

Residual Deterrence*

Francesc Dilmé

University of Bonn

fdilme@uni-bonn.de

Daniel F. Garrett

Toulouse School of Economics, University of Toulouse Capitole

daniel.garrett@tse-fr.eu

Spring 2017

Abstract

Successes of law enforcement in apprehending offenders are often publicized events. Such events have been found to result in temporary reductions in offending, or “residual deterrence”. We provide a theory of residual deterrence which accounts for the incentives of both enforcement officials and potential offenders. We do so by introducing to a standard inspection framework costs that must be incurred to commence enforcement. Such costs in practice include hiring specialized staff, undertaking targeted research and coordinating personnel. We illustrate how our model can be used to address a number of policy questions regarding the optimal design of enforcement authorities.

JEL classification: C73, K42

Keywords: deterrence, reputation, switching costs

*We would like to thank Heski Bar-Isaac, Bruno Biais, Simon Board, Daniel Chen, Nuh Aygün Dalkıran, Jan Eeckhout, Marina Halac, Johannes Horner, George Mailath, Moritz Meyer-ter-Vehn, Stephan Laueremann, Alessandro Pavan, Patrick Rey, William Rogerson, and seminar participants at Bilkent University, the CESifo in Munich, the CSIO/IDEI Joint Workshop at TSE, ESSET at Gerzensee, the London School of Economics, the NASM ES 2016 in Philadelphia, the Paris School of Economics, SITE at Stanford University, the 1st Triangle Microeconomics Conference at UNC Chapel Hill, the Universitat Autònoma in Barcelona, and the University of Edinburgh for helpful discussions.

1 Introduction

An important rationale for the enforcement of laws and regulations concerns the deterrence of undesirable behavior. The illegality of actions, by itself, may not be enough to dissuade offenders. Instead, the perceived threat of apprehension and punishment seems to play an important role (see Nagin, 2013, for a review of the evidence).

One factor that is salient in determining the perceived risk of punishment is past enforcement decisions, such as the extent of past convictions, fines or arrests. Evidence can be found in a range of settings, including: reductions in mark-ups of bread manufacturers in response to Department of Justice price fixing prosecutions (Block, Nold and Sidak, 1981); reductions in water pollution violations by paper mills in response to Environmental Protection Agency fines against other nearby mills (Shimshack and Ward, 2005); reductions in the extent of “aggressive” financial reporting in industries subject to recent Securities and Exchange Commission enforcement actions (Jennings, Kedia and Rajgopal, 2011, Schenk, 2012, and Brown et al., 2014); and reductions in drink driving and other personal offending following police crackdowns on specific crimes (Sherman, 1990, Taylor, Koper and Woods, 2013, and Banerjee et al., 2014). These observations fit a pattern which Sherman terms “residual deterrence”. Residual deterrence occurs when reductions in offending follow a phase of active or intensive enforcement, and persist even after this enforcement phase has ended.

In the above examples, the possibility of residual deterrence seems to depend, at first instance, on the perceptions of potential offenders about the likelihood of detection. It is then important to understand: How are potential offenders’ perceptions determined? What affects the extent and duration of residual deterrence (if any)? This paper aims at an equilibrium explanation of residual deterrence based on both the motives of enforcement officials (for concreteness, the “regulator” in our model) and potential offenders (the “firms”). In particular, we provide a model in which convictions sustained by the regulator against offending firms are followed by prolonged periods of low offending, which we equate with residual deterrence.

Our theory posits a self-interested regulator which gains by apprehending offending firms but finds inspections costly. Since firms are deterred only if the regulator is likely to be inspecting, the theory must explain why the regulator continues to monitor firms even when they are unlikely to offend. Our explanation hinges on the additional costs a regulator faces when commencing *new* inspections. We show how such costs can manifest in the episodes of residual deterrence that follow the apprehension of an offending firm.

There are often a myriad of costs a regulator faces when beginning to monitor a particular industry or to enforce a particular regulation; in our model, when inspecting following at least one period of inactivity. For instance, when the Federal Trade Commission chose to crack down on modern privacy offenses by large corporations in the late 2000s, it bore the start-up costs of a new forensics laboratory, employing new experts and purchasing new equipment.¹ More generally, a regulatory authority investigating new offenses or industries must engage in specialized research (even if by existing personnel) to come up to date with industry dynamics and relevant case law (Kovacic and Hyman, 2016, describe such research as a form of regulatory “R&D” that is essential for successful enforcement). Even in instances where the authority possesses the relevant knowledge and expertise, coordinating new enforcement activities requires time and planning that is often costly.²

We study a dynamic version of a simple workhorse model – the inspection game. In this game, a long-lived regulator faces a sequence of short-lived firms. Committing an offense is only worthwhile for a firm if the regulator is not inspecting, while inspecting is only worthwhile for the regulator if an offense is committed. In our baseline model (in Section 3), the only public information is the history of previous “convictions”; that is, the periods where the regulator was inspecting and the firm committed an offense.

¹See Schectman (2014, January 22).

²Related, there may be costs stemming from the uncertainty of undertaking a new activity. Or, there may be “psychological costs” of changing the current pattern of activity (see, e.g., Klemperer, 1995).

This corresponds to a view that the most salient action an enforcement authority can take is to investigate and penalize offending. It is through convictions that firms learn that the regulator has been active (say, investigating a particular instance of price fixing or cracking down on financial mis-statements by one of its peers).

In the repeated version of the inspection game described above, equilibrium follows a repetition of static play; i.e., past convictions do not affect the rate of offending. Things are different once we introduce the additional cost of commencing inspections. We show that equilibrium then features reputational effects: a conviction is followed by several periods during which firms believe the probability of inspection is relatively high and they are therefore less likely to offend. We identify this pattern as residual deterrence. Over time, in the absence of a conviction, firms gradually update their beliefs and the perceived likelihood of inspection falls, eventually reaching its lowest (the original, or “baseline”) level. Offending likewise rises to the baseline level. This pattern corresponds to what Sherman (1990) has termed “deterrence decay”.

The equilibrium pattern we uncover might be described in terms of “reputation cycles”. Each cycle is characterized by a conviction, a subsequent reduction in offending, and finally a resumption of offending at the baseline level. We show that the additional costs of commencing inspections are necessary to generate these cycles, since the episodes of residual deterrence disappear as the initiation costs shrink. The role of these costs is that the regulator remains willing to inspect while firms are deterred simply to avoid incurring the same costs when offending resumes at the baseline level. Indeed, continuing to vigilantly monitor for offenses is a natural way to *maintain* the initial investment in information or personnel that facilitates inspections.³

Our model can be used to shed light on a number of policy questions, such as determinants of the length of residual deterrence and the overall offense rate. We are also able to address the important question of optimal information disclosure: should a designer of the regulatory authority require its activities to be disclosed?

³Since a positive level of offending persists even during periods of residual deterrence, the chance to obtain additional convictions is always a further motivation for continuing inspections.

Analysis of such questions is enabled by the simplicity of our approach, and the fact that the equilibrium process for offenses and convictions is uniquely determined. The only source of multiple equilibria is that the regulator may condition its switching on privately-held and payoff-irrelevant past information.

It is worth emphasizing up front that we focus on a regulator concerned with obtaining convictions, as opposed, for instance, to deterrence itself. This specification seems to make sense in many settings, since the allocation of enforcement resources often rests on the discretion of personnel influenced by organizational incentives. Kovacic and Hyman (2016) argue that competition authorities face strong incentives to prosecute and take new cases; for instance, he notes that the *Global Competition Review's* (2015) evaluation of “top antitrust authorities” focuses on successful prosecutions. Similarly, Benson, Kim and Rasmussen (1994, p 163) argue that police “incentives to watch or patrol in order to prevent crimes are relatively weak, and incentives to wait until crimes are committed in order to respond and make arrests are relatively strong”. While financial rewards for law enforcers to catch offending are often controversial or illegal, even quite explicit incentives can arise. Perhaps the best-known example in recent times was the Ferguson Police Department’s focus on generating revenue by writing tickets.⁴

While we believe that a direct concern for obtaining convictions is often the most relevant objective, we are able to extend our model to settings where the regulator is concerned instead with deterrence. Again, we exhibit equilibria featuring reputation cycles. We find that the regulator in this case may be incentivized to inspect precisely because it anticipates residual deterrence following a conviction. However, we also find several differences in the equilibrium predictions for regulators motivated by deterrence rather than directly by a concern for convictions.

The rest of the paper is as follows. We next briefly review literature on the

⁴See the United States Department of Justice Civil Rights Division (2015). Note that the model we introduce below can explicitly account for this revenue-raising motive for inspections, for instance by setting the regulator’s reward for a conviction equal to the firm’s penalty.

economic theory of deterrence, as well as on reputations. Section 2 introduces the environment. Section 3 solves for equilibrium and discusses comparative statics. Section 4 provides further discussion, contrasting our theory with alternative explanations of residual deterrence. Section 5 examines the role of information disclosure, and Section 6 considers a regulator motivated by a desire to deter offending. Section 7 concludes.

1.1 Literature Review

At least since Becker (1968), economists have been interested in the deterrence role of policing and enforcement. Applications include not only criminal or delinquent behavior, but also the regulated behavior of firms such as environmental emissions, health and safety standards and anticompetitive practices. This work typically simplifies the analysis by adopting a static framework with full commitment to the policing strategy. The focus has then often been on deriving the optimal policies to which governments, regulators, police or contracting parties should commit (see, among others, Becker, 1968, Townsend, 1979, Polinsky and Shavell, 1984, Reinganum and Wilde, 1985, Mookherjee and Png, 1989 and 1994, Lazear, 2006, Bond and Hagerty, 2010, and Eeckhout, Persico and Todd, 2010).

In practice, however, there are limits to the ability of policy makers to credibly commit to the desired rate of policing. First, policing itself is typically delegated to agencies or individuals whose motives are not necessarily aligned with the policy maker's. Second, announcements concerning the degree of enforcement or policing may not be credible (see Reinganum and Wilde, 1986, Khalil, 1997, and Strausz, 1997, for settings where the principal cannot commit to an enforcement rule, reflecting the concerns raised here). Potential offenders are thus more likely to form judgments about the level of enforcement activity from past observations. To our knowledge, formal theories of the reputational effects of policing are, however, absent from the enforcement literature.

Our paper is related to the literature on reputations with endogenously switching

types; see for instance Mailath and Samuelson (2001), Iossa and Rey (2014), Board and Meyer-ter-Vehn (2013, 2014) and Halac and Prat (2016). Closest methodologically to our paper is the work by Dilmé (2014). Dilmé follows Mailath and Samuelson and Board and Meyer-ter-Vehn by considering firms that can build reputations for quality (see also Iossa and Rey in this regard), but introduces a switching cost to change the quality level. The present paper also features a switching cost for the long-lived player (we posit a cost only in case switching from “not inspect” to “inspect”, although our working paper version, Dilmé and Garrett, 2015, considers costs of switching in both directions). However, the stage game is different to Dilmé’s, requiring a separate analysis.

A key novelty of our setting relative to the various papers on seller reputation is that the public information depends on the actions of all players, both the regulator and firms. This feature is in common with Halac and Prat (2016), who analyze the deterioration of manager-worker relationships.⁵ They find an equilibrium with similar features to ours in the so-called “bad news” case, where the worker increases his effort immediately after being found shirking, since he believes that a monitoring technology is likely to be in place. However, there are several important differences to our paper. First, as we discuss further in Section 3.2, our regulator has the *ability* to cease inspecting, whereas the monitoring technology breaks exogenously (and randomly) in Halac and Prat’s model. This permits us to tackle directly the question of *why* the regulator may continue inspecting although firms are temporarily deterred from offending. We also find that the regulator actively chooses to stop inspecting in equilibrium, to avoid the costs of inspection, and this is the source of “deterrence decay” in our model. In Halac and Prat’s equilibrium, a related pattern of decay results instead from exogenous failure of the inspection technology. Second, we make a range of different modeling choices, motivated by applications of regulation and enforcement: we study a regulator directly motivated by convictions (rather than reputation) and

⁵This work arose independently and (as far as we know) simultaneously to our own.

firms with heterogeneous preferences for offending. Apart from requiring a novel analysis, these choices have implications for our key equilibrium predictions. Third, we study questions of interest for policy, such as the determinants of residual deterrence and the optimal choice of information disclosure policy.

Finally, our model (as well as Dilmé’s, 2014, above) is related to models of dynamic games with switching costs and public actions, such as Lipman and Wang (2000, 2009) and Caruana and Einav (2008). Relative to these papers, we focus on a setting with incomplete information regarding the long-run player’s actions. While the earlier works emphasize the possibility that small switching costs result in high persistence of equilibrium actions, persistence in our model is determined by the combination of switching costs and incomplete information.

2 Environment

2.1 The stage game

We study a long-lived regulator interacting at discrete dates $t = 0, 1, 2, \dots$ with an infinite number of firms, one per period. While each firm survives only a single period, the regulator is forward-looking with discount factor $\delta \in [0, 1)$.

Actions. At the beginning of any period t , the firm independently (and privately) draws a value π_t from a continuous distribution F with full support on a finite interval $[\underline{\pi}, \bar{\pi}]$, $0 < \underline{\pi} < \bar{\pi}$. Then, the firm chooses an action $a_t \in \{O, N\}$ where O denotes “offend” and N denotes “do not offend”. Simultaneously, the regulator chooses an action $b_t \in \{I, W\}$, where I denotes “inspect” and W denotes “wait”.

Somewhat abusively, we let $I = O = 1$ and $W = N = 0$. Thus $a_t b_t = 1$ if the firm offends while the regulator inspects at date t , while $a_t b_t = 0$ otherwise. If $a_t b_t = 1$, we say that the regulator “obtains a conviction” at date t .⁶

⁶We will assume that a firm can only be convicted in the period it takes its action a_t . One way to interpret this is that evidence of an offense lasts only one period. This seems unambiguously the

Payoffs. Period- t payoffs are determined according to a standard inspection game. If the firm offends without a conviction ($a_t = O$ and $b_t = W$), then it earns a payoff π_t . If it offends and is convicted ($a_t = O$ and $b_t = I$), then it sustains a penalty $\gamma > 0$, which is net of any benefits from the offense. Otherwise, its payoff is zero.

If the regulator inspects at date t , it suffers a cost $(1 - \delta)i > 0$, where the factor $(1 - \delta)$ represents a normalization based on the “length” of a period. It incurs no cost if waiting.⁷ In the event of obtaining a conviction, the regulator earns an additional lump-sum payoff of $y > (1 - \delta)i$. Later, we consider the possibility that the regulator cares about deterring firms rather than obtaining convictions.

In addition to the costs and benefits specified above, the regulator sustains a cost $S > 0$ if commencing inspection at period t .⁸ Let $\mathbf{1}(b_{t-1}, b_t)$ take value one if $(b_{t-1}, b_t) = (W, I)$ and equal zero otherwise. Period t payoffs are then summarized in the following table.

		firm	
		$a_t = N$	$a_t = O$
regulator	$b_t = W$	0, 0	0, π_t
	$b_t = I$	$-(1 - \delta)i - S \mathbf{1}(b_{t-1}, b_t), 0$	$y - (1 - \delta)i - S \mathbf{1}(b_{t-1}, b_t), -\gamma$

To ensure that inspections occur in equilibrium, we assume that inspection and

right assumption where punishment requires the offender to be “caught in the act”. More generally, it seems a reasonable simplification, one which has often been adopted, for instance, by the literature on leniency programs for cartels (see, e.g., Spagnolo (2005) and Aubert, Rey and Kovacic (2006)). One way to relax the assumption would be to assume that while firms take only one action, they can still be convicted for a limited time subsequently. We expect residual deterrence would continue to arise in equilibrium in this model.

⁷In general, “waiting” corresponds to a broad range of alternative actions the regulator could devote time to in a given period. For simplicity, we do not explicitly model these alternatives in the present paper.

⁸Using a similar argument as in Dilmé (2014) it is easy to see that the assumption that there is no cost of stopping inspection is without loss of generality, as otherwise it can be renormalized to zero.

commencement costs are not too large in the following sense.

Assumption 1 *The stage game is always an inspection game; i.e., $y - (1 - \delta)i - S > 0$.*

A brief analysis of the stage game allows us to anticipate the role of commencement costs in determining the equilibrium offense rate, as we explain in the following remark.

Remark 1 (Static analysis) *By Assumption 1, the stage game is a standard static inspection game in which the regulator faces a cost $(1 - \delta)i$ from inspecting if it inspected in the previous period, or a cost $(1 - \delta)i + S$ otherwise. As is well understood, these costs determine the equilibrium offense rate in the one-shot game (where we ignore the effect of the regulator's actions on its future payoffs). In particular, if it is common knowledge that the regulator played "wait" in the previous period, then the equilibrium probability of offending is $\frac{(1-\delta)i+S}{y}$. Conversely, if the regulator is commonly known to have inspected, then the equilibrium offense probability is $\frac{(1-\delta)i}{y}$. These offense probabilities ensure the regulator is indifferent between its two actions, wait and inspect.*

The same predictions continue to hold if instead the firm is commonly known to believe that the regulator previously played, respectively, "wait" or "inspect" with sufficiently high probability in the previous period. This indicates that it is the firm's beliefs regarding the regulator's costs which determine the probability of offending. When we study repeated play, these beliefs will be determined by the convictions that occur in equilibrium (for instance, a conviction will be taken as evidence that the regulator inspected in a given period).

While the static analysis is suggestive, it fails to shed light on at least two aspects of equilibrium play. First, past information regarding inspections emerges endogenously through past convictions that ought to be explicitly modeled. Second, and crucially, the regulator's decision to inspect at a given date lowers the total inspection costs at future dates. Whether to inspect is therefore a dynamic decision that must account for the future evolution of offense probabilities.

We now complete the dynamic specification of the model, defining the available information, player strategies and equilibrium. Because changes in the regulator's actions affect payoffs, it is necessary to specify the regulator's action in the period before the game begins. For concreteness we let $b_{-1} = W$, although no results hinge on this assumption.

Information. In each period t , a public signal may be generated providing information on the players' actions. If a signal is generated, we write $h_t = 1$; otherwise, $h_t = 0$. Motivated by the idea that the activity of an enforcement agency becomes known chiefly through enforcement actions themselves, we focus on the case where a signal is generated on the date of a conviction. That is, for each date t , we let $h_t = a_t b_t \in \{0, 1\}$. Players perfectly recall the signals so that, at the beginning of period t , the date- t firm observes the "public history" $h^t \equiv (h_0, \dots, h_{t-1}) \in \{0, 1\}^t$. We find it convenient to let $\mathbf{0}^\tau = (0, 0, \dots, 0)$ denote the sequence of τ zeros. Thus, for $j > 1$, $(h^t, \mathbf{0}^j) = (h_0, \dots, h_{t-1}, 0, \dots, 0)$ is the history in which h^t is followed by j periods without a conviction.

The regulator observes both the public history and its private actions. Thus a private history for the regulator at date t is $\tilde{h}^t \equiv (h^t, b^t)$, where $b^t = (b_0, \dots, b_{t-1}) \in \{I, W\}^t$ is the sequence of regulator actions up to date t .

Strategies, equilibrium and continuation payoffs. We anticipate that firms will choose cut-off strategies in equilibrium, with the date- t firm offending if and only if $\pi_t \geq \pi(h^t)$, where $\pi(h^t)$ is the threshold at public history h^t . The cut-off $\pi(h^t)$ then implies a probability of offending $\alpha(h^t) \in [0, 1]$ at history h^t , and we find it useful to describe each firm's strategy in terms of this probability. We use α_t to denote $\alpha(h^t)$ when there is no risk of confusion. A (behavioral) strategy for the regulator assigns to each private history $\tilde{h}^t \in \{0, 1\}^t \times \{I, W\}^t$ the probability that the regulator inspects at \tilde{h}^t , $\beta(\tilde{h}^t)$. We study perfect Bayesian equilibria of the above game.

For a fixed strategy β of the regulator, we find the following abuse of notation convenient. For each public history h^t , let $\beta(h^t) \equiv \mathbb{E}[\beta(\tilde{h}^t)|h^t]$ be the equilibrium

probability that the regulator inspects at time t as determined according to the strategy β , where the expectation is taken with respect to the distribution over private histories \tilde{h}^t with public component h^t . We use β_t to denote $\beta(h^t)$ when there is no risk of confusion. Probabilities $\beta(h^t)$ (which we term the firms’ “perceived probability of inspection”) are particularly useful since (i) a date- t firm’s payoff is affected by \tilde{h}^t only through $\beta(h^t)$, (ii) these probabilities will be determined uniquely across equilibria of our baseline model, and (iii) in many instances, we might expect an external observer to have data only on the publicly observable signals (that is, convictions). In contrast, equilibrium strategies for the regulator, as a function of private histories, will not be uniquely determined.

It is now useful to define the continuation payoff of the regulator at any date t and for any strategies of the firm and regulator. For a regulator history \tilde{h}_t , this is

$$V_t(\beta, \alpha; \tilde{h}^t) = \mathbb{E}_{\beta, \alpha} \left[\sum_{s=t}^{\infty} \delta^{s-t} (yb_s a_s - (1 - \delta)ib_s - S \mathbf{1}(b_{s-1}, b_s)) | \tilde{h}^t \right].$$

Under an optimal strategy for the regulator and for a fixed public history, the regulator’s payoffs must be independent of all but the last realization of $b \in \{I, W\}$. We thus denote equilibrium payoffs for the regulator following public history h^t and date $t - 1$ choice b_{t-1} by $V_{b_{t-1}}(h^t)$.

3 Equilibrium Characterization

At a given period t , if the probability that the regulator inspects (conditional on the public history) is β_t , a firm only offends if $\frac{\pi_t}{\pi_t + \gamma} \geq \beta_t$. This implies that the probability of offending is given by

$$\alpha_t = \Pr \left(\frac{\pi_t}{\pi_t + \gamma} \geq \beta_t \right) = \Pr \left(\pi_t \geq \frac{\beta_t}{1 - \beta_t} \gamma \right) = 1 - F \left(\frac{\beta_t}{1 - \beta_t} \gamma \right) \equiv \alpha(\beta_t), \quad (1)$$

where our definition of $\alpha(\cdot)$ involves an obvious abuse of notation. Let $\underline{\beta} \equiv \frac{\pi}{\pi + \gamma}$ and $\bar{\beta} \equiv \frac{\bar{\pi}}{\bar{\pi} + \gamma}$, that is, $\alpha(\underline{\beta}) = 1$ and $\alpha(\bar{\beta}) = 0$. Our first result is that the equilibrium inspection probability at any history h^t lies above $\underline{\beta}$ and below $\bar{\beta}$, so that firms never

offend with probability one ($\alpha(\beta_t)$ is never one) and deterrence is never perfect ($\alpha(\beta_t)$ is never zero).

Lemma 1 *For all equilibria, for all h^t , $\beta(h^t) \in (\underline{\beta}, \bar{\beta})$.*

To help understand this result, note that the proof in the Appendix argues the following. Assume, for the sake of contradiction, that $\beta(h^t) \geq \bar{\beta}$ for some history h^t , so either firms are completely deterred in every future period, or full deterrence lasts finitely many periods. In the first case, the regulator strictly gains by switching to “wait” from period t onwards (thus saving on inspection costs). In the second, there is some history h^s , with $s \geq t$, such that $\beta(h^s) \geq \bar{\beta}$ but $\beta(h^s, 0) < \beta(h^s)$. At history h^s , the firm does not offend. Also, since the regulator switches in equilibrium to wait with some probability at history $(h^s, 0)$, we must have

$$V_I(h^s) = -(1 - \delta)i + \delta V_W(h^s, 0).$$

Nevertheless, if the regulator switches to wait at time s and history h^s , it obtains $\delta V_W(h^s, 0)$. By ceasing inspection early, the regulator saves $(1 - \delta)i$, so this is a profitable deviation. This is a contradiction, so necessarily $\beta(h^t) < \bar{\beta}$. We can conclude that the regulator never has strict incentives to switch to inspect, which implies $V_W(h^t) = 0$ at any public history h^t . So now, again for the sake of contradiction, assume $\beta(h^t) \leq \underline{\beta}$ for some history h^t . Then, the firm offends with probability one, so if the regulator waited in the previous period and switches to inspect at time t it obtains at least $y - S - (1 - \delta)i$. This is positive by Assumption 1, so the regulator strictly prefers to inspect, a contradiction.

The above implies that the regulator never has a strict incentive to switch from wait to inspect (otherwise, we would have $\beta_t = 1$ at some t), nor a strict incentive to switch from inspect to wait (otherwise, we would have $\beta_t = 0$). Hence, for all histories h^t , $V_I(h^t) - V_W(h^t)$ lies between 0 and S . It takes value 0 in periods in which the regulator switches to wait, and takes value S in periods in which the regulator switches to inspect.

Next, notice that some switching must then occur in equilibrium. Following a conviction, firms must believe that the regulator inspected in the previous period with probability one. Hence, the regulator must switch from inspect to wait immediately following a conviction. Suppose that a conviction occurs at history h^t (i.e., suppose $h_t = 1$). Then, at any time s , from date $t+1$ onwards, if the regulator is known not to change its action, and absent any conviction, the probability of inspection conditional on the public history evolves according to Bayes' rule. In particular,

$$\frac{\beta(h^t, 0^{s-t})}{1 - \beta(h^t, 0^s)} = \frac{\beta(h^t, 0^{s-t-1})}{1 - \beta(h^t, 0^{s-t-1})} (1 - \alpha(\beta(h^t, 0^{s-t-1}))). \quad (2)$$

This implies that the probability of inspection, as perceived by firms, gradually declines over time. Hence, Lemma 1 implies there must eventually come a time when the regulator switches from wait to inspect. This pattern turns out to define equilibrium play, which we summarize next.

Equilibrium can be understood as consisting of two main phases: a “stationary phase” in which the probability of offending remains at a baseline level, and a “residual deterrence phase” which follows a conviction and during which the probability of offending is reduced relative to the baseline in the stationary phase.

Proposition 1 *An equilibrium exists. Furthermore, there exist unique values $\beta^{\min}, \beta^{\max} \in (\underline{\beta}, \bar{\beta})$ with $\beta^{\min} < \beta^{\max}$, and $T \in \mathbb{N}$, such that the following holds in any equilibrium:*

Step 1. (Stationary phase) If the regulator inspects in the previous period, then it keeps inspecting, and if, instead, the regulator waits in the previous period, then it switches to inspect so that the perceived probability of inspection remains equal to β^{\min} . If there is a conviction, the play moves to Step 2; and otherwise, it stays in Step 1.

Step 2. (Residual deterrence, following a conviction) In the period following a conviction, the regulator switches with probability $1 - \beta^{\max}$ to wait. As long as there is no conviction, in the subsequent $T - 1$ periods, the regulator does not switch, so the

perceived probability of inspection decreases over time until it reaches β^{\min} . If there is no conviction, play moves to Step 1; and otherwise, it reinitializes at Step 2.

While Proposition 1 characterizes the two main phases of equilibrium, there is also an initial phase. This is the first period, in which the regulator switches from “wait” to “inspect” with probability β^{\min} . Play then continues in the stationary phase if there is no conviction, and proceeds to the residual deterrence phase after the first conviction. It is worth noting that the values β^{\min} and β^{\max} , and hence the stochastic process for inspections and offenses, as viewed by the outside observer who has public information alone, are uniquely determined in equilibrium. The only reason for multiple equilibria is that the regulator’s decision to inspect may depend on payoff-irrelevant components of its private history; i.e., the decisions to inspect prior to the previous period.

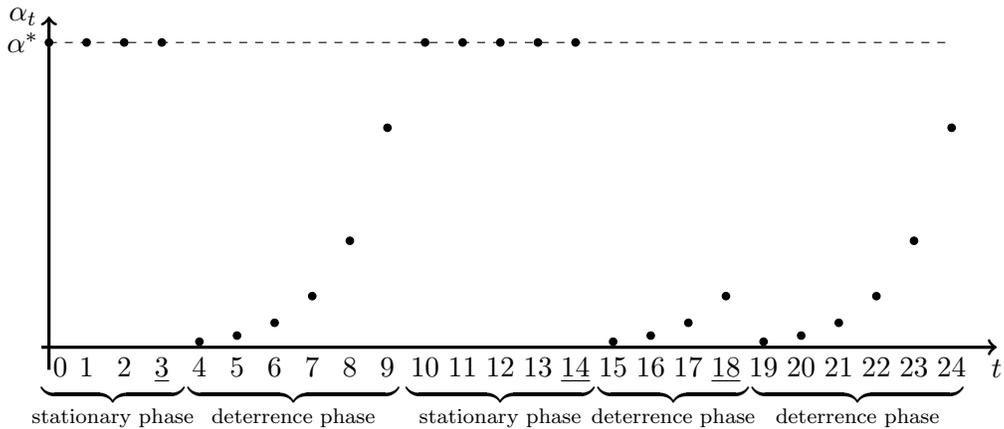


Figure 1: Example of dynamics of α_t . In the graph, the deterrence length is $T = 6$, and there are conviction in periods 3, 14 and 18. Each stationary phase ends with a conviction, while deterrence phases either end after 6 periods or are re-initialized by a conviction.

The equilibrium pattern fits well the examples described in the Introduction: “residual deterrence” follows a conviction, as there are several periods of low offending. Firms’ beliefs about whether the regulator is inspecting gradually deteriorate in the

absence of a further conviction, and this corresponds to what Sherman (1990) terms “deterrence decay”. These patterns of offending are illustrated in Figure 1.

The regulator’s incentives to switch are illustrated in Figure 2. First recall that, as we argued before, the regulator’s expected continuation payoff, having played wait in the previous period, is $V_W(h^t) = 0$ at any public history h^t . The incentive for the regulator to change its action can then be understood by studying $V_I(h^t)$, the continuation payoff from being in the “inspect” state at the beginning of period t . During the stationary phase, the regulator is willing to commence inspecting, if not already, and this requires $V_I(h^t) = S$. Conversely, immediately following a conviction, i.e. if $h_{t-1} = 1$, we must have $V_I(h^t) = 0$ so the regulator is willing to cease inspections. The deterrence phase can be understood as several periods of low offending such that the regulator’s payoff from inspecting passes from zero immediately following a conviction to S at the beginning of the next stationary phase (assuming that phase is reached without a further conviction). In particular, these periods of low offending reduce the regulator’s payoff from continuing to inspect immediately following a conviction.

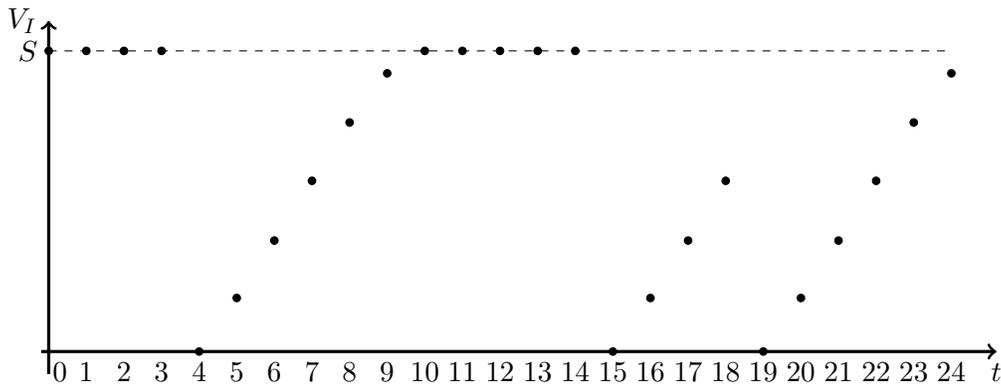


Figure 2: Example of dynamics of V_I as a function of time. In the graph, there is conviction in periods 3, 14 and 18.

3.1 Some comparative statics

With our central result in hand, let us consider some determinants of equilibrium offending. First, let $\alpha^* \equiv \alpha(\beta^{\min})$ give the baseline offense rate, i.e., the rate of offending during the stationary phase. This offending is such that, during the stationary phase, the regulator is indifferent between commencing inspections and not; that is,

$$\underbrace{S}_{\text{Cost of commencing inspection}} = \underbrace{-i(1-\delta) + \alpha^*y + (1-\alpha^*)\delta S}_{\text{Net benefit of commencing inspection}}. \quad (3)$$

This equation balances the cost of commencing inspection, S , against the benefit from doing so net of the inspection cost $i(1-\delta)$. The benefit of inspecting includes the expected payoff from a conviction in the first period of inspection α^*y , and the discounted continuation payoff in case inspection yields no conviction, which equals $(1-\alpha^*)\delta S$. The latter follows because, if h^t is a history at which the regulator is in the stationary phase, then $V_I(h^t, 0) = S$, since in the absence of a conviction the regulator remains in the stationary phase and hence willing to commence inspecting at $(h^t, 0)$ if it did not do so at h^t .

Equation (3) yields the following result:

Corollary 1 *The probability of offending in the stationary phase satisfies*

$$\alpha^* = \frac{(1-\delta)(i+S)}{y-\delta S} \quad (4)$$

and hence is increasing in S .

The comparative statics in Corollary 1 can be easily understood from (3). While increasing S increases both the cost of commencing inspection and the continuation payoff absent a conviction, the former effect dominates the regulator's payoff, so α^* must rise to keep the regulator indifferent to commencing inspection.

While increasing the initiation cost S increases the stationary-phase offense rate α^* , it also increases the length of residual deterrence, as explained in the following result.

Corollary 2 *The length of the residual-deterrence phase T is increasing in the cost of initiating inspections S .*

The reason for this result is simple. As the cost of initiating inspections grows so does the regulator’s continuation payoff $V_I(h^t)$ at any history h^t in the stationary phase; indeed, we have $V_I(h^t) = S$. Hence, the number of periods of low offending needed to ensure the regulator’s continuation payoff immediately after a conviction equals zero grows with the initiation cost S .

Since an increase in the initiation cost S both increases the length of residual deterrence and the probability of offending in the stationary phase, the overall impact on offending in the long-run is not immediately clear. One way to proceed is to define the “long-run average offense rate” $\bar{\alpha}$ by finding the ex-ante expected average rate of offending over the first τ periods, and then take the limit as $\tau \rightarrow \infty$; i.e.,

$$\bar{\alpha} = \lim_{\tau \rightarrow \infty} \frac{1}{\tau} \mathbb{E} \left[\sum_{s=0}^{\tau-1} \alpha(h^s) \right]. \quad (5)$$

We then obtain the following result when δ is close to one, which one may interpret as the regulator’s interactions with firms being sufficiently frequent.

Corollary 3 *If δ is large enough, then the long-run average offense rate, $\bar{\alpha}$, increases continuously in the initiation cost S .*

The result in Corollary 3 could be interpreted in different ways. On the one hand, it suggests that policy makers may want to promote regulatory flexibility, to reduce the cost of initiating new inspections. This might be achieved by defining broad and flexible organizational goals (allowing the regulator the flexibility to go after new offenses), by building enforcement authorities with the flexibility to hire new personnel, or by making long-term investments in organizational capabilities (such as maintaining a research team which can turn quickly to new topics or offenses). Alternatively, to the extent that initiation costs are difficult to evaluate directly, we predict that those enforcement authorities perceiving the most extensive residual deterrence may

themselves be the *least* nimble (i.e., they may have higher initiation costs S) and hence have the highest *long-run average* offense rates.

4 Alternative explanations for residual deterrence

We now discuss the two broad alternative explanations for residual deterrence that can be discerned from the literature; cognitive biases of offenders, and offender learning about an exogenous inspection technology.

Recency bias. One possible view is that residual deterrence results from individuals assigning too much weight to recently observed regulatory activities, such as convictions against rival firms. This idea is advanced, for instance, by Chen (2016) in his study of executions of deserters from the British military in World War I (see also Hertwig et al., 2004, and Kahneman, 2011, for a discussion of these biases in more general contexts). This explanation requires deviations in probability assessments by offenders from the true probability of being apprehended and potentially explains both residual deterrence and deterrence decay. For instance, Chen finds that executions of deserters do appear to modestly reduce desertions by English soldiers, although this response does not appear to have a rational motivation (there is no evidence that a deserter is more likely to be executed following a recent desertion; if anything, the relationship is the opposite). Similarly, Kastlunger et al. (2009) find that individuals who are audited earlier in a laboratory audit game exhibit higher compliance than those audited later on; the authors suggest possible psychological explanations.

The key difference between theories of recency bias and our proposed explanation for residual deterrence is that ours is based on the rational behavior of offenders. In other words, the probability of monitoring in the periods following a conviction (as perceived by firms whose only information is the public information regarding convictions) *is* in fact higher in the equilibrium of our model. This is not without empirical implications, even for a researcher whose only information is the public information regarding convictions. In particular, note that the probability of a conviction at date

t when the probability of inspection is β_t is given by

$$\beta_t \alpha(\beta_t) = \beta_t \left(1 - F \left(\frac{\beta_t}{1-\beta_t} \gamma \right) \right)$$

which may be increasing or decreasing in β_t . Depending on the value of γ and the function F , subsequent convictions thus become either *more* or *less* likely in the periods immediately following a conviction. A formal theory of recency bias might instead posit a constant true inspection probability, while a conviction leads firms only to *perceive* a higher probability. In turn, firms are deterred and so convictions necessarily become *less* likely immediately after a conviction.⁹ Hence, a researcher can potentially distinguish the two theories by examining the (temporal) correlation structure of convictions.

To give an empirical example, consider the Georgetown area “crackdown” by police on minor offenses such as parking violations and street crimes, as documented by Sherman (1990). The crackdown resulted in much higher levels of arrests and fines over a six month period, together with a modest reduction in crime (for instance, a reduction in reported robberies and self-reports of offending). The increase in arrests are inconsistent with deterrence being driven purely by cognitive or psychological biases. This is true even though interviews with residents suggest they did believe the likelihood of being caught for various offenses had risen, and that this belief persisted *at least one month after* the police crackdown had ended. Our model explains such an evolution of beliefs by positing fully-rational agents and start-up costs for crackdowns or inspections, rather than cognitive biases. In our view, the (ex-post) incorrect belief that the crackdown remained in force is a simple consequence of imperfect information.

Exogenously determined inspections. A number of other papers suggest, either formally or informally, that either (i) the inspection technology remains fixed over time and cannot be adjusted, or (ii) the inspection technology can change or

⁹Another possibility with the same implication is that the enforcer reacts strategically to the potential offender’s bias, and hence reduces inspections during a period of deterrence.

“break”, but it does so exogenously. In the first class of explanations, Block, Nold and Sidak (1981) write (footnote 23) that residual deterrence can be explained by

“assuming that colluders use Bayesian methods to estimate the probability that they will be apprehended in a particular period. In this formulation, whenever colluders are apprehended, colluders estimate of their probability of apprehension increases, and that increase is dramatic if their a priori distribution is diffuse and has small mean.”

Banerjee et al. (2014) also propose a dynamic model in which offenders learn about the policing policy through their actions, but where such a policing policy is taken as given. While these models can render “residual deterrence” (in Banerjee et al., it may be that drivers are deterred from a specific location, although not from drink driving altogether)¹⁰, they do not provide an explicit theory for “deterrence decay”. In particular, if the monitoring technology does not change, then potential offenders may potentially learn if it is in place (and where), and hence avoid being caught at all future dates. This is clearly different from our theory, where offending and convictions re-emerge after enough time.

A second possible class of explanations is to remedy the absence of deterrence decay in the first by positing that the inspection technology changes with time. For instance, one could generate patterns of residual deterrence and deterrence decay by letting the probability of inspection follow some exogenous (say Markov) stochastic process. After enough time, potential offenders would be willing to experiment to “test” the state of the technology; convictions would manifest in a heightened belief that monitoring is intense, and hence phases of residual deterrence. One difficulty with this view is that it leaves unexplained why monitoring changes and how fast. Naturally, the answer to

¹⁰Note that firms relocating their offending to another unmonitored location is also a possible interpretation of “deterrence” in our model. Under this view, deterrence at a given location may not provide deterrence overall, so the social welfare implications of monitoring may be moot.

the latter determines the answer to important questions, such as those regarding the duration of the residual deterrence.

Relatedly, Halac and Prat (2016) suppose that the inspector (the “manager” in their model) can invest in the monitoring technology, but that it breaks exogenously. While the length of deterrence in their “bad news” model *is* endogenously determined, it also depends on the rate at which the monitoring technology exogenously decays. For instance, if decay occurs ever more slowly, then the length of the “deterrence phase” grows without bound. In our model, the regulator instead has full control over whether it inspects or not in each period, and the length of deterrence depends on the costs of inspecting and commencing inspections (as well as the distribution over firm preferences F).

5 Disclosure of inspections

An important question in enforcement relates to the optimality of disclosing the authority’s activities. Several papers, notably Lando and Shavell (2004), Lazear (2006) and Eeckhout, Persico and Todd (2010), have explored the optimal disclosure of “groups” to be targeted by the enforcement authority. In those settings, disclosing that certain groups are more likely to be monitored can be an effective way to lower the offense rate; in particular, a uniformly low probability of monitoring across all groups can be an ineffective deterrent. Here, in contrast, we will argue that offending is often diminished by disclosing less information about the regulator’s activities rather than more.

In general, analyzing information disclosure is challenging in our framework because of the apparent need to characterize equilibria of our game for all possible information disclosure policies. While our working paper version, Dilmé and Garrett (2016), shows how to characterize equilibria for richer information structures, we have not attempted an exhaustive treatment. Here, we will simply contrast the case where the regulator’s activities are fully disclosed to the case where only convictions are disclosed.

Suppose then that the regulator's inspection choices are publicly disclosed. In this case, a public history h^t is an element of $(\{I, W\} \times \{0, 1\})^t$, and describes both inspection choices and convictions. Given a public history h^t , let $b_{t-1}(h^t) \in \{I, W\}$ denote the inspection decision made at date $t - 1$. We find the following.

Proposition 2 *There is a unique equilibrium of the game in which the regulator's inspections are publicly disclosed. In such an equilibrium,*

$$\alpha(h^t) = \alpha_{b_{t-1}(h^t)} \equiv \begin{cases} \frac{(1-\delta)i}{y} & \text{if } b_{t-1}(h^t) = I, \\ \frac{(1-\delta)i+S}{y} & \text{if } b_{t-1}(h^t) = W. \end{cases}$$

The proof of Proposition 2 shows that the regulator is indifferent at any date t between continuing its action from date $t - 1$ and changing this action. As such, the regulator's continuation payoff at any date must be equal to zero (indeed, the regulator is willing to play "wait" indefinitely). For this reason, the regulator's incentives at any date are the same as in the one-shot game discussed in Remark 1. Balancing the regulator's incentives to wait or inspect in each period then pins down the equilibrium probability of offending.

We can now compute the long-run average offense rate in the current setting with full disclosure. This is given by an average of α_I and α_W , each of them weighted by the likelihood of the corresponding actions of the inspector:¹¹

$$\bar{\alpha}^{\text{fd}} = \frac{\alpha^{-1}(\alpha_W)}{1 - \alpha^{-1}(\alpha_I) + \alpha^{-1}(\alpha_W)} \alpha_I + \frac{1 - \alpha^{-1}(\alpha_I)}{1 - \alpha^{-1}(\alpha_I) + \alpha^{-1}(\alpha_W)} \alpha_W .$$

We can compare the full-disclosure long-run offense rate $\bar{\alpha}^{\text{fd}}$ with the long-run offense rate of our base model.

Proposition 3 *There exists $\bar{\delta} < 1$ such that if $\delta > \bar{\delta}$ then $\bar{\alpha}^{\text{fd}} > \bar{\alpha}$.*

¹¹Here, $\frac{\alpha^{-1}(\alpha_W)}{1 - \alpha^{-1}(\alpha_I) + \alpha^{-1}(\alpha_W)}$ is the unique stationary probability of the regulator inspecting, while $\frac{1 - \alpha^{-1}(\alpha_I)}{1 - \alpha^{-1}(\alpha_I) + \alpha^{-1}(\alpha_W)}$ is the unique stationary probability of the regulator waiting, given that inspections follow a first-order Markov process with transitions determined by $\alpha^{-1}(\cdot)$.

When δ is close to one, i.e. when interactions between the regulator and firms are sufficiently frequent, long-run average offending is higher with disclosure of inspections than when the only public information concerns convictions. Intuitively, when decisions to inspect are public, the probability of offending following the regulator’s decision to “wait” must be high enough to incentivize the regulator to pay the initiation cost S , notwithstanding that the continuation payoff following “inspect” is zero. Conversely, when only convictions are observable, the offense rate in the stationary phase, α^* , is smaller. This is because, when the regulator switches to inspect, the regulator either obtains a conviction or enjoys a positive continuation payoff. This positive continuation payoff derives from the fact the regulator’s inspection is not detected absent a conviction (so the regulator continues to inspect until the next conviction without affecting equilibrium offending).

6 General Payoffs

So far we studied a regulator motivated directly by its concern for apprehending offenses. As noted in the Introduction, such an assumption seems reasonable in settings where regulatory officials are motivated by implicit rewards or career concerns. More generally, however, the regulator may have preferences for deterring offenses. To examine this possibility, we consider a more general payoff structure as follows.

		firm	
		$a_t = N$	$a_t = O$
regulator	$b_t = W$	$0, 0$	$-L, \pi_t$
	$b_t = I$	$-(1 - \delta)i - S \mathbf{1}(b_{t-1}, b_t), 0$	$y - L - (1 - \delta)i - S \mathbf{1}(b_{t-1}, b_t), -\gamma$

Here L is the regulator’s loss as a result of an offense, while $y \geq 0$ is its reward for apprehending an offense (when $L = 0$, the model is identical to that in Section

3). Continue to assume that the costs i and S , as well as the penalty γ , are strictly positive, and fix F to be the continuous distribution described above. The following result describes a sufficient condition under which an equilibrium as in Proposition 1 exists.

Proposition 4 *Fix $L > 0$. In the model with general payoffs, there is an equilibrium which permits the same characterization as in Proposition 1 if either (i) Assumption 1 holds, or (ii) δ is sufficiently close to one.*

Proposition 4 establishes that equilibrium dynamics with residual deterrence can be found also when the regulator cares about deterring offenses. A sufficient condition is that δ is sufficiently close to one, which guarantees the regulator stands to gain enough from future deterrence to justify inspections.

One case of particular interest is where the regulator is motivated by deterrence alone; that is, a regulator whose payoff is lowered by $L > 0$ whenever there is an offense, but for which $y = 0$. In this case, the unique equilibrium of the stage game without switching costs is (W, O) ; that is, the presence of the regulator does not deter crime. Still, with repeated interactions and positive switching costs, if L is positive and δ is large enough, there exist equilibria where the regulator inspects only to obtain residual deterrence following a conviction. It is worth noting that, here, positive switching costs are essential for any deterrence to occur in equilibrium, as the next result shows.

Corollary 4 *If $y = 0$ and $L > 0$, then the unique equilibrium in the repeated game with $S = 0$ is (W, O) after all histories.*

Conversely, when switching costs are positive, there can be a multiplicity of equilibria. As we saw, when δ is large enough and $S > 0$, there is an equilibrium with residual deterrence. However, there is also another equilibrium in which the firms always offend with probability one, and the regulator never inspects (and where, following any conviction, the regulator switches to “wait” with probability one). This suggests the existence of still further, non-stationary, equilibria, where after some histories there is

a “crime wave” and firms offend with probability one at all future dates. We make no attempt to characterize the multiplicity of equilibria.

7 Conclusions

We have studied a dynamic version of the inspection game in which an inspector (a regulator, police, or other enforcement official) incurs a resource cost to commencing inspections. We showed that this cost gives rise to “reputational” effects: following a conviction, offending is reduced for several periods before resuming at a steady level. This effect is present whether the inspector is motivated by obtaining convictions (as in our baseline model of Section 3) or “socially motivated” in the sense that it values deterrence itself (as in Section 6).

Our model sheds light on how frictions from reallocation of resources shape the incentives of both firms and the regulator in a market. Its tractability and uniqueness of its predictions allows for interesting comparative statics results. While the length of the deterrence phase increases when reallocating resources is more costly, the subsequent offense rate is also higher. We showed that, when the regulator is patient, revealing previous activities tends to increase offending, since a higher rate of offending is needed for the regulator to be willing to commence inspections.

Admittedly, the framework described in this paper addresses a much simplified setting relative to many applications seen in practice. A key simplification was our modeling of the authority’s decision as a binary choice — whether or not to inspect. In many settings, an authority might be expected to rotate its inspections around different possible targets, incurring costs to commence inspection at a new target (say a particular location, or a particular set of offenses). The authority then effectively faces an *endogenous outside option* in deciding whether to focus on any given target; the value of this option is the payoff obtained by focusing inspections on the next best target. At the same time, potential offenders face choices concerning not only whether to offend, but potentially where to offend or what kind of offense to commit. Such

decisions depend, in turn, on perceptions of where the authority's attention is focused at the relevant moment. While we aimed in this paper at exhibiting residual deterrence in the simplest possible setting, these additional issues call for a richer model. We expect this will be the subject of future work.

References

- [1] Aubert, Cecile, Patrick Rey and William Kovacic (2006), 'The Impact of Leniency and Whistle-blowing Programs on Cartels,' *International Journal of Industrial Organization*, 24, 1241-1266.
- [2] Banerjee, Abhijit V., Esther Duflo, Daniel Keniston, and Nina Singh, 'The Efficient Deployment of Police Resources: Theory and New Evidence from a Randomized Drunk Driving Crackdown in India,' mimeo MIT, Yale and Rajasthan Police.
- [3] Becker, Gary S. (1968), 'Crime and Punishment: An Economic Approach,' *Journal of Political Economy*, 76, 169-217.
- [4] Bénabou, Roland and Jean Tirole (2012), 'Laws and Norms,' Discussion Paper series 6290, Institute for the Study of Labor (IZA), Bonn.
- [5] Benson, Bruce L., Iljoong Kim and David W. Rasmussen (1994), 'Estimating Deterrence Effects: A Public Choice Perspective on the Economics of Crime Literature,' *Southern Economic Journal*, 61, 161-168.
- [6] Block, Michael K., Frederick K. Nold and Joseph G. Sidak (1981), 'The Deterrent Effect of Antitrust Enforcement,' *Journal of Political Economy*, 89, 429-445.
- [7] Board, Simon and Moritz Meyer-ter-Vehn (2013), 'Reputation for Quality,' *Econometrica*, 81, 2381-2462.

- [8] Bond, Philip and Kathleen Hagerty (2010), ‘Preventing Crime Waves,’ *American Economic Journal: Microeconomics*, 2, 138-159.
- [9] Brown, Stephen V., Xiaoli Tian and Jennifer W. Tucker (2014), ‘The Spillover Effect of SEC Comment Letters on Qualitative Corporate Disclosure: Evidence from the Risk Factor Disclosure,’ mimeo Arizona State U., Ohio State U., and U. of Florida.
- [10] Caruana, Guillermo and Liran Einav (2008), ‘A Theory of Endogenous Commitment,’ *Review of Economic Studies*, 75, 99-116.
- [11] Chen, Daniel (2016), ‘The Deterrent Effect of the Death Penalty? Evidence from British Commutations During World War I,’ mimeo Toulouse School of Economics.
- [12] DeMarzo, Peter, Michael Fishman and Kathleen Hagerty (1998), ‘The Optimal Enforcement of Insider Trading Regulations,’ *Journal of Political Economy*, 106, 602-632.
- [13] Dilmé, Francesc (2014), ‘Reputation Building through Costly Adjustment,’ mimeo University of Pennsylvania and University of Bonn.
- [14] Eeckhout, Jan, Nicola Persico and Petra E. Todd (2010), ‘A Theory of Optimal Random Crackdowns,’ *American Economic Review*, 100, 1104-1135.
- [15] *Global Competition Review* (2015), ‘Rating Enforcement – The Annual Ranking of the World’s Leading Competition Authorities’.
- [16] Halac, Marina and Andrea Prat (2016), ‘Managerial Attention and Worker Performance,’ *American Economic Review*, 106, 3104-3132.
- [17] Hertwig, Ralph, Greg Barron, Elke U. Weber and Ido Erev (2004), ‘Decisions from Experience and the Effect of Rare Events in Risky Choice,’ *Psychological Science*, 15, 534-539.

- [18] Iossa, Elisabetta and Patrick Rey (2014), ‘Building Reputation for Contract Renewal: Implications for Performance Dynamics and Contract Duration,’ *Journal of the European Economic Association*, 12, 549-574.
- [19] Jennings, Jared, Simi Kedia and Shivaram Rajgopal (2011), ‘The Deterrence Effects of SEC Enforcement and Class Action Litigation,’ mimeo University of Washington, Rutgers University and Emory University.
- [20] Kahneman Daniel (2011), *Thinking, Fast and Slow*. New York, NY: Macmillan.
- [21] Kastlunger, Barbara, Erich Kirchler, Luigi Mittone and Julia Pitters (2009), ‘Sequences of audits, tax compliance, and taxpaying strategies,’ *Journal of Economic Psychology*, 30, 405-418.
- [22] Khalil, Fahad (1997), ‘Auditing without Commitment,’ *RAND Journal of Economics*, 28, 629-640.
- [23] Klemperer (1995), ‘Competition when Consumers have Switching Costs: An Overview with Applications to Industrial Organization, Macroeconomics, and International Trade,’ *Review of Economic Studies*, 62, 515-539.
- [24] Kovacic, William E. and David A. Hyman (2016), ‘Consume or Invest: What Do/Should Agency Leaders Maximize,’ *Washington Law Review*, 91, 295-322.
- [25] Lazear, Edward P (2006), ‘Speeding, Terrorism, and Teaching to the Test,’ *Quarterly Journal of Economics*, 121, 1029-1061.
- [26] Lipman, Barton and Ruqu Wang (2000), ‘Switching Costs in Frequently Repeated Games,’ *Journal of Economic Theory*, 93, 149-190.
- [27] Lipman, Barton and Ruqu Wang (2009), ‘Switching Costs in Infinitely Repeated Games,’ *Games and Economic Behavior*, 66, 292-314.
- [28] Mookherjee, Dilip and IPL Png (1989), ‘Optimal Auditing, Insurance and Redistribution,’ *Quarterly Journal of Economics*, 104, 399-415.

- [29] Mookherjee, Dilip and IPL Png (1994), 'Marginal Deterrence in Enforcement of Law,' *Journal of Political Economy*, 76, 1039-1066.
- [30] Nagin, Daniel S. (2013), 'Deterrence in the Twenty-First Century,' *Crime and Justice*, 42, 199-263.
- [31] Nosenzo, Daniele, Theo Offerman, Martin Sefton and Ailko van der Veen (2014), 'Encouraging Compliance: Bonuses versus Fines in Inspection Games,' *Journal of Law, Economics and Organization*, 30, 623-648.
- [32] Polinsky, Mitchell A. and Steven Shavell (1984), 'The Optimal Use of Fines in Imprisonment,' *Journal of Public Economics*, 24, 89-99.
- [33] Reinganum, Jennifer and Louis Wilde (1985), 'Income Tax Compliance in a Principal-Agent Framework,' *Journal of Public Economics*, 26, 1-18.
- [34] Reinganum, Jennifer and Louis Wilde (1986), 'Equilibrium Verification and Reporting Policies in a Model of Tax Compliance,' *International Economic Review*, 27, 739-760.
- [35] Schectman, Joel (2014, January 22), 'Q&A David Vladeck, Former Director of FTC Consumer Unit,' *The Wall Street Journal*, <http://blogs.wsj.com/riskandcompliance/2014/01/22/qa-david-vladeck-former-director-of-ftc-consumer-unit/>.
- [36] Schenk, Kristy M. (2012), 'The Deterrence Effects of SEC Enforcement Actions,' dissertation at Pennsylvania State U..
- [37] Shavell, Steven (1992), 'A Note on Marginal Deterrence,' *International Review of Law and Economics*, 12, 343-355.
- [38] Sherman, Lawrence (1990), 'Police Crackdowns: Initial and Residual Deterrence,' *Crime and Justice*, 12, 1-48.

- [39] Shimshack, Jay P. and Michael B. Ward (2005), ‘Regulator reputation, enforcement, and environmental compliance,’ *Journal of Environmental Economics and Management*, 50, 519-540.
- [40] Spagnolo, Giancarlo (2005), ‘*Divide et Impera*: Optimal Leniency Programs,’ mimeo Stockholm School of Economics.
- [41] Strausz, Roland (1997), ‘Delegation of Monitoring in a Principal-Agent Relationship,’ *Review of Economic Studies*, 64, 337-357.
- [42] Taylor, Bruce, Christopher S. Koper and Daniel J. Woods, ‘A randomized controlled trial of different policing strategies at hot spots of violent crime,’ *Journal of Experimental Criminology*, 7, 149-181.
- [43] United States Department of Justice Civil Rights Division (2015), ‘Investigation of the Ferguson Police Department,’ https://www.justice.gov/sites/default/files/opa/press-releases/attachments/2015/03/04/ferguson_police_department_report.pdf.

Appendix A: Proofs of the Main Results

Proof of Lemma 1. Fix, for the rest of the proof, an equilibrium. We show first that $\beta(h^t) < \bar{\beta}$ at all h^t . Assume for a contradiction that there is a public history h^t with $\beta(h^t) \geq \bar{\beta}$, so $\alpha(h^t) = 0$. Let $\tau > 0$ be the first time such that, $\beta(h^t, 0^\tau) < \beta(h^t)$, that is, the first time the regulator switches to wait with a positive probability at history $(h^t, 0^\tau)$. Note that, necessarily, τ is finite since, otherwise, the continuation value at history h^t of inspecting forever would be $\sum_{\tau=0}^{\infty} -(1-\delta)\delta^\tau i = -i$, so the regulator would strictly prefer to switch to wait at time t and ensure a continuation payoff of 0, which would imply $\beta(h^t) = 0$. Note also that, switching to wait at history $(h^t, 0^{\tau-1})$ (instead of at history $(h^t, 0^\tau)$) gives the regulator a continuation value of $0 + \delta V_W(h^t, 0^\tau)$, while the payoff of inspecting one more period and switching at $(h^t, 0^\tau)$ is $-(1-\delta)i + \delta V_W(h^t, 0^\tau) < \delta V_W(h^t, 0^\tau)$. So, the regulator has a strict incentive to switch to wait at time $\tau - 1$, which leads to a contradiction.

We now prove that $V_W(h^t) = 0$ for all h^t . Since $V_W(h^t) \geq 0$ for all h^t (since waiting forever gives a continuation payoff equal to 0) assume, for the sake of contradiction, that $\bar{V}_W > 0$, where \bar{V}_W is the supremum of $V_W(h^t)$ among all public histories h^t . Consider a history h^t where $V_W(h^t) > \delta \bar{V}_W$. At this history, if the regulator waited in period $t - 1$, it has a strict incentive to switch to inspect, since waiting implies a utility of no higher than $\delta \bar{V}_W$ and, by assumption, $V_W(h^t) > \delta \bar{V}_W$, so $V_W(h^t) = V_I(h^t) - S$. Also, if the regulator inspected in period $t - 1$, it prefers to keep inspecting, since switching to wait gives it a continuation payoff of at most $\delta \bar{V}_W$, while inspecting gives $V_I(h^t) = V_W(h^t) + S > \delta \bar{V}_W + S > \delta \bar{V}_W$. So, since (independently of its previous action) the regulator has a strict incentive to inspect at history h^t , we have that $\beta(h^t) = 1$, which contradicts the first part of the proof.

Finally, we show that $\beta(h^t) > \underline{\beta}$ at all h^t . Assume for a contradiction that $\beta(h^t) \leq \underline{\beta}$ for some history h^t , so $\alpha(h^t) = 1$. If the regulator did not inspect at time $t - 1$, then the value of keeping waiting is 0. If, instead, it switches to inspect, it obtains

$-S - (1 - \delta)i + y + \delta V_I(h^t, 1)$. Since the regulator can switch to wait in period $t + 1$, we have that $V_I(h^t, 1) \geq 0$, so by Assumption 1 the regulator has a strict incentive to switch to inspect. Conversely, if the regulator inspected in period $t - 1$, its value for keeping inspecting is $-(1 - \delta)i + y$, which is higher than the value of switching to wait, which is 0. This implies a contradiction. ■

Proof of Proposition 1. Fix, for this proof, an equilibrium, assuming it exists. We will determine its properties and finally establish its existence.

Recall that in the proof of Lemma 1 we show that $V_W(h^t) = 0$ for all h^t , so $V_I(h^t) \in [0, S]$ for all h^t . In particular, if there was a conviction at time $t - 1$, since the regulator has to be weakly willing to switch to wait, we have that $V_I(h^t) = 0$. We use this observation to prove the following.

Lemma 2 *For all histories h^t , $\alpha(h^t) \leq \alpha^* \equiv \frac{(1-\delta)(i+S)}{y-\delta S}$. Also, the regulator is willing to switch to inspect at history h^t if and only if $\alpha(h^t, 0^s) = \alpha^*$ for all $s \geq 0$.*

Proof. We begin by proving the second part of the claim. Fix a history h^t and assume that the regulator is willing to switch to inspect at time t , so $V_I(h^t) = S$. Then

$$S = V_I(h^t) = -(1 - \delta)i + \alpha(h^t)(y - \delta 0) + (1 - \alpha(h^t))\delta V_I(h^t, 0).$$

Since, by Lemma 1 we have $V_I(h^t, 0) \leq S$, the above expression implies that $\alpha(h^t) \geq \alpha^*$, or, alternatively, $\beta(h^t) \leq \alpha^{-1}(\alpha^*)$ (recall that $\alpha(\cdot)$ is defined in equation (1)). Assume that $V_I(h^t, 0) < S$, and let $t + \tau$ be the first time after t such that, if there is no conviction from t to $t + \tau$, the regulator switches to inspect with positive probability at time $t + \tau + 1$ (note that it exists since, otherwise, the absence of convictions would drive $\beta(h^t, 0^\tau)$ below $\underline{\beta}$, for some $\tau \in \mathbb{N}$, which would contradict Lemma 1). Then, since between $t + 1$ and $t + \tau$ the regulator does not switch to inspect (if there is no conviction), the lack of convictions makes firms increasingly convinced that the regulator is not inspecting, so the perceived probability of apprehension decreases over time and $\alpha(h^t) < \alpha(h^t, 0^\tau)$. Nevertheless, we have that

$$V_I(h^t, 0^\tau) = -(1 - \delta)i + \alpha(h^t, 0^\tau)(y - \delta 0) + (1 - \alpha(h^t, 0^\tau))\delta S.$$

Hence, $V_I(h^t, 0^\tau) > S$, contradicting Lemma 1, so necessarily $V_I(h^t, 0) = S$ and $\alpha(h^t) = \alpha^*$. The same argument above implies $\alpha(h^t, 0) = \alpha^*$, and the argument can be iteratively use to prove that $\alpha(h^t, 0^\tau) = \alpha^*$ for all $\tau \geq 0$. The converse implication, that is, the fact that if $\alpha(h^t, 0^\tau) = \alpha^*$ for all $\tau \geq 0$ then the regulator is willing to switch to inspect at history h^t is easily obtained using analogous arguments.

Now consider the first part of the claim. For the sake of contradiction, assume that there exists a history h^t such that $\alpha(h^t) > \alpha^*$. By the previous part of the proof, the regulator is strictly willing to keep waiting at time t . Let $t + \tau + 1 > t$ be the first time where the regulator is willing to switch to inspect (which exists by the argument above). Notice that $\alpha(h^t, 0^s) > \alpha(h^t, 0^{s-1})$ for all $s = t + 1, \dots, t + \tau$, so $\alpha(h^t, 0^\tau) > \alpha^*$. Therefore,

$$\begin{aligned} V_I(h^t, 0^\tau) &= -(1 - \delta)i + \alpha(h^t, 0^\tau)(y - \delta 0) + (1 - \alpha(h^t, 0^\tau))\delta S \\ &> -(1 - \delta)i + \alpha^*(y - \delta 0) + (1 - \alpha^*)\delta S = S. \end{aligned}$$

Hence there is a strict incentive to switch to inspect, contradicting Lemma 1. ■

For each $\hat{\beta}_0 \in (\underline{\beta}, \bar{\beta})$ define the following sequence for $t \geq 0$:

$$\hat{\beta}_{t+1}(\hat{\beta}_0) = \frac{\left(1 - \alpha(\hat{\beta}_t(\hat{\beta}_0))\right) \hat{\beta}_t(\hat{\beta}_0)}{\left(1 - \alpha(\hat{\beta}_t(\hat{\beta}_0))\right) \hat{\beta}_t(\hat{\beta}_0) + 1 - \hat{\beta}_t(\hat{\beta}_0)}. \quad (6)$$

This is a decreasing sequence that reproduces the evolution of beliefs β_t when there are no convictions. In particular, we interpret $\hat{\beta}_0$ as a guess for the perceived probability of apprehension in the period immediately after a conviction, and our goal will be to find one that satisfies all equilibrium conditions. Notice that, since $\hat{\beta}_0 < \bar{\beta}$, then there exists some $T(\hat{\beta}_0)$ such that $\hat{\beta}_{T(\hat{\beta}_0)-1}(\hat{\beta}_0) > \alpha^{-1}(\alpha^*) \geq \hat{\beta}_{T(\hat{\beta}_0)}(\hat{\beta}_0)$. The value $T(\hat{\beta}_0)$ is interpreted as the length of the deterrence phase, that is, the time it takes for the regulator to be willing to switch to inspect and, as a result, by Lemma 2, the stationary phase to begin. The value $\alpha^{-1}(\alpha^*)$ will be the lowest level of the perceived probability of apprehension β^{\min} in the proposition.

For a fixed $\hat{\beta}_0$, define the sequence $(V_{I,s}(\hat{\beta}_0))_s$ from 0 to $T(\hat{\beta}_0)$ backwards by setting $V_{I,T(\hat{\beta}_0)}(\hat{\beta}_0) = S$ and, for all $s = 0, \dots, T(\hat{\beta}_0) - 1$,

$$V_{I,s}(\hat{\beta}_0) = -(1 - \delta)i + \alpha(\hat{\beta}_s(\hat{\beta}_0))y + (1 - \alpha(\hat{\beta}_s(\hat{\beta}_0)))\delta V_{I,s+1}(\hat{\beta}_0). \quad (7)$$

Here, $V_{I,s}(\hat{\beta}_0)$ is interpreted as the continuation value of the regulator in the period $s + 1$ after a conviction, if it inspected in the period s after a conviction. We can then see that there exists a unique $\beta^{\max} \in (\underline{\beta}, \bar{\beta})$ such that $V_{I,0}(\beta^{\max}) = 0$; that is, such that after a conviction the regulator is indifferent between switching to wait or not (which, by Lemma 1, is required to be the case in equilibrium). This follows from noticing that (i) if $\hat{\beta}_0 \rightarrow \bar{\beta}$, then $T(\hat{\beta}_0) \rightarrow \infty$ and $\hat{\beta}_t(\hat{\beta}_0) \rightarrow \bar{\beta}$ for all $t \geq 0$, so also $V_{I,0}(\hat{\beta}_0) \rightarrow -i$, while (ii) if $\hat{\beta}_0 = \alpha^{-1}(\alpha^*)$ then $T(\hat{\beta}_0) = 0$ and $V_{I,0}(\hat{\beta}_0) = S$, and (iii) the following lemma holds.

Lemma 3 $V_{I,0}(\cdot)$ is continuous and strictly decreasing in $(\alpha^{-1}(\alpha^*), \bar{\beta})$.

Proof. We can rewrite Equation (6) as follows:

$$\frac{\hat{\beta}_{t+1}(\hat{\beta}_0)}{1 - \hat{\beta}_{t+1}(\hat{\beta}_0)} = \frac{\hat{\beta}_t(\hat{\beta}_0)}{1 - \hat{\beta}_t(\hat{\beta}_0)} (1 - \alpha(\hat{\beta}_t(\hat{\beta}_0))).$$

Since $\alpha(\cdot)$ is strictly decreasing, it is clear that a higher value for $\hat{\beta}_t(\hat{\beta}_0)$ implies a higher value for $\hat{\beta}_{t+1}(\hat{\beta}_0)$, for all $t < T(\hat{\beta}_0)$. This implies that $T(\cdot)$ is both left-continuous and increasing.

We now show that $V_{I,0}(\cdot)$ is continuous. It is clear that if $T(\cdot)$ is continuous (and therefore locally constant) at $\hat{\beta}_0$ then $V_{I,0}(\cdot)$ is continuous at $\hat{\beta}_0$, since both $\hat{\beta}_t(\hat{\beta}_0)$ and $\alpha(\hat{\beta}_t(\hat{\beta}_0))$ are continuous in $\hat{\beta}_0$. Assume then that $\hat{\beta}_0$ is such that $T(\cdot)$ is not continuous at $\hat{\beta}_0$, so $T(\hat{\beta}_0) + 1 = \lim_{\beta \searrow \hat{\beta}_0} T(\beta)$. This implies that $\hat{\beta}_{T(\hat{\beta}_0)}(\hat{\beta}_0) = \alpha^{-1}(\alpha^*)$, and that $\lim_{\beta \searrow \hat{\beta}_0} \beta_{T(\hat{\beta}_0)}(\beta) = \alpha^{-1}(\alpha^*)$. So, we have that

$$V_{I,T(\hat{\beta}_0)}(\beta) = -(1 - \delta)i + \alpha(\hat{\beta}_{T(\hat{\beta}_0)}(\beta))y + (1 - \alpha(\hat{\beta}_{T(\hat{\beta}_0)}(\beta)))\delta S \rightarrow_{\beta \searrow \hat{\beta}_0} S,$$

where we use that

$$\alpha(\hat{\beta}_{T(\hat{\beta}_0)}(\beta)) \rightarrow_{\beta \searrow \hat{\beta}_0} \alpha^*.$$

As a result, $V_{I,0}(\cdot)$ is continuous at $\hat{\beta}_0$.

We finally show that $V_{I,0}(\cdot)$ is strictly decreasing. Notice first that, for a fixed $\hat{\beta}_0$, $V_{I,t}(\hat{\beta}_0)$ is an increasing sequence. Indeed, the right hand side of Equation (7) is both increasing in $V_{I,t+1}(\hat{\beta}_0)$ and $\alpha(\hat{\beta}_t(\hat{\beta}_0))$ when $V_{I,t+1}(\hat{\beta}_0) < y/\delta$. As a result, $V_{I,T(\hat{\beta}_0)-1}(\hat{\beta}_0) < S < y/\delta$ and, iteratively, since the sequence $\alpha(\hat{\beta}_t(\hat{\beta}_0))$ is increasing, $V_{I,t}(\hat{\beta}_0) < V_{I,t+1}(\hat{\beta}_0)$ for all $t = 0, \dots, T(\hat{\beta}_0) - 1$. Furthermore, if $T(\cdot)$ is continuous (and locally constant) at $\hat{\beta}_0$, it is also clear that $V_{I,0}(\hat{\beta}_0)$ is strictly decreasing in $\hat{\beta}_0$, given that $\alpha(\hat{\beta}_t(\hat{\beta}_0))$ is decreasing in $\hat{\beta}_0$ for all t . Assume then that $\hat{\beta}_0$ is such that $T(\hat{\beta}_0)$ is not continuous. Since $V_{I,0}(\cdot)$ is continuous over $[\hat{\beta}_0 - \varepsilon, \hat{\beta}_0 + \varepsilon]$ and strictly decreasing over $(\hat{\beta}_0 - \varepsilon, \hat{\beta}_0) \cup (\hat{\beta}_0, \hat{\beta}_0 + \varepsilon)$ for a small $\varepsilon > 0$, it is also strictly decreasing at $\hat{\beta}_0$. ■ ■

Proof of Corollary 1. Proven in the proof of Proposition 1. ■

Proof of Corollary 2. For each fixed $\hat{\beta}_0 \in (\beta^{\min}, \bar{\beta})$, let $T(\hat{\beta}_0, S)$ be defined as in the proof of Proposition 1, now making explicit its dependence on S through the fact that $\alpha^{-1}(\alpha^*)$ is a decreasing function of S (since $\alpha^{-1}(\cdot)$ is decreasing and α^* given in (4) is increasing in S). Consider first the generic case where $T(\cdot, S)$ is continuous (and therefore locally constant) at $\hat{\beta}_0$. Then, a small change in S does not change $T(\hat{\beta}_0, S)$ (since $\alpha^{-1}(\alpha^*)$ is continuous in S), and so

$$\frac{\partial V_{I,0}(\hat{\beta}_0, S)}{\partial S} > 0 ,$$

where $V_{I,0}(\hat{\beta}_0, S)$ corresponds to $V_{I,0}(\hat{\beta}_0)$ in the proof of Proposition 1, again making explicit its dependence on S . As a result, if the equilibrium perceived probability of inspection after a conviction is $\hat{\beta}_0$, a small increase in S requires that $\hat{\beta}_0$ slightly increase, maintaining $V_{I,0}(\hat{\beta}_0, S) = 0$, and therefore the length of the deterrence phase does not change.

Now assume that $T(\hat{\beta}_0, S)$ is not continuous at $\hat{\beta}_0$, so $T(\hat{\beta}_0, S) + 1 = \lim_{\beta \searrow \hat{\beta}_0} T(\beta, S)$ and $\hat{\beta}_{T(\hat{\beta}_0, S)} = \beta^{\min}$. In this case, fixing $\hat{\beta}_0$, a small increase in S increases α^* and so decreases $\alpha^{-1}(\alpha^*)$, and therefore $\lim_{S' \searrow S} T(\hat{\beta}_0, S') = T(\hat{\beta}_0, S) + 1$. As a result, if $\varepsilon > 0$

is small enough

$$\begin{aligned}
V_{I,0}(\hat{\beta}_0, S+\varepsilon) &= V_{I,0}(\hat{\beta}_0, S) \\
&\quad - \delta^{T(\hat{\beta}_0, S)} S + \delta^{T(\hat{\beta}_0, S)} (-(1-\delta)i + \alpha^* y + (1-\alpha^*)\delta(S+\varepsilon)) \\
&= V_{I,0}(\hat{\beta}_0, S) + \delta^{T(\hat{\beta}_0, S)+1} \varepsilon .
\end{aligned}$$

This implies, again, if the equilibrium perceived probability of apprehension after a conviction is $\hat{\beta}_0$, a small increase in S necessitates a slight increase in $\hat{\beta}_0$, in order that $V_{I,0}(\hat{\beta}_0, S) = 0$. Hence, the length of the deterrence phase increases by one. ■

Proof of Corollary 3. Let $T(\delta)$ denote the equilibrium length of the deterrence phase (now making the dependence of T on δ explicit). Let $A(\delta)$ be the probability of an offense in the $T(\delta)$ periods since a conviction.

Lemma 4 As $\delta \rightarrow 1$, $(1-\delta)T(\delta) \rightarrow \log(1+S/i)$ and $A(\delta) \rightarrow 0$.

Proof. Using a similar notation as in the proof of Proposition 1, let $\hat{\beta}_t$ denote the equilibrium perceived probability of apprehension $t+1$ periods after a conviction. Note that, as $\delta \rightarrow 1$ we have $\alpha^* \rightarrow 0$ (by Equation (4)), and so $\beta^{\min} = \alpha^{-1}(\alpha^*) \rightarrow \bar{\beta}$. Since α^* is the highest probability of an offense in any period, it is easily verified from Bayes' rule that $A(\delta) \rightarrow 0$ when $\delta \rightarrow 1$. Then, notice that

$$\begin{aligned}
& - \sum_{t=0}^{T(\delta)-1} \delta^t (1-\delta)i + \delta^{T(\delta)} S \\
& \leq 0 \\
& \leq -(1-\delta)i + A(\delta)y + (1-A(\delta))\delta \left(- \sum_{t=0}^{T(\delta)-2} \delta^t (1-\delta)i + \delta^{T(\delta)-1} S \right) .
\end{aligned}$$

The left-hand side of the first inequality is a lower bound on the regulator's expected payoff following a conviction. The inequality holds because the regulator prefers obtaining a conviction over continuing to inspect. The right-hand side of the second inequality is an upper bound on the regulator's expected payoff following a conviction.

The inequality holds because the regulator prefers convictions early rather than later. When $\delta \rightarrow 1$, the two bounds converge to zero. As a result, we have

$$\lim_{\delta \rightarrow 1} (-(1 - \delta^{T(\delta)})i + \delta^{T(\delta)}S) = 0 \quad \Rightarrow \quad \lim_{\delta \rightarrow 1} \delta^{T(\delta)} = \frac{i}{i + S} .$$

This implies that $\lim_{\delta \rightarrow 1} T(\delta) (1 - \delta) = \log(1 + S/i)$. ■

Now, let us turn our focus on the stationary phase. Let $\bar{T}(\delta)$ be the expected time until the first conviction during the stationary phase. The following result obtains its limit value:

Lemma 5 As $\delta \rightarrow 1$, $(1 - \delta)\bar{T}(\delta) \rightarrow \frac{1}{\beta} \frac{y-S}{i+S}$.

Proof. The expected time until the first conviction is given by

$$\bar{T}(\delta) = \sum_{t=0}^{\infty} t \alpha^{-1}(\alpha^*) \alpha^* (1 - \alpha^{-1}(\alpha^*) \alpha^*)^t = \frac{1 - \alpha^{-1}(\alpha^*) \alpha^*}{\alpha^{-1}(\alpha^*) \alpha^*} .$$

Using the expression for α^* in (4), $\lim_{\delta \rightarrow 1} \alpha^* = 0$ and hence $\lim_{\delta \rightarrow 1} \alpha^{-1}(\alpha^*) = \bar{\beta}$, and we have that $\lim_{\delta \rightarrow 1} (1 - \delta)\bar{T}(\delta) = \frac{y-S}{\bar{\beta}(i+S)}$. ■

Let's finally compute $\bar{\alpha}$ for δ close to one. To do this, notice (using Lemma 4) that

$$\begin{aligned} \bar{\alpha} &= \frac{\bar{T}(\delta) \alpha^*}{\bar{T}(\delta) + T(\delta)} + o(1 - \delta) \\ &= \frac{1 - \delta}{\frac{y-S}{i+S} + \bar{\beta} \log(\frac{i+S}{S})} + o(1 - \delta) \end{aligned}$$

where $o(1 - \delta)$ represents terms such that $o(1 - \delta) / (1 - \delta)$ approaches zero as $\delta \rightarrow 1$.

We can now take the first derivative with respect to S and obtain

$$\frac{\partial \bar{\alpha}}{\partial S} = (1 - \delta) \frac{y + i + \bar{\beta}i + \bar{\beta} \frac{i^2}{S}}{(i + S)^2 \left(\frac{y-S}{i+S} + \bar{\beta} \log(\frac{i+S}{S}) \right)^2} + o(1 - \delta),$$

which is positive if δ is close enough to one. ■

Proof of Proposition 2. Fix an equilibrium, and let $V(h^t)$ denote the expected discounted continuation payoff of the regulator at date t following public history h^t (note that $V(h^t) \geq 0$ at every history h^t , since the inspector can choose to play “wait”

forever). The result follows from three lemmas that characterize such an equilibrium and establish that it is unique. The first obtains that, like in our baseline model (see the proof of Lemma 1), the continuation value of the regulator when it waited in the previous period is equal to zero.

Lemma 6 *For all h^t such that $b_{t-1}(h^t) = W$, $V(h^t) = 0$.*

Proof. Assume, for the sake of contradiction, that the statement of the lemma does not hold. Fix an equilibrium and let $\bar{V}_W = \sup_{\{h^t \in H | b_{t-1}(h^t) = W\}} V(h^t)$, and assume $\bar{V}_W > 0$. Then, there is a history h^t with $b_{t-1}(h^t) = W$ and such that $V(h^t) > \delta \bar{V}_W$. If the regulator weakly prefers keeping waiting at history h^t then we have

$$\delta \bar{V}_W < V(h^t) = \delta V(h^t, (W, 0)) \leq \delta \bar{V}_W ,$$

a contradiction. This implies that the regulator switches to inspect (for sure) in period t . Let $t' > t$ be the first time where the regulator has the incentive to switch back to wait. Then, we have

$$\delta \bar{V}_W < V(h^t) = -S - (1 - \delta^{t'-t})i + \delta^{t'-t}V(h^{t'}) < \delta \bar{V}_W .$$

This, again, is a contradiction. ■

The second lemma establishes that, if the regulator inspected in the previous period, its current continuation value is 0. Notice that, in our baseline model, this was only true when it was common knowledge that the regulator inspected the previous period, that is, after a conviction.

Lemma 7 *For all h^t such that $b_{t-1}(h^t) = I$, $V(h^t) = 0$.*

Proof. Take a history h^t such that $b_{t-1}(h^t) = I$ and assume that $V(h^t) > 0$. Let $t' > t$ be the first time where the inspector is willing to switch to wait. We have

$$0 < V(h^t) = -(1 - \delta^{t'-t})i + \delta^{t'-t}0 < 0 ,$$

which is a contradiction. ■

Finally, the next lemma shows that the regulator is indifferent between choosing W and I in each period. This is different from our baseline model where, during the deterrence phase, the regulator has strict preferences in equilibrium (i.e., to continue with the same action).

Lemma 8 *At all histories h^t , the regulator is indifferent between choosing W and I .*

Proof. Take a history h^t with $b(h^t) = W$. Choosing $b_t = W$ gives a continuation payoff of 0. Alternatively, choosing $b_t = I$ gives

$$-S - (1 - \delta)i + \alpha_t(h^t)y + \delta 0 .$$

By Lemma 6 this expression cannot be strictly positive. Assume, for the sake of contradiction, that the previous expression is strictly negative. In this case, $\beta_t(h^t) = 0$, and therefore $\alpha_t(h^t) = 1$. By Assumption 1 the previous expression is strictly positive when $\alpha_t(h^t) = 1$, which leads to a contradiction.

Take a history h^t with $b(h^t) = I$. Choosing $b_t = W$ gives a continuation payoff of 0. Alternatively, choosing $b_t = I$ gives

$$-(1 - \delta)i + \alpha_t(h^t)y + \delta 0 . \tag{8}$$

By Lemma 7, the expression (8) cannot be larger than 0. If it is less than 0, then $\beta_t(h^t) = 0$, and therefore $\alpha_t(h^t) = 1$, implying, by Assumption 1, that in fact (8) is strictly positive, a contradiction. ■

The previous results together with arguments in the main text show the offense probabilities are given according to the statement of the proposition. ■

Proof of Proposition 3. In our base model, as δ approaches one, α^* vanishes (see Equation (4)), and hence $\bar{\alpha}$ vanishes as well. Nevertheless, $\bar{\alpha}^{\text{fd}}$ does not vanish as δ approaches α^* . Indeed, notice that, as $\delta \rightarrow 1$, while $\alpha_I \rightarrow 0$, we have that $\alpha_W \rightarrow \frac{S}{y} \in (0, 1)$. As a result,

$$\lim_{\delta \rightarrow 1} \bar{\alpha}^{\text{fd}} = \frac{1 - \bar{\beta}}{1 - \bar{\beta} + \alpha^{-1}\left(\frac{S}{y}\right) y} \frac{S}{y} > 0 .$$

Therefore, if δ is high enough, $\bar{\alpha}^{\text{fd}} > \bar{\alpha}$. ■

Proof of Proposition 4. We posit a continuation value \bar{V}_I for the regulator following a conviction, and show that a value exists which coincides with an equilibrium sharing the description of equilibrium in Proposition 1.

Let α^* denote the probability of offending in the stationary phase of our putative equilibrium, as in our base model. The regulator's continuation value in the stationary phase is $\frac{\alpha^*L}{1-\delta}$ if it waited the previous period, since it is (weakly) willing to never switch to inspect. The indifference condition in the stationary phase imposes that, if it inspected the previous period, its continuation value is $\frac{\alpha^*L}{1-\delta} + S$. So, α^* solves the following equation

$$-\frac{\alpha^*L}{1-\delta} + S = -(1-\delta)i + \alpha^*(-L + y + \delta\bar{V}_I) + (1-\alpha^*)\delta\left(-\frac{\alpha^*L}{1-\delta} + S\right).$$

The solution is then given by

$$\alpha^*(\bar{V}_I) = \frac{(1-\delta)\left(\sqrt{(y-\delta S + \delta\bar{V}_I)^2 + 4\delta L(i+S)} - y - \delta\bar{V}_I + \delta S\right)}{2\delta L}.$$

It is easy to show (differentiating the previous expression) that $\alpha^*(\bar{V}_I)$ is decreasing in \bar{V}_I , it is always positive and that $\alpha^*(\bar{V}_I) = 1$ for $\bar{V}_I^\dagger \equiv -\frac{L}{1-\delta} - \frac{y-(1-\delta)i-S}{\delta}$. So, as long as $\bar{V}_I \geq \bar{V}_I^\dagger$, $\alpha(\bar{V}_I) \in (0, 1]$.

We want to apply a similar argument as the one in the second part of Proposition 1, that is, obtain a perceived probability of inspection $\hat{\beta}_0$ such that the regulator is indifferent to switching to wait after a conviction. However, we cannot mimic the argument because now $\alpha^*(\bar{V}_I)$ is not known before solving for \bar{V}_I (in our base model α^* is given by equation (4), and $\bar{V}_I = 0$).

We therefore fix $\bar{V}_I \in [\bar{V}_I^\dagger, \infty)$ and define, for each $\hat{\beta}_0 \in (\underline{\beta}, \bar{\beta})$, $\hat{\beta}_t(\hat{\beta}_0)$ in the same way as in the proof of Proposition 1. Now, we use $T(\hat{\beta}_0; \bar{V}_I)$ to denote the time satisfying

$$\hat{\beta}_{T(\hat{\beta}_0; \bar{V}_I)-1}(\hat{\beta}_0) > \beta(\alpha^*(\bar{V}_I)) \geq \hat{\beta}_{T(\hat{\beta}_0; \bar{V}_I)}(\hat{\beta}_0).$$

We can now define $(V_{I,t}(\hat{\beta}_0; \bar{V}_I))_{t=0}^{T(\hat{\beta}_0; \bar{V}_I)}$ analogously to the proof of Proposition 1 and, additionally, we can now define the sequence $(V_{W,t})_{t=0}^{T(\hat{\beta}_0; \bar{V}_I)}$ backward defining, for all

$t = 0, \dots, T(\hat{\beta}_0; \bar{V}_I),$

$$V_{I,t}(\hat{\beta}_0; \bar{V}_I) = -(1-\delta)i + \alpha(\hat{\beta}_t(\hat{\beta}_0))(y - L + \delta\bar{V}_I) + (1 - \alpha(\hat{\beta}_t(\hat{\beta}_0)))\delta V_{I,t+1}(\hat{\beta}_0; \bar{V}_I)$$

$$\text{with } V_{I,T}(\hat{\beta}_0; \bar{V}_I) = -\frac{\alpha^*(\bar{V}_I)L}{1-\delta} + S,$$

$$V_{W,t}(\hat{\beta}_0; \bar{V}_I) = -\alpha(\hat{\beta}_t(\hat{\beta}_0))L + \delta V_{W,t+1}(\hat{\beta}_0; \bar{V}_I)$$

$$\text{with } V_{W,T}(\hat{\beta}_0; \bar{V}_I) = -\frac{\alpha^*(\bar{V}_I)L}{1-\delta} .$$

Notice that $\lim_{\hat{\beta}_0 \nearrow \bar{\beta}} [V_{I,0}(\hat{\beta}_0; \bar{V}_I) - V_{W,0}(\hat{\beta}_0; \bar{V}_I)] = -i < 0$, and if $\hat{\beta}_0 = \alpha^{-1}(\alpha^*(\bar{V}_I))$ then $V_{I,0}(\hat{\beta}_0; \bar{V}_I) - V_{W,0}(\hat{\beta}_0; \bar{V}_I) = S$. Then, using the same arguments as in the proof of Proposition 1, we have that there exists some $\beta^{\max}(\bar{V}_I)$ such that $V_{I,0}(\beta^{\max}(\bar{V}_I); \bar{V}_I) = V_{W,0}(\beta^{\max}(\bar{V}_I); \bar{V}_I)$. It is then only left to show that there exists some $\bar{V}_I^* \geq \bar{V}_I^\dagger$ such that $V_{I,0}(\beta^{\max}(\bar{V}_I^*); \bar{V}_I^*) = \bar{V}_I^*$.

Assume first that Assumption 1 holds. Since $\lim_{\bar{V}_I \rightarrow \infty} \alpha^*(\bar{V}_I) = 0$, we have $\lim_{\bar{V}_I \rightarrow \infty} V_{I,0}(\beta^{\max}(\bar{V}_I); \bar{V}_I) = 0$. Analogously, $\lim_{\bar{V}_I \searrow \bar{V}_I^\dagger} V_{I,0}(\beta^{\max}(\bar{V}_I); \bar{V}_I) > \bar{V}_I^\dagger$ since we have that, for each \bar{V}_I ,

$$V_{I,0}(\beta^{\max}(\bar{V}_I), \bar{V}_I) = V_{W,0}(\beta^{\max}(\bar{V}_I), \bar{V}_I) \geq \frac{-L}{1-\delta} .$$

Then, if Assumption 1 holds, there exists some $\bar{V}_I^* \in (\bar{V}_I^\dagger, \infty)$ such that $V_{I,0}(\beta^{\max}(\bar{V}_I^*); \bar{V}_I^*) = \bar{V}_I^*$.

To prove the result when Assumption 1 does not hold we take δ sufficiently large. Notice that, as $\delta \rightarrow 1$, $\alpha^* \rightarrow 0$. The same arguments as in the proof of Lemma 4 ensure that, as $\delta \rightarrow 1$, the probability of an offense during the deterrence phase decreases to 0. Now, for a given guess \bar{V}_I , we have that

$$V_{I,0}(\beta^{\max}(\bar{V}_I); \bar{V}_I) = -(1-\delta^{T(\delta; \bar{V}_I)})i + \delta^{T(\delta; \bar{V}_I)} \left(-\frac{\alpha^*(\bar{V}_I)L}{1-\delta} + S \right) + o(1), \text{ and} \quad (9)$$

$$V_{W,0}(\beta^{\max}(\bar{V}_I); \bar{V}_I) = \delta^{T(\delta; \bar{V}_I)} \left(-\frac{\alpha^*(\bar{V}_I)L}{1-\delta} \right) + o(1) ,$$

where, with some abuse of notation, $T(\delta; \bar{V}_I)$ denotes the length of the deterrence phase for $\hat{\beta}_0 = \beta^{\max}(\bar{V}_I)$ under the guess \bar{V}_I as found above, making explicit its dependence on δ . The difference is given by

$$V_{I,0}(\beta^{\max}(\bar{V}_I); \bar{V}_I) - V_{W,0}(\beta^{\max}(\bar{V}_I); \bar{V}_I) = -(1-\delta^{T(\delta; \bar{V}_I)})i + \delta^{T(\delta; \bar{V}_I)}S + o(1) .$$

As in the proof of Lemma 4, we find that, as $\delta \rightarrow 1$, $(1 - \delta)T(\delta; \bar{V}_I) \rightarrow \log(1 + S/i)$. It is then easy to see (using (9)) that a value of \bar{V}_I^* exists satisfying $V_{I,0}(\beta^{\max}(\bar{V}_I^*); \bar{V}_I^*) = \bar{V}_I^*$; indeed, \bar{V}_I^* satisfies

$$\bar{V}_I^* = -\frac{\sqrt{(y - S)^2 + 4LS} - y + S}{2S/i} + o(1) .$$

■

Proof of Corollary 4. For the sake of contradiction, assume that the continuation value of the regulator is above $\frac{-L}{1-\delta}$ at some history, and let $\bar{V} \equiv \sup_{h^t} V(h^t) > \frac{-L}{1-\delta}$. Fix a history h^t such that $V(h^t) > \delta\bar{V}$. If at such a history there is offending at date t with probability zero, then the regulator benefits from not inspecting, which implies a contradiction. Therefore, the regulator should be willing to wait, and therefore we have

$$\delta\bar{V} < V(h^t) = -\alpha(h^t)L + \delta V(h^t, 0) \leq \delta\bar{V} .$$

This is, again, a contradiction. ■