

Residual Deterrence*

Francesc Dilmé

University of Bonn

fdilme@uni-bonn.de

Daniel F. Garrett

Toulouse School of Economics

dfgarrett@gmail.com

December 2015

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. Our theory rests on the costs of reallocating enforcement resources. In light of these costs, we study the determinants of offending such as the role of public information about enforcement and offending.

JEL classification: C73, K42

Keywords: deterrence, reputation, switching costs

*We would like to thank Bruno Biais, Daniel Chen, Nuh Aygün Dalkıran, Jan Eeckhout, Marina Halac, George Mailath, Moritz Meyer-ter-Vehn, Alessandro Pavan, Patrick Rey, Stephan Lauer mann and seminar participants at Bilkent University, the CSIO/IDEI Joint Workshop on Industrial Organization at Toulouse School of Economics, and at 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, for instance, Nagin (2013) for a review of the evidence).

One factor that is salient in determining the perceived risk of punishment is past enforcement decisions, especially the extent of past convictions. Block, Nold and Sidak (1981) provide evidence that bread manufacturers lower mark-ups in response to Department of Justice price fixing prosecutions in their region. Jennings, Kedia and Rajgopal (2011) provide evidence that past SEC enforcements among peer firms deter “aggressive” financial reporting. Sherman (1990) reviews anecdotal evidence that isolated police “crackdowns”, especially on drink driving, lead to reductions in offending that extend past the end of the period of intensive enforcement.¹ All these instances appear to fit a pattern which Sherman terms “residual deterrence”. Residual deterrence occurs when reductions in offending follow a period of active enforcement.

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 inspecting, the theory must explain why the regulator continues to monitor firms even when they are unlikely to offend. Our explanation hinges on the regulator’s costs of allocating and reallocating resources to inspection activities, that is, on switching costs. We show how such costs can manifest in the episodes of residual deterrence that follow the apprehension of an offending firm.

¹In further examples, Shimshack and Ward (2005) find that a fine against a paper mill for non-compliance with water polluting standards deters non-compliance of firms in the same state in the following year. Chen (2015) finds some limited evidence that recent executions deterred English desertions in World War I (although executions of Irish soldiers *spurred* Irish desertions).

There are several reasons to expect switching costs in practice. It is costly for a regulator to understand a given market, costs which are, however, mitigated by recent experience monitoring that market (hence, monitoring a given market for two consecutive periods should be less costly than for two non-consecutive ones). Switching an enforcement agency's focus may require personnel changes, or technology changes,² which can be costly both to instigate and to reverse. On a shorter time scale, switching activities requires coordination of various personnel and this coordination imposes administrative costs.³

We study a dynamic version of a simple workhorse model – the inspection game. In this model, a long-lived regulator faces a sequence of short-lived firms. Committing an offense is only worthwhile for a firm if the regulator is “inactive”, while being active is only worthwhile for the regulator if an offense is committed. In our baseline model (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. This corresponds to a view that the most salient action an enforcement agency 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 model 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 cost of reallocating resources. We show that equilibrium then features reputational effects driven by the switching costs: a conviction is followed by several periods during which the firms do not offend. We identify this pattern as residual deterrence. Thus, equilibrium in our model involves reputation cycles, with each cycle characterized by a conviction, a subsequent reduction or total cessation of offending, and finally a resumption of offending at a steady level. We also show that the switching costs are necessary to generate residual deterrence, since the episodes of residual deterrence disappear as switching costs shrink.

The model permits a rich analysis of comparative statics. This is facilitated by the fact that the equilibrium process for convictions, as well as firms' strategies, are uniquely determined. The only source of multiple equilibria is that the regulator may condition its switching on privately-held and payoff-irrelevant past information.

We illustrate how the model may be used to evaluate the effect of the penalty incurred

²Consider, for instance, a police force switching its focus from speed infractions to drink driving violations.

³There may also be “psychological costs” of changing the current pattern of activity (see, e.g., Klemperer, 1995), or reputational concerns that frequent switching might suggest a “lack of focus” or lack of confidence about the appropriate allocation of regulatory resources.

by a convicted firm on firms' offending. The direction of this effect need not be obvious and can depend on the horizon of interest. While in the short run, the direction of the effect is as expected, it can be reversed in the long run. In particular, a permanent increase in the penalty leads to a higher long-run average rate of offending. This perhaps counterintuitive result is a consequence of the equilibrium behavior of the regulator. As in the static inspection game, the regulator responds to an increase in penalties by inspecting less often (in order that firms are still willing to offend). However, this means fewer convictions, and hence fewer reputation cycles and fewer episodes of residual deterrence. While this finding could potentially explain the difficulty in discerning a deterrence effect from increased penalties in empirical work (see Nagin (2013) for a review), we also point out how our prediction is sensitive to details of the model.

Further comparative statics are possible with respect to the costs of switching to and from inspection. The duration of residual deterrence is increasing in these switching costs, while the rate of offending after residual deterrence has ceased is also increasing in these costs. The long-run average rate of offending is in turn determined jointly by the length of deterrence and rate of offending after deterrence. We find that the long-run average offense rate is always increasing in the cost of switching *to* inspection, whereas it may increase or decrease in the cost of switching *from* inspection. Hence, it can be either higher or lower than in the model without switching costs.

Before presenting the model, it is worth clarifying up front a few important modeling choices. First, our baseline model posits that firms have identical preferences for offending. This simplifies the analysis and leads to the stark conclusion that the firms' beliefs as to the probability of inspection are constant in every period and equal to the equilibrium inspection probability in the static inspection game. This observation highlights that episodes of residual deterrence in our model need not involve fluctuations in firms' beliefs. However, we generalize (in Section 4.1) by allowing for ex-ante identical firms to have heterogeneous (and privately observed) preferences for offending. In this case, firms' beliefs do fluctuate in a predictable fashion over the cycle: the perceived probability of detection following a conviction is high, while this probability falls in the absence of convictions in accordance with Bayes' rule. That is, in the absence of convictions, firms place increasing weight on the possibility that the regulator has switched to being inactive. Our extended model is thus able to explain time-varying perceptions of the likelihood of detection (the intuition is nonetheless close to our baseline model, which provides a useful starting point).

A second important modeling choice is that the only public information about past play

is the history of convictions. In practice, additional information may be available about the offending of firms and the inspection activities of the regulator. We extend the model to consider both possibilities separately (Section 4.2 studies the case where additional noisy signals of firm offending are available, while Section 4.3 allows for noisy signals of the regulator’s inspecting). Allowing for noisy signals of firm behavior can bring our model closer to the applications mentioned above. When there is a possibility of price fixing, for instance, high prices might be a noisy signal of firm offending. We would then predict that firms set lower prices following a conviction, simply because they are less likely to fix prices. In this sense, the model directly accounts for the findings of Block, Nold and Sidak (1981). Moreover, our findings are an important robustness check; residual deterrence persists in the settings with richer information, although the implications for equilibrium offense and switching probabilities must of course be accounted for.

Another important modeling choice is the regulator’s concern for 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. For instance, 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 explicit incentives for law enforcers to catch offending are often controversial or illegal, even quite explicit incentives seem relatively common.⁴ Nonetheless, we are also able to extend our baseline model to settings where the regulator is concerned directly with deterrence, rather than convictions (see Section 4.4). Again, we exhibit equilibria featuring reputation cycles. The regulator in this case may be incentivized to inspect precisely because it anticipates residual deterrence following a conviction.

The rest of the paper is as follows. We next briefly review literature on the economic theory of deterrence, as well as on reputations. Section 2 introduces the baseline model, Section 3 solves for the equilibrium and provides comparative statics, and Section 4 provides extensions. Section 5 concludes. The Appendix contains the proofs of results in Section 3, and of Proposition 2 in Section 4.1, while Online Supplementary Material proves all other results.

⁴Perhaps the best-known example in recent times is the Ferguson Police Department’s focus on generating revenue by writing tickets; see, for instance, the Department of Justice Civil Rights Division 2015 report ‘Investigation of the Ferguson Police Department’. 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.

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, 1994), Lazear (2006), Bassetto and Phelan (2008), 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. Block, Nold and Sidak (1981) do informally suggest a dynamic theory. They view enforcement officials as committed to playing a fixed inspection policy over time, with potential offenders updating their beliefs about this policy based on enforcement actions against peer firms.⁵⁶ Relative to Block, Nold and Sidak, our theory allows the regulator to choose its enforcement policy strategically over time.

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, more recently (and independently of our own work) Halac and Prat

⁵They suggest (footnote 23) “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.”

⁶A similar view is taken by Chen (2015), who suggests that recent executions could temporarily change the perceived likelihood of being executed for desertion. He argues such effects would be temporary due to recency biases in decision making. Executions could also affect other parameters in the decision problems of would-be deserters (possible parameters are presented formally by Bénabou and Tirole (2012)).

(2014). 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, but the stage game is different to Dilmé’s, requiring a separate analysis.⁷

A key novelty of our setting relative to the various papers on seller reputation is that the generation of public information depends on the actions of all players, both the regulator and firms. This feature is in common with Halac and Prat (2014), who analyze the deterioration of manager-worker relationships. In their model, a manager can invest in a monitoring technology, which breaks down stochastically. 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 the monitoring technology is unlikely to be broken. Given our focus on modeling the behavior of enforcement agencies, our analysis deviates from theirs in several directions. First, our focus is on a regulator whose payoff is a function of the convictions, not the offense rate. Second, if the crime rate is low, our regulator has the incentive (and the ability) to stop inspecting due to its (opportunity) cost, and it does so in equilibrium. Third, we consider the implications of permitting a heterogeneous population of firms and a rich signal structure for the public information about inspections (see Sections 4.1-4.3). Finally, we analyze the comparative statics with respect to the short and long-run rates of offending (Section 3.1), which may be of interest for policy.

2 Baseline Model

Timing, players and actions. Time is discrete and infinite. There is a single regulator and a sequence of short-lived firms who are potential offenders, one per period.

In each period $t \geq 0$, the regulator chooses an action $b_t \in \{I, W\}$, where I denotes “inspect” and W denotes “wait”. The history of such decisions is denoted $b^t = (b_0, \dots, b_{t-1}) \in \{I, W\}^t$. For each period t , the firm simultaneously chooses an action $a_t \in \{O, N\}$ where O denotes “offend” and N denotes “does not offend”. 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

⁷Other papers featuring repeated interactions and switching costs are Lipman and Wang (2000, 2009) and Caruana and Einav (2008). One key difference with respect to our paper is their focus on settings with complete information.

t .⁸

Payoffs. Per-period payoffs are determined at first instance 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 $\pi > 0$. 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 $i > 0$. It incurs no cost if waiting. In the event of obtaining a conviction, the regulator earns an additional payoff of $y > i$. Later, we consider the possibility that the regulator cares about deterring the firm rather than obtaining convictions.⁹

In addition to the costs and benefits specified above, the regulator sustains a cost if switching action at period t . Hence, the switching cost in period t is $S_{b_{t-1}b_t}$; without loss of generality we assume $S_{b_{t-1}b_t} = 0$ when $b_{t-1} = b_t$. Payoffs are then summarized in the following table.

		firm	
		$a_t = N$	$a_t = O$
regulator	$b_t = W$	$-S_{b_{t-1}W}, 0$	$-S_{b_{t-1}W}, \pi$
	$b_t = I$	$-i - S_{b_{t-1}I}, 0$	$y - i - S_{b_{t-1}I}, -\gamma$

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.

The regulator's discount factor is $\delta \in (0, 1)$, while each firm is short-lived and hence myopic. That firms are short-lived excludes as a motivation for offending possible learning about the choices of the regulator.

Information. In each period t , a public signal may be generated providing information

⁸We 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 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.

⁹Note that, as is well known, there is a unique equilibrium of the stage game without switching costs. In this equilibrium, the regulator chooses I with probability $\pi/(\pi + \gamma)$, while the agent chooses O with probability i/y . These probabilities ensure players are indifferent between their two actions (W and I for the regulator and N and O for the agent).

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 his private actions. Thus a private history for the regulator at date t is $\tilde{h}^t \equiv (h^t, b^t)$. A total history of the game is the private history of the regulator and the actual choices of the firm, $\hat{h}^t \equiv (\tilde{h}^t, a^t)$, where $a^t \in \{O, N\}^t$.

Strategies, equilibrium and continuation payoffs. We let the strategy of a date- t firm be given as follows: For each $h^t \in \{0, 1\}^t$, let $\alpha(h^t) \in [0, 1]$ be the probability that the date- t firm offends (at date t). 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 $\beta(h^t)$, are particularly useful since (i) the 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.

Before beginning our analysis, it is 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

$$V_t(\beta, \alpha; \tilde{h}^t) = \mathbb{E}_{\beta, \alpha} \left[\sum_{s=t}^{\infty} \delta^{s-t} (y b_s a_s - i b_s - S_{b_{s-1} b_s}) \mid \tilde{h}^t \right].$$

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

We restrict attention to parameters such that equilibrium involves infinitely repeated switching, as described in the Introduction.

(A1) Double-switching is costly, i.e., $S_{IW} + S_{WI} > 0$.

(A2) The regulator switches to wait if no offending occurs in the future, i.e., $S_{IW} < \frac{i}{1-\delta}$.

(A3) The regulator has incentives to switch to inspect if the firm offends for sure, i.e., $S_{WI} + \delta S_{IW} < y - i$.

Assumption A1 implies a friction in the regulator's switching. Assumptions A2 and A3 ensure this friction is not too large, permitting repeated switching to emerge in equilibrium.¹⁰

Since firms are myopic, a firm offends at date t only if $(1 - \beta_t)\pi - \beta_t\gamma \geq 0$. Let $\beta^* \equiv \frac{\pi}{\pi+\gamma}$ be the belief which keeps a firm indifferent between offending and not. We begin by using Assumption A2 to show that the probability of inspection, conditional on the public history, is never higher than β^* in equilibrium.

Lemma 1 *For all equilibria, at all h^t , $\beta(h^t) \leq \beta^*$.*

The above property is the dynamic analogue of the equivalent observation for the static inspection game. If $\beta(h^t) > \beta^*$ at some history h^t , the firm would not offend, but this would undermine the regulator's incentive to inspect, yielding a contradiction. We can use this (together with Assumption A1) to provide an important observation regarding the regulator's expected equilibrium payoffs.

Lemma 2 *For all equilibria, for all h^t with $h_{t-1} = 0$, $V_W(h^t) = 0$.*

Thus, a regulator who waits in a given period weakly prefers to wait in all subsequent periods (since this yields a payoff zero). Indeed, the proof proceeds by showing that the regulator must never have strict incentives to switch to inspect; such incentives would imply a contradiction to Lemma 1. We use the result, together with Assumption A3, to show the following.

¹⁰If $S_{IW} > \frac{i}{1-\delta}$, the regulator never finds it optimal to switch from inspect to wait (irrespective of firms' offending decisions). If $S_{WI} + \delta S_{IW} > y - i$, one can show that the regulator never switches from wait to inspect in equilibrium (since $b_{-1} = W$, the regulator never switches).

Lemma 3 *For all equilibria, for all h^t , $\beta(h^t) \geq \beta^*$.*

The idea is simply that if $\beta(h^t) < \beta^*$ at any history h^t , the date- t firm would offend with certainty at date t , but (given A3 and Lemma 2) this would give the regulator a strict incentive to inspect, irrespective of the action taken in the previous period. But this contradicts $\beta(h^t) < \beta^*$.

Lemmas 1 and 3 together imply that firms are necessarily indifferent between offending and not, i.e. $\beta_t = \beta^*$ for all t . The indifference of the firm is analogous to the finding for the stage game without switching costs (where the regulator's equilibrium inspection probability is β^*). As we find in Section 4.1, the indifference property is particular to a model where firms are homogeneous in their payoffs and penalties for offending. A more realistic model which allows for firm heterogeneity not only yields firms with strict incentives to offend or not, but also inspection probabilities β_t which vary with time. The latter seems to offer a more plausible evolution of perceptions concerning the risk of detection. However, the baseline model of this section is easier to solve while yielding many of the key insights, and is therefore a natural place to start.

Given that firms are indifferent to offending, our task is to find any collection of offense probabilities $\alpha(h^t)$ which yield the optimality of a regulator strategy consistent with $\beta(h^t) = \beta^*$, for all h^t . We will show that this collection is unique. We begin by determining the range of expected continuation payoffs for the regulator $V_I(h^t)$ following an inspection, as summarized in the following lemma.

Lemma 4 *For all equilibria, for all h^t , $V_I(h^t) \in [-S_{IW}, S_{WI}]$. If $h_{t-1} = 1$, then $V_I(h^t) = -S_{IW}$. If $h_{t-1} = 0$ and $\alpha(h^{t-1}) > 0$, then $V_I(h^t) = S_{WI}$.*

Lemma 4 follows after noticing that the regulator cannot have strict incentives to switch, either from inspect to wait or wait to inspect. The regulator is willing to switch to wait after a conviction, and to inspect at a history h^t such that $h_{t-1} = 0$ and $\alpha(h^{t-1}) > 0$. In particular, following any h^t such that $h_{t-1} = 0$ and $\alpha(h^{t-1}) > 0$, the regulator must switch from wait to inspect with a probability such that date- t beliefs remain at β^* . Using Bayes' rule to find the posterior probability that the regulator inspected at $t-1$, we find that the probability of switching must be ξ satisfying

$$\underbrace{\frac{\beta^*(1 - \alpha(h^{t-1}))}{\beta^*(1 - \alpha(h^{t-1})) + 1 - \beta^*}}_{\Pr(b_{t-1}=I|h_{t-1}=0)} + \xi \underbrace{\frac{1 - \beta^*}{\beta^*(1 - \alpha(h^{t-1})) + 1 - \beta^*}}_{\Pr(b_{t-1}=W|h_{t-1}=0)} = \beta^*. \quad (1)$$

Here, ξ is the probability of switching from wait to inspect conditional on the public history (the regulator's switching probability may, however, vary with the payoff-irrelevant components of its private history; i.e., with its decision to inspect at dates before the previous period).

We can now determine the firms' strategies. Consider h^t such that $h_{t-1} = 0$ and $\alpha(h^{t-1}) > 0$. That the regulator is willing to inspect at t and then follow an optimal continuation strategy implies

$$V_I(h^t) = -i + (1 - \alpha(h^t)) \delta V_I(h^t, 0) + \alpha(h^t) (y + \delta V_I(h^t, 1)). \quad (2)$$

The right-hand side is the expected value of continuation payoffs at t given that the regulator inspects. Using Lemma 4, we have equivalently

$$S_{WI} = -i + (1 - \alpha(h^t)) \delta S_{WI} + \alpha(h^t) (y - \delta S_{IW}). \quad (3)$$

Thus, we must have $\alpha(h^t) = \alpha^*$, where $\alpha^* = \frac{i + S_{WI}(1 - \delta)}{y - \delta(S_{IW} + S_{WI})}$ is the value solving (3). By Assumptions A1-A3, $\alpha^* \in (0, 1)$.

Next, we show that the probability of offending can never exceed α^* . Because the regulator never has strict incentives to switch, Equation (2) holds for *any* h^t , so that (using Lemma 4)

$$\begin{aligned} S_{WI} &\geq -i + (1 - \alpha(h^t)) \delta S_{WI} + \alpha(h^t) (y - \delta S_{IW}) \\ &= -i + \delta S_{WI} + \alpha(h^t) (y - \delta(S_{WI} + S_{IW})). \end{aligned} \quad (4)$$

Given Assumptions A1-A3, $y - \delta(S_{WI} + S_{IW}) > 0$, i.e. the right-hand side is increasing in $\alpha(h^t)$. Thus (4) can hold only if $\alpha(h^t) \leq \alpha^*$.

Finally, consider any public history h^t with $h_{t-1} = 1$. Suppose that the next date the firm offends with positive probability is $t + T$ with $T \geq 0$. With an abuse of notation, denote this probability by $\alpha_0^* = \alpha(h^{t-1}, 1, \mathbf{0}^T)$. Since the regulator must be willing to continue inspecting at time t (it never strictly prefers to switch to wait), we must have

$$V_I(h^t) = - \sum_{j=0}^T \delta^j i + (1 - \alpha_0^*) \delta^{T+1} V_I(h^t, \mathbf{0}^{T+1}) + \alpha_0^* \delta^T (y + \delta V_I(h^t, \mathbf{0}^T, 1)).$$

Equivalently, given Lemma 4, we have

$$-S_{IW} = - \sum_{j=0}^T \delta^j i + (1 - \alpha_0^*) \delta^{T+1} S_{WI} + \alpha_0^* \delta^T (y - \delta S_{IW}). \quad (5)$$

It is straightforward to verify that there is a unique solution to this equation such that $\alpha_0^* \in$

$(0, \alpha^*]$ and T is integer-valued.¹¹ In particular, we have

$$T = \left\lceil \log \left(\frac{i - S_{IW}(1 - \delta)}{i + S_{WI}(1 - \delta)} \right) / \log(\delta) \right\rceil. \quad (6)$$

We now give our equilibrium characterization. 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. There is also an “initialization” period, which reflects our assumption that the game starts with the regulator waiting rather than inspecting.

Proposition 1 *An equilibrium exists. Furthermore, under Assumptions A1-A3, the following holds in any equilibrium:*

- Step 0. (Initialization) At time 0, the regulator switches with probability β^* to inspect and the firm offends with probability α^* . If there is no conviction the play moves to Step 1, and otherwise it moves to Step 2.*
- Step 1. (Stationary phase) If the regulator inspects in the previous period, then she keeps inspecting, and if, instead, the regulator waits in the previous period, then she switches to inspect with probability ξ given by (1). The firm randomizes, playing O with probability α^* . 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 $1 - \beta^*$ to wait. In the subsequent T periods, the regulator does not switch. Firms do not offend in the T periods following the conviction. Then, in the period $T + 1$ after the conviction, the firm offends with probability $\alpha_0^* \in (0, \alpha^*]$. If there is no conviction, the play moves to Step 1; and otherwise, it reinitializes at Step 2.*

It is worth reiterating that the equilibrium process for convictions is uniquely determined: the probability of inspection conditional on the public history is fixed at β^* , and firms’ strategies (which depend only on the public history) are unique. 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. Equilibrium

¹¹See the proof of Proposition 4 in the Supplementary Material.




Figure 1: Example of dynamics of the perceptions of the firms about the likelihood of detection (β , upper figure) and probability of a crime being committed (α , lower figure) for our base model. In the upper figure, empty dots correspond to the posterior about the action taken in the previous period. In both graphs, there is a conviction in periods 3 and 16.

is unique if we restrict attention to switching probabilities that depend only on the public history.

Equilibrium dynamics are depicted in Figure 1. In the residual deterrence phase which follows a conviction, firms stop offending for $T = 9$ periods. There is then one period of offending with a probability α_0^* . If there is no conviction, the stationary phase begins, in which firms offend with the highest probability, α^* . This gradual resumption of offending coincides with what Sherman (1990) terms “deterrence decay” — the re-emergence of offending at a baseline level, in this case with probability α^* .

The evolution of the regulator’s incentives in equilibrium can be understood by considering how the continuation payoff $V_I(h^t)$ changes over time. Figure 2 plots this continuation payoff for an example. Following a conviction, which occurs at $t = 3, 16$ in the example, the payoff from inspecting is equal to $-S_{IW}$, and the regulator is thus indifferent between inspecting and switching to wait. Given that T in the example, offending ceases at the following dates and the continuation payoff $V_I(h^t)$ grows as the resumption of offending grows nearer. While $V_I(h^t) \in (-S_{IW}, S_{WI})$, the regulator strictly prefers not to switch irrespective of whether it played wait or inspect in the previous period. Finally, for any history h^t with $h_{t-1} = 0$ and $\alpha_{t-1}(h^{t-1}) > 0$, we have $V_I(h^t) = S_{WI}$, and the regulator is indifferent between remaining at

wait and switching to inspect.

The length of the deterrence phase, and reduced level of offending α_0^* , ensure the regulator is precisely indifferent between continuing to inspect and switching to wait following a conviction. By continuing to inspect following a conviction at date $t - 1$, the regulator continues to incur the per-period cost i , but it obtains a conviction at date $t + T$ with probability $\alpha_0^* > 0$, and expects a continuation payoff S_{WI} if reaching date $t + T + 1$ without any new conviction. The payoff from continuing to inspect must be balanced against the cost S_{IW} of switching to wait, explaining why T is increasing in both switching costs.

If we take the regulator to be very patient, the expression for T in (6) becomes more parsimonious. In particular,

$$\lim_{\delta \nearrow 1} \left\{ \log \left(\frac{i - S_{IW}(1 - \delta)}{i + S_{WI}(1 - \delta)} \right) / \log(\delta) \right\} = \frac{S_{IW} + S_{WI}}{i},$$

so as $\delta \nearrow 1$, the length of deterrence is determined simply by the ratio of the (sum of) switching costs to the per-period inspection cost. Good information about these costs hence permits a straightforward prediction about the length of deterrence.¹²




Figure 2: Example of dynamics of V_I as a function of time. In the graph, there is conviction in periods 3 and 16.

3.1 Comparative Statics

A natural question is how the firms' rate of offending depends on the level of penalties. We find that the effect of a (permanent) change in the penalties depends on the horizon considered. In

¹²Admittedly, such costs would often be difficult to measure. For instance, it is difficult to know how regulatory officials perceive the opportunity cost of their time, and this might well differ from easily-observed measures of personnel costs such as wages.

the *short run*, raising penalties can be expected to reduce offending, while lowering penalties increase it. To illustrate, we consider an unanticipated and permanent change in the penalty at an arbitrary date $t > 1$. The effects of an increase and decrease are asymmetric, so they must be considered separately.

Corollary 1 *Suppose there is an unforeseen, but public and permanent, change in the penalty γ to γ_{new} at the beginning of period t . Then*

1. *if $\gamma_{new} > \frac{1}{1-\alpha_{t-1}}\gamma$ or $h_{t-1} = 1$, the equilibrium play enters the “residual deterrence phase” at date t , i.e. there is no offending until period $t + T$, and*
2. *if $\gamma_{new} < \frac{1}{1-\alpha_{t-1}}\gamma$ and $h_{t-1} = 0$, then the equilibrium play is in the “stationary phase” at date t , i.e. the date- t firm offends with probability α^* (and offending continues at rate α^* until after the next conviction).*

The result can be illustrated by considering a history h^t such that the date $t - 1$ firm offends with probability α^* (i.e., play has entered the “stationary phase”). Consider then the effect of a change in the penalty from γ to γ_{new} announced at the beginning of date t . After observing no conviction at date $t - 1$, the date- t firm assigns a probability $\frac{\beta^*(1-\alpha^*)}{\beta^*(1-\alpha^*)+1-\beta^*}$ to the regulator having inspected at date $t - 1$. If $\gamma_{new} > \frac{1}{1-\alpha^*}\gamma$, this probability is higher than $\beta_{new}^* = \frac{\pi}{\pi+\gamma_{new}}$. Lemmas 1-3 above imply that the regulator must be willing to switch from inspect to wait, in order to guarantee that $\beta_t = \beta_{new}^*$. This switching is, in turn, only incentive compatible for the regulator if the change in penalty is followed by a sequence of T periods without offending. If, instead, $\gamma_{new} < \frac{1}{1-\alpha^*}\gamma$, the reverse applies: the regulator must be willing to switch from wait to inspect, so equilibrium play must be in the stationary phase.

We next turn our focus to long-run effects, and in particular what we term the “long-run average offense rate”. One finds the ex-ante expected average rate of offending over the first τ periods, and then takes the limit as $\tau \rightarrow \infty$. To calculate this, we use that the expected duration between convictions is $T + 1 + \frac{1-\alpha_0^*\beta^*}{\alpha^*\beta^*}$. The long-run average offense rate is then

$$\bar{\alpha} = \frac{\alpha_0^* + \frac{1-\alpha_0^*\beta^*}{\alpha^*\beta^*}\alpha^*}{T + 1 + \frac{1-\alpha_0^*\beta^*}{\alpha^*\beta^*}} = \frac{\alpha^*}{1 + \beta^*(\alpha^*(T + 1) - \alpha_0^*)}. \quad (7)$$

We find the following.

Corollary 2 *The long-run average offense rate is increasing in the penalty γ .*

While Corollary 2 may seem counterintuitive, the reason for this result is straightforward. A higher value of the penalty γ (or a lower value of π) reduces the probability of monitoring β^* required such that a firm is indifferent to offending. Provided the regulator's mixing probabilities are adjusted appropriately to maintain firm beliefs at the lower level, firms are content to play the same strategies specified in Proposition 1 to maintain the regulator's incentives. Hence, the duration of complete deterrence (T) remains the same, as does the crime rate $T + 1$ periods after the previous conviction (α_0^*) and subsequently (α^*). Nevertheless, due to the lower value of β^* , convictions occur less frequently. Episodes of residual deterrence are thus less frequent.

Corollary 2 provides a new answer to an old question raised by Becker (1968) regarding why maximal penalties may not be optimal.¹³ In our model, a planner who is concerned only with the long-run average offense rate, gains by reducing the penalty. As noted in the Introduction, one should, however, be careful to interpret this result in light of the partial equilibrium nature of our analysis. In particular, it seems important to explicitly account for the alternative actions of the regulator when not inspecting a given group of potential offenders. For instance, consider an alternative model where the regulator must always inspect, but can rotate inspections (at a cost) among different industries. In such a model, we expect that raising the penalty uniformly across industries would increase long-run deterrence, since (contrary to the model of our paper) the total time inspecting remains unchanged. While such a model seems realistic, it also introduces considerable complexity to the analysis, which would tend to obfuscate the simple intuitions of the paper.¹⁴

Our finding should also be related to the result for a static setting. Tsebelis (1989), for instance, highlights that, in the one-shot inspection model, the probability of inspection offsets any change in penalties exactly, so that an increase in the penalty does not affect offending.¹⁵

¹³An alternative theory which has dominated the literature is the idea of “marginal deterrence”, as discussed, for instance, by Stigler (1970) and Shavell (1992). In this view, not all penalties should be set at their maximum level. Instead, it may be desirable to set penalties for less harmful acts below the maximum to entice offenders away from the most harmful acts (which should indeed receive the maximal penalty).

¹⁴In the richer model, the regulator's incentives to switch to inspecting a given industry would depend on the history of convictions (and hence subsequent deterrence) in other industries. In other words, one would need to account for the evolution of the value of the regulator's outside option of monitoring other industries. By focusing on the decision *whether* to inspect (rather than where to inspect), the model of the present paper abstracts from this difficulty.

¹⁵The empirical frequencies of choices by players of inspection and matching pennies games in the laboratory typically respond to the players' own payoffs, in an apparent contradiction to the Nash prediction (see Nosenzo et al. (2014) for a recent experiment involving the inspection game). This principle would suggest that higher

Our result goes further by suggesting a reason why the rate of offending may actually increase. Our finding hinges on the positive switching cost: if $S_{WI} = S_{IW} = 0$, then equilibrium play in the repeated game simply involves a repetition of the static inspection and offense rates. As one might expect, this is also the limiting behavior as S_{WI} and S_{IW} approach zero.¹⁶

Next, we investigate the effects of changes in the switching costs on the long-run average offense rate.

Corollary 3 *The values α^* and T are increasing in both S_{IW} and S_{WI} . The long-run average offense rate $\bar{\alpha}$ is continuous in (S_{IW}, S_{WI}) . Moreover:*

1. *For generic values of (S_{IW}, S_{WI}) ,¹⁷ $\bar{\alpha}$ is increasing in S_{IW} if $\beta^* < \delta^{T+1}$ and decreasing in S_{IW} if $\beta^* > \delta^{T+1}$; hence $\bar{\alpha}$ is quasi-concave in S_{IW} .*
2. *$\bar{\alpha}$ is increasing in S_{WI} .*

The effect of switching costs on T is discussed above, while the reason α^* is increasing in both switching costs is as follows. For a higher value of S_{WI} , the level of offending α^* must be greater to incentivize switching from wait to inspect. For a higher value of S_{IW} , the regulator's continuation payoff at the date of a conviction (i.e., $y - i - \delta S_{IW} > 0$) is lower. Hence, again, α^* must be higher for the regulator to be willing to switch to inspect.

The long-run average offense rate $\bar{\alpha}$ is, naturally, increasing in α^* but decreasing in the length of deterrence T . Consider then Part 1 of the corollary. If β^* is small, i.e. the probability of inspection in each period is low, then deterrence phases occur only rarely. If, in addition, T is not too large, then the long-run average offense rate is close to α^* . Since α^* is increasing in S_{IW} , it is unsurprising that the long-run average offense rate $\bar{\alpha}$ is then also increasing in S_{IW} . Conversely, when β^* is close to 1, deterrence phases occur frequently, and the length of the deterrence phase T has a greater impact on the long-run average offense rate. Since T is increasing in S_{IW} , the long-run average offense rate then decreases in S_{IW} .

In Part 2 of the corollary, we find that the long-run average offense rate is necessarily increasing in S_{WI} . In other words, the effect of S_{WI} on α^* always dominates. This finding

penalties reduce the frequency of offending, at least in the lab. It is difficult to make predictions about possible play of our dynamic game in hypothetical experiments, but we expect some features of our equilibrium would be robust. In particular, to the extent an increase in penalties reduces inspections, the effect should be to reduce the episodes of residual deterrence, with implications for the overall rate of offending.

¹⁶See Appendix B in the Supplementary Material.

¹⁷More precisely, for all values (S_{IW}, S_{WI}) except possibly points of discontinuity of T .

is perhaps of interest for policy: it suggests that lowering a regulator’s cost of commencing inspection activities may, by itself, reduce the rate of offending.

Corollary 3 indicates that positive switching costs may increase or decrease the long-run average offense rate relative to the rate for $S_{IW} = S_{WI} = 0$ (i.e., compared to i/y). This conclusion is perhaps surprising given that residual deterrence arises only when the switching cost is positive. The reason for this result is that, as explained above, higher switching costs increase α^* , the equilibrium rate of offending in the stationary phase. Nonetheless, in case the regulator’s reward y from a conviction is sufficiently large (precisely, if $y > \frac{i}{1-\delta} + S_{WI}$), $S_{IW} \nearrow \frac{i}{1-\delta}$ implies $T \rightarrow +\infty$, driving the long-run average offense rate to zero.

4 Extensions

4.1 Heterogeneous Firms

Our baseline model establishes a tractable framework for understanding deterrence. A deterrence phase provides, in equilibrium, a (weak) incentive for the regulator to stop inspecting after a conviction. Indeed, reduced offending after a conviction induces the regulator to withdraw from inspecting in order to save on inspection costs.

The model implies that the perceived threat of apprehension and punishment, measured by β_t , is constant and always equal to β^* . The threat of apprehension is hence the same irrespective of whether play is in the deterrence phase. Another observation is that deterrence phases are characterized by a complete cessation of offending, which means that firms learn nothing about the regulator’s activities during these phases. In this section, we introduce heterogeneity in firm payoffs and show that both conclusions are overturned.

We show that, when firms’ payoffs are heterogeneous, deterrence phases are characterized by a high perceived risk of apprehension. The resulting offense rate is low, but still positive. Hence, the absence of a conviction is evidence that the regulator is not inspecting, and the perceived likelihood of inspection falls with the time since the last conviction. This means that deterrence decay in this model coincides with a gradual reduction in the chances of being caught. We thus provide a rational basis on which beliefs about the likelihood of punishment decline with the time since the last observed punishment.¹⁸

¹⁸Compare this, for instance, to Chen’s (2015) suggestion that such changes in beliefs are to be expected on the basis of so-called “recency biases”, whereby individuals incorrectly overweight recent events in determining their beliefs.

We introduce payoff heterogeneity as follows. At the beginning of each period t , the firm independently (and privately) draws a value π_t from a continuous distribution F with full support on a finite interval $[\underline{\pi}, \bar{\pi}]$, $0 < \underline{\pi} < \bar{\pi}$. We maintain Assumptions A1-A3 precisely as in Section 3.

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

$$\alpha_t = \Pr\left(\frac{\pi_t}{\pi_t + \gamma} \geq \beta_t\right) = \Pr\left(\pi_t \geq \frac{\beta_t}{1 - \beta_t} \gamma\right) = 1 - F\left(\frac{\beta_t}{1 - \beta_t} \gamma\right) \equiv \alpha(\beta_t), \quad (8)$$

where our definition of $\alpha(\cdot)$ involves an obvious abuse of notation. Let $\bar{\beta} \equiv \frac{\bar{\pi}}{\bar{\pi} + \gamma}$, that is, $\alpha(\bar{\beta}) = 0$.

Before providing our characterization result, it is useful to describe how equilibrium is similar to our baseline model. After a conviction, using an argument similar to that for Lemma 1, it is easy to see that the regulator must have (weak) incentives to switch to wait. As before, the regulator prefers to make any switch to wait sooner rather than later, in order to avoid inspection costs. Also as before, the incentive to switch to wait must result from some periods of deterrence. Nevertheless, different from our base model, the crime rate is positive in the deterrence phase and beliefs are updated following the absence of a conviction. The probability assigned to the regulator inspecting falls gradually until the crime rate is high enough that the regulator is (weakly) willing to switch from waiting to inspect. The crime rate then remains fixed until a new conviction occurs. Figure 3 depicts the implied dynamics of the perceptions of the firms about the likelihood of detection and probability of a crime being committed.

Proposition 2 *In the model with heterogeneous firms, there exists a unique $\beta^{\max} \in (0, \bar{\beta})$ and $T \in \mathbb{N} \cup \{0\}$ such that, in any equilibrium:*

Step 0. (Initialization) At time 0 the regulator switches to inspect with probability $\beta^{\min} \equiv \alpha^{-1}(\alpha^)$ (with $\alpha^* = \frac{i + S_{WI}(1 - \delta)}{y - \delta(S_{WI} + S_{IW})}$ as in our baseline model) and the firm offends with probability α^* . If there is no conviction the play moves to Step 1, and otherwise it moves to Step 2.*

Step 1. (Stationary phase) If the regulator inspects in the previous period, then she keeps inspecting, and if, instead, the regulator waits in the previous period, then she switches to inspect with a positive probability such that the probability of inspection equals β^{\min} . The firm randomizes, playing O with probability α^ . If there is a conviction, the play moves to Step 2; otherwise, it stays in Step 1.*

Step 2. (Residual deterrence, following a conviction) *The regulator switches with probability $1 - \beta^{\max}$ to wait. In the subsequent T periods, the regulator does not switch, and, as long as there is no conviction, the firms' beliefs evolve according to Bayes' rule. That is, if there is a conviction at time $t - 1$ and no conviction in $\{t, \dots, t + j\}$, for $j \in \{0, \dots, T\}$, then $\beta_t = \beta^{\max}$ and for each $k \in \{0, \dots, j - 1\}$ we have that β_{t+k+1} is given recursively by*

$$\beta_{t+k+1} = \frac{\beta_{t+k}(1 - \alpha(\beta_{t+k}))}{\beta_{t+j}(1 - \alpha(\beta_{t+k})) + 1 - \beta_{t+j}} < \beta_{t+k} . \quad (9)$$

In this case, the firm in period $t + j$ offends with probability $\alpha(\beta_{t+j})$. If there is a conviction between t and $t + T$, Step 2 is reinitialized; if there is no conviction, the play moves to Step 1 at time $t + T + 1$.




Figure 3: Example of dynamics of the perceptions of the firms about the likelihood of detection (β , upper figure) and probability of a crime being committed (α , lower figure) for the random firm payoffs case. In the upper figure, empty dots correspond to the posterior about the action taken in the previous period. In both graphs, there is a conviction in periods 3 and 16.

We now briefly describe how one obtains the value of β^{\max} and β^{\min} (see the Appendix for details). First note that, as in our baseline model, Lemma 2 applies: The regulator's continuation payoff following wait is zero. This must be the case, since the regulator can obtain zero by simply continuing to wait. A strictly positive payoff would imply a strict preference for switching to inspect, which, as we explain above, is inconsistent with firm preferences for offending (such offending is, of course, essential for the switch to inspect to be

profitable). This pins down the continuation value of inspecting after a conviction (equal to $-S_{IW}$) and during the stationary phase which follows the deterrence phase (equal to S_{WI}). The same argument as for our baseline model shows that the rate of offending during the stationary phase is equal to α^* (again equal to $\frac{i+S_{WI}(1-\delta)}{y-\delta(S_{WI}+S_{IW})}$), which implies $\beta^{\min} = \alpha^{-1}(\alpha^*)$.

Now consider the payoff following a conviction; i.e., consider $V_I(h^t)$ with $h_{t-1} = 1$. We examine the effect on this payoff of changing β^{\max} , keeping β^{\min} fixed. For each value of β^{\max} , consider beliefs β_{t+j} updated according to Bayes' rule, as in (9). As $\beta^{\max} \rightarrow \bar{\beta}$, the number of periods j for β_{t+j} to reach β^{\min} after a conviction tends to infinity, and the crime rate during the deterrence phase tends to 0. This implies that $V_I(h^t)$ approaches $\frac{-i}{1-\delta} < -S_{WI}$. If, alternatively, $\beta^{\max} \rightarrow \beta^{\min}$ (so the number of periods of deterrence tends to 0), the value of inspecting $V_I(h^t)$ converges to S_{IW} . Thus, it is easy to see that there is a unique intermediate value of β^{\max} such that $V_I(h^t) = -S_{WI}$, and thus such that the regulator is indifferent between continuing to inspect and switching to wait.

Finally, it is worth reiterating that the forces in the model with heterogeneous firms are closely related to those in our baseline model. In fact (as we show in Appendix B of the Supplementary Material), equilibrium converges to that in the baseline model as the distribution F over the rewards π becomes degenerate. This also establishes that comparative statics for the baseline model continue to apply to heterogeneous firms, provided this heterogeneity is not too great.

4.2 Imperfectly Observed Offenses

A notable feature of our motivating examples is that firms' actions generate noisy information about whether they have broken the law. For instance, a firm engaging in price fixing would be expected to set higher prices than a non-price fixing firm, but the legality of a firm's pricing cannot be perfectly inferred based on prices alone. The simplest way to model this possibility is to introduce a noisy signal of firm offending which emerges after the regulator and firm's simultaneous actions. We then find that equilibrium is much the same as characterized in Proposition 1.

In addition to the public signals of past convictions, we assume that, at the end of each period t , there is some public signal $\mu_t \in (\underline{\mu}, \bar{\mu})$ about the action of the firm. The signal is distributed according to F_{a_t} , for $a_t \in \{N, O\}$. F_{a_t} is absolutely continuous with PDF f_{a_t} , which is positive on the entire support. Moreover, $f_O(\cdot)/f_N(\cdot)$ is strictly increasing and has \mathbb{R}_{++} as its image.

Firms' beliefs as to the probability of inspection must now be calculated accounting for

signals both of convictions and the signal μ_t . Let β'_t be the posterior about the action of the regulator at time $t - 1$ being I calculated at the beginning of period t (i.e., after period $t - 1$ signals are observed, but before the regulator switches in period t). Following a conviction at date $t - 1$ we have $\beta'_t = 1$. Otherwise, β'_t is given by Bayes' rule, that is

$$\frac{\beta'_t}{1 - \beta'_t} = \frac{\beta_{t-1}}{1 - \beta_{t-1}} \frac{(1 - \alpha_{t-1})f_N(\mu_{t-1})}{(1 - \alpha_{t-1})f_N(\mu_{t-1}) + \alpha_{t-1}f_O(\mu_{t-1})}. \quad (10)$$

We will find that equilibrium play is as specified in Proposition 1, with the exception of the regulator's randomizations. These must be adjusted to maintain the probability of inspection equal to β^* following every history. Thus, the regulator still switches to wait with probability $1 - \beta^*$ following a conviction. Its switching following no conviction is determined by (10), taking $\beta_{t-1} = \beta^*$. Note then that (given no conviction at $t - 1$) we have $\beta'_t \leq \beta^*$, with a strict inequality if and only if $\alpha_{t-1} > 0$. Hence, either $\alpha_{t-1} = 0$ and the regulator does not switch, or $\alpha_{t-1} > 0$ and it switches from wait to inspect with probability $\xi(\beta'_t)$, where $\xi(\beta'_t)$ satisfies

$$\beta'_t + \xi(\beta'_t)(1 - \beta'_t) = \beta^*. \quad (11)$$

We summarize these observations in the following corollary to Proposition 1.

Corollary 4 *Proposition 1 remains the same in the model with signals of the firms' offending, with the exception of the switching probability specified in Step 1 (which is now given by (11)).*

It is worthwhile to point out the relationship between the strength of the signal that the firm offended at date $t - 1$, μ_{t-1} , and the probability the regulator switches to inspect at date t (assuming no conviction occurred at date $t - 1$). When μ_{t-1} is large (it is likely the firm offended), $\frac{f_N(\mu_{t-1})}{f_O(\mu_{t-1})}$ is small, and so the posterior belief β'_t falls by more. In essence, the absence of a conviction is more informative about the regulator's failure to inspect when it is more likely that the firm has offended. This means that the probability the regulator switches from wait to inspect, $\xi(\beta'_t)$, is larger when μ_{t-1} is large.

4.3 Imperfectly Observed Inspections

To date we assumed that inspections become public only when the short-lived firm is offending. However, information about inspection activities sometimes becomes available from sources other than convictions. For instance, regulators are often required to disclose information about their activities which may provide noisy information to firms.¹⁹ We thus now consider

¹⁹For instance, the SEC is required to present a Congressional Budget Justification each year, which indicates levels of expenditure on different activities as well as some broad performance indicators.

a setting where there are public signals which are (partly) informative about the regulator's actions, in addition to the public signals which are (perfectly) informative about convictions.

As in the previous section, in addition to the public signals of past convictions, we assume that, at the end of each period t , there is some public signal $\mu_t \in (\underline{\mu}, \bar{\mu})$ about the action of the regulator. The signal is distributed according to F_{b_t} , for $b_t \in \{I, W\}$. F_{b_t} is absolutely continuous with PDF f_{b_t} , which is positive on the entire support. Moreover, $f_I(\cdot)/f_W(\cdot)$ is strictly increasing and has \mathbb{R}_{++} as its image.

As before, β'_t is the updated belief following the date $t-1$ signals but before any switching at date t . So, if there is a conviction in period $t-1$, we have $\beta'_t = 1$. Instead, if there is no conviction in period $t-1$, β'_t is given by

$$\frac{\beta'_t}{1 - \beta'_t} = \frac{\beta_{t-1}}{1 - \beta_{t-1}} (1 - \alpha_{t-1}) \frac{f_I(\mu_{t-1})}{f_W(\mu_{t-1})},$$

where α_{t-1} is the probability of offending at $t-1$, and μ_{t-1} is the realized signal of the firm's offending.

For simplicity, in this section we focus our analysis on Markov perfect equilibria (MPE), with the Markov state at time t being b_{t-1} and β'_t . Let $\beta_t = \beta(\beta'_t)$ be the probability of inspection at date t given the beginning-of-period- t posterior β'_t . We say that an MPE is *monotone* if $\beta(\cdot)$ is weakly increasing, which seems to us an appealing restriction on the class of equilibria.

Proposition 3 *In the model with imperfectly observed inspections, there exists a unique monotone Markov perfect equilibrium. In such an equilibrium, there exist $\bar{\beta} \geq \beta^*$ such that*

1. *if $\beta'_t \leq \beta^*$ then the regulator keeps inspecting if $b_{t-1} = I$ and, if $b_{t-1} = W$, switches to inspect with a probability such that $\beta_t = \beta^*$. In this case the firm offends with some probability α^{**} .*
2. *if $\beta'_t \in (\beta^*, \bar{\beta}]$ then the regulator does not switch, so $\beta_t = \beta'_t$, and the firm does not offend.*
3. *if $\beta'_t > \bar{\beta}$ then the regulator keeps waiting if $b_{t-1} = W$ and, if $b_{t-1} = I$, it switches to wait with a probability such that $\beta_t = \bar{\beta}$.*

(a) *If $\bar{\beta} = \beta^*$ the date- t firm offends with some probability $\alpha_0^{**} \in [0, \alpha^{**})$, and*

(b) *if $\bar{\beta} > \beta^*$, then the date- t firm does not offend.*




Figure 4: Example of dynamics of the perceptions of the firms about the likelihood of detection (β , upper figure) and probability of a crime being committed (α , lower figure) for the imperfectly observed inspections case. In the upper figure, empty dots correspond to the posterior about the action taken in the previous period (β'). In both graphs, there is a conviction in periods 3 and 16.

Figure 4 depicts the equilibrium dynamics in the presence of public signals. When the posterior β'_t is very high (either because of a signal (in period 6) and or because of a conviction (periods 4 and 17)), the payoff from inspecting is very low, since the expected time needed for crime to resume (and have the possibility of obtaining a payoff y) is large. As a result, the regulator switches to wait (with some probability) in order to save inspection costs. When, instead, β'_t is small, the probability of crime is positive, so inspecting becomes attractive. In this case, the regulator switches to inspect so that the perception of the firms about the likelihood of detection is β^* , in order that the probability of crime is positive but not 1. For intermediate values of β'_t , the difference in the continuation payoffs of inspecting and waiting is not enough to compensate switching in either direction, so the regulator does not change its action.

4.4 General Payoffs

So far we have considered a regulator motivated directly by its concern for apprehending offenses. As noted in the Introduction, such an assumption may be reasonable in light of the regulator's implicit rewards or career concerns. A more socially minded regulator, however,

might 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	$b_t = W$	$-S_{b_{t-1}W}, 0$	$-L - S_{b_{t-1}W}, \pi$
	$b_t = I$	$-i - S_{b_{t-1}I}, 0$	$y - L - i - S_{b_{t-1}I}, -\gamma$

Here L is the regulator's loss as a result of an offense, while y again is its reward for apprehending an offense (when $L = 0$, the model is identical to that in Section 3). We make no a priori restriction on the signs of L and y . We continue to impose Assumptions A1 and A2.

The following proposition describes a sufficient condition under which an equilibrium as in Proposition 1 exists (by this, we mean that there exist values T , α_0^* and α^* such that the play described in Proposition 1 is an equilibrium).

Proposition 4 *In the model with general payoffs, there is an equilibrium which permits the same characterization as in Proposition 1 for some $\alpha^* \in (0, 1)$, $\alpha_0^* \in (0, \alpha^*]$ and $T \in \{0\} \cup \mathbb{N}$ if*

$$L > -(y - i - S_{WI} - \delta S_{IW}) \frac{1 - \delta}{\delta \left(1 - \frac{i - (1 - \delta) S_{IW}}{i + (1 - \delta) S_{WI}}\right)}. \quad (12)$$

Furthermore, if (12) holds, the values α^ , α_0^* and T are unique.*

Condition (12) is analogous to Assumption A3 for the general payoffs case (note that it is the same when $L = 0$). Intuitively, it ensures that the regulator has sufficient incentives to switch to inspect. When L is large, this incentive is greater, because the deterrence motive is greater. Thus we indeed find that the conditions for the existence of equilibria with residual deterrence are relaxed when $L > 0$ relative to our baseline model.

Proposition 4 establishes that equilibrium dynamics with residual deterrence can be found also when the regulator cares about deterring offenses. This is not merely a robustness exercise, and extends our equilibrium construction to a different class of economically relevant settings. A particularly interesting case 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 $y = 0$. Note that, here, 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, in the repeated game

with switching costs, if L is large enough, there exist equilibria where the regulator inspects only because after a conviction there are some periods of residual deterrence.

Note that, unlike the analysis for the baseline model (and many of the other results above), the uniqueness result in Proposition 4 follows only after restricting attention to equilibria of the form in Proposition 1 and to values such that $\alpha_0^* \leq \alpha^* < 1$. To see why this restriction is important, consider the case above where $y = 0$. In this case, there is 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 1). This suggests the existence of still further, non-stationary, equilibria, although we make no attempt to characterize the multiplicity.

Finally, it is interesting to note that there is a certain equivalence between equilibria in the setting where the regulator is motivated by deterrence, and the one in our baseline model (where $L = 0$).

Corollary 5 *Fix parameter values $(\pi, \gamma, i, S_{IW}, S_{WI})$ as specified in the model set-up. Fix a pair of values (y, L) satisfying Equation (12). The values α^* , α_0^* and T for the equilibrium specified in Proposition 4 also characterize an equilibrium in the model with parameters $(y', 0)$, where*

$$y' = y + \delta L \left(\sum_{s=1}^T \delta^{s-1} \alpha^* + \delta^T (\alpha^* - \alpha_0^*) \right),$$

and with remaining parameters $(\pi, \gamma, i, S_{IW}, S_{WI})$ unchanged.

The result shows that any equilibrium described in Proposition 4 corresponds to an equilibrium of the baseline model, described in Proposition 1, once the reward for convicting a firm is correctly specified. In particular, this reward should be comprised of two terms: the direct benefit y from a conviction, and the reduction in offending that occurs in the deterrence phase following the conviction (this is α^* in the T periods following the conviction, while it is $\alpha^* - \alpha_0^*$ in period $T + 1$ after the conviction). The result justifies our initial focus on the model where the regulator is not motivated by deterrence, and is also suggestive of how one might extend many of our results above to accommodate a regulator’s preference for deterrence (at least after appropriately restricting the class of equilibria).

5 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 switching between the two

activities, inspect and wait. We showed that this switching cost gives rise to “reputational” effects. Following a conviction, offending may cease for several periods before resuming at a steady level. This effect may be 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 4.4). In an extension to the baseline model where potential offenders are heterogeneous (Section 4.1), the risk of apprehension follows a plausible pattern, being the highest immediately after a conviction and then decaying gradually. We thus provide a fully-fledged theory of deterrence decay.

While our model is stylized — the inspector faces a sequence of myopic offenders — we believe it is empirically relevant. The model provides predictions for both the enforcement authority’s and the firms’ behavior, although neither of these may be directly observable to an empirical researcher. For this reason, current empirical studies have often sought to measure deterrence effects by focusing on what are presumably noisy signals of offending (see, e.g., Block, Nold and Sidak, 1981, and Jennings, Kedia and Rajgopal, 2011). Our findings are in line with these studies. For instance, Section 4.2 presents a model where noisy signals of firm offending are public information. An observer of these signals would conclude that offending is less likely following a conviction.

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] Bassetto, Marco and Christopher Phelan (2008), ‘Tax Riots,’ *Review of Economic Studies*, 75, 649-669.
- [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] Caruana, Guillermo and Liran Einav (2008), ‘A Theory of Endogenous Commitment,’ *Review of Economic Studies*, 75, 99-116.
- [10] Chen, Daniel (2015), ‘The Deterrent Effect of the Death Penalty? Evidence from British Commutations During World War I,’ mimeo Toulouse School of Economics.
- [11] DeMarzo, Peter, Michael Fishman and Kathleen Hagerty (1998), ‘The Optimal Enforcement of Insider Trading Regulations,’ *Journal of Political Economy*, 106, 602-632.
- [12] Dilmé, Francesc (2014), ‘Reputation Building through Costly Adjustment,’ mimeo University of Pennsylvania and University of Bonn.
- [13] Eeckhout, Jan, Nicola Persico and Petra E. Todd (2010), ‘A Theory of Optimal Random Crackdowns,’ *American Economic Review*, 100, 1104-1135.
- [14] Halac, Marina and Andrea Prat (2014), ‘Managerial Attention and Worker Engagement,’ mimeo Columbia University and University of Warwick.
- [15] 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.
- [16] 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.
- [17] Khalil, Fahad (1997), ‘Auditing without Commitment,’ *RAND Journal of Economics*, 28, 629-640.
- [18] 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.

- [19] Lazear, Edward P (2006), ‘Speeding, Terrorism, and Teaching to the Test,’ *Quarterly Journal of Economics*, 121, 1029-1061.
- [20] Lipman, Barton and Ruqu Wang (2000), ‘Switching Costs in Frequently Repeated Games,’ *Journal of Economic Theory*, 93, 149-190.
- [21] Lipman, Barton and Ruqu Wang (2009), ‘Switching Costs in Infinitely Repeated Games,’ *Games and Economic Behavior*, 66, 292-314.
- [22] Mookherjee, Dilip and IPL Png (1989), ‘Optimal Auditing, Insurance and Redistribution,’ *Quarterly Journal of Economics*, 104, 399-415.
- [23] Mookherjee, Dilip and IPL Png (1994), ‘Marginal Deterrence in Enforcement of Law,’ *Journal of Political Economy*, 76, 1039-1066.
- [24] Nagin, Daniel S. (2013), ‘Deterrence in the Twenty-First Century,’ *Crime and Justice*, 42, 199-263.
- [25] 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.
- [26] Polinsky, Mitchell A. and Steven Shavell (1984), ‘The Optimal Use of Fines in Imprisonment,’ *Journal of Public Economics*, 24, 89-99.
- [27] Reinganum, Jennifer and Louis Wilde (1985), ‘Income Tax Compliance in a Principal-Agent Framework,’ *Journal of Public Economics*, 26, 1-18.
- [28] Reinganum, Jennifer and Louis Wilde (1986), ‘Equilibrium Verification and Reporting Policies in a Model of Tax Compliance,’ *International Economic Review*, 27, 739-760.
- [29] Shavell, Steven (1992), ‘A Note on Marginal Deterrence,’ *International Review of Law and Economics*, 12, 343-355.
- [30] Sherman, Lawrence (1990), ‘Police Crackdowns: Initial and Residual Deterrence,’ *Crime and Justice*, 12, 1-48.
- [31] Shimshack, Jay P. and Michael B. Ward (2005), ‘Regulator reputation, enforcement, and environmental compliance,’ *Journal of Environmental Economics and Management*, 50, 519-540.

- [32] Spagnolo, Giancarlo (2005), ‘*Divide et Impera: Optimal Leniency Programs*,’ mimeo Stockholm School of Economics.
- [33] Strausz, Roland (1997), ‘Delegation of Monitoring in a Principal-Agent Relationship,’ *Review of Economic Studies*, 64, 337-357.

Appendix A: Proofs of the Main Results

Proof of Lemma 1. We argue by contradiction. Suppose that there exists a date $t \geq 0$ and a public history h^t such that the firm strictly prefers not to offend, i.e. $\beta(h^t) > \beta^*$. We let $\tau \geq 1$ be the value such that each firm strictly prefers not to offend at dates t up to $t + \tau - 1$, but (weakly) prefers to offend at date $t + \tau$. Suppose first that $\tau = +\infty$, i.e., that all firms that arrive after t abstain from offending. In this case, by Assumption A2, the regulator optimally plays W from date t onwards. This contradicts the optimality of firm abstention, so it must be that $\tau < +\infty$. Next, note that $\beta(h^t, 0^{\tau-1}) > \beta^*$ and $\beta(h^t, 0^\tau) \leq \beta^*$. This implies that there exists a private history $\tilde{h}^{t+\tau-1} = (h^t, \mathbf{0}^{\tau-1}; b^{t+\tau-1})$ for the regulator where it plays inspect at date $t + \tau - 1$ and then switches to wait at date $t + \tau$ with positive probability (i.e., $\beta(h^t, 0^\tau; b^{t+\tau-1}, I) < 1$). However, the alternative choice of waiting at both dates saves $i - S_{IW}(1 - \delta)$ in costs without affecting convictions or the firms’ information. By Assumption A2, this saving is strictly positive; i.e. there exists a strictly profitable deviation for the regulator, contradicting sequential optimality of the regulator’s strategy. ■

Proof of Lemma 2. Assume otherwise, so there is a history such that $h_{t-1} = 0$ and $V_W(h^t) \neq 0$. Since the option of waiting forever gives the regulator a payoff of 0, it is necessarily the case that $V_W(h^t) > 0$. By assumption (i.e., by definition of V_W), the regulator waited in period $t - 1$. The fact that $V_W(h^t) > 0$ implies that, after this history, the regulator strictly prefers to switch from wait to inspect at some date $s \geq t$ after a public history $h^s = (h^t, \mathbf{0}^{s-t})$. This is the case because the payoff of the regulator is bounded, and if it waits forever it obtains a payoff of 0. By Assumption A1, the regulator must also strictly prefer to continue inspecting (rather than switch to wait) after h^s ,²⁰ i.e., inspecting is strictly

²⁰Suppose instead that the regulator is willing to switch from inspect to wait following h^s . This implies

$$-S_{IW} + \delta V_W(h^s, 0) \geq -i + (1 - \alpha(h^s)) \delta V_I(h^s, 0) + \alpha(h^s) (y + \delta V_I(h^s, 1)).$$

preferred to waiting at h^s irrespective of the regulator's private actions b^s . Hence, we must have $\beta(h^s) = 1 > \beta^*$, contradicting the finding in Lemma 1. ■

Proof of Lemma 3. Lemma 2 implies the following. In equilibrium, the regulator's continuation payoff, gross of present-period switching costs (i.e., ignoring present-period switching costs), is equal to zero whenever it plays the action wait. This follows because the regulator receives zero (gross of switching costs) in the period it plays wait, while its continuation payoff for the following period is also zero by Lemma 2.

Now, suppose for a contradiction that there exists a date $t \geq 0$ and a public history h^t such that $\beta(h^t) < \beta^*$, so that the firm offends with certainty at date t . There must exist a private history for the regulator $\tilde{h}^t = (h^t; b^t)$ such that it waits at date t . By the previous argument, its continuation payoff from date t onwards, ignoring any date- t switching costs, must equal zero. However, by instead inspecting at date t and then waiting at $t + 1$ the regulator can guarantee itself a payoff at least $y - i - S_{WI} - \delta S_{IW}$, which is strictly positive by Assumption A3. ■

Proof of Lemma 4. Suppose first that $V_I(h^t) > S_{WI}$ for some h^t . Given Lemma 2, the regulator must have strict incentives to switch from wait to inspect following h^t in case $b_{t-1} = W$, or (given A1) strict incentives to continue inspecting at date t in case $b_{t-1} = I$. Such strict incentives are inconsistent with $\beta(h^t) = \beta^* < 1$. Similarly, $V_I(h^t) < -S_{IW}$ implies strict incentives to wait at date t , which is inconsistent with $\beta(h^t) = \beta^* > 0$.

Next consider the payoff following a conviction, i.e. at h^t with $h_{t-1} = 1$. In this case, the regulator must be indifferent between inspect and wait. This indifference is necessary for the regulator to randomize – at h^t , it switches from inspect to wait with probability $1 - \beta^*$ (as noted above, this switching probability is conditional on the public history alone; the switching probability conditional on the regulator's private history will not be uniquely determined). Together with Lemma 2, this implies that the regulator's expected payoff from date t onwards is $V_I(h^t) = -S_{IW}$.

Finally, following h^t such that $\alpha(h^{t-1}) > 0$ and $h_{t-1} = 0$, the regulator switches from wait to inspect with positive probability (as explained in the main text). Thus, the regulator must be indifferent between switching from wait to inspect. This, together with Lemma 2, then

However, because the regulator strictly prefers to switch from wait to inspect, we have

$$-S_{WI} - i + (1 - \alpha(h^s)) \delta V_I(h^s, 0) + \alpha(h^s) (y + \delta V_I(h^s, 1)) > \delta V_W(h^s, 0).$$

Adding the previous expressions, we find $-S_{IW} - S_{WI} > 0$, in violation of Assumption A1.

implies that $V_I(h^t) = S_{WI}$. ■

Proof of Proposition 1. Follows from Lemmas 1-4 and the arguments in the text. ■

Proof of Corollary 1. If $h_{t-1} = 1$ the result is clear, given that the continuation play is unique in our model. So, consider a history with $h_{t-1} = 0$. The posterior of a firm at time t about the action of the regulator at time $t - 1$ is $\beta'_t \equiv \frac{(1-\alpha_t)\beta^*}{(1-\alpha_t)\beta^*+1-\beta^*}$. Suppose that the penalty changes from γ to γ_{new} at the beginning of date t , with all other parameters unchanged. Let $\beta_{new}^* \equiv \frac{\pi}{\pi+\gamma_{new}} < \beta^*$.

1. If $\gamma_{new} > \frac{1}{1-\alpha_{t-1}}\gamma$ then simple algebra shows that $\beta'_t > \beta_{new}^*$. Proposition 1 prescribes that $\beta_t = \beta_{new}^*$, so the regulator should be willing to switch to wait. As a result, $V_I(h^t) = -S_{IW}$, and equilibrium play enters the “residual deterrence phase” at date t .
2. If $\gamma_{new} < \frac{1}{1-\alpha_{t-1}}\gamma$ then $\beta'_t < \beta_{new}^*$. In this case, since $\beta_t = \beta_{new}^*$ by Proposition 1, the regulator should be willing to switch to inspect. This implies that $V_I(h^t) = S_{WI}$, so equilibrium must be in the stationary phase at date t .

■

Proof of Corollary 2. An increase in γ (or a decrease in π) decreases β^* and therefore the probability that the regulator inspects. The values α^* and α_0^* remain unchanged. Also, the duration T of the phase after a conviction during which firms abstain remains unchanged. These observations, together with (7), imply the result. ■

Proof of Corollary 3. Let $\vec{S} \equiv (S_{WI}, S_{IW})$. Let $\alpha_0^*(\vec{S})$, $\alpha^*(\vec{S})$, $\bar{\alpha}(\vec{S})$ and $T(\vec{S})$ denote, respectively, the corresponding (unique) values of α_0^* , α^* , $\bar{\alpha}$ and T for a given value of \vec{S} satisfying Assumptions A1-A3. Note from their expressions that $\alpha^*(\cdot)$ is continuous and increasing in both S_{IW} and S_{WI} , while $T(\cdot)$ and $\alpha_0^*(\cdot)$ are right-continuous in S_{IW} and S_{WI} .

Let's first show that $\bar{\alpha}(\cdot)$ is continuous. If $T(\cdot)$ is continuous at \vec{S} , then it is easy to see that $\alpha_0^*(\cdot)$ and $\alpha^*(\cdot)$ are also continuous at \vec{S} , so $\bar{\alpha}(\cdot)$ is continuous at \vec{S} . For a generic function f , let $f(\vec{S}^-)$ denote its limit from the left at \vec{S} .²¹ Note that if \vec{S} is such that

²¹By this we need the limit of any strictly increasing sequence $(\vec{S})_n$ converging to \vec{S} using the usual order in \mathbb{R}^2 , that is, $(x_1, y_1) > (x_2, y_2)$ if and only if $x_1 \geq x_2$, $y_1 \geq y_2$ and at least one of the inequalities is strict.

$T(\vec{S}) = T(\vec{S}^-) + 1$, then $\alpha_0^*(\vec{S}^-) = 0$, while $\alpha_0^*(\vec{S}) = \alpha^*$. Therefore

$$\begin{aligned}\bar{\alpha}(\vec{S}^-) &= \frac{\alpha^*(\vec{S}^-)}{1 + \beta^*((T(\vec{S}^-) + 1)\alpha^*(\vec{S}^-) - \alpha_0^*(\vec{S}^-))} \\ &= \frac{\alpha^*(\vec{S})}{1 + \beta^*T(\vec{S})\alpha^*(\vec{S})} \\ &= \frac{\alpha^*(\vec{S})}{1 + \beta^*((T(\vec{S}) + 1)\alpha^*(\vec{S}) - \alpha_0^*(\vec{S}))} \\ &= \bar{\alpha}(\vec{S}).\end{aligned}$$

Assume that \vec{S} is such that $\alpha_0^*(\vec{S}) \in (0, \alpha^*(\vec{S}))$. Then, there exists $\varepsilon > 0$ such that $T(\vec{S}')$ is constant for all \vec{S}' in an ε -neighborhood of \vec{S} . This implies that $\frac{\partial}{\partial S_w'} \bar{\alpha}(\vec{S}')|_{\vec{S}'=\vec{S}}$ can be computed holding T constant, for $w \in \{IW, WI\}$. In this case we have

$$\frac{\partial \bar{\alpha}(\vec{S})}{\partial S_{IW}} = A(i + (1 - \delta)S_{WI})(\delta^{T+1} - \beta^*), \quad (13)$$

$$\frac{\partial \bar{\alpha}(\vec{S})}{\partial S_{WI}} = A\left(i - S_{IW}(1 - \delta)\right)\left(\delta^{T+1} + \frac{\delta^T y}{\frac{i}{1-\delta} - S_{IW}} - \beta^*\right). \quad (14)$$

where $T \equiv T(\vec{S})$ and $A > 0$ is a (long) algebraic expression. Equation (13) implies Part 1 of the corollary: $\frac{\partial \bar{\alpha}(\vec{S})}{\partial S_{IW}} > 0$ if $\beta^* < \delta^{T+1}$ and $\frac{\partial \bar{\alpha}(\vec{S})}{\partial S_{WI}} < 0$ if $\beta^* > \delta^{T+1}$.

To see Part 2 of the corollary, note that, by Equation (6), $\delta^T \geq \frac{i - S_{IW}(1-\delta)}{i + S_{WI}(1-\delta)}$. Then, using this and Assumption A3, we have

$$\delta^{T+1} + \frac{\delta^T y (1 - \delta)}{i - S_{IW}(1 - \delta)} > 1,$$

so that $\frac{\partial \bar{\alpha}(\vec{S})}{\partial S_{WI}} > 0$ irrespective of $\beta^* \in (0, 1)$. The result then follows because $\bar{\alpha}$ is continuous in \vec{S} . ■

Proof of Proposition 2.

In order to prove Proposition 2, we first note that Lemmas 1-4 apply similarly in the random-payoffs case. Lemma 1 can be reformulated to: “For all equilibria, at all h^t , $\beta(h^t) < \bar{\beta}$.” The reader can check that all arguments in the proof remain the same in this case. As a consequence, Lemma 2 is also valid (since its proof only requires $\beta(h^t) < 1$ from Lemma 1). In a similar fashion, Lemma 3 can be rephrased as: “For all equilibria, for all h^t , $\beta(h^t) > \underline{\beta} \equiv \frac{\pi}{\pi + \gamma}$.” Finally, the first two claims in Lemma 4 trivially remain the same.

We next establish that the smallest probability of inspection at any public history is β^{\min} , attained when the continuation payoff from inspecting is S_{WI} .

Lemma 5 *In any equilibrium, $\beta(h^t) \geq \beta^{\min} \equiv \alpha^{-1}(\alpha^*)$ for all h^t , with equality when $V_I(h^t) = S_{WI}$.*

Proof. Given that $\beta_t \in (\underline{\beta}, \bar{\beta})$ for all t , after any history there is a time where, if there has still been no conviction, the regulator switches with positive probability from wait to inspect. That is, for each h^t , there exists a finite τ such that the regulator switches from wait to inspect at $(h^t, \mathbf{0}^\tau)$.

We can now prove the first part of our result by induction. Consider an equilibrium and a history h^t where the regulator switches with positive probability from wait to inspect at $t + \tau$ for the first time after t if there is no conviction at dates t up to $t + \tau - 1$. If $\tau = 1$ we have

$$\begin{aligned} V_I(h^t) &= -i + \alpha_t(y + \delta V_I(h^t, 1)) + (1 - \alpha_t)\delta V_I(h^t, 0) \\ &= -i + \alpha_t(y - \delta S_{IW}) + (1 - \alpha_t)\delta S_{WI}. \end{aligned}$$

(Here, $V_I(h^t, 0) = S_{WI}$ follows because (by assumption that $\tau = 1$) the regulator is willing to switch to inspect at $(h^t, 0)$, and because $V_W(h^t, 0) = 0$ (by the claim in Lemma 2).) Note that α^* solves the previous equation when $V_I(h^t) = S_{WI}$. If $V_I(h^t) < S_{WI}$, we have that $\alpha_t < \alpha^*$, that is, $\beta_t > \beta^{\min}$. Now consider $\tau > 1$. As is obvious from Bayes' rule (Equation (9)), when there is no switch from wait to inspect at dates t up to $t + \tau - 1$ and there are no convictions, $\beta_{t+s+1} < \beta_{t+s}$ for any $s \in \{0, \dots, \tau - 1\}$.²² This proves the first part of our result.

Now consider any h^t for which $V_I(h^t) = S_{WI}$. In this case, we have:

$$S_{WI} = -i + \alpha(h^t)(y - \delta S_{IW}) + (1 - \alpha(h^t))\delta V_I(h^t, 0).$$

Note that α^* solves this equation when $V_I(h^t, 0) = S_{WI}$. If $V_I(h^t, 0) < S_{WI}$, we have that $\alpha(h^t) > \alpha^*$, that is, $\beta(h^t) < \beta^{\min}$. However, this contradicts our previous argument, so $\beta(h^t) = \beta^{\min}$. ■

Lemma 5 establishes the following. Consider any history h^t such that the regulator switches with positive probability from wait to inspect at $t + \tau$ for the first time after t if there is no conviction at dates t up to $t + \tau - 1$. Then the regulator must switch with positive probability from wait to inspect at all histories $(h^t, \mathbf{0}^s)$ for all $s \geq \tau$ (at these histories, $\beta(h^t, \mathbf{0}^s) = \beta^{\min}$; and so the absence of switching would imply beliefs that fall below β^{\min}).

²²Equation (9) applies to histories where there is no switching and no convictions. One can appropriately modify the updating rule to account for switches from inspect to wait, concluding still that $\beta_{t+s+1} < \beta_{t+s}$ for any $s \in \{0, \dots, \tau - 1\}$.

Now consider any history h^t with $h_{t-1} = 1$. Given that $\beta_t \in (\underline{\beta}, \bar{\beta})$, the regulator must switch from inspect to wait at t , and hence $V_I(h^t) = -S_{IW}$. As argued above, there exists τ such that the regulator switches from wait to inspect for the first time at $t + \tau$ for $(h^t, \mathbf{0}^\tau)$. For $s = \tau - 1$, we have

$$V_I(h^t, \mathbf{0}^s) = -i + \alpha_{t+s}(y - \delta S_{IW}) + (1 - \alpha_{t+s})\delta S_{WI}$$

and $\alpha_{t+s} \leq \alpha^*$, implying that $V_I(h^t, \mathbf{0}^s) \leq S_{WI}$. Considering there are no switches from wait to inspect until $t + \tau$, and using updating, note that, for $s \in \{0, \dots, \tau - 2\}$, we have $\alpha_{t+s} = \alpha(\beta_{t+s}) < \alpha(\beta_{t+s+1}) = \alpha_{t+s+1}$. Similarly, we can show inductively that $V_I(h^t, \mathbf{0}^s) < V_I(h^t, \mathbf{0}^{s+1})$. It follows that $V_I(h^t) = -S_{IW}$ only if $h_{t-1} = 1$; i.e., switches from inspect to wait occur only immediately after a conviction.

The only part of the proposition left to prove pertains to the uniqueness of β^{\max} and T . Let's first define, for each $\beta^{\max} \geq \beta^{\min}$, the function $\tilde{V}_I(\beta^{\max})$ to be the expected payoff of the regulator following a conviction in a putative equilibrium in which the probability of inspecting following a conviction is β^{\max} . That is, let

$$\begin{aligned} \tilde{V}_I(\beta^{\max}) \equiv & \sum_{t=0}^{T(\beta^{\max})} A_t(\beta^{\max}) \left(-i + \alpha(\tilde{\beta}_t(\beta^{\max}))(y - \delta S_{IW}) \right) \delta^t \\ & + A_{T(\beta^{\max})}(\beta^{\max}) (1 - \alpha(\tilde{\beta}_{T(\beta^{\max})}(\beta^{\max}))) \delta^{T(\beta^{\max})+1} S_{WI} \end{aligned} \quad (15)$$

where $\tilde{\beta}_t(\beta^{\max})$ is the result of applying t times the Bayes' rule (using equation (9) with $\beta_0 = \beta^{\max}$), $T(\beta^{\max}) \equiv \min\{s | \tilde{\beta}_s(\beta^{\max}) < \beta^{\min}\} - 1$ indicates the number of periods that it takes for β to pass β^{\min} (less one), and $A_t(\beta^{\max}) \equiv \prod_{s=0}^{t-1} (1 - \alpha(\tilde{\beta}_s(\beta^{\max})))$ is the probability of reaching date t without any conviction (recalling that date 0 is the first date following a conviction).

Given that, for each t , $\tilde{\beta}_t(\beta^{\max})$ is increasing in β^{\max} , each $\alpha(\tilde{\beta}_t(\beta^{\max}))$ is decreasing in β^{\max} . It is thus easy to see that $\tilde{V}_I(\cdot)$ is strictly decreasing. It is also easy to see that $\tilde{V}_I(\cdot)$ is continuous. Finally, given that $\lim_{\beta^{\max} \rightarrow \bar{\beta}} \tilde{V}_I(\beta^{\max}) = -\frac{i}{1-\delta} < -S_{IW}$ and $\lim_{\beta^{\max} \rightarrow \beta^{\min}} \tilde{V}_I(\beta^{\max}) = S_{WI}$, there is a unique β^{\max} which guarantees the indifference condition $\tilde{V}_I(\beta^{\max}) = -S_{IW}$. ■