Preventing crime waves[1]

Philip Bond
University of Pennsylvania

Kathleen Hagerty
Northwestern University

December 2009

[1]We are grateful for helpful comments from Itay Goldstein, Kathryn Spier, Kenneth Hendricks (the editor), two referees, and from audiences at the Institute for Advance Study and the American Law and Economics Association meetings. Bond thanks the Institute for Advanced Study for hospitality and financial support (in conjunction with Deutsche Bank). Any remaining errors are our own.

Abstract

We study the design of enforcement mechanisms when enforcement resources are chosen *ex ante* and are inelastic *ex post*. Multiple equilibria arise naturally. We identify a new answer to the old question of why non-maximal penalties are used to punish moderate actions: "marginal" penalties are much more attractive in the Pareto inferior crime wave equilibrium. Specifically, although marginal penalties have both costs and benefits, the net benefit is strictly positive in the crime wave equilibrium. In contrast, marginal penalties frequently have a net cost in the non-crime wave equilibrium. We also show that increasing enforcement resources may worsen crime.

An important constraint faced by tax inspectors, regulators, police forces, etc. is that investigation resources are to a large extent fixed in the short-run. Human employees constitute the key resource in these organizations, and cannot be increased at short notice. Because of this, even if, for example, a tax inspector is in principle committed to auditing all returns claiming deductions above \$100,000, if an unusually large number of returns fall into this category he is unable to audit them all. For the most part, extant formal models of enforcement have ignored this constraint.[1]

In this paper we explicitly account for the *ex post* inelasticity of enforcement resources, and revisit one of the oldest and most enduring questions in law and economics, namely the optimality, or otherwise, of punishing proscribed behavior by using the severest sanction available — that is maximal penalties versus "marginal deterrence." Our main argument is that the *ex post* inelasticity of enforcement resources greatly strengthens the case for marginal deterrence.

As articulated by Becker (1968), there is a simple and compelling argument in favor of maximal penalties: given a non-maximal penalty, one can always achieve the same expected penalty by simultaneously reducing resources spent on detection

[1]The typical assumption is that an enforcement authority can *ex ante* commit to investigate any number of individuals whose signal *ex post* fits the criteria for investigation (see, for example, Mookherjee and Png 1992, 1994, DeMarzo, Fishman and Hagerty 1998). Bar-Gill and Harel (2001) discuss a number of reasons for why the probability of investigation may vary *ex post* with the number of offenders. A contemporaneous paper by Bassetto and Phelan (2008) analyzes a model in which the total number of *ex post* investigations is constrained. The most important difference between our paper and Bassetto and Phelan is that, as we discuss below, we are primarily interested in the relative penalties imposed for moderate and severe offenses (i.e., marginal deterrence). Bassetto and Phelan study a model in which there is just one possible offense, namely under-reporting income (since their model has two income levels, there is no choice of how much to under-report by). As such, the question of whether and how to differentially punish different offenses does not arise in their model.

and increasing the penalty. This argument is troubling because, of course, almost all real-world penalty schedules instead mandate different penalties for different actions. Arguably the leading counterargument to the optimality of maximal penalties relies on the need for marginal deterrence: in terms of a commonly given example, a penalty schedule needs to ensure not just that a potential offender prefers no crime to armed robbery, but also that he prefers armed robbery to murder.[2]

Although persuasive in many respects, the proposition that the need for marginal deterrence implies the optimality of non-maximal penalties is subject to at least two important caveats. First, as observed by Mookherjee and Png (1992), Shavell (1992), and Wilde (1992), it is often possible to monitor different actions at different rates (with the assumption that there is no limitation on the number of *ex post* investigations that can occur). Mookherjee and Png are the most specific in this regard, and point out that if the enforcement authority is able to observe a signal that is correlated to the action selected by the potential offender, it can vary its monitoring effort according to the signal observed. As such, the standard argument for maximal penalties still applies: the maximal penalty should be mandated for each action, with marginal deterrence across actions provided by varying the monitoring intensity.

Second, although marginal deterrence has the potential benefit of reducing the crimes of the worst offenders, it also has the cost of increasing the crimes of other individuals. In terms of the example above, reducing the penalty for armed robbery increases the attractiveness of this crime. This drawback of marginal deterrence is noted by Wilde (1992, page 334), and exists in Mookherjee and Png (1994). In their analysis, all crimes are monitored at the same rate regardless of the seriousness of the crime. This implies that marginal deterrence can only be accomplished through the

[2] While this argument is widely associated with Stigler (1970), it has considerably older antecedents — see Shavell's (1992) discussion.

penalty. They show that at the optimal penalty schedule, there are typically agents who cause a strictly positive amount of harm who would instead be fully deterred under the first best.[3] The harmful acts committed by these agents constitute the cost of marginal deterrence - were marginal deterrence not a concern, the expected penalty for these acts would be increased.

In this paper we show that in a model with limited *ex post* investigation resources, both issues can be simply and simultaneously addressed. First, this assumption immediately implies that an enforcement authority cannot reduce its expenditure by implementing marginal deterrence via varying monitoring intensities. Since monitoring resources must be decided ahead of time, there is nothing to be gained by reducing monitoring intensity *ex post*.

Second, as many authors have observed, the assumption that enforcement resources cannot be varied *ex post* leads to multiple equilibria in crime levels for standard congestion reasons.[4] When other individuals commit more crime, enforcement resources are stretched thin, and so the expected cost of crime falls.[5] For our purposes, the importance of multiple equilibria in crime levels is that they raise the question of which equilibrium the enforcement authority cares about. Our main result is that the balance of the costs and benefits of marginal deterrence differs dramatically across equilibria. Specifically, we show that when the high crime equilibrium is played — a situation we refer to as a *crime wave* — then there is always a level of marginal deterrence at which the benefits outweigh the costs. In contrast, when the low crime equilibrium is played, there are many parameter values for which the costs outweigh the benefits, and maximal penalties are optimal.

[3] Specifically, in Mookherjee and Png (1994) the agent types in the lower interval of panel b of Figures 1, 3 and 4 choose harmful actions which could be deterred at low cost.

[4] See, e.g., Schrag and Scotchmer (1997), Tabarrok (1997), Fender (1999), and Jost (2001).

[5] That is, the crime decisions of different individuals are strategic complements when enforcement resources cannot be changed *ex post*.

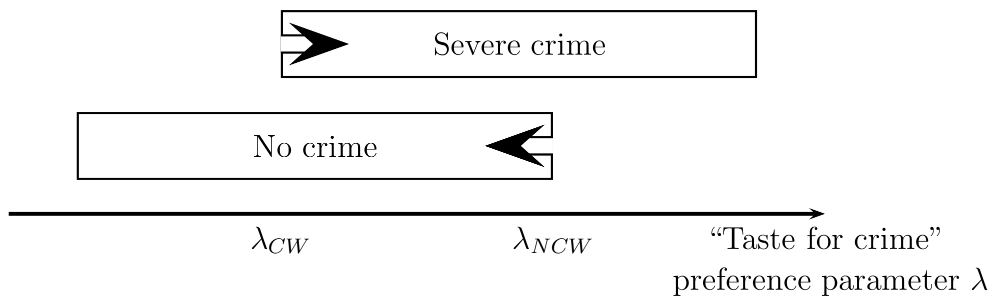

Figure 1: The diagram shows equilibrium crime levels when maximal penalties are used. No crime is an equilibrium for all sufficiently small realizations of the preference parameter λ. Severe crime is an equilibrium for all sufficiently large realizations. For moderate realizations of λ there is both a no crime equilibrium and a severe crime equilibrium. The arrows show the effect of introducing marginal deterrence.

The intuition for this result is easiest to give using the simple diagram of Figure 1. In our model, as in the papers referenced above, individuals are heterogeneous in their taste for crime, which we denote by λ. For the reasons discussed, for a given penalty schedule there are generally multiple equilibria in crime levels. More specifically, when extremal penalties are used both no crime and severe crime are equilibria for agents with a moderate taste for crime: in Figure 1 this is the range of taste parameters λ from λ_{CW} to λ_{NCW}.

Starting from the extreme of maximal penalties for moderate crime, the introduction of non-maximal penalties has two effects. On the one hand, severe crime is less likely to be an equilibrium outcome, since moderate crime is now a more attractive alternative. This is the familiar benefit of marginal deterrence noted by Stigler (1970). In terms of Figure 1, the value λ_{CW} increases. On the other hand, since moderate crime is now more attractive, no crime becomes harder to support as an equilibrium. This is the cost of marginal deterrence. In terms of Figure 1, the value λ_{NCW} decreases.

If one assumes that crime waves occur (that is, the highest crime equilibrium is played) then the equilibrium outcomes under maximal penalties are no crime when the taste parameter is less than λ_{CW}, and severe crime otherwise. So in this case, the benefits of marginal deterrence outweigh the costs: the relevant "boundary" in Figure 1 is the severe crime boundary λ_{CW}, which shifts right, while because of the equilibrium selection assumption the leftwards shift of the no crime boundary λ_{NCW} has no effect.

Conversely, if one instead assumes that the lowest crime equilibrium is played then the equilibrium outcomes under maximal penalties are no crime when the taste parameter is less than λ_{NCW}, and severe crime otherwise. In this case the costs of marginal deterrence outweigh the benefits, by a parallel argument. The relevant boundary in Figure 1 is now the no crime boundary λ_{NCW}, and this shifts leftwards.

The remainder of the paper is as follows. Section 1 presents the model and some preliminary analysis. Section 2 characterizes the benchmark outcomes when maximal penalties are used. Section 3 establishes our main result by formalizing the intuitive discussion above, and by dealing with the complication that marginal penalties introduce additional equilibria as well as destroying existing equilibria. Section 4 considers in more detail the optimal choice of the penalty for moderate crime, and presents several comparative static results. Section 5 explores the consequences of increasing per capita enforcement resources. Finally, Section 6 concludes.

1 Model and preliminary results

There are two agents, labelled i and j. (All our main results hold for $N \geq 2$ agents — see the discussion on page 21.) Each agent chooses between three possible action levels: $a \in \{0, a_M, 1\}$, where $a_M \in (0,1)$. The social costs of these actions are 0, C_M and C_1 respectively, where $C_1 > C_M > 0$. We will often refer to actions $a = 0, a_M, 1$

respectively as *no crime*, *moderate crime* and *severe crime*.

An enforcement authority oversees the two agents with the aim of minimizing the social cost of their actions. However, the enforcement authority does not observe the actions of agents i and j directly. Instead, it observes only noisy signals of these actions, $a^i + \frac{\varepsilon^i}{h}$ and $a^j + \frac{\varepsilon^j}{h}$, where $h > 0$ is a constant measuring the precision of the signal. The error terms ε^i and ε^j are identically and independently distributed, with distribution and density functions F and f. The signals observed by the enforcement authority satisfy the monotone likelihood ratio property (MLRP); equivalently, the density function f is log-concave.[6]

The enforcement authority can impose penalties on the two agents, but must expend resources in order to do so. A central assumption in our analysis is that the enforcement authority's resources are fixed before agents choose their actions, and cannot be increased *ex post*. This appears to be a reasonable description of most real-world enforcement authorities: human employees typically constitute the key resource, and tax collection agencies, regulatory inspection agencies, police forces etc. cannot increase their staff at short notice. Throughout, we assume that the enforcement authority has resources to penalize just one of the two agents.[7]

Formally, based on the pair of signals $a^i + \frac{\varepsilon^i}{h}$ and $a^j + \frac{\varepsilon^j}{h}$ the enforcement authority chooses whether to investigate agent i or agent j. That is, an investigation policy is

[6] Let $\sigma^i = a^i + \frac{\varepsilon^i}{h}$ be the signal observed by the enforcement authority. MLRP is defined as

$$\frac{\partial}{\partial \sigma^i}\left(\frac{1}{\Pr\left(\sigma^i|a^i\right)}\frac{\partial \Pr\left(\sigma^i|a^i\right)}{\partial a^i}\right) = \frac{\partial}{\partial \sigma^i \partial a^i}\ln \Pr\left(\sigma^i|a^i\right) > 0.$$

Since $\Pr\left(\sigma^i|a^i\right) = f\left(\frac{\sigma^i - a^i}{h}\right)$, this condition is equivalent to log-concavity of the density function f. Many common distributions, including the normal distribution, satisfy log-concavity.

[7] Of course, if society cares enough about preventing actions $a = a_M, 1$ it will endow the enforcement authority with resources to penalize both agents. In this case our environment corresponds to the single-agent problem analyzed by previous authors.

a mapping $\mu : \Re^2 \rightarrow \{i,j\}$.[8] We assume that the investigation policy is anonymous, in the sense that it is independent of the identity of the agent (if $\mu(x,x') = i$ then $\mu(x',x) = j$). Define $p(a^i, a^j)$ as the probability that agent i is investigated given actions a^i and a^j, i.e., $p(a^i, a^j) = \Pr\left(\mu\left(a^i + \frac{\varepsilon^i}{h}, a^j + \frac{\varepsilon^j}{h}\right) = i\right)$. For use throughout, observe that anonymity of the investigation policy implies that, for any pair of actions a^i and a^j,

$$p\left(a^i, a^j\right) = 1 - p\left(a^j, a^i\right). \tag{1}$$

Moreover, (1) implies that whenever both agents take the same action, each is investigated with probability 1/2, i.e., $p(a,a) = 1/2$ for any action a.

Investigating an agent allows the enforcement authority to perfectly observe the action chosen by that agent. We assume that the penalty technology is such that only an agent who has been investigated can be penalized. The maximum feasible penalty is S. A general penalty specification is a triple $(s_0, s_M, s_1) \in [0,S]^3$. For simplicity, we assume that penalties impose no social cost (for example, they are wealth transfers).

The enforcement authority does not know how much agents benefit from the socially costly actions $a = a_M, 1$. Each agent's payoff to action a is given by λa, where λ is unobserved by the enforcement authority and is drawn from an atomless distribution with support $\left[0, \bar{\lambda}\right]$. For simplicity we assume that both agents share the same taste parameter λ. This assumption allows us to work with symmetric equilibria

[8] Since the enforcement authority's resources are fixed in advance, it is natural to assume that it will always use these resources *ex post* by investigating one agent. If the enforcement authority derives any payoff from successfully penalizing an agent, then it will always choose to investigate one agent *ex post.* Qualitatively, we do not believe our results would change if instead the enforcement authority sometimes refrained from investigating anyone. On a technical level, the main issue is that our proof that there is always a symmetric pure-strategy equilibrium (Lemma 2) relies on the assumption that the enforcement authority always investigates one of the two agents.

throughout, and our results would be qualitatively unchanged by the introduction of a small amount of preference heterogeneity between the two agents. Economically, the assumption can be thought of as reflecting the widely held notion that cultural norms against wrongdoing vary geographically, and may also change over time. Alternatively, the assumption can be motivated by variations in community enforced social sanctions; or by variations in the marginal utility of income.

We focus throughout on pure-strategy symmetric equilibria.[9] No crime ($a_i = a_j = 0$) is an equilibrium if and only if

$$-s_0 p\,(0,0) \geq \lambda a_M - s_M p\,(a_M, 0) \tag{IC0-M}$$

$$-s_0 p\,(0,0) \geq \lambda - s_1 p\,(1,0) \tag{IC0-1}$$

Moderate crime ($a_i = a_j = a_M$) is an equilibrium if and only if

$$\lambda a_M - s_M p\,(a_M, a_M) \geq -s_0 p\,(0, a_M) \tag{ICM-0}$$

$$\lambda a_M - s_M p\,(a_M, a_M) \geq \lambda - s_1 p\,(1, a_M) \tag{ICM-1}$$

Severe crime ($a_i = a_j = 1$) is an equilibrium if and only if

$$\lambda - s_1 p\,(1,1) \geq -s_0 p\,(0,1) \tag{IC1-0}$$

$$\lambda - s_1 p\,(1,1) \geq \lambda a_M - s_M p\,(a_M, 1)\,. \tag{IC1-M}$$

We assume that the social cost of actions $a = a_M, 1$ exceeds the direct private benefit, i.e., $C_1 - \bar{\lambda} > C_M - \bar{\lambda} a_M > 0$. So taking the enforcement authority's *ex post* investigative capacity as fixed at one investigation,[10] its objective is to choose an

[9] Lemma A-1, stated and proved in the appendix, shows that the probability of the taste parameter λ being such that an asymmetric pure-strategy equilibrium exists is zero. Lemma 2 below shows that at least one pure-strategy symmetric equilibrium always exists. As standard in games with strategic complementarities, whenever two distinct pure-strategy symmetric equilibria exist, there is also a mixed strategy symmetric equilibrium. However, this mixed strategy equilibrium is unstable.

[10] It would be straightforward to expand the problem to one in which the investigative capacity is also determined optimally.

investigation policy μ and a penalty specification (s_0, s_M, s_1) so as to minimize the expected equilibrium crime level.

1.1 Preliminary results

The remainder of this section develops several preliminary results that allow us to simplify the equilibrium characterization.

Consider a pair of offsetting incentive conditions — (IC0-M) and (ICM-0), for example. Condition (IC0-M) says that agent i prefers no crime to moderate crime, conditional on agent j choosing no crime. Condition (ICM-0) says that agent i prefers moderate crime to no crime, conditional on agent j choosing the moderate crime. Since the probability that agent i is investigated depends on agent j's action, these two conditions are not symmetric. This raises the possibility that both (IC0-M) and (ICM-0) may simultaneously fail, a complication that would potentially lead to the non-existence of a pure-strategy symmetric equilibrium. However, (1) is enough to ensure that this possibility does *not* arise:

Lemma 1 *Suppose that both the penalty schedule and investigation policy are monotone, i.e., $s_1 \geq s_M \geq s_0$ and p is increasing in its first argument. Then at least one of (IC0-M) and (ICM-0) holds; at least one of (IC0-1) and (IC1-0) holds; and at least one of (ICM-1) and (IC1-M) holds.*

The existence of at least one pure-strategy symmetric equilibrium follows easily from Lemma 1:

Lemma 2 *Suppose that both the penalty schedule and investigation policy are monotone, i.e., $s_1 \geq s_M \geq s_0$ and p is increasing in its first argument. Then there exists at least one pure-strategy symmetric equilibrium.*

Having established that a pure-strategy symmetric equilibrium exists for any realization of λ, we turn now to simplifying the problem. Holding the investigation policy μ fixed, setting $s_0 = 0$ always increases the range of λ realizations for which no crime is an equilibrium; and setting $s_1 = S$ always increases the range of λ realizations for which severe crime is *not* an equilibrium. In a similar fashion, the investigation policy "investigate the agent with the higher signal" is the investigation policy that minimizes the range of the severe crime equilibrium and maximizes the range of the no crime equilibrium. Formally:

Lemma 3 *The probability that no crime is an equilibrium is maximized, and the probability that severe crime is an equilibrium is minimized, by choosing $s_0 = 0$, $s_1 = S$, and using the investigation policy "investigate the agent with the higher signal."*

Given Lemma 3, for the remainder of the paper we assume $s_0 = 0$, $s_1 = S$, and that the "investigate the agent with the higher signal" policy is used. Define $q\left(a^i - a^j\right) = p\left(a^i, a^j\right)$, since under this investigation policy only the *difference* in the actions of the two agents affects the investigation probability.

Lemma 4 *The investigation probability function $q : [-1, 1] \to [0, 1]$ is:*
(I) Symmetric about 0: $q(a) - q(0) = q(0) - q(-a)$.
(II) Increasing (decreasing) in signal precision for positive (negative) values.
(III) Concave (convex) over positive (negative) values.

For use in the remainder of the paper, the equilibrium conditions simplify to: no crime ($a = 0$) is an equilibrium if and only if

$$0 \geq \lambda a_M - s_M q\left(a_M\right) \tag{IC0-M}$$

$$0 \geq \lambda - S q\left(1\right); \tag{IC0-1}$$

moderate crime ($a = a_M$) is an equilibrium if and only if

$$\lambda a_M - s_M q(0) \geq 0 \tag{ICM-0}$$

$$\lambda a_M - s_M q(0) \geq \lambda - Sq(1 - a_M); \tag{ICM-1}$$

and severe crime ($a = 1$) is an equilibrium if and only if

$$\lambda - Sq(0) \geq 0 \tag{IC1-0}$$

$$\lambda - Sq(0) \geq \lambda a_M - s_M q(a_M - 1). \tag{IC1-M}$$

1.2 Extremal vs marginal penalties

Observe that in light of Lemma 3, the enforcement authority's problem has reduced to choosing the penalty s_M to impose on an agent who chooses action a_M. This is the focus of the paper. We refer to the choice $s_M = s_1 = S$ as an *extremal penalty*, and to any choice $s_M < s_1 = S$ as a *marginal penalty.*[11]

1.3 Assumptions

We assume that the supremum of the support of the taste parameter λ satisfies $\bar{\lambda} > S\frac{q(1-a_M)}{1-a_M}$. This guarantees that severe crime is the only equilibrium when the taste parameter is sufficiently high, regardless of the choice of penalty s_M.[12] Moreover, to make our analysis as transparent as possible, we assume that signal precision is sufficiently poor such that

$$q(0) > q'(0). \tag{2}$$

[11]Clearly one could also term the choice $s_M = 0$ as an extremal penalty. This is purely an issue of terminology.

[12]It is immediate that (ICM-1) does not hold at $\bar{\lambda}$. By the concavity of q over positive values, (IC0-1) does not hold either: $q(a) \geq aq(1) + (1-a)q(0) > aq(1)$ for any $a \in (0,1)$, and so $S\frac{q(1-a_M)}{1-a_M} > Sq(1)$. Hence $a = 1$ is the only equilibrium at $\lambda = \bar{\lambda}$.

This assumption ensures that moderate crime is never an equilibrium under extremal penalties. However, we stress that it is not essential for our main results, which hold under much weaker conditions. In particular, our results hold whenever a_M is not too different from 1/2, regardless of whether or not (2) is satisfied (details are available on the authors' webpages).

2 A benchmark: equilibrium outcomes under extremal penalties

Our main object of enquiry is when extremal penalties are, and are not, optimal. Accordingly, we begin our analysis by characterizing the equilibrium outcomes when extremal penalties are used. No crime ($a = 0$) is an equilibrium when (IC0-M) and (IC0-1) hold, i.e., if $\lambda \leq \frac{Sq(a_M)}{a_M}$ and $\lambda \leq Sq(1)$. Since q is concave over positive values,

$$q(a_M) \geq a_M q(1) + (1 - a_M) q(0) > a_M q(1). \tag{3}$$

As such, for extremal penalties (IC0-M) holds whenever (IC0-1) does, and so no crime is an equilibrium if and only if $\lambda \leq Sq(1)$.

At the other extreme, severe crime is an equilibrium when (IC1-0) and (IC1-M) hold, i.e., $\lambda \geq Sq(0)$ and $\lambda \geq S\frac{q(0)-q(a_M-1)}{1-a_M}$. Under assumption (2), for extremal penalties (IC1-M) holds whenever (IC1-0) does.[13]

Assumption (2) says that the investigation probability does not change very quickly as an agent changes his action choice. This means that any agent with preferences for crime that are strong enough for him to prefer moderate crime to no

[13] We must show that $q(0) \geq \frac{q(0)-q(a_M-1)}{1-a_M}$, or equivalently, $q(a_M - 1) - a_M q(0) \geq 0$. Convexity of q over negative values implies $q(0) - q'(0)(1 - a_M) \leq q(a_M - 1)$. From assumption (2) it follows that $q(0) a_M = q(0) - q(0)(1 - a_M) < q(a_M - 1)$.

crime would derive an even higher payoff from severe crime, given that under extremal penalties s_M and s_1 coincide. As such, under extremal penalties moderate crime is never an equilibrium. Formally, moderate crime ($a = a_M$) is an equilibrium if (ICM-0) and (ICM-1) hold, i.e., if $\lambda \geq \frac{Sq(0)}{a_M}$ and $\lambda \leq S\frac{q(1-a_M)-q(0)}{1-a_M}$, and assumption (2) implies that both inequalities cannot be satisfied at once.[14]

Corresponding to Figure 1, we define the critical values of the taste parameter identified above by $\lambda_{CW} = Sq(0)$ and $\lambda_{NCW} = Sq(1)$.

Proposition 1 *When extremal penalties are used, no crime ($a = 0$) is an equilibrium if $\lambda \leq \lambda_{NCW}$; severe crime ($a = 1$) is an equilibrium if $\lambda \geq \lambda_{CW}$; moderate crime ($a = a_M$) is never an equilibrium.*

As an immediate implication:

Corollary 1 *Both no crime and severe crime are equilibria when $\lambda \in [\lambda_{CW}, \lambda_{NCW}]$.*

The source of multiple equilibria is, or course, the enforcement authority's limited resources *ex post*. This means that the action choices of the two agents are strategic complements: when agent j chooses a more socially costly action, it reduces the investigation probability faced by agent i, and so increases the utility gain to agent i of choosing a more socially costly action.

Given the existence of multiple equilibria, an equilibrium selection rule is needed to select the optimal choice of the penalty s_M. Our main results characterize how the equilibrium selection rule affects the optimal choice. We consider two selection rules:

[14] Specifically, assumption (2) implies $\frac{q(0)}{a_M} > \frac{q(1-a_M)-q(0)}{1-a_M}$, or equivalently, $a_M q(1-a_M) < q(0)$, as follows. Concavity of q over positive values and assumption (2) together imply

$$q(1-a_M) \leq q(0) + (1-a_M)\, q'(0) < (2-a_M)\, q(0).$$

Since $a_M(2-a_M) \leq 1$, the result follows.

(CW), a *crime wave* selection rule, in which whenever multiple equilibria exist we assume the one with the highest crime level is played; and (NCW), a *no crime wave selection rule*, in which whenever multiple equilibria exist we assume the one with the lowest crime level is played. Note that whenever multiple pure-strategy symmetric equilibria exist, agents i and j strictly prefer the higher crime equilibrium.[15] Consequently, the crime wave selection rule corresponds to the commonly used device of selecting the Pareto dominant (from the perspective of players i and j) equilibrium.

3 The costs and benefits of marginal penalties

In this section we present our main result, namely that the introduction of marginal penalties always lowers crime under the crime wave selection rule, but for many parameter values raises crime under the no crime wave selection rule. To establish this result as directly as possible, we postpone until the next section a full characterization of equilibrium outcomes under the crime wave selection rule.[16]

We consider first the crime wave selection rule. By Proposition 1, under extremal penalties the equilibrium is severe crime ($a = 1$) for taste realizations $\lambda \geq \lambda_{CW}$, and no crime otherwise. The potential benefit of introducing marginal penalties — i.e., lowering s_M below S — is that doing so destroys the severe crime equilibrium for some $\lambda \geq \lambda_{CW}$, by making the deviation to moderate crime attractive. This benefit is first experienced at the taste realization λ_{CW} — which, by definition, is such that that agents have a zero payoff in the severe crime equilibrium. Hence the benefit of marginal penalties is first experienced when s_M is lowered beyond the point at which an agent with taste λ_{CW} gets a zero payoff from deviating from a severe crime

[15] The only exception is the zero-probability realization of the taste parameter λ in which, in equilibrium, the agents are indifferent between a positive action level, and action $a = 0$.

[16] Related, nowhere in our analysis do we need a full characterization of the set of equilibria.

equilibrium to commit moderate crime. The agent's probability of being investigated in this case is $q(a_M - 1)$, and so the benefit of marginal penalties is first experienced at the marginal penalty s_M defined by $\lambda_{CW} a_M - s_M q(a_M - 1) = 0$.

Conversely, the potential cost of marginal penalties is that they introduce a moderate crime equilibrium for some $\lambda < \lambda_{CW}$. Since a zero payoff is always attainable by taking action $a = 0$, a necessary condition for such an equilibrium to exist is that an agent derive a non-negative payoff from committing moderate crime, given that the other agent also commits moderate crime, and so each faces an investigation probability of $q(0)$. Consequently, it is only possible for marginal penalties to have a cost once s_M is lowered beyond the level defined by $\lambda_{CW} a_M - s_M q(0) = 0$.

Since $q(a_M - 1) < q(0)$, we have established our first main result:

Proposition 2 *Under the crime wave selection rule, as marginal penalties are adopted by reducing s_M away from S, equilibrium crime levels are at first unaffected and then reduced.*

The intuition is straightforward. Both the cost and benefit under the crime wave selection rule occur when the type λ_{CW} first gets a zero payoff from moderate crime. The difference is that the benefit occurs when this payoff arises from a downwards deviation from a severe crime equilibrium, with an associated investigation probability below 1/2, while the cost occurs when this payoff arises in equilibrium, with an associated investigation probability of 1/2.

It is important to note that Proposition 2 says only that when marginal penalties first have an impact on crime levels, that impact is positive. As will be clear from Proposition 5 below, it does not follow that the penalty s_M should be reduced all the way to 0.

We turn now to the opposite case of Proposition 2 in which the no crime wave selection rule is used. As we will see, for many parameter values the introduction of marginal penalties actually increases the crime level in this case.

By Proposition 1, under extremal penalties the equilibrium is severe crime for taste realizations $\lambda > \lambda_{NCW}$, and no crime otherwise. The potential benefit of introducing marginal penalties is now that doing so may introduce a moderate crime equilibrium for some $\lambda > \lambda_{NCW}$. This occurs when (ICM-0) and (ICM-1) are both satisfied. Note that for $\lambda > \lambda_{NCW}$, the payoff to deviating to severe crime from a moderate crime equilibrium, i.e., $\lambda - Sq(1 - a_M)$, is positive. Consequently, a moderate crime equilibrium is introduced whenever (ICM-1) is satisfied. As marginal penalties are adopted, this first occurs when s_M is lowered beyond the level defined by $\lambda_{NCW} a_M - s_M q(0) = \lambda_{NCW} - Sq(1 - a_M)$ (i.e., (ICM-1) holds with equality).

Conversely, under the no crime wave selection rule the potential cost of marginal penalties is that they destroy the no crime equilibrium for some $\lambda \leq \lambda_{NCW}$. This occurs when an agent can obtain a positive payoff by deviating from a no crime equilibrium to moderate crime, and facing an investigation probability of $q(a_M)$. Consequently, the cost of marginal penalties is experienced once s_M is lowered beyond the level defined by $\lambda_{NCW} a_M - s_M q(a_M) = 0$.

In contrast to the case under the crime wave selection rule, it is no longer clear whether the benefits of marginal penalties are reached before or after the costs. The reason is that the costs now arise when the type λ_{NCW} gets a zero payoff by deviating to moderate crime from a no crime equilibrium; while the benefits arise when the same type gets a strictly positive payoff from moderate crime in equilibrium. Because both the payoffs and investigation probabilities differ across these two cases, the comparison depends on the specific shape of the investigation probability function q. However, one case in which the comparison is clear is when the deviation to moderate crime from a no crime equilibrium does not affect the investigation probability very much. This is the case when the precision h of the enforcement authority's signal is low. Formally:

Proposition 3 *Under the no crime wave selection rule, as marginal penalties are adopted by reducing s_M away from S, equilibrium crime levels are at first unaffected and then increased if*

$$\frac{\lambda_{NCW} a_M}{q(a_m)} > \frac{Sq(1 - a_M) - \lambda_{NCW}(1 - a_M)}{q(0)}, \tag{4}$$

but decreased otherwise. (Recall $\lambda_{NCW} = Sq(1)$.) Consequently, extremal penalties are locally optimal when (4) holds. Moreover, condition (4) holds whenever the precision h of the enforcement authority's signal is sufficiently low.

(Remark: It is readily verified that (4) also holds for all a_M sufficiently small.)

Loosely speaking, the difference in the effect of marginal penalties across equilibrium selection rules reflects the fact that, in our model, all else being equal it is easier to destroy an existing equilibrium than to create a new equilibrium. Under the crime wave selection rule, the benefits stem from destroying an existing crime wave equilibrium; while under the no crime wave selection rule, the benefits stem from creating a new moderate crime equilibrium.

4 Comparative statics under crime waves

In the previous section we showed that under the crime wave selection rule, some use of marginal penalties is always optimal. As we showed, this contrasts sharply with the situation under the no crime wave selection rule. There remains the question of what level of marginal penalties is optimal under the crime wave selection rule. This is the subject of the current section.

We begin by comprehensively characterizing the equilibrium outcome under the crime wave selection rule. It is possible to do this using just three of the six equilibrium conditions, namely (IC1-0), (IC1-M) and (ICM-0):

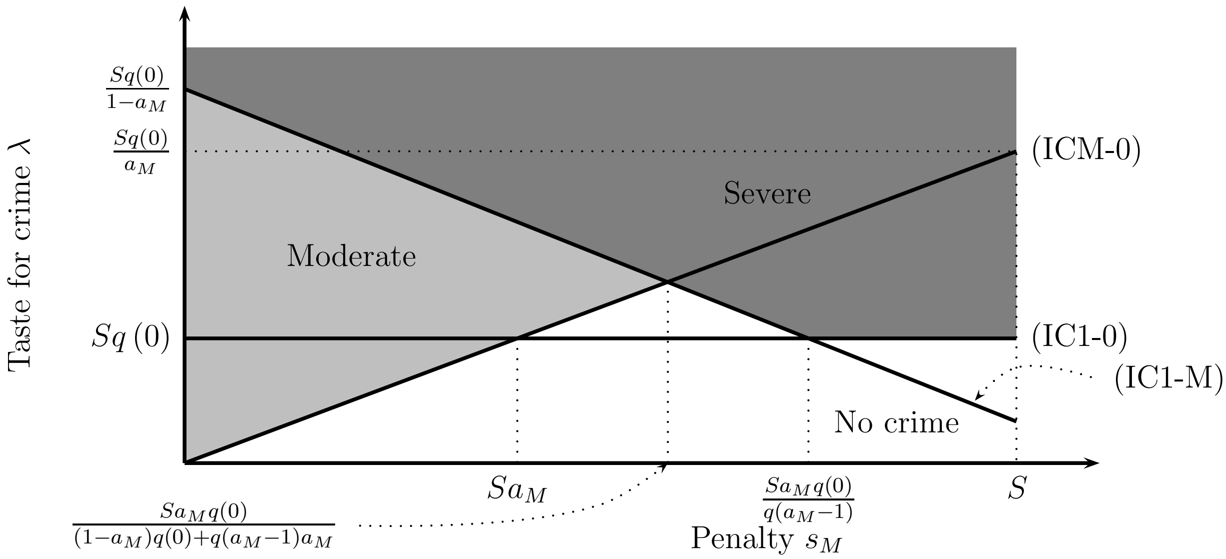

Figure 2: Equilibrium outcomes under the crime wave selection rule. The graph shows the highest crime equilibrium for each moderate crime penalty s_M and taste-for-crime parameter λ.

Proposition 4 *Under the crime wave selection rule, the equilibrium is severe crime ($a = 1$) whenever (IC1-0) and (IC1-M) both hold; the equilibrium is moderate crime ($a = a_M$) whenever (IC1-M) fails but (ICM-0) holds; and the equilibrium is no crime ($a = 0$) otherwise.*

Using Proposition 4, we can graphically represent the equilibrium outcomes under the crime wave selection rule. To this end, it is useful to rewrite the incentive constraints (IC1-0), (IC1-M) and (ICM-0) as bounds on λ:

$$\lambda \geq Sq(0) \tag{IC1-0}$$

$$\lambda \geq \frac{Sq(0) - s_M q(a_M - 1)}{1 - a_M} \tag{IC1-M}$$

$$\lambda \geq \frac{s_M q(0)}{a_M}. \tag{ICM-0}$$

Figure 2 plots these three constraints as functions of the taste parameter λ and the extent of marginal penalties, s_M. The lines (IC1-0) and (IC1-M) intersect at $s_M = \frac{Sa_M q(0)}{q(a_M - 1)}$, while the lines (IC1-0) and (ICM-0) intersect at $s_M = Sa_M$. Since $q(0) > q(a_M - 1)$ the (IC1-0) line lies below the intersection of (IC1-M) and (ICM-0). Economically, and as earlier established in Proposition 2, under the crime wave selection rule the introduction of marginal penalties always reduces the crime level.

Strikingly, the adoption of marginal penalties reduces crime levels by switching the equilibrium outcome from severe crime to no crime for some realizations of the taste parameter λ. That is, marginal penalties are useful even though no agent commits the moderate crime a_M in equilibrium. This is only possible because of the interdependency between different agents' action choices, which in turn stems from the *ex post* inelasticity of enforcement resources. Given this interdependency, it is possible for marginal penalties to eliminate the severe crime equilibrium without introducing moderate crime as an equilibrium. In contrast, absent *ex post* inelasticity of enforcement resources each agent's optimization problem would be independent of other agents' decisions, and marginal penalties would only affect outcomes if they resulted in moderate crime in equilibrium.

Let g denote the probability density of the taste parameter λ. Using Figure 2 one can easily calculate social welfare for each penalty level s_M, as follows. First, note that the penalty for moderate crime should be set below the intersection point of (ICM-0) and (IC1-M), i.e., below $\frac{Sa_M q(0)}{(1-a_M)q(0)+q(a_M-1)a_M}$, since above this point marginal penalities generate only benefits, and no cost. For penalties s_M below this level, social welfare under crime waves is given by

$$SW \equiv \int_{s_M q(0)/a_M}^{\frac{Sq(0)-s_M q(a_M-1)}{1-a_M}} (\lambda a_M - C_M)\, g(\lambda)\, d\lambda + \int_{\frac{Sq(0)-s_M q(a_M-1)}{1-a_M}}^{\bar{\lambda}} (\lambda - C_1)\, g(\lambda)\, d\lambda,$$

where the two terms represent the expected net social cost of moderate and severe crime respectively.[17] From this expression it is straightforward to show:

[17]One can also calculate social welfare without including the private benefits λa_M and λ from

Proposition 5 *Under the crime wave selection rule, the optimal penalty for moderate crime lies in the interval* $\left[0, \frac{S a_M q(0)}{(1-a_M)q(0)+q(a_M-1)a_M}\right]$, *and is decreasing in* C_1 *and increasing in* C_M.

The intuition behind Proposition 5 is clear. When the social cost of severe crime is high, society should do everything it can to curtail severe crime. Setting s_M low achieves this objective because it induces agents to switch from the severe crime to the moderate crime equilibrium. Of course, setting s_M low also engenders more moderate crime, but when C_M is low relative to C_1 this is a price worth paying.

The second comparative static we consider is the effect of signal precision h on the optimal choice of s_M. Signal precision affects social welfare via the investigation probability $q(a_M - 1)$, which is decreasing in precision h: when signal precision is high, the investigation probability faced by an agent who takes action a_M while the other takes action $a = 1$ is low. The investigation probability $q(a_M - 1)$ measures the strength of incentives provided by the investigation policy for agents to abstain from severe crime: if agent j chooses severe crime, agent i is punished with probability $q(0) = 1/2$ if he also chooses severe crime, but with probability $q(a_M - 1)$ if instead he chooses moderate crime. As such, when $q(a_M - 1)$ is low, the investigation policy alone delivers considerable "marginality" in expected penalties even when s_M is close to the maximum penalty S. If instead $q(a_M - 1)$ is high, the investigation policy delivers little marginality, and it is worthwhile setting s_M much lower than S in order to increase the difference in expected penalties for severe and moderate crime.

The effect of $q(a_M - 1)$, and hence of signal precision h, on social welfare depends in part on the density g of the taste parameter λ. To abstract from these effects, we consider the special case in which λ is uniformally distributed:[18]

actions $a = a_M, 1$. Doing so has no qualitative effect on our results.

[18] More generally, Proposition 6 would hold whenever the derivative of the density function g' is sufficiently small.

Proposition 6 *Suppose the taste parameter is distributed uniformally. Under the crime wave selection rule, the optimal penalty for moderate crime is decreasing in the investigation probability* $q(a_M - 1)$ *and so is increasing in signal precision* h.

5 Varying per capita enforcement resources

Thus far we have assumed that the enforcement agency has the resources to conduct one investigation for each two potential offenders. We have focused on this case purely for expositional convenience: our main results apply equally when instead the enforcement authority oversees $N \geq 2$ agents, and has the resources to investigate and penalize just one of them. Details are available from the authors' webpages. In this section, we analyze the effect of changes in N, that is, of changes in per capita enforcement resources.

With N agents, the probability of investigation faced by an agent choosing action $\tilde{a}$ while the other $N-1$ agents choose action a is given by $q_N(\tilde{a} - a)$, for some increasing function q. Clearly as the number of agents increases, the probability of investigation $q_N(x)$ decreases for any value of x. Moreover, it is possible to show that the ratio $q_N(x)/q_N(0)$ is increasing (respectively, decreasing) in the number of agents N if $x > 0$ (respectively, $x < 0$). That is, as N increases the probability of investigation decreases faster for an agent who commits a lesser crime, holding the actions of other agents fixed.

We can use this result to consider the effect of changes in per capita enforcement resources on crime levels. Suppose that enforcement resources increase, i.e., N is reduced. This generates some unambiguously positive effects. For example, it increases the range of λ realizations for which (IC0-1) and (IC0-M) hold, implying that no crime ($a = 0$) is an equilibrium more often. However, for $s_M > 0$ a decrease in N has a tendency to make (ICM-1) less likely to hold, since as N decreases the ratio $q_N(1-a_M)/q_N(0)$ decreases. Economically, a decrease in N increases $q_N(0)$ by a larger amount than $q_N(1-a_M)$, and this makes it harder to prevent an agent deviating from moderate crime to severe crime. In this case, the increase in resources actually has a negative effect on the distribution of crime.

The above discussion gives the main intuition, but is somewhat loose in that (ICM-1) involves the terms $q_N(1-a_M)$ and $q_N(0)$ separately, as opposed to in ratio. The following result gives one reasonably concise set of sufficient conditions for a decrease in N to have the negative effect described:[19]

Proposition 7 *Suppose the no crime wave selection rule is used, and $a_M = 1/2$. Then for any moderate crime penalty $s_M \in (0, S/(1+1/N))$ and any number of agents $N > 2$, there exists a signal precision $\bar{h}$ such that whenever $h \geq \bar{h}$, an increase in per capita enforcement (i.e., a reduction in N) increases the probability of severe crime (though it also increases the probability of no crime).*

In general, an enforcement agency that experiences an increase to its budget can choose to spend these extra funds in one of two ways. On the one hand, the enforcement agency can expand its capacity to investigate and penalize agents (i.e., decrease N). On the other hand, the enforcement agency can seek to increase the precision of its pre-investigation information, that is, of the information upon which it decides whom to investigate (i.e., increase h). Proposition 7 shows that the former use of funds may actually increase crime under the no crime wave selection rule. In contrast, an immediate implication of part (II) of Lemma 4 is that an increase in signal precision reduces crime under the no crime wave selection rule.[20] As such, using the increase in funds to increase signal precision is often the more attractive

[19] We emphasize that the effect described holds much more generally than the sufficient conditions of Proposition 7. Additionally, given that Proposition 7 is stated for the case of signal precision h being sufficiently high, it is worth noting that this result makes no use of assumption (2).

[20] Specifically, an increase in precision increases $q(a_M)$, $q(1-a_M)$ and $q(1)$, while leaving $q(0)$

alternative, since it leads to an unambiguous decrease in crime, while increasing investigation capacity has the potential to actually increase the crime level under some circumstances.

To some extent this prediction is consistent with the recent rise in "community policing." While the term encompasses a variety of distinct ideas, one important element is that police officers should spend more time patrolling streets by foot, and less time in patrol cars and responding to emergency calls.[21] This, it is often argued, will engender much better relations between police officers and the communities they oversee.[22] In terms of our analysis, this aspect of community policing can be thought of as corresponding to an increase in signal accuracy achieved at the cost of a decrease in investigation resources (increase in N). Enabling better communication between a community and its police is akin to increasing the amount of information a police department has on which to base its more formal investigations. However, taking police out of patrol cars reduces their ability to arrive promptly at a crime scene and possibly apprehend a criminal immediately.

unchanged; and this increases the range of parameters for which no crime is an equilibrium, and for which moderate crime is an equilibrium. Given the no crime wave selection rule, it follows that equilibrium crime declines for all realizations of the taste parameter λ.

[21]Perhaps the aspect of community policing to have attracted most comment is the "broken windows" theory, which emphasizes the importance of the eliminating small crimes (or at least their effects) for the control of more serious crimes. Although logically distinct from increasing the number of police officers on foot-patrol, in practice the two ideas are closely related.

[22]For example, in their much-cited article, Wilson and Kelling (1982) emphasize the awkwardness faced by an individual who wants to communicate with a police officer seated in a patrol car.

6 Conclusion

Our analysis identifies a new answer to the old question of why non-maximal penalties are used to punish moderate actions. In environments in which multiple equilibria arise naturally, marginal penalties are much more attractive under the crime wave selection rule, in which the highest crime equilibrium is played. Specifically, although marginal penalties have both costs and benefits (as identified by prior authors — see the introduction), the net benefit is strictly positive under the crime wave selection rule. In contrast, for a wide range of parameter values marginal penalties have a net cost under the no crime wave selection rule, in which the lowest crime equilibrium is played. The economic reason for this difference is that under the crime wave selection rule, the benefits of marginal penalties stem from the destruction of a severe crime equilibrium, while the cost stems from the possible creation of a moderate crime equilibrium (displacing a no crime equilibrium). Destroying an existing equilibrium is easier than creating a new equilibrium, and so the benefits arrive first. Under the no crime wave selection rule, the situation is reversed: the benefits of marginal penalties now stem from the creation of a moderate crime equilibrium (displacing a severe crime equilibrium), while the costs stem from the destruction of a no crime equilibrium. In addition, our analysis has implications for the comparative efficacy of increasing resources devoted to investigative capacity versus pre-investigation information

Bar-Gill, Oren, and Alon Harel. 2001. "Crime rates and expected sanctions: The economics of deterrence revisited." *Journal of Legal Studies*, 30(2): 485–501.

Bassetto, Marco, and Christopher Phelan. 2008. "Tax riots." *Review of Economic Studies*, 75(3): 649–669.

Becker, Gary. 1968. "Crime and punishment: An economic approach." *Journal of Political Economy*, 76(2): 169–217.

DeMarzo, Peter M., Michael J. Fishman, and Kathleen M. Hagerty.

1998. "The optimal enforcement of insider trading regulations." *Journal of Political Economy*, 106(3): 602–632.

Fender, John. 1999. "A general equilibrium model of crime and punishment." *Journal of Behavior and Organization*, 39(4): 437–453.

Jost, Peter J. 2001. "Crime, coordination, and punishment: an equilibrium analysis." *International Review of Law and Economics*, 21(1): 23–46.

Mookherjee, Dilip and Ivan Png. 1992. "Monitoring vis-a-vis investigation in enforcement of law." *American Economic Review*, 82(3): 556–565.

Mookherjee, Dilip and Ivan Png. 1994. "Marginal deterrence in enforcement of law." *Journal of Political Economy*, 102(5): 1039–1066.

Joel Schrag and Suzanne Scotchmer. 1997. "The self-reinforcing nature of crime." *International Review of Law and Economics*, 17(3): 325–335.

Shavell, Steven. 1992. "A note on marginal deterrence." *International Review of Law and Economics*, 12(3): 345–355.

Stigler, George J. 1970. "The optimum enforcement of laws." *Journal of Political Economy*, 78(3): 526–536.

Tabarrok, Alexander. 1997. "A simple model of crime waves, riots, and revolutions." *Atlantic Economic Journal*, 25(3): 274–88.

Wilde, Louis L. 1992. "Criminal choice, nonmonetary sanctions, and marginal deterrence: A normative analysis." *Internationl Review of Law and Economics*, 12(3): 333–344.

Wilson, James Q., and George L. Kelling. 1982. "Broken windows." *Atlantic Monthly*, March, 29–38.

A Appendix

For use below, observe that (1) implies that p is increasing in its first argument if and only if it is decreasing in its second argument.

Lemma A-1 *Suppose that both the penalty schedule and investigation policy are monotone, i.e., $s_1 \geq s_M \geq s_0$ and p is increasing in its first argument. The probability of the taste parameter λ being such that an asymmetric pure-strategy equilibrium exists is zero.*

Proof of Lemma A-1: We consider an asymmetric equilibrium in which agent i takes action a_M and agent j takes action 1. (The other two possibilities for asymmetric pure-strategy equilibrium can be dealt with in an identical way.) Suppose to the contrary that such an equilibrium exists. Given anonymity of the investigation policy μ, $\lambda a_M - s_M p(a_M, 1) \geq \lambda - s_1 p(1,1)$ and $\lambda - s_1 p(1, a_M) \geq \lambda a_M - s_M p(a_M, a_M)$. Hence

$$s_1 p(1,1) - s_M p(a_M, 1) \geq \lambda (1 - a_M) \geq s_1 p(1, a_M) - s_M p(a_M, a_M). \qquad \text{(A-1)}$$

Since $s_1 \geq s_M$ and, by (1) and monotonicity, $p(1,1) - p(1, a_M) = p(a_M, 1) - p(a_M, a_M) \leq 0$, this is possible only if both inequalities in (A-1) hold with equality. For any given penalty schedule and investigation policy, this can occur at at most a single realization of λ. ∎

Proof of Lemma 1: We will show that at least one of (ICM-1) and (IC1-M) holds — the other two claims follow by parallel arguments. Suppose that (ICM-1) does not hold, i.e.,

$$\lambda (1 - a_M) > s_1 p(1, a_M) - s_M p(a_M, a_M).$$

To show that (IC1-M) holds it suffices to show

$$s_1 p(1, a_M) - s_M p(a_M, a_M) \geq s_1 p(1,1) - s_M p(a_M, 1),$$

or equivalently,

$$s_1 \left(p\left(1,a_M\right)-p\left(1,1\right)\right) \geq s_M \left(p\left(a_M,a_M\right)-p\left(a_M,1\right)\right). \tag{A-2}$$

The result then follows from $s_1 \geq s_M$ and, by (1) and monotonicity, $p\left(1,a_M\right) - p\left(1,1\right) = p\left(a_M,a_M\right) - p\left(a_M,1\right) \geq 0$. ∎

Proof of Lemma 2: Suppose that contrary to the claimed result no pure-strategy symmetric equilibrium exists. Since $a=0$ is not an equilibrium, at least one of (IC0-M) and (IC0-1) must fail to hold.

First, suppose that (IC0-M) fails to hold. From Lemma 1, (ICM-0) holds, and so (ICM-1) fails to hold (else $a = a_M$ is an equilibrium); and so (IC1-M) holds, and (IC1-0) fails to hold (else $a=1$ is an equilibrium). So

$$\begin{aligned} -s_0 p\left(0,1\right) &> \lambda - s_1 p\left(1,1\right) \geq \lambda - s_1 p\left(1,a_M\right) \\ &> \lambda a_M - s_M p\left(a_M,a_M\right) \geq \lambda a_M - s_M p\left(a_M,0\right) > -s_0 p\left(0,0\right), \end{aligned}$$

a contradiction. (The strict inequalities follow from the failure of (IC1-0), (ICM-1) and (IC0-M) respectively; the weak inequalities and contradiction follow from the fact that p is decreasing in its second argument.)

Second, suppose that (IC0-1) fails to hold. Then (IC1-0) holds, and so (IC1-M) must fail to hold (else $a=1$ is an equilibrium). This in turn implies that (ICM-1) holds, and (ICM-0) fails to hold (else $a = a_M$ is an equilibrium). So

$$-s_0 p\left(0,a_M\right) > \lambda a_M - s_M p\left(a_M,a_M\right) \geq \lambda - s_1 p\left(1,a_M\right) \geq \lambda - s_1 p\left(1,0\right) > -s_0 p\left(0,0\right),$$

a contradiction. (The inequalities follow from the failure of (ICM-0), (ICM-1), p decreasing in its second argument, and the failure of (IC0-1).) ∎

Proof of Lemma 3: The statement about the choice of penalties is established in the main text. For the statement regarding the investigation policy, it suffices to show that if any other investigation policy is used then shifting to the "investigate the agent with the higher signal" policy strictly increases $p(a_M,0)$ and $p(1,0)$, and strictly decreases $p(0,a_M)$ and $p(0,1)$. The proof is as follows. For an arbitrary investigation policy μ and action $a>0$,

$$p(a,0)=\int\int\mu\left(a+\frac{\varepsilon^i}{h},\frac{\varepsilon^j}{h}\right)f\left(\varepsilon^i\right)f\left(\varepsilon^j\right)d\varepsilon^i d\varepsilon^j.$$

Changing variables,

$$p(a,0)=\int\int\mu\left(\frac{\varepsilon^i}{h},\frac{\varepsilon^j}{h}\right)\frac{f\left(\varepsilon^i-ah\right)}{f\left(\varepsilon^i\right)}f\left(\varepsilon^i\right)f\left(\varepsilon^j\right)d\varepsilon^i d\varepsilon^j.$$

Anonymity of μ implies that

$$\begin{aligned}p(a,0)=1-p(0,a)&=\int\int\left(1-\mu\left(\frac{\varepsilon^i}{h},a+\frac{\varepsilon^j}{h}\right)\right)f\left(\varepsilon^i\right)f\left(\varepsilon^j\right)d\varepsilon^i d\varepsilon^j\\&=1-\int\int\mu\left(\frac{\varepsilon^i}{h},\frac{\varepsilon^j}{h}\right)\frac{f\left(\varepsilon^j-ah\right)}{f\left(\varepsilon^j\right)}f\left(\varepsilon^i\right)f\left(\varepsilon^j\right)d\varepsilon^i d\varepsilon^j.\end{aligned}$$

Thus

$$p(a,0)=\frac{1}{2}+\frac{1}{2}\int\int\mu\left(\frac{\varepsilon^i}{h},\frac{\varepsilon^j}{h}\right)\left(\frac{f\left(\varepsilon^i-ah\right)}{f\left(\varepsilon^i\right)}-\frac{f\left(\varepsilon^j-ah\right)}{f\left(\varepsilon^j\right)}\right)f\left(\varepsilon^i\right)f\left(\varepsilon^j\right)d\varepsilon^i d\varepsilon^j.$$

Log-concavity of f implies that $\frac{f\left(\varepsilon^i-ah\right)}{f(\varepsilon^i)}\geq\frac{f\left(\varepsilon^j-ah\right)}{f(\varepsilon^j)}$ if and only if $\varepsilon^j\leq\varepsilon^i$. As such, the investigation policy μ that maximizes $p(a,0)$ sets $\mu\left(x^i,x^j\right)=1$ whenever $x^i>x^j$ and $\mu\left(x^i,x^j\right)=0$ whenever $x^i<x^j$ — that is, the "investigate the agent with the higher signal" policy. ∎

Proof of Lemma 4: Part (I) follows from rewriting (1) in terms of q, which gives $q(a)+q(-a)=1=q(0)+q(0)$ for any $a\in[0,1]$. For Part (II), we use the explicit expression for the function q:

$$q(a)=\int\Pr\left(a+\frac{\varepsilon^i}{h}\geq\frac{\varepsilon^j}{h}\right)f\left(\varepsilon^j\right)d\varepsilon^j=\int\left(1-F\left(\varepsilon-ha\right)\right)f(\varepsilon)\,d\varepsilon.$$

Clearly $q(a)$ is increasing (decreasing) in precision h when a is positive (negative). For Part (III), we first evaluate the derivatives of q:

$$q'(a) = h \int f(\varepsilon - ha) f(\varepsilon) d\varepsilon > 0$$
$$q''(a) = -h^2 \int f'(\varepsilon - ha) f(\varepsilon) d\varepsilon.$$

Integration by parts implies

$$q''(a) = h^2 \int f(\varepsilon - ha) f'(\varepsilon) d\varepsilon.$$

Since f is log-concave, f'/f is decreasing. As such, $f'(\varepsilon - ha)/f(\varepsilon - ha) > f'(\varepsilon)/f(\varepsilon)$ for any $a > 0$, which in turn implies

$$\int f(\varepsilon - ha) f'(\varepsilon) d\varepsilon < \int f'(\varepsilon - ha) f(\varepsilon) d\varepsilon = -q''(a).$$

Thus we have established that $q''(a) < -q''(a)$, and so $q(a)$ must be concave over positive values of a. The symmetry property of Part (I) implies that q is likewise convex over negative values. ∎

Proof of Proposition 3: The statement regarding (4) is established in the main text prior. Here, we formally establish the statement relating to signal precision. Condition (4) is equivalent to

$$Q(a_M) \equiv (1 - a_M) q(a_M) q(1) + a_M q(0) q(1) - q(a_M) q(1 - a_M) > 0.$$

Evaluating the derivatives of Q gives

$$Q'(a_M) = q(0) q(1) + q(a_M)(q'(1 - a_M) - q(1)) + q'(a_M)((1 - a_M) q(1) - q(1 - a_M))$$
$$Q''(a_M) = 2q'(a_M)(q'(1 - a_M) - q(1))$$
$$+q''(a_M)((1 - a_M) q(1) - q(1 - a_M)) - q(a_M) q''(1 - a_M).$$

Observe that $Q(0) = Q(1) = 0$. As such, it suffices to show that $Q''(a_M) < 0$ for all a_M when precision h is low. To do so, we show that $\frac{Q''(a_M)}{q'(a_M)} \to -2q(1)$ as $h \to 0$.

We make use of the expressions for q' and q'' derived in Lemma 4. First,

$$\frac{q''(a)}{q'(a)} = -h \frac{\int f'(\varepsilon - ha) f(\varepsilon) d\varepsilon}{\int f(\varepsilon - ha) f(\varepsilon) d\varepsilon} \to 0 \text{ as } h \to 0,$$

since at $h = 0$, $\int f(\varepsilon - ha) f(\varepsilon) d\varepsilon > 0$ while $\int f'(\varepsilon - ha) f(\varepsilon) d\varepsilon = 0$.[23] Second,

$$\frac{q''(1 - a_M)}{q'(a_M)} = \frac{q'(1 - a_M)}{q'(a_M)} \frac{q''(1 - a_M)}{q'(1 - a_M)} \to 1 \times 0 = 0 \text{ as } h \to 0$$

Finally, $q'(1 - a_M) \to 0$ as $h \to 0$. ∎

Proof of Proposition 4: As a preliminary, observe that (IC1-0) holds whenever both (IC1-M) and (ICM-0) do. To see this, suppose to the contrary that for some penalty s_M and taste parameter λ (IC1-M) and (ICM-0) both hold but (IC1-0) fails. On the one hand, by straightforward algebra (ICM-0) can hold while (IC1-0) fails only if $s_M < Sa_M$. On the other hand, and again by straightforward algebra, (IC1-M) can hold while (IC1-0) fails only if $s_M > \frac{Sa_M q(0)}{q(a_M - 1)}$. Since $\frac{Sa_M q(0)}{q(a_M - 1)} > Sa_M$ this delivers the required contradiction.

The Proposition's statement regarding severe crime is immediate. The statement regarding moderate crime is immediate given that (ICM-1) holds whenever (IC1-M) fails (see Lemma 1). For the remaining case, suppose that neither (IC1-0) and (IC1-M) both hold, nor that (IC1-M) fails but (ICM-0) holds. There are two remaining possibilities. If (IC1-M) and (ICM-0) both fail, the only remaining equilibrium candidate is no crime. Alternatively, if (IC1-M) holds then (IC1-0) must fail. The preliminary above then implies that (ICM-0) fails. Again, the only remaining equilibrium candidate is no crime. ∎

Proof of Proposition 5: Differentiating the social welfare SW with respect to s_M gives:

$$\frac{\partial SW}{\partial s_M} = -(\lambda_M a_M - C_M) \frac{q(0)}{a_M} g(\lambda_M) - (C_1 - C_M - \lambda_1 (1 - a_M)) \frac{q(a_M - 1)}{1 - a_M} g(\lambda_1),$$

[23] Observe that $\int f'(\varepsilon) f(\varepsilon) d\varepsilon = \left[\frac{1}{2} f(\varepsilon)^2\right]_{-\infty}^{\infty} = 0$.

where $\lambda_M = \frac{s_M q(0)}{a_M}$ and $\lambda_1 = \frac{Sq(0) - s_M q(a_M - 1)}{1 - a_M}$. Substituting in,

$$\frac{\partial SW}{\partial s_M} = (C_M - s_M q(0)) \frac{q(0)}{a_M} g(\lambda_M) - (C_1 - C_M - Sq(0) + s_M q(a_M - 1)) \frac{q(a_M - 1)}{1 - a_M} g(\lambda_1).$$

Since the derivative is increasing in C_M and decreasing in C_1, the result follows. ∎

Proof of Proposition 6: From the proof of Proposition 5, the derivative $\partial SW/\partial s_M$ is decreasing in $q(a_M - 1)$, which implies the result. ∎

Proof of Proposition 7: First, observe (by straightforward algebra) that whenever s_M lies below $\frac{Sa_M q_N(1-a_M)}{(1-a_M)q_N(a_M) + a_M q_N(0)}$, condition (ICM-1) holds whenever (IC0-M) does.

Second, for any $x > 0$ and any N, $q_N(x) \to 1$ as signal precision $h \to \infty$. Consequently, since $a_M = 1/2$ the penalty level

$$\frac{Sa_M q_N(1 - a_M)}{(1 - a_M) q_N(a_M) + a_M q_N(0)} \to \frac{S}{1 + \frac{1}{N}} \text{ as signal precision } h \to \infty.$$

As such, whenever precision is high enough the penalty s_M does indeed lie below $\frac{Sa_M q_N(1-a_M)}{(1-a_M)q_N(a_M) + a_M q_N(0)}$.

Third, observe that for a penalty s_M below $\frac{Sa_M q_N(1-a_M)}{(1-a_M)q_N(a_M) + a_M q_N(0)}$, severe crime is the equilibrium outcome whenever (ICM-1) fails: in this case, (IC0-M) also fails, and hence the only equilibrium is severe crime.

Fourth, we claim that whenever precision h is high enough, a reduction in N (i.e., an increase in enforcement resources) leads (ICM-1) to fail for a larger set of taste realizations. We must show that, for all h sufficiently large,

$$Sq_N(1 - a_M) - s_M q_N(0) > Sq_{N-1}(1 - a_M) - s_M q_{N-1}(0).$$

Rearranging, and substituting in $q_N(0) = 1/N$,

$$\frac{s_M}{N(N-1)} > S(q_{N-1}(1 - a_M) - q_N(1 - a_M)).$$

The righthand side converges to 0 as signal precision $h \to \infty$, while the lefthand side is clearly independent of h.

Together, these four observations establish the result. ∎