} \mid \tilde{x} \leq F^{-1}\left(\frac{\beta}{1+\beta}\right)\right],$$

which reduces to inequality (17). Finally, recall that when I derived the police's best response, I temporarily ignored the condition that the search probability ρ^* can be at most 1. This condition is satisfied at equilibrium because ρ^* satisfies $\rho^* \leq \mathbb{E}_F[\tilde{x}|\tilde{x} \geq F^{-1}(1-r^*)] \leq 1$. □

E.4 Solving the Original Problem with $\alpha^* > 0$: Example

Proposition 8. *Consider the case of endogenous search capacity. Suppose that $F = U[0,1]$ and $C(P) = \frac{L}{2}P^2$ with $L \geq \frac{3+\sqrt{5}}{4} \approx 1.31$. In the original problem, the equilibrium crime rate is minimized by a signal structure that reveals a fraction $\alpha^* = \max(0, 2-\sqrt{2L})$ of the agents with type below c^* and discloses information according to the truth-or-noise signal structure with cutoff c^* for the rest of agents. Here, α^* and c^* are parameters in the solution to the relaxed problem characterized by Lemma 2.*

Proof. Substituting $F = U[0,1]$ and $\beta = 1$ into the solution to the relaxed problem shown in the proof of Proposition 4, I obtain the following:

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

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 $I(r,\alpha(r)) = \frac{1-r}{2}$, I 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 $c^* = 1 - r^*$.

I now show that if $L \ge L^* = \frac{3+\sqrt{5}}{4}$, crime rate r^* can be implemented in the original problem. To do so, I modify the signal structure in Theorem 1 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 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, $L \ge L^*$ ensures that the highest search probability $\frac{c^*}{1-\alpha^*}$ is below 1, so the police's strategy is feasible. 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 $L \ge 2$, we have $\alpha^* = 0$, so the police's signal reduces to the truth-or-noise signal structure. $\square$

Online Appendix
(For Online Publication Only)

A Proof of the Claim in Footnote 16

The argument in Footnote 16 rests on the following result:

Proposition 9. *Consider any deterministic signal structure* (S, π)*, i.e., for each type* $x \in [0, 1]$*, the distribution* $\pi(\cdot|x)$ *over signals is degenerate at some* $s \in S$*. Then in any equilibrium, the crime rates are equal across almost all signals, and thus the police are indifferent between any search strategies that exhaust search capacity* $\overline{P}$*.*

Proof. If the signal structure is deterministic, each agent, who knows their type, also knows their signal. Take any equilibrium, and let $r(S')$ be the crime rate conditional on a set of signals $S' \subseteq S$. Suppose to the contrary that $r(S_2) > r(S_1)$ for some sets of signals, S_1 and S_2, that can arise with positive ex ante probabilities under (S, π). The inequalities imply $r(S_1) < 1$ and $r(S_2) > 0$. Inequality $r(S_1) < 1$ means that the police allocate positive search mass to signals in S_1, and $r(S_2) > 0$ implies that the search rate for signals in S_2 is strictly below 1.[28] Then the police would profitably deviate by shifting search mass from signal S_1 to S_2. This is a contradiction, so the crime rate must be equalized across almost all signals. □

B Omitted Materials for Section 4.3

B.1 Example for Proposition 3

In the following example, I use Proposition 3 to solve the designer's allocation problem.

[28] This argument could fail under a stochastic signal structure. For example, if some agents receive signal s_1 or s_2 with a positive probability, it is possible that, even though the police search agents with signal s_2 with probability 1, these agents still commit a crime because the search rate for s_1 may be low, and the agents may be uncertain about whether their signals will be s_1 or s_2. If we assume that the realized signals are publicly observable, Proposition 9 holds under any signal structure.

Example 2. There are two equally likely pre-signals, 1 and 2, such that $F_1(x) = F_2(x) = F(x) = x^\beta$ with $\beta > 0$, i.e., the pre-signals divide the population into two groups of equal size and type distribution. Since $F_1 = F_2$, I omit the dependence of F_t and c_t on t. For simplicity, assume that $0.5\mathbb{E}_F[x] < \overline{P} < \mathbb{E}_F[x]$. Because $m(1) = m(2) = 0.5$, the designer's problem (16) is

$$\min_{\overline{a}(1),\overline{a}(2)\geq 0} 0.5\left[1 - F\left(c\left(2\overline{a}(1)\right)\right)\right] + 0.5\left[1 - F\left(c\left(2\overline{a}(2)\right)\right)\right]$$
$$\text{subject to} \quad \overline{a}(1) + \overline{a}(2) = \overline{P},$$

or equivalently,

$$\max_{\overline{a}(1),\overline{a}(2)\geq 0} F\left(c\left(2\overline{a}(1)\right)\right) + F\left(c\left(2\overline{a}(2)\right)\right) \tag{O.1}$$
$$\text{subject to} \quad \overline{a}(1) + \overline{a}(2) = \overline{P}.$$

Direct calculation reveals that for each $t \in \{1, 2\}$,[29]

$$c\left(2\overline{a}(t)\right) = \frac{2(1+\beta)}{\beta}\overline{a}(t) \quad \text{and} \quad F\left(c\left(2\overline{a}(t)\right)\right) = \left(\frac{2(1+\beta)}{\beta}\right)^{\beta} \overline{a}(t)^{\beta}.$$

The solution to (O.1) then becomes as follows: If $\beta > 1$, each term $F\left(c\left(2\overline{a}(t)\right)\right)$ of the objective in (O.1) is strictly convex in $\overline{a}(t)$, which implies that the designer should first allocate as much search capacity as possible to one pre-signal and, once that pre-signal has a crime rate of 0 (or $F\left(c\left(2\overline{a}(t)\right)\right) = 1$), allocate the remaining search capacity to the other pre-signal. Specifically, for one pre-signal, say 1, the designer sets $\overline{a}(1) = 0.5\mathbb{E}_F[x]$ and provides the police with full information, so that the police can search each type x with probability x and attain a crime rate of 0. For pre-signal 2, the designer sets $\overline{a}(2) = \overline{P} - 0.5\mathbb{E}_F[x]$ and provides the police with the truth-or-noise signal structure with cutoff

$$c\left(2\overline{a}(2)\right) = \frac{2(1+\beta)}{\beta}\left(\overline{P} - 0.5\mathbb{E}_F[x]\right).$$

[29] It holds that $\mathbb{E}[x|x \leq c] = \int_0^c \frac{x \cdot \beta x^{\beta-1}}{c^\beta}\,\mathrm{d}x = \frac{\beta}{1+\beta}c$, so $\mathbb{E}[x|x \leq c] = 2\overline{a}(t)$ implies $c\left(2\overline{a}(t)\right) = \frac{2(1+\beta)}{\beta}\overline{a}(t)$.

In contrast, if $\beta \in (0,1)$, each term $F\left(c\left(2\overline{a}(t)\right)\right)$ of the objective in (O.1) is strictly concave in $\overline{a}(t)$, which implies that the designer should set $\overline{a}(1) = \overline{a}(2) = 0.5\overline{P}$ and provide the police with the truth-or-noise signal structure with cutoff

$$c\left(2\overline{a}(t)\right) = \frac{1+\beta}{\beta}\overline{P}$$

for both pre-signals.

B.2 Extending Theorem 0

Next, I study the case in which the designer can control an allocation policy based on pre-signals as in Section 4.3 but must provide the police with full information. We begin with a lemma that extends Theorem 0:

Lemma 5. *Suppose that signal structure (S_t, π_t) is fully informative for every $t \in \overline{S}$. Take any pre-signal t. If $\overline{a}(t) \geq m(t)\int_0^1 x\,\mathrm{d}F_t(x)$, then in the game between the police and the agents with pre-signal t, there exists an equilibrium with a crime rate of 0. If $\overline{a}(t) < m(t)\int_0^1 x\,\mathrm{d}F_t(x)$, any equilibrium has a crime rate of 1.*

Proof. To simplify notation, without loss of generality, we set $m(t) = 1$. Suppose that $\overline{a}(t) \geq \int_0^1 x\,\mathrm{d}F_t(x)$. First, we construct an equilibrium with a crime rate of 0. Suppose that the police adopt search strategy $p^*(x) = x, \forall x \in [0,1]$. This search strategy is feasible because it induces the total search capacity of at most $\int_0^1 p^*(x)\,\mathrm{d}F_t(x) \leq \overline{a}(t)$. Under p^*, it is optimal for every agent to not commit a crime . The police then find it optimal to choose any search strategy because the mass of successful searches is 0 regardless of the search strategy. Thus we obtain an equilibrium with crime rate 0. The second part follows from Theorem 0. $\square$

I now consider the designer's allocation problem when the police have full information. First, minimizing crime is equivalent to maximizing the mass of agents who do not commit a crime (which we call the mass of innocents) in equilibrium. Thus, we assume that the designer's payoff equals the mass of innocents. By Lemma 5, the designer's payoff from pre-signal t is $m(t)$ if $\overline{a}(t) \geq m(t)\int_0^1 x\,\mathrm{d}F_t(x)$ and 0 otherwise. It is then without loss of generality to assume that for each pre-signal t, the designer either allocates $m(t)\int_0^1 x\,\mathrm{d}F_t(x)$ or nothing. The designer's problem can then be restated as follows:

$$\max_{D\subseteq\overline{S}}\sum_{t\in D} m(t)$$
$$\text{subject to}\quad \sum_{t\in D} m(t)\int_0^1 x\,\mathrm{d}F_t(x)\leq\overline{P}.$$

This is a knapsack problem in which the set of items is $\overline{S}$, the weight of each item $t\in\overline{S}$ is $m(t)\int_0^1 x\,\mathrm{d}F_t(x)$, the value of item t is $m(t)$, and the maximum capacity is $\overline{P}$.

C A Continuum of Officers

In this appendix, I formalize the idea that the simultaneous-move assumption arises when a continuum of individual officers searches agents. Suppose that instead of the police, there is a unit mass of *officers*, $j\in[0,1]$. All officers face the same signal structure (S,π) and observe the same realized signal for each agent.

The timing of the game is as follows. First, each officer chooses a search strategy, $p_j : S\to[0,1]$. The profile of search strategies, $(p_j)_{j\in[0,1]}$, is publicly observed by all agents. Second, nature draws the type and the signal for each agent. Finally, each agent observes their type but not their realized signal, and decides whether to commit a crime.

Each officer j chooses a search strategy p_j to maximize the mass of successful searches, defined as

$$\sigma_j\triangleq\int_0^1\int_S p_j(s)\,\mathrm{d}\pi(s|x)a(x)\,\mathrm{d}F(x),\tag{O.2}$$

where $a(x)$ is the probability with which type x commits a crime. As in the baseline model, we obtain identical results if each officer's payoff is strictly increasing in σ_j and arbitrarily dependent on the crime rate. Note that the probability $a(x)$ of type x committing a crime may depend on the chosen profile of search strategies, but I do not write this dependency explicitly. Each officer faces the same search capacity constraint, $\overline{P}<\int_0^1 x\,\mathrm{d}F(x)$, which is analogous to Assumption 1. Identical results hold if there is a mass $\overline{P}$ of officers, each of whom has a search capacity of 1.

Given the search strategies $(p_j)_{j\in[0,1]}$, define *the aggregate search rate for signal* s as

$$\hat{p}(s) \triangleq \int_0^1 p_j(s)\,\mathrm{d}j. \tag{O.3}$$

Abusing notation, define *the aggregate search rate for type* x as

$$\hat{p}(x) \triangleq \int_S \hat{p}(s)\,\mathrm{d}\pi(s|x). \tag{O.4}$$

An agent's expected payoffs of committing a crime and not are $x - \hat{p}(x)$ and 0, respectively.

C.1 Extending Theorem 0

Proposition 10. *Suppose there is a continuum of officers and they have full information. An equilibrium exists, and in any equilibrium, almost every agent commits a crime with probability* 1*.*

Proof. A unilateral deviation by an individual officer does not affect the aggregate search rate (O.4) for each type. Thus, for any fixed profile of search strategies, the set of best responses by each agent does not change with or without an individual officer's deviation. Then we can construct an equilibrium in which all officers adopt the same search strategy p^* as the police in the first half of the proof of Theorem 0.

Second, take any equilibrium. There is no profile of search strategies such that the induced aggregate search rate satisfies $\hat{p}(x) \geq x$ for (almost) every type x, because it would violate the search capacity constraint of a positive mass of officers. Thus, the set $X \triangleq \{x \in [0,1] : x > p(x)\}$ has a positive mass, and any type in X commits a crime with probability 1. By the same logic as the proof of Theorem 0, the set Y of types that commit a crime with probability strictly below 1 must have measure zero. □

C.2 Extending Theorem 1

I characterize the crime-minimizing signal structure with a continuum of officers. First, to extend Lemma 1, we define the relaxed problem in this new setup: The information of the agents and the officers is now determined by a joint signal structure, (S_P, S_A, π). The timing of the game is as follows. First, each officer chooses a search strategy, $p_j : S_P \to [0,1]$. The profile of search strategies, $(p_j)_{j\in[0,1]}$, is publicly observed by all agents. Second, nature draws the type x_i and signals (s_i^P, s_i^A) for each agent i. Finally, each agent observes s_i^A but not (x_i, s_i^P), and decides whether to commit a crime.

Proposition 11. *Suppose there is a continuum of officers. In the relaxed problem, the following joint signal structure minimizes a crime rate: The officers 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*

$$\mathbb{E}_F[\tilde{x}|\tilde{x} \le \hat{c}] = \overline{P}. \tag{O.5}$$

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

Proof. The proof follows that of Lemma 1. Take any joint signal structure (S_P, S_A, π) and equilibrium. Let $\hat{p} : S_P \to [0,1]$ be the aggregate search rate function, and let r be the crime rate in equilibrium. First, replace the agents' signals with action recommendations so that the obedience constraints hold (see the proof of Lemma 1). This step is valid even though the agents move after observing the chosen search strategies, because a single officer's deviation does not change the aggregate search rates and thus does not affect the agents' incentives. In other words, after replacing the agents' signals with action recommendations, the agents' behavior remains the same both on-path and after an officer's deviation. Second, by the same logic as Lemma 1, providing the officers with no information only relaxes the obedience constraints. The resulting problem is to choose a disclosure policy from which the agents receive action recommendations. The rest of the proof is the same as that of Lemma 1. □

Proposition 12. *In the model with a continuum of officers, Theorem 1 holds verbatim.*

The proof of this claim is identical to that of Theorem 1 and thus omitted.