Three Essays in Applied Microeconomic Theory

by

Dan C. Simundza

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in The University of Michigan
2011

Doctoral Committee:

    Professor Daniel S. Silverman, Chair
    Professor Tilman M. Borgers
    Professor James Jondall Prescott
    Assistant Professor Stephan Lauermann
    Professor Lones A. Smith, University of Wisconsin

Acknowledgements

I want to thank my wife Danna, my family, and my advisors.

Table of Contents

List of Figures

List of Tables

Chapter 1

Criminal Registries, Community Notification, and Optimal Avoidance

Introduction

Criminal registry notification laws provide information about offenders to at-risk neighbors with the intent of protecting the community. This paper investigates the effect of notification laws on the behavior of criminals and their law abiding neighbors and derives optimal notification policies. In my model, informed neighbors 1) can practice costly avoidance to protect themselves, and 2) better recognize and report criminal activity to the authorities, thereby increasing the probability of catching repeat offenders. Notification therefore generates opposing externalities: protecting oneself often comes at the cost of exposing one's neighbor, while increasing the probability of detection helps to deter criminal activity. Put simply, informed neighbors face a choice: remain outside in harm's way and help to deter crime, or retreat to the relative safety of their house at a cost.

Modeling the neighborhood as a game in which each informed family independently chooses its avoidance level (i.e. fraction of the day to stay indoors), I first study how notification policies affect neighborhood behavior. I

show that avoidance obeys strategic complementarities: the more the neighbors stay inside, the higher the incentives for each family to remain indoors. I prove that equilibria exist and are necessarily symmetric. Equilibrium avoidance may increase or decrease in the notification rate depending on whether informing the marginal citizen mainly serves to deter crime or decrease the average amount of time spent outside.

With a better understanding of the neighborhood's response, I turn to the issue of designing optimal notification policies. The government controls both the penalty on convicted felons and the notification rate, and the main results of the paper highlight the importance of getting these two policy levers working together.[1] I show that there always exists a penalty large enough to ensure that equilibrium avoidance decreases in the fraction of the population informed. Whenever this is the case, social welfare is necessarily increasing in the notification rate, and therefore maximized by a "scarlet letter" policy which informs the entire neighborhood. The higher notification rate leads to lower per-family expenditures on avoidance and higher deterrence, since the probability of detection is larger when more informed people are outside.

But this sword cuts both ways because notification with too small of a penalty is worse than useless: it is harmful. The informed alter their behavior at non-negligible costs and impose negative externalities on their neighbors. In this case, notification entails a cost but no benefit: the criminal is insufficiently deterred while residents suffer from staying indoors. If the penalty is too small, the government is better off keeping the criminals' identities secret. By not releasing any information the government ensures society does not waste energy on costly avoidance.

---

[1]The penalty could take many forms but all that matters in the analysis is the criminal's loss in utility from punishment. For instance, a fine of $1,000 may reduce a criminal's utility by the same amount as a year in prison. These penalties would be equivalent in the model. The model is not dynamic so incarceration is not an issue.

Background

The most well-publicized criminal registry notification policy is "Megan's law," which requires states to notify the public of registered sex offenders in their neighborhood.[2] But there are registries and notification policies for many other types of criminals. Some examples of approved or proposed registries and notification policies are for criminals convicted of elderly abuse, animal abuse, hate crimes, gang crimes and drug dealing [Welch [2008]]. Part of the rationale for community notification is that these types of offenders are especially prone to recidivism and often commit their crimes close to home.[3] Notification is thought to help at-risk neighbors protect themselves by monitoring and avoiding the potential threats in their community. If the community is made aware of criminals' presence, they can take actions to protect themselves.

In practice, governments do not directly choose what fraction of the neighborhood to inform. But they do choose the method and intensity of notification, which together determine to a large extent the fraction informed. Popular methods of notification include online databases searchable by anyone with an internet connection, individual letters sent in the post, holding a registry at the police station which can be accessed by those deemed "at risk", and door-to-door notification.[4] Each of these methods entails different costs. I ignore the physical costs of notification in this paper and instead focus on the costs generated by the strategic interaction of the informed neighbors. My goal is to clarify our understanding of the costly behavioral

---

[2]All 50 states currently have versions of Megan's law in place.

[3]Hanson et al. [2003] report that recidivism rates amongst sex offenders are at least 14% after 5 years, 20% after 10 years, and 30-40% after 20 years (these are understatements because some crimes go undetected). While criminal behavior for sex offenders declines in age of the offender, it does so more slowly than for other crimes (Hanson [2002]).

[4]Surveys described in Philips [1998], Kernsmith et al. [2009], and Lieb and Nunlist [2008] provide information about notification methods and rates.

response to notification.

The reasoning supporting notification is that informed neighbors can better avoid potential threats, or at least limit their exposure. But as this paper shows, information can be a burden to the community. Informed families alter their behavior in order to reduce their chances of being victimized. Once a family learns there is a potential danger, it can no longer be as carefree: doors must remain locked, children must be supervised, and outdoor activity in general is curtailed. Such avoidance behaviors are costly. Along with the administrative costs and potential harm to the criminal, the cost of the neighborhood response should be a consideration in any cost-benefit analysis for notification laws.

I model avoidance as displacing crime. When i stays inside more often, the probability j is attacked increases because criminals are less likely to select i from the crowd. The actual mechanism might be more subtle, but the results are the same. With the informed inside more often, there are simply less potential witnesses on the street, making crimes easier to commit. Further, when the informed are inside they cannot warn their neighbors and protect them as dangerous situations arise. If, in contrast, the informed stay outside more often, deterrence increases because they know who the criminals are and can watch over their neighbors. Criminals will be caught more often and will therefore be deterred from attacking unless the opportunity is especially good.

Literature Review

Since Becker [1968], the study of how criminals respond to incentives has flourished. Recent studies have investigated the deterrence effects of police and incarceration [Levitt [1997], Di Tella and Schargrodsky [2004]], conditions in prisons [Katz et al. [2003]], and gun ownership [Lott [1998], Mialon

and Wiseman [2006]]. Other papers study the relationship between reputation concerns and violent crimes [Silverman [2004]], incarceration's effect on the supply of crime [Freeman [1996]], the interplay of crime and vigilance in a general equilibrium framework [Smith [2010]], and optimal law enforcement [Eeckhout et al. [2010], Polinsky and Shavell [2000]]. My research adds to this literature by studying an additional crime fighting tool available to the government: community notification policies.

Previous research on criminal registries and notification policies has focused on sex offenders. These studies were primarily concerned with the effect notification had on recidivism rates and offender well-being [Prescott and Rockoff [2011], Adkins et al. [2000], Walker et al. [2005], Agan [2008]]. Much of this literature has difficulty finding any effect of registration and notification on recidivism rates. Prescott and Rockoff [2011] show that registration programs reduce crime by providing police with better information on offenders' whereabouts, but the effect of notification programs is ambiguous. Notification programs reduce first time offenses by effectively increasing the penalty, but may increase recidivism rates by lowering the return to non-criminal activities and thereby the incentives for good behavior.

In my model the crime rate increases in the notification rate only when equilibrium avoidance increases sufficiently quickly. In such a situation, higher notification rates lead to a higher fraction of uninformed people on the street, making crime easier and more attractive.

The extent of the literature's focus on neighborhood well-being is to quantify the effect a nearby offender has on property values [Linden and Rockoff [2008], Pope [2008]]. The results generally show a negative relationship between the existence of a sex offender and housing values, with the drop in housing prices ranging from 2.3% to 18%. I take this as evidence that the law-abiding public cares about the registry. Avoidance is unobservable and difficult to verify but the correlation between house price and notification

implies that individuals' response to being informed is non-negligible.

This paper also contributes to a large literature on the value of publicly disclosing information, of which Jin and Leslie [2003] and Dranove et al. [2003] are particularly relevant. Jin and Leslie [2003] studies the restaurant hygiene report card program instituted in Los Angeles and Dranove et al. [2003] studies the health care provider report card program instituted in New York. While more information to the consumer increases social welfare in Jin and Leslie [2003], Dranove et al. [2003] find the opposite due to providers' ability to "game the system".

My result that notification can be welfare increasing or decreasing combines elements of Jin and Leslie [2003] and Dranove et al. [2003] since informed neighbors' incentives may not be aligned with society's. Once informed, a family faces private incentives to protect itself by staying off the street. From society's point of view, this is exactly the wrong thing to do. Informed families should be outside as much as possible, thereby increasing the deterrence effect and helping the entire neighborhood. Whether notification or secrecy is optimal depends precisely on how much deterrence informing provides relative to how much families keep off the streets. The government imposed penalty is needed to generate sufficient deterrence so that incentives are aligned and notification improves welfare.

Model

Motivation

Before going into detail, in this section I provide an intuitive story to motivate the model. Consider a neighborhood consisting of a past offender (i.e. a potential threat) and neighbors A and B. Suppose, due to low effort and insufficient resources devoted to notification, only A becomes informed of

the criminal's existence. Neighbor A responds by undertaking avoidance measures while B does not alter his behavior. This avoidance is costly, but it reduces A's probability of being attacked. If the offender wants to commit a crime, the probability that B is the one attacked therefore increases. In this way notification generates negative externalities.

The positive externality, deterrence, is generated by the following mechanism. When informed, A learns the offender's characteristics and prepares himself in the event of an attack. If A is attacked, he is more likely to identify the offender as the person he saw in the registry, whereas B might have trouble remembering the characteristics of his attacker. The criminal is more likely to be caught if he attacks an informed neighbor. The criminal knows the government's notification policy and computes the expected probability of detection. Since he cannot identify who is informed, the deterrence helps everyone in the neighborhood.

Model

The neighborhood has a mass of criminals, normalized to 1, and a mass $P > 1$ of families. A fraction β of the families are informed of the criminals' existence, while fraction $1 - \beta$ are uninformed and have no idea criminals might be nearby.[5] That mass βP of the families are informed is common knowledge between the informed families and the criminals, but the criminals cannot identify who is informed.

Informed families choose avoidance level $a \in [0, 1]$, where a represents the fraction of the day they "lay low" inside, away from harm. Avoidance level a costs $c(a)$, where $c : [0, 1] \mapsto [0, \infty)$ is increasing and convex and satisfies $c(0) = c'(0) = 0$. Uninformed families make no choice. I normalize the

---

[5]Later I will discuss the implications of relaxing this assumption and letting the uninformed do inference after observing a noisy signal of their neighbors' actions.

harm of being attacked to one; if a mass m of criminals attack i, his utility decreases by m.

Criminals attack when the expected benefits outweigh the expected costs. I assume the value criminal j derives from attacking a family is $\theta_j \sim U[0,1]$. The realization of θ_j is known to the criminal before he decides whether to attack. Attacking criminals randomly choose one (and only one) family to attack. If he attacks and does not get caught his payoff is θ_j. If he attacks and gets caught, his payoff is $\theta_j - \tau$, where τ is the penalty imposed by the government. If the criminal does not commit the crime, he gets nothing; he cannot target another family.

When the criminal attacks an uninformed family, he is caught with probability d, while attacking an informed family results in a detection rate of $D \geq d$. These values are known to the criminal and the informed families.

Family i's expected cost of attacks when all other informed families use avoidance level a and a mass of m criminals attack is

$$m \; \frac{1 - a_i}{P(\beta(1 - a) + (1 - \beta))} = m \; \frac{1 - a_i}{P(1 - \beta a)} \tag{1}$$

Equation (1) shows that each family's expected harm from attacks is equal to the amount of time it spends outside relative to the total amount of time all neighborhood families spend outside, scaled by the mass of attacking criminals (and the cost of being attacked, which is one).[6]

There is a potential problem when everyone is informed and staying inside all day because if family i considers spending any part of the day outside, it will be attacked infinitely often. The following assumption guarantees there will always be some families outside.

A1 The maximum amount of avoidance is $\bar{a} < 1$.

---

[6]Since I will prove that equilibria are necessarily symmetric, fixing the avoidance level of all other families at a is without loss of generality.

In order to guarantee interior equilibria, I also make the following assumption on the cost curve.

A2 Staying inside all day is dominated by staying outside all day; that is: $c(\bar{a}) > \frac{\bar{a}}{P(1-\bar{a})}$.

The qualitative results which follow do not depend on these two assumptions. Without them, the analysis would be a bit tedious, mainly due to worrying about corner solutions.

Comments

Since the neighborhood is composed of a continuum of agents, each individual's actions has negligible effect on the community. The analysis will focus on symmetric equilibria in which all informed agents take the same action. So even though each individual family's avoidance has negligible effect on the community, there is an externality generated by the avoidance of the informed as a *group*.

Avoidance generates strictly negative externalities in this paper. When A stays inside more often, not only does the conditional probability that B is attacked increase, but the level of deterrence decreases. This happens because A is informed and entails the higher detection rate D. When he stays inside more often, he is less likely to be randomly selected by the criminal, and so the criminal faces a lower expected probability of being caught. It is possible to imagine situations in which the action taken by the informed generates positive externalities. As an extreme example, suppose all informed families carry weapons at all times. This action would generate a positive externality for the uninformed because the criminal would worry that his potential victim is armed.

The literature on property values and notification laws reports significant results only for very small distances (i.e. less than 0.1 miles). Therefore,

I think of the mass of the law-abiding neighbors relative to the mass of criminals as "small". It is difficult to explicitly define "small" in this context, but P should be small enough to generate incentives for avoidance, and to make the externality effects relevant.

Non-Deterrable Criminals

I start by analyzing non-deterrable criminals to isolate one aspect of the model. Without deterrence, notification generates purely negative externalities. Informed neighbors practice avoidance to protect themselves, which shifts the expected cost of attacks onto their neighbors. This effect exists in the more general model with rational criminals, but is especially easy to see when criminals cannot be deterred by increases in the probability of detection.

The economics of crime literature has largely proceeded under the assumption that criminals respond to incentives. There is however evidence that some criminals cannot be deterred from committing crimes.[7] What is important for this section is that non-deterrable criminals always commit crimes; any opportunity is sufficient for them. The analysis here will be brief

---

[7]Lee and McCrary [2009] exploit the discontinuous increase in penalties as criminals turn 18 years old. They find an extremely small response: criminals a few days younger than 18 years old do not commit many more crimes than those a few days over 18 despite facing much shorter jail sentences. This, as they point out, could be evidence that the criminals are not responding to incentives, or simply myopia. Wright and Decker [1994] and Cromwell and Olson [2003] both study samples of burglars who have committed hundreds (and in some cases thousands) of thefts in their lifetimes. These criminals seem unable to resist committing a crime if the opportunity presents itself. That said, even these non-deterrable criminals regularly make "occupancy probes" in order to lower their probability of detection. An occupancy probe might take the form of consulting funeral announcements and/or contacting relevant houses by phone to make sure no one is home.

with the intent of simplifying one aspect of the more complicated general model.[8]

Informed family i takes as given the other informed families' avoidance levels and chooses a_i to minimize the expected costs of being attacked plus the costs of vigilance. That is, they solve:

$$\max_{a_i} \left\{ -\frac{1-a_i}{P(1-\beta a)} - c(a_i) \right\}$$

Optimal avoidance is always interior because marginal benefits exceed marginal costs at $a_i = 0$, and assumption A2 ensures $a_i = \bar{a}$ is dominated by $a_i = 0$. Optimal avoidance necessarily satisfies the first order condition

$$\frac{1}{P(1-\beta a)} = c'(a_i) \tag{2}$$

I write family i's best response when all other informed families use avoidance level a as $a_i(a)$. Family i's best response function is increasing because c' is increasing. The equation characterizing a symmetric equilibrium is

$$\frac{1}{P(1-\beta a)} = c'(a) \tag{3}$$

I will focus attention on equilibria that are locally dynamically stable. Such equilibria might be observed as the outcome of a dynamic process in which a small fraction of the informed neighbors adjust their avoidance each period by best responding to the average avoidance level observed in the neighborhood. The resting point of this dynamic process will be an equilibrium which is stable in that the best response to an avoidance level close to the equilibrium pulls the average avoidance level closer to the equilibrium. Equilibria which are locally dynamically unstable would never be observed as the outcome of this process. I will give a practical definition of local dynamic stability here in the text, and discuss the issue in more detail in the appendix.

---

[8]Important details like the existence of an equilibrium will be ignored for now but proved for the more general model later in the paper.

Figure 1.1 The Best Response Function
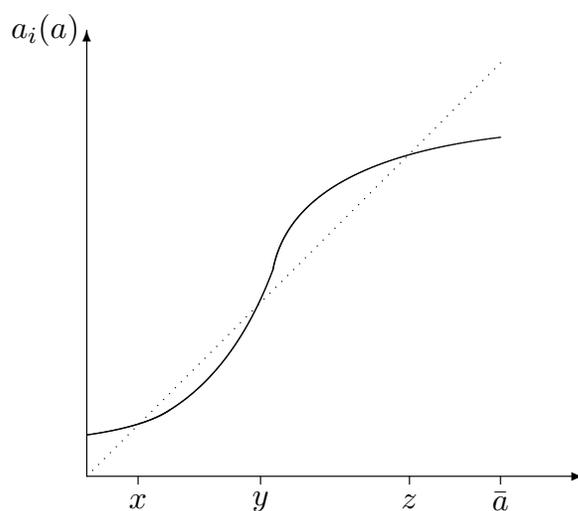
Definition A symmetric equilibrium is locally dynamically stable if the best response curve intersects the 45° line from above.

Figure 1.1 helps to illustrate this concept. In this figure, there are three symmetric equilibria, labeled x, y, and z. Avoidance levels x and z are locally dynamically stable because the best response curve intersects the 45° line from above, while avoidance level y is unstable.

We can now deduce equilibrium behavior as the fraction of neighbors informed of the criminals' existence increases. When there is no possibility of deterrence, as the notification rate increases and more neighbors take refuge in their house, the amount of time each stays inside increases.

Lemma 1. *For non-deterrable criminals, equilibrium avoidance increases in the notification rate.*

Proof. As β increases, the left hand side of family i's first order condition in (2) increases. This means the best response curve shifts up. Since the best response curve intersects the 45° line from above in any dynamically stable equilibrium, the equilibrium avoidance level must increase as the notification rate increases. $\square$

Since criminals cannot be deterred from crime by the higher probability of detection, the only externality at play here is the negative externality of avoidance. A higher notification rate leads to more people who practice avoidance. With more people practicing avoidance, the marginal benefit increases, which in turn causes everyone to practice yet more avoidance. Informed citizens face stronger incentives to stay inside, thereby increasing the expected harm to anyone who remains outside on the streets.

In this case, the government's optimal notification policy is to keep the criminals' identities secret. Define neighborhood welfare as the integral of the neighbor's utilities. Since notification entails costly avoidance but cannot deter crime, the government maximizes neighborhood welfare by not informing anyone of the criminals' existence.

Lemma 2. *For non-deterrable criminals, neighborhood welfare is decreasing in the notification rate. The optimal notification policy is to notify no one.*

Proof. Write equilibrium avoidance as a function of the notification rate as $a(\beta)$. Since the mass of attacks is always one, neighborhood welfare as a function of the notification rate is

$$w(\beta) = -1 - P\beta c(a(\beta))$$

As β increases, both the fraction of the population bearing the cost of avoidance *and* the amount of avoidance they use increases. Welfare decreases in the notification rate and is therefore maximized at $\beta = 0$. □

Notification policies, in the absence of deterrence, are harmful at the societal level because they cause families to change their behavior (at a cost) but do not affect the mass of criminals who attack. Notification merely displaces crime from one group of citizens to another. The total amount of crime stays constant while expenditures on avoidance increase. The welfare

maximizing policy is to notify no-one, and thus ensure society does not waste its effort on costly avoidance.

In order for notification to be good for the community, it therefore must be the case that criminals respond to incentives. The only way for the community to be made better off by notification is for the crime rate to decrease as more people learn of the criminals' existence. When more people know who is potentially dangerous, the criminals are more likely to get caught committing a crime. They therefore become more selective in determining when to attack, which decreases the danger to their neighbors.

Rational Criminals

I now turn to the more general model of rational criminals. The negative externality isolated in the analysis of non-deterrable criminals is still present, but it can be counteracted by the positive externality of deterrence. Criminals are more likely to be caught when attacking an informed neighbor, but they cannot identify who is informed and who is not. Holding everything else constant (including avoidance), higher notification rates therefore lead to more deterrence. Counteracting this is the fact that once informed, neighbors face a private incentive to practice avoidance which robs the neighborhood of the desirable deterrence effect.

Rational criminals commit crimes when the expected benefits outweigh the expected costs. Let δ represent the criminals' expected probability of detection, which will be determined by equilibrium behavior. The criminals attack when $(1 - \delta)\theta + \delta(\theta - \tau) \geq 0$, or $\theta \geq \tau\delta$. Since criminal values are distributed uniformly on the unit interval, the mass of attacking criminals is $1 - \tau\delta$.[9]

---

[9]Of course, this is only true if $\tau\delta \in [0, 1]$. I restrict attention to this case because when $\tau\delta > 1$, no crimes are committed.

A symmetric equilibrium is defined by a cutoff criminal value $\underline{\theta}$ at which criminals are indifferent between attacking and not, and an avoidance level a which is optimal for each informed family given that all other informed families use avoidance a. I will later show that any equilibrium is necessarily symmetric, so the analysis proceeds under this assumption.

Recall that the probability of being caught when attacking an uninformed family is d and the probability when attacking an informed family is $D > d$. Criminals' expected probability of detection in an equilibrium where the notification rate is β and all informed families use avoidance level a is

$$\delta(\beta, a) = \frac{(1 - \beta)d + \beta(1 - a)D}{1 - \beta a}$$

The cutoff criminal value is therefore

$$\underline{\theta} = \tau \frac{(1 - \beta)d + \beta(1 - a)D}{1 - \beta a}$$

Family i's problem is to minimize the expected costs of attacks plus the cost of avoidance, taking other informed families' avoidance level a and the cutoff criminal value described above as given:

$$\max_{a_i} \left\{ -\left(1 - \tau \frac{(1 - \beta)d + \beta(1 - a)D}{1 - \beta a}\right) \frac{1 - a_i}{P(1 - \beta a)} - c(a_i) \right\}$$

Optimal avoidance is interior because marginal benefits exceed marginal costs at $a_i = 0$, and assumption A2 implies $a_i = \bar{a}$ is dominated by $a_i = 0$. Family i's optimal avoidance necessarily satisfies the first order condition

$$\left(1 - \tau \frac{(1 - \beta)d + \beta(1 - a)D}{1 - \beta a}\right) \frac{1}{P(1 - \beta a)} = c'(a_i) \qquad (4)$$

I write family i's best response when all other informed families are using avoidance level a as $a_i(a)$. A symmetric equilibrium is characterized by the equation

$$\left(1 - \tau \frac{(1 - \beta)d + \beta(1 - a)D}{1 - \beta a}\right) \frac{1}{P(1 - \beta a)} = c'(a) \qquad (5)$$

Note that avoidance obeys strategic complementarities. The expected harm of attacks for a family who remains outside increases as the neighbors spend more time indoors for the simple reason that there are less people on the streets for the criminals to attack. Family i's incentives for avoidance are therefore increasing in the amount of avoidance their neighbors practice. This can be seen in family i's first order condition in (4) because both terms on the left hand side are increasing in a, implying a_i increases in a.

I start by studying the equilibrium behavior of criminals and neighbors with the goal of getting a better understanding of how equilibrium behavior depends on the notification rate. The following lemma shows that symmetric equilibria exist, and further, that *all* equilibria are necessarily symmetric. It does however allow for the possibility that multiple symmetric equilibria exist.

Lemma 3. *Locally dynamically stable equilibria exist and are necessarily symmetric.*

Proof. When $a = 0$, the left hand side of (5) is strictly positive while the right hand side is zero. When $a = \bar{a}$, assumption A2 guarantees the right hand side is larger than the left hand side. The intermediate value theorem guarantees there is some $a \in (0, \bar{a})$ where (5) holds with equality. Moreover, since $a_i(0) > 0$ and $a_i(\bar{a}) < \bar{a}$, the best response curve must cross the 45° line from above at least once.

Equilibria are necessarily symmetric because the marginal benefit of avoidance is independent of one's own avoidance while marginal costs are strictly increasing in own avoidance. It is impossible for marginal benefits to equal marginal costs at two or more distinct avoidance levels. □

Multiple symmetric equilibria may exist in this model. As in Section , I will focus attention on locally dynamically stable equilibria. These are the outcomes that we would expect to observe if the neighborhood arrived

at an equilibrium through a dynamic process where each period a fraction of the informed families best respond to their neighbors' average avoidance level.[10] Practically, locally dynamically stable equilibria occur where the best response curve intersects the 45° line from above.

The comparative static I am most interested in is how equilibrium avoidance depends on the notification rate. The unambiguous result that equilibrium avoidance increases in the notification rate for non-deterrable criminals in Lemma 1 derives from the fact that the term $1/(P(1 - \beta a))$, the marginal benefit of avoidance when criminals are non-deterrable, increases in β for all a. Rational criminals, however, can be deterred by notification; for fixed avoidance, the probability of detection increases in the notification rate. This means that the mass of attacking criminals, $1 - \tau \delta(\beta, a)$, decreases in the notification rate. Whether equilibrium avoidance increases or decreases in the notification rate depends on how much extra deterrence the marginal informed citizen provides relative to how much less time the informed spend outside.

Before answering this question analytically, it is useful to think about how changes in the notification rate affect the best response function. If an increase in β causes a sufficiently large increase in $1 - \tau \delta$ relative to the decrease in $1/(P(1 - \beta a))$, the left hand side of the first order condition in equation (4) increases. This causes agent i's best response function to shift upwards. Looking at Figure 1.1, we see that in any locally dynamically stable equilibrium an upward shift in the best response function corresponds to a larger equilibrium avoidance level. The opposite is true if the best response curve shifts downwards.

Lemma 4. *Equilibrium avoidance decreases in the notification rate if and only if, as the notification rate increases and avoidance is held constant, the*

---

effect on deterrence is sufficiently large relative to the effect on the average time spent outside per family.

Proof. Let $\pi(\beta, a) \equiv \frac{1}{P(1-\beta a)}$. Differentiating the left hand side of (4) with respect to β gives

$$-\tau \delta_\beta \pi + (1 - \tau \delta) \pi_\beta$$

The best response function shifts downwards whenever this expression is negative. Since the best response function intersects the 45° line from above in any locally dynamically stable equilibrium, a downwards shift in the best response function leads to a lower equilibrium avoidance level. Avoidance decreases in the notification rate if

$$\frac{\pi_\beta}{\pi} < \frac{\tau \delta_\beta}{1 - \tau \delta}$$

The left hand side of this inequality can be simplified to $\frac{a}{1-\beta a}$, so I re-write the above as

$$\frac{a}{1 - \beta a} < \frac{\tau \delta_\beta}{1 - \tau \delta}$$

Since $1 - \beta a = \beta(1 - a) + (1 - \beta)$, the left hand side of this inequality is the percentage decrease in the average amount of time each family spends outside. The right hand side is the percentage decrease in the mass of attacking criminals. $\qquad \square$

This result helps to clarify equilibrium behavior as the notification rate increases. As more of the neighborhood learns of the criminals' existence, two things happen. First, there are more people who know to look out for the criminals and can better recognize them if an attack happens. This raises the probability an attacking criminal is caught and helps to deter crime. Second, there are more people who can now practice avoidance. If the deterrence effect of notifying additional neighbors is large enough, equilibrium avoidance decreases; everyone can spend a bit more time outside because the mass of

criminals who have an incentive to attack decreases. If, on the other hand, notifying additional neighbors mainly serves to take them off the streets and increase the other families' probability of being targeted by criminals, then the equilibrium avoidance level will increase.

Building on the previous lemma, I derive a result on the shape of the equilibrium avoidance function. Let $a(\beta)$ be the equilibrium avoidance level as a function of the notification rate. The following lemma is primarily useful in proving the main results. Here I show that the equilibrium avoidance function is either always increasing, always decreasing, or first increasing and then decreasing in the notification rate.

Lemma 5. *The equilibrium avoidance function is single peaked in the notification rate.*

Proof. I show that whenever the slope of equilibrium avoidance equals 0, the second derivative is negative. The slope of equilibrium avoidance, given by the implicit function theorem, is

$$\frac{da}{d\beta} = \frac{\tau\delta_\beta\pi - (1-\tau\delta)\pi_\beta}{(1-\tau\delta)\pi_a - \tau\delta_a\pi - c''(a)} \tag{6}$$

The second derivative when $da/d\beta = 0$ is

$$\frac{d^2a}{d\beta^2} = \frac{(\tau\delta_{\beta\beta}\pi + 2\tau\delta_\beta\pi_\beta - (1-\tau\delta)\pi_{\beta\beta})((1-\tau\delta)\pi_a - \tau\delta_a\pi - c''(a))}{((1-\tau\delta)\pi_a - \tau\delta_a\pi - c''(a))^2} \ ,$$

because $\tau\delta_\beta\pi - (1-\tau\delta)\pi_\beta = 0$ by assumption. The denominator is always positive and $(1-\tau\delta)\pi_a - \tau\delta_a\pi - c''(a) < 0$ because the marginal cost curve intersects the marginal benefit curve from below in any equilibrium. Simple algebra gives

$$\tau\delta_{\beta\beta}\pi + 2\tau\delta_\beta\pi_\beta - (1-\tau\delta)\pi_{\beta\beta} = \tau\delta_\beta\frac{\delta_{\beta\beta}}{\delta_\beta}\pi - (1-\tau\delta)\pi_\beta\frac{\pi_{\beta\beta}}{\pi_\beta} + 2\tau\delta_\beta\pi_\beta \tag{7}$$

Since

$$\frac{\delta_{\beta\beta}}{\delta_\beta} = \frac{\pi_{\beta\beta}}{\pi_\beta} = \frac{2a}{1-\beta a} > 0 \ ,$$

19

the expression in (7) becomes

$$\frac{2a}{1 - \beta a}(\tau\delta_\beta\pi - (1 - \tau\delta)\pi_\beta) + 2\tau\delta_\beta\pi_\beta$$

Since $a'(\beta) = 0$ when $\tau\delta_\beta\pi - (1 - \tau\delta)\pi_\beta = 0$, second derivative $a''(\beta) < 0$ whenever $a'(\beta) = 0$. This means $a'(\beta)$ equals 0 at most once, and the equilibrium avoidance function is quasi-concave. $\square$

Lemma 5 implies that if equilibrium avoidance is decreasing at β_0, it is also decreasing for any $\beta_1 > \beta_0$. Similarly, if equilibrium avoidance is increasing at β_1, it is increasing for any $\beta_0 < \beta_1$. This result is useful because the behavior of equilibrium avoidance near the extreme values of β, which is often easiest to determine, can give information about the interior range. For instance, if $\lim_{\beta\to 0} a'(\beta) < 0$, then we know that equilibrium avoidance is decreasing in the notification rate at all possible notification rates. This fact will be used to prove the main results.

Welfare

The analysis up to this point has been mostly positive in focus. Lemma's 3, 4, and 5 help us to understand how the neighborhood responds to being notified an offender lives in their midst. In order to determine when notification is optimal a measure of community welfare is needed. I define welfare as the integral of the utilities of all of the neighbors living in the community, ignoring the criminals.[11] Then equilibrium welfare is

$$w(\beta) = -(1 - \tau\delta(\beta, a(\beta))) - P\beta c(a(\beta))$$

The next Lemma relates neighborhood welfare to equilibrium behavior when equilibrium avoidance is decreasing.

---

[11]Criminals only receive utility from successfully completing attacks. Including their utility in the welfare calculation would partially offset the lost utility of the neighbors.

Lemma 6. *Social welfare is increasing in the notification rate whenever equilibrium avoidance is decreasing in the notification rate.*

Proof. Differentiating equilibrium social welfare with respect to the notification rate gives

$$w'(\beta) = \tau(\delta_\beta + \delta_a a'(\beta)) - Pc(a(\beta)) - P\beta c'(a(\beta))a'(\beta)$$

The terms δ_β, $\delta_a a'$, and $-P\beta c' a'$ are positive, so it suffices to show $\tau\delta_\beta - Pc(a(\beta)) - P\beta c' a' > 0$. I note three facts used to show this sufficient condition holds.

1. Since $a(\beta)$ is equilibrium avoidance, it gives higher utility than using $a = 0$:

$$-(1 - \tau\delta)\frac{1 - a(\beta)}{P(1 - \beta a(\beta))} - c(a(\beta)) > -(1 - \tau\delta)\frac{1}{P(1 - \beta a(\beta))}$$

$$\Rightarrow \ (1 - \tau\delta)\frac{a(\beta)}{1 - \beta a(\beta)} > Pc(a(\beta))$$

2. From Lemma 4, we know that $a'(\beta) < 0 \Rightarrow \tau\delta_\beta > (1 - \tau\delta)\frac{\pi_\beta}{\pi}$.

3. From differentiating $\pi = 1/(P(1 - \beta a))$, $\frac{\pi_\beta}{\pi} = \frac{a(\beta) + \beta a'(\beta)}{1 - \beta a(\beta)}$.

The sufficient condition $\tau\delta_\beta - Pc(a(\beta)) - P\beta c' a' > 0$ holds because

$$
\begin{aligned}
\tau\delta_\beta - Pc(a) - P\beta c' a' \ &> \ \tau\delta_\beta - (1 - \tau\delta)\frac{a}{1 - \beta a} - P\beta c' a' \\
&> \ (1 - \tau\delta)\left(\frac{\pi_\beta}{\pi} - \frac{a}{1 - \beta a}\right) - P\beta c' a' \\
&= \ (1 - \tau\delta)\left(\frac{a + \beta a'}{1 - \beta a} - \frac{a}{1 - \beta a}\right) - P\beta c' a' \\
&= \ (1 - \tau\delta)\frac{\beta a'}{1 - \beta a} - P\beta c' a' \\
&= \ a'\beta P\left[(1 - \tau\delta)\frac{1}{P(1 - \beta a)} - c'\right] = 0
\end{aligned}
$$

21

where the first inequality follows from Fact 1, the second from Fact 2, and the first equality from Fact 3. The final equality holds because the first order condition is satisfied at the optimal avoidance level. I conclude that $a'(\beta) < 0 \Rightarrow w'(\beta) > 0$. □

This result shows that when equilibrium avoidance decreases in the notification rate, informing more neighbors makes the community better off. Notifying more neighbors lowers the per-family expenditure on avoidance. This increases deterrence because more people know *and* those who know are outside more often. The proof shows that, despite increasing the mass of families who use avoidance (and therefore adjust their behavior at a cost), a marginal increase in the notification rate is welfare improving.

With these results in place, I now state and prove the main results of the paper. These results demonstrate the importance of the government imposed penalty in determining optimal notification policies. Specifically, penalties and notification complement one another. I first show that for severe enough penalties, equilibrium avoidance decreases over the entire range of notification rates and, as Lemma 6 says, neighborhood welfare therefore increases. This means that notification can be worthwhile for communities if the government imposes sufficiently severe penalties on repeat offenders.

Proposition 1. *There always exists a severe enough penalty so that full notification is the optimal policy.*

Proof. As the notification rate approaches zero, the slope of equilibrium avoidance from (6) is

$$\lim_{\beta \to 0} a'(\beta) = \frac{a(1 - \tau d) - \tau(1 - a)(D - d)}{Pc''(a)} \; ,$$

which is negative whenever

$$\frac{a}{1 - a} < \frac{\tau(D - d)}{1 - \tau d} \tag{8}$$

22

The right hand side of (8) increases in the penalty τ and gets arbitrary large as $\tau \to 1/d$.

I now show that the left hand side of (8) approaches 0 for large enough penalties. Define κ as inverse to marginal costs, so that $\kappa(c'(a)) = a$. As $\beta \to 0$, equilibrium avoidance satisfies $\lim_{\beta \to 0} a(\beta) = \kappa((1 - \tau d)/P)$. By the properties of the cost function, κ is increasing and $\kappa(0) = 0$. As $\tau \to 1/d$, both avoidance and the function $a/(1 - a)$ fall to 0.

Then as $\tau \to 1/d$, the left hand side of (8) approaches 0 while the right hand side gets arbitrarily large. This means for large enough τ, $\lim_{\beta \to 0} a'(\beta) < 0$. By Lemma 5, equilibrium avoidance decreases over the entire range of notification rates and by Lemma 6 the neighborhood's welfare is maximized by notifying everyone. $\qquad\square$

Proposition 1 helps to understand both when and why community notification improves welfare. Severe penalties help the efficacy of notification. Large penalties on their own can deter crime, but the effect is magnified when coupled with community notification. The penalty and notification are complements. Informed neighbors entail a higher probability of detection, so the probability an attacking criminal has to pay the higher penalty increases with notification. This lowers the chance that a criminal will attack and also, importantly, the incentives for avoidance. The informed therefore spend more time outside, further generating deterrence.

The community is effectively empowered by public notification in conjunction with severe penalties. As more of the neighbors are informed, the probability of detection continues to rise and so the informed can spend more time outside instead of barricading themselves inside their house.

The next result provides a partial converse to Proposition 1 and is closely related to the analysis of non-deterrable criminals. If the penalty chosen by the government is too small, criminals will not be sufficiently deterred. In this case, informing generates a burden on society because the community spends

a lot on avoiding criminals but cannot completely insulate itself. Someone is always left outside, providing an easy target for the attacking criminals.

Proposition 2. *For small enough penalties, notifying no-one is optimal.*

Proof. I will show that when the penalty is lenient enough, neighborhood welfare is decreasing for all β, implying the optimal notification rate is $\beta = 0$. Equilibrium welfare is $w(\beta) = -(1 - \tau\delta(\beta, a(\beta))) - P\beta c(a(\beta))$, and so the slope is

$$w'(\beta) = \tau(\delta_\beta + \delta_a a'(\beta)) - Pc(a(\beta)) - P\beta c'(a(\beta))a'(\beta)$$

I first show that for small enough τ, $a'(1) > 0$. When $\beta = 1$, the slope of equilibrium avoidance equals

$$\left.\frac{da}{d\beta}\right|_{\beta=1} = \frac{P\tau(D - d) - (1 - \tau D)a(1)}{P^2(1 - \tau D)\beta - P^2(1 - a(1))^2 - c''(a(1))}$$

Fix Me! The denominator is negative in any equilibrium, and for small enough τ the numerator is also negative, implying $a'(1) > 0$. Lemma 5 ensures that $a'(\beta) > 0$ for all β whenever τ is small enough to make $a'(1) > 0$.

Then for τ small enough that $a'(1) > 0$, we know

$$w'(\beta) < \tau\delta_\beta - Pc(a(\beta))$$

because both $\delta_a a'(\beta)$ and $-P\beta c'(a(\beta))a'(\beta)$ are negative. The term δ_β is bounded above by $(D - d)/(1 - \bar{a})$, so as $\tau \to 0$, we know $\tau\delta_\beta \to 0$.

This means we can choose τ small enough so that both $\tau\delta_\beta < Pc(a(0))$ and $a'(1) > 0$ hold, guaranteeing that $w'(\beta) < 0$ for all β. $\square$

Very small penalties fail to deter crime, even if the probability of detection is high. In the extreme case where $\tau = 0$, even certain detection cannot deter crime. Notification cannot generate sufficient deterrence with small penalties,

despite the higher probability criminals are caught when attacking informed neighbors.

But notification imposes costs on the informed because they face high incentives to practice costly avoidance. From a welfare perspective, this avoidance is wasted because neighbors cannot isolate themselves from the criminals all of the time and crimes are committed. Notification has the effect of making the informed work harder without providing any benefit to the community.

As an extreme example, consider the limiting case of $\tau = 0$. The model is then identical to the model with non-deterrable criminals where avoidance increases in the notification rate and welfare decreases in the notification rate. Informing no-one therefore maximizes welfare.

Criminal registry and community notification programs are sometimes thought of as a ways to keep offenders in "prison" even after they are released. Registered criminals are watched over more closely by local law enforcement and the community than if they were anonymous members of society. But this additional "prison sentence" is not enough; it cannot act as a substitute for a penalty. The community needs a sufficiently large penalty to use as a threat against the criminal in combination with a notification policy.

These results highlight the importance of getting the government's two policy levers working together. Before deciding whether to notify, the government needs to determine if the penalty is such that notifying will help or harm the community.

The results in Propositions 1 and 2 result in zero or full-on notification as optimal policies. There are situations where the optimal notification policy is some interior fraction of the population. That is, informing some but not all of the neighborhood can maximize utilitarian social welfare. This occurs when then penalty is in some sort of "middle ground", neither severe nor lenient enough to generate unequivocal recommendations. Hence, Proposition

2 is only a *partial* converse to Proposition 1.

Discussion

Communication Amongst Neighbors

One might be concerned that the government does not have complete control over the notification rate. In particular, once the news that a potential threat exists is released, neighbors might communicate amongst themselves. In this way, even if the government notifies just a small fraction of the population, everyone may end up learning of the threat indirectly through their neighbors.

In this subsection I discuss the incentives neighbors face to communicate their information. I check when an informed agent's utility is increasing in the notification rate. If this is the case, the individual faces private incentives to communicate his information to his uninformed neighbors. If utility is decreasing in the notification rate, the informed do not want anyone else to learn of the potential threat.

The magnitude of the penalty is again important in determining the incentives informed neighbors face to share their information. Large penalties facilitate the sharing of information amongst the neighbors while small penalties encourage the informed to keep their information private.

Proposition 3. *The utility of the informed increases (decreases) in the notification rate when penalties are sufficiently large (small).*

Proof. I first show that the utility of an informed neighbor is decreasing when $\tau = 0$. I study utility at two notification rates, $\beta_1 < \beta_2$. Let a_k be optimal for β_k and let U_k be the utility of the informed in equilibrium when the notification rate is β_k. Lemma 1 shows that avoidance is increasing when

$\tau = 0$, so $a_1 < a_2$. Then

$$U_1 \equiv -\frac{1-a_1}{P(1-\beta_1 a_1)} - c(a_1) \geq -\frac{1-a_2}{P(1-\beta_1 a_1)} - c(a_2) > -\frac{1-a_2}{P(1-\beta_2 a_2)} - c(a_2) \equiv U_2$$

The first (weak) inequality holds because a_1 is optimal when the notification rate is β_1. The second (strict) inequality holds because $\beta_1 < \beta_2$, avoidance is increasing in the notification rate, and so $\beta_1 a_1 < \beta_2 a_2$. This shows that the utility of the informed is decreasing in the notification rate when $\tau = 0$. Since utility and avoidance are continuous in τ, the utility of the informed is decreasing in the notification rate for sufficiently small penalties.

I next show that an informed neighbor's utility is increasing if avoidance is decreasing, which Proposition 1 says occurs when the penalty is sufficiently large. Write the utility of an informed agent as $u(\beta) = -(1 - \tau\delta(\cdot))\frac{1-a}{P(1-\beta a)} - c(a)$. Then the derivative of utility in the notification rate is

$$
\begin{aligned}
u'(\beta) \quad &= \tau(\delta_\beta + \delta_a a')\frac{1-a}{P(1-\beta a)} - (1-\tau\delta)\frac{a-a^2+a'(\beta-1)}{P(1-\beta a)^2} - c'(a)a' \\
&= \tau(\delta_\beta + \delta_a a')\frac{1-a}{P(1-\beta a)} - (1-\tau\delta)\frac{a-a^2}{P(1-\beta a)^2} + a'\left[\frac{1-\tau\delta}{P(1-\beta a)}\frac{1-\beta}{1-\beta a} - c'(a)\right] \\
&= \frac{1-a}{P(1-\beta a)}\left[\tau\delta_\beta - (1-\tau\delta)\frac{a}{1-\beta a}\right] + \tau\delta_a a'\frac{1-a}{P(1-\beta a)} + a'\left[\frac{1-\tau\delta}{P(1-\beta a)}\frac{1-\beta}{1-\beta a} - c'(a)\right]
\end{aligned}
$$

Whenever $a' < 0$, Lemma 4 says that $\tau\delta_\beta > (1-\tau\delta)\frac{a}{1-\beta a}$, so the first term is positive. The second term is positive because both $\delta_a < 0$ and $a' < 0$. The term

$$\frac{1-\tau\delta}{P(1-\beta a)}\frac{1-\beta}{1-\beta a} - c'(a)$$

is negative because the first order condition says $\frac{1-\tau\delta}{P(1-\beta a)} - c'(a) = 0$ and $\frac{1-\beta}{1-\beta a} < 1$. Then the last term is positive because both the term in square brackets and a' are negative. This means that when penalties are sufficiently large, avoidance decreases in the notification rate and the utility of the informed increases in the notification rate. $\square$

This helps to show that the result that penalties and notification are complementary is not due to the fact that neighbors do not communicate. If the

informed are allowed to communicate their information to their neighbors, the government can facilitate this by using sufficiently large penalties. In this case, by telling their neighbors, the informed can decrease their probability of being attacked and their expenditure on costly avoidance.

Small penalties, on the other hand, cause the informed to keep their information private. Informing other people in the neighborhood causes more families to practice avoidance, raising costs for all. In the future, I will study situations where the informed face incentives to keep their information private but society would benefit from its transmission. Such a case would provide a clear justification for the role of governments in community notification policies.

Public Good vs. Public Bad

The goal of this paper is to enhance our understanding of the effects of community notification on neighborhood behavior and welfare, and to determine when notification is optimal. How agents respond to being notified that an offender lives nearby is largely an empirical question. The action I study, avoidance, generates negative externalities. Here I discuss the possibility and implications of neighbors undertaking actions which generate positive externalities.

Data on how people alter their behavior in response to the threat of crime is scarce. I have not found studies which specifically deal with how neighbors change their behavior upon being notified a criminal lives nearby through community notification policies. In a national survey of 1,101 individuals, Ferraro [1995] asks a more general question: what activities have respondents undertaken to reduce their risk of crime? Their responses are summarized in Table 1.1.

Table 1.1 shows that people take actions which have both positive and

Table 1.1 Actions Undertaken in Response to Fear of Crime

Externality	Action	% Yes
	Avoid Unsafe Areas - Day	51.7
	Avoid Unsafe Areas - Night	75.3
Negative	Additional Locks on House	57.3
	Buy Watchdog	24.5
Positive & Negative	Additional Outside Lighting	57.5
	Learn Self Defense	38.3
Positive	Keep Weapon in Home	40.7
	Carry Defensive Object	18.6

Source: Ferraro [1995]. Responses from a nationally representative survey of 1,101 individuals.

negative externalities in order to protect themselves from the threat of crime. For instance, avoiding unsafe areas only serves to make those areas more unsafe for the people who must go there. Avoidance is thus a public bad. On the other hand, criminals are less likely to attack if they are aware that some potential victims know self defense or might be carrying a weapon. This type of vigilance therefore creates a public good. Which effect dominates is an empirical question I do not attempt to answer.

It is worth pointing out that sometimes local governments pass laws limiting actions which generate positive externalities in this context. For instance, Washington D.C. had a ban on residents owning handguns from 1976-2008.

I will refer to any action which generates positive externalities for the neighborhood as "vigilance". The probability of a successful attack is lower against someone who is vigilant. Criminals are therefore less likely to attack people if they know that some are vigilant. As more people use vigilance, the deterrence effect increases.

If the actions people take entail positive as well as negative externalities,

I expect notification to be optimal more often. That is, the results in this paper are possibly pessimistic towards notification policies. While the policy recommendation in this paper may be to not notify, a study which takes account of the positive externality of vigilance behaviors may recommend some positive amount of notification. Said differently, any time my results advise notification, a model which also takes account of vigilance efforts would also advise notification, but not necessarily the other way around.

Inference by the Uninformed

In this paper the uninformed are constrained to take no action. They are unaware that something bad could happen to them. An interesting extension to this model would allow the uninformed to conduct inference after observing a noisy signal which is correlated with the neighborhood's behavior. Uninformed people might look around while outside and notice there is no one else out. The lack of people outside might be because everyone else is indoors avoiding the potential threat, or simply for other irrelevant reasons (away on vacation, out to lunch, etc.).

The uninformed would draw inference from the signal and, based on their beliefs, determine their optimal amount of avoidance. An equilibrium of this extended model would be the avoidance level of the informed, and the avoidance level of the uninformed as a function of their signal such that each is optimal given their beliefs and their neighbors behavior.

The level of avoidance in the neighborhood increases when the uninformed are allowed to conduct inference and change their behavior. Obviously, their avoidance level will be higher because previously it was zero. But the avoidance level of the informed also increases because avoidance obeys strategic complementarities; the marginal benefit of avoidance increases the more one's neighbors practice avoidance.

Conclusion

This paper studies the effect of criminal registry notification policies on neighborhood behavior. Informed neighbors practice costly avoidance measures in order to protect themselves, which generates negative externalities for the community. The negative externality of avoidance can be counterbalanced by the positive externality of notification: notified neighbors are more difficult for the criminal to successfully victimize, generating deterrence effects for the entire community. Whether notification is optimal or not depends, in part, on the magnitude of these two effects.

The main results of this paper highlight the complementary nature of penalties and notification. In particular, Proposition 1 shows that there always exists a severe enough penalty to make notification optimal. The more people know of the criminals' existence, the more criminals are deterred and the safer it is for neighbors to remain outdoors. Conversely, Proposition 2 shows that when penalties are lenient the government maximizes welfare by not informing the community of the criminals' existence. Notification leads to costly expenditures on avoidance but insufficient deterrence; community notification cannot act as a substitute for lenient penalties.

In addition to the physical cost of informing the neighborhood, notification policies entail costs for communities because the informed change their behavior. Policy makers need to consider these behavioral costs of notification when conducting cost-benefit analyses of community notification programs.

Appendix for Chapter 1

This appendix describes how equilibrium in the model could be reached by a dynamic process. The model is static, but since it aims to describe the

actions of a large group (continuum) of agents, it is helpful to think how outcomes can result dynamically. The analysis in the main body of the paper focuses on locally dynamically stable equilibria. As I will show, these are the equilibria we would expect to observe as the outcome of a dynamic process.

I will describe the dynamics first for discrete time, and then transform the difference equation to a differential equation for continuous time. Let a_t be the average avoidance of the informed at time t. Suppose that each period a small fraction s of the informed best respond to the average avoidance they observe in the neighborhood. Then the average avoidance level in period $t+1$ is

$$a_{t+1} = (1 - s)a_t + sa_i(a_t)$$

Re-arranging terms to get a difference equation gives

$$a_{t+1} - a_t = s\left[a_i(a_t) - a_t\right]$$

I now let the interval of time between periods shrink to ε. To keep the system consistent as the length of time between periods shrinks I also scale the fraction of the informed who best respond by ε.

$$a_{t+\varepsilon} - a_t = \varepsilon s\left[a_i(a_t) - a_t\right]$$

Dividing both sides by ε gives an expression of the time derivative of avoidance

$$\frac{da_t}{dt} = s\left[a_i(a_t) - a_t\right]$$

This equation says that the average avoidance level increases as time progresses if the best response to avoidance level a is above a. This makes intuitive sense. If the agents who change their avoidance level are best responding by choosing a level greater than the neighborhood average, they are bringing up the average. The opposite is true if the best response is below the neighborhood average.

32

We can now see why locally dynamically stable equilibria are characterized by best best response curves which intersect the 45° line from above. This type of equilibrium is represented by avoidance levels x and z in Figure 1.1. Suppose the average avoidance level is slightly less than x. Then the agents who best respond choose an avoidance level between the average and x, pushing the average closer to the equilibrium. If the average is instead slightly above x, then the best response is below the average but above the equilibrium, resulting in a new average avoidance level even closer to the equilibrium.

On the other hand, locally dynamically unstable equilibria have best response curves which intersect the 45° line from below. This can be see by looking at avoidance level y in Figure 1.1. If the average avoidance level is slightly below y, then the share of agents who best respond to this avoidance level choose an avoidance level even less than y. This drags the average level down, away from the equilibrium y. If the average avoidance level is slightly above y, then the best response is greater still, again pulling the average away from the equilibrium.

Then starting from any distribution of initial avoidance levels, the resting point of the dynamic system will be a dynamically stable equilibrium. We would never expect the system to rest at an unstable equilibrium because whenever the average avoidance level gets close to the unstable equilibrium, the best response dynamics push the average away from the equilibrium. For this reason, I focus only on locally dynamically stable equilibria in the paper. These are precisely the equilibria where the best response curve intersects the 45° line from above.

Chapter 2

The Economics of Performance Ratings

Introduction

In many economic environments, individuals' excellence is publicly rated: top chef, grandmaster in chess, employee of the month, etc. At the same time, the rating serves as the *de facto* currency that individuals strive to earn. Ratings therefore serve a dual purpose: they provide information and incentives. This paper studies how well ratings can accomplish these two tasks.

We develop a continuous time model of an agent whose performance is publicly rated. Ratings update continuously with new performance score realizations, while the individual's talent also evolves. The innovation is to allow for the agent to exert costly effort in order to improve his performance. Thus the problem becomes one of moral hazard; talent is unknown and only the agent knows how much effort was used. Since the performance score is stochastic, a good showing could be the result of high effort, high talent, or sheer luck.

Inspired by the characteristics of real-world rating systems, ratings in our model update in a linear fashion. That is, tomorrow's rating is a linear combination of today's rating and the most recent performance score. The

rating update rule is similar to a Kalman-Bucy filter. In a filtering problem, the "gain" (i.e. the weight placed on new information) adjusts over time in order to minimize the mean squared error of filtering. In our problem, the gain is chosen (and fixed forever) by the rating system designer in order to achieve some goal of providing incentives and information.

We study a model in which utility over ratings is linear and the cost of effort is quadratic. We show that by increasing the gain, that is, placing more emphasis on new information in the rating update process, the rating system designer can always elicit more effort. In contrast, there exists a finite accuracy-optimal gain; any deviation from this gain erodes the quality of the signal contained in the rating. If the gain is too low, ratings do not incorporate new information quickly enough; too high, and information contained in the old rating is wasted.

We then study the rating system designer's problem. We show that his optimal gain is always larger than the accuracy-optimal gain. If the gain is set below the accuracy-optimal level, the rating system designer can elicit greater effort *and* provide better information by increasing the gain. Further, we prove that the more the rating system designer cares about effort, the larger is his optimal gain. That is, ratings are more sensitive to new information when the rating system designer cares more about effort.

Background

Rating systems abound in practice. The Elo rating system deserves particular mention because it provided much of the original motivation for this paper. Arpad Elo invented his eponymous rating system for the United States Chess Federation in 1960, and today Elo ratings are used by multiple chess federations, professional tennis, golf, baseball and soccer, as well as various computer games. Elo ratings are very similar to both the Kalman filter and

our model in that ratings update in a *linear* fashion. The change in a player's rating is equal to a constant (called the "K-factor" in Elo ratings and the "gain" in the Kalman filter) multiplied by the difference between realized and expected performance. Our framework allows us to study the optimality of the Elo ratings as used by various chess organizations. We present evidence these ratings are sub-optimal; the K-factor is set too low. This means that Elo chess ratings fare poorly on both the incentive and information fronts.

Performance ratings are also prevalent outside of the realm of sports and games. One particularly well-known example is the Michelin Guide rating system for restaurants. The number of stars each restaurant receives acts as a signal of quality to the consumer. But it also acts as an incentive for chefs to expend effort because higher effort produces better food, which improves the rating, which leads to greater demand in the future.

Another practical example of performance ratings comes from online commerce websites such as eBay, where trading partners rate each other on the quality of their transaction. These ratings serve as signals for future potential trading partners: a higher rating for a seller signals a higher probability of a successful transaction to the buyer. At the same time, a higher rated seller might be able to command a higher price because buyers are less concerned about theft. This generates incentives for sellers to make each transaction as trouble-free as possible.

Ratings differ from rankings in that ratings are a cardinal concept, while rankings are ordinal. Rankings can be derived from any rating system: the highest rating is ranked number one, the second-highest rating is ranked number two, etc. College football, for instance, maintains ratings throughout the season, but the national championship game is played between the two teams with the highest rating at the end of the regular season.

Related Literature

This paper is an extension of Holmstrom [1999].[12] Both papers study how today's effort affects current and future beliefs about ability, and how these beliefs in turn generate incentives for effort. In Holmstrom [1999], wages depend on beliefs about the agent's ability, which are derived in a Bayesian fashion from observing past output. Greater effort today can favorably alter beliefs about the agents ability, but the magnitude of this effect is governed by the rules of Bayesian updating.

In our paper, the agent derives utility from his rating. In contrast to Holmstrom [1999], ratings may update in a non-Bayesian manner. For example, today's rating may depend solely on today's performance; information embodied in the old rating can be completely discarded. Alternatively, ratings may be very insensitive to recent performance and slow to adjust. We study how both the incentives for effort and the information content of ratings depend on the way in which ratings update.

In our model, ratings act as a store of the agent's reputation; the rating encompasses all known information about the agent's ability. There is a rich literature on reputation in repeated games that is primarily concerned with when and how agents can develop and maintain reputations (Kreps et al. [1982], Kreps and Wilson [1982], Milgrom and Roberts [1982], Mailath and Samuelson [2001]).

In the context of the chain store paradox model, Fudenberg and Levine [1992] show that when there is uncertainty over the long-run player's type and imperfect monitoring of his actions, he can achieve an expected average payoff arbitrarily close to what he could get if public commitment was credible. In contrast, Cripps et al. [2004] show that this result is ephemeral. They show that the long-lived player can maintain a permanent reputation for playing a

---

[12]In fact, our model subsumes the steady state of Holmstrom (1999). While Holmstrom [1999] deals with dynamics along the path to the steady state, we do not.

strategy in a game with imperfect monitoring only if that strategy eventually plays an equilibrium of the corresponding complete-information game. If the commitment types' strategies are not equilibria for the normal (i.e. utility maximizing) type long-lived player in the complete-information game, then in any Nash equilibrium of the incomplete-information game the short-lived players almost surely learn when the long-lived player is normal. For this result it is crucial that the short-lived players observe all past signals. In contrast, rating systems are often "forgetful"; that is, ratings discard old information.

Ekmekci [2010] studies the ability of rating systems to act as reputation devices. He shows that when the long-lived player's type is unknown, there exists a finite (i.e. forgetful) rating system and equilibrium of the resulting game such that the long-lived player's payoff is almost his Stackelberg payoff after every history. The focus of Ekmekci [2010], therefore, is to show that a rating system exists which satisfies certain properties. The focus of our paper is to study how the structure of the rating system affects the information and incentive content of its ratings.

While the settings and payoff functions are different, our result that optimal ratings adjust more quickly than Bayesian estimates of ability is similar to the main result in Dubey and Geanakoplos [2010]. They show that optimal grading schemes in games of status (e.g. a classroom in which students care only about their relative grade) trade-off informativeness for incentives. By grouping different performances into one grade (i.e. an "A" grade for all scores of 90, 91, . . . , 100, a "B" for all scores of 80, 81, . . . , 89, etc.), information is lost but greater incentives for effort are created. We similarly show that by over-weighting recent performance relative to the accuracy-optimal rating system, the rating system designer discards useful information contained in the old rating but creates greater incentives for effort.

Model

For simplicity, there is only one agent whose performance is rated. Time is continuous. The agent receives utility from ratings Utility is increasing and linear over ratings, and is written $u(r) = A + Br$ with $B > 0$.

The agent can use costly effort in order to improve his performance score. The cost of effort is quadratic: using effort level ε costs $\varepsilon^2/2$. The agent's discount rate is δ, and we write $\beta = \delta/(1-\delta)$ as the interest rate.

The agent's ability, which is unobserved, evolves in a mean-reverting manner. Formally, ability obeys an Ornstein-Uhlenbeck process:

$$dA_t = \theta(\mu - A_t)dt + \sqrt{L} \ dW(t) \ , \tag{9}$$

where $W(t) \sim N(0,t)$ is a Weiner process.

At each moment in time the agent produces output. Let the stock of output at time t be S_t. The performance score is the change in this stock of output; that is, $Q_t = dS_t$. The performance score is equal to ability plus effort plus noise:

$$dS_t = (A_t + \varepsilon) \, dt + \sqrt{M} \ dZ(t) \tag{10}$$

where ε is the agent's choice of effort and $Z(t) \sim N(0,t)$ is a Weiner process independent from $W(t)$.

The agent is assigned a rating at each moment in time. Provisional ratings are such that $R_0 \sim N(A_0, \sigma_0^2)$. If this were a filtering problem, ratings generated by the Kalman-Bucy filter would be Bayes optimal.[13] The Kalman-Bucy rating update rule would be of the form:

$$dR_t = \theta(\mu - R_t)dt + \frac{\xi_t}{M}(dS_t - R_t dt) \tag{11}$$

---

[13]The Kalman-Bucy filter (the continuous time version of the Kalman filter) is an algorithm for updating Bayesian priors after observing signals of the underlying quantity.

where the mean squared error of filtering $\xi_t = \mathrm{E}[(R_t - A_t)^2]$ solves

$$\dot{\xi}_t = -2\theta\xi_t + L - \frac{\xi_t^2}{M} \tag{12}$$

The fraction ξ_t/M is referred to in the filtering literature as the "gain".

We, however, are not solving a filtering problem. Our rating system designer picks a value γ and ratings update according to the Kalman-Bucy inspired equation:

$$dR_t = \theta(\mu - R_t)dt + \gamma\left(dS_t - \mathrm{E}\left[dS_t|R_t\right]dt\right) \tag{13}$$

In keeping with the filtering literature, we refer to the rating system designer's choice variable, γ, as the gain.

We make the simplifying assumption that the current rating is the only source of information available about the agent's talent. In particular, the agent is completely forgetful. He does not remember his past ratings, performance scores, or effort choices. In order to decide how much effort to use today, he infers his ability based solely on his most recent rating, and then maximizes lifetime expected utility.

Statistical Analysis

In this section we study the statistical properties of the model. We first show that ratings are unbiased signals of ability. This result simplifies the calculation of $E[dS_t|R_t]$ in equation (13).

Then, since ratings are unbiased for ability, we determine which gain minimizes the variance of the distribution of $A_t|R_t$. We call the gain which minimizes the variance of the distribution of $A_t|R_t$ the accuracy-optimal gain. The conditional distribution of ability given a rating is centered at the rating for all gains, but when ratings update using the accuracy-optimal gain, the variance of the distribution is minimized.

In order to show that ratings are unbiased estimates of ability, we first assume that $E[A_t|R_t] = R_t$ and then show that the distribution of $A_t - R_t$ has mean zero. Given our assumption that the agent is forgetful, his choice of effort depends on his current rating alone. Write the equilibrium effort function as $e(R)$. Assuming that ratings are unbiased for ability and that the agent uses the equilibrium effort level allows us to write

$$E[dS_t|R_t] = E[(A_t + \varepsilon_t)dt + \sqrt{M}dZ(t)|R_t] = (R_t + e(R_t))dt$$

Substituting this and equation (10) into equation (13) gives

$$dR_t = \theta(\mu - R_t)dt + \gamma((A_t + e(R_t))dt + \sqrt{M}dZ(t) - (R_t - e(R_t))dt)$$

This equation can be simplified to

$$dR_t = \theta(\mu - R_t)dt + \gamma(A_t - R_t)dt + \gamma\sqrt{M}dZ(t) \qquad (14)$$

Subtracting (14) from (9) gives

$$dA_t - dR_t = -\theta(A_t - R_t)dt - \gamma(A_t - R_t)dt + \sqrt{L}dW(t) - \gamma\sqrt{M}dZ(t) \quad (15)$$

We can write $\sqrt{L}\,dW(t) - \gamma\sqrt{M}\,dZ(t)$ as $\sqrt{L + \gamma^2 M}\,dY(t)$ where $Y(t)$ is a standard Weiner process independent from $W(t)$ and $Z(t)$.[14] Then defining $\delta_t = A_t - R_t$, equation (15) can be written

$$d\delta_t = -(\theta + \gamma)\delta_t dt + \sqrt{L + \gamma^2 M}\,dY_t \qquad (16)$$

---

[14]To see that this is true, note that

$$\mathrm{E}[\sqrt{L}\,dW(t) + \gamma\sqrt{L}\,dZ(t)] = \mathrm{E}[\sqrt{L + \gamma^2 M}\,dY(t)] = 0 \ ,$$

and

$$\mathrm{E}\left[\left(\sqrt{L}\,dW(t) - \gamma\sqrt{M}\,dZ(t)\right)^2\right] = \mathrm{E}\left[\left(\sqrt{L + \gamma^2 M}\,dY_t\right)^2\right] = (L + \gamma^2 M)dt \ .$$

Since $\sqrt{L}\,dW(t) + \gamma\sqrt{M}\,dZ(t)$ and $\sqrt{L + \gamma^2 M}\,dY(t)$ are both Gaussian, they are identically distributed because their first two moments are the same.

This last equation describes the evolution of an Ornstein-Uhlenbeck process. To see this more clearly, write (16) as

$$d\delta_t = (\theta + \gamma)(0 - \delta_t)dt + \sqrt{L + \gamma^2 M} \; dY_t \tag{17}$$

The first two moments for δ_t are

$$E[\delta_t] = 0 \;\; \text{and} \;\; E[(\delta_t)^2] = \frac{L + \gamma^2 M}{2(\theta + \gamma)}$$

But since $\delta_t = A_t - R_t$, we have

$$A_t | R_t \sim N \left(R_t, \frac{L + \gamma^2 M}{2(\theta + \gamma)} \right)$$

These calculations show that if we start with unbiased ratings, we end up with unbiased ratings. Since the provisional ratings are assumed to be unbiased ($R_0 \sim N(A_0, \sigma_0^2)$), ratings at any time t are also unbiased.

Since ratings are unbiased signals of ability, the accuracy-optimal rating system minimizes the variance of the conditional distribution. The following result relates the most accurate rating system to the Kalman-Bucy filter.

Proposition 4. *The accuracy-optimal rating system is identical to the estimator of the steady state Kalman-Bucy filter.*

Proof. The variance for the conditional distribution of $A_t | R_t$ is u-shaped in γ. To find the variance minimizing gain, differentiate w.r.t. γ to get

$$\frac{\partial}{\partial \gamma} \left[\frac{L + \gamma^2 M}{2(\theta + \gamma)} \right] = \frac{M\gamma^2 + 2\theta M\gamma - L}{2(\theta + \gamma)^2}$$

This term equals zero only when the numerator equals zero. Solving using the quadratic formula gives the accuracy-optimal gain, γ^a:

$$\gamma^a = \sqrt{\theta^2 + \frac{L}{M}} - \theta$$

42

Equation (12) gives the law of motion for the mean squared error of the Kalman-Bucy filter. Using the quadratic formula to solve for the steady state mean squared error gives

$$\xi = M\sqrt{\theta^2 + \frac{L}{M}} - \theta M$$

Then since the gain in the filter is ξ_t/M, we see that the steady-state Kalman-Bucy gain is the same as the accuracy-optimal gain, $\gamma^{KB} = \gamma^a$. $\qquad\square$

This proof shows that the accuracy-optimal gain exists and is finite. Any deviation from this gain results in less informative ratings. If the gain is set too low, ratings do not incorporate the new information contained in the performance score realizations quickly enough. If the gain is set too high, information contained in the old rating is wasted.

The accuracy-optimal rating system and Bayesian updating coincide in the steady state. On the path to the steady state, the Bayesian gain adjusts, but this is not how rating systems work. The rating system's gain is fixed from the begining and, after sufficient time, the system approaches the steady state where the conditional distribution of ability given a rating is unchanged over time. This is when the accuracy-optimal rating system and Bayesian updating are identical.

The last important step in the statistical analysis before turning to economic questions is to write the law of motion for ratings in terms of ratings alone. That is, we must get rid of terms involving ability and performance. This equation will be used to determine the agent's optimal effort.

Using the result on the distribution of δ_t, we can re-write equation (14). The steady state distribution of δ_t is Normal with mean zero and variance equal to $(L + \gamma^2 M)/2(\theta + \gamma)$. We can write $\delta_t dt = \sqrt{\frac{L+\gamma^2 M}{2(\theta+\gamma)}} d\hat{Z}(t)$, where $\hat{Z}(t)$ is an independent Weiner process. Then equation (14) becomes

$$dR_t = \theta(\mu - R_t) + \sigma(\gamma)\hat{W}(t) \tag{18}$$

where $\hat{W}(t)$ is an independent Weiner process and

$$\sigma(\gamma) = \gamma \sqrt{\frac{L + \gamma^2 M}{2(\theta + \gamma)} + M} \qquad (19)$$

Equation (18) gives us a stochastic differential equation governing the evolution of ratings in a steady state equilibrium. We will now use this equation to study the agent's problem.

Economic Analysis

We now turn to studying the agent's optimal choice of effort and the rating system designer's optimal rating system. With linear utility and quadratic cost of effort, the model can be solved analytically.

We first fix the rating system and solve for the agent's optimal effort. Once we understand the determinants of effort, we study the rating system designer's problem. The rating system designer has preferences over the accuracy of the ratings and the effort level they induce. We characterize optimal rating systems for different preferences of the rating system designer. The restriction to quadratic cost of effort allows us to describe precisely how optimal rating systems depend on the designer's preferences over providing information and incentives.

Optimal Effort

With the law of motion for ratings given, we can now turn to studying the agent's optimal effort function. Given his current rating, the agent chooses his effort in order to maximize his expected discounted lifetime utility from ratings less the cost of effort. The Hamilton-Jacobi-Bellman equation for the agent with rating r is

$$\beta v(r) = \max_{\varepsilon} \left\{ A + Br - \frac{\varepsilon^2}{2} + \left[\theta(\mu - r) + \gamma(\varepsilon - e(r)) \right] v'(r) + \frac{\sigma^2(\gamma)}{2} v''(r) \right\} \quad (20)$$

44

This equation gives the expected return to the agent having rating r. Think of $A + Br - c(\varepsilon)$, the instantaneous utility from ratings less the cost of effort, as the dividend, and the rest of the right hand side as the expected capital gain. We now show that with linear utility over ratings, the value function is linear and increasing in the agent's rating, and the agent's optimal effort level is constant for all ratings.

Lemma 7. *With linear utility over ratings, the value function is linear and optimal effort is constant.*

Proof. Let $v(r) = C + Dr$. Then $v'(r) = D$ and $v''(r) = 0$. Equation (20) can be re-written as

$$\beta(C + Dr) = \max_{\varepsilon} \left\{ A + Br - \frac{\varepsilon^2}{2} + [\theta(\mu - r) + \gamma(\varepsilon - e(r))] D \right\}$$

The first order condition for optimal effort is $\varepsilon = \gamma D$. The second order condition is met because the cost of effort function is convex. We find values of C and D satisfying the above value function to show that optimal effort is independent of rating. Let e^* denote optimal effort.

Then the equilibrium value function (i.e. when the rating system correctly predicts and accounts for optimal effort, so that $\varepsilon - e^* = 0$) solves

$$\beta(C + Dr) = A + Br - \frac{(e^*)^2}{2} + [\theta(\mu - r)] D \tag{21}$$

Taking the derivative in r and solving for the slope of the value function gives $D = B/(\beta + \theta)$. This means that optimal effort is $e^* = \gamma B/(\beta + \theta)$.

Making this substitution and then solving equation (21) for the constant C gives

$$C = \frac{1}{\beta} \left(A - \frac{1}{2} \left(\frac{\gamma B}{\beta + \theta} \right)^2 + \frac{\theta \mu B}{\beta + \theta} \right)$$

$\square$

The analysis with linear utility over ratings is very similar to the model in Holmstrom [1999]. In fact, when ratings are updated in a Bayesian manner (i.e. when $\gamma = \gamma^{KB}$), the models are identical in the steady state. While posterior beliefs over ability are formed in a Bayesian manner in Holmstrom (1999), our formulation allows for non-Bayesian updating of ratings. The rating encompasses all known information about the agent. Proposition 4 shows how the information content of ratings depends on the gain. The following result describes how the incentive content of ratings depends on the gain.

Proposition 5. *With linear utility over ratings, effort increases and value decreases in the gain. Further, effort increases in the slope of the utility function and decreases in the strength of mean reversion.*

Proof. From Lemma 7, optimal effort is $e^* = \gamma B/(\beta+\theta)$, and so is increasing in the gain. If the strength of mean reversion (θ) increases, optimal effort falls.

Value decreases in the gain because the constant C decreases as the gain increases because a larger gain leads to larger expenditures on effort for all ratings. $\square$

Value decreases when the gain increases because the agent optimally increases his effort level but he is in no way compensated for his effort. The rating system expects he will work harder and simply adjusts its inference on performance accordingly. The equilibrium law of motion of ratings is independent of the gain, even though the agent tries harder when the gain is larger. This is what Holmstrom [1999] very appropriately terms a "rat race".

In Proposition 4 we showed that the accuracy-optimal gain is a function of the strength of the mean reversion of ability and the underlying noise terms. Any gain above or below γ^a results in less accurate ratings, in the sense that the variance of the distribution of $A|R$ is larger the more γ deviates from

γ^a. In contrast, Proposition 5 shows that larger gains *always* induce greater effort.

The more quickly utility rises in rating the larger is the agent's optimal effort. This is because each unit of effort has greater return in utility. As the rate of mean reversion increases, optimal effort falls because the return to effort becomes more ephemeral.

Optimal Rating Systems

The rating system designer has preferences over the (steady state) accuracy of ratings and the effort his rating system induces. His only task is to set the gain and then let the system flow uninterrupted.

Our main result in this subsection establishes the tradeoff faced by the rating system designer. If the designer's preferences are such that more accuracy and effort are always better than less, the optimal gain is larger than the accuracy-optimal gain. The ratings are therefore less informative than those formed with a Bayesian rating system, but the equilibrium effort level is greater.

Since the optimal gain is set above the accuracy-optimal gain, ratings have more noise than Bayesian estimates of ability. This is because the ratings are relatively more sensitive to recent performance. This sensitivity is what gives the agent greater incentives for effort; as the gain increases, so does his ability to influence his rating via effort.

The rating system designer's preferences are assumed to place η weight on effort and $1-\eta$ weight on accuracy. We write the rating system designer's utility as

$$
\begin{aligned}
u\left(e^*, E[(A-R)^2]\right) &= \eta \, e^* - (1-\eta)E[(A-R)^2] \\
&= \eta\frac{\gamma B}{\beta+\theta} - (1-\eta)\left(\frac{L+\gamma^2 M}{2(\theta+\gamma)}\right)
\end{aligned}
$$

47

The rating system designer's utility is increasing in the equilibrium level of effort; he wants his agents to try hard. But he pays a penalty proportional to the variance of the distribution of ability given a rating; all else equal, he wants accurate ratings.

The following proposition characterizes optimal rating systems. If the rating system designer cares about both the accuracy of the ratings and the incentives for effort they provide, he trades off accuracy for effort.

Proposition 6. *With linear utility over ratings, optimal rating systems place at least as much emphasis on recent performance as the Bayesian rating system. Further, the emphasis placed on recent performance increases in the rating system designer's preference for effort.*

Proof. When $\eta = 0$, the rating system designer cares only for the accuracy of the ratings. In order to maximize his utility, he uses $\gamma = \gamma^{KB}$ because this is the gain that minimizes the variance of the distribution of ability given a rating.

For $\eta \in (0, 1)$, the first order condition for the rating system designer's choice of γ is

$$\eta \frac{B}{\beta + \theta} = (1 - \eta) \frac{M\gamma^2 + 2M\theta\gamma - L}{2(\theta + \gamma)^2} \tag{22}$$

Since the variance of the distribution of $A|R$ is u-shaped in γ, the fraction on the right hand side of (22) is negative if $\gamma < \gamma^{KB}$, so the optimal gain must be greater than γ^{KB} when $\eta > 0$.

To show that the emphasis placed on recent performance increases when the rating system designer cares relatively more about effort, it suffices to show that the γ which satisfies the first order condition in (22) increases in η. Let γ' be optimal for η'. Then, for $\eta'' > \eta'$, the two sides of equation (22) evaluated at η'' and γ' cannot be equal; the left hand side must be larger. Since γ' is optimal for η' and we know the rating system designer would never

48

optimally set $\gamma < \gamma^{KB}$, the fraction on the right hand side is increasing in γ. In order to restore equality, the gain must increase to some $\gamma'' > \gamma'$. $\qquad\square$

An important implication of this result is that optimal ratings always adjust more quickly than Bayesian ratings. If ratings are adjusted more slowly than Bayesian ratings, the rating system designer could increase the gain and generate more accurate ratings which also induce greater effort. At the Bayesian optimal gain, the cost of increasing the gain a small amount in terms of lost accuracy is second order, while the benefit in terms of increased effort is first order. Therefore, so long as the rating system designer cares for both effort and accuracy, he sets the gain strictly larger than the Bayesian optimal level. This generates less accurate ratings, but implements greater effort.

The more the rating system designer cares about effort, the larger is the optimal gain. If the rating system designer does not care about accuracy at all (i.e. $\eta = 1$), there is no optimal rating system. He would want to keep setting the gain larger and larger, and there is no upper bound. The problem is that effort increases in the gain, and the rating system designer could theoretically set the gain arbitrarily large. This generates both large rewards for exceeding expectations, and large penalties for failing to meet them.

Conclusion

This paper presents a model in which the incentive and information aspects of performance ratings can be studied. We show that ratings cannot be both statistically accurate and provide maximal incentives for effort. Accuracy-optimal rating systems update in a manner similar to the Kalman-Bucy filter, efficiently blending information contained in the old rating with new information contained in the recent performance score realization. We show that

the rating system designer can always elicit greater effort by placing more emphasis on the agent's recent performance score. When the rating system designer prefers more effort and information to less, he faces a tradeoff: ratings can elicit greater effort only at the cost of being less informative.

The world of competitive chess, as mentioned in the introduction, provides one practical application for our results. Interestingly, Sonas [2002] finds that the K-factor used by the World Chess Federation is *less* than the accuracy-optimal level.[15] Currently, the K-factor for players with ratings above 2400 is 10. Sonas claims that a K-factor of 24 maximizes the accuracy of the ratings in the sense that ratings with this K-factor provide the most predictive power. Supposing this claim is in fact true, our main result suggests that the World Chess Federation could increase both the information and the incentives generated by its ratings by increasing the K-factor. If greater effort leads to more exciting and creative play, these results indicate that the World Chess Federation is not using optimal ratings.

---

[15]Recall that the K-factor in chess ratings is analogous to the gain in our model.

Chapter 3
When Should Governments Reveal Criminal Histories?

Introduction

Governments often possess information about their citizens' criminal histories. *When* this information is released to the general public can affect the criminals' incentives for good behavior. This paper investigates the effects of the timing of the government's notification policy and determines optimal policies when the government is concerned with both the crime rate and the lost productivity of the criminal population.

Practical examples abound: should someone who was caught stealing once be required to reveal this fact on future job applications? When should a sex offender's identity be revealed to his neighbors? Publicly-known criminals may have difficulty getting jobs and might be ostracized from their communities. The penalty incurred by notification helps to deter first-time offenders but is costly because of the loss of productive members of society. Further, public notification after the first conviction diminishes the government's ability to deter repeat offenders. Since publicly known criminals are already underemployed and or ostracized, the government cannot threaten them with any additional penalties in the labor market.

I construct a discrete time, infinite horizon model where criminals observe opportunities each period and decide whether or not to commit the crime. The government chooses between a "strict" policy which publicly discloses criminal histories immediately upon first conviction, or a "lenient" policy which temporarily conceals this information and only reveals it after the criminal's second conviction. When the government is strict, there are only two types of criminals: the publicly known and the unknown. When the government is lenient, there are again the publicly-known and the unknown criminals, but there is also a third, intermediate type, with one prior conviction who is known to the authorities but not the public as a criminal. I study the behavior of criminals in these different situations, compute the steady state distribution of criminals under each notification policy, and then compare the overall crime rates.

I show that the lenient policy minimizes the crime rate when the following three conditions jointly hold: public notification has a large negative effect on wages, a small positive effect on detection probabilities, and criminals have a long expected lifetime. In order to prove this result, I show in a sequence of lemmas that criminals without any convictions always commit more crimes when the government is lenient, and that the behavior of publicly known criminals is independent of notification system. The behavior of the criminals with one conviction under the lenient system is therefore pivotal to the ordering of the crime rates. When public notification has a large negative effect on wages and a small positive effect on detection probabilities, criminals with one conviction under the lenient system commit few crimes because the expected benefit is small. When criminals also have a long expected lifetime, the steady state population consists of a sufficiently large fraction of well-behaved criminals with one conviction under the lenient system, leading to a low overall crime rate.

In addition to the crime rate, society may care about the productivity of

its citizens. Public notification of criminal histories can be harmful because those who are known as criminals have a hard time finding gainful employment and being a productive member of society. This provides an additional justification for leniency on the part of the government. Sometimes otherwise law abiding citizens are presented with an opportunity to commit a crime that they find difficult to ignore. If they are caught and the government's notification policy is strict, the public is immediately informed of their deed and the criminals cannot find gainful employment.

I show that there are always less publicly known criminals under the lenient notification system than under the strict system. This means that a larger fraction of the criminal population is employed under the lenient system, so productivity is higher. The more society cares about the productivity of its citizens, the more appealing the lenient notification system becomes.

Related Literature

Economists have been interested in how criminals respond to incentives since Becker [1968]. More recently, economists have investigated the deterrence effects of police and incarceration [Levitt [1997], Di Tella and Schargrodsky [2004]], conditions in prison [Katz et al. [2003]], and gun ownership [Lott [1998], Mialon and Wiseman [2006]]. These papers validate Becker's seminal work by showing that criminals respond to incentives in a rational manner. My research adds to this literature by studying an additional crime fighting tool available to the government: public notification policies.

The most well-known public notification policy is Megan's Law, a federal law in the U.S. which requires states to make information about sex offenders available to the public. But public notification policies also exist (or are currently being developed) for criminals convicted of homicide, gang crimes, animal abuse, elderly abuse, drug dealing and manufacturing, and drunk

driving [Goode [2011]]. As they are politically popular and becoming more prevalent in practice, it is important to understand when and how public notification policies can be helpful.

Neto [2006] studies how the public availability of criminal records affects behavior. He constructs a model in which agents can commit crimes in each of two periods. Agents are heterogeneous in their ability and their taste for criminal acts, and those who commit crimes in the first period learn about their proficiency. If the government makes criminal records publicly available, a criminal's payoff for good behavior decreases because a fraction of employers perform background checks and will not hire anyone with a record.

My work builds on Neto [2006] by exploring more fully how agents' dynamic incentives are affected by the government's notification policy. The government can control not just if, but also when, criminal records become publicly available. The model in Neto [2006] captures the intuition that, by making criminal records publicly available, criminals without any prior convictions face greater deterrence, but those with a record commit more crimes because their payoff to honest work decreases. In addition to this type of intuition, my model shows how the timing of the government's disclosure affects criminal behavior. If the government gives the criminals an extra chance before disclosing their criminal past, criminals without any prior convictions are less likely to be deterred while those with one conviction are more likely to be deterred. I also show how concern for the lost productivity of publicly known criminals can affect the government's choice of when to disclose criminal histories.

Economists interested in crime have also studied the optimal penalty size as a function of the number of convictions. While the punishment for many crimes is higher for recidivists than first time offenders in practice, the literature does not always find increasing penalty functions to be optimal.

Polinsky and Rubinfeld [1991], Rubinstein [1979], and Rubinstein [1980] all show that, under different model specifications, increasing penalties may be, but are not always, optimal. In contrast, Chu et al. [2000] show that if society has a distaste for wrongful convictions of the innocent, the penalty for repeat offenders is always greater than the penalty for first time offenders.

While public notification of criminal histories is a *de facto* penalty on convicted criminals, it is of a different sort from the penalties studied by the optimal penalty function literature. In the studies surveyed here, penalties are "one-dimensional" and are construed as any and all actions or policies which lower the utility of a convicted criminal. In contrast, public notification policies are seen as a way to keep criminals "behind bars" even after they are released from prison. That is, an informed community can help monitor convicted criminals and make it more difficult for them to get away with crimes. Public notification penalizes criminals not just by hurting their job market prospects, but also by making it less likely to commit crimes without detection.

This paper contributes to the literature on penalty functions by studying how the implicit penalty associated with public notification, combined with statutory penalties, determines the overall shape of the penalty function. Holding statutory penalties constant, the strict policy generates a decreasing penalty function because the immediate notification leads to a large penalty for first time offenders but does not affect the penalty for repeat offenders. The lenient policy generates a first increasing and then decreasing penalty function because criminals are penalized by public notification only after the second conviction. Depending on its objectives, the government will sometimes use relatively small penalties on first time offenders in order to hold a larger penalty in reserve to deter repeat offenders.

Economists also empirically study the effects of criminal registration and public notification policies. Prescott and Rockoff [2011] presents evidence

that notification policies can deter crime by first time offenders and increase the crime rate for known criminals. They also find evidence that the detection rate for intermediate types (i.e. criminals known to the police but not the public) is larger than the detection rate for those without any prior convictions. If this detection rate is high enough, their findings and the theoretical results in this paper imply that society might be better off temporarily concealing information. Governments could adopt a two-step program where first time offenders are registered with the police, but their identities are made public only if they commit another crime. In such a case, separating the registration and notification policies generates a lower crime rate.

My results indicate that empirical research might productively focus in two areas. First, how do detection probabilities depend on the transmission of information from governments to citizens? How much does the detection probability increase as the criminal transitions from unknown, to known only to the police, to a publicly known criminal? The answer to this question can help us predict how useful notification policies will be for different crimes.

Secondly, how much do publicly known criminals suffer in terms of lost wages? Several studies report mixed results on the effect of arrests on employment and earnings [Grogger [1995], Holzer et al. [2004], Waldfogel [1994]], but none of these studies deals specifically with public notification policies. If the effect of public notification on employment is large and negative, my results indicate that governments should consider temporarily delaying public notification and allowing criminals a second chance before outing them to their neighbors.

Model

The initial mass of criminals is one, and they face a constant probability of death equal to $1 - \beta$. Each period a mass $1 - \beta$ of new criminals without

any criminal history enter the population to replace the dead, so their mass stays constant at one.[16]

Time is discrete and the horizon is infinite. Each period criminals encounter a criminal opportunity $\theta \sim U[0,1]$, iid, and must decide whether to behave or to commit the crime. If the agent commits the crime he receives utility of θ. If the criminal is caught committing the crime, he pays a penalty of τ. Thus, a criminal committing a crime who is not caught receives θ, while a criminal who is caught receives $\theta - \tau$.

There are many possible ways in which criminal histories, known by the authorities, can be conveyed to the public. The two information systems studied here are denoted S for "strict" and L for "lenient". Let $h \in \{0, 1, \ldots\}$ be the criminal history where $h = k$ stands for k convictions; the public does not observe h. The information given to the public is a "criminal rating" of $r = g$ or $r = b$ for "good" and "bad", respectively. Each information system has a message function m which maps history to ratings.

Under the strict system, agents get good ratings only if they have a perfect history. That is:

$$m_S(h) = \begin{cases} g & \text{if } h = 0; \\ b & \text{if } h \geq 1. \end{cases}$$

This system perfectly reveals the government's information by identifying everyone convicted of a crime.

Under the lenient system, criminal status is temporarily concealed from the public. I study the case where:

$$m_L(h) = \begin{cases} g & \text{if } h \leq 1; \\ b & \text{if } h \geq 2. \end{cases}$$

---

[16]This normalization of the mass of criminals is purely for convenience. Since the notification policy does not affect the size of the criminal population and the relevant comparison is across notification policies, the size of the criminal population will not affect the ranking notification systems.

The authorities know the agent is a criminal if $h = 1$, but the lenient system hides this fact from the public. Only if they are convicted of another crime will their public rating change to $r = b$.

Independent of criminal activity, agents collect any income they earn in the labor market at the start of the period.[17] Labor income is a function of public rating where

$$\ell(r) = \begin{cases} 0 & \text{if } r = b; \\ \ell & \text{if } r = g. \end{cases} \tag{23}$$

The probabilities of detection are given by $\delta_S(h)$ and $\delta_L(h)$. Agents without any history of crimes face the lowest probability of detection because neither the public nor the authorities know they are criminals. Agents with a public bad rating face the highest probability of detection because everyone knows they are criminals and can keep a close eye on them. Agents with $h = 1$ and $r = g$ under the lenient system face an intermediate probability of detection because the authorities, but not the public, know they are criminals. For $0 < \underline{\delta} < \delta_1 < \overline{\delta} < 1$, the detection probability functions are

$$\delta_S(h) = \begin{cases} \underline{\delta} & \text{if } h = 0; \\ \overline{\delta} & \text{if } h \geq 1. \end{cases}$$

$$\delta_L(h) = \begin{cases} \underline{\delta} & \text{if } h = 0; \\ \delta_1 & \text{if } h = 1; \\ \overline{\delta} & \text{if } h \geq 2. \end{cases}$$

Criminal Behavior

The ultimate goal of this paper is to answer the question: "When should the government reveal criminal status?" In order to answer this question,

---

[17]As noted, the interpretation could instead be utility from maintaining good social relationships.

I start by studying how criminals' behavior depends on the timing of the government's disclosure policy.

As previously stated, I restrict attention to two policies for the government: the strict policy which reveals criminal status upon first conviction, and the lenient policy which temporarily conceals the criminal's past and reveals it only upon the second conviction. I begin the analysis by studying the strict system because there are only two types of criminals, and the mechanics for studying the lenient system are very similar. Once we know the behavior of individuals in different situations under each notification system we can derive the steady state crime rate for society.

The Strict System

At the start of a period, agents collect any labor income $\ell(m_S(h))$ and observe θ and then decide whether to behave or commit the crime. If the agent behaves, tomorrow's history equals today's and the continuation value is known with certainty. Committing the crime has an expected immediate return of $\theta - \delta_S(h)\tau$ and the continuation value depends on whether or not the agent is caught. If he is not caught, tomorrow's history equals today's while getting caught entails a longer criminal record and a possibly worse public rating.

Using primes to represent the value of a variable next period, I write the value function for an agent with history h and opportunity θ under the strict system as

$$v(h, \theta) = \max \{ \quad \ell(m_S(h)) + \beta E_\theta[v(h, \theta')],$$
$$\ell(m_S(h)) + \theta - \delta_S(h)\tau + \beta E_{h,\theta}[v(h', \theta')]\}$$

The criminal decides between committing the crime or not. By not committing the crime (the first argument inside the max operator), the criminal

59

forgoes θ, but guarantees that his public rating does not change. By committing the crime (the second argument inside the max operator), the criminal guarantees himself θ, but runs the risk of getting caught and worsening his public rating.

As makes intuitive sense, the following lemma shows that criminals' expected lifetime utility decreases in their criminal history and increases in the attractiveness of the opportunity at hand.

Lemma 8. *The value function for criminals under the strict notification system is decreasing in the length of criminal history and increasing in the attractiveness of the criminal opportunity at hand.*

All proofs will be collected in the appendix.

Publicly Known Criminals   I begin by studying the behavior and value of publicly known criminals. This is done for two reasons. First, we must understand the publicly known criminals' cutoff opportunity in order to determine the overall crime rate. Second, the publicly known criminals' value is needed in order to study the problem of an unknown criminal whose misdeeds could result in a bad public rating.

Under the strict system, agents with $h \geq 1$ have $r = b$ for the rest of their lives. Since they will be publicly known as criminals for as long as they live, their continuation value is independent of their action and the outcome of any crime committed. Their optimal action therefore depends only on the immediate return; they behave as if totally myopic.

If a publicly known criminal chooses to behave he gets an immediate payoff of zero. If he chooses to commit the crime, he gets an immediate expected payoff of $\theta - \overline{\delta}\tau$ because with probability $\overline{\delta}$ he is caught and must pay the penalty of τ. Publicly known criminals commit crimes if and only if the expected immediate return is positive.

Lemma 9. *Publicly known criminal's optimal behavior is defined by a cutoff opportunity $\underline{\theta}_b$ such that the criminal commits crimes if and only if $\theta \geq \underline{\theta}_b$. The cutoff opportunity $\underline{\theta}_b$ is increasing in the detection rate and the penalty if caught. The publicly known criminal's value is decreasing in the detection rate and the penalty if caught.*

That the publicly known criminals behave myopically makes the algebraic derivation of their cutoff opportunity and value function very simple. The cutoff opportunity does not depend on any endogenous objects (such as value) and so finding the cutoff opportunity simply requires solving a 1 equation – 1 unknown "system". I now turn to analyzing an unknown criminal's problem where, unfortunately, such a simplification is not possible.

Unknown criminals   Now I derive the optimal behavior of agents with $h = 0$ under the strict system. The analysis here is complicated by the fact that unknown criminals behave in a forward looking manner. Their optimal behavior is defined by a cutoff opportunity which depends on endogenous objects. In order to solve for the cutoff value and the value function, a system of equations is needed.

Unknown criminals receive wage ℓ at the start of the period (regardless of any criminal activity) because firms believe they might be good citizens. If he behaves, the criminal gets an immediate return of ℓ and the expected continuation value of an unknown criminal because his criminal record will remain unchanged. If he chooses to commit the crime, he gets an immediate return of ℓ plus the expected value of the committing the crime, equal to $\theta - \underline{\delta}\tau$. The continuation value in this case depends on whether or not the agent is caught. With probability $\underline{\delta}$ he will be caught and his continuation value will be that of a publicly known criminal. With probability $1 - \underline{\delta}$ he will not be caught, his public rating will remain the same, and his continuation value will be that of an unknown criminal.

Lemma 10. *Unknown criminals' optimal behavior is defined by a cutoff opportunity $\underline{\theta}_g$ which is increasing in labor income, the detection rate, and the penalty if caught. The ordering of $\underline{\theta}_g$ and $\underline{\theta}_b$ depends on the parameters of the model. The unknown criminal's value is decreasing in the detection rate and the penalty if caught.*

When $\ell = 0$, the unknown criminal's cutoff opportunity is less than that of the publicly known criminal; there are no future wages to lose and the probability of detection is lower so they commit more crimes. As ℓ increases, the unknown criminal's cutoff opportunity could be less than or greater than the known criminal's, depending on the value of $\overline{\delta} - \underline{\delta}$. When $\underline{\delta}$ is close to $\overline{\delta}$, the cutoff opportunity for the unknown criminal is higher; he is more careful deciding which crimes to commit because he doesn't want to get caught and lose future earnings. When $\underline{\delta}$ is small enough, his cutoff is lower because he can get away with more.

Ultimately, criminals' behavior is important only insofar as it determines the overall, steady state crime rate. In the Crime Rates section I will compute the crime rate induced by the strict system, but next I study how adding leniency to the notification policy changes the incentives for the criminals.

The Lenient System

In this section I will discuss the optimal behavior of criminals under the lenient system. The discussion is less detailed because the analysis is mechanically very similar to the previous section. The main difference is due to the intermediate type who does have a criminal history but is allowed to withhold this information from the public.

I first show that the publicly known criminal's problem under the lenient system is the same as under the strict system. Then, the problem for criminals known only to the authorities is similar to that of the unknown criminal

under strict system except the probability of detection is greater (δ_1 instead of $\underline{\delta}$). Lastly, I show that agents with $h = 0$ commit more crimes under the lenient system. To declare that leniency is better for society requires that the intermediate type with a good public rating and one prior conviction have a low crime rate and that criminal types without any convictions do not commit *too* many more crimes when the government is lenient.

Similar to the previous section, write the value function for an agent with history h and opportunity θ under the concealing system as

$$w(h, \theta) = \max\{ \quad \ell(m_L(h)) + \beta E_\theta[w(h, \theta')],$$
$$\ell(m_L(h)) + \theta - \delta_L(h)\tau + \beta E_{h,\theta}[w(h', \theta')]\}$$

As under the strict system, criminals' expected lifetime utility is decreasing in the number of times they have been caught committing crimes and increasing in the attractiveness of the opportunity at hand.

Lemma 11. *The value function under the lenient system is decreasing in the length of criminal history and increasing in the attractiveness of the criminal opportunity at hand.*

The proof is the same as in Lemma 8. The next lemma shows that the behavior and value of publicly known criminals is independent of notification system.

Lemma 12. *Publicly known criminals have the same value and optimal cut-off opportunities under the lenient system as they do under the strict system.*

The government's notification policy has no bearing on the publicly known criminals' decision; their identity has already been disclosed to the public. I use $\underline{\theta}_2$ to represent the cutoff opportunity for a publicly known criminal under the lenient system. Lemma 12 says that $\underline{\theta}_2 = \underline{\theta}_b$.

I now study the problem of criminals who have been convicted of one crime under the lenient system. These criminals are known to the authorities,

but not the public. Their problem is very similar to that of the unknown criminals under the strict system. Each is able to receive a wage because they are not known to the public as criminals, and each will transition to being a publicly known criminal with one more conviction. Lemma 12 implies that the consequences of a conviction are the same for each type of criminal, but those on their last strike under the lenient system are more likely to be caught because the authorities are aware of their identity and hence the probability of detection is higher.

Lemma 13. *Agents convicted of one crime under the lenient system have cutoff opportunity $\underline{\theta}_1$ which is increasing in labor income, the detection rate and the penalty if caught. Importantly, these agents have a higher cutoff opportunity, and hence a lower crime rate, than the unknown criminals under the strict system.*

The proof is in the appendix but, intuitively, agents convicted of one crime under the lenient system and agents without any convictions under the strict system have almost identical problems. The only difference is in the detection rate. Since the cutoff is increasing in the detection rate, agents convicted of one crime under the lenient system have a higher detection rate.

The last lemma on criminal behavior shows that leniency deters less crime by unknown criminals than a strict public notification policy.

Lemma 14. *Agents without a criminal history have higher values and commit more crimes under the lenient system than under the strict system.*

This lemma makes intuitive sense; the difference between the two situations is in the consequences if caught. If the unknown criminal under the lenient system mimicked the optimal behavior of an unknown criminal under the strict system, he would be caught just as often but the consequences would be less dire. His value must be higher using this mimicking strategy.

64

He can do even better, though, by using a lower cutoff opportunity and committing more crimes. Since his continuation value decreases less when he is convicted, the consequences are lower and it becomes worth his while to commit crimes which the agent under the strict system would not.

In summary, known criminals behave the same way under either information system. Agents with $h = 1$ and $h = 0$ have behavior defined by optimal cutoff values $\underline{\theta}_1$ and $\underline{\theta}_0$, resp., which are increasing in the probability of detection, the penalty if caught, and the size of the drop in expected value if caught. Importantly, the order of these cutoff values is $\underline{\theta}_0 < \underline{\theta}_g < \underline{\theta}_1$, meaning that unknown criminals under the lenient system commit more crimes than unknown criminals under the strict system, who in turn commit more crimes than criminals with one conviction under the lenient system. If temporarily concealing criminal status is to induce a lower crime rate, it will be because the higher crime rate of the criminal types without any convictions will be more than offset by the lower crime rate of criminals with one conviction.

Crime rates

So far we have studied the behavior (i.e. crime rate) of each type of criminal. In order to determine the overall crime rate, however, we need to know how many criminals of each type exist. That is, we must find the steady state composition of the criminal population.

I will first derive the composition under the strict system, and then state the analogous results for the lenient system. I use graphs of value functions in order to illustrate when and why each information system minimizes the crime rate.

Population Composition Under the Strict System
The probability a criminal of type $k \in \{0, 1, 2, g, b\}$ gets an opportunity

worth exercising is $1 - \underline{\theta}_k$. Each period the total mass of criminals is one because the mass is initially one and the mass of criminals dying each period is exactly offset the mass entering the model.

Let ξ_g^k be the mass of unknown criminals (i.e. $r = g$) under the strict system at the start of period k. Then fraction $1 - \underline{\theta}_g$ get an attractive enough opportunity and fraction $1 - \underline{\delta}$ of these agents do not get caught and maintain a good rating. The mass of unknown criminals after the "commit-or-not" decision is then $\xi_g^k \left(\underline{\theta}_g + (1 - \underline{\theta}_g)(1 - \underline{\delta}) \right)$. Of these, fraction $1 - \beta$ die, and then a mass $1 - \beta$ new criminals enter without any criminal history, and are therefore unknown criminals. So at the start of period $k+1$ the mass of unknown criminals is $\xi_g^{k+1} = 1 - \beta + \beta \xi_g^k (1 - \underline{\delta}(1 - \underline{\theta}_g))$. The steady state is where $\xi_g^k = \xi_g^{k+1}$. The value which solves this equation is

$$\xi_g = \frac{1 - \beta}{1 - \beta + \beta\underline{\delta} - \beta\underline{\delta}\underline{\theta}_g}$$

Since under the strict system there are only two relevant histories for criminal types, the steady state mass of known criminals is $\xi_b = 1 - \xi_g$. The steady state crime rate under the strict system is then $\rho_S = \xi_g(1 - \underline{\theta}_g) + \xi_b(1 - \underline{\theta}_b)$.

Population Composition Under the Lenient System

The calculations are very similar, so I will simply state the results. The only additional complication is that some agents transition into, and some transition out of, the $h = 1$ history under the lenient system. The steady state mass of unknown criminals and criminals who are known only to the authorities are:

$$\xi_0 = \frac{1 - \beta}{1 - \beta + \beta\underline{\delta} - \beta\underline{\delta}\,\underline{\theta}_0} \ ; \ \xi_1 = \frac{\beta(1 - \beta)(1 - \underline{\theta}_0)\underline{\delta}}{(1 - \beta + \beta\underline{\delta} - \beta\underline{\delta}\,\underline{\theta}_0)(1 - \beta + \beta\delta_1 - \beta\delta_1\underline{\theta}_1)}$$

Since there are only three types of criminals under the temporarily concealing system, $\xi_2 = 1 - \xi_0 - \xi_1$. The steady state crime rate under the lenient system is $\rho_L = \xi_0(1 - \underline{\theta}_0) + \xi_1(1 - \underline{\theta}_1) + \xi_2(1 - \underline{\theta}_2)$.

Which System Minimizes the Crime Rate?

The first main result of this paper is that sometimes leniency in the government's public notification policy can result in a lower crime rate. While those publicly known to be criminals behave the same under both information systems, criminals who are unknown to the public face different incentives. Agents with one conviction under the lenient system can be induced to behave if their probability of detection is close to that of a publicly known criminal (i.e. $\bar{\delta} - \delta_1$ is small) and the labor income they stand to lose if caught one more time is large. Agents without a criminal history commit more crimes under the lenient system, but when criminals have long expected lifetimes, the proportion of criminal types without any criminal record is small, and so they do not contribute much to the overall crime rate.

Proposition 7. *The lenient policy minimizes the crime rate when public notification has a large negative effect on wages, a small positive effect on detection probabilities, and criminals' expected lifetime is long.*

The results from the Criminal Behavior section imply that criminals with one conviction under the lenient system behave better than those without any criminal history. Additionally, criminals without any history under the strict system behave worse than those with one conviction, but better than those with no convictions, under the lenient system. Expressed in terms of cutoff opportunities, $\underline{\theta}_1 > \underline{\theta}_g > \underline{\theta}_0$. Intuitively then, the lenient system results in a lower crime rate if criminals with $h = 1$ make up a sufficiently large portion of the population and their crime rate is sufficiently low relative to criminals with $h = 0$.

In the remainder of this section, I graphically illustrate how ℓ, δ_1, and β affect behavior and the composition of the criminal population. Each figure shows the effects of changing one of the parameters of the model on criminal behavior. In each case, the strict notification system is optimal before the change, while the lenient system is optimal after.

Figure 3.1 illustrates the effect of a wage increase on the behavior of the different criminals. Each panel shows four value functions. The legend underneath the graphs associates line-styles to criminal types, but the value functions are always ordered, from lowest to highest, as: publicly known criminals, criminals with $h = 1$ under L, criminals with $h = 0$ under S, and criminals with $h = 0$ under L. The kink point in the value function represents the criminal's cutoff opportunity: it is the opportunity at which value transitions from being independent of opportunity to increasing in opportunity.

The graphs show that as ℓ increases, the cutoff opportunity for agents with $h = 1$ under the lenient system increases to nearly one; i.e. the crime rate for these agents drops to nearly zero. Since very few crimes are committed by criminals with one conviction under the lenient system, very few transition on to be publicly known criminals. As these criminals have a relatively high crime rate, this further helps lower the crime rate of the lenient system. The increase in the wage also affects the behavior of criminals without any convictions, but the change is less dramatic.

Figure 3.2 shows the effect of increasing δ_1. Initially, δ_1 is close to $\underline{\delta}$, which is why the value functions for criminals with one conviction under the lenient system and criminals with no convictions under the strict system are so close. Since the consequences when caught are the same and the probabilities of being detected are so close, these agents face similar problems.

After the increase, δ_1 is close to $\bar{\delta}$, and the crime rate for criminals with one conviction under the lenient system falls. It takes a very appealing opportunity for these criminals to risk getting caught and transitioning to a publicly known criminal, and thereby losing any future earnings.

The criminals without any convictions under the lenient system also commit fewer crimes because they correctly forecast the lower value associated

Figure 3.1 The Effect of an Increase in ℓ

a) before	b) after
—— unknown criminal	—— $h = 0$ under S
⋯⋯ $h = 1$ under L	⋯⋯ $h = 0$ under L
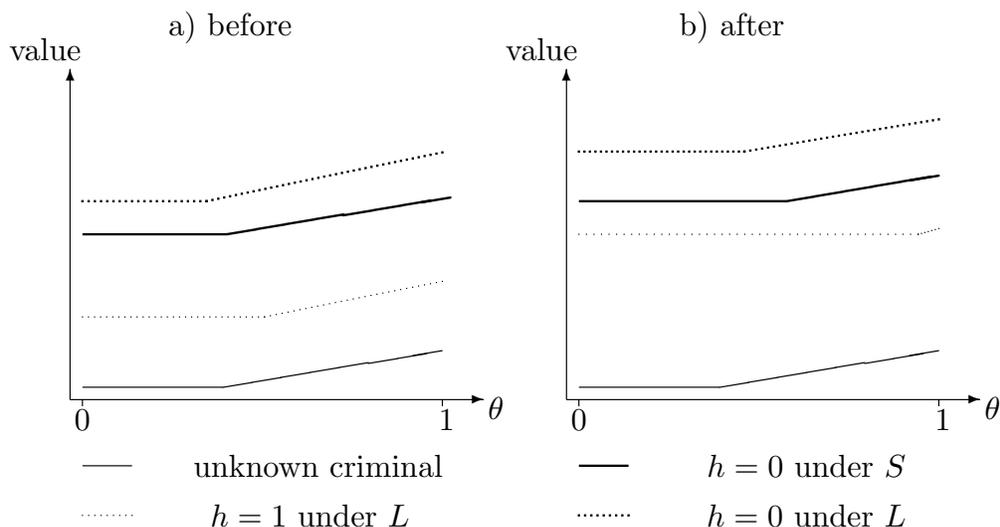
with having one conviction after δ_1 increases. The two types of criminals under the strict system do not change their behavior because δ_1 does not affect either of their problems.

Since $\delta_1 < \bar{\delta}$, public notification is assumed to affect detection probabilities in ways that police work alone cannot. How large this effect is in practice is an empirical question. If having the public know a criminal's identity has little benefit beyond having the police know, then $\bar{\delta} - \delta_1$ is small. All else equal, the lenient policy looks better when public notification has little effect on detection probabilities. If, after the first conviction, the detection probability increases almost to the maximum level, there is little to gain by notifying immediately. Instead, the government can withhold punishing the criminal with public notification for the time being, and use this punishment as a deterrent against future crimes, which are now detected at a high rate.

More than the other parameters, the probability a criminal lives from one period to the next affects the overall crime rate through two channels: the

Figure 3.2 The Effect of an Increase in δ_1

a) before

b) after

——	unknown criminal
········	$h = 1$ under L

——	$h = 0$ under S
········	$h = 0$ under L
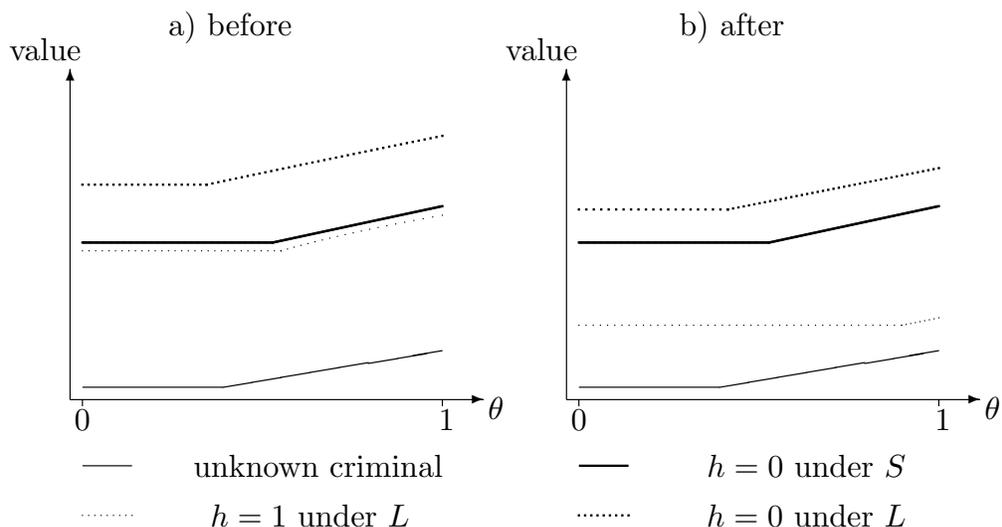
criminals' behavior and the steady state composition of criminal types. To see how the probability of living can affect the steady state composition of criminals directly (i.e. not necessarily through changes in criminal behavior), consider extreme values of β. When β is close to zero, even if the crime rates are high, there are few publicly known criminals in the steady state simply because criminals do not live long enough to be caught committing enough crimes to become publicly known. Alternatively, when β is close to one, the steady state criminal population consists of many publicly known criminals, even if the crime rates are low, simply because criminals live so long that they are eventually caught committing a crime.

Of course, the probability of living also directly affects the behavior of non-publicly known criminals, as the lemmas in the Criminal Behavior section describe. As β increases, criminal types who are not publicly known as criminals expect to live longer and thus require better opportunities in order to commit a crime.

70

Table 3.1 details the composition of the criminal population when β increases from 0.8 to 0.97 for the environment where $\ell = 2$, $\underline{\delta} = 0.1$, $\delta_1 = 0.19$, $\bar{\delta} = 0.3$, and $\tau = .8$. The fraction of the criminal population without a conviction is roughly equal across the two information systems at the low value of β. When the probability of living increases, the fraction of criminals with at least one conviction increases more under the lenient system.

Table 3.1 Criminal Population Composition

	Low β	High β
$h = 0$ under IR:	0.76	0.42
$h \geq 1$ under IR:	0.24	0.58
$h = 0$ under TC:	0.75	0.37
$h = 1$ under TC:	0.17	0.25
$h \geq 2$ under TC	0.08	0.38

In this example, the strict system is optimal for the low β while the lenient system is optimal for the high β. The reason this is the case can be seen in Figure 3.3. After β increases, the probability a criminal with one conviction under the lenient system commits a crime decreases so that a large enough fraction of the criminal population gets "stuck" with one conviction and dies before transitioning to being a publicly known criminal who commits a lot of crimes.

That criminal activity depends on age is one of the most accepted theories in criminology [Hirschi and Gottfredson [1983]]. The vast majority of crimes ar committed by the young. But the length of criminals' lifetimes varies across the type of crime. There is evidence that criminal behavior by sex offenders declines much more slowly in age than for other types of crimes [Hanson [2002]].

In the context of this paper, this evidence suggests that the β for sex

Figure 3.3 The Effect of an Increase in β

a) before

b) after

| | unknown criminal | | $h = 0$ under S |
| | $h = 1$ under L | | $h = 0$ under L |
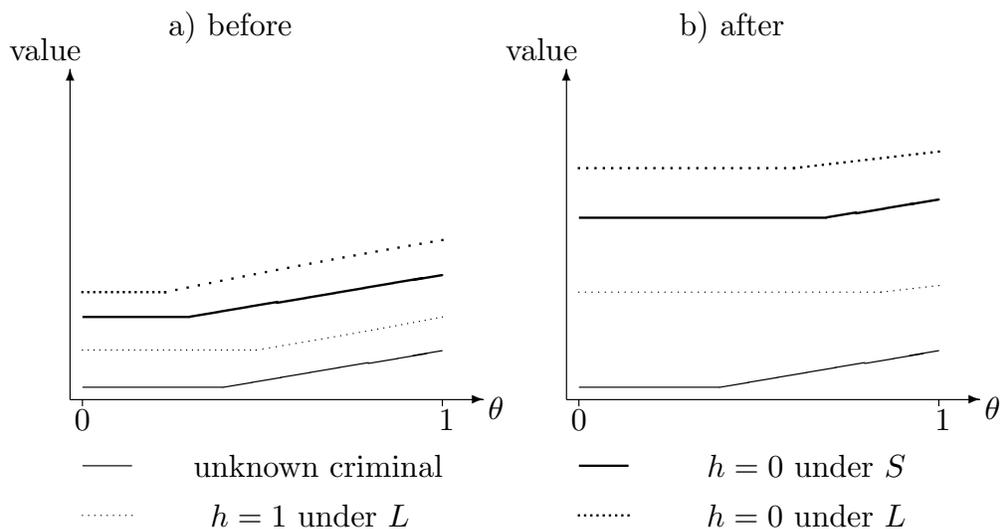
offenders might be larger than the β for other types of criminals. If this is the case, governments might consider leniency in their public notification system. By holding the notification punishment in reserve, once-convicted sex offenders might face strong enough incentives so that they do not commit any more crimes.

These examples help to illustrate that the crime rate depends on both criminals' behavior and the number of each type of criminal in the population. A lenient policy which temporarily conceals criminal status from the general public can minimize the overall crime rate if it provides strong incentives for criminals on their last strike, and these criminals make up a sufficiently large percentage of the population.

Productivity

The threat of public notification acts as a deterrent against criminal acts, but its use is also costly for society. When the government makes a criminal's identity public, the criminal has difficulty finding work in the future. Public notification of criminal histories therefore costs society a productive member.

While the crime rate can be minimized by either the strict or lenient public notification policy depending on the environment, productivity is always higher under the lenient notification policy.

Proposition 8. *The productivity of the criminal population is always larger when the government is lenient.*

The productivity is higher under the lenient system because there are more criminals with good public ratings. That is, they have not yet been publicly outed as a criminal, either because they have not yet been caught or because they have been caught only once. Since the public perceives them as good, they are able to find gainful employment and remain productive members of society.

The fact that it takes two convictions to become publicly known under the lenient system as opposed to just one under the strict system is not enough by itself to prove Proposition 8. The reason is that if the crime rates of criminals with good public ratings induced by the lenient system were very large compared to the crime rate of unknown criminals under the strict system, there would be more publicly known criminals when the government is lenient. The fact that criminals with one conviction under the lenient policy have a higher cutoff opportunity, and hence a lower crime rate, than unknown criminals under the strict policy keeps this from happening.

Proposition 8 shows how public notification can be costly for society. The threat of public notification does act as a deterrent, but the lost productivity of publicly known criminals hurts society. This provides some justification

for leniency on the government's behalf. Rather than hit first time offenders with the full penalty (statutory plus public notification) immediately, the government can hold some of its punishment in reserve and give the criminal a second chance. This type of leniency always leads to a more productive criminal population.

Optimal Public Notification Policies

Propositions 7 and 8 measure the efficacy of public notification policies on two different metrics: minimizing the crime rate and maximizing productivity. In practice governments likely care about both.

If the government doesn't care about the productivity of its criminals at all, then Proposition 7 specifies how to derive the optimal public notification system. The more the government cares about lost productivity, the better the lenient public notification policy looks. In the limit, when the government only cares about productivity, the lenient policy is always optimal.

My results add to the literature on the shape of optimal penalty functions. The government's policy concerning the timing of public notification affects the shape of the penalty function. In my model, with constant statutory penalties, the strict notification system generates a decreasing penalty function and the lenient system generates a first increasing and then decreasing penalty function. The lemmas in the Criminal Behavior section show how the penalty generated by public notification depends on detection probabilities, length of expected lifetime, and wages.

The public notification of criminal histories creates a penalty which can only be used once. Depending on its objectives, the government can sometimes do better by penalizing first time offenders lightly and waiting to use the notification penalty on repeat offenders. This type of policy can minimize the crime rate, but it always leads to greater productivity.

Conclusion

The way in which governments reveal criminal status can affect criminals' incentives for good behavior. This paper studies a strict and lenient notification policy. The strict policy reveals criminal status after the first conviction, while the lenient policy conceals criminal status until the criminal is caught committing a second crime, at which time the criminal's identity is made public. When the government's public notification policy is strict, the unknown (i.e. not yet convicted) criminals face relatively greater incentives to behave. In contrast, under the lenient policy, the unknown criminals face low incentives while the once-convicted criminals who are known only to the authorities face high incentives to behave. The lenient policy minimizes the crime rate when public notification has a large negative effect on wages, a small positive effect on detection probabilities, and a criminal's expected lifetime is long.

In contrast to that qualified result, the lenient policy always maximizes the productivity of the criminal population. When the government is lenient, fewer criminals are publicly identified as such, and more criminals are therefore able to remain productive members of society. Public notification is effective as a deterrent against crime, but it is also costly to society in terms of productivity. An optimal public notification policy balances these costs and benefits according to the government's preference for a low crime rate and its distaste for the lost productivity of the publicly known criminals.

Appendix for Chapter 3 – Proofs

Proof of Lemma 8. Let $\mathcal{V}$ be the space of functions $f : \mathbb{N} \times [0, \theta^{max}] \mapsto \mathbb{R}$. Define an operator

$$\mathbb{T}f(h, \theta) = \max \{ \quad \ell(m_S(h)) + \beta E_\theta[f(h, \theta')],$$
$$\ell(m_S(h)) + \theta - \delta_S(h)\tau + \beta E_{h,\theta}[f(h', \theta')]\}$$

Operator $\mathbb{T} : \mathcal{V} \mapsto \mathcal{V}$ is a contraction because it satisfies Blackwell's sufficient conditions (i.e. it is monotone and it discounts).

Let $f_0(h, \theta)$ be decreasing in h and increasing in θ. Then $\mathbb{T}f_0(h, \theta)$ is increasing in θ because the only time θ appears within the max is with a positive coefficient. Also, $\mathbb{T}f_0(h, \theta)$ is decreasing in h because each component within the max operator is decreasing in h, and so the maximum is as well. Since these weak properties are preserved by the operator, they must obtain in the unique fixed point, v. $\qquad\square$

Proof of Lemma 9. Define $V_b \equiv E_\theta[v(h, \theta)|h \geq 1]$ as the expected value of having rating $r = b$ under S before the draw of θ is realized and behaving optimally once it is. Then the immediately revealing value function for publicly known criminals can be written

$$v(h, \theta) = \max \left\{ \beta V_b, \theta - \bar{\delta}\tau + \beta V_b \right\} \tag{24}$$

The publicly known criminal optimally commits the crime if $\theta \geq \underline{\theta}_b \equiv \delta\tau$, which is increasing in the probability of detection and the penalty if caught.

The publicly known criminal's value function can be written as a piecewise function:

$$v(h, \theta) = \begin{cases} \beta V_b & \text{if } \theta < \underline{\theta}_b; \\ \theta - \bar{\delta}\tau + \beta V_b & \text{if } \theta \geq \underline{\theta}_b. \end{cases}$$

Since $\theta \sim U[0, 1]$,

$$E\left[\theta - \bar{\delta}\tau + \beta V_b | \theta \geq \underline{\theta}_b\right] = \frac{\underline{\theta}_b + 1}{2} - \bar{\delta}\tau + \beta V_b$$

The equation defining V_b is then:

$$V_b = \underline{\theta}_b \beta V_b + (1 - \underline{\theta}_b)\left(\frac{\theta_b + 1}{2} - \bar{\delta}\tau + \beta V_b\right)$$

Solving for V_b gives

$$V_b = \frac{1 - \underline{\theta}_b}{1 - \beta}\left(\frac{\theta_b + 1}{2} - \bar{\delta}\tau\right)$$

which is decreasing in $\bar{\delta}$ and τ. It is then easy to see that all components of equation 24 are decreasing in $\bar{\delta}$ and τ. $\qquad\square$

Proof of Lemma 10. Define $V_g \equiv E_\theta[v(0, \theta)]$ as the expected value of having rating $r = g$ under S before the opportunity is realized and behaving optimally once it is. Then the value function for unknown criminals can be written

$$v(0, \theta) = \max\left\{\ell + \beta V_g, \ell + \theta - \underline{\delta}\tau + \beta\left((1 - \underline{\delta})V_g + \underline{\delta}V_b\right)\right\} \qquad (25)$$

The unknown criminal optimally commits the crime if

$$\theta \geq \underline{\theta}_g \equiv \beta\underline{\delta}\left(V_g - V_b\right) + \underline{\delta}\tau \qquad (26)$$

An implication of Lemma 8 is that $V_g - V_b > 0$, so $\underline{\theta}_g > 0$. Equation 26 gives one equation in two unknowns: V_g and $\underline{\theta}_g$. Since they are determined jointly, a two equation system is needed to solve for the expected continuation value and the cutoff opportunity. In order to come up with another equation, I write V_g as the expectation of a piecewise linear function, as in the proof of Lemma 9.

$$V_g = \underline{\theta}_g(\ell + \beta V_g) + (1 - \underline{\theta}_g)\left(\ell + \frac{\theta_g + 1}{2} - \underline{\delta}\tau + \beta\left((1 - \underline{\delta})V_g + \underline{\delta}V_b\right)\right) \qquad (27)$$

Equations (26) and (27) form a system of two equations and two unknowns and can be solved for $\underline{\theta}_g$ and V_g. The solution is suppressed due to space constraints.

77

Comparative statics on $\underline{\theta}$ cannot be computed directly because the cutoff opportunity depends on the endogenous value:

$$\frac{\partial \underline{\theta}_g}{\partial \underline{\delta}} = \beta(V_g - V_b) + \tau + \beta\underline{\delta}\frac{\partial V_g}{\partial \underline{\delta}} \quad ; \quad \frac{\partial \underline{\theta}_g}{\partial \tau} = \underline{\delta} + \beta\underline{\delta}\frac{\partial V_g}{\partial \tau}$$

In order to compute $\partial V_g/\partial \underline{\delta}$ and $\partial V_g/\partial \tau$, I use the fact that $\underline{\theta}$ is the opportunity which maximizes the right hand side of (27). Differentiating equation 27 with respect to $\underline{\delta}$ and applying the Envelope Theorem gives

$$\frac{\partial V_g}{\partial \underline{\delta}} = \frac{-(1 - \underline{\theta}_g)(\beta(V_g - V_b) + \tau)}{1 - \beta + \beta\underline{\delta} - \beta\underline{\delta}\underline{\theta}_g}$$

Differentiating equation 27 with respect to τ and again applying the Envelope Theorem gives

$$\frac{\partial V_g}{\partial \tau} = \frac{-(1 - \underline{\theta}_g)\underline{\delta}}{1 - \beta + \beta\underline{\delta} - \beta\underline{\delta}\underline{\theta}_g}$$

I note that these partial derivatives of V_g are negative. Substituting these expressions back into $\partial \underline{\theta}_g/\partial \underline{\delta}$ and $\partial \underline{\theta}_g/\partial \tau$ gives

$$\frac{\partial \underline{\theta}_g}{\partial \underline{\delta}} = (\beta(V_g - V_b) + \tau)\left(1 - \frac{\beta\underline{\delta} - \beta\underline{\delta}\underline{\theta}_g}{1 - \beta + \beta\underline{\delta} - \beta\underline{\delta}\underline{\theta}_g}\right)$$

$$\frac{\partial \underline{\theta}_g}{\partial \tau} = \underline{\delta}\left(1 - \frac{\beta\underline{\delta} - \beta\underline{\delta}\underline{\theta}_g}{1 - \beta + \beta\underline{\delta} - \beta\underline{\delta}\underline{\theta}_g}\right)$$

Each of these expressions is positive because the fraction is less than one, implying that the optimal cutoff opportunity increases in both the detection rate and the penalty if caught.

Lastly, since V_g is decreasing in $\underline{\delta}$ and τ, it is easy to see that the value function from (25) is also decreasing in $\underline{\delta}$ and τ. $\qquad \square$

Proof of Lemma 12. A publicly known criminal's problem under the TC system is

$$w(h, \theta) = \max\left\{\beta W_2, \theta - \bar{\delta}\tau + \beta W_2\right\}$$

This equation is identical to the the value function of the publicly known criminal under the S system from equation 24. The maximum and the maximizer are therefore the same. $\qquad\square$

Proof of Lemma 13. The derivation of $\underline{\theta}_1$ is identical to that of $\underline{\theta}_g$ except using δ_1 as the detection probability. From Lemma 10, we know that the cutoff is increasing in the detection rate. Since $\delta_1 > \underline{\delta}$, we get $\underline{\theta}_1 > \underline{\theta}_g$. $\qquad\square$

Proof of Lemma 14. Agents with $h = 0$ receive the same wage and face the same probability of detection under each information system. To show $w(0, \theta) > v(0, \theta)$, let the agent under L mimic the agent under S. Then he gets caught just as often, but the circumstances are less dire. If caught, he becomes a $h = 1$ agent under L which has a higher value than being a known criminal under S because w is decreasing in h and $w(2, \theta) = v(1, \theta)$.

To show that $\underline{\theta}_0 < \underline{\theta}_g$ is more complicated, but also uses arguments based on mimicking behavior. Let W_x be an unknown criminal's expected lifetime utility before observing the opportunity θ under L while using the cutoff x when $h = 0$ and behaving optimally once convicted. Then W_x solves

$$W_x = \ell + x\beta W_x + (1 - x)\left(\frac{1 + x}{2} - \tau\underline{\delta} + \beta((1 - \underline{\delta})W_x + \underline{\delta}W_1)\right) \qquad (28)$$

Denote by $W_g \equiv W_x|_{x=\underline{\theta}_g}$; that is, W_g is the expected utility of an unknown criminal under the lenient system who mimicks the unknown criminal under the strict system until convicted, and then behaves optimally afterwards.

As an intermediary step, I first show that $W_g - W_1 < V_g - V_1$. This fact will be useful later in the proof. Subtracting W_1 from both sides of equation 28 evaluated at $x = \underline{\theta}_g$ gives

$$W_g - W_1 = \ell + \beta\underline{\theta}_g(W_g - W_1) + (1 - \underline{\theta}_g)\left(\frac{1 + \underline{\theta}_g}{2} - \underline{\delta}\tau + \beta(1 - \underline{\delta})(W_g - W_1)\right) - (1 - \beta)W_1$$

Solving for $W_g - W_1$ gives

$$W_g - W_1 = \frac{\ell + (1 - \underline{\theta}_g)\left(\frac{1+\underline{\theta}_g}{2} - \underline{\delta}\tau\right) - (1 - \beta)W_1}{1 - \underline{\theta}_g\beta - (1 - \underline{\theta}_g)\beta(1 - \underline{\delta})}$$

Analogous operations on V_g give

$$V_g - V_b = \frac{\ell + (1 - \underline{\theta}_g)\left(\frac{1+\underline{\theta}_g}{2} - \underline{\delta}\tau\right) - (1 - \beta)V_b}{1 - \underline{\theta}_g\beta - (1 - \underline{\theta}_g)\beta(1 - \underline{\delta})}$$

Since $W_1 > V_b$ and the other terms in the fractions are identical, we see that $W_g - W_1 < V_g - V_b$.

Differentiating equation 28 in x gives the rate of change of value as the cutoff changes. The formula is

$$\frac{\partial W_x}{\partial x} = \frac{-x + \underline{\delta}\tau + \beta\underline{\delta}(W_x - W_1)}{1 - \beta + \beta\underline{\delta} - \beta\underline{\delta}x}$$

Evaluating this derivative at $x = \underline{\theta}_g$ gives

$$\left.\frac{\partial W_x}{\partial x}\right|_{x=\underline{\theta}_g} = \frac{-\underline{\theta}_g + \underline{\delta}\tau + \beta\underline{\delta}(W_g - W_1)}{1 - \beta + \beta\underline{\delta} - \beta\underline{\delta}\underline{\theta}_g} < \frac{-\underline{\theta}_g + \underline{\delta}\tau + \beta\underline{\delta}(V_g - V_b)}{1 - \beta + \beta\underline{\delta} - \beta\underline{\delta}\underline{\theta}_g} = 0$$

The inequality holds because $W_g - W_1 < V_g - V_b$, and the equality holds because $\underline{\theta}_g$ is optimal for unknown criminals under S and the first order condition gives $\underline{\theta}_g = \underline{\delta}\tau + \beta\underline{\delta}(V_g - V_b)$.

This means the unknown criminal under the lenient system who is mimicking the unknown criminal under the strict system could increase his expected utility by decreasing his cutoff opportunity; that is, he could use some $x < \underline{\theta}_g$ and increase his utility. But this is still a local argument; it is possible some $x > \underline{\theta}_g$ maximizes utility. To make the argument global, assume towards a contradiction that $\underline{\theta}_0 > \underline{\theta}_g$. Then since lifetime expected utility is continuous and differentiable in the cutoff opportunity, it must be the case that there is some $\theta' > \underline{\theta}_g$ at which lifetime expected utility attains a local

minimum. If this is true, the slope of lifetime expected utility in the cutoff is necessarily zero. But the second order condition holds whenever the first order condition holds, so there cannot be a local minimum at θ'. That is, whenever the slope of the objective function is zero, the second derivative is negative and we are at a local maximum, not a local minimum. This proves that $\underline{\theta}_0 < \underline{\theta}_g$, and therefore unknown criminals commit more crimes under L than they do under S. $\qquad\square$

Proof of Proposition 8. I find the value of $\underline{\theta}_g$ which solves the equation $\xi_0 + \xi_1 = \xi_g$. The cutoff value for unknown criminals under S which equates the mass of unknown criminals under the two notification systems is

$$\underline{\theta}_g = \frac{1 - \beta + \beta\underline{\delta}(1 - \underline{\theta}_0) + \underline{\theta}_0\beta\delta_1(1 - \underline{\theta}_1)}{1 - \beta + \beta\underline{\delta}(1 - \underline{\theta}_0) + \beta\delta_1(1 - \underline{\theta}_1)} \equiv h(\underline{\theta}_1)$$

I will show that $h(\underline{\theta}_1 \geq \underline{\theta}_1$, and since Lemma 13 shows that $\underline{\theta}_g < \underline{\theta}_1 \leq h(\underline{\theta}_1)$, it cannot be the case that $\underline{\theta}_g = h(\underline{\theta}_1$. The mass of unknown criminals across rating systems cannot be equal. Since the mass of unknown criminals under S is decreasing in $\underline{\theta}_g$, this proves that the mass of criminals with a good public rating is always larger when the government is lenient.

First note that

$$h(1) = \frac{1 - \beta + \beta\underline{\delta}(1 - \underline{\theta}_0)}{1 - \beta + \beta\underline{\delta}(1 - \underline{\theta}_0)} = 1$$

But since $\underline{\theta}_g < \underline{\theta}_1$ by Lemma 13, it cannot be the case that the mass of unknown criminals is equal across notification policies if $\underline{\theta}_1 = 1$.

I next show that the slope of the function h is less than one. This, along with the fact that $h(1) = 1$ proves that $h(\underline{\theta}_1) \geq \underline{\theta}_1$. Differentiating the function h gives

$$h'(\underline{\theta}_1) = \frac{\beta\delta_1(1 - \underline{\theta}_0)(1 - \beta + \beta\underline{\delta} - \beta\underline{\delta}\underline{\theta}_0)}{(1 - \beta + \beta\underline{\delta} - \beta\underline{\delta}\underline{\theta}_0) + \beta\delta_1(1 - \underline{\theta}_1)} < 1$$

Then at any value $\underline{\theta}_1 \in [0, 1)$, $h(\underline{\theta}_1) > \underline{\theta}_1$.

But Lemma 13 says that $\underline{\theta}_g < \underline{\theta}_1$. This means that $\underline{\theta}_g < \underline{\theta}_1 < h(\underline{\theta}_1)$, and so the mass of unknown criminals cannot be equal across the two notification systems.

All that remains is to show that the mass of unknown criminals is larger under the lenient system, and hence productivity is higher. The mass of unknown criminals under S decreases in $\underline{\theta}_g$ because the higher is the cutoff, the less crimes get committed, and less criminals are caught and then become publicly known criminals. Then since $\underline{\theta}_g < h(\underline{\theta}_1)$, the mass of unknown criminals is larger under the strict system. $\square$

References

G. Adkins, D. Huff, and P. Stageberg. The Iowa sex offender registry and recidivism. *Report of the Iowa Department of Human Rights, Division of Criminal and Juvenile Justice Planning and Statistical Analysis Center*, 2000.

A. Agan. Sex offender registries: Fear without function? *SSRN Working Paper*, 2008.

G. Becker. Crime and punishment: An economic approach. *Journal of Political Economy*, 76:169–217, 1968.

C. Y. C. Chu, S.-C. Hu, and T.-Y. Huang. Punishing repeat offenders more severely. *International Review of Law and Economics*, 20:127–140, 2000.

M. Cripps, G. Mailath, and L. Samuelson. Imperfect monitoring and impermanant reputations. *Econometrica*, 72:407–432, 2004.

P. F. Cromwell and J. N. Olson. *Breaking and Entering: Burglars on Burglary*. Wadsworth Publishing Company, Address, 2003.

R. Di Tella and E. Schargrodsky. Do police reduce crime? estimates using the allocation of police forces after a terrorist attack. *American Economic Review*, 94(1):115–133, 2004.

D. Dranove, M. Satterthwaite, D. Kessler, and M. McClellan. Is more information better? the effects of "report cards" on health care providers. *Journal of Political Economy*, pages 555–588, 2003.

P. Dubey and J. Geanakoplos. Grading exams: 100, 99,...,1 or a, b, c? incentives in games of status. *Games and Economic Behavior*, 69:72–94, 2010.

J. Eeckhout, N. Persico, and P. Todd. A theory of optimal random crackdowns. *American Economic Review*, 100(3):1104–1135, 2010.

M. Ekmekci. Sustainable reputations with rating systems. 2010.

K. F. Ferraro. *Fear of Crime: Interpreting Victimization Risk*. State University of New York Press, Albany, NY, 1995.

R. B. Freeman. Why do so many young american men commit crimes and what might we do about it? *Journal of Economic Perspectives*, 10(1): 24–42, 1996.

D. Fudenberg and D. K. Levine. Maintaining a reputation when strategies are imperfectly observed. *Review of Economic Studies*, 59:561– 579, 1992.

E. Goode. States seeking new registries for criminals. *The New York Times*, May 20, 2011.

J. Grogger. The effect of arrests on the employment and earnings of young men. *The Quarterly Journal of Economics*, 110(1):55–71, 1995.

R. Hanson. Recidivism and age: Follow-up data from 4,673 sexual offenders. *Journal of Interpersonal Violence*, 17(10):1046–1062, 2002.

R. K. Hanson, K. E. Morton, and A. J. R. Harris. Sexual offender recidivism risk: What we know and what we need to know. *Annals of the New York Academy of Sciences*, 989:154–166, Jun 2003.

T. Hirschi and M. Gottfredson. Age and the explanation of crime. *The American Journal of Sociology*, 89(3):552–584, 1983.

B. Holmstrom. Managerial incentive problems - a dynamic perspective. *Review of Economic Studies*, 66(1):169–182, 1999.

H. Holzer, S. Raphael, and M. Stoll. How willing are employers to hire ex-offenders? *Focus*, 23(2):40–43, 2004.

G. Jin and P. Leslie. The effects of information on product quality: Evidence from restaurant hygiene grade cards. *Quarterly Journal of Economics*, 118: 409–451, 2003.

L. F. Katz, S. D. Levitt, and E. Shustorovich. Prison conditions, capital punishment and deterrence. *American Law and Economics Review*, 5(2): 318–343, 2003.

P. D. Kernsmith, E. Comartin, S. W. Craun, and R. M. Kernsmith. The relationship between sex offender registry utilization and awareness. *Sexual Abuse*, 21(2):181–193, 2009.

D. Kreps and R. Wilson. Reputation and imperfect information. *Journal of Economic Theory*, 27:253–279, 1982.

D. Kreps, P. R. Milgrom, D. J. Roberts, and R. Wilson. Rational cooperation in the finitely repeated prisoner's dilemma. *Journal of Economic Theory*, 27:245– 252, 1982.

D. S. Lee and J. McCrary. The deterrence effect of prison: Dynamic theory and evidence. Working Papers 1168, Princeton University, Department of Economics, Center for Economic Policy Studies, Jul 2009.

S. D. Levitt. Using electoral cycles in police hiring to estmate the effect of police on crime. *American Economic Review*, 87(3):270–290, 1997.

R. Lieb and C. Nunlist. Community notification as viewed by washington's citizens: A 10-year follow-up (document 08-03-1101). *Washington State Institute for Public Policy*, Mar 2008.

L. Linden and J. E. Rockoff. Estimates of the impact of crime risk on property values from megan's laws. *American Economic Review*, 98(3):1103–1127, 2008.

J. Lott. *More Guns, Less Crime.* University of Chicago Press, Chicago, 1998.

G. Mailath and L. Samuelson. Who wants a good reputation? *Review of Economic Studies*, 68:415–441, 2001.

H. M. Mialon and T. Wiseman. The impact of gun laws: A model of crime and self defense. *Economics Letters*, 88:170–175, 2006.

P. Milgrom and D. Roberts. Predation, reputation and entry deterrence. *Journal of Economic Theory*, 27:280–312, 1982.

J. R. Neto. Should criminal records be publicly available? Dissertation Chapter, 2006.

D. M. Philips. Community notification as viewed by washington's citizens (document 98-03-1101). *Washington State Institute for Public Policy*, Mar 1998.

A. M. Polinsky and S. Shavell. The economic theory of public enforcement of law. *Journal of Economic Literature*, 38(1):45–76, 2000.

M. Polinsky and D. Rubinfeld. A model of fines for repeat offenders. *Journal of Public Economics*, 46:291–306, 1991.

J. C. Pope. Fear of crime and housing prices: Household reactions to sex offender registries. *Journal of Urban Economics*, 64(3):601–614, 2008.

J. Prescott and J. E. Rockoff. Do sex offender registration and notification laws affect criminal behavior? *Journal of Law and Economics*, 54, 2011.

A. Rubinstein. An optimal conviction policy for offenses that may have been committed by accident. In S. Brams, A. Schotter, and G. Schwödiauer, editors, *Applied Game Theory*, pages 406 – 413. Physica-Verlag, 1979.

A. Rubinstein. On an anomaly of the deterrent effect of punishment. *Economic Letters*, 6:89–94, 1980.

D. Silverman. Street crime and street culture. *International Economic Review*, 45(3):761–786, 2004.

L. Smith. Crime and vigilance as a large game. Unpublished Notes, 2010.

J. Sonas. The Sonas rating formula – better than elo? *www.chessbase.com*, 2002.

J. Waldfogel. The effect of criminal conviction on income and the trust 'reposed in the workmen'. *The Journal of Human Resources*, 29(1):62–81, 1994.

J. Walker, S. Maddan, B. Vàsquez, A. VanHouten, and G. Ervin-McLarty. The influence of sex offender registration and notification laws in the united states. *Arkansas Crime Information Center Working Paper*, 2005.

O. D. Welch. Good enough for skinheads: How the proliferation of non-sexual criminal registries makes us less safe. Unpublished Manuscript, 2008.

R. Wright and S. Decker. *Burglars on the Job: Streetlife and Residential Break-Ins.* Northeastern University Press, Boston, 1994.