I Z A Institute
of Labor Economics
Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 15053

Socially Optimal Crime and Punishment

Eduardo Ferraz
Rodrigo R. Soares

JANUARY 2022

IZA Institute of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 15053

Socially Optimal Crime and Punishment

Eduardo Ferraz
Universidad del Rosario

Rodrigo R. Soares
Insper and Columbia University and IZA

JANUARY 2022

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Socially Optimal Crime and Punishment*

This paper develops a dynamic life-cycle equilibrium model of crime with hetero-geneous agents and human capital accumulation. Agents decide at each point in time whether to commit crimes by comparing potential gains from crime to the expected cost of punishment (determined from the probability of apprehension, the utility cost of incarceration, and reduced future wages in the legal labor market). Public security policies are defined as pairs of a size of the police force and an average length of sentences. We propose an original micro-founded police production function linking the level of police expenditures to the probability of apprehension. The structural model, estimated using 2000s US data and causal parameters from the empirical literature, allows us to evaluate the global optimality of policies in a way that would not be possible with reduced form estimates or traditional partial equilibrium, static models of crime. Equilibrium effects can be particularly relevant when studying crime, given the interactions across individuals' decisions and policies. We also extend the model to include investments in schooling and explore the potential complementarities across public security and educational policies.

JEL Classification: K42, I38

Keywords: crime, welfare, police, sentence length, socially optimal policy

Corresponding author:
Rodrigo R. Soares
School of International and Public Affairs
Columbia University
420 W. 118th Street
New York, NY 10027
USA
E-mail: r.soares@columbia.edu

* We are grateful to Braz Camargo, Bernardo Guimarães, Lance Lochner, Steven Raphael, David Rivers and seminar participants at Sao Paulo School of Economics-FGV, Universidad del Rosario, University of Western Ontario, Brazilian Econometrics Society Webinar Series, Economic Research Southern Africa Webinar Series, the 2019 Workshop on the Economics of Risky Behaviors, the 2019 ALCAPONE-LACEA Workshop, the 2018 LACEA-LAMES Annual Meeting, the 2018 Transatlantic Workshop on the Economics of Crime, and the 2016 Brazilian Econometrics Society Meeting for valuable comments and discussions.

1 Introduction

What are the socially optimal crime rate and associated public security policy? Is the US really overspending on police forces and over punishing minor crimes? Questions such as these, recurrent in the contemporary public security debate (PEW 2016, Hayes 2020), were the original motivations for the classic contributions to the economics of crime (Beccaria 1764, Bentham 1781, Becker 1968, Stigler 1970, Ehrlich 1981). Nevertheless, economics still has little to say about the optimality of public security policies. Until today, the discipline does not share an accepted theoretical framework to quantitatively assess interactions across policy instruments, equilibrium responses, and their implications for the social optimality of policies.

This paper revisits the theoretical discussion on the global optimality of public security policies to start filling the wide gap that currently exists between the empirical and theoretical literature in economics of crime. This is particularly relevant in light of the large costs that crime imposes on society and of the amount of evidence currently available on the effectiveness of specific interventions. In the US, for example, expenditures on the criminal justice system added up to $210 billion in 2012, with estimates suggesting that the aggregate welfare loss experienced by victims is at least as large.[1] At the same time, a vast array of evidence indicates that crime is very responsive to policy, be it related to police, incarceration, schooling, or social protection.[2] In fact, among the different dimensions of public policy, public security is somewhat exceptional in the extent to which it has historically relied on empirical evidence (see, for example, Sherman et al. 1998, Cordner 2020). Virtually all of this evidence, though, comes from randomized controlled trials and reduced-form estimates, with few, if any, equilibrium or dynamic considerations.

Crime depends heavily on equilibrium responses, so different policies are likely to interact in non-trivial ways. Interventions may also have dynamic implications that are themselves relevant in determining the optimality of policies. Longer sentences, for example, reduce the number of potential criminals on the streets in the short run, increasing the effectiveness of a given police force, which in turn affects the optimal size and allocation of police. But longer sentences also reduce the future employability of inmates, increasing their likelihood of recidivism. Are these trade-offs relevant from a policy perspective? If so, which force dominates in the long run? Cost-benefit analyses based on reduced-form estimates cannot hope to answer questions such as these. The potential gain from incorporating evidence from different interventions into a unified framework is therefore very large. Some of these trade-offs have surely been conceptually identified, in one way or another, dating back to the early contributions to the economics of crime. But, since then, theoretical work on the topic has remained highly stylized, providing little guidance on how to consolidate different

[1] Our calculations based on data from the FBI and NCVS, combined with estimates of the costs of crime from Cohen (2000). We estimate an yearly cost of victimization of the order of $240 billion.

[2] For example: on police, see Levitt (1997), Di Tella and Schargrodsky (2004), and Draca et al. (2008); on incarceration, see Levitt (1996), Owens (2009), and Buonanno and Raphael (2013); and on schooling and welfare transfers, see Hannon and DeFranzo (1998), Lochner (2004), Foley (2011), and Jacob and Ludwig (2010).

pieces of evidence in a consistent way.

We contribute to the literature by proposing a structural equilibrium model to study the optimality of public security policies. Our theory incorporates interactions across policies and equilibrium responses in a tractable way, allowing the model to be estimated and used for counterfactual policy simulations. Public security policies in our context are understood as pairs describing a size of the police force and an average length of sentences. We consider a dynamic equilibrium life-cycle model with heterogeneous agents and human capital accumulation, where agents decide at each point in time whether to commit crimes. This decision is based on the comparison between the potential gain from crime and the expected loss due to a positive probability of being caught and the associated punishment (duration of incarceration). A crime is modelled as an encounter between two agents implying: (i) a transfer of resources from victims to criminals (the good stolen); (ii) a non-monetary utility loss to victims (the psychological and physical costs of being victimized); and (iii) a net loss of resources (the potential destruction or use of resources during the perpetration of a crime). We assume that individuals are heterogeneous in their subjective (psychological or moral) costs to commit crimes, such that there is a mass of individuals who contemplate committing crimes and a mass of individuals who, irrespective of the circumstances, never commit crimes.

The potential criminals' problem has a simple and intuitive solution: at each point in time, conditional on a given public security policy and on the decisions of other agents, there is a cut-off level of labor earnings below which potential criminals engage in crime. This cut-off evolves over time as individuals grow older and their earnings, costs to commit crimes, and horizons change. We prove the existence of a rational expectations equilibrium in this economy and characterize it.

A key challenge in writing down this problem, and maybe one of the reasons why quantitative theoretical work in economics of crime has lagged behind, is how to model the *police production function* available to policy makers. In other words, how do expenditures on police translate into a probability of apprehending criminals?[3] Our first contribution to the theoretical literature is to propose a simple micro-founded police production function. This function is the equilibrium outcome of a game played between police officers and criminals, in which criminals choose the location of crimes to avoid detection and policymakers choose where to deploy police forces in order to maximize the number of apprehensions. In our setting, this game can be interpreted as the within-period game played between criminals and police, conditional on criminal participation decisions and on the size of the police force. Assuming that each apprehension consumes police resources, and relying on a Poisson point process modelling strategy, we show that the probability of apprehension of any individual criminal can be characterized as a two-parameter closed-form function depending on the size of the police (positively) and on the crime rate (negatively). This function displays all the properties that would seem intuitively desirable. In addition, its parameters can be identified directly from the estimates of the elasticity of crime with respect to prison population

[3] According to the Bureau of Justice Statistics, approximately 45% of expenditures on the justice system in the US are allocated to police forces. Therefore the relevance of what we call the *police production function*.

and police expenditures available from the empirical literature (such as, for example, Levitt 1996, 2002).[4]

This police production function also has the appealing feature of naturally bringing into the picture social interactions and the possibility of multiple equilibria in crime rates. Since the probability of apprehension of any given criminal is a direct function of the number of active criminals, the possibility of social interactions arises endogenously. The potential relevance of social interactions in determining equilibrium crime rates has been highlighted before both theoretically and empirically (see, for example, Gleaser et al. 1996, Burdett et al. 2003, and Gaviria 2000, each one with a different rationalization for the presence of social interactions).

The incorporation of a life-cycle component is another key feature of our model. A realistic life-cycle earnings curve is essential, for example, for the model to be able to reproduce the widely documented age profile of participation into crime (e.g., see the discussion in Beirne 1993, Lochner 2004, and Munyo 2015). It is also indispensable to capture the implications of potential reductions in future legal earnings due to incarceration. Grogger (1995) suggests that one year of incarceration reduces wages by as much as 30% in the long run. This labor market channel amplifies the deterrence effect of incarceration in non-trivial ways and makes it highly age dependent. We summarize life-cycle productivity changes with a learning-by-doing human capital accumulation technology, somewhat similar to Flinn (1986). We assume that individuals accumulate skills as a by-product of work and whenever they are incarcerated their productivity depreciates. The interaction of public security policies and age leads to implications in terms of criminal behavior analogous to those from the standard Becker-Ben Porath model of investments in human capital.

The combination of the ingredients highlighted above introduces equilibrium interactions that would otherwise have been impossible to consider. For example, in principle, it is possible that increasing the length of sentences leads to an increase in the steady-state crime rate, due to the positive effect of incarceration on recidivism through reduced future earnings. Whether this mechanism is relevant depends on the equilibrium under consideration and on interactions across various margins, which cannot be assessed without an explicit life-cycle equilibrium model.

We estimate the model using US property crime data in 2004. Our estimation uses a minimum distance procedure with moments taken from the National Crime Victimization Survey, the National Corrections Reporting Program, the Current Population Survey, the National Longitudinal Survey of Youth, the Sourcebook of Criminal Justice Statistics, Abrams and Rohlfs (2007), and Levitt (1996, 2002). Some of these moments, as well as some

[4]Ehrlich (1973) was the first to make the apprehension probability an explicit function of public security expenditures. He proposed an arbitrary public security production function in which the probability of punishment is a function of police expenditures and the crime rate. The functional form is appealing at the local level, but can lead to probabilities that exceed 1 for large deviations. İmrohoroğlu et al. (2000) adopt a similar functional form that, instead, can lead to negative probabilities (they deal with this problem by truncating the apprehension probability whenever it is negative). To our knowledge, these are the only papers that explicitly model a public security production function. The limitations of their approaches lie on the arbitrariness of both the functional form and the assignment of numerical values to the parameters.

of the parameters that we borrow directly from previous work, rely on quasi-experimental evidence from the applied micro literature, tying our estimation and quantitative exercises closely together with causal evidence from the natural experiments literature. Since we consider rational criminals motivated by the economic return to crime, we focus on property crimes. We choose 2004 because it lies in the middle of a period ranging from 2000 to 2010 when crime trends remained relatively stable, coming closer to the stationarity hypothesis needed to numerically solve the model.

For a wide range of policy alternatives and starting points, our estimation leads to a single point in the parameter space and the equilibrium seems to be unique. The estimated model reproduces various targeted and untargeted moments related to crime level and frequency, inflow of prisoners by age, prison population, welfare loss from incarceration, and responses of crime to different dimensions of policy.

Our main quantitative exercise considers the implications of the model regarding the interaction between the two most widely studied dimensions of public security policy: police and sentencing. The estimated model indicates that the 2004 public security policy came very close to minimizing the crime rate subject to the expenditures actually observed in 2004. This suggests some rationality on the side of public security policy makers, access to information, and an exclusive focus on crime reduction.

But the goal of public security policy is not to minimize crime at all costs, but rather to minimize the total welfare loss from crime. We show that, under an additive welfare function, the optimal 2004 public security policy characterized by the model implies a higher crime rate, a smaller police force, and longer average sentences than those actually observed in 2004. Under this policy, public security expenditures would be lower by 11% and crime rates higher by 11.9%, corresponding to yearly welfare gains of the order of $1 billion (or 1.6% of the initial welfare cost of crime). The reduction in police expenditures by itself is enough to generate a much cheaper public security system, which translates into a non-trivial gain for taxpayers and more than compensates the increase in property crime. The increase in average sentence lengths, in turn, comes directly as a way to minimize the increase in crime from the reduction in police expenditures. The small increase in crime in the social optimum, though maybe counterintuitive at first sight, should not be seen as particularly surprising for an audience of economists (crime rates in the US were close to their historical minimum during the period analyzed).

We conduct other quantitative exercises acknowledging that policy makers sometimes face political constraints not incorporated in our theory, and these may limit the set of feasible instruments or objectives at their disposal. Our structural model allows us to consider any such policy and to determine the optimal policy under any welfare metric. So, for example, we also discuss optimal policies that do not increase the initial crime rate or that exclude part of criminals' welfare from the social welfare function. Reassuringly, for all exercises considered, we find that the optimal policy would weakly reduce expenditures on police and increase average sentences. This broad qualitative characterization of the inefficiency of the 2004 public security policy seems therefore to be quite general.

Our final set of quantitative exercises consider the complementarities between public security and educational policies. By affecting life-cycle productivity and increasing the opportunity cost of incarceration, education strengthens the deterrence effect and can lead to substantial reductions in criminal participation. We analyze the additional welfare gains that could be obtained by reallocating part of the expenditures in public security to educational policies. We conduct this analysis considering internal rates of return from investments in education of 0%—analogous to a cash transfer welfare program—and 15% (Heckman et al. 2008), and also varying the ability of the government to target the policy on certain income groups. We show that even with a 0% rate of return and no targeting (equal transfer to all individuals), shifting resources from public security to education increases the potential gains in welfare, though by a small magnitude. As the internal rate of return increases and targeting improves, these potential gains become very large, reaching \$5.6 billion per year—or 8.8% of the initial welfare cost of crime—for the scenario that maximizes targeting. Under these scenarios with improved education, public security policies would move in the direction of further reducing the size of the police and increasing average sentence lengths. The number of potential criminals, at the margin, is reduced, so the benefits from a large police force also fall. But the remaining criminals, after the policy change, are very inelastic to punishment and have low opportunity cost, so, for these, incarceration becomes a more cost-effective way to reduce crime. The model suggests that educational policies are indeed a potentially important instrument to reduce the costs of crime while at the same time minimizing the need for punishment. Incorporating potential long-term effects of education on aggregate productivity and growth would only reinforce this qualitative result.

Our quantitative model also sheds new light on the traditional decomposition of the effects of public security policies. First, we define theoretically the incapacitation, deterrence, and recidivism effects and their respective roles, something that, to our understanding, was never done before. And second, we show that, when considering equilibrium responses, there is always an additional effect that has not yet been considered in the literature: we call it the *load* effect. The *load* effect refers to the lower effectiveness of a given police force under higher crime rates, due to the dilution of human and physical resources when trying to catch a larger number of criminals. The *load* effect is, by nature, an equilibrium effect, therefore requiring an equilibrium framework to be conceptually defined. Our theoretical discussion also makes it clear that, in equilibrium, none of the four effects ever exists in isolation, each one being always intrinsically connected to the others.

The conceptual question of optimal law enforcement is by no means new, dating back at least to Hobbes in 1651, followed by Beccaria (1764) and Bentham (1781), and then by the more recent treatments from the economics of crime, including Becker (1968) and Ehrlich (1981). Despite the unquestionable value of these seminal contributions and the philosophical and theoretical insights they continue to provide, their use for actual policy analysis is limited. From this perspective, the papers closest to our work are Engelhardt et al. (2008), Fella and Gallipoli (2014), and Fu and Wolpin (2017). These papers develop equilibrium models of crime, simulate policies, and conduct normative analyses. Engelhardt et al. (2008) focus on the interaction between the labor market and criminal behavior. Fella and Gallipoli (2014) analyze the relationship between crime and education. Fu and Wolpin (2017), in turn, are mainly concerned about the optimality of police expenditures. Our contributions to this theoretical literature include a micro-founded public security technology, multiple dimensions of policy, the explicit incorporation of the life cycle of wages and productivity losses from incarceration, and a theoretically consistent—and intuitively appealing—definition of the welfare loss from crime. From an empirical perspective, we are the first ones to estimate the model, to link parameters directly to evidence from causal estimates, and to conduct counterfactual exercises exploring the interactions across different policies, namely, the size of the police force, average sentence lengths, and educational policies.

Needless to say, our model has many limitations. From the perspective of victims, we do not consider risk aversion nor the possibility of changes in behavior to avoid victimization. From the perspective of criminals, we do not consider non-economic crimes (such as emotionally motivated interpersonal violence), do not incorporate the possibility of accumulation of criminal capital, and do not explore the implications of behavioral biases, such as hyperbolic discounting. Finally, we are completely silent about the inner workings of the justice system, simply assuming a direct link between the apprehension of a criminal and the expected punishment. Our model has no police error nor bias and no wrongful convictions. These represent injustices derived from public security policy that can, by themselves, lead to major welfare losses, adding up to those discussed in the paper. The incorporation of most of these dimensions in our theoretical framework would be relatively straightforward (see, for example, Mocan et al. 2005 and Munyo 2015 for models with criminal human capital and Ferraz 2017 for a model with multiple types of crimes). The main challenge, though, would reside in the estimation. Data limitations already make it difficult for us to estimate the model as it is. We see our contribution therefore as a methodological first step towards building a manageable model of crime that can bring together different pieces of evidence—and equilibrium considerations—when evaluating the trade-offs across alternative policies. The points raised in this paragraph would be natural extensions in the future development of this research agenda.

The remainder of the paper is organized as follows. Section 2 presents the theory. Section 3 describes the estimation of the model. Section 4 develops the main quantitative exercises and discusses the optimal public security policy in the context of the model. Section 5 concludes the paper.

2 The Model

We consider a discrete-time dynamic equilibrium model with heterogeneous agents. Crime is the only source of inefficiency. We describe our theoretical framework in detail below.

2.1 Preferences

The economy is populated by a continuum of risk-neutral individuals who are ex ante heterogeneous with respect to productivity and the propensity to commit crimes. At each period t, an individual maximizes the expected discounted lifetime utility

$$\mathbb{E}\Big[\sum_{\tau=t}^{T}\beta^{\tau-t}c_{\tau}+\beta^{T+1-t}\alpha c_{T+1}\Big],$$

where β denotes the intertemporal discount factor and c_τ the consumption at age τ. Individuals can commit crimes only up to age T, after when they become too old to engage in criminal activities. The term αc_{T+1} captures the utility continuation value at age $T+1$.

We assume that some individuals, who we call *honest agents*, do not consider committing a crime under any circumstance. These can be interpreted as individuals who face a prohibitively high cost of committing crimes, possibly due to cultural, psychological or ethical reasons. *Dishonest agents* are defined as those who contemplate the possibility of committing crimes. At each period, a constant number of newborn individuals (age 0) of each type enter the economy. The number of dishonest individuals entering the economy is M_0. The productivity of each agent at age 0 is drawn from a distribution given by g_0 and, with abuse of notation, we set $\mathrm{d}G_0(w) = g_0(w)\mathrm{d}w$ for each w in the support of g_0.

Honest agents are entirely passive. They are targets of criminal actions and are important for social welfare considerations, but make no active choices. In addition, since agents are risk neutral, the (utilitarian) social loss that crime imposes directly on victims depends only on the aggregate crime rate.

2.2 Legal Sector and Crime

Dishonest agents younger than $T+1$ can be in one of three states: free and engaged in crime (state E), free and not engaged in crime (state $\tilde{E}$), and incarcerated (state P). After age $T+1$, all agents are necessarily in state $\tilde{E}$. At the beginning of a period, each dishonest agent who is not incarcerated chooses whether or not to engage in crime.

There is only one type of crime in the economy. A crime is an interaction between a dishonest agent (criminal) and an honest agent (victim), resulting in three outcomes: a financial transfer of value z from the victim to the criminal, representing the market value of the property stolen; a cost d^V inflicted on the victim due to physical or psychological damages; and an age-dependent cost d_t^C paid by the criminal. The term d_t^C captures the psychological and physical cost (or excitement) from committing a crime, as well as the material investment in and opportunity cost of the criminal enterprise.

If a dishonest individual i chooses to engage in crime at age t, this individual remains engaged until the end of the period. In this case, the number of criminal opportunities is determined by an intra-period continuous-time stochastic process. The agent receives successive opportunities of crime, with the interval between two (within period) opportunities defined as a random variable exponentially and independently distributed with parameter ν. By assumption, an individual i who engages in crime in a certain period commits a criminal act at every opportunity that is presented, until she/he is arrested or the period ends (in other words, the criminal participation decision is taken at the period level).

Each crime is detected with probability p. A criminal can be arrested for a given crime only at the instant when the crime is committed.[5] The exponentially distributed interval between two criminal opportunities implies that, if agent i at age t is not arrested for any crime, the number of crimes J_t^i within a period is a random variable distributed according to a Poisson with parameter ν. Therefore, conditional on J_t^i, the probability that an active criminal is arrested in a given period is $1-(1-p)^{J_t^i}$. This implies an expected probability of being arrested in a period, from the perspective of a criminal, given by

$$\mathbb{E}\left[1-(1-p)^{J_t^i}\right] = \sum_{j=0}^{\infty}\left((1-(1-p)^j)\frac{\nu^j e^{-\nu}}{j!}\right) = 1-e^{-\nu p}.$$

Criminals at age t consume an extra $z-d_t^C$ for each non-detected crime. In case a crime is detected, there is no financial transfer between victim and criminal, but the victim still suffers the damage d^V and the criminal still pays the cost d_t^C. Under these assumptions, the gain from committing crimes in a given period is a random variable, denoted by Π_t^i, depending on the age of the agent. For a given number of crimes J_t^i, if the criminal is apprehended for the j-th crime (chronologically), $1 \leq j \leq J_t^i$, the realized within-period criminal gain is $(j-1)z - jd_t^C$. The probability of being apprehended for this j-th crime is $(1-p)^{j-1}p$, so that

$$\mathbb{E}\left[\Pi_t^i | J_t^i\right] = \frac{(z(1-p)-d_t^C)(1-(1-p)^{J_t^i})}{p}.$$

Thus, the unconditional expected gain from crime in a given period is given by

$$\overline{\Pi}_t := \mathbb{E}[\Pi_t^i] = \frac{(z(1-p)-d_t^C)(1-e^{-\nu p})}{p}.$$

When a criminal at age $t<T$ is apprehended or when an agent starts a period in prison, there is a probability $1-\mu$ that she/he will remain in prison in the next period, and a probability μ that she/he will be released in the beginning of the next period. Note that this implies that a criminal can be released immediately after being arrested, capturing the fact that some criminals are incarcerated for very short periods of time. Agents at the terminal age T that were apprehended or started the period in prison are all set free at age $T+1$. The average sentence length is given by $\lambda = \frac{1}{\mu} - 1$.[6] One of the roles of the government in this economy is to choose the average sentence length λ, which is equivalent to choosing the probability μ.

Agents who are free earn a legal sector income equal to their productivity. For brevity, we call their legal sector income simply income. We denote by w_t^i the productivity of an

[5] We write that crimes are "detected" whereas criminals are "arrested" or "apprehended."

[6] Strictly, the average sentence length is lower than $\frac{1}{\mu}-1$ due to the fact that incarcerated individuals are set free at age $T+1$. However, if the number of agents at age T who are engaged in crime or in prison is small enough, the average sentence length becomes very close to $\frac{1}{\mu}-1$. This is the case in our estimated model and the purpose of defining the terminal age T.

agent i at age t. Each individual pays a lump-sum tax f. So, the expected net total earnings of an agent i at age t as a function of state s_t^i is given by

$$\mathbb{E}\left[c_t^i|s_t^i\right] = \begin{cases} w_t^i - f & \text{if } s_t^i = \tilde{E}, \\ w_t^i - f + \overline{\Pi}_t & \text{if } s_t^i = E, \\ b - f & \text{if } s_t^i = P, \end{cases}$$

where b is the pre-tax consumption equivalent of being in prison. Under these assumptions, and defining $\overline{p} := (1-\mu)(1-e^{-\nu p})$, we obtain the transition probabilities across states:

$$\begin{gathered} \Pr(s_{t+1}^i \in \{\tilde{E}, E\}|s_t^i = \tilde{E}) = 1 \\ \Pr(s_{t+1}^i \in \{\tilde{E}, E\}|s_t^i = E) = 1 - \overline{p} \\ \Pr(s_{t+1}^i = P|s_t^i = E) = \overline{p} \\ \Pr(s_{t+1}^i \in \{\tilde{E}, E\}|s_t^i = P) = \mu \\ \Pr(s_{t+1}^i = P|s_t^i = P) = 1 - \mu \end{gathered}$$

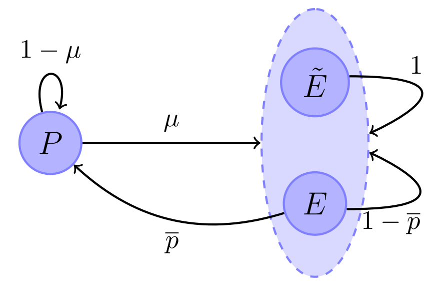

To capture life-cycle income profiles, we assume that the productivity of an individual at age t increases by a factor γ_t. By contrast, due to stigma and human capital depreciation, incarcerated individuals loose productivity by a factor θ per period spent in prison. We assume that $\theta < \gamma_t$ for every t. The dynamics of life-cycle productivity in the model can therefore be summarized as

$$w_{t+1}^i = \begin{cases} \gamma_t w_t^i & \text{if } s_t^i \in \{\tilde{E}, E\}, \\ \theta w_t^i & \text{if } s_t^i = P. \end{cases}$$

For individuals that were never incarcerated, we define the cumulative productivity gain at age t as Γ_t, where $\Gamma_0 = 1$ and $\Gamma_{t+1} = \gamma_t \Gamma_t$.

2.3 The Police Production Function

The technology available to detect and apprehend criminals is motivated as follows. Consider a region R where, for simplicity, criminal opportunities are homogeneously distributed. All crimes occur within R. Police units are deployed in this region. A police unit is able to detect, with probability q, any crime committed within a circle of radius r^p centered at its position. Figure 1 illustrates the spatial distribution of detection probabilities for a given configuration of police units. Upon detecting a crime, the government incurs in a cost ζ_0 to apprehend the criminal.

Let k^{patrol} be the total amount of police expenditures devoted to patrolling. The number of police units is assumed to follow a Poisson distribution with parameter ξk^{patrol}, where ξ is a technological constant. This distribution can be interpreted as capturing the bureaucratic, legal, and logistical uncertainties intervening in the relationship between expenditures and

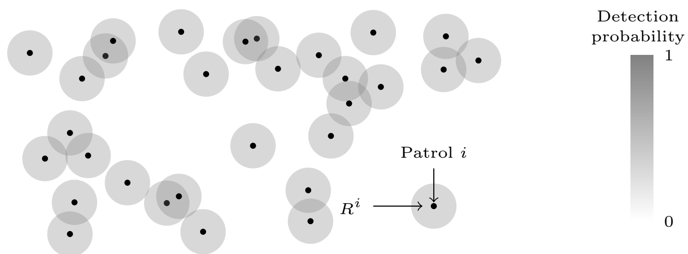

Figure 1: Apprehension probability as a function of the position for a given configuration of patrolling police units.

actual deployment of police forces. Conditional on the number of police units, each unit is deployed independently with a distribution D over R. The police command can choose the probability density D. While the police command chooses D in order to maximize the detection of crimes for a given k^{patrol}, each criminal chooses where to commit crimes in order to minimize the probability of being arrested.[7] This game has a unique Nash equilibrium where the police command sets the distribution D to be uniform—otherwise, all criminals would choose to commit more crimes in the areas where the density of patrols is the lowest—and criminals choose the location of crimes such that the density of crimes over R is constant—otherwise, the police command would choose a higher D in regions with higher incidence of crime. This equilibrium induces an apprehension probability given by the following expression:

$$p(v,k) = 1 - \frac{W_L(\zeta_2 v e^{-\zeta_1 k + \zeta_2 v})}{\zeta_2 v}, \tag{1}$$

where W_L is the Lambert-W function, v is the crime rate, k is the sum of patrolling and apprehension costs, and ζ_1 and ζ_2 are positive constants that depend on q, ζ_0, ξ, r^p, and the area of R.[8] The technical details related to the game describe above and to the derivation of this function and its properties are presented in Appendix A. Intuitively, we can interpret the parameters ζ_1 and $(\zeta_2)^{-1}$ as measures of the effectiveness of patrolling and of the procedures involved in the arrest of criminals, respectively.

This function displays a set of appealing properties. First, p is continuous and differentiable in the interior of its domain. This mirrors the intuition that small changes in police size and crime rate should be translated into small changes in the probability of apprehension. Second, for any v, $p(v,\cdot)$ is strictly increasing, with $p(v,0) = 0$ and $p(v,\infty) = 1$, which are properties imposed as starting points in previous papers using *ad hoc* apprehension technology functionals. Third, for any crime rate, the apprehension probability exhibits decreasing marginal returns, as intuitively proposed by İmrohoroğlu et al. (2000). Fourth, $p(\cdot,k)$ is decreasing for each k. Since a higher crime rate means more criminals to be caught and also diverts resources from patrolling to other procedures, the probability of apprehension of each individual crime is negatively related to the crime rate, as also suggested intuitively

[7] In equilibrium, the apprehension probability is the proportion of crimes that are detected.

[8] To ensure the continuity of p, we set $p(0,\cdot) = \lim_{v\to 0} p(v,\cdot) = 1 - e^{-\zeta_1(\cdot)}$.

by Ehrlich (1973) and Fu and Wolpin (2017).

2.4 Public Security Expenditures

We assume that the cost of incarceration per unit of time is linear on the number of prisoners. Thus, it costs κN^p per period to keep a mass of N^p agents incarcerated. This implies that, if expenditures on police are given by k, total expenditures on the public security system add up to $k + N^p\kappa$ per period. Public security expenditures are financed through the lump-sum tax f, raised over all agents in the economy.

2.5 The Criminal's Problem and the Stationary Equilibrium

We consider a stationary rational expectations equilibrium. Before defining it formally, though, we state the criminal's recursive problem from age $T + 1$ to age 0 and describe the intuition for its solution. By assumption, the value function of an agent past the terminal age $T + 1$ with productivity w is given by $\alpha(w - f)$. The state of an agent at age $t \leq T$ is determined by her/his productivity and whether she/he is free or in prison. A free agent chooses whether to engage in crime. Since agents at age T are always set free in the following period, the value function at that age for an agent with productivity w is $V_T^F(w) := w - f + \beta\alpha(\gamma_T w - f) + \max\{\overline{\Pi}_T, 0\}$ if she/he is free and $V_T^P(w) := b - f + \beta\alpha(\theta w - f)$ if she/he is incarcerated.

Now consider $t < T$. We define $V_t^F(w)$ as the value function of an agent at age t, with productivity w, who starts the period free, and $V_t^P(w)$ the analogous for an agent starting the period in prison. The dynamic programming problem of agent with productivity w at age t is, therefore

$$V_t^F(w) = w - f + \max\big\{\beta V_{t+1}^F(\gamma_t w), \overline{\Pi}_t + \beta(1 - \overline{p})V_{t+1}^F(\gamma_t w) + \beta\overline{p}V_{t+1}^P(\gamma_t w)\big\}. \tag{2}$$

Regardless of the decision to engage in crime, each free agent at age t earns income (from legal activities), pays the lump-sum tax, and experiences an increase in productivity by a factor γ_t. Agents who decide not to engage in crime will be certainly free at age $t + 1$. Agents engaging in crime have, in addition, an expected criminal gain of $\overline{\Pi}_t$, may evade apprehension and start next period free (with probability $1-\overline{p}$), or may be caught and start the next period in prison (with probability $\overline{p}$). Similarly, given b (the before-tax consumption equivalent of spending one period in prison), the value function for an incarcerated agent is

$$V_t^P(w) = b - f + (1 - \mu)\beta V_{t+1}^P(\theta w) + \mu\beta V_{t+1}^F(\theta w). \tag{3}$$

From this theoretical framework, one can fully characterize a dishonest agent's behavior. The solution for an agent at age T is trivial: engage in crime whenever $\overline{\Pi}_T > 0$. At age $t < T$, from Equation (2), an individual with productivity w strictly prefers to engage in crime if

$$\overline{\Pi}_t > \overline{p}(V_{t+1}^F(\gamma_t w) - V_{t+1}^P(\gamma_t w)).$$

The left-hand term expresses the expected utility gain from crime, which is independent of income. Expected losses from conviction include the foregone income in the legal sector and the permanent reduction in productivity. These losses are captured by $V^F_{t+1}(\gamma_t w) - V^P_{t+1}(\gamma_t w)$, so the right-hand term expresses the expected future utility losses from engagement in crime. Notice that both dimensions of expected losses from engaging in crime are increasing in the productivity level w. However, the net present value of marginal increases in productivity is larger for free agents than for those in prison. That is, the derivative of V^F_{t+1} is greater than the derivative of V^P_{t+1}, implying that $V^F_{t+1} - V^P_{t+1}$ is increasing in w. Thus, at each age, the agent's solution depends only on whether productivity is above an age-dependent cut-off level $w^\star_t$, determined implicitly from

$$\overline{\Pi}_t = \overline{p}(V^F_{t+1}(\gamma_t w^\star_t) - V^P_{t+1}(\gamma_t w^\star_t)).$$

This implies that an agent at age t strictly prefers to engage in crime if, and only if, her/his income is below $w^\star_t$. Agents at age t with income above $w^\star_t$ do not engage in crime.

Let $m^s_t(w)$, for $s \in \{\tilde{E}, E, P\}$, denote, respectively, the measures of dishonest agents free and not engaged in crime, free and engaged in crime, and incarcerated, for each productivity w and age t. In addition, recall that the average sentence length is given by $\lambda = 1/\mu - 1$.

Definition 2.1. *A stationary equilibrium, for a given choice of public security policy (k, λ), is a collection of value functions $V^P_t(w)$, $V^F_t(w)$, individual policy rules of engagement in crime by age and income, time-invariant measures of dishonest agents $m^s_t(w)$ for each w, $t \in \{0, 1, \ldots, T\}$ and state $s \in \{\tilde{E}, E, P\}$, an aggregate crime rate $v^\star$, an apprehension probability $p^\star := p(v^\star, k)$ (with p given by Equation 1), and a lump-sum tax f such that:*

(i) Individual and aggregate behavior are consistent, i.e., the aggregate crime rate is given by

$$\sum_{t=0}^{T} \int_0^\infty m^E_t(w)(1 - p^\star)(1 - e^{-\nu p^\star})/p^\star \, dw.$$

(ii) Individual and aggregate state changes are consistent.

(iii) Given the public security policy (k, λ), f and a belief in p, the individual policy rules of engagement in crime solve the individuals' dynamic program defined by equations (2) *and* (3) *(and the equivalent relations for age T).*

(iv) Expectations are rational, i.e., the common belief in the apprehension probability is $p^\star$.

(v) The government runs a balanced budget, so that revenues collected through the lump-sum tax f are fully used to finance the public security expenditures $k + \kappa \sum_{t=0}^{T} \int_0^\infty m^P_t(w)\, dw$.

Appendix B presents the formal details of the discussion developed in this section. It solves the criminal's problem, proves that this economy has at least one rational expectations equilibrium, and presents the idea behind the computation of the stationary equilibrium.

2.6 Welfare

Crime is the only source of inefficiency in this economy. It imposes costs on victims and criminals without creating value, so any first-best solution should imply zero crime and no expenditure on public security. The government can reduce the incidence of crime through the public security technology at its disposal. But public security expenditures also do not create net value, and can in principle be used for other purposes, so in reality they are part of the aggregate inefficiencies associated with the existence of crime. For this reason, the goal of public security policy should not be to minimize crime, but rather to minimize the aggregate welfare loss from crime.

Consider the stationary state and a social planner that puts equal weight on the consumption of all agents in this economy. Aggregate consumption is given by $\int c^i(k,\lambda)\mathrm{d}i$, where (k,λ) is the policy set by the social planner and $c^i(k,\lambda)$ is the consumption of agent i under this policy (remember that k represents expenditures on police and λ the average sentence length). Define $\widetilde{W}$ as the first-best outcome: the counterfactual aggregate consumption in a scenario where crime is not an option. The social planner's problem is to minimize the social welfare loss from crime, given by

$$L(k,\lambda) = \widetilde{W} - \int\limits_{\substack{\text{All}\\\text{agents}}} c^i(k,\lambda)\mathrm{d}i.$$

Implicitly, this means that the social planner seeks to minimize the loss of welfare caused by the criminal acts themselves, plus public security expenditures and punishment costs for criminals.[9] The loss directly caused by criminal acts is given by

$$L^{\text{crime}}(k,\lambda) := \sum_{t=0}^{T} v_t(k,\lambda)(d^V + d_t^C),$$

where $v_t(k,\lambda)$ is the equilibrium number of crimes committed by agents at age t, under policy (k,λ).

Total public security expenditures add up to $L^{\text{keep}}(k,\lambda) := \kappa N^P(k,\lambda)$, where $N^P(k,\lambda)$ is the equilibrium number of incarcerated agents under (k,λ). Given the balanced budget assumption, public sector expenditures are entirely financed by reduced consumption through the lump-sum tax.

In addition to increasing government expenditures, the punishment of criminals generates two other types of social loss: freedom deprivation and human capital depreciation. The former corresponds to the foregone consumption from being in prison. It is captured by the difference between each criminal's (legal sector) productivity and his consumption level while incarcerate (b). Integrating across all individual, this gives

$$L^{\text{FD}}(k,\lambda) := \int \sum_{t=0}^{T} (w-b)m_t^P(w)\mathrm{d}w = \sum_{t=0}^{T} \int w m_t^P(w)\mathrm{d}w - bN_t^P(k,\lambda),$$

[9]In Section 4, we consider alternative definitions of welfare.

where $N_t^P(k,\lambda)$ is the equilibrium number of prisoners at age t under (k,λ).

The human capital depreciation of an individual at age t, with productivity w, and who had productivity w_0 at age 0, is given by the difference between her/his maximum potential productivity had she/he never been incarcerated ($\Gamma_t w_0$) and her/his actual productivity at age t. Aggregating over all agents gives us

$$L^{\text{HK}}(k,\lambda) := \Big(\sum_{t=1}^{T}\Gamma_t + \Gamma_{T+1}\alpha\Big)\int wM_0 g_0(w)\mathrm{d}w - \int\sum_{t=1}^{T+1} w(m_t^{\tilde{E}}(w) + m_t^{E}(w) + m_t^{P}(w))\mathrm{d}w.$$

Combining all of these terms, the social planner's problem is to choose the policy

$$(k^\star,\lambda^\star) = \underset{(k,\lambda)}{\operatorname{argmin}}\left\{k + L^{\text{keep}}(k,\lambda) + L^{\text{crime}}(k,\lambda) + L^{\text{FD}}(k,\lambda) + L^{\text{HK}}(k,\lambda)\right\}.$$

3 Estimation

In this section, we describe our estimation strategy. Further details on specific numerical procedures are provided in Appendix D. Our theoretical model focuses on crimes with pecuniary motivations, so we consider only property crimes (defined as robberies, burglaries, motor vehicle thefts, and larcenies with value above 50 dollars).[10] We exclude low-value larcenies due to their low impact on social welfare—less then 0.1% of the social loss caused by property crimes—and more questionable strictly economic motivation. We also exclude fraud because white-collar crimes are very distinct from common street crimes, both in terms of their nature and detection technology.

In Section 3.1, we present a set of parameters taken directly from the literature and calculated from official datasets. We estimate the remaining parameters using a two stage GMM procedure, explained in Section 3.2. Section 3.3 assesses the fit of the model in terms of targeted and non-targeted moments.

We benchmark our model using 2004 as the reference year (or the year closest to 2004 for which specific variables are available). We choose this point in time because the number of prisoners and the crime rate were relatively stable in the first decade of the 2000s, coming closer to the stationarity hypothesis imposed in our numerical solution. Data from the National Prisoners Survey and the National Criminal Victimization Survey, discussed in Carson and Anderson (2016) and Truman and Langton (2014), for example, illustrate this point. We choose 2004 specifically due to a higher availability of data in this particular year.

3.1 Parameters Taken from the Literature and Official Sources

The unit of time in our estimated model is set to one year and monetary units are set to 2004 US$. Since individuals aged below 18 are mostly subject to correctional youth facilities,

[10] This is a different definition from the one given by the Bureau of Justice Statistics (BJS) and the Federal Bureau of Investigation (FBI). They consider robberies as a violent crimes and include all types of larceny as property crimes.

we assume that agents enter our economy at age 18 (corresponding to age 0 in the theoretical model). Due to the very low criminal activity and incarceration rate of individuals aged 65 and above, we set the terminal age T from the theoretical model to 47 ($= 65 - 18$). Based on the 2000 US Census, we set the number of individuals in each cohort between ages 18 and 65 to 3.6 million.

We set the same discount factor as Fella and Gallipoli (2014), 0.95, which falls in between the discount factors in İmrohoroğlu et al. (2004) and Mastrobuoni and Rivers (2017).

We recover the life-cycle productivity profile from a Mincer regression of log earnings on age, age squared, years of schooling, race, and gender, estimated using the 2004 Current Population Survey (CPS). The estimates for the coefficients of age and age squared are 0.0687 and -0.00121, respectively. Our estimate for the productivity $\widehat{\gamma}_t$ in the model is given by[11]

$$\widehat{\gamma}_t = \exp(\beta_{\text{age}} + 2\beta_{\text{age}^2} t + \beta_{\text{age}^2}).$$

We also use the CPS to calculate the productivity distribution of 18 year-old agents (age 0 in the theoretical model). We consider individuals who reported positive earnings and assume that income is distributed according to a lognormal with parameters μ_w (the expected value of the log of income) and σ_w (the standard deviation of the log of income). The maximum likelihood estimates for these parameters when yearly earnings are measured in hundreds of dollars are $\widehat{\mu}_w = 4.453$ and $\widehat{\sigma}_w = 0.7169$. These parameters imply that average yearly earnings for 18 year-old individuals in 2004 were \$12,290.

To estimate the multiplier α defining the continuation value at age $T+1$, we use the US 2004 life tables from the National Vital Statistics System. We assume that income remains constant after age T and that an agent with income w at age $T+1$ has a continuation value of αw. Then, using the age-specific mortality rates and the intertemporal discount rate of 0.95, we calculate the present discounted value of receiving a flow of income w for a 65 year-old individual. This leads to $\widehat{\alpha} = 11.38$.

Turning our attention to crime and public security policy, we first need to determine the total number of prisoners and their distribution according to type of crime. Prisoners in the US can be in State prisons, Federal prisons, or jails. According to the bulletin "Prison and Jail Inmates at Midyear 2005," a convicted prisoner can be in any of these facilities, but before adjudication individuals are mostly confined in jails. Using this bulletin and other BJS reports, we estimate the total number of incarcerated individuals, convicted and awaiting trial, for each type of property crime. Note that a fraction of inmates awaiting trial are eventually sentenced, which we must take into account in our calculations. Following Dobbie et al. (2018), we set this fraction to 0.578 and obtain the number of prisoners incarcerated for property crime in 2004: $\widehat{N}^P = 509,700$.

Next, we calculate the average cost of a property crime to victims. This includes the value of the goods stolen plus other costs, such as those associated with injuries, productivity losses,

[11]Given our Mincer regression, $\widehat{\ln w_t} = \text{constant} + \beta_{\text{age}} t + \beta_{\text{age}^2} t^2$. Therefore,

$$\widehat{\ln \gamma_t} = \ln(\widehat{\tfrac{w_{t+1}}{w_t}}) = \beta_{\text{age}}(t+1) + \beta_{\text{age}^2}(t+1)^2 - (\beta_{\text{age}} t + \beta_{\text{age}^2} t^2) = \beta_{\text{age}} + 2\beta_{\text{age}^2} t + \beta_{\text{age}^2}.$$

and psychological stress. Estimates of the total costs of crime to victims for different types of property crime are provided by Cohen (2000). To calculate the relative frequency of each type of property crime, we use data from the 2004 National Criminal Victimization Survey (NCVS). Combining the costs of crime from Cohen (2000) with these relative frequencies, we estimate the average property loss per property crime as \$848 and the average total loss as \$1,445.

In relation to expenditures, the Criminal Justice Expenditure and Employment Extract Program (CJEE) reports a total value of \$88.9 billion spent on police forces in 2004, while the value allocated to correctional facilities amounted to \$62 billion. Using the BJS 2004 prisoner count (2.385 million individuals), our estimate for κ is $\$26,000$ (cost per prisoner/year). This value is close to the $\$23,000$ used by Engelhardt et al. (2008).

Since we focus on property crimes, we need somehow to assign the fraction of total police expenditures devoted to detecting property crimes in particular (we denote this fraction by h). It is not clear that, in reality, police expenditures are separable into different types of crime. Nevertheless, our model needs a number that would capture the share of overall police resources allocated primarily to fight property crime. Since such a number is not available from official statistics, we assume some proportionality with the flow and stock of prisoners. BJS 2004 data on the universe of prisoners in more than 35 states indicates that 27.10% of prisoners were sentenced for property crimes, while 26.47% of new admission into the system are also related to property crimes (data from the National Corrections Reporting Program, ICPSR 36285). Thus, we set $\widehat{h} = 0.27$, which implies that expenditures on police associated with property crimes add up to $\widehat{k}_{2004} = 0.27 \times 88.9 = \$24$ billion. Given the uncertainties regarding this specific parameter, in our robustness exercises we re-estimate the model considering different values for h.

The same BJS dataset discussed in the previous paragraph allows us to estimate the average sentence of individuals convicted for property crimes in State prisons. Since this dataset does not include individuals who served time in jails, we make some additional assumptions. Once a prisoner awaiting trial is convicted and incarcerated, we assume that: (i) the probability of being sent to a prison equals the proportion of prisoners held in prisons over the total number of convicted prisoners; (ii) if sent to jail, the individual serves at most one year; (iii) if sent to a State prison, the time served is distributed as in the data for State prisons; and (iv) for all prisoners, time served is exponentially distributed. With these assumptions, we estimate an average sentence length $\widehat{\lambda}_{2004} = 1.482$ years (17.78 months), corresponding to a yearly probability of being released from prison $\widehat{\mu}_{2004} = 1/(1 + \widehat{\lambda}_{2004}) = 0.403$.

We take the human capital depreciation rate while incarcerated from Grogger (1995). He estimates the legal sector income loss for individuals between ages 18 and 25 who served time in prisons and jails (we weight these losses by the proportion of individuals in each type of facility).[12] Estimates suggest an average reduction in earnings of 21.6%. Since the average time in prison is 1.48 years, this corresponds to a 15% yearly reduction in earnings.

[12]We consider a long-term effect, since Grogger (1995) finds that income losses remain stable for at least seven years after release.

The average growth in earnings for free agents between ages 18 and 25 is approximately 6%, meaning that the average loss due exclusively to incarceration—as opposed to the loss of one year of learning by doing in the market— is $15\% - 6\% = 9\%$. Our estimate for θ is therefore given by $\widehat{\theta} = 1 - 0.09 = 0.91$.

Finally, we estimate the mass of potential criminals using Sampson and Laub (2003). They estimate the proportion of young individuals committing property crimes as a function of age using self-reported, parental-reported, and teacher-reported data. We make three assumptions to translate their numbers into a mass of individuals willing to commit crimes. First, most young boys committing property crimes face very low costs associated with punishment. Second, the direct cost of committing crimes is 0 at the age $\bar{t}$ at which the proportion of youngsters committing crimes is maximum. Note that these two assumptions imply that, at age $\bar{t}$, all dishonest agents are engaged in crime. Third, the ratio between the number of dishonest women and dishonest men is the same as the ratio between women and men being admitted to prison due to property crimes. Under these assumptions, the proportion of dishonest individuals is approximately 15%, which implies that $\widehat{M_0} = 0.15 \times$ (population at 18 years old) $= 540,000$.

3.2 Estimated Parameters

We assume that the age-dependent cost of committing crimes d_t^C is linear on t: $d_t^C = b^C t + c^C$. With this assumption, we have six parameters to estimate: b^C and c^C, the consumption-equivalent of one period in jail b, the average number of criminal opportunities per year ν, and the police production function parameters ζ_1 and ζ_2.

While, in principle, all six parameters could be estimated simultaneously, due to computational limitations we break the estimation into two stages. In the first stage, we estimate b^C, c^C, b, ν, and the detection probability, which we call p. Therefore, rather than estimating ζ_1 and ζ_2 directly together with the other parameters in the first stage, we estimate the value of p, which is a function of ζ_1 and ζ_2. Then, in the second stage, $\widehat{p}$ becomes one of the moments used as targets to estimate $\widehat{\zeta_1}$ and $\widehat{\zeta_2}$. With this two stage estimation procedure, we avoid working at the same time with a large space of parameters and a computationally demanding score: either the space of parameters is large (first stage) or evaluating the score is slow (second stage).

3.2.1 Moments and Estimation of the First-stage GMM

To estimate the first-stage parameters, we use as targets the following 19 moments: the number of prisoners for property crimes, the total number of property crimes, the value of 90 days of freedom, the proportion of 18-19 year-old individuals committing crimes, and 15 moments from the age distribution of prison inflow (from the proportion of admissions in prisons for each age group). For a more efficient estimation of the parameters, we also use the (estimated) variances of the moments. Since we take the moments from various distinct sources, we assume that the covariance matrix is diagonal.

In Section 3.1, we estimate the number of prisoners as 509,700, using 57.8% (from Dobbie et al. 2018) as the probability of a detained defendant being found guilty. We assume that this probability is independent for each inmate awaiting trial, so the distribution of the number of prisoners is binomial with a variance of $96,500 \times 0.578 \times (1 - 0.578) = 23,538$.

The total number and the variance of property crimes are taken from the NCVS 2004 Criminal Victimization report. The total number of property crimes is 15.032 million and, adding the squares of the standard errors of the estimate for each property crime, we obtain 187,683.71 as its variance.

We use the estimated value of 90 days of freedom reported by Abrams and Rohlfs (2007). They find a point estimate of \$1,050, with standard error of \$631. Therefore, the expected value of not getting caught is $1.48 \times 4 \times \$1,050 = \$6,222$, since the average sentence is 1.48 years (which implies a standard error of $\$3,736$). The equivalent of this moment in our model is the average, among people engaged in crime (below the cut-off productivity), of the difference between the value of being free and the value of being in prison:

$$\frac{\sum_{t=0}^{T} \int_0^{w_t^\star} (V_t^F(w) - V_t^P(w)) m_t^E(w) \mathrm{d}w}{\sum_{t=0}^{T} \int_0^{w_t^\star} m_t^E(w) \mathrm{d}w}.$$

To estimate the proportion of 18-19 year-old individuals committing crimes, we use the NLSY97. Individuals are asked whether they committed at least one crime in the last year. If we assume that all missing values in the questionnaire refer to individuals who did not commit crimes, we have a lower bound of 2.645% for the proportion of active criminals. If we assume that missing values refer to individuals who committed crimes, we have an upper bound of 3.353%. We assume that the number of active criminals follows a uniform distribution in the interval $[2.645\%, 3.353\%]$, so our estimate is the middle point in this interval: 2.999% (the uniform distribution means that we assume that we have as little information as possible for a random variable inside the interval). The variance of this uniform distribution is 3.5358×10^{-5}. The counterpart of this moment in the model is

$$\frac{1 - e^{-\nu p_{2004}}}{p_{2004}} \frac{\int_0^{w_0^\star} m_0^E(w) \mathrm{d}w + \int_0^{w_1^\star} m_1^E(w) \mathrm{d}w}{2 \times \text{population of a cohort}}.$$

The first term is the probability of receiving a positive number of criminal opportunities in a given year, while the second represents the proportion of individuals aged 18 and 19 who are willing to engage in crime.

The remaining moments are the fractions of prison admissions for 15 age groups. The Sourcebook of Criminal Justice Statistics provides number of arrests by type of crime and age for 2004. This information is available for each cohort up to 24 years old and, after that, for five-year age brackets (25-29, 30-34, etc.). We calculate the proportion of individuals arrested by age group (over all arrests of individuals aged 18 to 65). In the model, the proportions of arrests by age group correspond to the proportions of admissions into the system by age group. Given that, in the model, detection probabilities and criminal opportunities are constant across ages, the theoretical counterpart for each age group in this case is given by

$$\frac{\sum_{t \in \text{age group}} \int_0^{w_0^\star} m_0^E(w) \mathrm{d}w}{\sum_{t=0}^{T} \int_0^{w_0^\star} m_0^E(w) \mathrm{d}w}.$$

To estimate the variance of these fractions, we use the temporal variation between 2002 and 2011.

We assume that the covariance matrix Ξ_1 is diagonal with the variances mentioned above. We use the particle swarm numerical optimization to find the minimum score of the GMM using as weighting matrix $(\Xi_1)^{-1}$. To calculate the standard errors, we first compute the Jacobian matrix $\widehat{J_1}$ of the 19 moments with respect to the parameters, evaluated at the estimated parameters. Then, the covariance matrix of the estimated parameters is calculated as $\widehat{\boldsymbol{V}}_1 = (\widehat{J}_1'(\Xi_1)^{-1}\widehat{J}_1)^{-1}$.

3.2.2 Moments and Estimation of the Second-stage GMM

In the second stage of the GMM, we estimate the parameters of the police production function: ζ_1 and ζ_2. The moments we use are the estimated detection probability $\widehat{p}$ (point estimate of 2.30% and standard error of 0.06% from the first stage), the elasticity of property crime with respect to police size from Levitt (2002) (point estimate of -0.501, standard error of 0.235), and the elasticity of property crime with respect to prison population from Levitt (1996) (point estimate -0.261, standard error of 0.117).

To evaluate the elasticity of crime with respect to police expenditures k for any given pair (ζ_1, ζ_2), we calculate the numerical derivative of the crime rate with respect to k at $\widehat{k}_{2004}$, say $\widehat{\beta}_1(\zeta_1, \zeta_2)$. The elasticity of crime is given by $\widehat{\beta}_1(\zeta_1, \zeta_2)\widehat{k}/\widehat{v}$, with $\widehat{v}$ being the number of property crimes predicted by the model using the parameters estimated in the first-stage GMM.

The elasticity of property crimes with respect to prison population is the moment that remains to be evaluated in the model. We assume that if a few agents in prison were set free, most of them would commit crimes. In this case, the sum of agents engaged in crime and in prison is approximately constant, and the derivative of property crimes with respect to prison population is the opposite of the derivative of property crimes with respect to the number of agents engaged in crime. Suppose that the number of agents that could engage in crime increased by a multiplicative factor of $(1 + \epsilon)$. This would cause, as a first-order change, an increase of the crime rate by an additive factor of ϵv. In addition, there would be a second-order change due to a lower probability of detection caused by the load effect. The total approximate effect is the sum of the first- and second-order effects (equilibrium effects are negligible for small values of ϵ). This allows us to calculate the derivative of the crime rate on the number of criminals, which we define as $\widehat{\beta}_2(\zeta_1, \zeta_2)$. The elasticity we seek is thus $-\widehat{\beta}_2(\zeta_1, \zeta_2)\widehat{N}^P/\widehat{v}$, where $\widehat{N}^P$ is our prediction for the number of prisoners in the first stage of the GMM.

Let Ξ_2 be the diagonal matrix with elements given by the variances of the second-stage moments calculated in Section D.2.3. The weighting matrix of the second stage of the GMM is $(\Xi_2)^{-1}$. We use the particle swarm optimization algorithm to find the minimum score of the GMM. The covariance matrix $\widehat{\boldsymbol{V}}_2$ is obtained using the same procedure we used in the first stage.

3.2.3 Estimation Results

Table 1 reports the parameter estimates for our first-stage GMM, with respective standard errors in parentheses. Within the theoretical structure of the model, the estimates in

Parameter	Estimate	Unit
b^C	0.2031	hundreds of dollars/year
	(0.0080)	
c^C	3.2669	hundreds of dollars
	(0.3741)	
ν	10.0797	opportunities per year
	(0.6512)	
b	82.9433	hundreds of dollars/year
	(8.1641)	
p	2.30	%
	(0.06)	

Table 1: Parameters obtained from the first stage of the GMM. The standard error appears in parentheses below the estimate.

the table are those that best rationalize the patterns observed in the data. These estimates imply that the average subjective cost that 18 year-old individuals face to commit a crime has a monetary equivalent of roughly \$300. This cost increases by \$20 per year during the life-cycle. An individual engaged in crime receives, on average, 10 criminal opportunities per year. Also, according to the estimated model, one month in prison is comparable, in terms of welfare, to being free with monthly earnings of \$600. Finally, the probability that a property crime is detected is estimated to be 2.3%.

In the second stage, we estimate $\widehat{\zeta}_1 = 0.969$ (with standard error 0.023) and $\widehat{\zeta}_2 = 0.918 \times 10^{-6}$ (with standard error 0.420×10^{-6}). Though it is much more difficult to give an intuitively appealing interpretation to the magnitude of these two parameters, they lead to a police production function that, in equilibrium, reproduces the elasticities of crime estimated in the empirical literature (see discussion in the next subsection).

3.3 Model Fit

Table 2 presents the targeted moments with their model counterparts.[13] For the sake of conciseness, some of the moments related to the age-specific inflow of prisoners are aggregated into broader age groups. The table shows that the estimated model comes quite close to reproducing the targeted moments.

Moment	Data	Model	Unit
Inflow in prison 18-24	39.5	36.0	%
	(4.3)		
Inflow in prison 25-34	26.4	32.3	%
	(3.7)		
Inflow in prison 35-44	22.0	23.0	%
	(1.9)		
Inflow in prison 45+	12.1	8.7	%
	(2.9)		
Crime level	15032	15068	thousand crimes
	(433)		
Value of Freedom	62.2	63.2	hundred dollars
	(37.4)		
Proportion of criminals 18-19	3.00	3.16	%
	(0.61)		
Prisoners	509.7	505.9	thousand prisoners
	(0.0235)		
Elast. Police	−0.50	−0.58	–
	(0.24)		
Elast. Prisoners	−0.26	−0.28	–
	(0.12)		

Table 2: Comparison between targets and moments obtained from the model. The standard errors of a target is in parentheses between the target.

We also assess the performance of the model in terms of some untargeted moments. Using the NLSY97, the average number of property crimes committed in a given year by individuals that committed at least one crime is 7.3. The equivalent moment in our model is 9.0 crimes

[13] The moments generated by the model after the first and second stages are slightly different. This occurs because the estimated probability of detection $\widehat{p}$ is used as a moment in the second stage, but we do not force the model to set the $\widehat{p}$ implicit in the second-stage estimates to be identical to that estimated in the first stage. In fact, while in the first stage $\widehat{p} = 2.30\%$, we obtain implicitly $\widehat{p} = 2.27\%$ from the second-stage estimates. The model moments in the table are those obtained after our second-stage estimation.

per year.[14] This 20% difference may come, among other reasons, from the under-reporting of criminal activity by NLSY97 respondents.

In addition, we check how the model fits the distribution of prisoners across four age groups: 18-14, 25-34, 35-44, and 45^{+}. We can calculate these numbers directly from BJS prisoners data. Table 3 presents the numbers generated by the model and the data. The model reproduces the age-distribution of prisoners well. This may not seem particularly surprising, since we use the age-specific inflows of prisoners as targeted moments. But it actually shows that our assumptions of common probabilities of apprehension and average length of sentences, independent of individual characteristics, do not impose significant limitations. The model is flexible enough to reproduce a realistic age distribution of the stock of prisoners, even when only the inflows of prisoners are used as targets.

Age group	Data	Model
18-24	22.2	25.6
25-34	33.6	35.5
35-44	27.2	25.2
> 45	17.0	13.7

Table 3: Share of prisoners for property crimes by each age group, data versus model.

Finally, we estimate the lifetime probability of going to prison (due to property crimes) among all individuals in the economy. This probability can be estimated in our model as follows. First, take a productivity w_0 for an agent at age 0. Since we know the cutoff productivity and potential productivity gain at each age, we can calculate the number of years $Y_{w_0} = \sum_{t=0}^{T} \mathbb{1}_{\{\Gamma_t w_0 < w_t^\star\}}$ this agent engages in crime if never caught. Therefore, the lifetime probability that this individual is caught at least once is $1 - (1 - \overline{p})^{Y_{w_0}}$ (where our estimated $\overline{p}$ is 12.2%). As a consequence, the proportion of dishonest individuals who will be incarcerated at least once is $\int_0^\infty (1 - (1 - \overline{p})^{Y_w}) \mathrm{d}G_0(w)$, or 11.8%. This corresponds to 1.8% of the total population, including honest and dishonest agents.

According to BJS estimates (see Bonczar 2003), the implicit lifetime probability of going to prison calculated from 2001 numbers was 6.6%. However, this estimate includes prisoners for all types of crime and only Federal and State prisons (not jails). To obtain a number that is more directly comparable to our model, we do some back-of-the-envelope calculations, adjusting the BJS estimate by the fraction of prisoners incarcerated for property crimes (27%) and including jail inmates (which increases the number of prisoners in State and Federal prisons by 28%). If we assume that no convicted prisoner in jail ever ended up in

[14] From Section 2.2, our model predicts that the number of crimes committed by individuals engaged in crime, conditional on having at least one criminal opportunity, is $(1 - e^{-\nu p})/p(1 - e^{-\nu})$.

prison, the adjusted number is then $6.6\% \times 0.27 \times 1.28 = 2.3\%$. If, by contrast, we assume that most of the convicted individuals in jail were at some point incarcerated in State prisons, this adjusted number becomes closer to $6.6\% \times 0.27 = 1.8\%$. Both numbers are close to our estimate, and the latter matches it within one decimal place.

4 Quantitative Exercises

4.1 Socially Optimal Public Security Policy

This section presents the results from our policy experiments. Once estimated, our theoretical model can be used to evaluate any policy change, under whatever objective function is deemed adequate. We start by focusing on the most general welfare metric: aggregate surplus measured in consumption equivalent units, without distinction between criminals and victims. Later, we consider alternative objective functions, treating criminals differently from victims or positing crime minimization (subject to a public security budget) as the social goal. The objective of this section, rather than to focus on specific welfare metrics or to defend particular policies, is to illustrate the potential use of an equilibrium model to analyze the optimality of public security policies under different scenarios. Given the estimation strategy adopted, we conduct this analysis using parameters that are directly informed by the causal evidence from the applied micro literature.

Table 4 summarizes the 2004 socially optimal public security policy according to the estimated model. The rows list the key statistics, while the the first column presents the 2004 allocation, the second column presents the optimal policy according to the estimated model, and the third column calculates the percentage change in each variable when comparing the optimal policy to the 2004 allocation. For consistency, we use the 2004 values of the fitted model in this comparison rather than the actual data.

	2004 policy	Optimum	$\Delta\%$
Average time in prison	17.8 mth	22.3 mth	25.3
Expenditures on PSP	37.2 \$bn	33.1 \$bn	−11.0
Expenditures on police	24.0 \$bn	18.8 \$bn	−21.7
Expenditures on prison	13.2 \$bn	14.4 \$bn	9.1
Number of crimes	15.1 mi	16.9 mi	11.9
Total loss	63.9 \$bn	62.9 \$bn	−1.6

Table 4: Variables of interest generated by our model under the benchmark policy of 2004 (first column) and under the socially optimal policy (second column). The third column displays the percentage change between second and first column.

The optimal policy predicted by the model corresponds to a reduction of 11.0% in public security expenditures when compare do the 2004 baseline. This overall reduction masks movements in opposite directions along the two dimensions of policy. Police expenditures are reduced by 21.7%, while prison expenditures increase by 9.1% (due to an increase in average sentence lengths of 25.3%).[15] The optimal policy leads, maybe counterintuitively, to an increase of 11.9% in the property crime rate. In net welfare terms, the change in policy reduces the aggregate welfare loss from crime by 1.6%, or $1 billion.

The socially optimal policy predicted by the model illustrates the non-trivial interaction between the different dimensions of policy. Though the crime rate increases slightly in the optimum, this happens through a combination of apparently contradictory movements in police and sentencing. This comes from the fact that, at the margin, police expenditures in 2004 were inefficiently high, opening the possibility of significant savings that, in terms of aggregate welfare, more than compensate for the small increase in crime. At the same time, since the probability of catching a criminal is reduced, it becomes optimal to hold criminals in prison for longer when they are caught. The optimal policy also implies an increase in the 2004 prison population (for property crimes) of approximately 50,000 extra inmates (9.1%) and a roughly constant share of individuals being incarcerated at some point over their lifetimes (1.9% of each cohort against 1.8% under the 2004 policy).

Though our results related to sentencing and prison population may seem surprising in light of the high incarceration rate in the US, remember that we focus on property crimes, which accounted for only 20% of the total inmate population in the early 2000s. The strong reduction in police expenditures is an important driver of the increase in prison population under the optimal policy, since it incentivizes many more individuals to commit crimes. In fact, globally, the optimal safety policy is milder than that observed in 2004, which is attested by the fact that the equilibrium level of crime increases. In other words, among active criminals, fewer go through prison under the optimal policy, but there are more active criminals in the economy, so the overall number of individuals going through prison stays roughly constant and the prison population increases slightly.

Figure 2 illustrates the solution to the optimal public security policy in the (k, λ) plane. Each point on this plane represents a public security policy. The black lines represent social indifference curves according to our welfare function (the kinks in the indifference curves are due to the numerical approximations used to construct them). Notice that these indifference curves are closed, indicating that for each average sentence length there are two values of police expenditures resulting in the same welfare loss from crime (and vice-versa): one where expenditures on public security are low and the direct cost of crime is high, and another

[15] Notice that the reduction in police expenditures contrasts with the policy recommendation of hiring more police officers in Levitt (1997). Levitt uses his estimation of the elasticity of total crime on the number of police officers to conclude that the yearly welfare gain (due to fewer crimes) in dollars caused by hiring one police officer is greater than the yearly wage of an officer. It is clear that Levitt's methodology differs from ours in several aspects, but a sharp distinction is the fact that he considers all types of crimes, not only property crimes. Indeed, Levitt's estimates also lead to a negative net benefit of hiring an extra officer to fight *property* crimes (assuming an officer spends 27% of the total working time dealing with property crimes), as in our model.

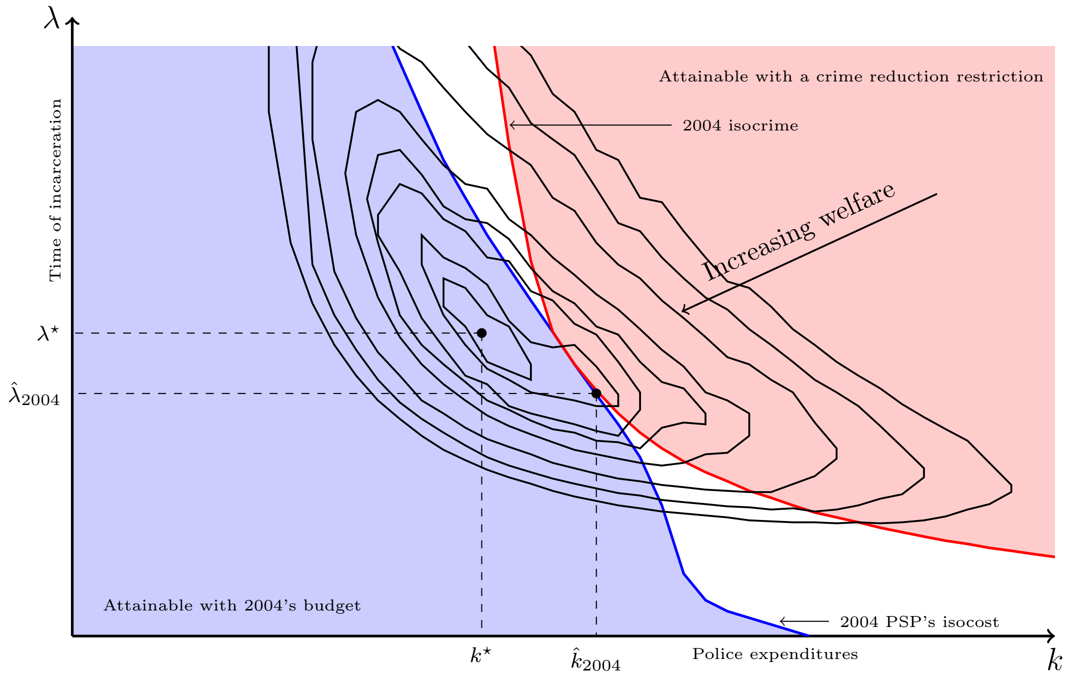

Figure 2: Contour lines and the political constraints of tax collection and crime level. The red region represents the set of policies leading to a crime reduction. The blue region is the set of policies leading to a tax-burden reduction.

where the reverse is true. As a result, there is a global optimum, denoted by $(k^\star, \lambda^\star)$. In the figure, the blue region indicates the set of policy pairs that are attainable under expenditure levels weakly smaller than that observed in 2004, while the pink region indicates policy pairs leading to crime rates weakly smaller than that observed in 2004 (the frontiers of these two sets, drawn in stronger colors, correspond to the respective equalities).

The figure indicates that, as discussed before, the socially optimal policy would lead to lower expenditures on police, longer average sentence, and higher crime rates than those observed in 2004. Interestingly, the figure also shows that, despite not maximizing social welfare, the 2004 allocation came close to solving another problem: minimizing the crime rate subject to the 2004 level of public security expenditures. The solution to this problem would be characterized by the tangency between the two sets. Later in this section, we characterize this point precisely. For now, it is enough to highlight that, despite not corresponding exactly to the solution to this problem, the 2004 allocation comes very close to it. This is remarkable since there is nothing in the structure of the model or in the estimation procedure that would mechanically push the quantitative implications in this direction. This result suggests that, despite not necessarily focusing on global welfare, the 2004 policy was based on some understanding of the determinants of crime and on some consideration of efficiency at the margin.

We also calculate confidence intervals for the optimal public security policy. To do that, we first use the Delta method to write down the covariance matrix of the optimal policy as a function of the covariance matrix of the parameters and the Jacobian of the optimal policy with respect to the parameters. Then, we apply the Implicit Function Theorem to evaluate this Jacobian using the fact that the optimal policy is a solution to the first-order conditions. Once we have the covariance matrix for the optimal policy, finding the error ellipse is straightforward. Appendix Figure A1 shows the optimal policy and the 2004 allocation in the (k, λ) plane, together with the estimated confidence interval. It shows that the estimated optimal policy is indeed significantly different from the 2004 allocation.

4.1.1 Decomposing the Changes in Welfare

In order to shed light on the sources of the gain in welfare as we move from the 2004 allocation to the social optimum, Table 5 presents a quantitative breakdown of the gains in welfare attributable to the different components of the social cost of crime. We classify the different sources of welfare costs into those associated with expenditures on the public security system, with the criminal event itself, and with the cost of punishing to criminals. There are a few points worth highlighting in this table. First, the key driver of the change in allocation is the high inefficiency of the size of the police force observed in 2004. The reduction in police expenditures is the main source of the welfare gains, being partly offset by losses along other dimensions. Second, the increased cost of maintaining the prison system with a larger population is modest in comparison to the revenues saved by the reduction in police expenditures. And third, criminals lose in the new equilibrium because of the increased human capital depreciation induced by the optimal policy. The considerably lower police size pushes agents with higher productivity to commit crimes. Thus, under the new optimal policy, prisoners with higher productivity would be subject to longer sentence lengths. In the end, the increased welfare cost directly associated with the criminal event itself is the main balancing force pinning down the social optimum. The model confirms the idea, popular with part of the political spectrum, that expenditures on police in the US are indeed above the optimal point, while suggesting that, in relation to property crime, average sentences are probably slightly shorter than they should be. A small increase in crime, accompanied by simultaneous adjustments along these two margins could, in principle, generate welfare gains by allowing for non-trivial reductions in overall expenditures in the public security system. In Subsection 4.1.3, we consider different welfare metrics and constraints on policy makers.

4.1.2 Comparative Statics on Police Expenditures and Average Sentence Length

In this subsection, we consider comparative static exercises where each dimension of policy is changed one at a time, while leaving the other dimension constant at its 2004 level. The goal of these exercises is to shed light on the mechanisms through which public security policies affect this economy and on the interaction across policies.

Figure 3 (a) shows how each component of the social cost of crime varies when we change police expenditures (k), holding the average sentence length constant at its 2004 value. As expected, the direct cost of criminal events is decreasing in k, while expenditures on the

All sources	1.0
Public Safety	4.0
Police	5.2
Prison	−1.2
Criminal events	−2.1
Punishment to criminals	−0.9
Freedom deprivation	0.2
HK depreciation	−1.1

Table 5: Welfare gain breakdown, in billion dollars.

public security system are increasing.[16] Notice as well that the criminals' welfare loss from punishment is non monotonic. Recall that the criminals' loss comes from human capital depreciation and freedom deprivation and is increasing in productivity. When k is low, the apprehension probability is also low, so criminals are better off because they are almost never incarcerated. In an intermediate region of k, there is still a moderately large number of agents with relatively high productivity committing crimes and being incarcerated, implying a high welfare loss. As k increases beyond this intermediate range, though, there are fewer agents committing crimes, and those who are committing crimes have very low productivity, so the loss becomes smaller.

In Figure 4, we conduct an analogous exercise looking at changes in average sentence length (λ), holding police expenditures constant at their 2004 level. The patterns and interpretation are similar to those in Figure 3 (a). Regardless of the strategy adopted to strengthen law enforcement—either through increases in the severity or certainty of punishment—, the components of the welfare loss from crime follow qualitatively similar patterns.

Figures 3 (b) and 4 (b) show the responses of the equilibrium crime rate to changes in policing and sentencing, respectively. The estimated model predicts that, at the margin, crime in 2004 was more responsive to changes in the probability than in the severity of punishment. This is captured in the figures by the much flatter curve around the 2004 allocation in the case of sentencing. This prediction agrees with a widely held belief in the economics of crime literature (see, for example, Grogger 1991, Polinsky and Shavell 2000). In the model, considering the 2004 allocation, this happens because the marginal effect of increases in sentence length on criminals' expected loss is relatively small in comparison to the effect of increases in the probability of punishment. Property crime rates are already substantially low and incarcerated individuals already have very low productivities, so in-

[16] Changes in police expenditures affect the number of criminals apprehended. So, for a constant average length of sentences, changes in police expenditures generally bring together changes in prison expenditures.

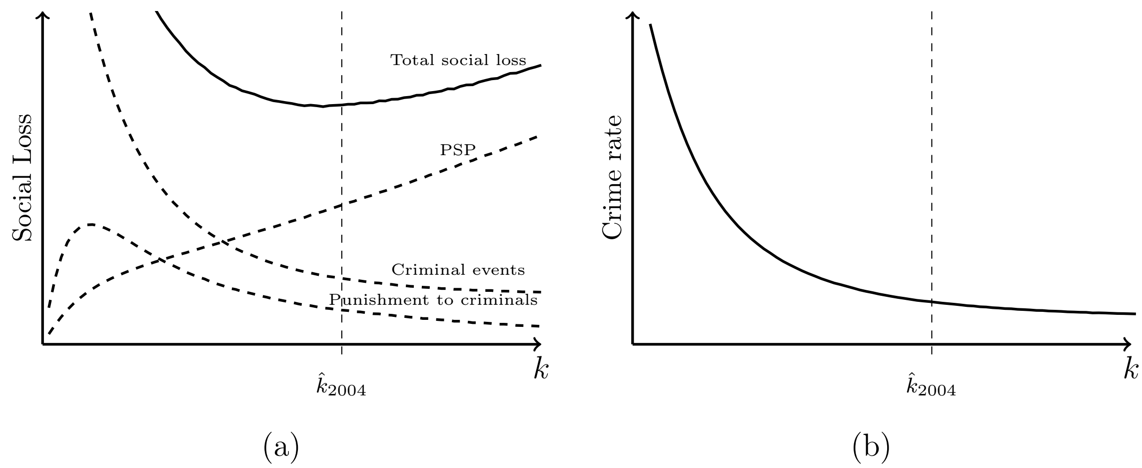

Figure 3: The behavior of some variables as k varies keeping the average sentence length at the value of 2004. (a) Components of the social loss and the total social loss. (b) Crime rate.

creases in sentence length do not alter the decision to participate in crime in a substantial way.

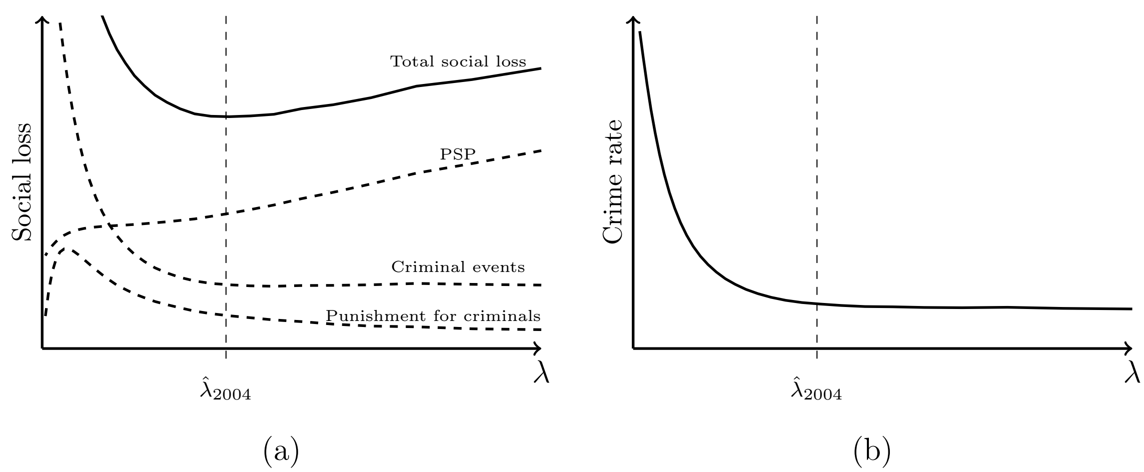

Figure 4: The behavior of some variables as μ varies keeping the average sentence length at the optimal level. (a) Components of the social loss and the total social loss. (b) Crime rate.

The main insight provided by these comparative static exercises, nevertheless, is related to the relevance of equilibrium considerations and the interaction across policies. Despite the fact that marginal increases in the probability and severity of punishment lead to reductions in crime (Figures 3 (b) and 4 (b)), once all relevant social costs are considered, the comparative static exercises recommend reductions along both dimensions (Figures 3 (a) and 4 (a)). This is clearly the case for police expenditures, but is also true for the average sentence length, though its 2004 level is much closer to the partial optimum. Looking at each dimension of policy separately, it would be optimal to slightly reduce the size of the police force and the length of sentences, and let crime increase a bit.

Now consider the globally optimal policy, as discussed in the previous subsection. Once the interaction across policies and equilibrium responses are taken into account, the optimum suggests reducing police expenditures but increasing the average sentence length. The best way to save on public security expenditures without letting the crime rate increase too much is through reduced police expenditures, partly compensated by increased sentences. The reduced probability of punishment increases the effectiveness of sentencing since it changes the pool of potentially active criminals (individuals at the margin of choosing crime). Therefore, the optimal combined response leads to an increase in the severity of punishment, contrary to the comparative statics exercise. This type of discussion requires an equilibrium model and could not possibly be conducted based exclusively on the marginal estimates available from the reduced-form empirical literature.

4.1.3 Robustness: Alternative Objective Functions and Police Expenditures

In this subsection, we assess how our quantitative and qualitative results change when we vary the social objective function and the imputation of the share of police resources allocated to fighting property crime.

Regarding the latter, in our benchmark estimation we set the share of police resources allocated to property crime proportionally to the share of inmates incarcerated for this type of crime. This proportionality is unlikely to hold precisely, so we consider variations of h around our benchmark value of 27%: from 21% to 33%, in 3 p.p. intervals.

The results from these exercises are presented in Table 6. When considering alternative values of h, we also re-estimate the 2004 model, given that expenditures on police in the benchmark economy are different. The key implication of variations in h refers to the implicit effectiveness of the police. Lower values of h mean that the police was more effective and cheaper in the benchmark economy, and therefore that the optimal policy would be closer to the 2004 allocation along the police expenditure dimension. The reverse is true for higher values of h. Nevertheless, the qualitative prescription that the socially optimal policy should slightly reduced police expenditures and increase sentence length remain true for all values of h in the table. The magnitude of these changes, though, is strongly increasing in h. With $h = 0.21$, for example, the 2004 allocation is much closer to the social optimum.

On the social objective function, our benchmark exercise considered a welfare metric that added up the monetary value of all the social costs of crime. This approach does not consider the good stolen as a net social loss, since it represents a transfer of resources from victims to criminals. It also incorporates direct utility losses to criminals associated with punishment as part of the social welfare function. In addition, it ignores any potential political constraint faced by policy makers. For example, policymakers may face political backlash from organized and politically powerful groups if they propose a policy that increases the equilibrium crime rate. If constraints such as this play a decisive role in the implementation of policies, and are taken as a fact of political life, we might want to analyze the second best solution that minimizes social losses while satisfying the political constraints.

h		2004 policy	Optimal policy
	Police expenditures (\$bn)	18.7	17.8
0.21	Sentence length (mth)	17.8	18.1
	Number of crimes (mi)	15.1	13.6
	Police expenditures (\$bn)	21.3	18.9
0.24	Sentence length (mth)	17.8	19.6
	Number of crimes (mi)	15.1	14.3
	Police expenditures (\$bn)	24.0	18.8
0.27	Sentence length (mth)	17.8	22.3
	Number of crimes (mi)	15.1	16.9
	Police expenditures (\$bn)	26.7	21.0
0.30	Sentence length (mth)	17.8	22.1
	Number of crimes (mi)	15.1	16.7
	Police expenditures (\$bn)	29.3	20.6
0.33	Sentence length (mth)	17.8	25.2
	Number of crimes (mi)	15.1	16.7

Table 6: Estimates for optimal policies for several values of h.

To illustrate how different social welfare metrics can be easily incorporated into the model, we look at two alternative cases. First, we consider a social welfare function that does not include most components of criminals' welfare. This metric ignores the value of goods stolen to criminals, the cost of committing crimes, and the direct utility cost from punishment, and includes only the criminals' lost legal labor market productivity due to incarceration. More specifically, for a policy (k, λ), define $L^{\text{prop}}(k, \lambda)$ as the loss due to the property stolen and $L^{\text{work}}(k, \lambda)$ as the loss due to a lower productivity of agents that are not incarcerated (see Appendix C for details of the derivation of L^{work}). Using the notation from Subsection 2.6, the welfare loss under this metric is given by $k + L^{\text{keep}}(k, \lambda) + L^{\text{crime}}(k, \lambda) + L^{\text{prop}}(k, \lambda) + L^{\text{work}}(k, \lambda)$.

Second, we consider an alternative social goal, defined by the minimization of crime subject to the 2004 public security budget.

	2004 policy	Social optimum (benchmark)	Social optimum (no criminals' welfare)	Minimum crime
Police expenditures ($bn)	24.0	18.8	21.2	23.4
Sentence length (mth)	17.8	22.3	19.5	19.2
Number of crimes (mi)	15.1	16.9	16.0	15.0

Table 7: Estimates for optimal safety policies (police expenditures in billion dollars and average sentence length in months) and the level of crime (in millions) for the three metrics we consider.

The results are presented in Table 7. First, as was already clear from Figure 2, the policy that minimizes crime subject to the 2004 public security budget is in fact quite close to the 2004 allocation. The two points are not identical, though, and the crime minimizing policy displays only slightly lower police expenditures and slightly longer sentences. The optimal policy according to the welfare function that ignores most dimensions of criminals' welfare is also closer to these two. This should not be surprising, since this metric ends up placing more weight on reductions in crime, since it increases the direct costs of criminal events and ignores part of the costs of punishment. Interestingly, the different welfare functions lead to the same qualitative recommendation of reducing police expenditures and increasing average sentence length, though the extent of these changes varies according to the optimality criterion used.

Notice that there is nothing built into the model forcing the optimal policy, in general, towards lower police expenditures and longer average sentences. Sufficiently low values of h, for example, generate optimal policies leading to increases in police expenditures and reductions in the average length of sentences, for any of the three objective functions considered before (results not shown in the table, but available from the authors upon request). Not surprisingly, the effectiveness of expenditures on police is key to determine the trade-off along these two dimensions of policy. Still, the qualitative result from our benchmark specification holds true within a large neighborhood of $h = 0.27$ and for different definitions of optimality. This indicates that, from the perspective of property crimes, the 2004 allocation seems indeed to reflect some degree of inefficiency from over-investment in police.

4.2 Incapacitation, Deterrence, Recidivism, and Load Effects

The empirical literature in economics of crime is often interested in the decomposition of the effects of public security policies into the specific channels affecting crime rates. Traditionally, this discussion has focused on the direct effect of incarceration on the physical capacity of individuals to commit crimes—the incapacitation effect—, on the effect of policies on the incentives of free individuals to commit crimes—the deterrence effect—, and on the effect of punishment on the likelihood that individuals commit additional crimes in the future—the recidivism effect.[17] In this subsection, we define each of these effects theoretically and illustrate their roles at different margins of public security policy. In addition, we argue that none of these effects can explain variations in crime due to the dilution of police resources when there is a larger number of active criminals. We define this new effect theoretically and call it the *load* effect.

Appendix E presents the theoretical definition of each of the four effects listed in the previous paragraph. Two points on these definitions deserve attention. First, to calculate the variation in crime, we look at the stationary states before and after a given policy change. Second, when comparing two different sets of policies, we calculate the role of a given effect by integrating the marginal changes through a path linking the initial policy to the new policy. This procedure is adequate if each policy leads to exactly one stationary equilibrium.

We illustrate the idea behind the definitions of the four effects using changes in police expenditures k (analogous arguments hold for small changes in sentence length). Following a small increase in k, the crime rate under the new equilibrium is a result of the combination of: (i) new beliefs of potential criminals about the number of police officers, holding constant their previous effectiveness (the deterrence effect); (ii) a different number of potential criminals incarcerated (the incapacitation effect); (iii) a new profile of productivity in the legal labor market among potential criminals due to the changed incarceration rate (the recidivism effect); and (iv) a different effectiveness of each police officer in catching criminals, given the equilibrium change in the number of active criminals (the load effect).

We explain intuitively how the model can be used to identify these four effects. For further details and a more precise discussion, refer to Appendix E. For the sake of exposition, assume for a moment that there is no load effect.[18] Consider a small change in policy from k to $k + \epsilon$, when ϵ is small and, to simplify, positive. This change induces small variations in the cutoff productivities, the mechanics of apprehension, and the productivity profile.

The deterrence effect refers to how criminals' beliefs about the policy affect participation decisions, which are captured in the model by the cutoff productivities. Hence, to measure

[17] We call it recidivism because incarceration depreciates productivity, increasing the likelihood of future criminal participation. In a context where incarceration increased human capital, it could be called rehabilitation effect.

[18] This implies that the probability of detecting a crime is $1 - \tilde{\zeta} e^{-\zeta_1 k}$ for some constant $\tilde{\zeta}$ and any level of police expenditures k.

deterrence in the model, we calculate the mass of free agents between the new and the old cutoff productivities under the old productivity profile. To find the corresponding variation in crime, we multiply this mass by the expected number of crimes each agent commits. Notice that we consider the old productivity profile to evaluate deterrence, even though the productivity distribution changes under the new policy. But since ϵ is small, the (small) change in productivity in the (small) group of individuals between the new and old cutoff productivities is a second-order change and has a negligible impact when evaluating the deterrence effect. A similar consideration can be made regarding the variation in the mass of free agents among this group of individuals. To evaluate the recidivism effect, we calculate the difference in the mass of individuals below the productivity cutoffs and multiply it by the expected number of crimes agents commit in a period (recall that the productivity distribution changes under the new policy).

To calculate the incapacitation effect, note first that increasing the police size affects incapacitation both intratemporally and intertemporally: it reduces the number of crimes a criminal commits before being caught; and it stops incarcerated individuals with low productivity from committing further crimes in future periods. Keeping fixed the productivity distribution and cutoff productivities (which capture the deterrence and recidivism effects), it is straightforward to calculate these two terms.

Following, we discuss the new *load effect*. This effect arises from the fact that an increase in crime rates reduces police effectiveness, which in turn affects crime. To calculate this effect, we first obtain the crime level under the new policy. The new crime rate increases the probability of detection to, say, p^{load}. Then, looking at the old economy and ignoring equilibrium effects (since they are of second order and we are considering small variations in k), we calculate the crime rate that a detection probability p^{load} would generate.

We decompose numerically the effect of changes in k into its four underlying components according to the logic outlined above. Figure 5 presents the results. The figure shows the marginal deterrence, incapacitation, recidivism, and load effects on crime rates. Negative values indicate marginal reductions in crime.

The signs of most marginal effects are as expected, indicating reductions in crime coming from the four different channels. For lower values of k, though, the marginal recidivism effect is positive. On the one hand, increasing k results in more arrested criminals among those not deterred by the higher k, which translates into lower future earnings for this group, pushing more agents under the productivity cut off. On the other hand, agents who were deterred do not suffer anymore the productivity losses from incarceration, which pushes more agents above the productivity cut off. The first effect dominates for lower value of k, since the probability of apprehension is extremely low, leading to a negative recidivism effect.

Overall, the marginal effects become smaller in magnitude as k increases and the marginal efficacy of police expenditures is reduced. The deterrence and load marginal effects, in particular, tend to be the strongest when starting from very low apprehension probabilities (low value of k), but become virtually null for a high enough k. When the size of the police force is sufficiently large, and the associated crime rate low, there is little impact on police effectiveness from changes in crime. Combined with the distribution of productivity in society, this can also eventually lead to a margin over which potential criminals do not respond much to small changes in the probability of apprehension. Or, in other words, from this point on there are few marginal criminals, so deterrence is low and all that remains is the incapacitation effect.

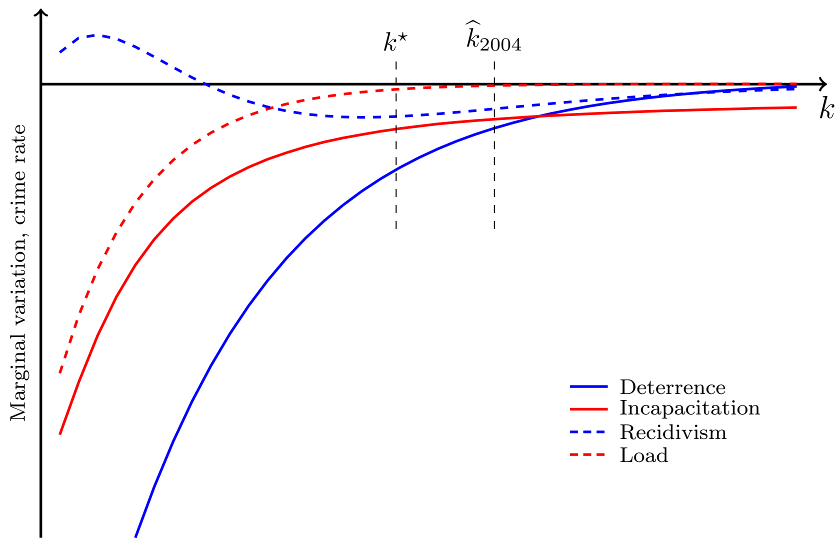

Figure 5: Marginal effect decomposition of the crime rate in k.

The previous discussion also highlights a point that is somewhat obvious theoretically, but that is worth mentioning explicitly. In reality, when considering specific policy changes, none of the classic effects ever exists in isolation. In equilibrium, any policy change that affects crime rates inevitably carries all the different components at the same time, though with different weights depending on the policy margin in question.

4.3 Schooling as an Instrument for Crime Control

A large academic literature in economics points to the role of education and social policies as protective factors against criminal involvement (e.g., see Lochner 2004, 2007, Foley 2011). Similarly, in the public debate, it is often argued that reallocating part of the public security budged to schooling and social protection would be a more effective and humane way to control crime. Nevertheless, once again, without an equilibrium framework capable of bringing together these different dimensions, it is impossible to analyze the potential trade-offs present in alternative policy choices.

In this subsection, we analyze the potential welfare gains from incorporating schooling as an additional dimension of policy and allowing part of the original 2004 public security budget to be reallocated to it. For simplicity, we assume that the direct effects of additional investments in education can be fully characterized by two sufficient statistics: the average internal rate of return and the degree of focalization. These two parameters determine the net increase in aggregate income from investments in schooling and the degree to which this increase is concentrated on the bottom of the productivity distribution (or spread more evenly across productivity levels).

We assume that educational policies generate parallel shifts in the lifecycle earnings profile in all ages. Consider an individual with productivity w at age 0 (i.e., 18 years old in the data). We assume that additional investments in education increase productivity at age 0 to $w(1 + \eta e^{-\psi w})$, where $\psi > 0$ is a free parameter and $\eta > 0$ is a constant to be obtained from the data (as a function of ψ and the internal rate of return). This functional form has the convenient feature of making the relative productivity gain of less productive individuals always weakly larger than that of more productive individuals. In particular, the relative productivity gain for the least productive individuals is approximately η. The parameter ψ captures the degree of focalization: how concentrated on the less productive individuals the benefits of the policy are (possibly due either to political or technological reasons). We call ψ the focalization of the educational policy. When focalization is high, individuals at the bottom of the distribution of productivity benefit more from investment in education. In the extreme case of $\psi = 0$, the relative productivity gain of all agents is the same, while when ψ tends to infinity gains become increasingly concentrated on the least productive individual.

Regarding the average internal rate of return (IRR), we consider two alternative scenarios. The first one assumes an internal rate of return equal to zero, meaning that education works simply as a cash transfer to individuals, mimicking a social protection program. The second one assumes an internal rate of return of 15%, roughly the midpoint among the estimates presented in Heckman et al. (2008). We calibrate the parameter η in the function above to match these two internal rates of return.

More specifically, let e be the additional investment in education at each period and r the internal rate of return. In other words, in the stationary state, investing e in education increases aggregate productivity by $(1 + r)e$. Given a focalization ψ, the parameter η is pinned down by this aggregate productivity gain $(1+r)e$. Appendix F explains in detail the calculations and specific assumptions needed for us to numerically arrive at this number.

With ψ, r, and the η determined implicitly from them in hand, we can simulate the aggregate effect of any increase in educational expenditures e. Our quantitative exercises in this section consider the triplet (k, λ, e), with a non-negative e, as the policy space, rather than simply the pair (k, λ). In Subsection 4.1, we calculated the welfare gain from moving the 2004 policy from $(k_{2004}, \lambda_{2004}, 0)$ to the social optimum $(k^\star, \lambda^\star, 0)$. We now ask, for given ψ and r, what the additional welfare gain would be from moving from $(k_{2004}, \lambda_{2004}, 0)$ to $(k^{\star\star}, \lambda^{\star\star}, e^{\star\star})$, a new social optimum that allows for additional educational expenditures $e^{\star\star}$. This new policy $(k^{\star\star}, \lambda^{\star\star}, e^{\star\star})$ is subject to the constraints that $(k^{\star\star}, \lambda^{\star\star})$ must cost less than $(k_{2004}, \lambda_{2004})$ and that $e^{\star\star}$ must be financed entirely by the difference between these public security costs. When considering a positive internal rate of return from investments in education, we compute the additional welfare gain ignoring the direct effect of increased productivity on welfare through increased consumption (we consider only its indirect effect on the welfare cost of crime).

Additional investments in education can reduce crime directly by improving the earning profiles of individuals and therefore increasing the opportunity cost of crime. The extent to which this mechanism is strong enough to significantly reduce crime depends on whether the benefits are concentrated on marginal agents, those who are close to indifference regarding criminal involvement. Further welfare gains from investments in education can, in principle, be also obtained from adjustments in the public security policy (k, λ), given the reduction in the criminogenic potential of the population. Notice that, in this setting, it might be possible to generate welfare gains from investments in education e, while at the same time reducing public security expenditures in (k, λ), without actually changing the public security policy pair (k, λ). This could happen because the additional investment in education can reduce crime, therefore potentially reducing the number of prisoners and the budget allocated to prisons even with a constant λ.

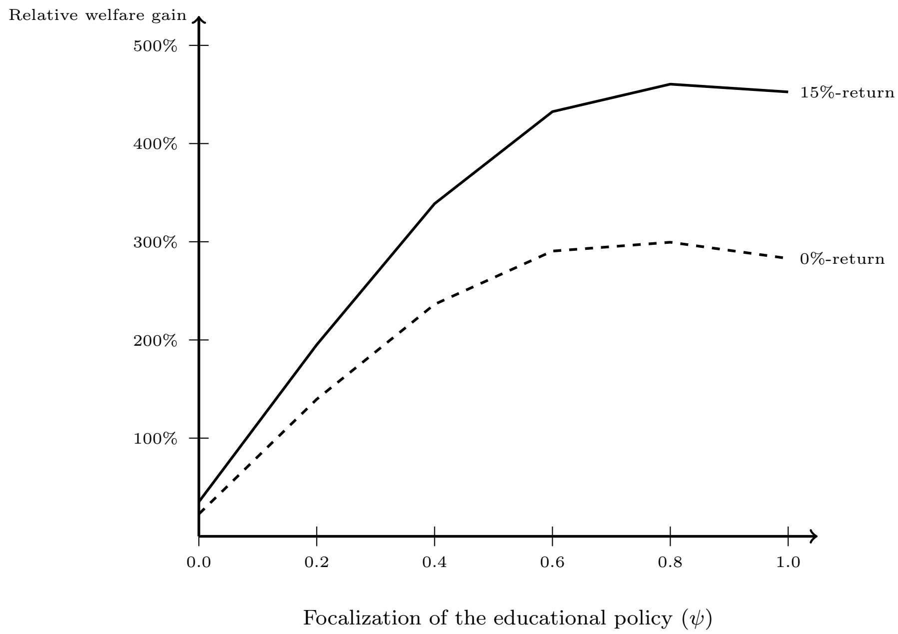

Figure 6: Relative welfare gain when we compare optimal policies with possibility of PSP defunding (to fund education) and without this possibility, as in Section 4.1. The dashed curve represents an IRR in education of 0%, while the solid curve represents an IRR of 15%. We exclude the direct gains from returns with education.

Figure 6 presents the main results from our exercises for various focalizations levels and for the two internal rates of return. The figure shows that, even when schooling works simply as a cash transfer program ($IRR = 0$) and under very poor focalization, reallocating some expenditures from public security to schooling can be welfare enhancing. But these gains are limited, of the order of 23% in terms of the welfare gains from Subsection 4.1 (\$1.2 billion gain, or 1.92% reduction in the welfare cost of crime, in comparison to the 2004 allocation). With poor focalization, gains are limited even under a 15% internal rate of return, reaching 35% more than the values calculated before (\$1.35 billion gain, or 2.12% reduction in the welfare cost of crime, in comparison to the 2004 allocation). This is the case because educational investments are somewhat ineffective as a crime reducing tool in this scenario, since expenditures have to be spread over the entire population in order to generate relatively small increases in the opportunity cost of crime for marginal individuals. Still, it remains true that some improvement upon the previous optimum is possible.

The main power of education, though, comes from the possibility of focalization. As educational investments become more concentrated on marginal individuals at the lower bottom of the productivity distribution, potential welfare gains increase strongly, reaching close to 300% for the case with $IRR = 0\%$ and 460% for $IRR = 15\%$. In comparison to the 2004 allocation, these correspond to welfare gains of the order of $4 billion and $5.6 billion, respectively, and to reductions in the welfare loss from crime of 6.3% and 8.8%. These large increases in welfare gains come entirely from improved effectiveness under better focalization. In this situation, it becomes much cheaper to increase the opportunity cost of crime for the marginal individuals, therefore maximizing the crime reducing impact of a given increase in educational expenditures. The last part of the curves in Figure 6 are non-monotonic because, from some point on, the educational policy becomes too concentrated on infra-marginal individuals at very bottom of the productivity distribution (who participate in crime anyway), therefore reducing a bit its effectiveness.

To illustrate what is happening to the optimal policy at different levels of focalization and internal rates of return, Figure 7 plots the change in some key variables at the optimal policy (compared to the optimal policy from Subsection 4.1). First, in order to fully characterize the educational policy, Panel (a) plots the relative productivity gain of the least productive individual under the different scenarios. Panel (b) shows that more focalized educational policies are associated with substantial reductions in the optimal crime rate. These reductions tend to be also higher, for most part, under higher internal rates of return. Panel (c), in turn, documents that reductions in public security expenditures under the new optimum are, not surprisingly, increasing in focalization and in the internal rate of return. Panel (d) shows that these changes are associated with major reductions in the probability that individuals go through the penal system at some point in their lives, particularly so for the case with $IRR = 15\%$, reaching above 20% for the highest levels of focalization.

Increased educational investments in this economy, when combined with adequate focalization, greatly increase the opportunity cost of marginal individuals. This reduces the pool of both active and potential criminals (around the relevant cut-offs), reducing crime rates. Panel (e) in Figure 7 shows that this leads to major reductions in the optimal expenditures in police, an important source of the additional welfare gain obtained in this setting. Nevertheless, as ψ and r increase, only very inelastic (infra-marginal) individuals remain engaged in crime—dishonest agents with very low productivity—, so the new optimum predicts a non-trivial increase in average sentence lengths (Panel (f)). This is the case because these individuals are likely to continue committing crimes in the near future and, at the same time, have low social cost of punishment (given their low productivity levels).

In short, this section lends support to the view that, by combining public security and educational policies in a smart way, it may be possible to reduce crime rates and the welfare costs of crime while at the same time reducing the size of police forces and punishing a smaller

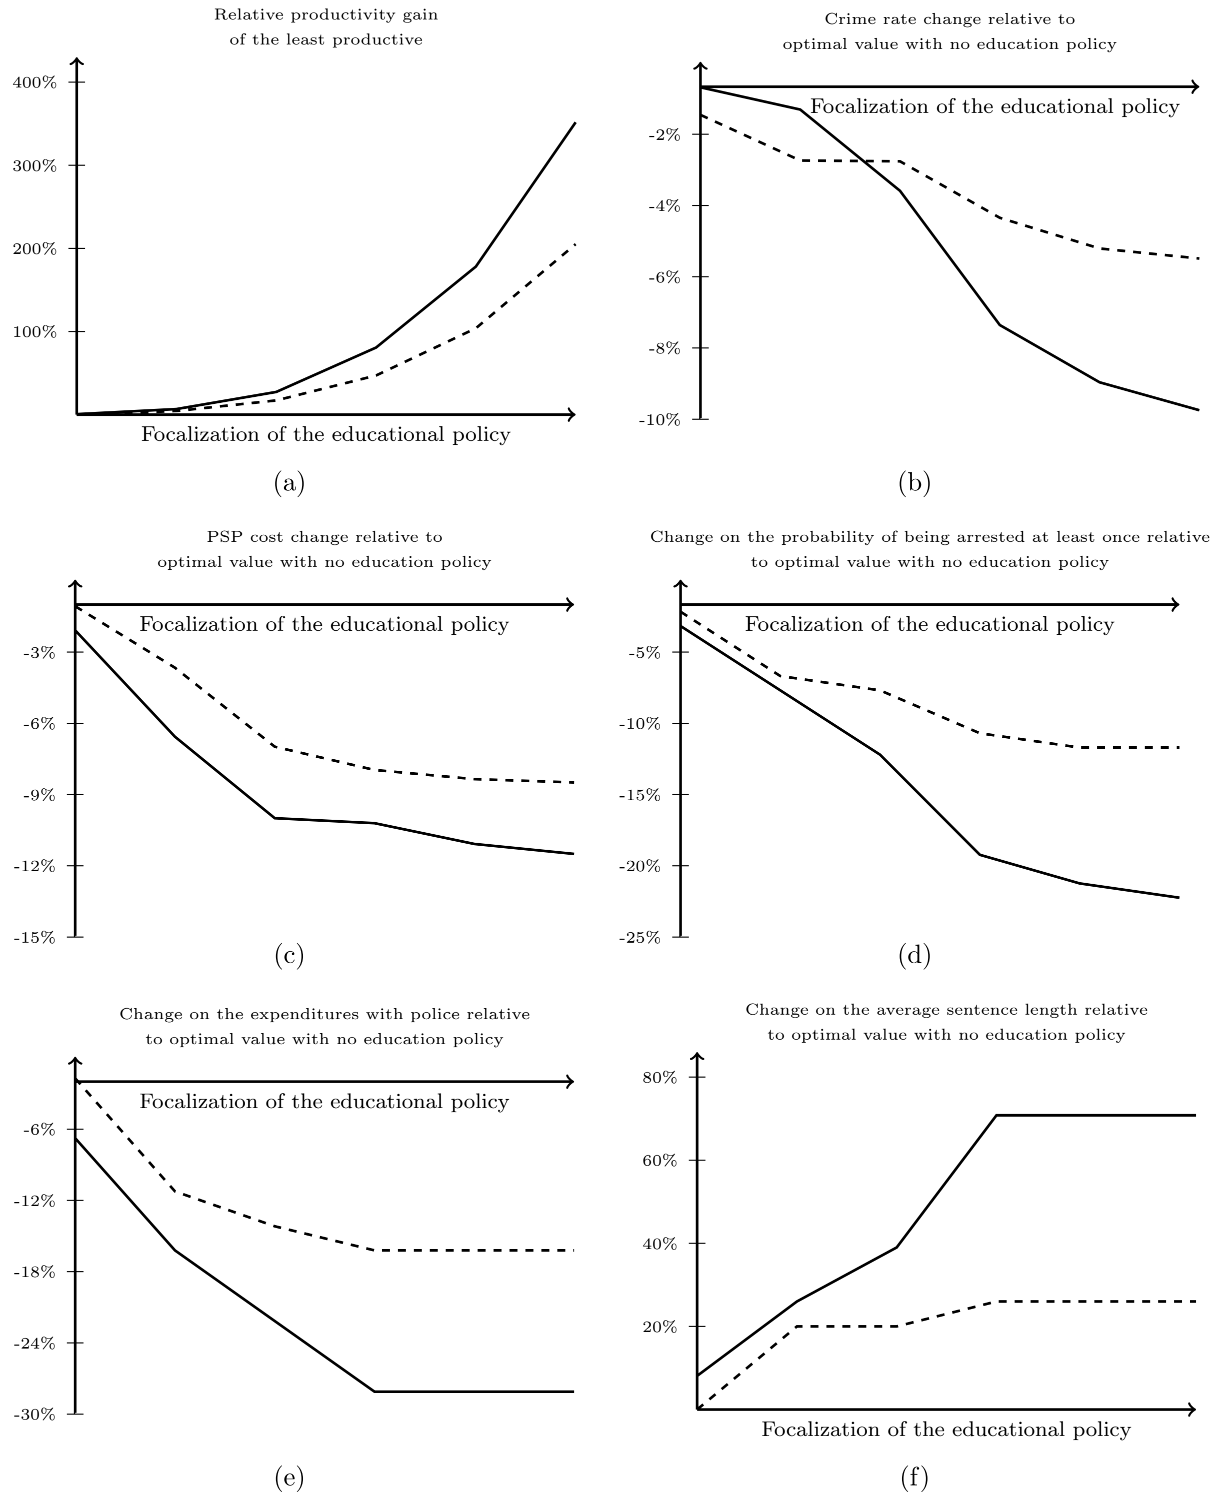

Figure 7: Relative changes in selected variables when we compare optimal policies with possibility of PSP defunding (to fund education) and without this possibility, as in Section 4.1. The dashed curves represent an IRR in education of 0%, while the solid curve represents an IRR of 15%. We exclude the direct gains from returns with education.

number of individuals. The potential benefits from such combination of policies depend to a great extent on the rate of return to investments in education and on the focalization of the educational policy. If educational investments have large returns and can be focalized on marginal individuals, potential welfare gains are large. Still, even under the best-case scenario, the new optimum displays longer average sentences, given that individuals who remain active criminals are very inelastic and have low productivities.

5 Conclusion

This paper proposes a framework that can be used to perform a comprehensive analysis of the welfare impact of public security policies. We rely on a life-cycle equilibrium model to capture the various non-trivial responses and dynamic implications of crime policies. We estimate the model using 2004 US data and solve numerically for the socially optimal policy, considering different welfare metrics and political constraints.

A key novelty of our model is the way we deal with the police production function. Whereas previous papers in the literature have used ad hoc functional forms, we give micro-foundations to a functional representation of the apprehension probability. The resulting function possesses all the properties deemed important by the crime literature.

In addition, our estimation strategy establishes a close link between the theory and the empirical evidence from natural experiments. The parameters in our apprehension technology are directly identified from the causal estimates of the effect of police and prison population on crime available from this literature.

Though our main focus is methodological, our quantitative results also provide some insight. They suggest that the 2004 expenditures on police in the US were too high and crime rates were too low. Saving on police expenditures, even with some increase in crime, could therefore lead to increases in social welfare. However, the estimated model suggests that the 2004 allocation came close to attaining another objective: minimizing the crime rate subject to the 2004 public security budget. These different results signal to the inescapable importance of a normative discussion on the social objective of criminal justice systems. Our model is flexible enough to incorporate any welfare dimension considered relevant in the evaluation of public security policies. We therefore see our quantitative exercises more as illustrations of the possible uses of our framework, rather than as definitive policy recommendations. Subjective considerations of fairness, for example, do not appear in any of the exercises developed in the paper, but still may be relevant in designing an optimal public security system.

The estimated model is also used to shed light on the potential role of schooling as a crime reducing policy. We show that, when educational investments have high returns and can be targeted to individuals at the margin of engaging in crime, welfare gains can be substantial. In these situations, shifting resources from public security to educational investments can lead to reductions in crime together with reductions in police expenditures and in the probability of lifetime incarceration.

Some limitations of our analysis suggest potential avenues for further research. First, due to data challenges, we restrict our estimation to property crimes. Including other types of crime could capture substitution effect across alternative crimes. Second, individual-level georeferenced data could be used to push forward the understanding of the micro-founded police production function proposed here. Third, other components of human capital could be incorporated in the model, such as productivity gains in the criminal sector and the possibility of public expenditures targeted at reducing recidivism. Finally, it would be interesting to incorporate an explicit model of the functioning of the justice system and the sentencing process, particularly regarding the possibility of wrongful convictions and the effect of the celerity of punishment. These are important topics that we hope will be addressed in future contributions to the literature.

References

Abrams, D. and Rohlfs, C. (2007), Optimal bail and the value of freedom: Evidence from the philadelphia bail experiment. John M. Olin Program in Law and Economics Working Paper No. 343.

Beccaria, C. (1963), *On crimes and punishments*, (introduction by H. Paolucci, Trans.) New York: Macmillan. (Original work published 1764).

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

Beirne, P. (1993), *Inventing Criminology: Essays on the Rise of Homo Criminalis*, State University of New York Press, Albany.

Bentham, J. (2000), *An Introduction to the Principles of Morals and Legislation*, number v. 1 *in* 'An Introduction to the Principles of Morals and Legislation', W. Pickering. (Original work published 1781).

Bonczar, T. P. (2003), Prevalence of imprisonment in the u.s. population, 1974-2001, Technical report, U.S. Department of Justice, Office of Justice Programs, Bureau of Justice Statistics.

Buonanno, P. and Raphael, S. (2013), 'Incarceration and incapacitation: Evidence from the 2006 italian collective pardon', *The American Economic Review* **106**(6), 2437–65.

Burdett, K., Lagos, R. and Wright, R. (2003), 'Crime, inequality and unemployment', *American Economic Review* **5**(93), 1764–77.

Carson, E. A. and Anderson, E. (2016), Prisoners in 2015, Technical report, U.S. Department of Justice, Office of Justice Programs, Bureau of Justice Statistics.

Cohen, M. A. (2000), 'Measuring the costs and benefits of crime and justice', *Journal of Criminal Justice* **4**, 263–315.

Cordner, G. (2020), *Evidence-based Policing in 45 Small Bytes*, U.S. Department of Justice, Office of Justice Programs, National Institute of Justice.

Di Tella, R. and Schargrodsky, E. (2004), 'Do police reduce crime? estimates using the allocation of police forces after a terrorist attack', *The American Economic Review* **94**(1), 115–33.

Dobbie, W., Goldin, J. and Yang, C. S. (2018), 'The effects of pretrial detention on conviction, future crime, and employment: Evidence from randomly assigned judges', *American Economic Review* **108**(2), 201–40.

Draca, M., Machin, S. and Witt, R. (2008), Panic on the streets of london: Police, crime and the july 2005 terror attacks. CEP Discussion Paper no. 852.

Ehrlich, I. (1973), 'Participation in illegitimate activities: A theoretical and empirical investigation', *Journal of Political Economy* **81**, 521–65.

Ehrlich, I. (1981), 'On the usefulness of controlling individuals: An economic analysis of rehabilitation, incapacitation and deterrence', *American Economic Association* **71**(3), 307–22.

Engelhardt, B., Rocheteau, G. and Rupert, P. (2008), 'Crime and labor market: A search model with optimal contracts', *Journal of Public Economics* pp. 1876–91.

Fella, G. and Gallipoli, G. (2014), 'Education and crime over the life cycle', *Review of Economic Studies* **81**, 1484–517.

Ferraz, E. (2017), Socially optimal crime and punishment: on optimal penal codes. Work in progress.

Flinn, C. (1986), *Criminal Careers and 'Career Criminals'*, Vol. 2, National Academy Press, Washington DC.

Foley, C. F. (2011), 'Welfare payments and crime', *The Review of Economics and Statistics* **93**(1), 97–112.

Fu, C. and Wolpin, K. I. (2017), 'Structural Estimation of a Becker-Ehrlich Equilibrium Model of Crime: Allocating Police Across Cities to Reduce Crime', *The Review of Economic Studies* **85**(4), 2097–2138.

Gaviria, A. (2000), 'Increasing returns and the evolution of violent crime: the case of colombia', *Journal of Development Economics* **61**(1), 1–25.

Gleaser, E., Sacerdote, B. and Scheinkman, J. A. (1996), 'Crime and social interactions', *Quarterly Journal of Economics* **111**(2), 507–48.

Grogger, J. (1991), 'Certainty vs. severity of punishment', *Economic Inquiry* **29**(2), 297–309.

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

Hannon, L. and DeFranzo, J. (1998), 'Welfare and property crime', *Justice Quarterly* **15**, 273–87.

Hayes, T. O. (2020), Assessing calls to defund the police: Police budgets and employment levels, Technical report, American Action Forum.

Heckman, J., Lochner, L. and Todd, P. (2008), 'Earnings functions and rates of return', *Journal of Human Capital* **2**(1), 1–31.

İmrohoroğlu, A., Merlo, A. and Rupert, P. (2000), 'On the political economy of income redistribution and crime', *International Economic Review* **41**(1), 1–25.

İmrohoroğlu, A., Merlo, A. and Rupert, P. (2004), 'What account for the decline in crime?', *International Economic Review* **45**(3), 707–29.

Jacob, A. B. and Ludwig, J. (2010), The effects of family resources on children's ourcomes. Working Paper, University of Michigan.

Laub, J. H. and Sampson, R. J. (2003), Criminal careers and crime control: a matched-sample longitudinal research design, phase i, Technical report. Final Report (87-IJ-CX-0022).

Levitt, S. D. (1996), 'The effect of prison population size on crime rates: Evidence from prison overcrowding litigation', *The Quarterly Journal of Economics* **111**(2), 319–51.

Levitt, S. D. (1997), 'Using electoral cycles in police hiring to estimate the effect of police on crime', *The American Economic Review* **87**(3), 270–90.

Levitt, S. D. (2002), 'Using electoral cycles in police hiring to estimate the effects of police on crime: Reply', *The American Economic Review* **92**(4), 1244–50.

Lochner, L. (2007), 'Individual perceptions of the criminal justice system', *The American Economic Review* **97**(1), 444–60.

Lochner, L. J. (2004), 'Education, work and crime: a human capital approach', *International Economic Review* **45**, 958–79.

Mastrobuoni, G. and Rivers, D. A. (2017), Criminal discounted and deterrence.

Mocan, H. N., Billups, S. C. and Overland, J. (2005), 'A dynamic model of differential human capital and criminal activity', *Economica* **72**(288), 655–681.

Munyo, I. (2015), 'The juvenile crime dilemma', *Review of Economic Dynamics* **18**(2), 201–211.

Owens, E. G. (2009), 'More time, less crime? estimating the incapacitive effect of sentence enhancements', *The Journal of Law and Economics* **52**, 551–79.

PEW (2016), The punishment rate—new metric evaluates prison use relative to reported crime, Technical report, The PEW Charitable Trusts.

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

Sampson, R. J. and Laub, J. H. (2003), 'Life-course desisters? trajectories of crime among deliquent boys followed to age 70*', *Criminology* **41**(3), 555–592.

Sherman, L. W., Gottfredson, D. C., Mackenzie, D. L., Eck, J. E., Reuter, P. and Bushway, S. D. (1998), Preventing crime: What works, what doesn't, what's promising. research in brief. national institute of justice.

Stigler, G. J. (1970), 'The optimum enforcement laws', *Journal of Political Economy* **78**, 526–36.

Stoyan, D., Kendall, W. S. and Mecke, J. (1987), *Stochastic Geometry and Its Applications*, John Wiley & sons.

Truman, J. L. and Langton, L. (2014), Criminal victimization, 2013, Technical report, U.S. Department of Justice, Office of Justice Programs, Bureau of Justice Statistics.

U.S. Department of Justice, Office of Justice Programs, Bureau of Justice Statistics (2016), 'National corrections reporting program, 2000-2013: Selected variables'. ICPSR36285-v1. Ann Arbor, MIL Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/ICPSR36285.v1.

Appendix

A Apprehension Probability

In this Appendix, we discuss in detail the derivation of Equation (1), which represents our police production function. In order to simplify the calculations, we set the region R to be rectangular, with area $S(R)$, and ignore border effects.[19]

Criminals can commit crimes at each point in R and they seek to minimize the probability of apprehension. By assumption, the number of patrolling units is distributed as a Poisson with parameter $\Omega = \xi k^{\text{patrol}}$ and, conditional on the number of patrolling units, each unit is independently and identically distributed with a continuous distribution D over R, where D is chosen by the police command in order to maximize the number of apprehensions. The game is played simultaneously. According to Stoyan et al. (1987), the stochastic process of deployment of points is a Poisson point process with intensity function $\omega : R \to \mathbb{R}_+$, $x \mapsto D(x)\Omega$. We call this process Φ_D.

If a crime is committed at an arbitrary position $x \in R$, the probability $\pi(x; D)$ of detecting this crime is equal to the probability of having at least one police unit in the set $\mathbb{B}(x) = \{y \in \mathbb{R}^2 \text{ s.t. } \|x - y\| \leq r^p\}$. Thus we have that

$$\pi(x; D) = 1 - \mathbb{E}\Big(\prod_{y \in \Phi_D} (1 - q)\mathbb{1}_{y \in \mathbb{B}(x)}\Big),$$

and using the generatrix functional property of Φ_D we obtain[20]

$$\pi(x; D) = 1 - e^{-\int_{\mathbb{B}(x)} q\omega(y)\mathrm{d}y}.$$

Since each criminal wants to minimize the probability of being caught, all crimes are committed in the set $\operatorname{argmin}_{x \in R}\{\pi(x; D)\}$.

The police command's problem is then to choose D such that

$$\underset{D \in C^0}{\operatorname{argmax}}\Big\{\min_{x \in R}\{\pi(x; D)\} \,|\, \textstyle\int_R D(y)\mathrm{d}y = 1\Big\},$$

[19]There are several ways to eliminate border effects. We choose to identify the four corners of R to the same point. This means that the city is a flat torus. Roughly, an agent in the easternmost border of the city is close to the westernmost border. This might cause a distortion since part of the patrolling region of a patrol in the east of the city might fall on the west. This distortion is negligible as long as the relationship between the patrolling radius r^p and the smallest dimension of R is small.

[20]The generatrix functional property of a Poisson point process states that, for a measurable function $u : R \to \mathbb{R}$, $\mathbb{E}\Big(\prod_{y \in \Phi_D} u(x)\Big) = e^{-\int(1-u(x))\omega(\mathrm{d}y)}$.

which is equivalent to[21]

$$\underset{D \in C^0}{\operatorname{argmax}} \{ \min_{x \in R} \{ \int_{\mathbb{B}(x)} D(y)\mathrm{d}y \} \}$$
$$\text{s.t.} \int_R D(y)\mathrm{d}y = 1.$$

Because D is continuous and $\int_R D(y)\mathrm{d}y = 1$, any point $x \in R$ with $\int_{\mathbb{B}(x)} D(y)\mathrm{d}y > 1/S(R)$ implies the existence of a point $x' \in R$ such that $\int_{\mathbb{B}(x')} D(y)\mathrm{d}y > 1/S(R)$. Thus, the unique solution to this problem is $D^\star = 1/S(R)$ and we can write the apprehension probability as

$$1 - e^{-\xi k^{\text{patrol}} \int_{\mathbb{B}(x)} \frac{1}{S(R)} \mathrm{d}y} = 1 - e^{-\zeta_1 k^{\text{patrol}}},$$

where $\zeta_1 = q\xi \frac{\pi r^{p2}}{S(R)}$. Given the distribution $D^\star$, criminals are indifferent regarding where to commit crimes. Thus, the only Nash equilibrium occurs when crime is uniformly distributed over R and the social planner chooses the distribution $D^\star$. Any other strategy profile is ruled out as a Nash equilibrium by contradiction.

So far, we have expressed the apprehension probability as a function of k^{patrol}. We rewrite this relationship to express it as a function of the crime rate v and total police expenditures k. By assumption, the cost of each apprehension is ζ_0 and the total number of apprehensions given v and k is $vp(v, k)$, so total police expenditures are given by $k = k^{\text{patrol}} + \zeta_0 vp(v, k)$. This means that we can write the apprehension probability implicitly as a function of k and v as

$$p(v, k) = 1 - e^{-\zeta_1(k - \zeta_0 vp(v,k))} = 1 - e^{-\zeta_1 k + \zeta_2 vp(v,k)},$$

where $\zeta_2 = \zeta_0\zeta_1$. Solving this functional equation, we obtain for all $v, k \geq 0$

$$p(v, k) = \begin{cases} 1 - \frac{W_L(\zeta_2 v e^{-\zeta_1 k + \zeta_2 v})}{\zeta_2 v} & \text{if } v > 0 \\ 1 - e^{-\zeta_1 k} & \text{if } v = 0. \end{cases}$$

We show that p respects several interesting properties that we expect from an apprehension probability function. Since W_L is continuous and differentiable, p is continuous for $k \geq 0$ and $v > 0$ and differentiable in $k, v > 0$. The implicit function of p guarantees the continuity also at $v = 0$. Differentiating implicitly $p(v, k) = 1 - e^{-\zeta_1 k + \zeta_2 vp(v,k)}$ with respect to k and v shows that p is strictly increasing in k, strictly concave in k and strictly decreasing in v. Differentiating implicitly the number of apprehensions $a(v, k) = v - ve^{-\zeta_1 k + \zeta_2 a(v,k)}$ with respect to k and v shows that a is strictly increasing in k and v. Using the implicit function of p, we can show that: (i) $p(v, k) \to 1$ when $k \to \infty$; (ii) $p(v, k) \to 0$ when $k \to 0$; and (iii) $p(v, k) \to 0$ when $v \to \infty$.

[21] To see why, notice that

$$\max_{D \in C^0} \{ \min_{x \in R} \{ 1 - e^{-\Omega \int_{\mathbb{B}(x)} D(y)\mathrm{d}y} \} | \int_R D(y)\mathrm{d}y = 1 \} = \max_{D \in C^0} \{ 1 - e^{-\Omega \min_{x \in R} \{ \int_{\mathbb{B}(x)} D(y)\mathrm{d}y \}} | \int_R D(y)\mathrm{d}y = 1 \}$$
$$= 1 - e^{-\Omega \max_{D \in C^0} \{ \min_{x \in R} \{ \int_{\mathbb{B}(x)} D(y)\mathrm{d}y \} | \int_R D(y)\mathrm{d}y = 1 \}}.$$

B Equilibrium Existence and Computation

We show that our economy yields an equilibrium and outline the procedure used to compute it. Corollary B.1 refers to condition *(iii)* of the equilibrium definition presented in the text. Corollary B.2 incorporates conditions *(i)* and *(ii)*. Conditions *(i)*, *(ii)*, and *(iii)* together allow us to calculate all stationary state variables of this economy for any given belief regarding the crime rate v. The restriction imposed by condition *(iv)* pins down the equilibrium crime rate $v^\star$. The existence of this equilibrium crime rate is shown in Proposition B.3. In the particular case of linear utilities and non-distortionary taxation, condition *(v)* does not have any impact on the individual or aggregate behavior and is simply an accounting identity, respected by setting f accordingly—condition *(v)* only matters to compute the welfare loss.

For a given belief v, we define

$$\overline{\Pi}_t(v) := \frac{(z(1-p(v,k)) - d_t^C)(1-e^{-p(v,k)\nu})}{p(v,k)}$$

and $\overline{p}(v) := (1-e^{-p(v,k)\nu})(1-\mu)$. We also denote $V_t^F(w;v)$ and $V_t^P(w;v)$ as the value functions of, respectively, free and incarcerated individuals with productivity w, at age t, as a function of v.

Corollary B.1. *There exists a sequence of cut-off productivities $w_t^\star(v)$ satisfying condition (iv) of the equilibrium, such that an agent at age t engages in criminal activities if, and only if, his productivity is below $w_t^\star(v)$. Moreover, there are recursive expressions for $V_t^F(w;v)$, $V_t^P(w;v)$ and $w_t^\star(v)$.*

Proof. First, we note that $V_t^F(\cdot;v)$ and $V_t^P(\cdot;v)$ are strictly increasing. We find each $w_t^\star(v)$ by backward induction. At age T we have

$$V_T^F(w;v) = w - f + \beta\alpha(\gamma_T w - f) + \max\{\overline{\Pi}_T(v), 0\}.$$

Therefore, $w_T^\star(v) = -\infty$ if $\overline{\Pi}_T(v) > 0$ and $w_T^\star(v) = +\infty$ otherwise. In addition, $V_T^P(w;v) = b - f + \beta\alpha(\theta w - f)$. As a second step, we use the (backward) induction principle to prove that $V_t^F(\cdot;v)$ and $V_t^P(\cdot;v)$ respect the following set P_V of properties: they are continuous, convex, non-differentiable at a finite number of points, and linear by parts. In addition, we show they respect the property P_D that $V_t^F(\cdot;v) - V_t^P(\cdot;v)$ is strictly increasing.

Since $V_T^F(\cdot;v)$ and $V_T^P(\cdot;v)$ are linear with $V_T^{F'}(\cdot;v) = 1+\beta\alpha\gamma_T$ and $V_T^{P'}(\cdot;v) = \beta\alpha\theta$, they respect P_V and P_D. Suppose $V_{t+1}^F(\cdot;v)$ and $V_{t+1}^P(\cdot;v)$ also respect P_V and P_D for $t \le T-1$. It is trivial to see that P_V also holds for $V_t^P(\cdot;v)$. To show the same for $V_t^F(\cdot;v)$, it suffices to observe that if H respects P_V while being increasing and non-differentiable at N_{ND} points, then $\max\{0, H(\cdot)\}$ also respects P_V, is non-differentiable at, at most, $N_{ND}+1$ points, and is linear by parts.

It remains to prove that P_D holds. Let $V_t^{F'}(\cdot;v)$ and $V_t^{P'}(\cdot;v)$ be the derivatives of, respectively, $V_t^F(\cdot;v)$ and $V_t^P(\cdot;v)$ at all differentiable points. The convexity of $V_t^P(\cdot;v)$

implies that $\mathrm{d}V_t^P(w;v)/\mathrm{d}w \leq \mathrm{d}V_t^P(\gamma_t w/\theta;v)/\mathrm{d}w$ for each w (recall that $\gamma_t/\theta > 1$), so

$$\frac{\mathrm{d}(V_t^F(w;v) - V_t^P(w;v))}{\mathrm{d}w} \geq \frac{\mathrm{d}(V_t^F(w;v) - V_t^P(\gamma_t w/\theta;v))}{\mathrm{d}w}.$$

If we prove that the right-hand side term is positive, we prove P_D at age t. Using equations 2 and 3, and since $\mathrm{d}(V_{t+1}^F(\gamma_t w;v) - V_{t+1}^P(\gamma_t w;v))/\mathrm{d}w > 0$ (due to P_D at age $t+1$), we obtain

$$\frac{\mathrm{d}(V_t^F(w;v) - V_t^P(\gamma_t w/\theta;v))}{\mathrm{d}w} \geq 1 + \beta(1 - \overline{p}(v) - \mu)\frac{\mathrm{d}(V_{t+1}^F(\gamma_t w;v) - V_{t+1}^P(\gamma_t w;v))}{\mathrm{d}w}.$$

Recalling that $\overline{p}(v) = (1-\mu)(1 - e^{-\nu p(v,k)})$, we have

$$1 - \overline{p}(v) - \mu = (1-\mu)e^{-\nu p(v,k)} > 0,$$

which proves that $V_t^F(\cdot;v) - V_t^P(\cdot;v)$ is strictly increasing wherever it is differentiable. Since $V_t^F(\cdot;v)$ and $V_t^P(\cdot;v)$ are continuous and non-differentiable only at a finite number of points, it follows straightforward that P_D holds at age t, concluding the induction.

We just proved that $V_t^F(\cdot;v) - V_t^P(\cdot;v)$ is continuous, strictly decreasing and (since it is linear by parts) ranges from $-\infty$ to $+\infty$. From Equation 2, since $\overline{\Pi}_t(v)$ is constant, there exists $w_t^\star(v)$ such that $V_{t+1}^F(w_t^\star(v);v) - V_{t+1}^P(w_t^\star(v);v) = \overline{\Pi}_t(v)$.

Once we have $w_t^\star(v)$, the expressions for $V_t^P(w;v)$ and $V_t^F(w;v)$ for each w are

$$V_t^F(w;v) = \begin{cases} w - f + \beta V_{t+1}^F(\gamma_t w;v), & \text{if } w \geq w_t^\star(v), \\ w - f + \overline{\Pi}_t + (1-\overline{p}(v))\beta V_{t+1}^F(\gamma_t w;v) + \overline{p}(v) V_{t+1}^P(\gamma_t w;v), & \text{if } w < w_t^\star(v), \end{cases}$$

and

$$V_t^P(w;v) = b - f + (1-\mu)\beta V_{t+1}^P(\theta w;v) + \mu\beta V_{t+1}^F(\theta w;v).$$

This completes the proof. $\square$

Since we know the behavior of the agents as a function of v, how productivity evolve as a function of their state, and the initial distribution of productivity, we can find the measure of productivity at all ages. Suppose that we know the measure of agents that are free and the measure of agents that are incarcerated at age $t-1$. The general principle to obtain the measure of agents with productivity w in cohort t is to find the measure of agents at age $t-1$ that ends up with productivity w at age t and add these measures (adjusting properly the measures by the Jacobian of the productivity changes). Since we have the measure of productivity for cohort 0, we can use induction to find the measure of productivity for each cohort.

To formalize this idea, we introduce the following variables and functions. For a given crime rate and for each $s \in \{\tilde{E}, E, P\}$, let $m_t^s(w;v)$ be the measure of agents at age t with productivity w in state s.[22] Then, denote $M_t^s(w;v) := \int_0^w m_t^s(w';v)\mathrm{d}w'$, where $M_t^s(\infty;v) := N_t^s(v)$ (so $N_t^E(v)$ is the mass of agents at age t engaging in criminal activities and $N_t^P(v)$ is

[22]Notice that, if an equilibrium exists with crime rate $v^\star$, $m_t^s(w;v^\star) = m_t^s(w)$.

the mass of prisoners at age t). Let also $v_t(v)$ be the number of crimes committed by agents at age t. Recall that agents at age 0 are free, so $M_0^{\tilde{E}}(\cdot;v)+M_0^E(\cdot;v)=M_0G_0$, $M_0^P(\cdot;v)=0$, and we can obtain $w_0^\star(v)$ by Corollary B.1.

Corollary B.2. *For a given crime rate belief v and $2<t<T+1$, we can obtain recursive expressions for $m_t^F(\cdot;v)$, $m_t^P(\cdot;v)$ and $v_t(v)$ as functions of $m_{t-1}^F(\cdot;v)$ and $m_{t-1}^P(\cdot;v)$, satisfying conditions (i) and (iii) of the equilibrium.*

Proof. We use induction. For age 0, recall that agents are free when they enter the economy, so $N_0^P(v)=0$, $N_0^E(v)=M_0G_0(w_0^\star(v))$ and

$$v_0(v)=\frac{(1-e^{-\nu p(v,k)})}{p(v,k)}N_0^E(v)=\frac{\overline{p}(v)}{(1-\mu)p(v,k)}N_0^E(v).$$

Besides, we have, for all w, $m_0^P(w;v)=0$, $m_0^E(w;v)=M_0g_0(w)\mathbb{1}_{[w\leq w_0^\star(v)]}$, and $m_0^{\tilde{E}}(w;v)=M_0g_0(w)\mathbb{1}_{[w>w_0^\star(v)]}$, and the mass of new incarcerated agents at age 1 is $(1-e^{-\nu p(v,k)})(1-\mu)N_0^E(v)$.

Consider now the measures for age t and assume we know them for $t-1$. There are two possibilities allowing an individual to be in prison at age t with productivity w. The first one is that, at age $t-1$, he/she was free with productivity w/γ_{t-1}, engaged in crime, and was then apprehended and incarcerated. The second one is that, at age $t-1$, he/she was in prison with productivity w/θ and was not released between ages $t-1$ and t. Since a proportion $(1-\mu)$ of agents at age $t-1$ remains in prison at age t, and a proportion $\overline{p}(v)$ of agents engaged in crime at age $t-1$ become prisoners at age t, we have that

$$m_t^P(w;v)=\left(\overline{p}(v)\frac{m_{t-1}^E(w/\gamma_{t-1};v)}{\gamma_{t-1}}+(1-\mu)\frac{m_{t-1}^P(w/\theta;v)}{\theta}\right).$$

We calculate $m_t^E(w;v)+m_t^{\tilde{E}}(w;v)$ in a similar way. The following groups of agents at age $t-1$ start period t with productivity w: all agents not engaged in crime with productivity w/γ_{t-1}; a proportion $(1-\overline{p}(v))$ of agents engaged in crime with productivity w/γ_{t-1}; and a proportion μ of agents in prison with productivity w/θ. So

$$m_t^E(w;v)+m_t^{\tilde{E}}(w;v)=\Big(\frac{(1-\overline{p}(v))m_{t-1}^E(w/\gamma_{t-1};v)+m_{t-1}^{\tilde{E}}(w/\gamma_{t-1};v)}{\gamma_{t-1}}+\mu\frac{m_{t-1}^P(w/\theta;v)}{\theta}\Big).$$

Then, according to Corollary B.1,

$$m_t^E(w;v)=(m_t^E(w;v)+m_t^{\tilde{E}}(w;v))\mathbb{1}_{[w<w_t^\star(v)]}\text{ and }m_t^{\tilde{E}}(w;v)=(m_t^E(w;v)+m_t^{\tilde{E}}(w;v))\mathbb{1}_{[w\geq w_t^\star(v)]}.$$

In addition, notice that, by assumption, $m_{T+1}^E(w;v)=m_{T+1}^P(w;v)=0$, and

$$m_{T+1}^{\tilde{E}}(w;v)=\alpha\left(\frac{m_T^E(w/\gamma_T;v)+m_T^{\tilde{E}}(w/\gamma_T;v)}{\gamma_T}+\frac{m_T^P(w/\theta;v)}{\theta}\right).$$

From these differential relations, we obtain the integral ones. The total number of prisoners at age t is $N_t^P(v) = \int_0^\infty m_t^P(w; v)\mathrm{d}w$ and the total number of agents engaged in crime at this age is $N_t^E(v) = \int_0^\infty m_t^E(w; v)\mathrm{d}w$. The total number of crimes committed by agents at age t is then given by

$$v_t(v) = \frac{\overline{p}(v)}{(1-\mu)p(v,k)} N_t^E(v)$$

and the total number of apprehensions is $\overline{p}(v)N_t^E(v)$. The total number of crimes is $\sum_{t=0}^T v_t(v)$. $\square$

Using the previous results, regularity conditions and the Intermediate Value Theorem, we can show the existence of a stationary equilibrium. We highlight this result in the following theorem.

Proposition B.3. *There exists a rational expectations stationary equilibrium in the economy.*

Proof. We have $0 \leq \sum_{t=0}^T v_t(v) \leq \overline{M}$, where $\overline{M}$ is the crime rate when the apprehension probability is 0, i.e., all dishonest agents engage in crime. If each v_t is continuous, then $\sum_{t=0}^T v_t(\cdot)$ is continuous and, thus, has a compact image (in $[0, \overline{M}]$), so we can use the Intermediate Value Theorem to show that there is at least one fixed point.

It remains to show that $v_t(\cdot)$ is continuous for each t. First, by the Implicit Function Theorem, the function $w_t^\star(\cdot)$ is continuous and differentiable. Using induction on the recursive relation of $m_t^s(w; v)$, we can see that $m_t^s(w; \cdot)$ is differentiable for each s and t. This, in turn, implies that $N_t^s(\cdot)$ is continuous for each s and t. The continuity of $N_t^E(\cdot)$, $\overline{p}(\cdot)$ and $p(\cdot, k)$ leads to a continuous $v_t(\cdot)$ for each t. $\square$

C Expression for L^{work}

This Appendix explains how we arrive at the expression for L^{work}. This loss is defined as the difference between a counterfactual aggregate productivity of free dishonest agents and their actual aggregate productivity. Agents in the counterfactual scenario use the decision rule of agents in the actual equilibrium, but do not suffer productivity depreciation due to incarceration. We can calculate the counterfactual distributions until age T, defined as $\dot{m}_t^P(w)$, $\dot{m}_t^E(w)$, and $\dot{m}_t^{\tilde{E}}(w)$, in the same way we have obtained $m_t^P(w)$, $m_t^E(w)$, and $m_t^{\tilde{E}}(w)$ in Appendix B. So, for all $w \geq 0$, $\dot{m}_0^P(w) = 0$, $\dot{m}_0^E(w) = M_0 g_0 \mathbb{1}_{[w<w_0^*]}$, $\dot{m}_0^{\tilde{E}}(w) = M_0 g_0 \mathbb{1}_{[w \geq w_0^*]}$,

$$\dot{m}_t^P(w) = \frac{1}{\gamma_{t-1}} \left(\overline{p} \dot{m}_{t-1}^E(w/\gamma_{t-1}) + (1-\mu) \dot{m}_{t-1}^P(w/\gamma_{t-1}) \right),$$

and

$$\dot{m}_t^E(w) + \dot{m}_t^{\tilde{E}}(w) = \frac{1}{\gamma_{t-1}} \Big((1-\overline{p}) \dot{m}_{t-1}^E(w/\gamma_{t-1}) + \dot{m}_{t-1}^{\tilde{E}}(w/\gamma_{t-1}) + \mu \dot{m}_{t-1}^P(w/\gamma_{t-1}) \Big).$$

Besides, $\dot{m}_t^E(w) = (\dot{m}_t^E(w) + \dot{m}_t^{\tilde{E}}(w)) \mathbb{1}_{[w \leq w_t^*]}$ and $\dot{m}_t^{\tilde{E}}(w) = (\dot{m}_t^E(w) + \dot{m}_t^{\tilde{E}}(w)) \mathbb{1}_{[w > w_t^*]}$ for $t \leq T$. We also have $\dot{m}_{T+1}^E(w; v) = \dot{m}_{T+1}^P(w; v) = 0$, and

$$\dot{m}_{T+1}^{\tilde{E}}(w) = \frac{\alpha}{\gamma_T} \Big(\dot{m}_T^E(w/\gamma_T) + \dot{m}_T^{\tilde{E}}(w/\gamma_T) + \mu \dot{m}_T^P(w/\gamma_T) \Big).$$

Thus

$$L^{\text{work}} = \sum_{t=0}^{T+1} \left(\int_0^\infty w (\dot{m}_t^{\tilde{E}}(w) + \dot{m}_t^E(w) - (m_t^{\tilde{E}}(w) + m_t^E(w))) \mathrm{d}w \right).$$

D Estimation

This Appendix explains the technical details related to various estimation procedures mentioned in the main text.

D.1 Pre-GMM Stage

The pre-GMM stage maps each parameter with a statistics from a dataset or some estimate obtained from the literature.

Population of each cohort

We use the 2000 US Census, the decennial census that is closest to 2004. According to this census, the 2000 US population between ages 18 and 65 corresponded to 176.1 million people. Since we assume a constant cohort size, the population in each cohort is 3.67 million people.

Productivity gain

The CPS sample used to estimate the age-specific productivity gain for each year in freedom is taken from the 2004 CPS March Supplement. We run a regression of the logarithm of the reported weekly income on years of education, age minus 18, the square of age minus 18, race, and gender. We exclude individuals with missing information on any of these four characteristics, individuals under 18 years-old and over 65 years-old. We assume that prison population is small compared to the total population, so that the depreciation for having served prison does not significantly impact the aggregate estimates obtained from this regression. In addition, we assume that the productivity gain of dishonest agents when they spend one year in freedom is, on average, the same as the one of an honest agent at the same age.

We find a coefficient for age minus 18 $\beta_{\text{age}} := 0.0687$ and a coefficient for squared age minus 18 $\beta_{\text{age}^2} := -0.00121$. The average productivity gain γ_t is such that

$$\widehat{\ln \gamma_t} = \widehat{\ln(\tfrac{w_{t+1}}{w_t})} = \beta_{\text{age}}(t+1) + \beta_{\text{age}^2}(t+1)^2 - (\beta_{\text{age}}t + \beta_{\text{age}^2}t^2) = \beta_{\text{age}} + 2\beta_{\text{age}^2}t + \beta_{\text{age}^2},$$

Therefore

$$\widehat{\gamma_t} = \exp(\beta_{\text{age}} + 2\beta_{\text{age}^2}t + \beta_{\text{age}^2}).$$

The initial productivity distribution

To calculate the distribution of productivity of 18 year-old agents, we use the monthly CPS from January 2003 to December 2005. We pool the datasets using the combination of individual-time as the unit of observation, in order to increase sample size. Then, we consider the log of yearly income in hundreds of dollars (52/100 times the reported weekly earnings), restricting the sample to agents aged 18 that report positive weekly earnings. We assume that the productivity distribution is log-normal with parameters μ_w (the expected value of the log of income) and σ_w (the standard deviation of the log of income). Then, we estimate the two log-normal parameters by maximum likelihood, considering the weight of each observation. We find $\widehat{\mu}_w = 4.453$ and $\widehat{\sigma}_w = 0.7169$.

Multiplier α defining the continuation value at age $T+1$

We assume that the income of each agent remains constant after age T and that an agent with income w at age $T+1$ has a continuation value of αw. To estimate the multiplier α, we use the US 2004 life tables from the National Vital Statistics System. Note that

$$\alpha = \sum_{t=0}^{\infty} \beta^t \Pr(\text{dying age} > t+T+1 | \text{alive at } T+1).$$

The life tables give us $P_x^\ell := \Pr(\text{dying at age } x+1 | \text{alive at age } x)$. After a straightforward calculation, it follows that

$$\Pr(\text{dying age} > t+T+1 | \text{alive at } T+1) = \prod_{\tau=T+1}^{T+t} (1 - P_\tau^\ell).$$

Using this probability and the discount factor, we obtain our estimate for the multiplier $\widehat{\alpha} = 11.38$.

Total number and distribution of prisoners by type of crime

Prisoners in the US can be in State prisons, Federal prisons, or jails. In particular, prisoners in jails may be waiting for trial or already convicted (in this case, usually serving a short sentence). We start with the count of prisoners in jails in 2004. The BJS reports, in "Prison and Jail inmates at Midyear 2005," that there were 713,900 prisoners in jails in 2004, from which 39.7% (283,418) were convicted whereas the remainder (430,481) were awaiting trial.

We complement this information with the BJS Special Report "Profile of Jail inmates, 2002." This report presents the proportion of prisoners in jails by type of offense (conditional on the status of being convicted or not convicted). We assume that these proportions, presented in Table A1, were the same in 2004.

Using the number of convicted and non-convicted prisoners in jails in 2004 and the proportions from Table A1, we obtain the number of prisoners in jails in 2004 by type of offense and conviction status. The number of prisoners in State and Federal prisons is taken from the BJS Bulletin "Prisoners in 2006." Table A2 presents these numbers.

Assuming that the number of prisoners for larcenies of values lower than 50 dollars is negligible and that p_{conv} is the proportion of non-convicted prisoners that ended up convicted, the estimated number of convicted prisoners in 2004 is $453900 + 96500 p_{\text{conv}}$. According to Dobbie et al. (2018), 57.8% of detained defendants are found guilty, so we use this value as our estimate for p_{conv}. This implies that our estimated number of prisoners that were or would eventually be convicted in 2004 is $509,700$.

	Convicted (%)	Unconvicted (%)
Burglary	6.4	6.8
Larceny	7.6	5.3
Motor Theft	2.0	1.6
Robbery	3.9	8.7

Table A1: The first column shows the proportion of prisoners in jail convicted, by type of crime, with respect to the total number of convicted prisoners in jails. The second column is the analogous for prisoners awaiting trial in jails.

Most serious offense	Type of incarceration			
	Jail convicted	Jail unconvicted	State prison	Federal prison
Burglary	18100	29300	135700	500
Larceny	21500	22800	50400	0
Motor Theft	5700	6900	22300	0
Robbery	11100	37500	178900	9700
All	56400	96500	387300	10200

Table A2: Estimated number of prisoners by type of incarceration and most serious offense.

Average losses from property crimes

First, we calculate the relative frequency of each type of crime in 2004. To do so, we use the BJS report "Criminal Victimization, 2004." This report presents the number of victimizations by type of crime. Adding up the victimizations in our definition of property crimes (15.032 million property crimes), we obtain the first column of Table A3.

Then, from Cohen (2000), we take the estimated losses from property crimes in 1993 dollars. Column 2 of Table A3 shows the average property loss as a function of the type of property crime, while column 3 lists the value of other losses associated with victimization (including medical care, quality of life, productivity, police/fire services, and social services). All values in the table are converted into 2004 dollars. The average loss for a property crime is the sum of the losses for each type of offense weighted by the frequency of the offense.

Type of crime	Frequency	Losses		
	(%)	Property	Other	Total
Robbery	3.6	975	9425	10400
Burglary	22.7	1261	559	1820
Motor Theft	6.8	4290	650	4940
Theft	66.9	351	130	481
Average		848	597	1445

Table A3: The first column presents the estimated proportion of each property crime with respect to the total estimated number of property crimes. The second column shows the average property loss as a function of the type of crime. The third column depicts the average losses not related to the property. We average the losses, weighting by the frequency of the offense.

Expenditures with public safety

The 2004 Criminal Justice Expenditure and Employment Extract Program (CJEE) provides data on expenditures with police protection, \$88.9 billion, and correction, \$62.0 billion. Since, according to the BJS, the total number of prisoners in this same year was 2,385 million, our estimate for κ is $26,000$ dollars per prisoner.

We also estimate the fraction of the expenditures on police protection devoted to property crimes, which we call h. To do that, we use data from the Bureau of Justice Statistics (2016) containing information at the prisoner level for many states. We estimate, from 2004 to 2008, that: (i) the proportion of prisoners sentenced for property crimes was 27.10%; and (ii) the proportion of admissions of individuals sentenced for property crimes was 26.47%. For this reason, we set $\hat{h} = 0.27$.

Distribution of sentences

We start with State prisons and then incorporate jail inmates. We do not have data on release from Federal prisons, but it does not have a significant impact on our estimate since there are few inmates incarcerated in Federal prisons for property crimes (see Table A2). We use data from the Bureau of Justice Statistics (2016) that reports the time each prisoner served upon his release from State prisons (we only have the time interval in which the sentence was served). We restrict our sample to released prisoners that have some type of property crime as the most serious offense. In addition, we pool all prisoners from 2004 to 2008. Although the study does not include all State prisons, we assume that the variation in time served presented in this study, including 35^+ States, can be extrapolated to all State and Federal prisons. The distribution we estimate for prisoners in State and Federal prisons is shown in Table A4.

Time served	Proportion (%)
$<$ 1 year	56.24
1 - 1.9 years	18.77
2 - 4.9 years	16.55
5 - 9.9 years	5.77
$\geq$ 10 years	2.67

Table A4: Estimated distribution of the time served by inmates in State or Federal prisons.

To estimate the distribution of time served for all inmates, we make five assumptions. First, that time served is exponentially distributed. Second, that prisoners that are convicted and incarcerated in jails serve at most one year.[23] Third, we maintain the assumption that 58.7% of prisoners awaiting trial were eventually convicted. Forth, once prisoner awaiting trial are convicted, the probability of being sent to a prison equals the proportion of prisoners held in prisons over the total number of inmates (in State and Federal prisons or in jails). The fifth assumption is that, once a convicted prisoner is transferred from a jail to a State or Federal prison, the served time is distributed as in Table A4. Under these assumptions, we arrive Table A5.

Given that we parameterized the distribution of time served as an exponential and that we have the distribution for the five time served brackets, we estimate the parameter of the exponential through a standard maximum likelihood procedure. We find an estimate for the mean time served $\widehat{\lambda}_{2004} = 1.482$ years (17.78 months), corresponding to a probability of being released from prison at each year of $\widehat{\mu}_{2004} = 1/(1 + \widehat{\lambda}_{2004}) = 0.403$.

[23]See the box named "Jail populations" on page 7 of the bulletin "Prison and Jail Inmates at Midyear 2005." According to it, "Inmates sentenced to jail usually have a sentence of 1 year or less."

Time served	Proportion (%)
< 1 year	63.60
1 - 1.9 years	15.61
2 - 4.9 years	13.76
5 - 9.9 years	4.80
≥ 10 years	2.22

Table A5: Estimated distribution of the time served by all inmates.

Productivity depreciation due to incarceration

We take our estimate of the productivity depreciation due to incarceration from Grogger (1995). The average quarterly income in his sample is given by \$1,182 (in 1980s dollars) for individuals aged between 18 and 25. Grogger's paper describes the effect of incarceration in jails and in prisons on the legal incomes with 7 time lags (in the year of the release, one year after the release, and so on, until six years after the release). The simple average (across all years after the release) of the loss of legal income for having been in jail is \$161. The equivalent number for prison is \$303. The ratio of incarcerated individuals in prisons (State or Federal) and in jails is 1.973. Thus, the average productivity depreciation is \$255, or 21.6% of the sample mean. Since the average time in prison is 1.48 years, this corresponds to a yearly depreciation rate of 15%. The average gain for free agents between ages 18 and 25 is 6% (this is an approximation, we averaged $\hat{\gamma}_0$, $\hat{\gamma}_1$, up to $\hat{\gamma}_7$), so the average loss exclusively due to incarceration is 9%, which implies that our estimate for θ is $\widehat{\theta} = 0.91$.

Proportion of dishonest individuals

We estimate the proportion of dishonest individuals from Sampson and Laub (2003). They estimate the proportion of individuals committing property crimes as a function of age. They use self-reported, parental-reported, and teacher-reported data on whether boys committed crimes (as described in Laub and Sampson 2003). First, we make two assumptions: (i) boys committing property crimes face very low costs associated with punishment; and (ii) the age at which the proportion of boys committing crimes is the largest is an age at which the subjective cost of committing crimes is 0. These assumptions imply that, at the peak (in age) of the proportion of individuals committing crimes, all dishonest agents are engaged in crime. Sampson and Laub (2003) show two peaks, corresponding to 0.31 and 0.21, one for each level of childhood risk. We take the average in the population of boys as the simple average, 0.26, to obtain the proportion of dishonest males. Our last assumption is that the ratio between the proportion of dishonest women and men is the same as the ratio between women and men admitted into prison due to property crimes. This ratio is given by 0.1254. Therefore, the proportion of dishonest women is 0.0326 and the proportion of dishonest individuals in general is approximately $(0.26 + 0.0326)/2 \approx 15\%$.

D.2 GMM

For computational reasons, we perform our GMM estimation in two stages. We start by describing the derivative procedures used in the code.

D.2.1 Derivatives

In our code, we calculate elasticities and derivatives. Since the discretization generates functions of variables and parameters that are sometimes noisy, we perform linear regressions to obtain the derivatives. So, when we evaluate the derivative of Z with respect to X, we estimate via OLS the coefficient β_1 of the equation $Z = \beta_0 + \beta_1 X + \text{error}$ and the derivative is given by the estimate $\widehat{\beta_1}$ of β_1. To evaluate a Hessian, the procedure is similar. If Z is a function of X (the first argument) and Y, we estimate by OLS the coefficients of the equation $Z = \beta_0 + \beta_1 X + \beta_2 Y + \beta_3 X^2 + \beta_4 Y^2 + \beta_5 XY + \text{error}$. The Hessian is given by $\left(\begin{smallmatrix} \widehat{\beta}_3 & \widehat{\beta}_5 \\ \widehat{\beta}_5 & \widehat{\beta}_4 \end{smallmatrix}\right)$.

D.2.2 First Stage

In the first stage, we estimate the parameters defining the cost of committing crimes b^C and c^C, the average number of crime opportunities per year ν, the wage-equivalent utility of incarcerated individuals b, and the probability of being caught after committing a crime p. To estimate these parameters, we use 19 moments: 15 moments from the age distribution of prison inflows (proportion of admissions by age group), the total number of crimes, the value of 90 days of freedom, the proportion of 18-19 year-old individuals committing crimes, and the number of prisoners. The procedures used to estimate the mean and the variance of these moments is described as follows.

Number of property crimes

Mean The crime rate is taken from the 2004 NCVS. We consider all property crimes excluding thefts of less than \$50 (14.488 million crimes), plus robberies (0.548 million crimes), totaling 15.032 million property crimes.

Variance The standard errors of our estimate of the number of crimes come from the sampling errors discussed in the 2004 NCVS report. The standard deviations are 71,649 for robbery, 162,722 for burglary, 75,127 for auto-theft and 387,850 for theft. We assume they are independent, so the variance is given by the sum of the squares of the standard errors, which corresponds to 187683.71 (thousands of crimes)^2.

Distribution of admissions by age group

Mean We obtain, from the Sourcebook of Criminal Justice Statistics, the number of arrests as a function of the type of crime and age of the offenders, from 2002 to 2011. For any specific year, we have information for each age from 18 up to 24 and, after this age, we have information for five-year brackets (25-29, 30-34, etc.). So, we define age groups corresponding to one-year intervals until 24 and five-year intervals after 24. For each year from 2002 to 2011, we calculate the total number of arrests for property crimes (robbery, larceny, burglary, and auto theft) of individuals at each age group. Next, we calculate the proportion of individuals that were arrested at each age with respect to all arrests for property crimes for individuals aged between 18 and 65. We assume that the proportion of arrests by age group corresponds to the proportion of admissions in prisons/jails by age group.

The moment we use in our GMM to match the proportion of admissions of individuals by age group is the corresponding arrest shares by age group in 2004.

Variance We assume that the time fluctuations of this proportion capture most of the variation in this variable, so we calculate the variance of the time series from 2002 to 2011.

Value of freedom

Mean We take our estimate for the value of being free from Abrams and Rohlfs (2007). They estimate the value of 90 days of freedom in their sample as \$1,050 (in 2003 dollars). Since we estimate that a convicted criminal spends 1.48 years incarcerated, we approximate the value of not being in prison once as $1.48 \times 4 \times \$1,050 = \$6,222$.

Variance The variance of this moment also comes from Abrams and Rohlfs (2007). As with the mean, we multiply the standard error 6.31 hundred dollars by 4×1.48 to estimate the standard error of the willingness to pay to avoid being incarcerated. The variance is therefore given by 1395.4 (hundred dollars)^2.

Proportion of 18-19 year-old individuals committing crimes

Mean We use the NLSY97 to estimate the proportion of 18-19 year-old individuals commiting crimes. Individuals in this sample are interviewed yearly. In this survey, individuals self-report whether they committed at least one crime in the interview year.

We exclude instances of individual-year that were missing. Among interviewed individuals, there are also missing values on questions related to whether they committed crimes. We input 0 (resp. 1) for each missing answer to obtain a lower (resp. upper) bound for the number of crimes committed. We consider only individuals aged 18 or 19 at the time of the interview. The weighted averages of the lower and upper bound probabilities of engaging in crime are 2.645% and 3.353%, respectively. Our estimate for the proportion of individuals committing crimes is the middle point of the interval $[2.645, 3.353]$, 2.999%.

Variance We assume this proportion is uniformly distributed between the lower and upper bounds described above. Therefore, the variance is given by $(0.03352707^2 - 0.02645321^2)/12 = 3.5358 \times 10^{-5}$.[24]

Number of prisoners

Mean We calculate the number of prisoners for property crimes in Section D.1 to be 509,700.

Variance In Section D.1, we estimate that the probability that a detained defendant that has not been convicted is found guilty is 57.8%. We assume that this probability is independent for each inmate awaiting trial, so the distribution of the number of prisoners is binomial with a variance of $96,500 \times 0.578 \times (1 - 0.578) = 23538$ prisoners.

Results of the first stage of the GMM

The weight matrix of the GMM is defined as $(\Xi_1)^{-1}$, where Ξ_1 is the diagonal matrix with the variances calculated in Section D.2.2, which we assume to be the actual covariance matrix of the moments (since we do not have data to calculate the covariance between our moments, we assume they are null). We use the particle swarm optimization with 10,000 points in the space $[-1, -0.01] \times [0, 10] \times [0.01, 0.30] \times [5, 20] \times [20, 120]$ (in order, b^C, c^C, p, ν, b) to find the minimum score of the GMM.

To calculate the standard errors of the parameters, we use the covariance matrix. Let $\widehat{\varphi}$ be the vector of estimated parameters $(\widehat{b}^C, \widehat{c}^C, \widehat{p}, \widehat{\nu}, \widehat{b})$ and $m_i(\widehat{\varphi})$ the prediction of our model for the i-th moment. To estimate the covariance matrix of the parameters estimated in the first stage, we have to calculate the Jacobian of $\boldsymbol{m} = (m_1, m_2, \ldots, m_{19})$ evaluated at $\widehat{\varphi}$, defined as $J_\varphi(\boldsymbol{m}(\widehat{\varphi}))$. To evaluate $\partial_j m_i(\widehat{\varphi})$, we use the procedure explained in Section D.2.1. For each parameter φ_j, we set $X = (0.98\widehat{\varphi}_j, 0.9801\widehat{\varphi}_j, 0.9801\widehat{\varphi}_j, \ldots, 1.02\widehat{\varphi}_j)$ and $Y = m_i(X)$. The estimated covariance matrix of our parameters is given by $\widehat{\boldsymbol{V}}_1 = (J_\varphi(\boldsymbol{m}(\widehat{\varphi}))'(\Xi_1)^{-1} J_\varphi(\mathbf{m}(\widehat{\varphi})))^{-1}$.

D.2.3 Second Stage of the GMM

In the second stage of the GMM, we estimate the parameters of the detection technology: ζ_1 and ζ_2. The moments of our GMM are the estimated detection probability $\widehat{p}$ (point estimate of 2.30% and standard error of 0.06%). We take the elasticity of property crime in relation to police size from Levitt (2002) (point estimate of -0.501 and standard error of 0.235). The elasticity of property crime in relation to prison population is taken from Levitt (1996) (point estimate -0.261, standard error of 0.117).

[24]The uniform distribution assumption is conservative since it is the distribution with the greatest entropy when the support of a continuous random variable is a bounded interval.

Evaluation of the elasticity of property crimes in relation to police size

We estimate the expenditures on police for property crimes as $\widehat{k} = 0.27 \times 88.9\$\text{bn} = 24\$\text{bn}$. Let $X := (.99\widehat{k}, .991\widehat{k}, \ldots, 1.01\widehat{k})$. For a component of X, say X_i, the model can evaluate the crime rate, $\text{crime}(X_i, \zeta_1, \zeta_2)$. Given the vector $\text{crime}(X, \zeta_1, \zeta_2)$, we calculate the OLS estimate for $\beta_1(\zeta_1, \zeta_2)$ in the regression

$$\text{crime}(X, \zeta_1, \zeta_2) = \beta_0(\zeta_1, \zeta_2) + \beta_1(\zeta_1, \zeta_2)X + \text{error}$$

and associate $\beta_1(\zeta_1, \zeta_2)$ to the derivative of the crime rate on police size. The elasticity is simply $\widehat{\beta}_1(\zeta_1, \zeta_2)\widehat{k}/\widehat{v}$, with $\widehat{v}$ being the number of property crimes predicted by the model using the parameters estimated in the first stage of the GMM.

Estimation of the elasticity of property crimes in relation to prison population

We calculate an approximation for the variation of the number of crimes after a variation in the number of prisoners. First, we assume that, if some agents in prison were set free, most of them would commit crimes. Therefore, the sum of agents engaged in crime and incarcerated is, at the margin, approximately constant. This means that the derivative of property crimes on prison population is (approximately) the opposite of the derivative of property crimes in relation to the number of agents engaged in crime. Second, recall that our model predicts a crime rate v equal to

$$v = \frac{(1 - e^{-\nu p(k,v;\zeta_1,\zeta_2)})}{p(k, v; \zeta_1, \zeta_2)} \sum_{t=0} N_t^E,$$

where $p(k, v; \zeta_1, \zeta_2)$ is the probability that our model estimates under crime rate v, expenditures in police k, and police parameters (ζ_1, ζ_2). Assume that the number of agents that engage in crime increases by a multiplicative factor $(1 + \epsilon)$. This increases the crime rate through two channels. The first-order change is the mechanical increase of the crime rate by an additive factor of ϵv. The second-order change is the reduction in police effectiveness via the load effect. The new detection probability $\tilde{p}(\epsilon)$ would be

$$\tilde{p}(\epsilon) = p(k, (1 + \epsilon)v; \zeta_1, \zeta_2) = \frac{(1 - W_L(\zeta_2(1 + \epsilon)v) \exp(-\zeta_1 k + \zeta_2(1 + \epsilon)v))}{\zeta_2(1 + \epsilon)v}.$$

Therefore, adding these two factors (effects of equilibrium are negligible for small values of ϵ and, in addition, ζ_2, the main driver for these effects, is small), the crime rate would increase by approximately

$$\epsilon v + \sum_{t=0} N_t^E \frac{(1 - e^{-\nu \tilde{p}(\epsilon)})}{\tilde{p}(\epsilon)} - v.$$

Then, we make ϵ range from -0.01 to 0.01 with increments of 0.001 and for each ϵ we calculate the new crime rate. We regress the vector of number of new criminals on the vector of crime rates and the coefficient of this regression (say $\widehat{\beta}_1(\zeta_1, \zeta_2)$) is minus the derivative of crime in relation to the prison population. The elasticity of the crime rate on prison population is thus $-\widehat{\beta}_1(\zeta_1, \zeta_2)\widehat{N}^P/\widehat{v}$.

Results of the second stage of the GMM

Let Ξ_2 be the diagonal matrix with diagonal elements given by the variances calculated in Section D.2.3. The weight matrix of the second stage of the GMM is $(\Xi_2)^{-1}$. We use the particle swarm optimization with 100 points in the space $[-8,-6]\times[-16,-11]$ (for $\ln\zeta_1$ and $\ln\zeta_2$, respectively) to find the minimum score of the GMM. The covariance matrix $\boldsymbol{V}_2$ is obtained using the same procedure we used in the first stage.

D.3 Welfare Maximization

D.3.1 Optimal Policy

For each definition of welfare, we find the optimal police size and sentence length that maximize the welfare using the particle swarm optimization with 100 points in the space $[10, 40] \times [0.2, 0.8]$ (for police size k and probability of leaving jail μ, respectively). Recall that sentence length is given by $\lambda = 1/\mu - 1$.

D.3.2 Ellipse of Error

We detail how we obtained the covariance matrix for the vector of the estimated parameters $\widehat{\varphi}$ in Section D.2.2. We apply the same methodology to obtain the covariance matrix $\boldsymbol{V}_2$ of $(\widehat{\zeta}_1, \widehat{\zeta}_2)$. We perform the calculations using the optimal probability of being released from jail $\mu^\star$ instead of the optimal sentence length $\lambda^\star$. Then, in the end, we apply the delta method to calculate the covariance matrix of $(k^\star, \lambda^\star)$.

Let $\boldsymbol{V}_1$ be the covariance matrix of $\widehat{\varphi}$ excluding the line and the column corresponding to the estimate of $\widehat{p}$ (i.e., $\boldsymbol{V}_1$ is $\widehat{\boldsymbol{V}}_1$ after excluding the third line and third column of $\widehat{\boldsymbol{V}}_1$). We define $\phi = (b^C, c^C, \nu, b, \zeta_1, \zeta_2)$ and set the covariance matrix of the estimates $\widehat{\phi} = (\widehat{b}^C, \widehat{c}^C, \widehat{\nu}, \widehat{b}, \widehat{\zeta}_1, \widehat{\zeta}_2)$ as $\boldsymbol{V} = \left(\begin{smallmatrix} \boldsymbol{V}_1 & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{V}_2 \end{smallmatrix}\right)$. Then, applying the Delta method, we have that the (assymptotic) variance of the optimal policy is given by

$$\boldsymbol{V}_\mu = \begin{bmatrix} \nabla_\phi k^\star(\widehat{\phi}) \\ \nabla_\phi \mu^\star(\widehat{\phi}) \end{bmatrix} \mathbf{V} \begin{bmatrix} \nabla_\phi k^\star(\widehat{\phi}) \\ \nabla_\phi \mu^\star(\widehat{\phi}) \end{bmatrix}'.$$

Let the welfare loss be defined as $L(k, \mu; \phi)$, L_k be the first-order derivative of the social loss with respect to its first argument, and define L_μ, $L_{k\mu}$, $L_{\mu k}$, L_{kk}, and $L_{\mu\mu}$ analogously. Thus, from the first-order conditions, we obtain

$$L_k(k^\star(\widehat{\phi}), \mu^\star(\widehat{\phi}); \widehat{\phi}) = L_\mu(k^\star(\widehat{\phi}), \mu^\star(\widehat{\phi}); \widehat{\phi}) = 0.$$

Applying the implicit function theorem (and assuming that Clairaut's theorem applies), we have

$$\begin{bmatrix} \nabla_\phi k^\star(\widehat{\phi}) \\ \nabla_\phi \mu^\star(\widehat{\phi}) \end{bmatrix} = \begin{bmatrix} L_{kk}(k^\star(\widehat{\phi}), \mu^\star(\widehat{\phi}); \widehat{\phi}) & L_{k\mu}(k^\star(\widehat{\phi}), \mu^\star(\widehat{\phi}); \widehat{\phi}) \\ L_{k\mu}(k^\star(\widehat{\phi}), \mu^\star(\widehat{\phi}); \widehat{\phi}) & L_{\mu\mu}(k^\star(\widehat{\phi}), \mu^\star(\widehat{\phi}); \widehat{\phi}) \end{bmatrix}^{-1} \begin{bmatrix} \nabla_\phi L_k(k^\star(\widehat{\phi}), \mu^\star(\widehat{\phi}); \widehat{\phi}) \\ \nabla_\phi L_\mu(k^\star(\widehat{\phi}), \mu^\star(\widehat{\phi}); \widehat{\phi}) \end{bmatrix},$$

where $\nabla_\phi L$ is the gradient with respect to the parameters (from the third to the eighth components of L).

To calculate the second-order derivatives of the first matrix in the right-hand side, we create the vectors $X := (0.895k^\star, 0.902k^\star, \ldots, 1.105k^\star)$ and $Y := (0.895\mu^\star, 0.902\mu^\star, \ldots, 1.105\mu^\star)$. For each component of X and Y, say X_i and Y_j, we calculate the welfare loss $L(X_i, Y_j; \widehat{\phi})$. Then, we obtain the OLS estimation of the β's below

$$L(X_i, Y_j) = \beta_0 + \beta_1 X_i + \beta_2 Y_j + \beta_3 X_i^2 + \beta_4 Y_j^2 + \beta_5 X_i Y_j + \text{error}_{ij}.$$

for all pairs $(i, j) \in \{1, \ldots, 15\}^2$ (15 is the size of the vectors X and Y). We estimate the second-derivative of L with respect to its first argument as $\widehat{\beta}_3$, the second-derivative of L with respect to its second argument as $\widehat{\beta}_4$, and the cross derivative as $\widehat{\beta}_5$. That is, the first matrix is given by $\left(\begin{smallmatrix} \widehat{\beta}_3 & \widehat{\beta}_5 \\ \widehat{\beta}_5 & \widehat{\beta}_4 \end{smallmatrix}\right)$.

The procedure to calculate the components of the second matrix in the right-side hand is similar. Each component of $\nabla_\phi L_k$ is a cross derivative. For the first component, we define the $X := (0.895k^\star, 0.902k^\star, \ldots, 1.105k^\star)$ and $Y := (0.895b^C, 0.902b^C, \ldots, 1.105b^C)$. Let X_i be the i-th component of X and Y_j be the j-th component of Y. Defining $\phi^j := (Y_j, \widehat{c}^C, \widehat{\nu}, b, \widehat{\zeta}_1, \widehat{\zeta}_2)$, we set the first component of $\nabla_\phi L_k(k^\star(\widehat{\phi}), \mu^\star(\widehat{\phi}); \widehat{\phi})$ as the OLS estimate of β_5 from the model described below

$$L(X_i, \mu^\star; \phi_j) = \beta_0 + \beta_1 X_i + \beta_2 Y_j + \beta_3 X_i^2 + \beta_4 Y_j^2 + \beta_5 X_i Y_j + \text{error}_{ij}$$

for all pairs $(i, j) \in \{1, \ldots, 15\}^2$. We repeat this same procedure for all components of the vectors $\nabla_\phi L_k(k^\star(\widehat{\phi}), \mu^\star(\widehat{\phi}); \widehat{\phi})$ and $\nabla_\phi L_\mu(k^\star(\widehat{\phi}), \mu^\star(\widehat{\phi}); \widehat{\phi})$.

At this point, we have the covariance matrix $\boldsymbol{V}_\mu$ of optimal police expenditures, but sentence length is expressed by the hazard rate for leaving jail. We apply the Delta method once again to obtain the covariance matrix $\boldsymbol{V}_\lambda$ as a function of sentence length:

$$\boldsymbol{V}_\lambda = \begin{bmatrix} 1 & 0 \\ 0 & -\frac{1}{\mu^\star} \end{bmatrix}' \boldsymbol{V}_\mu \begin{bmatrix} 1 & 0 \\ 0 & -\frac{1}{\mu^\star} \end{bmatrix}.$$

Our estimation for $\boldsymbol{V}_\lambda$ is the following:

$$\boldsymbol{V}_\lambda = \begin{bmatrix} 0.2781 & -0.0016 \\ -0.0016 & 0.0046 \end{bmatrix}.$$

This implies that the standard error of our estimate for optimal police expenditures is \$1.33 billion. In the same way, the standard error of our estimate for the optimal average sentence length is 0.97 months. The error ellipses representing 99%, 95%, and 90% confidence sets are depicted in Figure A1.

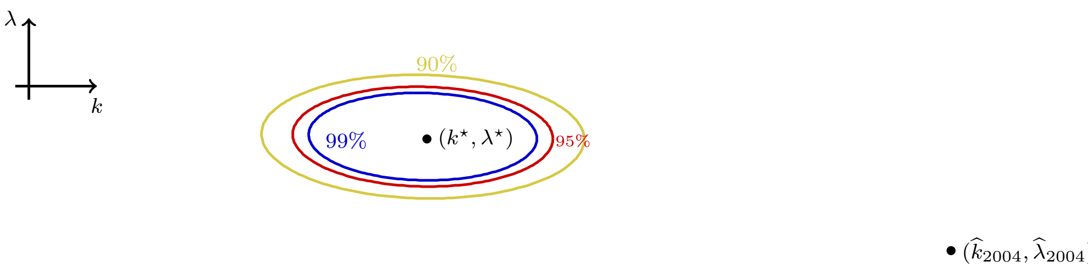

Figure A1: The error ellipses for our estimate of optimal policy. The yellow ellipse represents the boundary of the 90% confidence region, the red one 95%, and the blue one 99%. The optimal policy and the 2004 police are also indicated.

E Effects Decomposition

Four main channels can be used to decompose the equilibrium response of a given change in public security policy. First, the beliefs of dishonest agents regarding the policy, that is, deterrence. Second, the rate at which criminals are arrested and the time they are kept incarcerated, that is, incapacitation. Third, the effectiveness of the police at detecting crimes, that is, the load effect (in our model of detection, the higher is the crime rate, the less effective is the police). Finally, incarceration changes the productivity of individuals, which changes their future opportunity cost of committing crimes, that is, the recidivism effect.

These mechanisms interact mutually, but only first-order effects are relevant under small enough changes in public security policy. In particular, considering only changes in k (we can apply the same idea for generic changes in (k, λ)), the change in the crime rate v can be expressed by four components: (i) the belief regarding expenditures on police k^{b} (b for "belief"), determining the cut-off productivity for each age; (ii) actual expenditures on police k^{m} (m for "mechanical"), determining the rate of arrest and the hazard rate of leaving prison; (iii) the crime rate itself v, affecting the police effectiveness;[25] and (iv) the total mass of individuals $\overline{M}$ below the cut-off productivity for a given distribution of productivity.

In this theoretical construction, we consider that one channel can vary keeping the other channels fixed. For example, a change in k^{b} alters only the cut-off productivity at each age, but it does not change, for example, the productivity profile. Therefore, the mass of individuals committing crimes changes due to the change in the cut-off. By contrast, changing $\overline{M}$ means that the cut-offs are kept fixed, while the productivity varies—which also changes the mass of agents below the cut-off. Notice that if incarceration had no impact on productivity, then a change in k would not cause any change via recidivism. In the same way, a change in policy that maintained fixed the cut-off productivities would imply a null deterrence effect.

[25] It may seem odd to state that v is a function of v, but recall that v in equilibrium is a fixed point.

The total derivative of v with respect k can be written as:

$$\frac{\mathrm{d}}{\mathrm{d}k}v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M}) = \partial_1 v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M})\frac{\mathrm{d}k^{\mathrm{b}}}{\mathrm{d}k} + \partial_2 v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M})\frac{\mathrm{d}k^{\mathrm{m}}}{\mathrm{d}k} \\ + \partial_3 v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M})\frac{\mathrm{d}v}{\mathrm{d}k} + \partial_4 v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M})\frac{\mathrm{d}\overline{M}}{\mathrm{d}k}.$$

If, instead of derivatives, we consider small differences, we have that changing k to $k+\Delta k$ causes variations Δk^{b}, Δk^{m}, Δv, and $\Delta\overline{M}$. Define

$$\Delta v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M}) := v(k^{\mathrm{b}}+\Delta k^{\mathrm{b}},k^{\mathrm{m}}+\Delta k^{\mathrm{m}},v+\Delta v,\overline{M}+\Delta\overline{M}) - v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M}),$$

$$\Delta_1 v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M}) := v(k^{\mathrm{b}}+\Delta k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M}) - v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M}),$$
$$\Delta_2 v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M}) := v(k^{\mathrm{b}},k^{\mathrm{m}}+\Delta k^{\mathrm{m}},v,\overline{M}) - v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M}),$$
$$\Delta_3 v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M}) := v(k^{\mathrm{b}},k^{\mathrm{m}},v+\Delta v,\overline{M}) - v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M}),$$

and

$$\Delta_4 v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M}) := v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M}+\Delta\overline{M}) - v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M}).$$

Then, we have

$$\begin{aligned}\frac{\Delta v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M})}{\Delta k} &\approx \frac{\Delta_1 v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M})}{\Delta k^{\mathrm{b}}}\frac{\Delta k^{\mathrm{b}}}{\Delta k} + \frac{\Delta_2 v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M})}{\Delta k^{\mathrm{m}}}\frac{\Delta k^{\mathrm{m}}}{\Delta k} \\ &\quad + \frac{\Delta_3 v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M})}{\Delta v}\frac{\Delta v}{\Delta k} + \frac{\Delta_4 v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M})}{\Delta\overline{M}}\frac{\Delta\overline{M}}{\Delta k} \\ &= \frac{\Delta_1 v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M})}{\Delta k} + \frac{\Delta_2 v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M})}{\Delta k} + \frac{\Delta_3 v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M})}{\Delta k} + \frac{\Delta_4 v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M})}{\Delta k}.\end{aligned}$$

Each of the terms in the previous expression is a numerical derivative. To estimate them, we use the same procedure we applied when estimating other derivatives in the paper (see Section D.2.1).

For a given k, we create a vector $X := (-0.020k, -0.018k, -0.016k, \ldots, 0.020k)$. Then, we have a vector for several variations of k. For each variation of k, we find $\Delta_1 v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M})$, so we have a vector Y. Then, we find the OLS estimate for the coefficient β_1 of the regression $Y = \beta_0 + \beta_1 X + \text{error}$ and we repeat this procedure for the other three terms (Δ_2, Δ_3, and Δ_4).

We present now the details on how to calculate $\Delta_j v(k^{\mathrm{b}},k^{\mathrm{m}},v,\overline{M})$ for each $j = 1,2,3,4$. We assume that, in the equilibrium under policy (k,λ), the crime rate is v, the cut-off productivity at age t is $w_t^\star$, the detection probability is p, the mass of individuals at state s and age t with productivity lower than w is $M_t^s(w)$, and the total number of prisoners is N^P.

Deterrence

We have $\Delta k^{\mathrm{b}} = \Delta k$. In this case, the belief regarding the detection probability is given by

$$p^{\mathrm{b}} := 1 - \frac{W_L(\widehat{\zeta}_2 v e^{-\widehat{\zeta}_1(k+\Delta k)+\widehat{\zeta}_2 v})}{\widehat{\zeta}_2 v}.$$

Note that v is kept constant. This probability induces a belief of expected gains with crime for each age (say, $\overline{\Pi}_t^{\mathrm{b}}$) and a probability of incarceration (say $\bar{p}^{\mathrm{b}}$). Then, we calculate by backward induction the cut-off productivity for each age $w_t^{\star \mathrm{b}}$. Our Δ_1 is given by

$$\Delta_1 v(k^{\mathrm{b}}, k^{\mathrm{m}}, v, \overline{M}) = \frac{1 - e^{-\widehat{\nu} p}}{p} \sum_{t=0}^{T} (M_t^E(w_t^{\star \mathrm{b}}) - M_t^E(w_t^{\star})).$$

Incapacitation

We have $\Delta k^{\mathrm{m}} = \Delta k$. To simplify the explanation, we assume that $\Delta k > 0$ (the idea for $\Delta k < 0$ is analogous). This defines a detection probability

$$p^{\mathrm{m}} := 1 - \frac{W_L(\widehat{\zeta}_2 v e^{-\widehat{\zeta}_1(k+\Delta k)+\widehat{\zeta}_2 v})}{\widehat{\zeta}_2 v}$$

and induces a probability of incarceration $\bar{p}^{\mathrm{m}}$. This alters the crime rate due to both intratemporal and intertemporal reasons. First, it decreases the expected number of crimes each individual engaged in crime commits. This happens because criminals would be detected earlier in the period or, alternatively, criminals that are not detected under a detection probability p would be under p^{m}. Second, since more individuals would be arrested, more individuals with low productivity would be incarcerated in future periods. This means that $\Delta_2 v(k^{\mathrm{b}}, k^{\mathrm{m}}, v, \overline{M})$ is given by

$$\begin{aligned}\Delta_2 v(k^{\mathrm{b}}, k^{\mathrm{m}}, v, \overline{M}) = \sum_{\tau=0}^{T} M_\tau^E(w_\tau^\star) \left[\frac{1 - e^{-\widehat{\nu} p^{\mathrm{m}}}}{p^{\mathrm{m}}} - \frac{1 - e^{-\widehat{\nu} p}}{p} \right] \\ - (\bar{p}^{\mathrm{m}} - \bar{p}) \frac{1 - e^{-\widehat{\nu} p}}{p} \sum_{\tau=0}^{T} \sum_{t=\tau+1}^{T} M_\tau(w_t^\star / \Gamma(\tau, t)) \left((1-\mu)^{t-\tau-1} \right).\end{aligned}$$

The first (resp. second) term on the left-hand side represents the intratemporal (resp. intertemporal) crime change via incapacitation. To better understand the second term, consider an individual at age 0 (so $\tau = 0$) engaged in crime. His/her probability of being in prison at age 1 would increase by $(\bar{p}^{\mathrm{m}} - \bar{p})$. If he/she is caught, he/she will be in prison at age 1 (since $t = 1$, $t - \tau - 1 = 0$, and $(1-\mu)^{t-\tau-1} = 1$) and, if his/her productivity at age 1 would be below $w_1^\star$, he/she would not commit crimes at age 1 (which would amount to $\frac{1-e^{-\widehat{\nu} p}}{p}$ fewer crimes).

Load

We calculate the variation of v in the standard model when we vary k. This gives us our Δv. Then, we calculate the following detection probability

$$p^{\text{load}} := 1 - \frac{W_L(\widehat{\zeta}_2(v+\Delta v)e^{-\widehat{\zeta}_1 k+\widehat{\zeta}_2(v+\Delta v)})}{\widehat{\zeta}_2(v+\Delta v)}.$$

Ignoring equilibrium effects for small variations in k, as usual, we calculate the crime rate that a detection probability p^{load} would generate. Subtracting this crime rate from the equilibrium one gives us $\Delta_3 v(k^{\text{b}}, k^{\text{m}}, v, \overline{M})$.

Recidivism

First, we can calculate the measure of individuals with productivity w, for each $w > 0$, given by $\overline{m}^E(w) + \overline{m}^{\tilde{E}}(w) + \overline{m}^P(w)$ under the policy $(k + \Delta k, \lambda)$. Then, given this measure, we calculate the mass of criminals below the productivity cutoffs generated by the policy (k, λ) (recall that only deterrence changes these cutoffs). This is given by

$$\overline{A} := \sum_{t=0}^{T} \int_0^{w_t^\star} (\overline{m}^E(w) + \overline{m}^{\tilde{E}}(w) + \overline{m}^P(w))\mathrm{d}w.$$

So $\overline{A}$ is the mass of individuals that would commit crimes if they were free. However, we know that N^P of them are in prison (under the policy (k, λ)), which means that

$$\Delta_4 v(k^{\text{b}}, k^{\text{m}}, v, \overline{M}) := (\overline{A} - N^P)\frac{1 - e^{-\widehat{\nu} p}}{p} - v.$$

Computing the effects

We obtain the four effects for each value of k ranging from \$5 billion to \$40 billion with increments of \$1 billion. The plot of the marginal incapacitation effect as a function of k in Figure 5 shows precisely the numerical values we obtain from this procedure. For the three remaining effects, we plot smooth curves that fit the values obtained from our numerical procedure (minimizing the squared error; the values from our procedure are noisy for some of the calculations because of numerical approximations). The actual values and the fitted curves are shown in Figure A2. The function we use to fit the marginal deterrence effect has the form $a^{\text{det}} + b^{\text{det}} \exp(c^{\text{det}} k)$, and we find $a^{\text{det}} = 153$, $b^{\text{det}} = -9863$, and $c^{\text{det}} = -0.115$. We use the same functional form to represent the marginal load effect, and we find $a^{\text{load}} = 7.768$, $b^{\text{load}} = -831.2$, and $c^{\text{load}} = -0.168$. The functional form we use to represent the recidivism effect is $k(a^{\text{rec}} + kb^{\text{rec}} + k^2 c^{\text{rec}}) \exp(kd^{\text{rec}})$, and we find $a^{\text{rec}} = 140.5$, $b^{\text{rec}} = -17.84$, $c^{\text{rec}} = 0.08246$, and $d^{\text{rec}} = -0.1539$. Since the fitted curves are simply descriptive tools, in order to choose these functional forms we experimented with some alternatives and chose the best fits by visual inspection.

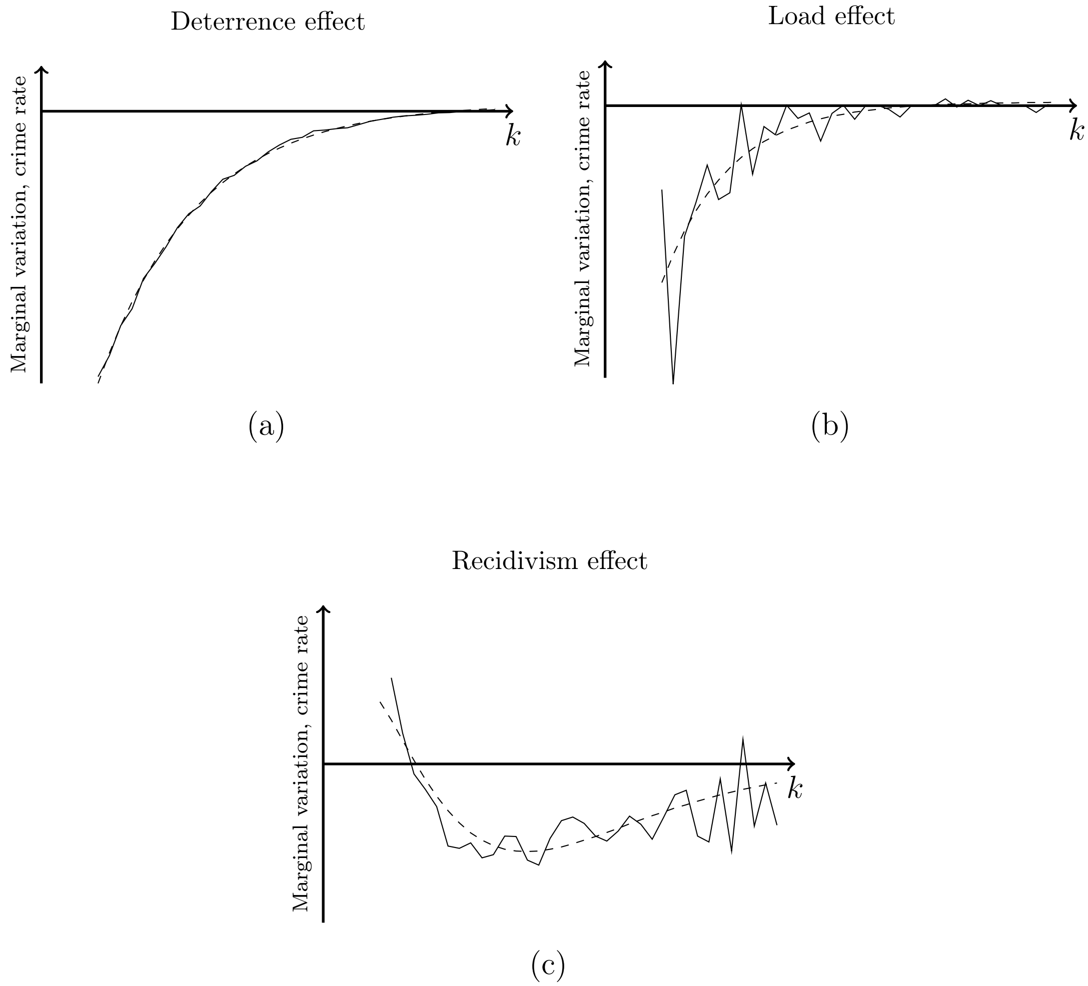

Figure A2: Each panel depicts two versions of a marginal effect as a function of k. The solid curves are the actual values we obtain in the simulations; the dashed ones are the fitted curves for specific functional forms.

F Education Exercises

F.1 Finding the New Distribution of Productivity at Age 0

Assume that additional investments in education are given by e at each period and that the internal rate of return (*IRR*) for educational investments is r. That is, in the stationary state, investing e in education increases the total mass of productivity by $(1+r)e$. To calculate η, we ignore the impact of incarceration on aggregate productivity. This is a harmless approximation for two reasons. First, under the 2004 policy, we estimate that around 1.8% of the total population were or would be convicted and incarcerated. Second, the productivity of prisoners is already quite low, so as they lose a fraction of their productivities, the relative loss (compared to the total mass of productivity) is negligible.

Now, we pin down η. Let w_t be the productivity of a random agent at age t (using the assumption discussed in the previous paragraph). The mass of productivity at age t is given by

$$M_0 \int \Gamma_t w \mathrm{d}G_0(w) = M_0 \mathbb{E}\left[w_t\right] = M_0 \Gamma_t \mathbb{E}\left[w_0\right].$$

Let $\tilde{w}_t$ be the productivity of a random individual at age t under the new educational policy. Then, we have

$$\begin{aligned}
e(1+r) &= M_0 \sum_{t=0}^{T} \beta^t (\mathbb{E}\left[\tilde{w}_t\right] - \mathbb{E}\left[w_t\right]) + \alpha M_0 \beta^{T+1} (\mathbb{E}\left[\tilde{w}_{T+1}\right] - \mathbb{E}\left[w_{T+1}\right]) \\
&= M_0 \sum_{t=0}^{T} \beta^t \Gamma_t (\mathbb{E}\left[\tilde{w}_0\right] - \mathbb{E}\left[w_0\right]) + \alpha M_0 \beta^{T+1} \Gamma_{T+1} (\mathbb{E}\left[\tilde{w}_0\right] - \mathbb{E}\left[w_0\right]) \\
&= M_0 (\mathbb{E}\left[w_0 + \eta w_0 e^{-\psi w_0}\right] - \mathbb{E}\left[w_0\right]) \Big(\alpha \beta^{T+1} \Gamma_{T+1} + \sum_{t=0}^{T} \beta^t \Gamma_t\Big) \\
&= M_0 \eta \mathbb{E}\left[w_0 e^{-\psi w_0}\right] \Big(\alpha \beta^{T+1} \Gamma_{T+1} + \sum_{t=0}^{T} \beta^t \Gamma_t\Big),
\end{aligned}$$

and

$$\eta = \frac{e(1+r)}{M_0 \mathbb{E}\left[w_0 e^{-\psi w_0}\right] \Big(\alpha \beta^{T+1} \Gamma_{T+1} + \sum_{t=0}^{T} \beta^t \Gamma_t\Big)}.$$

We take the *IRR* from Heckman et al. (2008). They estimate an *IRR* of 14% for white men and 18% for black men. We use $\widehat{r} = 0.15$ but also perform the same exercise assuming an internal rate of return equal to 0%.

With the values of η and r pinned down, we calculate the new productivity distribution at age 0 as a function of e, defined as g^e. We also define its cumulative distribution as G^e. Since the map $f^e : w \mapsto w(1 + \eta e^{-\psi w})$ has a well-defined inverse, we have

$$G^e(w) = G(f^{e-1}(w)).$$

We calculate $g^e(w)$ numerically by using

$$g^e(w) = \frac{G(f^{e-1}(w)) - G(f^{e-1}(w - \Delta w))}{\Delta w}.$$

F.2 Defunding the Public Security Policy

Fix a focalization ψ. Consider a policy now as a triplet, with the two first components being the public security policy (PSP) and the third one being a non-negative amount of additional investments in education. This triplet determines an equilibrium. Now, suppose we move from the policy $(k, \lambda, 0)$ to (k', λ', e). Define the equilibrium financial cost of the PSP as C_{PSP} for the first policy and C'_{PSP} for the second one. We say that the second policy results from a PSP defunding if $e \geq 0$ and $e = C_{\text{PSP}} - C'_{\text{PSP}}$, and we set e as the defund value.

The value of the defund is a fixed point, and we calculate it iteratively. Suppose we change the PSP from (k, λ) to (k', λ'). Our first guess for the (extra) investment in education is $e_0 := k - k'$. The policy (k', λ', e_0) implies a financial cost with PSP given by C^0_{PSP}. Then, we set $e_1 := C_{\text{PSP}} - C^0_{\text{PSP}}$, and, more generally, $e_n := C_{\text{PSP}} - C^{n-1}_{\text{PSP}}$. We stop when $|e_n - e_{n-1}|$ is smaller than some threshold.

F.3 Optimal Policy under PSP Defunding

Fix the values of ψ and r, and let the initial policy be $(k, \lambda, 0)$. Let $\mathcal{E}$ be the set of PSPs resulting from a PSP defunding of (k, λ). For any arbitrary PSP (k', λ'), we can calculate the defund value $e(k', \lambda')$ and the welfare loss of the policy $(k', \lambda', e(k', \lambda'))$, redefined with three arguments: $L(k', \lambda', e(k', \lambda'))$. In particular, we can calculate the optimal policy when it is possible to shift PSP resources towards education:

$$(k^{\text{ed}}, \lambda^{\text{ed}}) = \underset{(k', \lambda') \in \mathcal{E}}{\operatorname{argmax}} \{L(k', \lambda', e(k', \lambda'))\}.$$

We perform this exercise to find new optimal policies for values of focalization in the set $\{0.00, 0.02, 0.04, 0.06, 0.08, 0.10\}$ and *IRR* of 15% and 0%.