THE JOURNAL OF INDUSTRIAL ECONOMICS 0022-1821
Volume LXIII June 2015 No. 2

PROSECUTION AND LENIENCY PROGRAMS: THE ROLE OF BLUFFING IN OPENING INVESTIGATIONS*

JULIEN SAUVAGNAT†

This paper characterizes the optimal investigation and leniency policies when the Competition Authority is privately informed about the strength of a cartel case. I show that the Competition Authority can then exploit firms' uncertainty about the risk of conviction to obtain confessions even when the case is weak. More generally, I show that offering full leniency allows the Competition Authority to open more successful investigations (what I refer to as the 'activism effect' of leniency), which overall raises both cartel desistance and cartel deterrence. Finally, I discuss the policy implications of the model.

'Io so. Ma non ho le prove. Non ho nemmeno indizi.'[1]

Pier Paolo Pasolini. *Il Corriere della Sera.* 1974.

I. INTRODUCTION

THIS PAPER CHARACTERIZES THE OPTIMAL INVESTIGATION AND LENIENCY POLICIES when the competition authority has private information on how strong its case is against cartels. It is shown that the competition authority can then exploit firms' uncertainty about the probability of conviction to trigger reporting even when the case is weak, which improves both cartel desistance and deterrence.

Both the Department of Justice (DoJ) in the United States (U.S.) and the Commission in the European Union (EU) have made breaking up cartels a top priority. This is reflected in the recent sharp increase in the total amount of penalties inflicted on cartel members.[2] In the EU, the record fines of €1,470 million against computer monitor tubes producers in 2012 is an illustration.

*I am very grateful to Patrick Rey for his guidance from the early stages of this research. I also thank two anonymous referees, Philippe Choné, Chiara Fumagalli, Bruno Jullien, Kambiz Mohkam, Massimo Motta, Michele Polo, Jérôme Renault, Giancarlo Spagnolo and Thibaud Vergé for very useful suggestions and advice, as well as participants at the ENTER Jamboree meeting (2010) in Toulouse, the European Economic Association (EEA) Congress (2010) in Glasgow, the European Association Research in Industrial Economics (EARIE) Conference (2010) in Istanbul, the Association of Southern European Economic Theorists (ASSET) Meeting (2010) in Alicante and the CDSE Seminar in Mannheim for helpful comments.

†Author's affiliations: Toulouse School of Economics (TSE) and ENSAE-CREST, 92245 Malakoff, France.
e-mail: julien.sauvagnat@gmail.com

[1] I know. But I do not have the proof. I do not even have clues.

[2] In the U.S. corporate fines averaged $315 million per year between 1995 and 1999 and $628 million between 2005 and 2009. In the EU, the rise is even more impressive, going from

A path-breaking development in the U.S. antitrust policy was the revision of its corporate leniency program in 1993. In the initial version,[3] discretionary leniency was granted to confessors provided that the cartel was not already under investigation. The new program instead automatically guarantees full amnesty to the first firm who blows the whistle; in addition, some leniency may be granted when an investigation has begun. Finally, the DoJ added a complementary individual leniency program in 1994 that protects individual informants from pecuniary fines or criminal sanctions.[4]

For cartels that are too stable to be deterred, competition authorities must try to detect them and break them down. The conviction of the car-glass makers in 2008 highlights the role of post-investigation leniency in the prosecution of cartels. According to the European Commission:

> *The Commission started this investigation on its own initiative on the basis of reliable information provided by an anonymous informant. The information prompted the Commission to carry out surprise inspections in 2005 at several sites of car glass producers in Europe. After the inspections, the Japanese Asahi Glass Co. and its European subsidiary AGC Flat Glass Europe (formerly 'Glaverbel) filed an application under the 2002 Leniency Notice. . .. Asahi/Glaverbel cooperated fully with the Commission and provided additional information to help to expose the infringement and its fine was reduced by 50%.*[5]

The above decision suggests that it might be optimal for the competition authority to run an investigation against a cartel, pretending that conviction is likely, in the hope that firms will denounce the cartel. I build a model to explore this idea.

Formally, I consider a standard supergame in which firms decide whether to compete or collude on the market. The Competition Authority cannot condemn the cartel without gathering hard evidence, and receives a binary signal (either good or bad) which determines the probability of convicting cartel members if an investigation is launched, a bad signal meaning a low probability. As cartel members do not observe the signal, they may confess their illegal activities even when the Competition Authority opens an investigation after receiving a bad signal. In such a case, the cartel is condemned even though the Competition Authority was unlikely to convict firms by its own means.

€54 million per year between 1995 and 1999 to €1,951 million between 2005 and 2009 (data from the U.S. Antitrust Division and the European Commission).

[3] The DoJ introduced its first corporate leniency program in 1978.

[4] In the EU, violations of Article 81 and 82 are not punished with criminal fines. In particular, individuals are not subject to imprisonment.

[5] Available at *http://ec.europa.eu/competition/cartels/what_is_new/news.html*, reference: IP/08/1685.

This paper is related to Motta and Polo [2003]. They show that offering leniency increases firms' incentives to report information which enables the Competition Authority to save on prosecution costs. On the other hand, by reducing expected sanctions, leniency increases *ex ante* firms' incentives to form new cartels. As a result, the leniency program is only a second-best instrument: if the budget of the Competition Authority is large enough to deter cartel formation, leniency should not be offered. One important innovation in my paper is the incorporation of private information on the Competition Authority side. I show that this informational advantage, combined with leniency, raises the conviction rate and thereby enhances cartel enforcement. In the model, more leniency is desirable because it allows the Competition Authority to open more successful investigations.

Recent research incorporates private information in theoretical models of leniency. In this vein, my work is closely related to that of Harrington [2013], who studies the effect of leniency when each cartel member receives private information about the likelihood of conviction. This generates a *pre-emption* effect—in that a firm may apply for leniency because it fears another firm will apply.

Finally, this work is related to the literature on self-reporting (e.g. Malik [1993] and Kaplow and Shavell [1994]) as well as on plea bargaining. Reinganum [1988] considers for example a privately informed prosecutor offering plea bargains to a defendant; the prosecutor's offer then reveals the strength of the case against the defendant. In my paper, the Competition Authority is privately informed about the strength of a cartel case, and its investigation policy affects firms' decisions to report information.

This paper has important policy implications. In particular, unlike the recommendations of the U.S. and European leniency policies, the results suggest that leniency should be granted even when the risk of conviction is large.

The paper is organized as follows. Section II describes the model. Section III provides a benchmark where there is no information asymmetry between the Competition Authority and cartel members. Section IV studies the case where the Competition Authority has private information about the strength of the cartel case. In section V, I discuss the policy implications of the model. Finally, section VI concludes. All proofs are in the appendix.

II. THE MODEL

II(i). *Players*

Firms. Consider a continuum of industries with unit mass. In each industry, $N \geq 2$ symmetric and risk-neutral firms play an infinitely repeated game. In each period, each firm decides whether to compete or collude. The gross profit of a firm is:

- 0 if firms compete,
- $\Pi^C > 0$ if firms collude,
- $\Pi^D \geq \Pi^C$ if the firm deviates from collusion—that is, if it competes when the other firm(s) collude.

For example, in a standard Bertrand oligopoly in which N firms produce an homogenous good with the same unit variable cost, static price competition drives profits to 0, the benefit from collusion is equal to a share $\frac{1}{N}$ of the monopoly profits $\left(\Pi^C = \frac{\Pi^M}{N}\right)$ whereas a deviation brings the whole monopoly profits ($\Pi^D = \Pi^M$).

In order to analyze the impact of the antitrust policy on cartel formation, we follow Harrington and Chang [2009] and assume that industries are heterogenous with respect to the gains from deviating, Π^D.[6]

We focus on explicit collusion, which is based on communication, meetings and so forth.[7] We therefore assume that collusion generates hard evidence that can be used to condemn the cartel. Evidence of collusion lasts only one period, which implies that a cartel cannot be prosecuted for its past activity.[8] Finally, all firms have the same discount factor $\hat{\delta} \in (0,1)$ and maximize the expected discounted sum of their profits.

Competition Authority (CA). The CA cannot condemn a cartel without gathering hard evidence. Conditional on there being a cartel, the CA receives a binary signal, either good or bad, which determines the probability of finding hard evidence if it launches an investigation. The probability of receiving a good signal is ψ, in which case the CA knows that it will find hard evidence of collusion with probability $\mu \in (0, 1)$ if it runs an investigation. If instead the signal is bad, the CA has no chance of finding hard evidence on its own.[9] Finally, if firms do not collude, the CA receives a bad signal with probability $\beta > 0$, and no signal otherwise.

[6] As mentioned in Harrington and Chang [2009], this heterogeneity may reflect differences in the number of firms or in the elasticity of firm demand.

[7] In their study of the Sugar Institute Case, Genesove and Mullin [2001] show that communication helps firms collude. For theoretical works considering the role of communication in collusion, see e.g., Compte [1998], Athey and Bagwell [2001] and Aoyagi [2002].

[8] In practice, prosecution is possible even after firms stopped colluding. However, allowing this possibility would seriously complicate the analysis.

[9] The model delivers qualitatively similar results when the probability of finding hard evidence after receiving a bad signal is relatively low.

In each period, the CA decides whether or not to run an investigation based on its signal.[10] If the CA gathers evidence, the cartel is condemned and each member must pay a fine F, which is exogenously set by law.[11]

Leniency Program. For the sake of exposition, we will assume that a firm which deviates from collusion faces no risk of being convicted. This rules out any role for pre-investigation leniency (see Spagnolo [2004] and Rey [2003]).[12] The CA can however offer a leniency rate q *during* investigations. Eligible firms then pay only a reduced fine (equal to $(1 - q)F$) if they report information, in which case all cartel members are condemned with probability one. We assume that firms decide whether or not to report simultaneously, and restrict our attention to leniency rates lower than one.[13] Finally, we assume that firms compete following a condemnation.[14,15] Enforcing competition can for example be achieved through either close monitoring of the industry or higher fines for repeat offenders.[16]

The literature has shown that adopting a first informant rule—i.e., restricting eligibility to the first confessor—is preferred to granting leniency to all informants. This is also the case in this model.[17] To facilitate the exposition of the results, we therefore directly assume that leniency is granted only to the first informant.[18]

[10] In the absence of signal, the CA knows that firms compete and it is thus not optimal to open an investigation.

[11] Alternatively, F may be viewed as the maximal punishment allowed by the law. In this case, Becker's [1968] argument applies: it is optimal to set fines as high as possible.

[12] To see this, consider as in the previous version of the paper (Sauvagnat [2010]) that deviating firms still face a risk of being convicted and that the CA offers a pre-investigation leniency rate $q_0 \leq 1$, available before the beginning of any investigation. Note that reporting before an investigation is not a sustainable collusive strategy: since the continuation value of collusion is zero, as the cartel will be shut down forever, firms would deviate and compete rather than collude. It follows that pre-investigation leniency has no pro-collusive effects, and thus offering full amnesty in case of pre-investigation reports is always optimal ($q_0 = 1$). This allows defecting cartel members to report and avoid paying the fine.

[13] Well-designed reward schemes are very effective in fighting collusion (see Spagnolo [2004] and Aubert, Rey and Kovacic [2006]). Authorizing rewards in our setting would also deter collusion in all industries. However, in practice, it is generally politically unfeasible.

[14] This assumption, which follows Harrington [2008], is made for simplicity. The results are qualitatively unchanged if we assume instead that, following a condemnation, firms compete only for a finite length of time.

[15] Using event study techniques, Aguzzoni *et al.* [2013] find that EU surprise investigations and infringement decisions have a negative effect on a firm's share price. Moreover, they show that the antitrust fine accounts for less than one third of the share price decrease, suggesting that most of the loss is due to the cessation of the illegal activity. As their sample is composed mostly by cartels, their evidence supports the assumption that firms cease colluding once condemned.

[16] Higher fines for repeat offenders are incorporated in U.S. and European antitrust laws.

[17] The reason is the same as in Harrington [2008]: *'If the other firms are colluding then a firm that cheats and applies for amnesty will necessarily be the first firm to come forward. Offering leniency to more than the first firm does not then enhance the payoff to cheating [. . .]. However,*

Interpretation of the Signal. In practice, the signal may be obtained during a sector inquiry. According to the European Commission:

> *The Commission may decide to start a sector inquiry when a market does not seem to be working as well as it should. This might be suggested by evidence such as limited trade between Member States, lack of new entrants on the market, the rigidity of prices, or other circumstances suggest that competition may be restricted or distorted within the common market.*

Alternatively, the realization of the signal may be interpreted as whether the CA received some initial incriminating evidence from third parties such as internal employees, complaints or local agencies.[19] Accordingly, ψ is likely to vary in practice with the efficiency of existing schemes designed to encourage whistleblowing. In particular, sentencing individuals to imprisonment or offering bounties to informants strongly incentivize third parties' cooperation. The U.S. Amnesty Plus Program, which gives strong financial incentives to firms already under prosecution for denouncing cartels in other markets, is also an efficient instrument for raising the probability of obtaining initial evidence on separate cartels. As for μ, it is likely to vary with technological progress, for instance the use of digital forensics.

II(ii). *Strategies*

We focus on stationary strategies.[20]

Competition Authority. We assume that the CA can commit to an overall probability of investigation against cartels, $\tilde{\sigma}$. It will then find it optimal to open an investigation with probability one after receiving a good

it does boost the payoff to continuing to collude since, when all firms decide to discontinue colluding and apply for amnesty, allowing more than one firm to receive it reduces expected penalties [. . .].'

[18] As in the literature, we assume that if m firms (simultaneously) apply for leniency, they are equally likely to be the first informant and thus face an expected fine equal to $\left(1-\frac{q}{m}\right)F$.

[19] We could furthermore assume that the initial evidence is more or less reliable, which would justify the stationarity of the model. Suppose for example that a good signal means that the initial evidence is reliable with probability μ, in which case an investigation will be successful for sure, and is otherwise unreliable, in which case an investigation cannot succeed. If the CA launches an investigation, either the cartel is condemned and the game thus ends for that industry, or the investigation fails and the CA will then infer that its initial evidence was not reliable.

[20] However, it is important to stress that non-stationary strategies are theoretically very powerful in deterring cartels. See Frezal [2006].

signal, and to open an investigation with probability σ—such that $\psi+(1-\psi)\sigma=\tilde{\sigma}$—after receiving a bad signal.[21,22]

When the signal is public, cartel members will not report information if the CA launches an investigation after having received a bad signal, since they do not face any risk of conviction. In this scenario, the only policy variable is q. When instead the signal is privately observed by the CA, there is scope for bluffing—i.e., choosing $\sigma>0$—as firms may fear that the investigation could be successful. In that case, the policy variables are q and σ. We shall consider both scenarios below.

Firms. We focus on grim trigger strategies in which any deviation from collusion is punished by reverting forever to competition, which is here the minmax and thus constitutes the most severe punishment. We consider two modes of (symmetric) collusive equilibrium, in which firms either 'collude and remain silent,' or 'collude and report in case of investigation.'[23]

In both cases, firms collude in every period until a deviation occurs or the cartel is condemned, and if a firm deviates on the market, then firms revert to competition forever. If the CA opens an investigation, firms remain silent in 'collude and remain silent' and report in 'collude and report in case of investigation.' If the CA finds evidence, or one firm reports, the cartel is condemned and firms compete forever. Otherwise, firms go on colluding.

In order to be sustainable, both collusive strategies must therefore resist unilateral deviations on the market—i.e., the expected value of future collusion must exceed the gains from deviating on the market, Π_D. The strategy 'collude and remain silent' must moreover be incentive-compatible, that is, robust to unilateral reporting deviations: no firm should gain by reporting in case of investigation when the other firms remain silent. If both collusive strategies are sustainable, we assume that firms select the Pareto-dominant collusive strategy—i.e., the most profitable one.

Let us denote by $\gamma\in\{0,1\}$ firms' decision to report information in case of investigation,[24] and by ϕ the conviction rate faced by firms when they

[21] Commitment about $\tilde{\sigma}$ (and thus σ) may be made credible by legislation or disclosure rules, or by the CA's desire to establish a reputation in a repeated game. I discuss some commitment devices in section V.

[22] The model would deliver similar insights if we assume that the CA receives the signal only after the opening of an investigation, when handling the case (but before firms' reporting decisions). Intuitively, the CA would then open an investigation in every period and close the case with probability $1-\sigma$ if the signal is bad.

[23] This is without loss of generality. First, since the antitrust policy is stationary, if collusion is sustainable then the best collusive strategy consists in colluding in every period until a deviation occurs or the cartel is condemned. Second, if the CA offers leniency ($q>0$), the equilibrium reporting strategies are necessarily symmetric since a firm is better off reporting whenever at least one other firm reports.

[24] Assuming that the reporting decision is $\gamma\in\{0,1\}$ is without loss of generality. Colluding and randomizing between reporting or not in case of an investigation (even using a public lottery) is less profitable than either 'collude and remain silent' or 'collude and report in case

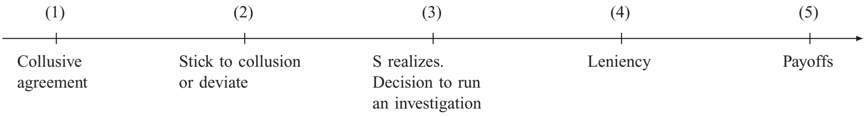

Figure 1
The Stage Game

collude. When firms collude, the CA runs an investigation with probability $\psi + (1 - \psi)\sigma$. Then either firms report, in which case the cartel is condemned with certainty, or no firm reports, and the cartel is condemned only if the investigation is successful—i.e., with probability $\psi\mu$. Therefore, the conviction rate is:

$$\phi(\sigma,\gamma) = (\psi + (1-\psi)\sigma)\gamma + \psi(1-\gamma)\mu.$$

Let us summarize the framework. The antitrust rules ψ, μ, β and F, which are taken as given by the CA, are common knowledge; the realization of the signal is also common knowledge in section III, and instead privately observed by the CA in section IV. At the beginning of the game, the CA announces the policy variables q and σ, and sticks to this policy afterwards. Then, in each period, the timing of the game, summarized in Figure 1, is as follows:

- Stage 1. Each firm chooses whether to enter into a collusive agreement. If at least one firm chooses not to collude, competition takes place and the game moves to the next period. If all firms enter into a collusive agreement, this decision leaves hard evidence of collusion and the game proceeds to stage 2.
- Stage 2. Each firm chooses whether to respect the agreement and collude, or to deviate and compete on the market. The game then proceeds to stage 3.
- Stage 3. The CA receives the signal and decides whether or not to run an investigation. This decision is publicly observed by firms. If the CA runs an investigation, the game proceeds to stage 4.
- Stage 4. Each firm decides simultaneously whether to apply to the leniency program. If at least one firm reports, the cartel is condemned; in that case the fine is reduced to $(1 - q)F$ for the eligible firms. Otherwise, the game proceeds to stage 5.
- Stage 5. If the CA received a good signal, it finds hard evidence and can thus condemn the cartel with probability μ. If the signal is bad, the CA has no chance of gathering evidence of collusion.

of investigation.' Moreover, colluding and randomizing between reporting or not in case of an investigation is not sustainable when 'collude and remain silent' is not sustainable (see also footnote 33).

II(iii). *Welfare*

The CA is benevolent and maximizes total welfare, or equivalently minimizes the social cost of collusion. Society incurs a per-period deadweight loss $L > 0$ when firms collude. We also assume that there is a fixed social cost, $c \geq 0$, from investigating an industry that proves to be competitive.[25] The cost c incorporates what is referred to as the 'chilling' of desirable behavior in Kaplow [2011a, 2011b].[26,27]

Let us denote by V the value of collusion, which depends on whether firms report or not in case of investigation. We will assume that firms choose to collude whenever collusion is sustainable, i.e., whenever $V \geq \Pi^D$. Gains from deviating, Π^D, are distributed across industries according to a strictly increasing and continuously differentiable cumulative distribution function G defined over the support $[\Pi^C, +\infty[$.

In order to generate a stationary proportion of collusive industries in the economy, we assume that in each period, an industry is replaced with probability κ by a new industry with the same Π^D.[28]

Denoting x_t the proportion of collusive industries in period t and y_t the proportion of (previously collusive) industries which were condemned by the CA and are thus competing in period t, we have:

$$x_t + y_t = G(V)$$

With probability κ, (previously collusive) industries condemned by the CA are replaced by new industries with identical gains from deviating, which then will be colluding in the following period; on the other hand, with probability $\phi(\sigma, \gamma)$, collusive industries are condemned by the CA and prevented from colluding in the following period. x_t and y_t thus also solve:

[25] Alternatively, we could consider c as the resources used to conduct an investigation (be it in a collusive or competitive industry). In that case, C (see below) would be equal to $\frac{\kappa G(V)}{\kappa+\phi(\sigma,\gamma)}(L+(\psi+(1-\psi)\sigma)c)+\left(1-\frac{\kappa G(V)}{\kappa+\phi(\sigma,\gamma)}\right)\beta\sigma c$. Qualitatively, the model would generate very similar results.

[26] Kaplow [2011a, p. 1,105] provides illustrations of the chilling of desirable behavior in several environments: '[. . .] antitrust, where misapplication of sanctions may discourage efficient, procompetitive behavior (e.g., promotional product pricing may look like predation); securities regulation, where the prospect of erroneous sanctions may increase the cost of capital [. . .]; medical malpractice, where worries about false positives may discourage cost-effective care [. . .]; and contract breach generally, where concern for misassessment of performance may deter efficient contracting or misdirect behavior governed by contracts.' See also Kaplow [2011b] for a discussion of chilling effects in the context of price-fixing.

[27] The cost c might also incorporate the resources and enforcement costs wasted by investigating an industry in which there is no cartel.

[28] This assumption is inspired by the modeling of the birth and death process for cartels in Harrington and Chang [2009], and allows me to generate, as in their paper, a stationary distribution of cartels in the economy.

$$x_t = (1-\phi(\sigma,\gamma))\,x_{t-1} + \kappa y_{t-1}; \text{ and:}$$

$$y_t = \phi(\sigma,\gamma)\,x_{t-1} + (1-\kappa)\,y_{t-1}.$$

For a given conviction rate $\phi(\sigma, \gamma)$, it follows that the proportion of collusive industries is stationary and equals $x_{t-1} = x_t = \frac{\kappa G(V)}{\kappa+\phi(\sigma,\gamma)}$. We are now able to express the (per-period) social cost of collusion in the economy, C, which equals:

$$C = \frac{\kappa G(V)}{\kappa+\phi(\sigma,\gamma)} L + \left(1 - \frac{\kappa G(V)}{\kappa+\phi(\sigma,\gamma)}\right)\beta\sigma c$$

In each period, firms collude in $\frac{\kappa G(V)}{\kappa+\phi(\sigma,\gamma)}$ industries, in which case society incurs the (per-period) loss L. In the other industries, firms compete and a bad signal is sent with probability β. In that case, the CA opens an investigation with probability σ, which imposes a fixed social cost, c.

An antitrust policy has a positive effect on cartel deterrence if it reduces the proportion of industries in which collusion is sustainable, $G(V)$. It has a positive effect on cartel desistance if it raises the conviction rate $\phi(\sigma, \gamma)$. By contrast, antitrust intervention is costly when it mistakenly targets a competitive industry.

III. OPTIMAL POLICY UNDER PUBLIC SIGNALS

As already mentioned, when cartel members observe the CA's signal, they will never report if the CA receives a bad signal and thus opening an investigation can only be socially costly in that case. The only policy variable is thus the leniency rate q. Cartel members may want to apply for leniency when the CA opens an investigation after receiving a good signal, since they then face a risk of being condemned.

'Collude and report in case of investigation' (R). If firms report when the CA runs an investigation after receiving a good signal ($\gamma = 1$), the cartel is then condemned with probability 1 instead of μ. *Ex ante*, the (expected) conviction rate is therefore $\phi(0, 1) = \psi$, the probability that the CA receives a good signal. If the leniency rate is q, each firm faces an expected reduction rate $\frac{q}{N}$, since leniency is granted only to the first informant. Therefore, denoting $\delta = \hat{\delta}(1-\kappa)$, the value of collusion, V_B^R (where B stands for Benchmark), solves $V_B^R = \Pi^C - \psi\left(1-\frac{q}{N}\right)F + (1-\psi)\delta V_B^R$, that is:

$$V_B^R(q) = \frac{\Pi^C - \psi\left(1 - \frac{q}{N}\right)F}{1 - \delta + \delta\psi}$$

'Collude and be silent' (S). If all firms choose to remain silent when the CA carries out an investigation ($\gamma = 0$), the cartel is dismantled only if, having received a good signal, the CA succeeds in uncovering hard evidence during the investigation. The conviction rate is then only $\phi(0, 0) = \psi\mu$ but, when convicted, firms pay the full fine F. Therefore, the value of collusion, V^S, is:

$$V^S = \frac{\Pi^C - \psi\mu F}{1 - \delta + \delta\psi\mu}$$

As already mentioned, the collusive path S faces an incentive-compatibility (IC_B) constraint: no firm should gain by instead reporting information when an investigation is ongoing, in which case it pays only a reduced fine $(1 - q)F$ but foregoes future collusion. If instead firms stick to S, with probability μ the investigation proves successful in which case cartel members pay F and then compete forever, whereas with probability $(1 - \mu)$ the investigation fails and cartel members' discounted continuation payoffs are δV^S. The incentive-compatibility constraint is therefore:[29]

$$(IC_B) \quad -\mu F + (1 - \mu)\delta V^S > -(1 - q)F$$

Firms choose S instead of R only if it is both incentive-compatible and more profitable. Lemma 1 derives the firms' decisions as a function of q.

Lemma 1. There exists a threshold, $\hat{q} \equiv (1 - \mu)\frac{\delta\Pi^C + (1-\delta)F}{(1-\delta+\delta\mu\psi)F}$, such that:

for $q < \hat{q}$,

firms collude and remain silent if $V^S > \Pi^D$, and compete otherwise.

for $q \geq \hat{q}$,

firms collude and report in case of investigation if $V_B^R(q) > \Pi^D$, and compete otherwise.

The threshold $\hat{q}$ is driven by the incentive compatibility constraint (IC_B), which ensures that no firm reports when the other firms remain silent. A reporting firm, in this case, would obtain leniency with probability one. When (IC_B) is not satisfied, firms choose the collusive path R, in which all firms report in case of investigation and receive leniency with probability $1/N$. When $q = \hat{q}$, firms do not switch from S to R because the collusive path R is more profitable, but because S is no longer

[29] For the sake of exposition, we assume that a firm decides to report whenever it is indifferent between reporting and remaining silent.

incentive-compatible. As a consequence, the value of collusion falls. This ensures that, when $q = \hat{q}$, both cartel desistance (a higher conviction rate due to reporting) and cartel deterrence (a lower value of collusion) are improved.

Note that this insight rests on the presence of the first informant rule and echoes the 'race to the courthouse effect' identified in Harrington [2008]. Under a first informant rule, the decision to report is similar to a prisoner dilemma: for $q = \hat{q}$, firms would be better off choosing S rather than R; however, S is not incentive compatible, and firms run to the courthouse in case of an investigation. As leniency is granted only to the first informant, expected penalties increase and, as a result, deterrence is improved.

The threshold $\hat{q}$ is lower than 1 if μ is higher than a threshold $\underline{\mu} < 1$. Figure 2(a) represents the conviction rate and Figure 2(b) the value of collusion as a function of q for the case $\mu > \underline{\mu}$.

The optimal antitrust policy consists in choosing the leniency rate q which minimizes the social cost of collusion C.

Cartel desistance is enhanced when firms report information, since in that case, conviction is obtained for sure during an investigation instead of being uncertain. The CA has to set the leniency rate above $\hat{q}$ in order to force cartel members to report. When $\mu < \underline{\mu}$, it is however impossible to induce reporting if rewards are ruled out (i.e., if $q \leq 1$). For the rest of the paper, we assume $\mu \geq \underline{\mu}$. Introducing a leniency rate $q \geq \hat{q}$ suffices to optimize desistance, and raises the conviction rate from $\psi\mu$ to ψ.

The value of collusion is minimized when the CA sets q equal to $\hat{q}$ (see Figure 2(b)). Offering more leniency is pro-collusive: this reduces expected sanctions without further increasing the conviction rate.

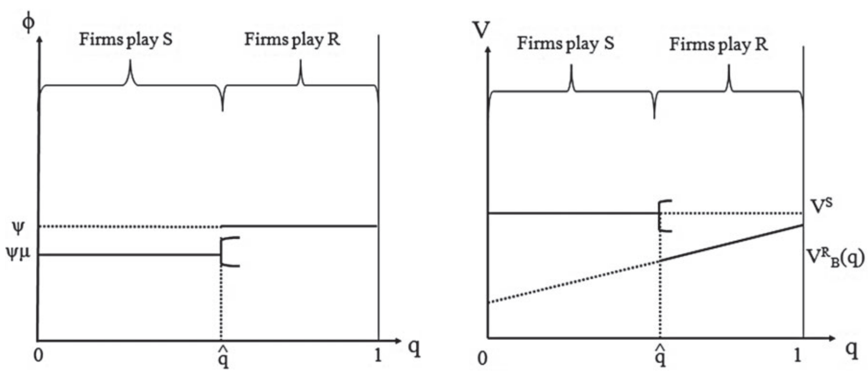

(a) Conviction rate. (b) Value of collusion.

Figure 2
Benchmark

Choosing q equal to $\hat{q}$ optimizes both cartel desistance and cartel deterrence and thus forms the optimal antitrust policy.[30] Proposition 1 summarizes the analysis:

Proposition 1 (Public signals). *Leniency policy*. The optimal leniency rate, granted only to the first informant, is to set q equal to $\hat{q}$.

Welfare. Introducing leniency is desirable both from a deterrence and a desistance perspective. The proportion of collusive industries in the economy drops from $\frac{\kappa G(V^S)}{\kappa+\psi\mu}$ to $\frac{\kappa G(V_B^R(\hat{q}))}{\kappa+\psi}$.

In what follows, we assume that collusion is not deterred in all industries in the absence of leniency. This implies that:[31]

$$\text{(H1)} \quad F < \frac{\delta(1-\psi\mu)\Pi^C}{\psi\mu}$$

IV. OPTIMAL POLICY UNDER PRIVATE SIGNALS

Suppose now that firms do not observe the signal. When the signal is bad, the CA has no chance of condemning the cartel through its own investigations. Still, the CA may want to run an investigation in the hope that cartel members will themselves denounce the cartel.

In this section, the policy variables are q and σ. If cartel members report in case of an investigation, 'bluffing'—that is, choosing $\sigma > 0$—enhances both the conviction rate and cartel deterrence. However, increasing σ has a shadow cost: it dilutes the risk of conviction if cartel members remain silent. By Bayes rule, the probability of prosecutorial success in case of an investigation becomes $\frac{\psi\mu}{\psi+(1-\psi)\sigma}$, which decreases in σ. As shown below, because of this *dilution* effect, the optimal investigation policy σ^* is strictly lower than 1, so as to keep inducing cartel members to report.[32]

Finally, when bluffing, the CA is not certain that there is a cartel. As we will see, a relatively high c—the social loss from investigating a competitive industry—then pushes the CA to reduce the probability of bluffing, σ, possibly to the point $\sigma = 0$.

[30] Formally, minimizing the social costs of collusion in the benchmark boils down to minimizing $\frac{\kappa G(V)L}{\kappa+\phi(0,\gamma)}$. Choosing $q = \hat{q}$ both triggers reporting ($\gamma = 1$) and minimizes the value of collusion V, and thus forms the optimal antitrust policy.

[31] Π^D is drawn from $[\Pi^C, +\infty]$ and thus, collusion is not deterred in all industries in the absence of leniency if $V^S > \Pi^C$. Replacing V^S by its value gives (H1).

[32] Commitment has value because of this dilution effect. In the absence of commitment, the CA may prefer bluffing with probability one in order to raise the conviction rate. However, this would undo cartel members' incentives to report.

As in the preceding section, let us compute the value of R and S.

'Collude and report in case of investigation' (R). When firms do not observe the CA's signal and report in case of investigation ($\gamma = 1$), they are condemned even when the CA received a bad signal. The conviction rate is therefore $\phi(\sigma, 1) = \psi + (1 - \psi)\sigma$, and the value of collusion, $V^R(\sigma, q)$, is then:

$$V^R(\sigma, q) = \frac{\Pi^C - \phi(\sigma,1)\left(1-\frac{q}{N}\right)F}{1-\delta+\delta\phi(\sigma,1)}$$

'Collude and be silent' (S). The value of collusion remains V^S because the CA can find hard evidence only after receiving a good signal. S is again subject to an incentive-compatibility (*IC*) constraint: no firm should be willing to betray the cartel by reporting during an investigation, which is the case if:

$$\text{(IC)} \quad \left(1-\frac{\psi\mu}{\psi+(1-\psi)\sigma}\right)\delta V^S - \frac{\psi\mu}{\psi+(1-\psi)\sigma}F > -(1-q)F$$

Again, firms choose S instead of R only if it is both incentive-compatible and more profitable. Lemma 2 derives the firms' decisions as a function of q and σ.[33]

[33] As already mentioned, focusing on the discrete strategy space $\gamma \in \{0, 1\}$ is not restrictive. Consider that firms collude and report (using public lotteries) with probability $\gamma < 1$ in case of investigation (the discussion below also applies when firms randomize using private lotteries), and denote V_γ the value of this collusive path. It is straightforward that $V_\gamma(\sigma, q) < \max(V^S, V^R(\sigma, q))$. Consider first $V_\gamma(\sigma, q) < V^R(\sigma, q)$; in that case, it is more profitable for firms to report with probability 1 in case of an investigation than with probability $\gamma < 1$. Consider now $V_\gamma(\sigma, q) < V^S$. Either S is incentive-compatible, in which case it is more profitable for firms to remain silent with probability 1; or S is not incentive-compatible, but then, remaining silent with probability $1 - \gamma < 1$ is not incentive-compatible either. The incentive-compatibility constraint (IC_γ) would be:

$$-\gamma\left(1-\frac{q}{N}\right)F+(1-\gamma)\left(\left(1-\frac{\psi\mu}{\psi+(1-\psi)\sigma}\right)\delta V_\gamma(\sigma,q)-\frac{\psi\mu}{\psi+(1-\psi)\sigma}F\right)$$
$$> -\gamma\left(1-\frac{q}{N}\right)F-(1-\gamma)(1-q)F$$

which is equivalent to:

$$\left(1-\frac{\psi\mu}{\psi+(1-\psi)\sigma}\right)\delta V_\gamma(\sigma,q)-\frac{\psi\mu}{\psi+(1-\psi)\sigma}F > -(1-q)F$$

which is not satisfied when (IC) is not satisfied (because $V_\gamma(\sigma, q) < V^S$).

Lemma 2. For $q < \hat{q}$, $\forall \sigma$,
firms collude and remain silent if $V^S > \Pi^D$, and compete otherwise.
For $q \geq \hat{q}$, there exists a threshold $\tilde{\sigma}(q) \in [0,1)$ such that:
for $\sigma \leq \tilde{\sigma}(q)$,
firms collude and report in case of investigation if $V^R(\sigma, q) > \Pi^D$, and compete otherwise;
for $\sigma > \tilde{\sigma}(q)$,
firms collude and remain silent if $V^S > \Pi^D$, and compete otherwise.

The result in Lemma 2 depends on how the two policy variables, the leniency rate q and the probability of opening an investigation when the case is weak, σ, affect the incentive compatibility constraint (IC). An increase in q makes reporting more attractive when the other firms remain silent, weakening the collusive path S. Conversely, an increase in σ dilutes the conviction rate and makes the collusive path S easier to sustain. For these reasons, firms report in case of an investigation only when the CA sets q sufficiently high and σ sufficiently low.

For $q < \hat{q}$, cartel members do not report information in the absence of bluffing (see the benchmark case). As bluffing dilutes the risk of being condemned when there is an investigation, they will not report *a fortiori* when $\sigma > 0$. Suppose now that $q \geq \hat{q}$. In that case, there exists a threshold $\tilde{\sigma}(q) \in [0,1)$ such that for $\sigma \leq \tilde{\sigma}(q)$, S is not incentive-compatible and therefore firms report in case of investigation. Figure 3(a) represents the conviction rate and Figure 3(b) the value of collusion as a function of σ when $q \geq \hat{q}$.

Intuitively, if the CA chooses σ low enough—i.e., below the threshold $\tilde{\sigma}(q)$—it is likely to have received a good signal when it launches an investigation; consequently firms prefer to report. In contrast, when the CA

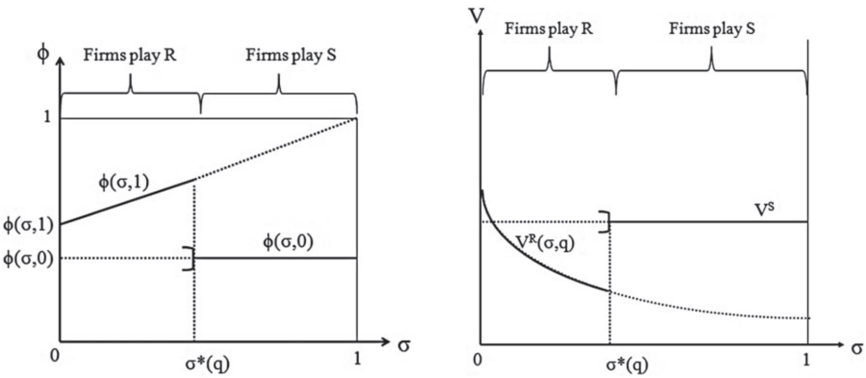

(a) Conviction rate. (b) Value of collusion.

Figure 3
Conviction Rate and Value of Collusion as a Function of the Investigation Policy, σ

chooses σ strictly above $\tilde{\sigma}(q)$, the risk of conviction is too low to induce firms to denounce the cartel. By bluffing, the CA dilutes the risk of conviction, possibly to the point of discouraging leniency applications.

Observe that bluffing is potentially optimal only if it triggers reporting. Choosing $\sigma > \tilde{\sigma}(q)$ is socially costly for two reasons. First, firms remain silent during investigations and the conviction rate drops to $\psi\mu$. Opening investigations too frequently with a bad signal creates a negative externality on the efficiency of investigations with a good signal because this discourages cartel members from reporting information. Secondly, when bluffing, social costs are inflicted if the investigation mistakenly targets a competitive industry.

As $\tilde{\sigma}(.)$ is increasing, we can invert it, defining $\tilde{q}(.) = \tilde{\sigma}^{-1}(.)$. $\tilde{q}(.)$ is the minimal leniency which triggers reporting for a given σ. Denoting $\sigma^M = \tilde{\sigma}(1)$ (σ^M for maximal bluffing), we have $\tilde{q}(0) = \hat{q}$ and $\tilde{q}\left(\sigma^M\right) = 1$. To characterize the optimal cartel policy, we can thus focus on the policy space (q, σ) such that $q \geq \tilde{q}(\sigma)$ (and $\sigma \leq \sigma^M$), in which firms collude and report in case of an investigation.[34]

The following lemma gives the optimal leniency rate in function of σ.

Lemma 3. For $\sigma \leq \sigma^M$, it is optimal to set q equal to $\tilde{q}(\sigma) \in [\hat{q}, 1]$.

Intuitively, fixing the investigation policy σ, the leniency program is too generous if $q > \tilde{q}(\sigma)$. By reducing q to $\tilde{q}(\sigma)$, the CA would increase expected sanctions and still trigger reporting. On the other hand, as already mentioned, when $q < \tilde{q}(\sigma)$, firms remain silent in case of investigation which has a negative effect on welfare.

Note that there is an increasing relationship between the optimal leniency and the investigation policy, σ ($\tilde{q}(\sigma)$ is increasing in σ). The leniency policy not only triggers applications as in the benchmark case but it also gives more room for bluffing. The model generates what I refer to as the 'activism effect' of leniency: by increasing the leniency rate, the CA can open more investigations while still ensuring that firms report, which raises the conviction rate.

Increasing σ and q also reinforces deterrence. In the region where firms report, increasing q would be pro-collusive if the investigation policy were unchanged—formally, $V^R(\sigma, q)$ is *increasing* in q. However, taking into account the 'activism effect' of leniency, increasing q (and σ) reduces the value of collusion—formally, $V^R(\tilde{\sigma}(q), q)$ is strictly *decreasing* in q.

[34] Formally, when $q < \tilde{q}(\sigma)$ or $\sigma > \sigma^M$, firms remain silent in case of investigation and the social costs of collusion equal $\frac{\kappa G(V^S)}{\kappa+\psi\mu} L + \left(1 - \frac{\kappa G(V^S)}{\kappa+\psi\mu}\right)\beta\sigma c$, which are for all σ strictly larger than $\frac{\kappa G(V_B^R(\hat{q}))}{\kappa+\psi} L$, the social costs of collusion when the CA sticks to the benchmark optimal policy $q = \hat{q}$ and $\sigma = 0$.

Note finally that $V^R(\tilde{\sigma}(q), q)$ is *strictly* decreasing in q because leniency is granted only to the first informant. As already mentioned, under a first informant rule, the reporting decision is a prisoner dilemma, which explains why bluffing outweighs the pro-collusive effect of offering more leniency.[35] By bluffing more, the CA increases the likelihood of the reporting decision, which, by virtue of the 'run to the courthouse effect,' *increases* expected sanctions and improves deterrence.

It follows that the proportion of collusive industries in the economy is minimized for $\sigma = \sigma^M$ and $q = 1$.[36] As mentioned below, this result contrasts with previous work in which full leniency is generally pro-collusive (see e.g., Motta and Polo [2003], Harrington [2008]).

The remaining issue is to find the optimal investigation policy, $\sigma \in [0, \sigma^M]$. As in the previous section, the CA considers the effect of its policy on the conviction rate and the value of collusion. However, when bluffing, the CA also takes into account that it is socially costly to open investigations in competitive industries. This may occur because the CA is not certain that there is a cartel when it receives a bad signal. The following proposition characterizes the optimal cartel policy in function of c, the social loss from investigating a competitive industry.

Proposition 2. The optimal investigation policy σ^* and the optimal leniency rate $q^*(=\tilde{q}(\sigma^*))$ are weakly decreasing in c. Specifically, there exists two thresholds $\underline{c}$ and $\overline{c}$ such that:

for $c \leq \underline{c}$, the optimal cartel policy involves maximal bluffing—i.e., $\sigma^* = \sigma^M$—and full leniency to the first informant—i.e. $q^* = 1$;

for $\underline{c} < c < \overline{c}$, the optimal cartel policy involves partial bluffing—i.e., $0 < \sigma^* < \sigma^M$—and partial leniency to the first informant—i.e. $\hat{q} < q^* < 1$;

for $\overline{c} \leq c$, the optimal cartel policy involves as in the benchmark, no bluffing—i.e., $\sigma^* = 0$—and partial leniency to the first informant— $q^* = \hat{q}$.

Intuitively, the optimal investigation policy, σ^*, solves a trade-off between cartel deterrence—i.e., reducing the proportion of collusive industries in the economy—and the chilling costs of investigating competitive industries.

When the social loss from investigating a competitive industry, c, is low—i.e., below $\underline{c}$ —, it is optimal for the CA to maximize the conviction rate, which is accomplished by choosing the highest probability of bluff, σ^M. To trigger reporting in that case, it is necessary to offer full leniency. In this region, the cartel policy then minimizes the proportion of collusive industries in the economy.

[35] If leniency were instead granted to all informants, the value of the collusive path R would be constant in q and equal to V^S. In that case, bluffing raises the conviction rate, but has no effect on the value of collusion.

[36] Rewards would raise further the probability of bluff. In particular, if the CA offers the reward $\tilde{q} > 1$ such that $\sigma^*(\tilde{q}) = 1$, it can bluff with probability one while still ensuring that cartel members report.

Observe that welfare is strictly improved compared to the benchmark only when c is not too large. When $c \geq \overline{c}$, the CA prefers to avoid bluffing since otherwise, the chilling costs of investigating competitive industries would outweigh the benefit from improving deterrence. In that case, offering $\hat{q}$ (as in the benchmark) suffices to trigger reporting and is thus the optimal leniency rate.

Finally, for intermediate values of c—between $\underline{c}$ and $\overline{c}$ —, it is optimal for the CA to bluff only partially ($\sigma^* \in (0, \sigma^M)$). The corresponding optimal leniency, $\tilde{q}(\sigma^*)$ is the lowest one which triggers reporting, and lies in that case between the benchmark level $\hat{q}$ and full amnesty.

Overall, note that a lower c pushes the CA to implement a more aggressive investigation policy, which in turn must be associated with more leniency in order to preserve firms' incentives to report.

The following proposition provides comparative statics.

Proposition 3 (Comparative statics). i) The probability of bluff, σ^*, is weakly increasing in the social loss from collusion, L, and weakly decreasing in the social loss from investigating a competitive industry, c, and the probability that a competitive industry sends a bad signal, β. Comparative statics with respect to the private gain from collusion, Π^C, the fine, F, the probability of having initial evidence, ψ, and the efficiency of investigations, μ, depend on the other parameter values and on the form of $G(.)$, the distribution across industries of the gains from deviating.

ii) Welfare is decreasing in the private gain from collusion, Π^C, the social loss from collusion, L, the social loss from investigating a competitive industry, c, and the probability that a competitive industry sends a bad signal, β, and is increasing in the fine, F, the probability of having initial evidence, ψ, and the efficiency of investigations, μ.

As already mentioned, a higher social loss from investigating a competitive industry pushes the CA to reduce σ^*. Similarly, a higher β increases the risk of facing a competitive industry when bluffing, which also pushes the CA to reduce σ^*. In contrast, an increase in the social loss inflicted by cartels, L, increases the benefit of interrupting collusion, which pushes the CA to raise σ^*.

Note that the comparative statics for σ^* with respect to the private gain from collusion, Π^C, the fine, F, the probability of having initial evidence, ψ, and the efficiency of investigations, μ, depend on the other parameter values and on the form of $G(.)$ only when $\underline{c} < c < \overline{c}$.[37] When $c \leq \underline{c}$, the optimal investigation policy, σ^M, is constrained by firms' incentives to report. In this case, a decrease in the private gain from collusion, Π^C, or a

[37] Formally, when $\underline{c} < c < \overline{c}$, the optimal investigation policy is an interior solution of (P') (see the proof of proposition 2 in the appendix) and the comparative statics are obtained from the first order condition.

harsher stick (i.e., a higher fine F, a higher μ or a higher ψ) reinforces firms' incentives to betray the cartel, and enables the CA to raise the maximal probability of bluff.

V. POLICY IMPLICATIONS

I discuss here the policy implications of the model.

Real World Leniency Policy. In the benchmark case, the optimal leniency rate tends to zero if μ tends to one (see Lemma 1). Offering leniency in this case would be pro-collusive. This result echoes the pro-collusive effect of leniency identified in Motta and Polo [2003]—referred to as the 'cartel amnesty effect' in Harrington [2008] —, and is the standard argument for refusing to grant leniency when the probability of winning the case in the absence of firms' confessions is high. This recommendation is implemented by competition authorities in the U.S. and in Europe. Section B of the Corporate Leniency Policy grants post-investigation leniency to the first informant provided that the DoJ,

> *'at the time the corporation comes in, does not yet have evidence against the company that is likely to result in a sustainable conviction'.*

Similarly, the latest (European) Commission Notice [2006] on immunity from fines specifies that:

> *In order to qualify [for reduction of a fine], an undertaking must provide the Commission with evidence of the alleged infringement which represents significant added value with respect to the evidence already in the Commission's possession.*

The analysis presented in this paper challenges the common view that leniency is pro-collusive when the risk of conviction is large.[38] Unlike previous research including Motta and Polo [2003] and Harrington [2008], this paper considers the investigation policy as a strategic variable. In this case, I have shown that more leniency does not only trigger reporting, but also pushes the CA to open more investigations, which raises both the conviction rate and cartel deterrence. This is what I refer to as the 'activism effect' of leniency. Contrary to what the above motion suggests, it might thus be optimal to maintain $q = 1$ even when μ is large, so as to raise the number of successful investigations.

[38] Harrington [2008] also provides examples where it is optimal to award amnesty even though the chances of a conviction are already quite high; however, in his model, amnesty '*should not be provided when the antitrust authority's case is sufficiently strong*' (see p.218).

Miscoordination Among Cartel Members. By assuming that cartel members can coordinate on the Pareto-dominant collusive path, we have chosen the *equilibrium selection* which is the most 'favorable' to firms, and by the same token, the most 'detrimental' to the CA. When cartel members cannot perfectly coordinate their reporting decisions, the collusive path S becomes more difficult to sustain. Depending on the value of the parameters, the CA would then raise the investigation policy σ^* or reduce the leniency rate q^*.

When the social loss from investigating a competitive industry c is large, I have shown that the optimal policy does not involve bluffing. Intuitively, in that case, the CA will maintain $\sigma^* = 0$ and exploit firms' miscoordination to reduce the leniency rate q^* (which increases expected sanctions and thus improves welfare). By contrast, when c is low enough, the CA will still offer full leniency and exploit firms' miscoordination to raise the investigation policy beyond σ^M.

Note that miscoordination might be explicitly introduced in the model by allowing cartel members to receive private information about the conviction probability. In that case, as shown in Harrington [2013], a firm might apply for leniency even after receiving a private signal indicating that the risk of conviction is relatively low, because it fears that another cartel member will apply for leniency.

Commitment Devices. We assumed that the CA could commit to a given investigation policy. A potential tool to credibly commit is *ex post* transparency: the CA should report cartel cases where it refrained from opening (or postponed) an investigation due to the insufficiency of initial evidence against the cartel. Alternatively, the CA budget can be used to put a cap on the number of investigations opened in each period.

VI. CONCLUDING REMARKS

This paper presents a model in which the CA is privately informed about the expected strength of its case.

I show that the CA can then obtain confessions even when it is unlikely to find evidence on its own. However, the CA should carefully choose its investigation policy as prosecuting a cartel when the success probability is low dilutes the average risk of conviction faced by cartel members and therefore lowers the likelihood of leniency applications. In other words, when the CA decides whether or not to launch an investigation with a low probability of success, there is a tradeoff between a desistance effect—i.e., if one firm reports, the investigation leads to the cartel's condemnation—and a dilution effect—i.e., this reduces firms' incentives to report. The dilution effect arises only if private information on the CA side is assumed since otherwise there is no linkage between investigations.

Harrington [2013] considers the case in which each cartel member receives a private signal about the probability of conviction, and shows that offering less than full leniency is sufficient to maximize the conviction rate. By contrast, in this paper, offering full leniency raises the conviction rate because this allows the CA to open more successful investigations when the cartel case is weak. Both studies are complementary in that they highlight the strategic role of the CA in the presence of private information. I hope that more work will be done to better understand how the CA can manipulate cartel members' beliefs in order to improve the efficiency of anti-cartel policies.

APPENDIX

Proof of Lemma 1. Firms choose S instead of R if and only if $V^S > V_B^R(q)$ and (IC_B) is satisfied. First, let us show that (IC_B) is satisfied if and only if $q < \hat{q}$. Note that (IC_B) rewrites:

$$qF < (1-\mu)\left(F + \delta V^S\right)$$

Plugging the value of V^S (see Section III) in the previous inequality and rearranging, we obtain that (IC_B) is equivalent to:

$$q < (1-\mu)\frac{\delta\Pi^C + (1-\delta)F}{(1-\delta+\delta\mu\psi)F} = \hat{q}$$

Finally, as $V_B^R(q)$ is increasing in q, it suffices to show that $V^S > V_B^R(\hat{q})$ to prove that $V^S > V_B^R(q)$ for $q < \hat{q}$.

$$V^S > V_B^R(\hat{q}) \Leftrightarrow (1-\delta+\delta\psi)\left(\Pi^C - \psi\mu F\right) > (1-\delta+\delta\psi\mu)\left(\Pi^C - \psi\left(1-\frac{\hat{q}}{N}\right)F\right)$$

Replacing $\hat{q}$ by its value, we obtain that:

$$V^S > V_B^R(\hat{q}) \Leftrightarrow N > 1$$

$N > 1$ holds by assumption. □

Proof of Lemma 2. Firms choose S instead of R if and only if $V^S > V^R(\sigma, q)$ and (IC) is satisfied. Let us define $\tilde{\sigma}(q)$ the threshold such that:

$$\text{(A1)} \qquad \left(1-\frac{\psi\mu}{\psi+(1-\psi)\tilde{\sigma}(q)}\right)\delta V^S - \frac{\psi\mu}{\psi+(1-\psi)\tilde{\sigma}(q)}F = -(1-q)F$$

Replacing V^S by its value, we obtain after some computations:

$$\tilde{\sigma}(q) = \frac{\psi}{1-\psi}\left(\frac{\mu\left(\delta\Pi^C + (1-\delta)F\right)}{\delta\Pi^C + (1-\delta)F - (1-\delta+\delta\psi\mu)qF} - 1\right)$$

(IC) is satisfied for $\sigma > \tilde{\sigma}(q)$. Observe that $\tilde{\sigma}(q)$ is strictly increasing in q and $\tilde{\sigma}(\hat{q}) = 0$. It follows that for $q < \hat{q}$, (IC) is satisfied for all $\sigma \geq 0$. Moreover, we showed in the proof of Lemma 1 that $V^S > V_B^R(q)$ for $q < \hat{q}$. As $V^R(0, q) = V_B^R(q)$ and $V^R(\sigma, q)$ is decreasing in $\sigma \left(\frac{\partial V^R}{\partial \sigma} = \frac{(1-\psi)\left(-\delta\Pi^C - (1-\delta)\left(1-\frac{q}{N}\right)F\right)}{(1-\delta+\delta\phi(\sigma,1))^2} < 0 \right)$, $V^S > V^R(\sigma, q)$ for all σ and therefore firms choose S instead of R when $q < \hat{q}$.

As $\tilde{\sigma}(q)$ is increasing in q, it suffices to show that $\tilde{\sigma}(1) < 1$ to prove that $\tilde{\sigma}(q) \in [0, 1)$ for $q \geq \hat{q}$.

$$\tilde{\sigma}(1) < 1 \Leftrightarrow \mu \frac{\delta\Pi^C + (1-\delta)F}{\delta\Pi^C - \delta\psi\mu F} - 1 < \frac{1}{\psi} - 1 \Leftrightarrow \mu\psi F < \delta\Pi^C(1 - \psi\mu)$$

The last inequality holds under (H1).

For $q \geq \hat{q}$, (IC) is not satisfied for $\sigma \leq \tilde{\sigma}(q)$ and therefore firms choose R instead of S. For $\sigma > \tilde{\sigma}(q)$, (IC) is satisfied. To complete the proof, let us show that in that case, S is more profitable than R. As $V^R(\sigma, q)$ is decreasing in σ, it suffices to show that $V^R(\tilde{\sigma}(q), q) < V^S$. Equation (A1) rewrites:

$$(\psi + (1-\psi)\tilde{\sigma}(q) - \psi\mu)\delta V^S - \psi\mu F = (\psi + (1-\psi)\tilde{\sigma}(q))\left(-\left(1 - \frac{q}{N}\right)F + \frac{N-1}{N}qF\right)$$

Adding Π^C on both sides, this is equivalent to:

$$\begin{aligned}
&(\psi + (1-\psi)\tilde{\sigma}(q) - 1)\delta V^S + \Pi^C - \psi\mu F + (1-\psi\mu)\delta V^S \\
&\quad = \Pi^C + (\psi + (1-\psi)\tilde{\sigma}(q))\left(-\left(1 - \frac{q}{N}\right)F + \frac{N-1}{N}qF\right)
\end{aligned}$$

Next, note that $V^S = \Pi^C - \psi\mu F + (1 - \psi\mu)\delta V^S$. The previous equation rewrites:

$$\begin{aligned}
(1 - \delta + \delta(\psi + (1-\psi)\tilde{\sigma}(q)))V^S &= \Pi^C - (\psi + (1-\psi)\tilde{\sigma}(q))\left(1 - \frac{q}{N}\right)F \\
&\quad + (\psi + (1-\psi)\tilde{\sigma}(q))\frac{N-1}{N}qF \\
\Leftrightarrow V^S &= \underbrace{\frac{\Pi^C - (\psi + (1-\psi)\tilde{\sigma}(q))\left(1 - \frac{q}{N}\right)F}{1 - \delta + \delta(\psi + (1-\psi)\tilde{\sigma}(q))}}_{V^R(\tilde{\sigma}(q), q)} \\
&\quad + \underbrace{\frac{(\psi + (1-\psi)\tilde{\sigma}(q))\frac{N-1}{N}qF}{1 - \delta + \delta(\psi + (1-\psi)\tilde{\sigma}(q))}}_{>0} \\
\Leftrightarrow V^S &> V^R(\tilde{\sigma}(q), q)
\end{aligned}$$

□

Proof of Lemma 3. Denoting $w(\sigma,q)=\frac{\kappa G\left(V^R(\sigma,q)\right)}{\kappa+\phi(\sigma,1)}$, the CA's objective is:

$$(P)\quad \underset{\sigma,q}{\text{Min}}\, f(\sigma,q)=w(\sigma,q)L+(1-w(\sigma,q))\sigma\beta c$$

$$\text{subject to}\quad 0\le\sigma\le\sigma^M,$$

$$\tilde{q}(\sigma)\le q\le 1,\text{ with }\tilde{q}(\sigma)\in[\hat{q},1].$$

(P) boils down to minimizing a continuous function on a compact set, and thus admits a solution. Let us show that any solution (q^*, σ^*) of (P) is such that $q^*=\tilde{q}(\sigma^*)$.

Suppose by contradiction that $q^*>\tilde{q}(\sigma^*)$. Observe first that:

$$\forall q,\frac{\partial f(\sigma^*,q)}{\partial\sigma}=\frac{\partial w(\sigma^*,q)}{\partial\sigma}(L-\sigma^*\beta c)+(1-w(\sigma^*,q))\beta c$$

Noting that $\frac{\partial w(\sigma^*,q)}{\partial\sigma}<0$, it follows that $\frac{\partial f(\sigma^*,q)}{\partial\sigma}>0$ if $\sigma^*\ge\frac{L}{c\beta}$. In that case, the objective function is not minimized as it can be reduced by marginally lowering σ and the constraint $\tilde{q}(\sigma^*)\le q^*$ will still be satisfied by continuity: $q^*>\tilde{q}(\sigma^*-\varepsilon)$.

Consider now the case $\sigma^*<\frac{L}{c\beta}$. Observe that:

$$\frac{\partial f(\sigma^*,q)}{\partial q}=\frac{\partial w(\sigma^*,q)}{\partial q}(L-\sigma^*\beta c)$$

As $\frac{\partial w(\sigma^*,q)}{\partial q}>0$, it follows that $\frac{\partial f(\sigma^*,q)}{\partial q}>0$ for $\sigma^*<\frac{L}{c\beta}$. The objective function is in that case strictly increasing in q, which contradicts $q^*>\tilde{q}(\sigma^*)$. □

Proof of Proposition 2. We have shown in the proof of Lemma 3 that $q^*=\tilde{q}(\sigma)$. Denoting $w(\sigma)=w(\sigma,\tilde{q}(\sigma))$, optimizing ($P$) thus boils down to optimizing the following program:

$$(P')\quad \underset{\sigma,q}{\text{Min}}\, g(\sigma)=w(\sigma)+(1-w(\sigma))\sigma\frac{\beta c}{L}$$

$$\text{subject to}\quad 0\le\sigma\le\sigma^M.$$

Let us first show that $w(\sigma)=\frac{\kappa G\left(V^R(\sigma,\tilde{q}(\sigma))\right)}{\kappa+\phi(\sigma,1)}$ is strictly decreasing in σ. From the proof of Lemma 2, we have:

$$V^R(\sigma,\tilde{q}(\sigma))=V^S-\frac{(\psi+(1-\psi)\sigma)\frac{N-1}{N}\tilde{q}(\sigma)F}{1-\delta+\delta(\psi+(1-\psi)\sigma)},$$

which is strictly decreasing in σ.

Moreover, as $\phi(\sigma, 1)$ is increasing in σ, it follows that $w(\sigma)$ is strictly decreasing in σ.

Note that $\sigma^* = 0$ optimizes (P') if and only if:

$$\forall \sigma \in \left[0, \sigma^M\right], w(0) \le w(\sigma) + (1 - w(\sigma))\sigma \frac{\beta c}{L},$$

which is equivalent to:

$$\frac{\beta c}{L} \ge \underset{\sigma \in \left[0,\sigma^M\right]}{\text{Max}} \frac{w(0) - w(\sigma)}{\sigma(1 - w(\sigma))} \left(\text{observe that } \lim_{\sigma \to 0} \frac{w(0) - w(\sigma)}{\sigma(1 - w(\sigma))} = \frac{-w'(0)}{1 - w(0)} \in \mathbf{R}^+ \right)$$

Defining $M = \underset{\sigma \in \left[0,\sigma^M\right]}{\text{Max}} \frac{w(0)-w(\sigma)}{\sigma(1-w(\sigma))}$, it follows that $\sigma^* = 0$ is optimal when $c \ge \overline{c} = \frac{L}{\beta} M$.

Note that $M \ge \frac{w(0)-w\left(\sigma^M\right)}{\sigma^M\left(1-w\left(\sigma^M\right)\right)}$.

On the other hand, $\sigma^* = \sigma^M$ optimizes (P') if and only if:

$$\forall \sigma \in \left[0, \sigma^M\right], w\left(\sigma^M\right) + \left(1 - w\left(\sigma^M\right)\right)\sigma^M \frac{\beta c}{L} \le w(\sigma) + (1 - w(\sigma))\sigma \frac{\beta c}{L},$$

which is equivalent to:

$$\frac{\beta c}{L} \le \underset{\sigma \in \left[0,\sigma^M\right]}{\text{Min}} \frac{w(\sigma) - w\left(\sigma^M\right)}{\left(1 - w\left(\sigma^M\right)\right)\sigma^M - (1 - w(\sigma))\sigma} \in \mathbf{R}^+$$

Defining $m = \underset{\sigma \in \left[0,\sigma^M\right]}{\text{Min}} \frac{w(\sigma)-w\left(\sigma^M\right)}{\left(1-w\left(\sigma^M\right)\right)\sigma^M-(1-w(\sigma))\sigma}$, it follows that $\sigma^* = \sigma^M$ is optimal when $c \le \underline{c} = \frac{L}{\beta} m$.

Note that $m \le \frac{w(0)-w\left(\sigma^M\right)}{\left(1-w\left(\sigma^M\right)\right)\sigma^M}$ and thus $m \le M$.

Finally, for $\underline{c} < c < \overline{c}$ (when $\underline{c} \ne \overline{c}$), σ^* is interior and satisfies the first order condition.

The optimal leniency rate follows from Lemma 3, that is:

for $c \le \underline{c}$, the optimal leniency rate is $q^* = \tilde{q}\left(\sigma^M\right) = 1$;

for $c \ge \overline{c}$, the optimal leniency rate is $q^* = \tilde{q}(0) = \hat{q}$;

for $\underline{c} < c < \overline{c}$, the optimal leniency rate is $q^* = \tilde{q}(\sigma^*)$ with $\sigma^* \in (0, \sigma^M)$ and thus $q^* \in (\hat{q}, 1)$. □

Proof of Proposition 3. Defining $\tilde{c} = \frac{\beta c}{L}$, let us show that $\sigma^*(\tilde{c})$ is decreasing in $\tilde{c}$. First note that $\frac{\partial^2 g}{\partial \sigma \partial \tilde{c}} > 0$. $\forall \sigma' > \sigma^*(\tilde{c}), \forall \tilde{c}' > \tilde{c}$, it follows that:

$$g(\sigma^*(\tilde{c}), \tilde{c}') - g(\sigma^*(\tilde{c}), \tilde{c}) < g(\sigma', \tilde{c}') - g(\sigma', \tilde{c})$$

As $\sigma^*(\tilde{c})$ optimizes (P'), $g(\sigma^*(\tilde{c}), \tilde{c}) \le g(\sigma', \tilde{c})$. Hence, $\forall \sigma' > \sigma^*(\tilde{c})$, $g(\sigma^*(\tilde{c}), \tilde{c}') < g(\sigma', \tilde{c}')$ which implies that $\sigma^*(\tilde{c}') \le \sigma^*(\tilde{c})$.

Moreover, if $\sigma^*(\tilde{c})$ is interior, it must satisfy the first order condition $\frac{\partial g}{\partial \sigma}(\sigma^*(\tilde{c}), \tilde{c}) = 0$. As $\frac{\partial^2 g}{\partial \sigma \partial \tilde{c}} > 0$, $\frac{\partial g}{\partial \sigma}(\sigma^*(\tilde{c}), \tilde{c}') > 0$ and thus $\sigma^*(\tilde{c}') \neq \sigma^*(\tilde{c})$. We thus have $\forall \tilde{c}' > \tilde{c}$, $\sigma^*(\tilde{c}') < \sigma^*(\tilde{c})$. It follows that $\sigma^*(\tilde{c})$ is strictly decreasing in $\tilde{c}$ when $\sigma^*(\tilde{c})$ is an interior solution of (P').

Note finally that the comparative statics for welfare are derived from the (general) envelope theorem.

REFERENCES

Aguzzoni, L.; Langus, G. and Motta, M., 2013, 'The Effect of EU Antitrust Investigations and Fines on a Firm's Valuation,' *Journal of Industrial Economics*, 61(2), pp. 290–338,

Aoyagi, M., 2002, 'Collusion in Dynamic Bertrand Oligopoly with Correlated Private Signals and Communication,' *Journal of Economic Theory*, 102, pp. 229–248.

Athey, S. and Bagwell, K., 2001, 'Optimal Collusion with Private Information,' *RAND Journal of Economics*, 32, pp. 428–465.

Aubert, C.; Rey, P. and Kovacic, W., 2006, 'The Impact of Leniency and Whistleblowing Programs on Cartels,' *International Journal of Industrial Organization*, 24, pp. 1241–1266.

Becker, G., 1968, 'Crime and Punishment: An Economic Approach,' *Journal of Political Economy*, 76, pp. 169–217.

Compte, O., 1998, 'Communication in Repeated Games with Imperfect Private Monitoring,' *Econometrica*, 66, pp. 597–626.

European Union, 2006, *Commission Notice on Immunity from Fines and Reduction of Fines in Cartel Cases*, Official Journal C 298 p. 17 (EU Commission, Brussels, Belgium).

Frezal, S., 2006, 'On Optimal Cartel Deterrence Policies,' *International Journal of Industrial Organization*, 24, pp. 1231–1240.

Genesove, D. and Mullin, W., 2001, 'Rules, Communication and Collusion: Narrative Evidence from the Sugar Institute Case,' *American Economic Review*, 91, pp. 379–398.

Harrington, J. E., Jr., 2008, 'Optimal Corporate Leniency Programs,' *Journal of Industrial Economics*, 56, pp. 215–246.

Harrington, J. E., Jr. and Chang, M. H., 2009, 'Modelling the Birth and Death of Cartels with an Application to Evaluating Antitrust Policy,' *Journal of the European Economic Association*, 7, pp. 1400–1435.

Harrington, J. E., Jr., 2013, 'Corporate Leniency with Private Information: The Push of Prosecution and the Pull of Pre-emption,' *Journal of Industrial Economics*, 61, pp. 1–27.

Kaplow, L., 2011a, 'On the Optimal Burden of Proof,' *Journal of Political Economy*, 119, pp. 1104–1140.

Kaplow, L., 2011b, 'An Economic Approach to Price Fixing,' *Antitrust Law Journal*, 77, pp. 343–449.

Kaplow, L. and Shavell, S., 1994, 'Optimal Law Enforcement with Self-Reporting of Behavior,' *Journal of Political Economy*, 3, pp. 585–606.

Malik, A., 1993, 'Self-Reporting and the Design of Policies for Regulating Stochastic Pollution,' *Journal of Environmental Economics and Management*, 24, pp. 241–257.

Motta, M. and Polo, M., 2003, 'Leniency Programs and Cartel Prosecution,' *International Journal of Industrial Organization*, 21, pp. 347–379.

Reinganum, J., 1988, 'Plea Bargaining and Prosecutorial Discretion,' *American Economic Review*, 78, pp. 713–728.

Rey, P., 2003, 'Towards a Theory of Competition Policy,' in Dewatripont, M.; Hansen, L. P. and Turnovsky, S. J. (eds), *Advances in Economics and Econometrics: Theory and Applications*, (Eight World Congress, Cambridge University Press, Cambridge, England).

Sauvagnat, J., 2010, 'Prosecution and Leniency Programs: A Fool's Game,' TSE working paper Series 10-188 (Toulouse School of Economics, Toulouse, France).

Spagnolo, G., 2004, 'Divide et Impera: Optimal Leniency Programmes,' Center for Economic Policy Research, discussion paper N.4840 (CEPR, Bastwick Street, London, England).