

Residual Deterrence

Francesc Dilme
University of Bonn
fdilme@uni-bonn.edu

Daniel F. Garrett
Toulouse School of Economics
dfgarrett@gmail.com

September 2014

Abstract

We study a dynamic version of the canonical inspection game in which the inspector incurs a cost when switching between “inspect” and “not inspect”. Myopic potential offenders learn about the past activity of the inspector via publicly observed convictions. Equilibrium involves reputation cycles, with each conviction followed by a phase of “residual deterrence” during which the rate of offending is lower. We show that (due to the equilibrium response of the inspector) the long-run average crime rate is increasing in the penalty imposed on a convicted offender. It may be increasing, decreasing or hump-shaped in the inspector’s cost of switching. In several extensions, residual deterrence is examined under different information and payoff structures.

JEL classification: C73, K42

Keywords: deterrence, reputation, switching costs

PRELIMINARY AND INCOMPLETE - COMMENTS WELCOME

1 Introduction

An important rationale for the enforcement of laws and regulations concerns the deterrence of the 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 deters “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. 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 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 and potential offenders.

We consider a dynamic version of a simple workhorse model – the inspection game. In this model, a long lived inspector faces a sequence of short-lived potential offenders. Committing an offense is only worthwhile for an offender if the inspector is “off duty”, while being on duty is only worthwhile for the inspector if an offense is committed. Potential offenders only observe the previous history of “convictions”; that is, the periods where the inspector was inspecting and the potential offender 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 potential offenders learn that the inspector 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 aforementioned setting, equilibrium follows a repetition of static play; i.e., past convictions do not affect the rate of offending. Things are different once we introduce to this setting an often overlooked friction, the cost of reallocating resources.¹ We show that the unique equilibrium public

¹There are several reasons why an enforcement agency may face additional costs when switching between activities. In the context of price fixing, the enforcement agency may expend costs to come up to date with events in a given market, as may be necessary for it to identify delinquent behavior. Thus actively monitoring a market over two periods may be less costly if these periods are consecutive; the information learned about the market in the first period remains relevant in the second. There may also be physical costs to reallocating resources, as for police reassigning their focus

outcome then features reputational effects driven by the switching costs: a conviction is followed by several periods during which the agents 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 show that the switching costs are necessary to generate residual deterrence, since the episodes of residual deterrence disappear as switching costs shrink.

It is then of interest to examine the average rate of offending, taken over a long horizon (we term this the “long-run average offense rate”). There are two intriguing comparative statics results. First, the average rate of offending is increasing in the penalties faced by a convicted offender. This perhaps counterintuitive result is a consequence of the equilibrium behavior of the inspector. As in the static inspection game, the inspector responds to an increase in penalties by inspecting less often (in order that offenders are still willing to offend). However, this means fewer convictions, and hence fewer reputation cycles and fewer episodes of residual deterrence. Taken by itself, our result suggests that a planner who is concerned alone with the long-run average crime rate, and who cannot directly influence the incentives of the inspector, would favor lower rather than higher penalties. This suggests a possible answer to Becker’s (1968) puzzle as to why penalties seem rarely to be set at their maximal level. (Of course, our analysis is “partial equilibrium” in the sense that it focuses on enforcement of a single activity. In practice, a planner should account for the range of activities that inspectors undertake.)

A second comparative statics result is the effect of the switching cost on the offense rate. For switching costs in the relevant region (the region where reputation cycles occur), the crime rate may be either increasing, decreasing or hump-shaped in the switching cost. The reason is that, on the one hand, the duration of residual deterrence is increasing in the switching cost, while on the other the rate of offending after residual deterrence has ceased is also increasing in this cost. (We provide the intuition for these two observations below.) Which effect dominates depends on the parameters of the model.

Before presenting the model, it is worth clarifying up front a few important modeling choices. First, our baseline model posits that potential offenders have identical preferences for offending. This simplifies the analysis and leads to the stark conclusion that the potential offenders’ 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 the beliefs of potential offenders. However, we generalize by allowing for ex-ante identical offenders to have heterogeneous (and privately observed) preferences for offending. In this case, the beliefs of potential offenders do fluctuate in a predictable fashion over the cycle: the perceived probability of detection following a conviction is high, while this

from one kind of offense to another (say drink driving to speed infractions), or from one location to another.

probability falls in the absence of convictions in accordance with Bayes' rule. That is, in the absence of convictions, potential offenders place increasing weight on the possibility that the inspector has switched to being off duty. 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 the inspector'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". Such incentives can exist notwithstanding that the mission of an organization itself may be achieving deterrence (for instance, consider the Security and Exchange Commission's directive from Congress to "achieve an appropriate level of deterrence in each case and thereby maximize the remedial effects of its enforcement actions" (H.R. Rep. No. 101-616, at 13, 1990)). Nonetheless, we are also able to extend our baseline model to settings where the inspector is concerned directly with deterrence, rather than convictions. Again, we exhibit equilibria featuring reputation cycles. The inspector in this case may be incentivized to inspect precisely because she anticipates residual deterrence following a conviction.

The rest of the paper is as follows. We next briefly review the economics literature on 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 the extensions highlighted above (further extensions can be found in Appendix B at the end of the document). Section 5 concludes.

1.1 Literature Review

At least since Becker (1968), economists have been interested in the deterrence role of policing and enforcement. Applications include not only criminal or delinquent behavior, but also the regulated behavior of firms such as environmental emissions, health and safety standards and anticompetitive practices. This work typically simplifies the analysis by adopting a static framework with full commitment to the policing strategy. The focus has then often been on deriving the optimal policies to which governments, regulators, police or contracting parties should commit (see, among others, Becker (1968), Townsend (1979), Polinsky and Shavell (1984), Reinganum and Wilde (1985), Mookherjee and Png (1989, 1994), 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 reputational effects are, however, absent from the literature. Block, Nold and Sidak (1981) do informally suggest a possible 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 to which peer firms are submitted.² Relative to Block, Nold and Sidak, our theory allows the inspector to choose its enforcement policy strategically over time.³

Our paper is related to the literature on reputations with endogenously switching types; see for instance Mailath and Samuelson (2001), Iossa and Rey (2012), Board and Meyer-ter-Vehn (2013, 2014) and Halac and Prat (2014). Closest methodologically to our paper is the work by Dilmé (2012). 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),⁴ but introduces a switching cost to change the quality level. The stage game in this paper is different to Dilmé’s, however, requiring a separate analysis.

Closest to our paper in terms of motivation is recent (and independent) work by Halac and Prat (2014). There, a manager builds his reputation for managerial feedback. The so-called “bad news” case is closest to our model, where the manager provides feedback and penalizes the worker if the worker shirks. Workers exert effort immediately after being caught shirking, a possibility that we relate in this paper to the notion of “residual deterrence”. However, there are important differences. First, the technology of the manager is different to our paper: the manager’s choice is the extent of a continuous investment in the monitoring technology. Instead, the inspector of our model either inspects or actively chooses not to do so. Second, the manager’s motivation in our baseline model is different, since the manager is directly rewarded for convictions (reflecting the inspector’s own career concerns). Finally, our paper has a very different focus; our comparative statics results regarding the average crime rate, for instance, are entirely new.

²They 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.”

³ Other dynamic models of deterrence seem scarce. See, however, Lui (1986).

⁴Mailath and Samuelson (2001) endogenize the replacements in Section 5 of their paper. Given that they focus on equilibria where the competent firm exert high effort always, the resulting model can be interpreted in a similar way as Board and Meyer-ter-Vehn (2013).

2 Baseline Model

Timing, players and actions. Time is discrete and infinite. For our baseline model, there is a single inspector and a sequence of short-lived agents who are potential offenders, one per period.

In each period $t \geq 0$, the inspector 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 agent 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 agent offends while the inspector inspects at date t , while $a_t b_t = 0$ otherwise. If $a_t b_t = 1$, we say that the inspector “obtains a conviction” at date t .

Payoffs. Per-period payoffs are determined at first instance according to a standard inspection game. If the agent offends without a conviction ($a_t = O$ and $b_t = W$), then he earns a payoff $\pi > 0$. If he offends and is convicted ($a_t = O$ and $b_t = I$), then he sustains a cost $\gamma > 0$, which is net of any benefits from the offense. Otherwise, his payoff is zero.

If the inspector inspects at date t she suffers a cost $i > 0$. She incurs no cost if waiting. In the event of obtaining a conviction, the inspector earns an additional payoff of $y > i$. Later, we consider the possibility that the inspector cares about deterring the agent rather than convicting him.⁵

In addition to the costs and benefits specified above, the inspector sustains a cost $S > 0$ if switching action at period t .⁶ Hence, the switching cost in period t is $S\mathbf{1}_t^S(b_{t-1}, b_t)$, where $\mathbf{1}_t^S(b_{t-1}, b_t)$ indicates a switch at date t (either $b_{t-1} = I$ and $b_t = W$ or $b_{t-1} = W$ and $b_t = I$). Payoffs are then summarized in the following table.

	N	O
W	$-S\mathbf{1}_t^S(I, W), 0$	$-S\mathbf{1}_t^S(I, W), \pi$
I	$-S\mathbf{1}_t^S(W, I) - i, 0$	$y - S\mathbf{1}_t^S(W, I) - i, -\gamma$

Because changes in the inspector’s actions affect payoffs, it is necessary to specify the inspector’s action in the period before the game begins. For concreteness we let $b_{-1} = I$, although no results hinge on this assumption.

The inspector discounts the future at rate $\delta \in (0, 1)$, while each agent is short-lived and hence myopic. That agents are short-lived excludes as a motivation for offending possible learning about the choices of the inspector.

⁵Note that, as is well known, there is a unique equilibrium of the stage game without switching costs. In this equilibrium, the inspector chooses I with probability $\pi/(\pi + \gamma)$, while the agent chooses O with probability i/y . These probabilities ensure the agent is indifferent between the two actions (W and I for the inspector and N and O for the agent).

⁶The assumption that the switching cost is symmetric (i.e., the same whether switching to or from “inspect”) is made without any loss of generality. Indeed, proceeding similarly to Dilmé (2012), one can show that, for general asymmetric switching costs, the value functions and flow payoffs can be renormalized so that switching costs are symmetric. Our key qualitative findings thus apply also to a model with asymmetric switching costs.

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

As discussed above, part of the novelty in our modeling is that information concerning the inspector's activity depends on actions taken also by the short-lived agents. This is distinct from the existing literature such as Board and Meyer-ter-Vehn (2013) and Dilme (2012) (although see Halac and Prat (2014)).⁷

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

Strategies, equilibrium and continuation payoffs. We let the strategy of a date- t agent 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 agent offends (at date t). A (behavioral) strategy for the inspector assigns to each private history $\tilde{h}^t \in \{0, 1\}^t \times \{I, W\}^t$ the probability that the inspector inspects at \tilde{h}^t , $\beta(\tilde{h}^t)$. We study perfect Bayesian equilibria (PBE) of the above game.

For a fixed strategy β of the inspector, we find the following abuse of notation convenient. For each public history h^t , $\beta(h^t) \equiv \mathbb{E}[\beta(\tilde{h}^t) | h^t]$, where the expectation is taken with respect to the distribution over private histories \tilde{h}^t with public component h^t , as determined according to the strategy β . 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 agent'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 inspector, 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 inspector at

⁷In these papers, in the so-called "good news" case, signals are generated with positive probability whenever the seller provides high quality. One interpretation is that these signals are exogenous to buyer decisions. Another is that they require the buyer to purchase, but buyers purchase with probability 1 at all dates in equilibrium. Which interpretation one prefers is thus irrelevant for outcomes.

any date t and for any strategies of the agent and inspector. For an inspector history \tilde{h}_t , this is

$$V_t(\beta, \alpha; \tilde{h}^t) = \mathbb{E}_{\beta, \alpha} \left[\sum_{s=t}^{\infty} \delta^{s-t} \left(y \tilde{b}_s \tilde{a}_s - i \tilde{b}_s - S \mathbf{1}_s^S \left(\tilde{b}_{s-1}, \tilde{b}_s \right) \right) | \tilde{h}_t \right].$$

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

3 Equilibrium Characterization

Since agents are myopic, an agent offends at date t only if $(1 - \beta_t)\pi - \beta_t\gamma \geq 0$. Let $\beta^* = \frac{\pi}{\pi + \gamma}$ be the belief which keeps an agent indifferent between offending and not. We restrict attention to parameters such that equilibrium involves infinitely repeated switching, as described in the Introduction.

Assumption 1: The inspector has incentives to switch to wait if no offending occurs in the future, i.e., $S < \frac{i}{1-\delta}$.⁸

Assumption 2: The inspector has incentives to switch to inspect if the agent offends for sure, i.e., $S < \frac{y-i}{1+\delta}$.⁹

We begin by showing the following result.

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

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

The above property of equilibrium makes intuitive sense. Indeed, otherwise an agent would find abstaining from offense strictly preferable in some period t , i.e. $\beta_t > \beta^*$. Nevertheless, since by

⁸Otherwise the inspector inspects at all dates in the unique equilibrium.

⁹Otherwise, one can show that the agent does not switch from wait to inspect in equilibrium.

Assumption 1 there is no equilibrium continuation play where no crime is ever committed again, we have $\beta_{t+s} \leq \beta^*$ for some $s > 0$, so the inspector would switch back to wait at some point. Also using Assumption 1 it is easy to show that she wants to do this as soon as possible, so $\beta_t \leq \beta^*$. Next, we use this result to provide an important observation concerning the inspector's payoffs.

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

Proof. Assume otherwise, so there is a history such that $h_{t-1} = 0$ and $V_W(h^t) > 0$. By assumption (i.e., by definition of V_W), the inspector waited in period $t-1$. The fact that $V_W(h^t) > 0$ implies that, after this history, the inspector strictly prefers to switch from wait to inspect at some date $s \geq t$ after a public history $h^s = (h^t, \mathbf{0}^{s-t})$. In particular, inspecting is strictly preferred to waiting at h^s irrespective of the inspector's private actions b^s . Hence, we must have $\beta(h^s) = 1 > \beta^*$, contradicting the finding in Lemma 1. ■

Thus, an inspector who waits in a given period weakly prefers to wait in all subsequent periods. Put another way, the inspector must never have strict incentives to switch to inspect, since such incentives are incompatible with the agents' choice to offend. A positive probability of offending is of course essential for the inspector to profit from switching to inspect.

An important consequence of Lemma 2 is the following. In equilibrium, the inspector's continuation payoff, net of present-period switching costs, is equal to zero whenever she plays the action wait (irrespective of her action in the previous period). This follows because the inspector receives zero (net of switching costs) in the period she plays wait, while her continuation payoff for the following period is also zero by Lemma 2. We use this to show the following result.

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

Proof. 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 agent offends at date t . There must exist a private history for the inspector $\tilde{h}^t = (h^t; b^t)$ such that she waits at date t . As explained above, her continuation payoff from date t onwards, net of any date- t switching costs, must equal zero. However, by instead inspecting at date t and then waiting at $t+1$ the inspector can guarantee herself a payoff at least $y - i - S(1 + \delta)$, which is strictly positive by Assumption 2. ■

Lemmas 1 and 3 together imply that agents are necessarily indifferent between offending and not, i.e. $\beta_t = \beta^*$ for all t . The indifference of the agent is analogous to the finding for the stage game without switching costs (where the inspector inspects with probability β^*). As we find in Section 4, the indifference property is particular to a model where agents are homogeneous in their payoffs and penalties for offending. A more realistic model which allows for agent heterogeneity not only yields agents 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 agents are indifferent to offending, our task is to find any collection of offense probabilities $\alpha(h^t)$ which yield the optimality of an inspector 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 inspector $V_I(h^t)$ following an inspection, as summarized in the following lemma.

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

Proof. Suppose first that $V_I(h^t) > S$ for some h^t . Given Lemma 2, the inspector must have strict incentives to switch from wait to inspect following h^t in case $b_{t-1} = W$, or 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, $V_I(h^t) < -S$ 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 inspector must be indifferent between inspect and wait. This indifference is necessary for the inspector to randomize – at h^t , she 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 inspector’s private history will not be uniquely determined). Together with Lemma 2, this implies that the inspector’s expected payoff from date t onwards is $V_I(h^t) = -S$.

Finally, following h^t such that $\alpha(h^{t-1}) > 0$ and $h_{t-1} = 0$, the inspector switches from wait to inspect with positive probability. This is necessary for agent date- t beliefs to remain at β^* . Indeed, after calculating the agent’s posterior belief of inspection following $h_{t-1} = 0$, we find that the probability of switching must be ξ satisfying

$$\frac{\beta^* (1 - \alpha(h^{t-1}))}{\beta^* (1 - \alpha(h^{t-1})) + 1 - \beta^*} + \xi \frac{1 - \beta^*}{\beta^* (1 - \alpha(h^{t-1})) + 1 - \beta^*} = \beta^*. \quad (1)$$

Thus, the inspector must be indifferent between switching from wait to inspect. This, together with Lemma 2, then implies that $V_I(h^t) = S$. ■

Lemma 4 follows after noticing that the inspector cannot have strict incentives to switch, either from inspect to wait or wait to inspect. The inspector is willing to switch to wait after a conviction, and to inspect at a history h^t such that $h_{t-1} = 0$ and $\alpha(h^{t-1}) > 0$.

We can now determine the agents’ strategies. Consider h^t such that $\alpha(h^{t-1}) > 0$ and $h_{t-1} = 0$. That the inspector 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 inspector inspects. Using Lemma 4, we have equivalently

$$S = -i + (1 - \alpha(h^t)) \delta S + \alpha(h^t) (y - \delta S). \quad (3)$$

Thus, we must have $\alpha(h^t) = \alpha^*$, where $\alpha^* = \frac{i+S(1-\delta)}{y-2\delta S} \in (0, 1)$ is the value solving (3).

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

$$\begin{aligned} S &\geq -i + (1 - \alpha(h^t)) \delta S + \alpha(h^t) (y - \delta S) \\ &= -i + \delta S + \alpha(h^t)(y - 2\delta S). \end{aligned} \quad (4)$$

Given Assumption 2, $y - 2\delta S > 0$, 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 agent offends with positive probability is $t + T$ with $T \geq 0$. With an abuse of notation, denote this probability by $\alpha_T = \alpha(h^{t-1}, 1, \mathbf{0}^T)$. Since the inspector must be willing to continue inspecting (she never strictly prefers to switch to wait), we must have

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

Equivalently, given Lemma 4, we have

$$-S = - \sum_{j=0}^T \delta^j i + (1 - \alpha_T) \delta^{T+1} S + \alpha_T \delta^T (y - \delta S). \quad (5)$$

It is straightforward to verify that there is a unique solution solution to this equation such that $\alpha_T \in (0, \alpha^*]$ and T is integer-valued. In particular, it is easy to show that we have

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

We now give our equilibrium characterization.

Proposition 1 *Under Assumptions 1 and 2, any equilibrium unfolds as follows.*

1. At date 0, and following h^t such that $h_{t-1} = 1$, the inspector switches with probability $1 - \beta^*$ to wait, where this switching probability is conditional on the public history h^t .
2. If $T \geq 1$, the inspector does not switch from date $t + 1$ up to date $t + T$. The agent does not offend from from t up to $t + T - 1$.

3. At $t+T$, the agent randomizes, playing O with probability $\alpha_T \in (0, \alpha^*]$. If the agent is convicted, play returns to Step 1, otherwise it continues.
4. At each date $s > t + T$ such that there has been no further conviction, if the inspector waits at $s - 1$, then she switches to inspect at date s with probability ξ given by (1) (where again this probability is conditional on the public history). The agent randomizes, playing O with probability α^* . In the event of a conviction, play returns to Step 1.

Equilibrium behavior exhibits the properties described in the Introduction. Following a conviction, offending ceases for T periods. This can be understood as the “residual deterrence” phase. Offending then resumes with a lower probability α_T for one period, and then resumes to its highest level α^* . The 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^* . Below, we will refer to the periods following deterrence decay as the “stationary phase”.

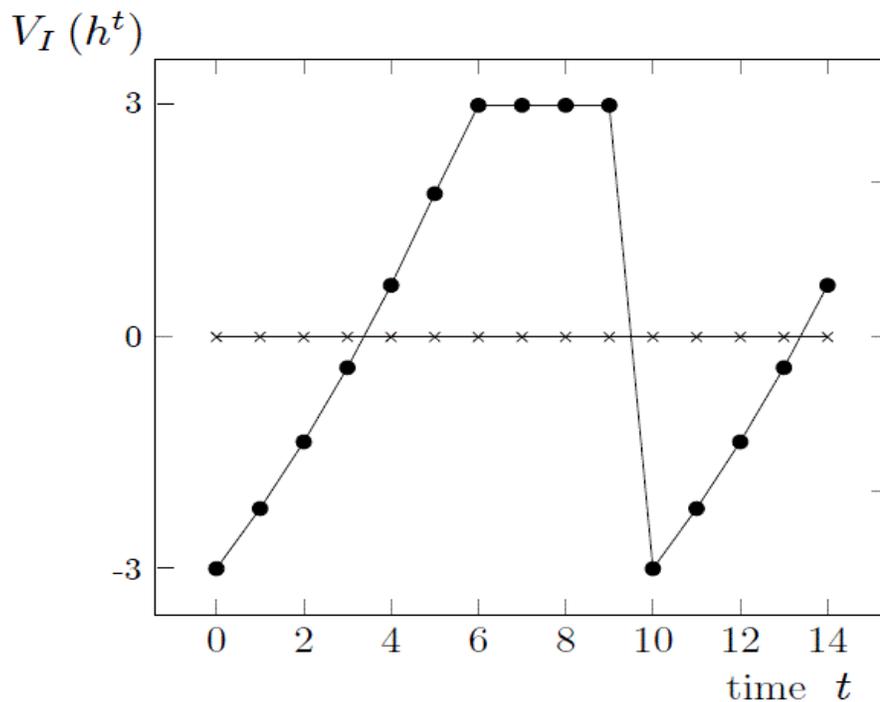


Figure 1: Inspector continuation payoff following inspection: $\delta = 0.9$, $y = 6$, $S = 3$, $i = 1$ (hence one calculates $T = 5$)

The evolution of the inspector’s incentives in equilibrium can be understood by considering how the continuation payoff $V_I(h^t)$ changes over time. Figure 1 plots this continuation payoff for an

example. Following a conviction, which occurs at $t = -1, 10$ in the example, the payoff from inspecting is equal to $-S$, and the inspector is thus indifferent between inspecting and switching to wait. Given that $T = 5$ 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 $V_I(h^t) \in (-S, S)$, the inspector strictly prefers not to switch irrespective of whether she 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 $V_I(h^t) = S$, and the inspector is indifferent between remaining at wait and switching to inspect.

3.1 Comparative Statics

We next turn to an analysis of comparative statics. Our focus is on 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 $T + 1 + \frac{1 - \alpha_T \beta^*}{\alpha^* \beta^*}$. The long-run average offense rate is then

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

Our first result is the following.

Corollary 1 *The long-run average offense rate is increasing in the penalty γ and decreasing in π .*

While Corollary 1 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 an agent is indifferent to offending. Provided the inspector's mixing probabilities are adjusted appropriately to maintain agent beliefs at the lower level, agents are content to play the same strategies specified in Proposition 1 to maintain the inspector'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_T) and subsequently (α^*). Nevertheless, due to the lower value of β^* , convictions occur less frequently. Episodes of residual deterrence are thus less frequent.

Corollary 1 provides a new answer to an old question raised by Becker (1968) regarding why maximal penalties may not be optimal.¹⁰ In our model, a planner who is concerned only with the rate of offending, gains by reducing the penalty. As noted in the Introduction, one must be careful to interpret this result in light of the partial equilibrium nature of our analysis. For instance, a planner may be concerned also with the opportunity cost of the inspector's time. Consider an inspector

¹⁰An 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).

who must be on duty at all periods, but who can monitor activities at several different locations (or, equivalently, who can inspect for several different violations among groups of unrelated offenders). Our intuition suggests that raising the penalty at a given location *raises* the rate of offending at that location. However, it seems reasonable to conjecture that a planner concerned with the average crime rate at all locations should set penalties uniformly high. We leave the analysis of such a “multi-tasking” model to future work.

Our finding should also be related to the result for a static setting. Tsebelis (1989), for instance, points out 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. 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 $S = 0$, then equilibrium play in the repeated game simply involves a repetition of the static inspection and offense rates. It is pertinent to note that this is also the limiting behavior as $S \searrow 0$.

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

Corollary 2 *The long-run average offense rate $\bar{\alpha}$ is a continuous function of S over $S \in [0, \min(\frac{y-i}{1+\delta}, \frac{i}{1-\delta})]$. There are two thresholds, $0 \leq \underline{\beta}^* \leq \bar{\beta}^* \leq 1$ such that $\bar{\alpha}$ is a decreasing function of S if $\beta^* \geq \bar{\beta}^*$ (in this case $T = 0$ for all S), hump-shaped if $\beta^* \in (\underline{\beta}^*, \bar{\beta}^*)$ and increasing if $\beta^* \leq \underline{\beta}^*$.¹¹*

Two key variables which determine the long-run average offense rate are (i) the length of the period in which agents do not offend following a conviction, T , and (ii) the rate α^* of offending once offenses resume. Both are increasing in S . To see why T increases in S , consider (5). Following a conviction, the inspector must face sufficient incentives to switch to wait: when S is large, this necessitates a longer horizon over which the agents do not offend. To see why α^* is increasing in S , we note from (3) that α^* must be large enough to incentivize the inspector to switch to inspect after no conviction; such incentives must be larger when S is larger. Which effect dominates as S increases (the increase in T or in α^*) depends on parameters.

Corollary 2 indicates that a positive switching cost may increase or decrease the offense rate relative to the rate for $S = 0$. 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, a higher switching cost raises the equilibrium rate of offending in the stationary phase, i.e. the phase after residual deterrence has passed. Nonetheless, in case the inspector’s reward y from a conviction is sufficiently large (in particular, if $\frac{y-i}{1+\delta} > \frac{i}{1-\delta}$), $S \nearrow \frac{i}{1-\delta}$ implies $T \rightarrow +\infty$, driving the long-run average offense rate to zero.¹²

¹¹It is easily checked that any of the three configurations are possible.

¹²We do not examine the case $S \geq \frac{i}{1-\delta}$. In this case, given the inspector’s date $t = -1$ action is to inspect, the inspector continues inspecting forever in equilibrium. However, note that the outcome is different if $S > \max\left\{\frac{i}{1-\delta}, \frac{y-i}{1+\delta}\right\}$ and

4 Extensions

4.1 Random Agent Payoffs

We now permit the inspector to face uncertainty concerning the agents' payoffs. Thus, at the beginning of each period t , the agent independently 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 1 and 2 precisely as in Section 3.

At a given period t , if the probability that the inspector inspects (conditional on the public history) is β_t , an agent 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 Lemma 1, it is easy to see that the inspector must have (weak) incentives to switch to wait. As before, if the inspector is thought to be inspecting, she prefers to switch to wait sooner than later, in order to avoid inspecting costs. Also as before, the incentive to switch to wait must result from some periods of deterrence; that is, periods with a low crime rate in which the inspector does not change her action. Nevertheless, different from our base model, the crime rate is positive in the deterrence phase. This means that beliefs are updated following the absence of a conviction (since the absence of a conviction is now informative about the inspector's actions). Hence, during the deterrence phase, the absence of convictions make the offenders progressively convinced that the inspector is not inspecting. This happens until the crime rate is high enough that the inspector is (weakly) willing to switch from waiting to inspect. The crime rate then remains fixed until a new conviction occurs.

Proposition 2 *There exists a unique $\beta^{\max} \in (0, \bar{\beta})$ and $T \in \mathbb{N} \cup \{0\}$ such that, in any equilibrium:*

1. *At date 0, and following h^t such that $h_{t-1} = 1$, the inspector switches with some probability $1 - \beta^{\max}$ to wait, where this switching probability is conditional on the public history h^t .*
2. *Subsequently, provided there is no conviction, the inspector does not switch for some number of periods $T \geq 0$. In these periods, the agents' beliefs evolve according to Bayes' rule. That is, $\beta_t = \beta^{\max}$ and, for $j \in \{0, \dots, T - 1\}$,*

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

the inspector's $t = -1$ action is to wait. For these parameters, there is an equilibrium in which the inspector continues to wait. This shows that, at least for $S > \max\left\{\frac{i}{1-\delta}, \frac{y-i}{1+\delta}\right\}$, the long-run average offense rate may depend on the inspector's initial action. Our focus on $S < \min\left\{\frac{i}{1-\delta}, \frac{y-i}{1+\delta}\right\}$ ensures this is not the case.

3. At each date $s > t + T$ such that there has been no further conviction, if the inspector waits at $s - 1$, then she switches to inspect at date s with probability such that $\beta_s = \beta^{\min} \equiv \alpha^{-1}(\alpha^*)$ (with $\alpha^* = \frac{i+S(1-\delta)}{y-2\delta S}$ as in our baseline model).

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 inspector's continuation payoff following wait is zero. This must be the case, since the inspector 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 agent 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 $-S$) and during the stationary phase which follows the deterrence phase (equal to S). 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 $\frac{i+S(1-\delta)}{y-2\delta S}$), 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 a value $V_I(h^t)$ approaching $\frac{-i}{1-\delta} < -S$. 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 . Thus, it is easy to see that there is an intermediate value of β^{\max} such that $V_I(h^t) = -S$, and thus such that the inspector is indifferent between continuing to inspect and switching to wait.

Next, we would like to understand how equilibrium behavior compares to that in our baseline model. In particular, does equilibrium behavior in the random payoff model approach that in the baseline model as payoff uncertainty becomes small? To answer this question, fix a baseline distribution F with support on $[\underline{\pi}, \bar{\pi}]$. Define $G(\pi; \kappa) = F\left(\underline{\pi} + \frac{\bar{\pi} - \underline{\pi}}{\kappa - \underline{\pi}}(\pi - \underline{\pi})\right)$, which has support on $[\underline{\pi}, \kappa]$. Thus, as κ approaches $\underline{\pi}$, the model approaches the baseline case with $\pi = \underline{\pi}$. We show the following:

Corollary 3 *Consider the distribution over rewards for offending π given by $G(\pi; \kappa)$. As $\kappa \searrow \underline{\pi}$, the probability of an offense at h^t , $\alpha(\beta(h^t))$, converges to its value for the model with deterministic payoffs, i.e. for $\kappa = \underline{\pi}$, as given in Proposition 1.*

The corollary shows a form of robustness of the equilibrium described in Proposition 1. In particular, the indifference of agents to offending is not crucial. In the model with random preferences for offending, agents have a strict preference for their chosen action with probability one and have pure equilibrium strategies. However, the limiting offense probabilities and inspection probabilities are the same as for the baseline model, where agents are always indifferent and play mixed strategies.

A further notable feature of equilibrium with random payoffs is that agent beliefs $\beta(h^t, \mathbf{0}^{s-t})$, $s \geq t$, evolve in a predictable fashion following a conviction at $t - 1$. In particular, agents correctly perceive the risk of apprehension as smaller the longer since the last conviction. Thus “deterrence decay” here coincides with a steady reduction in the actual risk of being caught.

4.2 Directly, if not imperfectly, observed inspections

To date we assumed that inspections become public only when the short-lived agent is offending. However, inspection activities are sometimes communicated directly, irrespective of convictions. Examples include patrols of street crimes, where police presence is apparent to observers independent of possible offending. To capture this possibility, we suppose that inspections become public information with some probability $\rho \in (0, 1]$, irrespective of agent actions. That is, if $b_t = 1$, then $h_t = 1$ with probability ρ ; otherwise, $h_t = 0$. We let $h^t = (h_0, \dots, h_{t-1}) \in \{0, 1\}^t$ as above. When $\rho = 1$ then the actions of the inspector are perfectly observable.

To ensure existence of an equilibrium with switching, we replace Assumption 1 above with a more restrictive condition, while maintaining Assumption 2.

Assumption 1a: The switching cost S is not too large: $S < \frac{i}{1+\delta(1-2\rho)}$.

Note that Assumption 1a coincides with Assumption 1 in case $\rho = 1$.

Even though the information structure of a model with imperfectly observable inspections is quite different to the one of our base model, our preliminary results still hold. We begin with an analogue to Lemma 2 above, which requires now a somewhat different proof (see the Appendix).

Lemma 5 *For all h^t with $h_{t-1} = 0$, $V_W(h^t) = 0$.*

Next we provide an analogue to Lemmas 1 and 3.

Lemma 6 *For all h^t , $\beta_t = \beta^*$.*

Finally, we have an analogue to Lemma 4.

Lemma 7 *For all equilibria, for all h^t , $V_I(h^t) \in [-S, S]$. If $h_{t-1} = 1$, then $V_I(h^t) = -S$. If $h_{t-1} = 0$, then $V_I(h^t) = S$.¹³*

Note that the only difference between Lemma 7 and Lemma 4 is that $V_I(h^t) = S$ following *any* history with $h_{t-1} = 0$; i.e., we do not require $\alpha(h^{t-1}) > 0$. The payoff S is the reward required for the inspector to switch from wait to inspect. The reason the inspector must switch to inspect with positive probability after *any* history with $h_{t-1} = 0$ is that h_{t-1} is determined independently of

¹³If $\rho = 1$, then for $h_{t-1} = 0$, $V_I(h^t)$ involves an abuse of notation. Here, we mean the inspector’s expected payoff following h^t after a switch has been made to inspect, but excluding the cost of this switch, and before the realization of the agent’s date- t action.

agent offending; if the inspector did not switch, agent beliefs necessarily fall below β^* in accordance with Bayes' rule (independent of the value of $\alpha(h^{t-1})$).

We can now determine the agents' strategies. Consider h^t with $h_{t-1} = 0$. Lemma 7 implies that the inspector is willing to inspect at t and then follow an optimal continuation strategy, so that

$$\begin{aligned} V_I(h^t) &= -i + \alpha(h^t)y + \delta \left((1 - \rho) V_I(h^t, 0) + \rho V_I(h^t, 1) \right) \\ \Leftrightarrow S &= -i + \alpha(h^t)y + \delta \left((1 - 2\rho) S \right). \end{aligned}$$

If instead $h_t = 1$, we have

$$\begin{aligned} V_I(h^t) &= -i + \alpha(h^t)y + \delta \left((1 - \rho) V_I(h^t, 0) + \rho V_I(h^t, 1) \right) \\ \Leftrightarrow -S &= -i + \alpha(h^t)y + \delta \left((1 - 2\rho) S \right). \end{aligned}$$

We therefore find that, for all h^{t-1} ,

$$\alpha(h^{t-1}, 0) = \frac{S(1 - \delta(1 - 2\rho)) + i}{y} \quad (10)$$

$$\alpha(h^{t-1}, 1) = \frac{-S(1 + \delta(1 - 2\rho)) + i}{y}. \quad (11)$$

With an abuse of notation, let $\alpha_W \equiv \alpha(h^{t-1}, 0)$ and $\alpha_I \equiv \alpha(h^{t-1}, 1)$. Note that $\alpha_W > \alpha_I$. We have $\alpha_W < 1$ by Assumption 2 and $\alpha_I > 0$ by Assumption 1a. We can summarize these observations as follows.

Proposition 3 *Suppose that any inspection becomes public information with probability ρ , and impose Assumption 1a and Assumption 2. Then equilibrium unfolds as follows.*

1. *At date 0, and following h^t such that $h_{t-1} = 1$, the inspector switches to wait with probability $1 - \beta^*$. The agent offends with probability α_I .*
2. *Following any h^t with $h_{t-1} = 0$, the inspector switches to inspect with positive probability (such that the updated probability of inspection is β^*). The agent offends with probability α_W .*

Our results in this subsection highlight the importance of the information structure for our residual deterrence result. When information about the inspector's actions are generated independently of the decision to offend, residual deterrence is limited. Following an inspection, offending continues, albeit at the lower rate given in (11) before resuming at rate (10) in the absence of a signal indicating inspection.

The reason for the difference in results is that in the present case, as noted above, the inspector must switch to inspect with positive probability at *any* date t following $h_{t-1} = 0$ (compare this to the baseline model, where, if $\alpha_{t-1} = 0$, β_t remains equal to β^* following $h_{t-1} = 0$ although the inspector does not switch). However, for the inspector to be willing to switch to inspect at t requires a positive probability of offending. That is, there must be a positive probability of offending following $h_{t-1} = 0$, unlike what we found for the baseline model.

4.3 General Payoffs

So far we have considered who is motivated directly by her concern for apprehending offenses. As noted in the Introduction, such an assumption may be reasonable in light of the inspector's implicit rewards or career concerns. A more socially minded inspector, however, might have preferences for deterring offenses in the first place. We therefore consider a more general payoff structure as follows.

	N	O
W	$-S\mathbf{1}_t^S(I, W), 0$	$-L - S\mathbf{1}_t^S(I, W), \pi$
I	$-S\mathbf{1}_t^S(W, I) - i, 0$	$y - S\mathbf{1}_t^S(W, I) - i, -\gamma$

Here $-L$ is the the inspector's loss (or gain if $L \leq 0$) as a result of offenses which are not apprehended, while y again in her reward for apprehending an offense. We assume that $S > 0$ and that Assumption 1 still holds.

The following proposition describes a sufficient condition under which equilibria as the one described in Proposition 1 exists.

Proposition 4 *An equilibrium as described in Proposition 1 exists if*

$$L > -\frac{(i + (1 - \delta)S)(y - i - (1 + \delta)S)}{i + (1 + \delta)S}. \quad (12)$$

Note that the condition (12) is equivalent to Assumption 2 when $L = 0$. So, as one might expect, the condition for existence of equilibria with residual deterrence is relaxed when $L > 0$. Indeed, in this case, the inspector has an additional incentive to switch to inspect, since her payoff decreases when an offense occurs when she is waiting.

Proposition 4 establishes that equilibrium dynamics with residual deterrence periods can be found generically in our setting. This is not only helpful as a robustness check for our results, but also proves the existence of residual deterrence in a variety of economically relevant settings. A particularly interesting case is where the inspector is motivated by deterrence alone; that is, an inspector whose payoff is lowered by $L > 0$ whenever there is a crime (in our notation $y = -L$). Note that here, the unique equilibrium of the stage game is (N, O) . Still, if L is large enough, there exist equilibria where the inspector inspects only because after a conviction there are some periods of residual deterrence. Note that, unlike the analysis for the baseline model, and for our other analysis above, we make no claim as to the uniqueness of equilibria here.

5 Conclusions

We have studied a dynamic version of the inspection game in which the inspector 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) or “socially motivated” in the sense that she values the deterrence of crime. In an extension to the baseline model where agents have private information as to their rewards to offending, we show that beliefs that the inspector is inspecting follow a plausible pattern, being the highest immediately after a conviction before decaying gradually. We thus provide a fully-fledged theory of deterrence decay.

A few realistic possibilities that our paper leaves unexplored are the following. First, in practice there may be a range of different offenses which compete for the inspector’s attention. While such settings may be more difficult to analyze, and are perhaps prone to multiple equilibria, their study might reveal something as to the role of differences in penalties across different crimes. Such a theory could give a richer view of the role of “marginal deterrence” (which considers the effect of the penalties for one crime on the offense rates for others) by taking an equilibrium view of the behavior of enforcement officials. Another question relates to the role of the media in publicizing information about offenses and convictions, and in turn the influence this has on enforcers of the law. These possibilities are, however, left to future work.

References

- [1] Andreozzi, Luciano (2004), ‘Rewarding policemen increases crime. Another surprising result from the inspection game,’ *Public Choice*, 121, 69-82.
- [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] 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.
- [5] Block, Michael K., Frederick K. Nold and Joseph G. Sidak (1981), ‘The Deterrent Effect of Antitrust Enforcement,’ *Journal of Political Economy*, 89, 429-445.
- [6] Board, Simon and Moritz Meyer-ter-Vehn (2013), ‘Reputation for Quality,’ *Econometrica*, 81, 2381-2462.
- [7] Bond, Philip and Kathleen Hagerty (2010), ‘Preventing Crime Waves,’ *American Economic Journal: Microeconomics*, 2, 138-159.

- [8] DeMarzo, Peter, Michael Fishman and Kathleen Hagerty (1998), ‘The Optimal Enforcement of Insider Trading Regulations,’ *Journal of Political Economy*, 106, 602-632.
- [9] Dilmé, Francesc (2012), ‘Building (and Milking) Trust: Reputation as a Moral Hazard Phenomenon,’ mimeo University of Pennsylvania and University of Bonn.
- [10] Eeckhout, Jan, Nicola Persico and Petra E. Todd (2010), ‘A Theory of Optimal Random Crack-downs,’ *American Economic Review*, 100, 1104-1135.
- [11] Halac, Marina and Andrea Prat (2014), ‘Managerial Attention and Worker Engagement,’ mimeo Columbia University and University of Warwick.
- [12] Iossa, Elisabetta and Patrick Rey (2012), ‘Building Reputation for Contract Renewal: Implications for Performance Dynamics and Contract Duration,’ mimeo Toulouse School of Economics and University of Rome Tor Vergata.
- [13] 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.
- [14] Khalil, Fahad (1997), ‘Auditing without Commitment,’ *RAND Journal of Economics*, 28, 629-640.
- [15] Lui, Francis (1986), ‘A Dynamic Model of Corruption Deterrence,’ *Journal of Public Economics*, 31, 215-236.
- [16] Mookherjee, Dilip and IPL Png (1989), ‘Optimal Auditing, Insurance and Redistribution,’ *Quarterly Journal of Economics*, 104, 399-415.
- [17] Mookherjee, Dilip and IPL Png (1994), ‘Marginal Deterrence in Enforcement of Law,’ *Journal of Political Economy*, 76, 1039-1066.
- [18] Nagin, Daniel S. (2013), ‘Deterrence in the Twenty-First Century,’ *Crime and Justice*, 42, 199-263.
- [19] Polinsky, Mitchell A. and Steven Shavell (1984), ‘The Optimal Use of Fines in Imprisonment,’ *Journal of Public Economics*, 24, 89-99.
- [20] Reinganum, Jennifer and Louis Wilde (1985), ‘Income Tax Compliance in a Principal-Agent Framework,’ *Journal of Public Economics*, 26, 1-18.
- [21] Reinganum, Jennifer and Louis Wilde (1986), ‘Equilibrium Verification and Reporting Policies in a Model of Tax Compliance,’ *International Economic Review*, 27, 739-760.

- [22] Shavell, Steven (1992), ‘A Note on Marginal Deterrence,’ *International Review of Law and Economics*, 12, 343-355.
- [23] Sherman, Lawrence (1990), ‘Police Crackdowns: Initial and Residual Deterrence,’ *Crime and Justice*, 12, 1-48.
- [24] Strausz, Roland (1997), ‘Delegation of Monitoring in a Principal-Agent Relationship,’ *Review of Economic Studies*, 64, 337-357.

A Appendix A: Omitted Proofs

Proof of Corollary 1. An increase in γ or a decrease in π decreases β^* and therefore the probability that the inspector inspects. The values α^* and α_T remain unchanged. The duration T of the phase after a conviction during which agents abstain remains unchanged. These observations, together with (7), imply the result. ■

Proof of Corollary 2. Let $\alpha_T(S)$, $\alpha^*(S)$, $\bar{\alpha}(S)$ and $T(S)$ denote, respectively, the corresponding (unique) values of α_T , α^* , $\bar{\alpha}$ and T for a given value of $S \in [0, \min\{\frac{i}{1-\delta}, \frac{y-i}{1-\delta}\})$. Note that $\alpha^*(\cdot)$ is continuous for all S , while $T(\cdot)$ and $\alpha_T(\cdot)$ are only right-continuous for all S .

Let’s first show that $\bar{\alpha}(\cdot)$ is continuous. If $T(\cdot)$ is continuous at S , then it is easy to see that $\alpha_T(\cdot)$ and $\alpha^*(\cdot)$ are also continuous at S , so $\bar{\alpha}(\cdot)$ is continuous at S . For a generic function f , let $f(S^-)$ denote the left limit. Then, note that if S is such that $T(S) = T(S^-) + 1$, then $\alpha_T(S^-) = 0$, while $\alpha_T(S) = \alpha^*$. Therefore

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

Note that $T(S)$ (given in equation (6)) is increasing in S , and $T(0) = 0$. Assume that S is such that $\alpha_T(S) \in (0, \alpha^*(S))$. Then, there exists $\varepsilon > 0$ such that $T(S')$ is constant for all $S' \in (S - \varepsilon, S + \varepsilon)$. This implies that $\frac{d}{dS'}\bar{\alpha}(S')|_{S'=S}$ can be computed holding T constant. In this case we have

$$\bar{\alpha}'(S) = \frac{(1 - \delta)^2 \delta^T (\delta^T (2\delta i - \delta y + y) - 2\beta^* i)}{(\beta^* i (\delta^T ((1 - \delta)T + 1) - 1) - (1 - \delta) (\delta^T (2\delta S - \beta^* ((1 - \delta)T + 1)S - y) - \beta^* S))^2} \quad (13)$$

where, in order to keep the expression short, we use T to denote $T(S)$. Note that the sign of the RHS of the previous expression only depends on S through T . Also, given that T is increasing in S ,

if $\bar{\alpha}'(S)$ is negative for some S , it is also negative for all $S' > S$. Similarly, if $\bar{\alpha}'(S)$ is positive for some S , it remains positive at all $S' < S$. This ensures that $\bar{\alpha}(\cdot)$ is quasi-concave.

The condition for $\bar{\alpha}'(S)$ to be non-positive for all S is that $\delta^0(2\delta i - \delta y + y) - 2\beta^*i \leq 0$, that is $\beta^* \geq \frac{(1-\delta)y}{2i} + \delta$. Hence we put $\bar{\beta}^* \equiv \min \left\{ \frac{(1-\delta)y}{2i} + \delta, 1 \right\}$. For the case where $\frac{(1-\delta)y}{2i} + \delta < 1$, we now show that $T = 0$ when $\beta^* \geq \bar{\beta}^*$. To see this, note from (6) that $T \geq 1$ iff $S \geq \frac{i}{1+\delta}$. Also, Assumption 2 implies that $S < \frac{i(2\beta^*-1-\delta)}{1-\delta^2}$ for $\beta^* \geq \bar{\beta}^*$. It is easy to see that, given $\beta^* < 1$, these two conditions cannot be met simultaneously; i.e., we must have $T = 0$.

Next, note that for $\beta^* < \bar{\beta}^*$ we have $\bar{\alpha}'(0) > 0$. If $\frac{i}{1-\delta} \leq \frac{y-i}{1-\delta}$ then $\lim_{S \rightarrow \min\{\frac{i}{1-\delta}, \frac{y-i}{1-\delta}\}} T(S) = \infty$, so clearly $\bar{\alpha}$ is hump-shaped; i.e., in this case $\underline{\beta}^* = 0$. If $\frac{i}{1-\delta} \geq \frac{y-i}{1-\delta}$ then $\bar{\alpha}$ is only hump-shaped if $\lim_{S \rightarrow \frac{y-i}{1-\delta}} \bar{\alpha}'(S) < 0$, that is, if $\beta^* > \frac{(1-\delta)y+2i\delta}{2i} \delta^{T(\frac{y-i}{1-\delta})} \equiv \underline{\beta}^*$. ■

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 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.

Lemma 8 *In any equilibrium, $\beta(h^t) \geq \beta^{\min} \equiv \alpha^{-1}(\alpha^*)$ for all h^t , with equality when $V_I(h^t) = S$.*

Proof. Trivially, 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 inspector switches with positive probability from wait to inspect. That is, for each h^t , there exists a finite τ such that the inspector switches from wait to inspect at $(h^t, \mathbf{0}^\tau)$.

We can now prove the result by induction. Consider an equilibrium and a history h^t where the inspector 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

$$\begin{aligned} V_I(h^t) &= -i + \alpha_t(y + \delta V_I(h^t, 1)) + (1 - \alpha_t)(\delta V_I(h^t, 0)) \\ &= -i + \alpha_t(y + \delta(-S)) + (1 - \alpha_t)(\delta S). \end{aligned} \tag{14}$$

(Here, $V_I(h^t, 0) = S$ follows because the inspector is willing to switch to inspect at $(h^t, 0)$, because $V_W(h^t, 0) = 0$ (by the claim in Lemma 2), and because we must have $V_I(h^t, 0) \in [0, S]$ (by the first claim of Lemma 4).) Note that α^* solves the previous equation when $V_I(h^t) = S$. If $V_I(h^t) < S$, 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 $s \in \{0, \dots, \tau - 1\}$.¹⁴ This proves the first part of our result.

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

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

Note that α^* solves this equation when $V_I(h^t, 0) = S$. If $V_I(h^t, 0) < S$, 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 8 establishes the following. Consider any history h^t such that the inspector 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 inspector 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}). Now consider any history h^t with $h_{t-1} = 1$. Given that $\beta_t \in (\underline{\beta}, \bar{\beta})$, the inspector must switch from inspect to wait at t , and hence $V_I(h^t) = -S$. As argued above, there exists τ such that the inspector switches from wait to inspect for the first time at $t + \tau$ for $(h^t, \mathbf{0}^\tau)$. For $s = \tau - 1$, we have

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

and $\alpha_{t+s} \leq \alpha^*$, implying that $V_I(h^t, \mathbf{0}^s) \leq S$. Considering there are no switches from wait to inspect until τ , 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 $V_I(h^t) = -S$ 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 $\tilde{V}_I(\beta^{\max})$ in the following way:

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

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 number of periods that it takes for β to pass β^{\min} (less one), and $A_t(\beta^{\max}) \equiv \prod_{s=0}^{t-1} (1 - \min \{ \alpha(\tilde{\beta}_s(\beta^{\max})), \alpha^* \})$ is the probability of reaching t without any conviction.

¹⁴Equation (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 $s \in \{0, \dots, \tau - 1\}$.

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 from (15) that $\tilde{V}_I(\cdot)$ is strictly decreasing. Also, since $\min\{\alpha(\tilde{\beta}_t(\beta^{\max})), \alpha^*\}$ is continuous in β^{\max} , $\tilde{V}_I(\cdot)$ is continuous. Finally, given that $\lim_{\beta^{\max} \rightarrow \bar{\beta}} \tilde{V}_I(\beta^{\max}) = -\frac{i}{1-\delta} < -S$ and $\lim_{\beta^{\max} \rightarrow \bar{\beta}} \tilde{V}_I(\beta^{\min}) = S$, there is a unique β^{\max} which guarantees the indifference condition $\tilde{V}_I(\beta^{\max}) = -S$. ■

Proof of Corollary 3. Let $\bar{T} \geq 0$ and $\bar{\alpha}_{\bar{T}} \in (0, \alpha^*]$ be the unique values satisfying (5), i.e.

$$-S = -\sum_{j=0}^{\bar{T}} \delta^j i + (1 - \bar{\alpha}_{\bar{T}}) \delta^{\bar{T}+1} S + \bar{\alpha}_{\bar{T}} \delta^{\bar{T}} (y - \delta S). \quad (16)$$

Consider the game with payoff uncertainty in Section 4.1.

Index believes for the game with uncertainty parameter κ as $\beta^\kappa(h^t)$. We first show that there exists $\bar{\kappa} > \underline{\pi}$ such that, for all $\kappa \in (\underline{\pi}, \bar{\kappa})$, for any h^t with $h_{t-1} = 1$, $\beta_{t+\bar{T}}^\kappa(h^t, \mathbf{0}^{\bar{T}}) > \alpha^{-1}(\bar{\alpha}_{\bar{T}})$. To see this, note that there is $\bar{\kappa}$ such that, for all $\kappa \in (\underline{\pi}, \bar{\kappa})$, all h^t with $h_{t-1} = 1$, if $\beta_{t+\bar{T}}^\kappa(h^t, \mathbf{0}^{\bar{T}}) \leq \alpha^{-1}(\bar{\alpha}_{\bar{T}})$, then $\beta_{t+\bar{T}+1}^\kappa(h^t, \mathbf{0}^{\bar{T}+1}) = \beta^{\min}$. It is thus easy to see that the probability of offending must be uniformly higher for histories $(h^t, \mathbf{0}^s)$, $s \in \{0, \dots, \bar{T}\}$, in the game with uncertainty ($\kappa > \underline{\pi}$) than in the game without ($\kappa = \underline{\pi}$). However, comparing (16) to (15), we see that the continuation payoff following a conviction $\tilde{V}_I(\beta^{\max})$ must exceed $-S$, which is inconsistent with equilibrium in the game with random payoffs.

It is then easy to see that, for any $\varepsilon > 0$, there exists $\kappa^\# \in (\underline{\pi}, \bar{\kappa})$ such that, for all $\kappa \in (\underline{\pi}, \kappa^\#)$, for all h^t with $h_{t-1} = 1$, all $s \in (0, \dots, \bar{T} - 1)$, we have $\alpha(h^t, \mathbf{0}^s) < \varepsilon$. Again, comparing (16) to (15), we must have $\beta_{t+\bar{T}}^\kappa(h^t, \mathbf{0}^{\bar{T}}) \rightarrow \alpha^{-1}(\bar{\alpha}_{\bar{T}})$ as $\kappa \rightarrow \underline{\pi}$. One can then check that we must have $\beta_{t+\bar{T}+1}^\kappa(h^t, \mathbf{0}^{\bar{T}+1}) \rightarrow \beta^{\min}$ as well. ■

Proof of Lemma 5. Note first that since the inspector has the option to wait forever after any history, so we must have $V_W(h^t) \geq 0$ at any history with $h_{t-1} = 0$. Suppose for a contradiction that $V_W(h^t) > 0$ for some h^t with $h_{t-1} = 0$. Let \bar{V}_W be the supremum of continuation values over all histories with $h_{t-1} = 0$. Take $\varepsilon \geq 0$ small, and consider a history with $h_{t-1} = 0$ where the inspector expects payoff $\bar{V}_W - \varepsilon$. Note then that the inspector must strictly prefer to switch to inspect at t . Indeed, if he instead waits at t , then he obtains a continuation payoff at $t+1$ no greater than \bar{V}_W ; i.e., his expected continuation payoff at t from this strategy is no greater than $\delta \bar{V}_W$, which is less than $\bar{V}_W - \varepsilon$ provided ε is sufficiently small.

Next, Assumption 1a implies that the inspector who switches to inspect at t does not continue to inspect forever with probability 1. The reason is that this would imply perfect deterrence, irrespective of the signals generated; the inspector could hence save $\frac{i}{1-\delta}$ in date $-t+1$ terms by incurring the switching cost S , and we have $S < \frac{i}{1-\delta}$ by Assumption 1a. Let $t + \tau^*$ be the first date following t at which the inspector switches back to wait with positive probability (for some

realization of the signals), where $\tau^* \geq 1$. By considering this date, we obtain an upper bound on the date- t continuation payoff of our inspector who switched to inspect at date t :¹⁵

$$-S - \sum_{s=0}^{\tau^*-1} \delta^s i - \delta^{\tau^*} S + \delta^{\tau^*+1} \bar{V}_W.$$

However, irrespective of the value $\tau^* \geq 1$, this is less than $\delta \bar{V}_W$, and hence less than $\bar{V}_W - \varepsilon$, a contradiction. ■

Proof of Lemma 6. Suppose first that $\beta(h^t) < \beta^*$ at some h^t . In this case, we must have $\alpha(h^t) = 1$. However, by switching to inspect at date t and then to wait at $t+1$, the inspector is guaranteed the payoff $y - i - S - \delta S$, which is positive by Assumption 2. Hence we must have $\beta(h^t) \geq \beta^*$.

first time at which the inspector switches to wait with positive probability. Let this date be $t + \tau^*$ where $\tau^* \geq 1$. Then (using that there is no offending until at least $t + \tau^*$) the continuation payoff at h^t is

$$- \sum_{s=0}^{\tau^*-1} \delta^s i - \delta^{\tau^*} S = -i \frac{1 - \delta^{\tau^*}}{1 - \delta} - \delta^{\tau^*} S.$$

This is less than $-S$ by Assumption 1a. However, the inspector could obtain at least $-S$ by switching to wait at t , a contradiction. ■

Proof of Lemma 7. The argument is similar to that for Lemma 4. The key difference is that if $h_{t-1} = 0$, the inspector must switch to inspect at t with probability ξ satisfying

$$\frac{\beta^*(1 - \rho)}{\beta^*(1 - \rho) + 1 - \beta^*} + \xi \frac{1 - \beta^*}{\beta^*(1 - \rho) + 1 - \beta^*} = \beta^*. \quad (17)$$

■

Proof of Proposition 3. The rates of inspector switching following $h_{t-1} = 1$ (namely, $1 - \beta^*$) and following $h_{t-1} = 0$ (namely ξ satisfying (17)) are the unique rates which maintain agent beliefs at $\beta_t = \beta^*$. Agents are thus indifferent between offending and not, and the unique rates of offending that maintain inspector incentives to switch are given by (10) and (11). Optimality of the inspector's switching strategy follows as usual from the one-shot deviation principle. ■

Proof of Proposition 4. We look for an equilibrium for which $V_W(h^t, \mathbf{0}^{T+1})$ and $V_I(h^t, \mathbf{0}^{T+1})$ are independent of h^t , and denote these values respectively by V_W^{T+1} and V_I^{T+1} . Similarly, we look for an equilibrium for which $V_W(h^t, 1)$ and $V_I(h^t, 1)$ are determined independently of h^t and denote these values respectively by V_W^0 and V_I^0 . Independence of these values from h^t allows us to write

$$V_W^{T+1} = \alpha^*(-L) + \delta V_W^{T+1}, \quad (18)$$

$$V_I^{T+1} = \alpha^*(y - i + \delta V_I^0) + (1 - \alpha^*)(-i + \delta V_I^{T+1}). \quad (19)$$

¹⁵This is obtained by observing that no agent offends between t and $t + \tau^* - 1$.

Also, the continuation values immediately following a conviction are given by:

$$V_W^0 = \delta^T \alpha_T (-L) + \delta^{T+1} V_W^{T+1}, \quad (20)$$

$$V_I^0 = \frac{1 - \delta^T}{1 - \delta} (-i) + \alpha_T \delta^T (y - i + \delta V_I^0) + (1 - \alpha_T) \delta^T (-i + \delta V_I^{T+1}). \quad (21)$$

Finally, the indifference conditions are

$$V_W^{T+1} = V_I^{T+1} - S \quad (22)$$

$$V_W^0 = S + V_I^0. \quad (23)$$

We want to show that there exists $(T, \alpha^*, \alpha_T) \in (\mathbb{N} \cup \{0\}) \times (0, 1) \times (0, \alpha^*]$ (and the corresponding V_W^0, V_I^0, V_W^{T+1} and V_I^{T+1}) that satisfy the equations (18)-(23).

Using all equations except (20) and (21), one can write V_W^0, V_I^0, V_W^{T+1} and V_I^{T+1} in terms of just α^* . Also, one can write

$$\alpha^*(L + y + \delta V_I^0 - \delta V_I^{T+1}) - (i + (1 - \delta)S) = 0 \quad (24)$$

Subtracting (21) from (20) and using (22) and (23) we find

$$\begin{aligned} S &= \frac{1 - \delta^T}{1 - \delta} i + \delta^T (i - \delta S) - \alpha_T \delta^T (L + y + \delta V_I^0 - \delta V_I^{T+1}) \\ &= \frac{1 - \delta^T}{1 - \delta} i + \delta^T (i - \delta S) - \frac{\alpha_T}{\alpha^*} \delta^T (i + (1 - \delta)S) \end{aligned}$$

where we used (24) to obtain the second equality. Finally, simple algebra manipulation implies

$$(1 + \delta^T)S - \frac{1 - \delta^T}{1 - \delta} i = (1 - \alpha_T / \alpha^*) \delta^T (i + (1 - \delta)S)$$

Given that $\alpha_T \in (0, \alpha^*]$, the previous equation implies that T is given by the expression (6), as in our main model (note that the expression does not depend on L). Also, it is trivial to find that that α_T / α^* is decreasing in S for fixed T , so the solution for α_T and T , for a given α^* , is unique.

It only remains to show that there exists $\alpha^* \in (0, 1)$ solving (18)-(23). Using all equations except (20), one can write $\alpha_T, V_W^0, V_I^0, V_W^{T+1}$ and V_I^{T+1} in terms only of model primitives and α^* . Plugging the corresponding expressions into equation (20) we find

$$\frac{-\frac{y(i - \delta S + S)}{i - \alpha^* L - \delta S + S} + \frac{i - \delta S + S}{\alpha^*} + 2\delta S}{\delta} = 0.$$

Multiplying by the denominators, the previous expression is transformed into

$$0 = f(\alpha^*) \equiv (i + 2\alpha^* \delta S - \delta S + S)(i - \alpha^* L - \delta S + S) - \alpha^* y (i - \delta S + S)$$

It is easy to verify that $f(\cdot)$ (which is a second order polynomial) satisfies $f(0) > 0$. So, a sufficient condition for a solution to $f(\alpha^*) = 0$ for $\alpha^* \in (0, 1]$ to exist is $f(1) < 0$, which leads to the equation stated in the proposition. ■

B Appendix B: Omitted extensions

B.1 Imperfect inspection

Suppose now that (I, O) leads to conviction only with some probability $\lambda \in (\beta^*, 1)$. All other components of the model remain unchanged. We keep Assumption 1 the same, and we now adjust Assumption 2 in the following way:

Assumption 2b: The inspector has incentives to switch to inspect if the agent offends for sure, i.e., $\lambda y - i > S(1 + \delta)$.

Under these assumptions, Lemmas 1-3 still hold adjusting the probability the inspector inspects in any period is equal to $\beta_a^* = \frac{\pi}{\lambda(\pi + \gamma)}$, thus ensuring agent indifference between O and N .

Now consider each agent's probability of offending. Let $\alpha_{a,T}$ be the probability of an offense at date t following a conviction, i.e. following $h_{t-1} = 1$. Let α_a^* denote the probability of an offense at date t following a date $t - 1$ such that the agent offends with positive probability and $h_{t-1} = 0$. Equations (3) and (5) remain as before, but with $\lambda\alpha_{a,T}$ and $\lambda\alpha_a^*$ substituted for α_T and α^* . That is,

$$S = \lambda\alpha_a^*(y - i - \delta S) + (1 - \lambda\alpha_a^*)(-i + \delta S)$$

and

$$-S = -\sum_{j=0}^T \delta^j i + (1 - \lambda\alpha_{a,T})\delta^{T+1}S + \lambda\alpha_{a,T}\delta^T(y - \delta S).$$

We thus find that $\alpha_a^* = \alpha^*/\lambda$ and $\alpha_{a,T} = \alpha_T/\lambda$.¹⁶ We continue to find that T does not depend on λ . The structure of equilibrium remains as above, but with these higher rates of offending and the higher probabilities the inspector attempts inspection.

Note that the above findings mirror the case for the static inspection game. In the static game a decrease in the probability of a conviction for the actions (I, O) must be met by an increase in the probability of offending to preserve the inspector's indifference between I and W . The same observation holds in our repeated setting.

B.2 Continuous Time Limit

We now investigate a natural continuous time limit. Let the period length be $\Delta > 0$ and let the discount rate satisfy $\delta = e^{-\rho\Delta}$ for $\rho > 0$. Let $i = i'\Delta$, where i' is the flow cost of inspection. We let $\pi = \pi'$, $\gamma = \gamma'$, $y = y'$ and $S = S'$ where π' , γ' , y' and S' are lump-sum payoffs for the agent and inspector respectively.

These choices will ensure that the probability of offending in each period shrinks to zero as Δ is taken to zero, but that the rate per unit of time converges to a positive constant. Indeed, the value

¹⁶ Assumption 2a ensures that $\alpha^*/\lambda < 1$ (and therefore $\alpha_T/\lambda < 1$).

for α^* solving (3) is now given by

$$\alpha' = \frac{i'\Delta + S'(1 - e^{-\rho\Delta})}{y' - 2S'e^{-\rho\Delta}}.$$

Thus the rate of offending $\frac{\alpha'}{\Delta}$ converges to

$$\alpha^{\text{lim}} = \frac{i' + S'\rho}{y' - 2S'}.$$

We next calculate the limiting value for the duration of residual deterrence $\bar{T} = T\Delta$. Note that (5) now becomes

$$-S' = -\sum_{j=0}^T e^{-\Delta\rho j} \Delta i' + e^{-\Delta\rho(T+1)} S' + \bar{\alpha}_T e^{-\Delta\rho T} (y' - 2e^{-\rho\Delta} S'). \quad (25)$$

The solution $\bar{\alpha}_T$ converges to zero as Δ goes to zero, and $\bar{T} = T\Delta$ converges to $\bar{T}^{\text{lim}} = \frac{1}{\rho} \log \frac{i' + \rho S'}{i' - \rho S'}$ as Δ goes to zero (the latter uses that the first two terms on the right-hand side of (25) converge to $-i' \frac{1 - e^{-\rho\bar{T}}}{\rho} + S' e^{-\rho\bar{T}}$).

Thus, in the limit, agents abstain from offending for duration \bar{T}^{lim} following a conviction. After this time, offending resumes at rate $\bar{\alpha}^{\text{lim}}$. The long-run average rate of offending is

$$\frac{\bar{\alpha}^{\text{lim}}}{\lambda \bar{T}^{\text{lim}} + 1} = \frac{\rho \bar{\alpha}^{\text{lim}}}{p^* \bar{\alpha}^{\text{lim}} \log \frac{i' - \rho S'}{i' + \rho S'} + \rho}.$$