16/10/2013

Corporate Leniency in a Dynamic World: The Preemptive Push of an Uncertain Future*

Dennis L. Gärtner†

First Version: May 2012
This Version: October 2013

Abstract

This paper explores the incentives for collusive offenders to apply for leniency in a dynamic setting, where the risk of being independently caught evolves stochastically over time. We show how such future uncertainty can push firms into preemptive application, and that these preemptive incentives may in fact unravel to the point where firms apply long before the risk of independent detection is in any way imminent. The analysis sheds light on factors and policy instruments which favor such an unraveling effect. These include: little discontinuity in time and state, firms' patience, and a relatively harsh treatment of firms which fail to preempt other whistleblowers. In contrast, the described effects do not necessarily require a very high absolute level of leniency reduction, or even rewards.

Keywords: cartel, collusion, leniency program, preemption, dynamics.

JEL Classification: D43; D84; K21; K42; L41.

*This paper has substantially benefited from discussions with Joe Harrington and Jun Zhou, as well as from very helpful comments by Eric van Damme, Hans-Theo Normann, and seminar participants at the Düsseldorf Institute for Competition Economics (DICE). All remaining errors are my own.

†University of Bonn, Department of Economics, Lennéstrasse 37, 53113 Bonn, Germany; Tel. +49-228-739478; dennis.gaertner@uni-bonn.de.

1 Introduction

Firms which have committed a joint legal offense by colluding on prices face a risk: that the offense will be detected, and firms prosecuted and fined by authorities. Leniency programs are meant to mitigate the problem that such detection is typically imperfect and costly: By granting self-reporters a substantial fine reduction, they aim to entice firms to self-report, thereby reducing prosecution costs and raising the probability of prosecution, which in turn reduces also the temptation to commit the offense in the first place.

The functioning of these programs is obvious if the fine reduction is so large that paying the reduced fine is more attractive to each firm than continuing to face the risk of being found out. Intuitively though, leniency programs are thought to introduce also another risk: that one of my partners in crime reports and turns me in. By this feature, leniency programs may have a bite even in situations where the risk of independent discovery alone does not make it profitable for a firm to self report—but the risk of others reporting does.

Theory has had a hard time putting its finger on this preemptive aspect. Essentially, this is because in the standard static full information setting, the equilibrium concept itself precludes any such strategic uncertainty: In equilibrium, firms are assumed to perfectly predict each others' behavior, which means we typically obtain multiple equilibria, including one in which all firms self report, and one in which no firm does. Thus, not only does the model fail to depict a preemptive motive, but as a result, we also face the issue of multiple equilibria.

One idea to resolve this and identify preemptive motives more clearly is to introduce asymmetric information. This idea is pursued in a recent paper by Harrington (2013), who argues that if firms obtain private signals on the chance of being caught, then the resulting uncertainty about others' signals (and consequent actions) introduces an additional 'push' for preemptive application.[1]

This paper points out a different plausible motive for preemption: symmetric uncertainty about the future in a dynamically evolving world. Our consideration of dynamics requires no motivation other than simply observing that if no firm reports today, and if the offense is not detected by authorities, then firms can always still report tomorrow. Moreover, it seems likely that fundamentals which affect this decision change over time. In particular, firms' perceived likelihood of being caught will evolve as all sorts of relevant information comes in: as offices are raided, as a sector inquiry is launched, as customer complaints are filed, as authorities hire new specialists, as related cartels are busted, etc.

It is obvious that this dynamic extension of the baseline setting can induce behavior which conditions on these fundamentals (and thereby on time), particularly

[1] Marshall et al. (2013) consider a similar informational issue in the context of multi-product collusion.

that firms again apply only if the detection rate is high. Perhaps a bit less obviously, this in turn can spur preemptive incentives of the following nature: If I am worried about the fundamentals moving into a region where reporting is profitable *irrespective of strategic considerations* (i.e., I want to report even if the other doesn't), then I may actually want to report *before* we get into that region, in order to preempt others. Moreover, if the process by which the environment evolves shows some persistence, then this incentive to preempt can unravel and cause firms to self report long before the state actually reaches a level where reporting is dominant, which can make leniency programs enormously powerful.

We first illustrate this in a simple two-firm *discrete-time* framework, in which we let the likelihood of detection (the 'state') follow a trendless random walk. We highlight the aforedescribed preemptive incentive, and we derive sufficient conditions for this incentive to unravel to the point where immediate reporting is the only equilibrium. Crucial factors are firms' patience, and how harshly the leniency program treats firms who fail to be the very first to report (i.e., who report together with the other firm, or later).

Moreover, the discrete analysis suggests that the incentive to preempt is not least driven by how coarsely we choose the discrete grid for time and state. To follow up on this, we next consider the limit case of *continuous* time and state. This limit case vividly illustrates the potential power of leniency programs in a changing environment: For essentially any leniency award structure and any level of firms' patience, the mere possibility, however small, of reaching a state in which reporting is dominant is enough to make immediate reporting the only equilibrium.

In terms of policy conclusions and optimal leniency design, our model stresses the importance of being *relatively* harsh on latecomers, and that this may in fact be more crucial than the *absolute* degree of amnesty given. It also shows the importance of creating an environment which comes as close as possible to a continuous setting—such as by clearly distinguishing firstcomers from latecomers even if the lag is very small. In terms of modeling, it illustrates that the size of a discrete grid can be quite crucial in these sorts of models of preemption.

Although our primary aim is a better understanding of the strategic risks created by leniency programs, a further motivation comes from an empirical observation by Gärtner and Zhou (2012), which is that leniency applications tend to be made with *significant delay*, not only relative to the installation or revision of leniency programs, but also relative to the collapse of the relevant cartel (see Figure 1). Generally, this hints at the importance of dynamic considerations, and that retaliation for deviations from collusive behavior may not be the only important cause for leniency applications. More specifically, casual intuition suggests that the observed delays may be driven by exogenous shocks to the environment, which is precisely the setting we consider. Our results indicate that such an explanation will crucially rely on discreteness in time or state, which should therefore

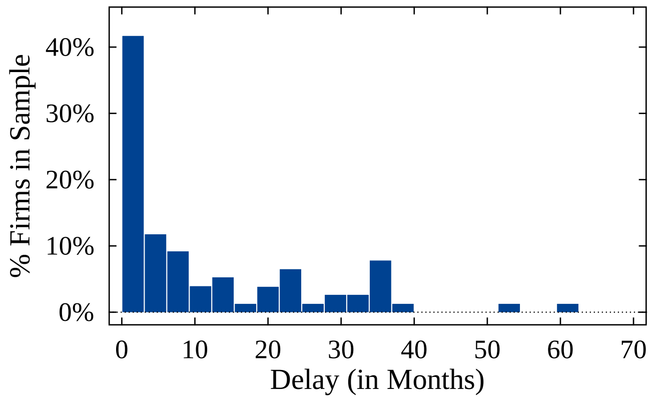

Figure 1: Delays in EC Leniency Applications (Cartel Collapse until First Application) from 1996–2008 (78 Obs.). Data: Gärtner and Zhou (2012).

be carefully explained rather than assumed *ad hoc*.

The rest of the paper is organized as follows: Section 2 reviews the related literature. Section 3 lays out the basic static model of leniency application. Section 4 embeds this model into a *dynamic* setting, in which both time and the state evolve in discrete steps. We isolate the key preemptive motive and discuss sufficient conditions for this to completely unravel. Section 5 picks up on the insight that these sufficient conditions are not least driven by *how finely* the discrete grid is chosen, by considering the (hypothetical) limit case of *continuous* time and state. We show that for a wide class of models, including but not limited to Section 4's model, this limit entails immediate reporting. Section 6 discusses implications, limitations and extensions, and draws some tentative policy conclusions. Section 7 concludes.

2 Related Literature

This paper adds to a rapidly developing strand of literature which looks at the mechanics by which leniency programs affect incentives to collude and to self report (see Spagnolo, 2008, for an excellent survey).[2] This strand has produced many interesting insights by considering various different facets, depending (among other things) on whether collusion is still ongoing at the time self reporting is considered (which can produce perverse effects in terms of stabilizing collusion), whether self-reporting is an on- or off-equilibrium phenomenon, whether the model includes investigations as a necessary precursor to prosecution, whether firms have (jointly) committed only one or numerous offenses, and, if offenses were commit-

[2]More recent additions include Harrington (2013), Marshall et al. (2013), Lefouili and Roux (2012), Chen and Rey (2012), and Choi and Gerlach (2009).

ted on multiple markets, whether leniency programs provide special treatment for multiple offenders (and multiple self-reporters).

We pick up a very basic strategic issue which, in some flavor, is found in virtually all of these models. As nicely described by Spagnolo (2004), the basic idea behind leniency programs is to create a strategic situation which brings wrongdoers as close as possible to a Prisoners' Dilemma in terms of inducing them to self report. This basic idea faces a caveat, though: While it is usually clear that firms want to self-report if others do, the optimal action if others *don't* will depend on factors such as the level of fine reduction, the perceived risk of being caught independently, and other opportunity costs of reporting (if, for instance, reporting causes collusion to break down). Ultimately, this can lead to a multiplicity in equilibria which is difficult to interpret.

As pointed out already by Harrington (2013), the literature to date has largely sidestepped this issue by focussing on cases in which reporting is optimal also in the latter case—not least by *suggesting* policies which achieve precisely this effect by means of generous fine reductions or even rewards for firms which self report (cf. Spagnolo, 2004; Aubert et al., 2006). In contrast, we know far less about situations in which fine reductions are not so high as to make self reporting a clear dominant choice. Technically, the resulting complementarities in best responses lead to multiple Nash-equilibrium predictions. Intuitively, these complementarities give rise to a situation of 'strategic risk' (Spagnolo, 2004), in that firms' actions are mainly driven by *beliefs* about how others behave. A better understanding of this strategic risk is desirable not least because it is a key aspect which distinguishes the leniency-application problem from the traditional literature on public law enforcement with *single* offenders (see the seminal contribution by Becker, 1968, and the more recent survey in Polinsky and Shavell, 2000).

Of the few papers that portray such strategic worry, Harrington (2013) develops a model in which firms hold private information on the chance of being caught. He shows that this private information—particularly, the uncertainty about *other* group members' information—can give firms an extra push to apply for leniency. Marshall et al. (2013) follow a similar approach in the context of multi-product collusion. Relatedly, Spagnolo (2004) previously used Selten and Harsanyi's (1988) concept of 'risk dominance' to select among multiple equilibria, which can be motivated by a similar story (see Carlsson and van Damme, 1993).

This paper, in contrast, describes a strategic worry which is not driven by asymmetric information, but by (symmetric) uncertainty about the future: Firms worry about the chance that the environment moves into a state in which the other will report, which gives them a reason to preempt. This argument links aspects which have been touched upon elsewhere: Not least, although not a main result, Motta and Polo's (2003) seminal paper actually points out that leniency reductions need not always be so strong as to make reporting altogether dominant, but it can suffice for this to be the case only *conditional on firms already being under investigation*: In a setting in which the environment evolves in a Markovian fashion over three possible states (the initial state, investigation, and prosecution) and where prosecution must be preceded by a previous investigation, this can lead to a one-step preemption incentive in which firms report even before being investigated, because they fear that the other will report once they *are* under investigation.[3] This paper can be seen as embellishing on this theme by 'adding finer steps' to this argument, and by showing how the argument can iteratively unravel across steps.

This iterative unraveling argument itself is related to an effect pointed out by Motchenkova (2004), who models collusion and leniency application as an optimal stopping problem in which firms weigh the continuation value of colluding against the risk of a fine which increases deterministically with the cartel's duration. There, the fact that firms *know* that the impending fine will eventually grow so large as to make reporting dominant can unravel so as to make firms report right at the outset, when the fine risked by collusion is still very small. Our analysis derives a related unraveling argument, but in an environment which is not deterministic, and in which there is no clear (upward) trend in the environment's evolution—reporting is not *certain* to become the dominant choice at any point in the future.

Our specific model setup is very similar to Harrington (2008), who also considers a dynamic setting in which the probability of detection varies stochastically over time. The crucial difference is that the latter assumes this stochastic process to be without memory, i.e. detection probabilities are independent over time. In contrast, persistence in the process is a key ingredient to the preemptive effect which we describe. And, as noted in Harrington (2008), persistence seems natural when allowing for shorter periods, as we do.

The flavor of the argument which this persistence allows us to develop in a *dynamic* setting, in turn, is very reminiscent of arguments in the global-games literature (Carlsson and van Damme, 1993; Morris and Shin, 1998) on *static* games of private information, where the mere possibility of an extreme signal which triggers a dominant action unravels to predict actions over the full spectrum of (more likely) signals. In both cases, the driving force for the unraveling is a complementarity in payoffs: The more prone the other firm is to report (the lower the critical signal or state), the more attractive it is for me to report (so I lower my critical level), and vice versa.[4]

[3] This can happen in Motta and Polo's model when the chance of investigation is sufficiently high (see the first item in Proposition 1).

[4] One noteworthy difference is that, in our context, the argument is *one-sided*, as not reporting is never a dominant strategy. Moreover, more recent papers in the global games literature (Angeletos et al., 2007, for instance) have challenged the stability of the uniqueness argument if *dynamics* are introduced into the game. This stands in contrast to our setting, where dynamics are crucial to the

Finally, the general point that embedding a stage game with multiple equilibria into a stochastic dynamic setting can resolve multiplicity has previously been made by Frankel and Pauzner (2000) in the context of a labor-market model (see also Burdzy et al., 2001), and by Mason and Weeds (2010) in the context of a real-options investment game. In spite of the very different contexts, their models share structural similarities with our continuous case.[5]

3 The Static Whistleblowing Game

We begin with a description of a static version of the basic leniency-application game, which we later extend into a dynamic version. It is identical to the baseline model in Harrington (2013), but a similar structure is found in literally any (more elaborate) model of leniency programs.

The game is played between two firms. To focus ideas, we consider a situation in which these two firms have already committed an offense, but this offense has not yet been detected or brought to trial. If the offense is detected and firms are prosecuted, the (full) fine imposed on each firm is $F > 0$. Firms simultaneously choose to report the offense to authorities ('blow the whistle' b) or remain silent ('don't blow' $\not b$).[6] If neither firm reports, firms are independently caught by authorities with probability $\rho \in [0,1]$, in which case each firm pays the full fine F. If only *one* firm blows the whistle, it pays a reduced fine θF, $\theta < 1$, whereas the other pays the full fine F. Finally, if *both* firms blow the whistle, each pays a reduced fine $\hat{\theta} F$, $\hat{\theta} \in (\theta, 1)$.[7] Normalizing the full fine to $F = 1$ (and assuming firms to be risk neutral), the symmetric normal-form game is that shown in Figure 2.

For $\rho > \theta$, $\not b$ is strictly dominated, so the unique equilibrium has both firms blow the whistle. For $\rho \leqslant \theta$ in turn, there are multiple equilibria: The two pure-strategy equilibria (b, b) and $(\not b, \not b)$, and a mixed-strategy equilibrium in which firms each blow the whistle with probability $(\theta - \rho)/(1 - \rho + \theta - \hat{\theta})$.

This simple model has various shortcomings. *First*, predictions are unique only for $\rho > \theta$. In this case, the model predicts whistleblowing because the fine

results.

[5]Apart from the difference in context and remaining technical dissimilarities (for instance, in Frankel and Pauzner, 2000, either action becomes dominant for certain states, much like in the global-games literature) issues, a key feature of our analysis is that we analyze both the discrete *and* continuous case, and it is precisely the combination of the two which yields the most useful insights in our setting.

[6]We use the terms '(self-) reporting,' 'applying for leniency' and 'whistleblowing' interchangeably, as our model will not distinguish leniency applications by legal entities from applications by private individuals.

[7]In the context of US leniency programs (where only one single firm can receive leniency), for instance, Harrington (2013) lets $\hat{\theta} = (1 + \hat{\theta})/2$, thereby implicitly assuming that if both apply, firms receive full lenience or none with equal probability.

		F2	
		b	$\not b$
F1	b	$-\hat{\theta},-\hat{\theta}$	$-\theta,-1$
	$\not b$	$-1,-\theta$	$-\rho,-\rho$

Figure 2: The Static Whistleblowing Game.

reduction is so high (relative to the chance of being caught) that blowing the whistle becomes optimal *independently of what the other does*. Thus, the prediction is unique only in cases where strategic considerations actually play no role.[8]

Second, when strategic considerations *do* play a role (for $\rho < \theta$), predictions are rather weak, in that they include both the case of neither blowing the whistle and both blowing the whistle.[9]

Finally, in the latter case, predictions heavily rely on the coordination of beliefs implicit in the concept of Nash equilibrium: Firms prefer to blow the whistle if the other does, and prefer not to if the other doesn't. Owing to the nature of Nash equilibrium (specifically, its assumption that beliefs about the other's actions be correct), this complementarity in best responses translates into a multiplicity of equilibria, and which equilibrium obtains in turn is solely driven by *beliefs*—the formation of which is beyond the model.

The last point in particular suggests that the model may miss out on an important intuitive aspect of leniency programs: That I become enticed to blow the whistle because I am nervous about the other blowing the whistle.[10] The next section will show how such a story can be told in the context of a dynamically evolving environment.

[8]Formally, most leniency programs in fact award full leniency to solitary whistleblowers, implying $\rho = 0$, which would indeed make this the relevant case—and thereby predict the immediate dissolution of all cartels. Bear in mind, though, that blowing the whistle may entail various costs which are outside of the model, and which effectively increase θF: Legal costs, uncertainty about whether submitted evidence will suffice, and not least more implicit costs in terms of dimmer prospects for colluding in the future (due to stricter surveillance by authorities, or due to a loss of reputation as a reliable colluder).

[9]See Spagnolo (2004) for a discussion on various possible refinements for equilibrium selection in this context, such as Pareto- or risk-dominance.

[10]Strictly speaking, the mixed-strategy equilibrium can be interpreted as containing such a feature. However, that equilibrium involves the usual counterintuitive comparative statics, because whistleblowing in this equilibrium is driven by the aim of keeping the *other* indifferent, so that whistleblowing becoming more attractive leads both firms to blow the whistle with *lower* probability. In particular, the probability of reporting *decreases* in the extent of leniency awarded to first-in solitary whistleblowers.

Before we proceed to that extension, it is worth observing that the basic strategic issue captured by the above game extends to all sort of elaborations: more than two firms, that collusion is still ongoing at the time the decision is taken (in which case blowing the whistle entails the added opportunity cost of foregoing collusion in the future), that the industry is already under investigation (which will raise ρ), that firms collude in multiple markets, or that blowing the whistle entails a direct cost (of assembling documents and evidence) or a risk (that this evidence may not suffice).

4 Preemptive Unraveling in Discrete Time: A Random-Walk Example

This section extends Section 3's static baseline model by two aspects: (i) That, so long as no one blew the whistle previously, firms can always reconsider, and (ii) that parameters of the environment may change over time. For specificity, like Harrington (2008, 2013), we focus on the detection probability ρ as the relevant parameter. We first lay out a (discrete) model with these features, and then discuss how the interplay between the two can create incentives for preemptive whistle-blowing which are not present in the static model.

4.1 Model Setup

We let firms play the above game over discrete periods $t = 0, 1, 2, \ldots$, where firms discount the future exponentially at rate $\delta \in (0, 1)$. In every period, firms simultaneously decide whether or not to blow the whistle (we now refer to the latter as 'waiting,' at times). As soon as a firm blows the whistle, the game ends, with payoffs as above: A firm which didn't blow the whistle pays F, a firm which did pays θF if it did so alone, and $\hat{\theta} F$ if whistleblowing was joint.

As long as no firm blows the whistle, they run the risk of independent detection by antitrust authorities, which entails a fine of F for both firms. We allow this risk to vary over time: In any period t, the probability of detection is $\rho_t \in [0, 1]$.

In any period t, firms mutually know the current risk of detection ρ_t (and ρ_t's past trajectory), but face symmetric uncertainty about its future development.[11] We let the stochastic evolution of ρ_t be described by a transformation of a simple symmetric random walk: $\rho_t = f(X_t)$, where $f : \mathbb{Z} \to [0, 1]$ is a nondecreasing function, and where X_t denotes a simple symmetric random walk over the integers $\mathbb{Z}$, which starts at zero, which can only increase or decrease by one from each period

[11] The important feature is not that firms literally *correctly* assess all probabilities ρ_t to date, but that they share the *same* assessment of the present and the uncertain future.

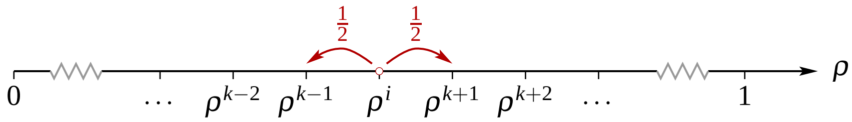

Figure 3: Random-Walk Evolution of ρ_t over the Grid $\{\rho^k\}_{k\in\mathbb{Z}}$.

to the next, and does so with equal probability. The transformation by f makes sure that $\rho_t \in [0,1]$ for all t.

The stochastic evolution of ρ_t can equivalently be understood as a simple (Markovian) random walk over an infinite countable grid $R = \{\rho^k\}_{k\in\mathbb{Z}} \subset [0,1]$ with $\rho^k \leqslant \rho^{k+1}$ for all k, and where

$$\Pr\big(\rho_{t+1} = \rho^{k'}\,|\,\rho_t = \rho^k\big) = \begin{cases} 1/2, & \text{for } k' \in \{k-1, k+1\}, \\ 0, & \text{otherwise} \end{cases}$$

for any initial $\rho_0 \in R$ (see the illustration in Figure 3).[12]

As outlined above, such a random walk can be motivated by the idea that the end of each period brings news regarding the independent discovery rate ρ, where this news can be either good or bad (from firms' point of view).[13]

We let firms play Markovian, possibly mixed strategies. Firm i's strategy is thus represented by a mapping $\sigma_i : R \to \Delta\{b, \not b\}$, which describes the probability that the firm blows the whistle given any current state $\rho_t \in R$. The restriction to Markovian strategies involves no loss of generality given that the stochastic process for ρ_t is Markovian, and given that there is no scope for conditioning current behavior on past actions: In any t, firms only face a choice of blowing the whistle or not for a unique history of previous actions, which is that both played $\not b$ in all previous periods.

[12] Being infinite but contained in $[0,1]$, the grid obviously cannot be equispaced.

[13] A few details about the process ρ_t are worth noting. *First*, when interpreting candidate paths for ρ_t, bear in mind that ρ_t represents the hazard of being detected *conditional on not having been detected by* t. Thus, a plausible response for ρ_t to a single piece of (bad) news coming in might in fact be that ρ_t rises in the short run, but drops again in the long run, indicating that authorities have failed to prosecute based on the new piece of evidence. Contrary to the above motivation, we may thus in fact expect ρ_t to fall in response to no news. This is little more than a matter of framing, though. *Second*, we assume the underlying random walk X_t to be driftless. We relax this assumption in our continuous case below. For the moment, observe that discounting introduces a sort of 'downward drift' as far as payoffs are concerned, not least in the sense that the discount factor can be interpreted as incorporating the probability that the game ends without prosecution (authorities drop the case for sure). *Finally*, an important property of this process is that the probability of ρ_t staying the same from one period to the next is zero. We discuss the role of this assumption further below.

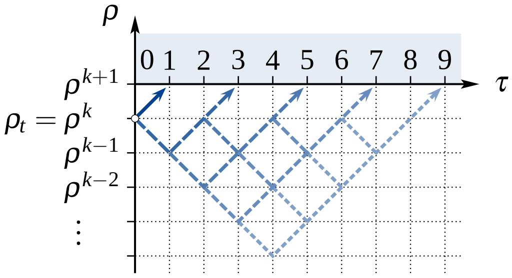

Figure 4: Illustration of First Crossing for the Process ρ_t.

4.2 The Incentive to Preempt: A Lower Bound

This section's key argument for preemption evolves around the following question (illustrated in Figure 4): Suppose we are currently in state $\rho_t = \rho^k$, and that I *know* that the other player is sure to blow the whistle in the future as soon as the state hits the next notch ρ^{k+1}. What are my incentives then to preemptively blow the whistle in the current state already?

To assess this, for any current period t and state $\rho_t = \rho^k$, we let

$$\gamma(\tau) \equiv \Pr\Big(\rho_{t+\tau} = \rho^{k+1} \text{ and } \rho_{t+1}, \rho_{t+2}, \ldots, \rho_{t+\tau-1} < \rho^{k+1} \Big| \rho_t = \rho^k\Big)$$

denote the probability that, starting from $\rho_t = \rho^k$ in t, ρ_t will have 'moved up to the next notch' τ periods from now (but not before). This probability is independent of t and k by ρ_t's Markov properties. Moreover, it simply corresponds to the probability that the underlying simple symmetric random walk X_t, starting at zero in $X_0 = 0$, will first hit $X_t = 1$ by any future period t. As illustrated in Figure 4, this can obviously only happen in odd periods, so $\gamma(\tau) = 0$ for even τ. The process will be a notch up after *one* period with probability $\gamma(1) = 1/2$. The process will be one notch up after *three* periods with probability $\gamma(3) = 1/8$, as there is one path by which this can happen ('down, up, up'), out of $2^3 = 8$ possible paths. Eventually, elementary combinatorics yields:

Lemma 4.1. *The probabilities $\gamma(\tau)$, $\tau \in \{1, 2, 3, \ldots\}$, are given by*

$$\gamma(\tau) = \begin{cases} (-1)^{(\tau-1)/2} \dbinom{1/2}{(\tau+1)/2}, & \tau \text{ odd}, \\ 0, & \tau \text{ even}. \end{cases} \tag{1}$$

Figure 5's panel (a) illustrates these probabilities, and the rate at which they decline in (odd) τ. Building on this, we next let

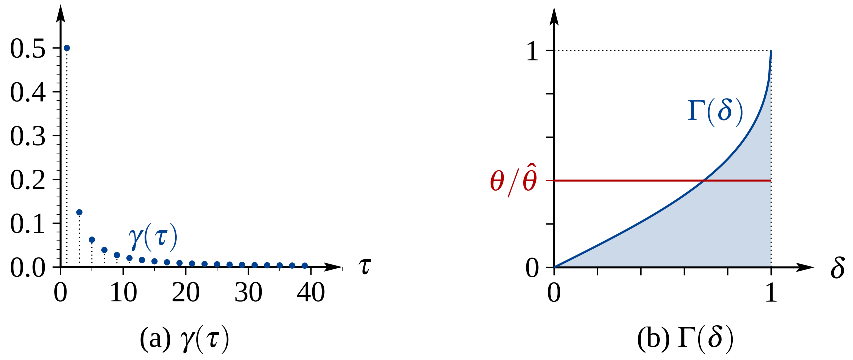

Figure 5: First-Passage Probabilities $\gamma(\tau)$ and their Discounted Sum $\Gamma(\delta) \equiv \sum_{\tau=1}^{\infty} \delta^\tau \gamma(\tau)$.

$$\Gamma(\delta) \equiv \sum_{\tau=1}^{\infty} \delta^\tau \gamma(\tau) \tag{2}$$

compound the *discounted* probability that ρ_t will (first) be up one notch any time in the future. This sum, which has no convenient algebraic form, is plotted in panel (b) of Figure 5. It approaches zero for $\delta \to 0$ for obvious reasons. At the other extreme, it approaches 1 for $\delta \to 1$ because a random walk is *sure* to pass 1 *eventually*, i.e. at some point in the future (this in fact holds for *any* finite number).

With this notation in place, we can formulate this section's key preemption argument:

Proposition 4.2. *If $\Gamma(\delta) > \theta/\hat{\theta}$, then any equilibrium in which some firm $j \in \{1,2\}$ blows the whistle for sure in some state $\rho^k \in R$ must have firm $i \neq j$ blowing the whistle for sure in state ρ^{k-1}.*

See the Appendix for a proof. The key rationale for this results lies in considering a hypothetical situation in which (i) we are in state ρ^{k-1}, (ii) both firms blow the whistle *only* in state ρ^k (symmetric 'cutoff-strategies,' essentially), and (iii) we ignore the fact that firms may be caught independently (i.e., no fine without a report). Firms' discounted expected payoff is then $-\sum_{\tau=1}^{\infty} \delta^\tau \gamma(\tau) \hat{\theta} F = -\Gamma(\delta)\hat{\theta} F$, which represents the sole risk of running into state ρ^i some time in the future, and then paying the fine $-\hat{\theta} F$ from joint reporting. This situation will be inferior to immediately blowing the whistle (alone) whenever $-\theta F > -\Gamma(\delta)\hat{\theta} F$, which is equivalent to the condition required in Proposition 4.2.

The next step is to establish that this hypothetical situation in fact puts a lower bound on how attractive it is for firms to blow the whistle at the outset. This is indeed the tedious part of the formal proof, but the intuition is as follows: Incorporating the probability of independent detection ρ_t serves only to lower the expected payoff from waiting (i.e., from not blowing the whistle immediately). Likewise, raising firm j's propensity to blow the whistle will only make it more attractive for firm i to blow the whistle immediately. Finally, it is an immediate consequence of the game's structure that if it is not optimal for a firm to blow the whistle in state ρ^{k-1} (and given that the other doesn't blow in this or any lower state), then it is also not optimal in any of the states below.

Thus, $\Gamma(\delta) > \theta/\hat{\theta}$ is sufficient to ensure that whistleblowing in any state ρ^k will prompt preemptive whistleblowing already in state ρ^{k-1}. This condition in turn is satisfied for δ high enough (recall that $\Gamma(\delta)$ increases in δ from 0 to 1, whereas $\theta/\hat{\theta} \in (0,1)$), or $\theta/\hat{\theta}$ low enough—which (inversely) represents the relative benefits in fine reduction from blowing the whistle alone rather than together (or, more figuratively speaking, of being the first to knock at the enforcer's door). Both fuel incentives for preemption in an obvious way.

4.3 Preemption Unravels

The condition in Proposition 4.2 bounds preemptive incentives *from below* in a useful way: Because it ignores the risk of independent detection ρ_t, it is altogether independent of the current state of the environment. As such, *if* the preemptive motive kicks in, it kicks in across *all* states. Indeed, straightforward iteration of Proposition 4.2 yields:

Corollary 4.3. *If $\Gamma(\delta) > \theta/\hat{\theta}$, then the only equilibrium in which some firm blows the whistle for sure in* some *state $\rho^k \in R$ is that in which both firms blow the whistle for sure in* all *states.*[14]

That the required condition derives by ignoring the risk of independent detection also implies, though, that the argument thus far is completely detached from any real fundamentals: Rather than the probability of detection, the state may just as well represent meaningless correlating devices such as outside temperature, or even calendar time—the implication being that *if* the unraveling condition is met, the equilibrium must satisfy Corollary 4.3.[15]

All but one of these remaining equilibria disappear, however, if, as in our case, the state variable *does* have a real impact, and if it can reach a state in which blowing the whistle becomes a *dominant* strategy. Formally, letting $v_i(\rho|\sigma_i,\sigma_j)$ denote player i's continuation payoff in any subgame beginning in state $\rho \in R$ given (Markovian) strategies σ_i, σ_j, we say that whistleblowing is a strictly dominant

[14] Note that, allowing for mixed-strategy equilibria, Proposition 4.2 leaves not just equilibria in which both firms always or never blow the whistle, but also equilibria in which they mix—but with a strictly positive probability of *not* blowing the whistle in *all* states.

[15] Of course especially the latter does not follow a random walk. After appropriate recalculation of $\Gamma(\delta)$, though, a very similar (and in fact stronger) unraveling argument applies.

strategy in state ρ if $v_i(\rho\,|\,\sigma_i^b,\sigma_j) > v_i(\rho\,|\,\sigma_i^b,\sigma_j)$ for all σ_j, all σ_i^b with $\sigma_i^b(\rho) = b$, and all σ_i with $\sigma_i(\rho) \neq b$. The whistle will then be blown for sure in such a state, which in turn triggers preemptive whistleblowing in *all* states by simple iteration of Corollary 4.3:

Corollary 4.4. *Suppose that $\Gamma(\delta) > \theta/\hat{\theta}$, and that the grid $R = \{\rho^k\}_{k\in\mathbb{Z}}$ contains a ρ^k for which blowing the whistle is a dominant strategy. Then the unique equilibrium has firms blow the whistle in* all *states.*

We refrain from explicitly characterizing the level of ρ above which whistleblowing becomes dominant.[16] Suffice to say that (i) the second requirement in Corollary 4.4 is clearly met if the grid R contains a $\rho^k > \theta$ (where whistleblowing is obviously dominant), (ii) more generally, using the argument in the proof of Proposition 4.2, the set of states for which whistleblowing is dominant can be identified as those in which a firm finds it optimal to blow the whistle *given that the other never does* (in any state), and (iii) the resulting set is of an obvious cutoff type (containing only high enough states in R).

On a final more technical note, Corollary 4.4 can be strengthened in terms of equilibrium concept. For expositional reasons, we have focussed our discussion on Nash equilibria, but the proof of Proposition 4.2 actually establishes that, whenever firm k blows the whistle for sure in state ρ^k, it is *strictly dominant* for firm $j \neq i$ to blow in state ρ^{k-1}. Thus, Corollary 4.4 can be strengthened in that, under the stated precondition, not only is 'always blowing the whistle' the only Nash equilibrium, but it is in fact the only strategy profile which survives the *iterated elimination of strictly dominated strategies*.

4.4 Discussion

A noteworthy feature about Corollary 4.4 is that the unraveling argument does not require a very *immediate* threat of being independently caught: The state in which whistleblowing becomes dominant may be very distant (as represented by a situation in which the initial state is far lower)—the preemptive unraveling argument can nonetheless completely unfold from there into the present, potentially very 'safe' state.[17] This is reminiscent of arguments in the literature on *global games*, where the mere possibility (no matter how faint) of an extreme signal can trigger a similar cascade across the full range of possible signals (cf. Carlsson and van Damme, 1993; Morris and Shin, 1998).

[16]An explicit characterization is neither simple nor instructive due to the fact that it involves expectations of a (transformed) random walk conditional on random walk not having crossed a certain level yet.

[17]The threat is anything but faint, though, in one respect: even if it may take very long, a random walk is *sure* to hit any level *eventually*. Discounting counters this feature to some extent. This issue can be further defused by introducing a (downward) drift, as we do in the continuous setting below.

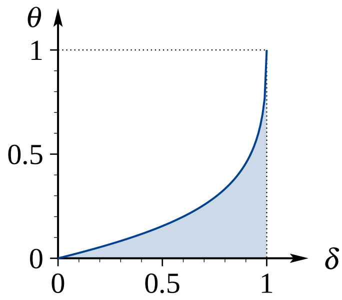

Figure 6: Set of (δ,θ) for which $\Gamma(\delta) > \theta/\hat{\theta}$, given that $\hat{\theta} = (1+\theta)/2$.

What matters instead is that $\Gamma(\delta) > \theta/\hat{\theta}$—that firms view the threat of the state 'moving up a notch' as sufficiently immediate relative to the losses incurred from blowing the whistle together with the other firm rather than alone.

As in the static model (see Footnote 8), any leniency policy with $\theta = 0$ and $\hat{\theta} > 0$ (full amnesty to solitary whistleblowers but incomplete amnesty to joint whistleblowers) will do the trick. If firms are sufficiently worried about the future, though, a less costly policy of *incomplete* amnesty may be just as effective. In this sense, and in contrast to much of the previous literature, the model at hand shifts the focus away from *absolute* levels of leniency reduction towards the importance of the *relative* benefit from blowing the whistle alone rather than together.

Note in this respect, though, that while the model may suggest setting $\hat{\theta}$ close to one (almost no leniency reduction whenever there is more than one applicant) and θ just a bit below $\hat{\theta}\Gamma(\delta)$, this may be problematic on legal grounds. Indeed, a common modeling approach (cf. Harrington, 2013) posits that simultaneous applicants face an equal chance of being considered first or second, and hence $\hat{\theta} = (1+\theta)/2$. The set of (δ,θ) which satisfy $\Gamma(\delta) > \theta/\hat{\theta}$ under this additional assumption are shown in Figure 6. Since $\theta/\hat{\theta}$ still increases in θ under this constraint (even if at a rate lower than one), stronger leniency (lower θ) will still decrease the critical δ above which unraveling takes place. As leniency becomes full ($\theta \to 0$), the critical δ again approaches zero.[18]

Finally, whether unraveling takes place or not will crucially depend on how finely we choose the grid of the discrete model, in the following sense: Suppose we reduce the length of periods, raise δ accordingly, and similarly make the grid for ρ_t finer (so as to make sure that the stochastic process retains a similar variability over time). If we proceed to make the grid sufficiently fine in this manner, then the condition for unraveling will eventually hold. This is troublesome for two interrelated reasons: First, there is no clear natural rationale for how small intervals

[18]Extending the model to $n > 2$ firms (and, consequently, letting $\hat{\theta} = (1+\theta)/n$) will obviously strengthen this preemption effect, so the two-firm case gives a lower bound.

in time and state should be, and this may even reflect policy choices (such as if authorities consider applications which drop in within x days of each other 'joint'). Second, we would not want predictions to depend crucially on an *arbitrary* choice of grid.

The next section picks up on this point by considering the case of *continuous* time and state. Not only does this represent the natural limit case of the above argument of 'making the grid finer'—the continuous-time setting is also much more amenable to more general processes for ρ_t. In particular, it allows us to conveniently consider processes which, unlike the random walk considered above, are *less* than certain to *eventually* hit a region in which whistleblowing is dominant.

Before proceeding to the continuous case, let us still put into perspective two issues in the discrete case. First, what if the condition $\Gamma(\delta) > \theta/\hat{\theta}$ is violated? In this case, the described effect can still lower the threshold ρ above which firms necessarily blow the whistle well below the level above which it is dominant, but the argument need not unravel all the way to the lowest possible ρ on the grid. Second, what about other processes for ρ_t? For instance, we might want to consider processes by which ρ_t can jump more than a notch (or stay constant) from one period to the next, or in which there is a downward drift in that ρ_t is more likely to fall than to rise? Formally, this will result in different (and less conveniently derivable) probabilities $\gamma(\tau)$ and $\Gamma(\delta)$, but leave the remaining argument and results unchanged. In terms of applicability of these results, however, we can no longer be sure that $\lim_{\delta\to 1}\Gamma(\delta) = 1$, and, consequently, that there exists a $\theta/\hat{\theta} < 1$ such that complete unraveling takes place. Again, this motivates the following continuous analysis, where such generalizations are easier to handle.

5 Continuous Time and State: A Limit Result

We assume in this section that the game described above is played in continuous time $t \in [0,\infty)$, where both firms discount the future exponentially at rate $r > 0$. At any t, the risk of independent detection by antitrust authorities is now captured by an instantaneous *hazard rate* $\rho_t \in [0,\infty)$, so that the probability of having been independently detected by any t (the 'failure probability') is $1 - e^{-\int_0^t \rho_t d\tau}$.

We let ρ_t be governed by a transformed random walk $\rho_t = f(X_t)$, where X_t is Brownian motion, and f is a smooth and strictly increasing mapping from $\mathbb{R}$ into $[0,\infty)$. This is a natural continuous-limit extension of the discrete-time process considered in the previous section (and, by the functional central limit theorem, of a much wider class of discrete processes). In contrast to Section 4, we now allow for drift in X_t, so so $X_t = W_t + \mu t$, where W_t is a standard Wiener process. Figure 7 illustrates this class of processes by plotting sample realizations for processes with and without drift (all somewhat arbitrarily originating at $\rho_0 = 0.2$).

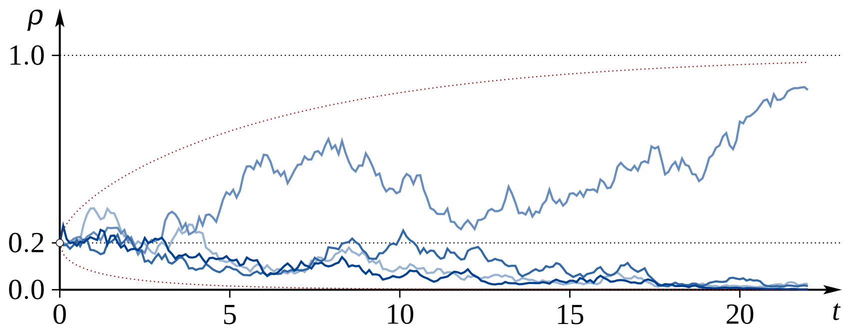

(a) no drift ($\mu = 0$), transformed by $f(x) = \Phi(x/4 - 0.84)$.

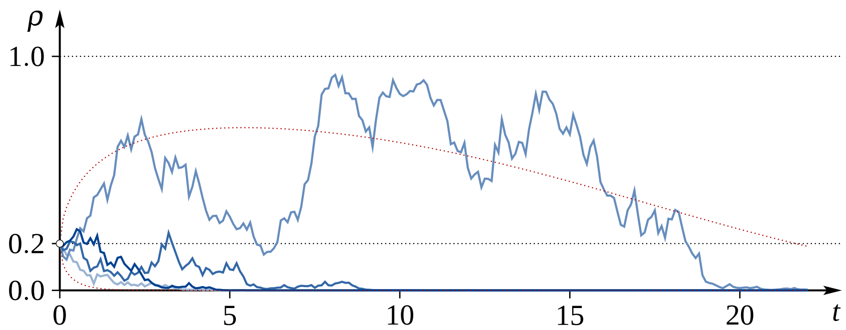

(b) drift $\mu = -1/2$, transformed by $f(x) = \Phi(x/2 - 0.84)$.

Figure 7: Sample Realizations of the Process for ρ_t (4 Draws Each, Including 1%- and 99%-Percentile Bands; $\Phi(\cdot)$ denotes standard-normal cdf).

To avoid cumbersome technicalities, we follow Harrington (2013) in restricting firms to pure (and stationary) strategies of a cut-off type: Each firm i is assumed to blow the whistle if and only if the state ρ_t exceeds a certain threshold $\hat{\rho}_i$.

Clearly, equilibria cannot be asymmetric: Since $\hat{\theta} < 1$, given that one firm blows the whistle in a certain state, it can never be optimal for the other firm not to blow in that same state. All remaining non-degenerate cutoff equilibria, in turn, can be ruled out by a preemption argument similar to that in the discrete case, which yields:

Proposition 5.1. *The only (cutoff-) equilibria are those in which both firms always or never blow the whistle (i.e., in all states or in no states).*

See the Appendix for a proof. Importantly, in contrast to the discrete-time analogue in Proposition 4.2, this result holds for any $\theta/\hat{\theta} < 1$, and for any discount rate $r > 0$.

The intuition for Proposition 5.1 lies in a key property of Brownian motion: That, albeit on a small scale, such a process is sufficiently 'nervous' even over very short time intervals. Specifically, a process starting at zero is certain to take on strictly positive values within any (arbitrarily small) $[0,t]$. Therefore, for any symmetric candidate cutoff-equilibrium with threshold $\hat{\rho}$, there exists a state $\rho < \hat{\rho}$ (sufficiently close) such that the probability of the whistle being jointly blown in the very near future is arbitrarily close to one. For ρ close enough to the threshold $\hat{\rho}$, firms therefore have an incentive to preempt which parallels that in the discrete case—only that it is now altogether independent of firms' patience and the magnitude of $\theta/\hat{\theta}$.

Again as above, the mere possibility of reaching a state ρ in which whistleblowing is dominant is enough to eliminate the 'never report'-equilibrium:

Corollary 5.2. *If $f(\mathbb{R})$ contains an interval of ρ over which whistleblowing is dominant, then the only equilibrium has firms blow the whistle in all states.*

This result provides a strong illustration of the potential power of leniency programs in a dynamic environment: For any $\theta/\hat{\theta} < 1$ and any level of firms' patience, the mere possibility of reaching a state in which whistleblowing is dominant is enough to make immediate whistleblowing the only equilibrium.[19] Interpreting the continuous case as a limit case of Section 4's discrete model, the message of Corollary 5.2 is that discreteness, particularly the size thereof, matters a lot, in that the same holds for a sufficiently fine grid.

Moreover, the result in Corollary 5.2 is more general because it represents a continuous-time limit to a larger class of discrete models than just the example of Section 4, not least by allowing for (downward) drift in the process. And it is easy to see that the result in fact extends to an even broader class of processes than transformed Brownian motion with drift: The key rationale is a *local* one which, roughly speaking, requires that the probability of ρ_t increasing by $\rho + \varepsilon$ within the next instant τ comes close to one for small τ, ε. As such, the argument extends to a broad class of processes which retain this random-walk property in a local sense, such as: (i) the introduction of some kind of mean reversion, or (ii) the addition of random occasional discrete jumps (see also Frankel and Pauzner, 2000, for a discussion of extensions along these lines).

Relative to Section 4's analysis, some technical caveats apply. *First*, the derivation of Corollary 5.2 focussed on stationary, pure strategies of a cutoff type. While stationarity and the cutoff-property are easily seen to be nonessential using the same arguments as in Section 4, an explicit treatment of mixed strategies would be technically cumbersome.[20] Because our discrete analysis *did* allow for mixed

[19] Again, this result is reminiscent of results in the global-games literature, where the uniqueness argument is typically shown to hold when firms' individual signals of the common parameter are perturbed by an *arbitrarily small* amount.

[20] We would need *two* functions to describe generic mixed strategies: A cumulative distribution function which describes the probability that the player has moved by a given time t, and an in-

strategies, we can be sure though that mixed strategies might in the worst case weaken our prediction *in*, but not *around* the limit (i.e., for a discrete, but sufficiently fine grid). *Second*, in contrast to Corollary 4.4, the unique Nash prediction obtained in Corollary 5.2 cannot be strengthened using iterated dominance arguments. This comes from the fact that there generally exists no best response to the other blowing above a certain level $\hat{\rho}_j$: Given almost sure continuity of Brownian motion, a firm would optimally want to choose its cutoff level below, but arbitrarily close.[21] Again, this is a mere technical issue, though, if we view the continuous case as a hypothetical limit case which approximates an arbitrarily fine grid.[22]

6 Discussion and Policy Implications

The analyses of the discrete and the continuous case offer complementary insights on preemptive whistleblowing in an evolving environment: The former illustrates key drivers for preemption when discreteness matters, whereas the latter provides a telling illustration of the effect's potential when this discreteness becomes sufficiently immaterial. This section first discusses crucial limitations and extensions to the two models, and the discusses insights and implied policy recommendations.

6.1 Limitations and Extensions

Particularly the continuous analysis above should certainly be taken with a grain of salt. Not only does it predict immediate dissolution of all cartels, but continuity also takes some assumptions to the extreme—such as that preempting the rival by a literal split second leads to drastically different payoffs than being a split second late. Moreover, by letting the state follow Brownian motion, the continuous case relies on a strong ‘nervousness’ in the state: The probability that *no* news comes in the next instant is *zero*. This feature becomes more crucial for preemption and unraveling the more strongly firms discount and the stronger the downward drift. These caveats in mind, the continuous case is perhaps indeed best interpreted as

tensity function which measures the intensity of atoms in the interval $[t, t+dt]$, where the latter replicates discrete-time results which are lost in passing to the continuous-time limit (see Fudenberg and Tirole, 1985).

[21] A similar issue arises with homogenous Bertrand competition, where marginal-cost pricing is the unique Nash equilibrium, but not obtainable by iterated elimination of dominated strategies (Ståhl, 1972, discretizes the action space to obtain a unique prediction).

[22] Such ‘reverse conclusions’ require a bit of caution, as the continuous model represents the limit to a *larger* class of models than that of Section 4. There is little reason to believe, though, that the corresponding changes to the latter model (adding drift, allowing the discrete process to jump by more than one, etc.) will change these qualitative features (i.e., introduce mixed-strategy equilibria or impede the iterated dominance argument).

a hypothetical but instructive limit case which takes precisely these features to an extreme.

A methodological caveat which concerns both the discrete and continuous case is that the unraveling argument in both cases relies heavily on backward induction arguments. Indeed, the game shares structural similarities with Rosenthal's (1981) centipede game, which raises the question: If players deviate by *not* blowing the whistle at the outset, is it reasonable to assume that they return to subgame-perfection in the ensuing subgame? While this is a valid worry, there is no obvious candidate for an alternative equilibrium concept. Moreover, similarly strong assumptions underlie also the iterated dominance arguments in global games.

A further limitation lies in the fact that, while we assumed firms to be uncertain about the future, we assumed their assessment thereof to be completely symmetric. And indeed, it is easy to imagine situations where firms' assessment of the future (and present) is asymmetrically tainted by private bits of information (or, relatedly, by different judgement of public facts). In a static setting, Harrington (2013) finds that such private information leads to a similar preemptive push, and it seems likely that the two effects complement each other in a dynamic setting. As a simple example, consider a setting where, at each instant t, firm i observes not ρ_t, but $S_{it} \equiv \rho_t + \varepsilon_i$, where the unobservable idiosyncratic errors $\varepsilon_1, \varepsilon_2$ are drawn once at the beginning of the game. This setting will involve intricate updating: As time evolves, firms will update their beliefs on ε_1 and ε_2 by the fact that the other has not yet blown the whistle and that firms have not yet been caught. However, it is clear that there again exist signals for which whistleblowing is dominant. Moreover, as signals move in perfect tender, the unraveling argument will again unfold from there. In fact, we should expect the preemptive unraveling motive to be even stronger: Uncertainty about how close the other is to his threshold introduces the additional risk that the other blows the whistle alone (and I pay the full fine)—much like in the global-games literature.

Finally, our analysis has focussed on one particular source of an uncertain future, namely firms' perceived chance of being caught independently. It seems clear though that the key argument will extend to many other dimensions of the environment, provided that this dimension has the potential to make whistleblowing dominant. For instance, blowing the whistle might reduce a firm's prospect of future collusion (because of higher distrust from other industry members or due to a more alert antitrust authority), which creates an opportunity cost of blowing the whistle. By this channel, anything that affects the future value of collusion (GDP, industry profitability, the interest rate, etc.) can trigger the same kind of preemption.[23] Moreover, such factors might reasonably evolve in a quite continuous fashion—arguably even more so than firms' perceived chance of being caught.

[23] This holds *a fortiori* if collusion is still ongoing when firms decide whether to blow the whistle.

6.2 Implications for Policy and Modeling

The above caveats notwithstanding, the analysis produces useful policy insights. First, the discrete model shows that $\theta/\hat{\theta}$, the ratio of fines granted to a solitary firstcomer and joint whistleblowers, respectively, plays a key role for whether leniency programs can spark the preemptive 'panic' described by the unraveling argument. Generally, this suggests being harsh on firms that fail to preempt fellow colluders. In particular, using a slightly broader interpretation of the model, it proposes a harsh treatment of firms whose leniency application *postdates* the first application. This illustrates an aspect by which leniency programs which essentially preclude fine reductions for latecomers (such as in the US, Israel, and Brazil) may be more effective than programs in which awarding latecomers is not only permitted (such as in the EU), but where this option is frequently exercised.[24]

In contrast to both to the literature on single offenders (the literature following Becker, 1968) and much of the literature on leniency programs (cf. Spagnolo, 2004; Aubert et al., 2006), the model also illustrates that being *relatively* lenient on solitary firstcomers ($\theta/\hat{\theta}$ in our model) may be more important than the *absolute* level of leniency θ. In particular, our models show that it is possible to bust cartels at much lower cost than by setting $\theta < \rho$ (the static requirement): At the continuous-time limit, setting $\hat{\theta}$ close to one and θ just an epsilon below will do the trick.

A bit less obviously, the model's key mechanism relies on firms being eligible for leniency also for high levels of ρ: Whistleblowing being dominant when firms are likely to get caught is a necessary starting point for the preemption argument to unravel also into calmer states. This resonates with policy recommendations in Motta and Polo (2003) and Chen and Rey (2012), who suggest that leniency should be accessible also to firms that are already under investigation by authorities (corresponding to a higher ρ in our model).[25]

Further insights for policy come from the result that the discreteness of the grid is a decisive driver, as illustrated by the continuous limit case. In broad terms, this suggests that policy should do whatever it can to remove obstacles to continuity. As far as discreteness in *time* is concerned, this adds another angle to the above theme of discriminating harshly between first- and latecomers, in terms of differentiating not just with regard to fine levels, but also with regard to time: An effective policy should award firms for being even *just a bit* early, and be harsh on

[24] Needless to say, there are aspects which are outside of the current model which may reinvoke a rationale for awarding late whistleblowers, such as if the latter delivers additional information which significantly reduces the costs of prosecution.

[25] On a related note, to the extent that our literal model considers the problem of self reporting a completed rather than an ongoing crime, policy should impose a time limit neither on firms' liability nor on their access to leniency, as this may similarly jeopardize the (possibly quite distant) anchor for the unraveling argument.

those who are even just a bit late.[26]

As far as discreteness in the *state* is concerned, the model somewhat more vaguely suggests that policy should do whatever it can to let firms' perceived assessment of the environment evolve in many small as opposed to few big steps—firms should essentially expect an update 'any minute.' More specifically, insofar as information dispersed by authorities themselves will influences firms' assessed risk of being caught, this suggests that authorities might want to be rather communicative in terms of spreading any new information in small pieces rather than in accumulated chunks.[27]

Finally, besides policy advice, the analysis also conveys useful modeling advice regarding models of leniency programs in a discrete stochastic environment (and discrete preemption games, more generally). It shows that how finely the discrete grid is chosen can be anything but irrelevant. Arguments which may go through in a two-state model (such as: firms only blow the whistle in the high state, which might explain the type of delays found in Gärtner and Zhou, 2012) may well collapse as soon as we introduce an intermediate third state. A meaningful analysis should thus carefully choose and explain these discontinuities rather than assume them *ad hoc*, as the results may well be driven by precisely these discontinuities.

In conclusion, perhaps the following thought experiment nicely brings out our analysis' basic punchline: Suppose that antitrust authorities *had* the means to convict any offenders if they only put sufficient (possibly very high, but finite) cost and effort into it (by collecting all sorts of data and evidence, forming large expert teams, etc.).[28] Then an effective (and arbitrarily cheap) leniency policy could proceed as follows: For any industry, construct a public indicator which follows a random walk. As soon as this indicator hits a certain predefined level, investigate the industry with all possible rigor (so that, if there *was* an offense it is found with probability one). By the above analysis, the trigger levels for the indicators could be chosen arbitrarily high, meaning a rigorous investigation is very unlikely to ever occur—but the unraveling effect would kick in. Thus, with very low discounted and expected costs of investigation (coming from the arbitrarily unlikely event of

[26] This is also related to a point made in Motchenkova (2004), by which leniency programs should be 'confidential' in that authorities should hold back on publicly announcing leniency applicants, so as to keep fellow offenders from immediately following suit. In Motchenkova's continuous-time model, 'non-confidential' leniency programs give firms the option of *instantaneously* responding to whistleblowing by the rival, thus qualifying for joint leniency. Our continuous-time analysis abstracts from this possibility, which seems natural given our interpretation thereof as a limit to the discrete case.

[27] A wholesome argument along these lines would of course need to explicitly model the informational asymmetries which arise if authorities *withhold* information, and the ensuing repercussions on firms' beliefs.

[28] This might be justified along the lines of Aubert et al.'s (2006) Assumption 1, whereby collusion always generates hard evidence which can be found by authorities.

an investigation), authorities could achieve an extreme effect.

7 Conclusion

This paper has presented an analysis of corporate leniency application in the context of a dynamically evolving environment. We have shown how firms' symmetric uncertainty about the environment's future evolution can trigger strategic incentives for preemption. At the extreme, these incentives can unravel to the point where firms blow the whistle long before the state of the environment itself gives them any real reason to do so. We feel that this aspect captures an important aspect of the distrust and strategic uncertainty which leniency programs are thought to create. As such, it also leads to clear advice regarding policy and optimal design of leniency programs, by suggesting that a strong discrimination between first- and latecomers may be more crucial than offering high fine reductions in absolute terms, and by suggesting that authorities might be better off spreading information on its investigation efforts in small bits rather than in lumps.

As regards directions for further research, a key insight of our analysis is that, in an evolving world, discontinuities in time and state are essentially the only obstacle to a maximally effective leniency policy which elicits immediate self reporting at arbitrarily low cost. We have loosely discussed some examples of factors which might affect the magnitude of discontinuities, such as how authorities differentiate applications that are made at different but very close points in time. But given their very crucial role, it should be worth further and more rigorously investigating candidate reasons for discontinuities and how policy might address them.

Concerning discontinuities in the state (or, more precisely, firms' assessment thereof), it should further be interesting to explicitly endogenize authorities' information policy. Our preliminary insights suggest that a rather 'talkative' policy might be better in terms of eliminating discontinuities, but how this peters out in a full-fledged model remains to be seen—specifically if we explicitly allow for private information not only on authorities' behalf, but if we also allow authorities to spread asymmetric, and possibly even misleading, information among firms. Eventually, such an approach should be able to refine our somewhat crude preliminary insights on how antitrust authorities might optimally disseminate their information among potential colluders so as to create a maximal degree of distrust, and thereby preemptive whistleblowing.

Appendix A. Proofs

Proof of Proposition 4.2. For any pair of (Markovian) strategies $\sigma = (\sigma_1, \sigma_2)$ and any state $\rho \in R$, let $v_i(\rho \,|\, \sigma_i, \sigma_j)$ denote player i's expected continuation payoff in any subgame beginning in state ρ.

We prove the result in a sequence of steps. The first step argues that ignoring the possibility of independent detection only makes blowing the whistle less attractive. Formally, let Γ denote the original game, and let Γ° denote the altered game in which we ignore independent detection (i.e., firms are caught with a probability of zero rather than ρ_t). Then we have:

Lemma A.1. *If in any state $\rho \in R$ it is strictly dominant for firm i to blow the whistle in Γ°, then this must hold also in Γ.*[29]

To see this, let $v_i^\circ(\rho \,|\, \sigma_i, \sigma_j)$ denote i's expected continuation payoff in Γ°, and observe that

$$v_i^\circ(\rho \,|\, \sigma_i, \sigma_j) \geqslant v_i(\rho \,|\, \sigma_i, \sigma_j), \qquad \text{for all } \sigma_i, \sigma_j \tag{A.1}$$

and any state $\rho \in R$. This is because, being a convex combination of zero and the three possible fines $-F < -\hat{\theta}F < -\theta F$, both $v_i(\rho \,|\, \sigma_i, \sigma_j)$ and $v_i^\circ(\rho \,|\, \sigma_i, \sigma_j)$ are bounded from below by F. Hence, in any period, the continuation value from *not* being caught cannot be lower than the payoff from being caught, which implies that lowering the latter to zero in any period can only increase payoffs.

Now in both games Γ and Γ°, a strategy of whistleblowing in the concurrent state will yield the same payoff: $v_i^\circ(\rho \,|\, \sigma_i, \sigma_j) = v_i(\rho \,|\, \sigma_i, \sigma_j) = -[\sigma_j(\rho)\hat{\theta} + (1 - \sigma_j(\rho))\theta]F$ for $\sigma_i(\rho) = b$. Combined with (A.1), this establishes Corollary A.1: For any two strategies σ_i, σ_i^b with $\sigma_i^b(\rho) = b$, $v_i^\circ(\rho \,|\, \sigma_i^b, \sigma_j) > v_i^\circ(\rho \,|\, \sigma_i, \sigma_j)$, $\forall \sigma_j$, implies $v_i(\rho \,|\, \sigma_i^b, \sigma_j) > v_i(\rho \,|\, \sigma_i, \sigma_j)$ $\forall \sigma_j$. Note that Corollary A.1 allows us to establish Proposition 4.2 by showing that, under the stated conditions, blowing the whistle is strictly dominant in state ρ^{k-1} of the altered game Γ°.

While strict dominance requires us to establish optimality against essentially arbitrary strategies of the other player (the only restriction being that the proposition assumes the other player to blow at ρ^k), we argue next that it suffices to show optimality against one particular strategy. To this end, for any player i, let Σ^k denote the set of strategies σ_i for which $\sigma_i(\rho^k) = b$ (the player blows the whistle in state ρ^k), and let $\hat{\sigma}^k \in \Sigma^k$ denote the strategy for which $\sigma_i(\rho^{k'}) = \not{b}$ for all $k' \neq k$ (the whistle is not blown in any other states). Then the next step can be formulated as follows:

[29] Precisely speaking, 'blowing the whistle in state ρ' describes the *set* of strategies σ_i with $\sigma_i(\rho) = b$ (and any possible action in other states $\rho' \neq \rho$). As such, 'blowing the whistle being strictly dominant' means that any strategy in this set strictly dominates any strategy in the set's complement, or, equivalently, that any strategy without certain whistleblowing in state ρ is strictly dominated. These technicalities notwithstanding, the former formulation seems more accessible.

Lemma A.2. *In the game Γ°, if it is strictly optimal for player i to blow the whistle in state ρ^{k-1} given that the other player plays $\sigma_j = \hat{\sigma}^k$, then this is also strictly optimal given any other strategy $\sigma_j \in \Sigma^k$ for j.*

To see this, observe first that for any state $\rho^{k'} \in R$,we have

$$v_i^\circ(\rho^{k'}|\sigma_i,\hat{\sigma}^k) \geqslant v_i^\circ(\rho^{k'}|\sigma_i,\sigma_j), \qquad \text{for any } \sigma_i \text{ and all } \sigma_j \in \Sigma^k, \tag{A.2}$$

i.e. given that the other player j blows the whistle at ρ^k, assuming that he doesn't blow the whistle in any other states can only increase i's payoff. This follows simply because player j blowing the whistle in any period unanimously lowers i's payoff from $-\theta F$ to $-\hat{\theta}F$ if i blows the whistle as well, and from 0 to $-F$ if i doesn't.

Now player i's benefit from blowing vs. not blowing the whistle in ρ^{k-1} given $\sigma_j = \hat{\sigma}^k$ (and any continuation strategy σ_i for player i) is

$$-\theta F - \delta\Big[\tfrac{1}{2}v_i^\circ(\rho^k|\sigma_i,\hat{\sigma}^k) + \tfrac{1}{2}v_i^\circ(\rho^{k-2}|\sigma_i,\hat{\sigma}^k)\Big]. \tag{A.3}$$

In contrast, i's benefit from blowing vs. not blowing given any other strategy σ_j is

$$\begin{aligned} &\sigma_j(\rho^{k-1})(1-\hat{\theta})F + (1-\sigma_j(\rho^{k-1}))\Big[-\theta F - \delta\big[\tfrac{1}{2}v_i^\circ(\rho^k|\sigma_i,\sigma_j) + \tfrac{1}{2}v_i^\circ(\rho^{k-2}|\sigma_i,\sigma_j)\big]\Big] \\ &\quad \geqslant (1-\sigma_j(\rho^{k-1}))\Big[-\theta F - \delta\big[\tfrac{1}{2}v_i^\circ(\rho^k|\sigma_i,\hat{\sigma}^k) + \tfrac{1}{2}v_i^\circ(\rho^{k-2}|\sigma_i,\hat{\sigma}^k)\big]\Big], \end{aligned} \tag{A.4}$$

where the inequality uses the fact that the first term in the left-hand side is positive (recall $\hat{\theta} < 1$), and bounds the second term from below using (A.2). Hence, if (A.3) is positive, then so is the left-hand side of (A.4), which establishes Lemma A.2. By Lemma A.2, we can establish Proposition 4.2 by establishing optimality of whistleblowing for player i in state ρ^{k-1} of Γ° *given* that the other player j *only* blows the whistle in ρ^k.

The next step in our argument is the following:

Lemma A.3. *Consider the game Γ°, and assume that player j plays $\sigma_j = \hat{\sigma}^k$. If it is* not *optimal for i to blow the whistle in state ρ^{k-1}, then it must be optimal for him to play $\sigma_i = \hat{\sigma}^k$.*

To see this, notice first that the former strategy yields $-\theta F$, so that it can only be improved upon by a strategy σ_i for which $v_i^\circ(\rho^{k-1}|\sigma_i,\hat{\sigma}^k) \geqslant -\theta F$. Notice next that in the subgame beginning at ρ^k, j blowing the whistle for sure in this state implies that i's optimal strategy will be to do the same, so $\sigma_i(\rho^k) = b$, which implies $v_i^\circ(\rho^k|\sigma_i,\hat{\sigma}^k) = -\hat{\theta}F$. Since $\theta < \hat{\theta}$, we thus have $v_i^\circ(\rho^{k-1}|\sigma_i,\hat{\sigma}^k) > v_i^\circ(\rho^k|\sigma_i,\hat{\sigma}^k)$ for any strategy which beats blowing the whistle in ρ^{k-1}.

Next, we use the following iterative argument: Fix any strategies σ_i and $\sigma_j \in \Sigma^k$, and suppress them to ease notation. Then, if σ_i maximizes player i's payoff in $\hat{\Gamma}$, then for any state k' such that $\sigma_i(\rho^{k'}) = \not b$ and $v_i^\circ(\rho^{k'}) \geqslant v_i^\circ(\rho^{k'+1})$, we must have $v_i^\circ(\rho^{k'-1}) \geqslant v_i^\circ(\rho^{k'})$ and $\sigma_i(\rho^{k'-1}) = \not b$. To see this, notice that $\sigma_i(\rho^{k'}) = \not b$ implies that the payoff in state ρ^k is given by

$$v_i^\circ(\rho^{k'}) = \delta\left[\tfrac{1}{2}v_i^\circ(\rho^{k'-1}) + \tfrac{1}{2}v_i^\circ(\rho^{k'+1})\right],$$

which, since $\delta < 1$, can be rearranged to give

$$v_i^\circ(\rho^{k'-1}) > v_i^\circ(\rho^{k'}) + \left[v_i^\circ(\rho^{k'}) - v_i^\circ(\rho^{k'+1})\right].$$

Hence, $v_i^\circ(\rho^{k'}) \geqslant v_i^\circ(\rho^{k'+1})$ and $\sigma_i(\rho^{k'}) = \not b$ imply that $v_i^\circ(\rho^{k'-1}) > v_i^\circ(\rho^{k'})$. But $\sigma_i(\rho^{k'}) = \not b$ being optimal for i implies $v_i^\circ(\rho^{k'}) \geqslant -\theta F$, and hence $v_i^\circ(\rho^{k'-1}) > -\theta F$, so that it must be optimal for i not to blow the whistle also in state $k' - 1$ (i.e., $\sigma_i(\rho^{k'-1}) = \not b$). Taken together, we thus have: For any optimal strategy with $\sigma_i(\rho^{k'}) = \not b$ and $v_i^\circ(\rho^{k'}) \geqslant v_i^\circ(\rho^{k'+1})$, we must have $v_i^\circ(\rho^{k'-1}) \geqslant v_i^\circ(\rho^{k'})$ and $\sigma_i(\rho^{k'-1}) = \not b$.

This argument can now be iterated downward beginning at $k' = k-1$: There, we already know that *if* the strategy is optimal and has $\sigma_i(\rho^{k-1}) = \not b$, then $v_i^\circ(\rho^{k-1}) > v_i^\circ(\rho^k)$, which implies by the above result that the strategy must have $\sigma_i(\rho^{k-2}) = \not b$ and $v_i^\circ(\rho^{k-2}) > v_i^\circ(\rho^{k-1})$, so that the same argument can be made at $k' = k-2$, and so forth. We thus find that the only strategy which can yield player i a higher payoff than blowing the whistle straightaway in ρ^{k-1} is σ^k, which establishes Lemma A.3.

Taken together, Lemma A.1 through Lemma A.3 imply that in state ρ^{k-1} of the original game, given that player j blows the whistle at ρ^k, it can only be optimal for i not to blow the whistle if $v_i^\circ(\rho^{k-1}|\hat{\sigma}^k,\hat{\sigma}^k) \geqslant -\theta F$. But because $v_i^\circ(\rho^{k-1}|\hat{\sigma}^k,\hat{\sigma}^k) = -\sum_{\tau=1}^{\infty}\delta^\tau\gamma(\tau)\hat{\theta}F = -\Gamma(\delta)\hat{\theta}F$, this is equivalent to the condition stated in the proposition. □

Proof of Proposition 5.1. It is useful to introduce the following notation: For any path ρ_t with starting point $\rho_0 = \rho$ and any $\hat{\rho}$, let $\tau_{\rho_t}(\hat{\rho}\,|\,\rho) \equiv \inf_{t>0}\{\rho_t \geqslant \hat{\rho}\}$ denote the instant t where the path first passes $\hat{\rho}$. Let $\hat{v}(\rho_t\,|\,\hat{\rho})$ denote each firm's continuation payoff given that the current state is ρ_t, and given that firms each blow the whistle at any t if and only if $\rho_t \geqslant \hat{\rho}$. Now, as in the discrete case, we obtain an upper bound on $\hat{v}(\rho_t\,|\,\hat{\rho})$ by ignoring the risk of independent detection, as this only leads to *earlier* fine payments (of at least the same magnitude). Thus,

$$\hat{v}(\rho\,|\,\hat{\rho}) \leqslant -\hat{\theta}F\int_0^\infty e^{-rt}d\,\mathrm{Pr}(\tau_{\rho_t}(\hat{\rho}\,|\,\rho) \leqslant t), \qquad \text{for } \rho < \hat{\rho}, \tag{A.5}$$

where $d\,\mathrm{Pr}(\tau_{\rho_t}(\hat{\rho}\,|\,\rho) \leqslant t)$ is simply the instantaneous probability that ρ_t crosses $\hat{\rho}$ at any instant t, in which case a fine of $-\hat{\theta}F$ results.

Now at any $\rho < \hat{\rho}$, instead of adhering to the prescribed action of *not* blowing the whistle, a firm can instead blow, which yields an instantaneous payoff of $-\theta F$ and ends the game, instead of receiving the continuation payoff $\hat{v}(\rho\,|\,\hat{\rho})$. Combined with (A.5), a necessary condition for equilibrium is therefore that

$$\int_0^\infty e^{-rt} d\,\Pr(\tau_{\rho_t}(\hat{\rho}\,|\,\rho) \leqslant t) \leqslant \theta/\hat{\theta}, \qquad \text{for all } \rho < \hat{\rho}. \tag{A.6}$$

Now the left-hand side of (A.6) can be further bounded from below as follows: Fix any $\tilde{t} \geqslant 0$. Then

$$\begin{aligned}\int_0^\infty e^{-rt} d\,\Pr(\tau_{\rho_t}(\hat{\rho}\,|\,\rho) \leqslant t) &> \int_0^{\tilde{t}} e^{-rt} d\,\Pr(\tau_{\rho_t}(\hat{\rho}\,|\,\rho) \leqslant t) \\ &> e^{-r\tilde{t}} \int_0^{\tilde{t}} d\,\Pr(\tau_{\rho_t}(\hat{\rho}\,|\,\rho) \leqslant t) = e^{-r\tilde{t}} \Pr(\tau_{\rho_t}(\hat{\rho}\,|\,\rho) \leqslant \tilde{t}),\end{aligned}$$

where the first inequality follows because the integrand is strictly positive, and the second because e^{-rt} is strictly decreasing in t (so $e^{-rt} > e^{-r\tilde{t}}$ for all $t < \tilde{t}$). Thus, a necessary condition for (A.6) is

$$\Pr(\tau_{\rho_t}(\hat{\rho}\,|\,\rho) \leqslant \tilde{t}) \leqslant e^{r\tilde{t}}\theta/\hat{\theta}, \qquad \text{for all } \rho < \hat{\rho} \text{ and } \tilde{t} \geqslant 0. \tag{A.7}$$

Notice that the right-hand side approaches $\theta/\hat{\theta} < 1$ for $\tilde{t} \to 0$. Thus, for any $\varepsilon > 0$, we can find a $\tilde{t}$ such that the right-hand side of (A.7) falls strictly short of $1 - \varepsilon$. For any $\varepsilon > 0$, there must therefore exist a $\tilde{t}$ such that

$$\Pr(\tau_{\rho_t}(\hat{\rho}\,|\,\rho) \leqslant \tilde{t}) \leqslant 1 - \varepsilon, \qquad \text{for all } \rho < \hat{\rho}. \tag{A.8}$$

But this contradicts basic properties of Brownian motion, as described by the following auxiliary result:

Lemma A.4. *For any* $\tilde{t} > 0$ *and* $\varepsilon > 0$, *there exists* $\rho < \hat{\rho}$ *such that* $\Pr(\tau_{\rho_t}(\hat{\rho}\,|\,\rho) \leqslant \tilde{t}) > 1 - \varepsilon$.

To see this, recall that $\rho_t = f(X_t)$, where $X_t = W_t + \mu t$ is Brownian motion with drift μ, and $f(\cdot)$ is a smooth and strictly increasing function. Thus, letting $\tau_{X_t}(a) \equiv \inf_{t>0}\{X_t \geqslant a\}$ denote the time at which the process X_t (starting at 0) first crosses $a > 0$, we have

$$\Pr(\tau_{\rho_t}(\hat{\rho}\,|\,\rho) \leqslant \tilde{t}) = \Pr(\tau_{X_t}(f^{-1}(\hat{\rho}) - f^{-1}(\rho)) \leqslant \tilde{t}).$$

By standard properties of Brownian motion with drift, the probability of first passing any $a > 0$ by $\tilde{t} > 0$ is given by

$$\Pr(\tau_{X_t}(a) \leqslant \tilde{t}) = 1 - \Phi\left(\frac{a - \mu t}{\sqrt{t}}\right) + e^{2\mu a}\Phi\left(\frac{-a - \mu t}{\sqrt{t}}\right),$$

where $\Phi(\cdot)$ denotes the cdf of the standard normal distribution. For any $\tilde{t} > 0$, this probability continuously approaches 1 as $a \searrow 0$. Thus, for any $\tilde{t} > 0$, we can find an $a > 0$ such that $\Pr(\tau_{X_t}(a) \leqslant \tilde{t}) > 1 - \varepsilon$. But then choosing ρ such that $f^{-1}(\rho) = f^{-1}(\hat{\rho}) - a$ will satisfy both $\Pr(\tau_{\rho_t}(\hat{\rho}|\rho) \leqslant \tilde{t}) > 1 - \varepsilon$ and $\rho < \hat{\rho}$, as required by Lemma A.4. □

References

Angeletos, George-Marios, Christian Hellwig, and Alessandro Pavan (2007) 'Dynamic global games of regime change: Learning, multiplicity, and the timing of attacks.' *Econometrica* 75(3), 711–756

Aubert, Cécile, Patrick Rey, and William E. Kovacic (2006) 'The impact of leniency and whistle-blowing programs on cartels.' *International Journal of Industrial Organization* 24(6), 1241–1266

Becker, Gary S. (1968) 'Crime and punishment: An economic approach.' *Journal of Political Economy* 76(2), 169–217

Burdzy, Krzysztof, David Frankel, and Ady Pauzner (2001) 'Fast equilibrium selection by rational players living in a changing world.' *Econometrica* 69(1), 163–189

Carlsson, Hans, and Eric van Damme (1993) 'Global games and equilibrium selection.' *Econometrica* 61(5), 989–1018

Chen, Zhijun, and Patrick Rey (2012) 'On the design of leniency programs.' Mimeo, forthcoming in *Journal of Law and Economics*

Choi, Jay Pil, and Heiko Gerlach (2009) 'Multi-market collusion with demand linkages and antitrust enforcement.' Mimeo

Frankel, David, and Ady Pauzner (2000) 'Resolving indeterminacy in dynamic settings: The role of shocks.' *Quarterly Journal of Economics* 115(1), 285–304

Fudenberg, Drew, and Jean Tirole (1985) 'Preemption and rent equalization in the adoption of new technology.' *Review of Economic Studies* 52, 383–401

Gärtner, Dennis L. and Jun Zhou (2012) 'Delays in leniency application: Is there really a race to the enforcer's door?' Mimeo

Harrington, Jr., Joseph E. (2008) 'Optimal corporate leniency programs.' *Journal of Industrial Economics* 56(2), 215–246

_ (2013) 'Corporate leniency programs when firms have private information: The push of prosecution and the pull of pre-emption.' *Journal of Industrial Economics* 61(1), 1–27

Lefouili, Yassine, and Catherine Roux (2012) 'Leniency programs for multimarket firms: The effect of amnesty plus on cartel formation.' *International Journal of Industrial Organization* 30(6), 624–640

Marshall, Robert C., Leslie M. Marx, and Claudio Mezzetti (2013) 'Antitrust leniency with multi-product colluders.' Mimeo

Mason, Robin, and Helen Weeds (2010) 'Investment, uncertainty and preemption.' *International Journal of Industrial Organization* 28(3), 278–287

Morris, Stephen, and Hyun S. Shin (1998) 'Unique equilibrium in a model of self-fulfilling currency attacks.' *American Economic Review* 88(3), 1769–1787

Motchenkova, Evguenia (2004) 'Effects of leniency programs on cartel stability.' Mimeo

Motta, Massimo, and Michele Polo (2003) 'Leniency programs and cartel prosecution.' *International Journal of Industrial Organization* 21(3), 347–379

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

Rosenthal, Robert W. (1981) 'Games of perfect information, predatory pricing, and the chain store.' *Journal of Economic Theory* 25(1), 92–100

Selten, Reinhard, and John Harsanyi (1988) *A General Theory of Equilibrium Selection in Games* (Cambridge MA: MIT Press)

Spagnolo, Giancarlo (2004) 'Divide et impera: Optimal leniency programs.' Mimeo

_ (2008) 'Leniency and whistleblowers in antitrust.' In *Handbook of Antitrust Economics*, ed. Paolo Buccirossi (Cambridge, MA: MIT Press) pp. 259–304

Ståhl, Ingolf (1972) *Bargaining Theory* (Stockholm: Stockholm University)