Inspecting Cartels over Time: with and without Leniency Program*

Takako Fujiwara-Greve[†] Yosuke Yasuda[‡]

August 12, 2024

Abstract

Research on cartel inspection has focused on the dynamic behaviors of firms but not so much on the dynamic behavior of the regulator. This paper allows the antitrust authority to choose the level of cartel monitoring intensity and its varying patterns. Specifically, we compare constant monitoring policies with "stochastic" policies that randomize monitoring intensities over time. Under a simplified Bertrand competition, (i) without leniency, both policies have the same effect on cartel deterrence, and (ii) with leniency, for each constant policy, there are stochastic policies with the same mean probability of cartel detection that can prevent collusion strictly more effectively. Thus, (iii) stochastic policies can use lower amnesty rates (reduction of the fine) without compromising the effectiveness of cartel deterrence. The synergy between randomizing monitoring intensity and leniency arises because a deviating firm can use leniency to increase the deviation value only in high-intensity periods, which makes collusion more difficult.

Key words: dynamic regulation, collusion, leniency program, repeated game.

JEL classification: C73, L13, L41

*We thank Johannes Hörner and Takuo Sugaya for useful comments on an earlier draft. We also thank Laurent Linnemer, Catarina Marvão, Nisvan Erkal, and session participants at the EARIE 2022 Vienna for their helpful comments. We also thank the APIOC 2022 Sydney organizers for giving us the Graeme Woodbridge Prize for the previous version of this paper. The usual disclaimer applies. Yasuda gratefully acknowledges research support from Grant-in-Aid for Scientific Research (KAKENHI), 18KK0342. He would also like to thank the ISEG, the University of Lisbon, for providing an excellent research environment as a visiting researcher, which enabled him to proceed with much academic research, including this project.

[†]Department of Economics, Keio University, 2-15-45 Mita, Minato-ku, Tokyo 108-8345 Japan. E-mail: takakofg@econ.keio.ac.jp.

[‡]Department of Economics, Osaka University, 1-7 Machikaneyama, Toyonaka, Osaka 560-0043 Japan. E-mail: yasuda@econ.osaka-u.ac.jp.

1 Introduction

Research on cartel inspection has focused on the dynamic behaviors of firms but not so much on the dynamic behavior of the regulator (e.g., see surveys by Ivaldi et al. (2003), Marshall and Marx (2012), and Marvão and Spagnolo (2013)). It is often assumed that the detection probability is the same over time.[1] In this paper, we allow the antitrust authority (AA hereafter) to choose not only the monitoring intensity but also whether to keep the intensity constant or to randomize it over time.[2] The monitoring intensity determines the detection probability of firms' collusion[3], and if collusion is detected, both firms must pay a (prespecified amount of) fine. When the AA varies the monitoring intensities, firms can also choose whether to collude depending on the realized monitoring intensity of that period.

In the repeated game framework, it is often claimed (e.g., Rotemberg and Saloner (1986) and Dal Bó (2007)) that fluctuations in the environment make collusion more difficult. Thus, one may expect that, by creating a fluctuation of the probability of cartel detection (or market-monitoring intensities), the AA can deter collusion more effectively. The intuition is that firms give up colluding in high monitoring-intensity periods, which reduces the continuation value and incentives to collude even in low-intensity periods.

It turns out that this intuition may not be correct in a simple Bertrand competition model based on Chen and Rey (2013). Specifically, by comparing a constant monitoring policy with "stochastic" policies that have the same mean probability of cartel detection as the constant policy, we show that both policies have identical effects on cartel deterrence. The reason is that, even though firms learn whether the current period has a high probability of cartel detection, the continuation value of future collusion is based only on the *expected* detection probability. That is, the fluctuation of detecting probabilities does not affect the deviation incentives if firms

[1]Exceptions include Frezal (2006), Harrington (2008a), and Gärtner (2022). Frezal (2006) advocates rotation policies to investigate industries over time but does not consider leniency programs. Gärtner (2022) analyzes how firms utilize leniency in a dynamic setting where detection probability stochastically evolves over time.

[2]Another interpretation of our model is that the AA continuously varies the monitoring intensities, but it can choose whether to commit to keeping the actual intensity secret (so that firms only know the mean probability) or to announce the intensity to the firms before each period.

[3]As Harrington (2008b) points out, there are multiple ways that collusion is detected. In this paper, we focus on tacit collusion, and thus detection means that some firm's collusive action H is discovered. Even if the authority inspects, whether they can discover the collusion is not certain, which is reflected in the model.

choose to collude every period (which we call *full collusion*).[4] A possible drawback of stochastic policies is that firms may engage in cartels only in some states of monitoring intensities (which we call *partial collusion*). However, we also show that, for each constant policy, there is a class of stochastic policies with the same mean cartel-detection probability (i.e., the same effectiveness to prevent full collusion) such that any kind of partial collusion is deterred for the same range of market parameters. Therefore, both constant policies and some of its mean-preserving stochastic policies are equally effective.

In the same model with a leniency program installed, the effect of the two kinds of policies can differ. The leniency program is a system to incentivize members of a cartel to report to the AA to reduce the fine and terminate collusion.[5] The key insight is that a deviating firm can choose when to denounce the cartel to benefit from leniency. Under stochastic policies, a deviating firm is willing to use the leniency program only in periods of a high detection probability to improve the deviation payoff, and the reduction of the expected fine is greater than that under the constant policy with the "average" detection probability. Consequently, under stochastic policies, deviation is more attractive. Thus stochastic policies are more effective to prevent full collusion than the mean-probability, constant policy. We also show that the same class of stochastic policies continues to prevent any kind of partial collusion under a leniency program.

Our findings suggest that leniency programs and fluctuations of the intensities in cartel investigations complement each other. To our knowledge, literature has yet to address this type of synergy.

Harrington (2008a) addresses a complementary policy issue: given the detection probability structure and assuming that firms always engage in partial collusion, he characterizes the optimal amnesty rate. By contrast, we fix the amnesty rate and compare constant or stochastic detection policies that deter both full and partial collusion.

We organize the paper as follows. Section 2 analyzes the base model without leniency

[4] Still, varying the AA's monitoring intensity may be helpful for at least two cases. First, if the AA's monitoring cost is an inverted-S shape function of the intensity, making a very low probability state may reduce the expected cost of monitoring. Second, if the firms are global entities and the inspection requires international coordination, it may be challenging for the countries involved to coordinate on a constant inspection policy.

[5] For a survey of the leniency programs, see Marvão and Spagnolo (2013). Since we have only two firms, we focus on the most straightforward leniency program in which only the first informant gets amnesty. Landeo and Spier (2020) investigate the optimal design of a leniency program of multiple firms, which chooses the number of firms to get the amnesty and the amnesty rate.

and shows the equal effectiveness result. Section 3 compares constant policies with the mean-preserving stochastic policies under leniency programs. Section 4 concludes.

2 Base Model without Leniency

Our model is based on Chen and Rey (2013). Consider a duopoly market in which firms 1 and 2 operate over the discrete time horizon $t = 1, 2, \ldots$. They have the common discount factor $\delta \in (0, 1)$. The one-shot profit of each firm is 0 if both firms compete, B (> 0) if they collude, and $2B$ for a firm that deviates from the collusion, in which case the other firm gets 0. Such a situation can be formulated by a **reduced Bertrand game** whose payoff matrix is described by Table 1. Note that actions H and L correspond to collusive and defective behaviors, respectively. We regard (H, H) as a successful cartel and (L, L) as competition. The asymmetric case (H, L) corresponds to the situation where firms once entered into a cartel agreement, but firm 2 alone deviates to a **slightly lower** price. The other asymmetric case (L, H) means that firm 1 alone deviates from the cartel.[6]

	H	L
H	B, B	0, $2B$
L	$2B$, 0	0, 0

Table 1: Reduced Bertrand Game

As in Chen and Rey (2013), collusion (i.e., any action combination other than (L, L)) leaves some evidence that might be found by the AA. A monitoring policy of the AA in a period is represented by a probability p with which the cartel is detected if firms collude.[7] The evidence of collusion lasts only for one period. Hence, even if a cartel is detected, each firm must pay a fine F (constant across periods) only for that period and can restart collusion in the next period. Unlike Chen and Rey (2013), we assume that the detection probability sequence $\{p_t\}$

[6] In this reduced Bertrand game, (L, H) and (H, L) are also one-shot Nash equilibria. However, our intention is that the situations in which one firm slightly undercuts the other are not the focus of "collusion" because these are not a Nash equilibrium of the ordinary Bertrand game. In the name of "L", we have two meanings in the reduced game.

[7] To be more precise, p is the probability such that the AA investigates this market *and* succeeds in uncovering cartels. Chen and Rey (2013) distinguish these two events and denote the probability of investigating the market by α and the (conditional) probability of uncovering the cartel by p. Therefore, our p is equivalent to their αp.

may depend on the investigation strategy chosen by the AA.[8]

We compare the following two types of dynamic investigation policies. A *constant policy* implements the same detection probability over time, i.e., $p_t = p \in (0, 1)$ for all t. In contrast, a (stationary) *stochastic policy* randomizes over multiple detection probabilities, keeping the distribution over the possible detection probabilities stationary. For example, the AA can switch between an "intensive monitoring" period and a "normal monitoring" period, as sometimes done by the traffic police[9], and the probability that the intensive monitoring period will be realized is the same over time.

There are two interpretations of our stochastic policies. One is that the AA literally randomizes how intensively it monitors a market. Another is that the AA rotates inspections over many markets so that from each market point of view, the detection probability varies over time.

2.1 Effectiveness threshold under constant policies

We first derive the sufficient condition for cartel deterrence for each constant policy with $p_t = p \in (0, 1)$, for all $t = 1, 2, \ldots$. We interpret that p is **the status quo policy**, fixed by the institutional constraints. Suppose that firms engage in a cartel, i.e., play (H, H), for every period, which we call *full collusion.* Taking into account the possible fines, the expected profit for each firm, denoted by V, is formulated as follows.

$$V := B - pF + \delta(B - pF) + \delta^2(B - pF) + \cdots = \frac{B - pF}{1 - \delta}$$

Note that in our model, firms can restart collusion after paying the fine. The full collusion is sustainable in a subgame perfect equilibrium if and only if the following trigger strategy combination is a subgame perfect equilibrium (Abreu (1988)): firms play (H, H) as long as no firm deviates from it and will play (L, L) forever[10] once some firm deviates. The trigger strategy

[8] In related research, Gärtner (2022) analyzes a model where detection probability stochastically evolves over time. However, this fluctuation process is exogeneously given, hence the AA cannot choose its policy.

[9] In some countries (e.g., Japan and Finland), the traffic police announces "intensive control periods" during which the traffic violations are monitored with high intensity. (See for example, `https://poliisi.fi/en/-/intensive-speed-control-campaign-by-the-police-next-weekend`.)

[10] Since (L, L) is a Nash equilibrium of our reduced Bertrand game, playing (L, L) in every period (irrespective of the history) is a subgame perfect equilibrium. This equilibrium would generate 0 profit, which clearly serves

combination is a subgame perfect equilibrium if any one-step deviation is not beneficial. Given the above trigger strategy, the expected total payoff from a deviation to (L, H) in one period is

$$2B - pF + 0.$$

This is because the deviating firm earns $2B$ in that period but may need to pay the fine F with probability p because choosing L when the other firm is playing H is not considered a competitive behavior. Note that the detection probabilities at (H, H) and (L, H) are assumed to be the same. For relaxing this assumption, see Appendix.

Thus, full collusion is sustainable under a constant policy p if and only if the following incentive condition is satisfied.

$$V = \frac{B - pF}{1 - \delta} \geqq 2B - pF \iff \delta \geqq \frac{B}{2B - pF}. \tag{1}$$

When either p or F is 0, (1) reduces to $\delta \geqq \frac{1}{2}$. Throughout the paper, we assume that the following condition holds so that the firms have a strict incentive to sustain collusion if there is no antitrust enforcement.

$$\delta > \frac{1}{2}. \tag{2}$$

Given (2), the condition (1) can be rewritten as follows.

$$B \geqq \underline{B} := \frac{\delta pF}{2\delta - 1} \tag{3}$$

The condition (3) means that only markets that generate sufficient collusive payoff B can sustain the full collusion. Therefore, $\underline{B}$ can be interpreted as the effectiveness of the (constant) antitrust policy p.

2.2 Effectiveness threshold under binary stochastic policies

When the AA chooses a stochastic policy, it sets up a distribution function G over various investigation intensities, corresponding to various detection probabilities (the possible "states"

as the severest punishment for both firms.

from the viewpoint of the firms) and uses G every period. We focus on binary stochastic policies, i.e., $supp(G) = \{p_1, p_2\}$, where $p_1 < p_2$ and $p_k \in (0, \hat{p}]$ for $k = 1, 2$. The probability that the "state" p_k realizes is denoted by $x_k := G(p_k) > 0$. We assume that $p \leqq \hat{p}$: the temporarily highest detection probability for the AA may be higher than the status quo constant probability.

If firms cannot know the realized p_k before choosing their actions, the game is essentially the same as the one under the constant policy with the mean probability $x_1 p_1 + x_2 p_2$. Thus, an underlying assumption of stochastic policy implementation is that the AA announces the realized detecting intensity in each period, before the firms choose the stage-game actions.

There are at least two reasons that the AA may want to use a stochastic policy. One is the implementation cost of the policies. If the cost of monitoring is a function of the detection probability p and follows the standard inverted-S shape, it is possible that mixing different probabilities is cheaper than monitoring with the mean probability for sure. The other is the case that the AA is facing multiple markets to monitor. Then the AA may want to rotate the monitoring activities across markets instead of monitoring all markets every period. The latter benefit is advocated by Frezal (2006).

Under a binary stochastic policy, firms can try to collude for all realizations of p_k (essentially the same phenomenon as full collusion under a constant policy) or only in one state (which we call *partial collusion*). First, consider full collusion. For each $k = 1, 2$, let V_k be the total expected payoff of a firm that colludes in any state, starting in a period when p_k is realized. It is recursively formulated as follows:

$$V_k := B - p_k F + \delta\{x_1 V_1 + x_2 V_2\}, \ \forall k = 1, 2. \tag{4}$$

To explain, a firm earns B by (H, H) but may pay the fine F if the cartel is detected in that period, which occurs with probability p_k. In the next period, the firm may start in the state p_1 (with probability x_1) or p_2 (with probability x_2). In either case, the firms choose (H, H) by the full collusion agreement, so that the continuation value is V_k for $k = 1, 2$.

In order to sustain (H, H) in all states, the following incentive conditions must be simultaneously satisfied.

$$V_1 = B - p_1 F + \delta\{x_1 V_1 + x_2 V_2\} \geqq 2B - p_1 F \tag{5}$$

$$V_2 = B - p_2 F + \delta\{x_1 V_1 + x_2 V_2\} \geqq 2B - p_2 F \tag{6}$$

To compare the performance with a constant policy p, we should focus on the stochastic policies with the mean p:

$$x_1 p_2 + x_2 p_2 = p. \tag{7}$$

This is not only for mathematical comparison, but also by the underlying assumption that the AA can distribute its regulatory resources (e.g., labor forces and the budget) over various states. However, the allocation of resources must be constrained by the original constant investigation activities, i.e., $x_1 p_2 + x_2 p_2 \leqq p$, and the optimal allocation of the tasks results in (7).

Proposition 1 [Equal Effect without Leniency] *Assume there is no leniency program. Then, full collusion is sustained in a subgame perfect equilibrium under some constant policy if and only if full collusion is sustained in a subgame perfect equilibrium under* ***any*** *of its mean-preserving, binary stochastic policies.*

Proof. We prove the following: for a given p and any binary G such that (7) holds[11], the conditions (5) and (6) are both equivalent to (3). Recall also that $x_1 + x_2 = 1$.

Multiplying both sides of (4) by x_k for each $k = 1, 2$, adding up and using the mean-preservation (7), we have

$$\begin{aligned} x_1 V_1 + x_2 V_2 &= B - pF + \delta\{x_1 V_1 + x_2 V_2\} \\ \Longleftrightarrow x_1 V_1 + x_2 V_2 &= \frac{B - pF}{1 - \delta}. \end{aligned}$$

[11]Note that if $p = \hat{p}$, then there is no non-degenerate, mean-preserving stochastic policy. However, the proposition still holds trivially.

Hence (5) becomes

$$V_1 = B - p_1 F + \delta \frac{B - pF}{1-\delta} \geqq 2B - p_1 F$$
$$\iff B - \delta \frac{B - pF}{1-\delta} \geqq 2B \quad \iff \quad B \geqq \frac{\delta pF}{2\delta - 1} \ (= \underline{B}(p)).$$

(6) also becomes

$$V_2 = B - p_2 F + \delta \frac{B - pF}{1-\delta} \geqq 2B - p_2 F$$
$$\iff B - \delta \frac{B - pF}{1-\delta} \geqq 2B \quad \iff \quad B \geqq \frac{\delta pF}{2\delta - 1} (= \underline{B}).$$

Thus, the incentive conditions under a binary stochastic policy (5) and (6) are both equivalent to the incentive condition under the corresponding constant policy (3). ■

Proposition 1 illustrates that, in the absence of a leniency program, a constant policy and **any** mean-preserving binary stochastic policy are identically effective for full collusion deterrence.[12] This is because the continuation payoff of collusion $\delta\{x_1 V_1 + x_2 V_2\}$ depends only on the mean detection probability p and not on the realized p_k.

Under a stochastic policy, firms may engage in partial collusion such that they choose (H, H) in some states but (L, L) in other states. This is a possible drawback of a stochastic policy that allows firms to collude in various ways. Nonetheless, we can show that, for each stationary policy, there is a class of mean-preserving binary stochastic policies under which no partial collusion is sustainable. Then, stochastic policies in that class are as effective as the corresponding constant policy in cartel deterrence.

As preparation, we formulate the long-run payoff structure of partial collusion. For any $k = 1, 2$, suppose that firms choose (H, H) only when p_k is realized. The total expected payoff of a firm starting in state k and $j(\neq k)$ are recursively formulated as follows, where the

[12]This result can be generalized for non-binary stochastic policies.

superscript stands for the partial collusion state and the subscripts stand for the starting state.

$$W_k^k := B - p_k F + \delta\{x_k W_k^k + x_j W_j^k\};$$
$$W_j^k := 0 + \delta\{x_k W_k^k + x_j W_j^k\}.$$

(Note that firms compete when p_j realizes and thus each firm earns 0 in that period.)

Multiplying both sides of the above equations by x_k and x_j respectively and adding them up, we have

$$x_k W_k^k + x_j W_j^k = x_k(B - p_k F) + \delta\{x_k W_k^k + x_j W_j^k\};$$
$$\iff x_k W_k^k + x_j W_j^k = \frac{x_k(B - p_k F)}{1-\delta}.$$

Thus partial collusion only in state k is sustained if and only if

$$W_k^k = B - p_k F + \delta\frac{x_k(B - p_k F)}{1-\delta} \geqq 2B - p_k F. \tag{8}$$

Lemma 1 *Assume that there is no leniency program. If full collusion is deterred under a binary stochastic policy, then partial collusion in which the firms collude only in state* 2 *is deterred.*

Proof. The incentive condition (8) for $k = 2$ is equivalent to

$$B + \delta\frac{x_2(B - p_2 F)}{1-\delta} \geqq 2B, \tag{8'}$$

while the incentive condition of full collusion is

$$B + \delta\frac{(B - pF)}{1-\delta} \geqq 2B.$$

Without loss of generality, assume a non-trivial binary policy such that $x_1, x_2 > 0$. Since $p_1 < p_2$, the mean-preservation (7) implies that $p < p_2$. Hence the LHS of (8') is strictly smaller than the LHS of the full collusion condition. This means that, if full collusion is deterred, the partial collusion in state 2 must also be deterred. ■

Lemma 2 *Assume that there is no leniency program. The partial collusion only in state* 1 *is deterred for any* B *if and only if*

$$x_1 < \frac{1-\delta}{\delta}. \tag{9}$$

Proof. By rearrangements, the incentive condition (8) for $k = 1$ becomes

$$\delta \frac{x_1(B - p_1 F)}{1-\delta} \geqq B \iff (x_1 - \frac{1-\delta}{\delta})B \geqq x_1 p_1 F.$$

Since $x_1 p_1 F \geqq 0$ (p_1 can be 0), if $x_1 < \frac{1-\delta}{\delta}$, then the above inequality is violated for any $B(> 0)$, i.e., partial collusion in state 1 is not sustainable. ■

Proposition 2 [Equal Effect Class for All Possible Collusion, without Leniency] *Assume there is no leniency program. For any constant policy* p*, any of its mean-preserving binary stochastic policy with* $x_1 < \frac{1-\delta}{\delta}$ *prevents any type of collusion for the same range of* B *(i.e.,* $(0, \underline{B}(p))$*).*

Therefore, the AA can use any of the binary stochastic policies satisfying (9) without compromising the effectiveness of cartel deterrence. When the AA uses a stochastic policy such that $x_1 > \frac{1-\delta}{\delta}$, still firms cannot conduct partial collusion in p_1 if B is less than some bound. However, this will not expand the range of B that deters all kinds of collusion. Hence such stochastic policies are not more effective than those satisfying (9). See Figure 1 for intuition.

Let us call a policy *optimal* when it **maximizes** the lower bound to B at which the incentive condition holds. This reduces the set of markets that sustain collusion as much as possible. Without leniency, the optimal constant policy is $p = \overline{p}$. If $\overline{p} = \hat{p}$, there is no better binary stochastic policy than this. When $\overline{p} < \hat{p}$, there are more effective binary stochastic policies: set $p_2 = \hat{p}$ and (x_1, p_1) to satisfy $x_1 p_1 + (1 - x_1)\hat{p} > p$ and $x_1 < \frac{1-\delta}{\delta}$. This is because the average detection probability increases and not because of fluctuations. In other words, without leniency, the only case that a stochastic policy outperforms constant policies is when the mean detection probability is higher than the best constant policy's.

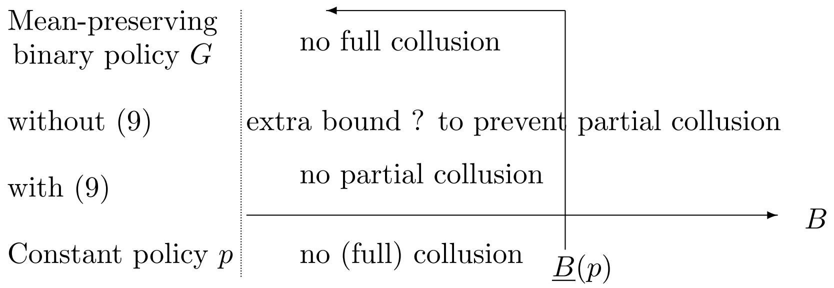

Figure 1: Effectiveness comparison without a leniency program

3 Policy comparison under a leniency program

Let us introduce a leniency program that allows the first (and only first) informant to benefit from a reduced fine qF where $1-q>0$ is the amnesty rate. Following Chen and Rey (2013), We assume that, in each period, firms simultaneously choose a stage game action from $\{H, L\}$ as well as whether to report (action R) the evidence of collusion to the AA or not (action N).[13] See Figure 2 for the outline of the new repeated game.

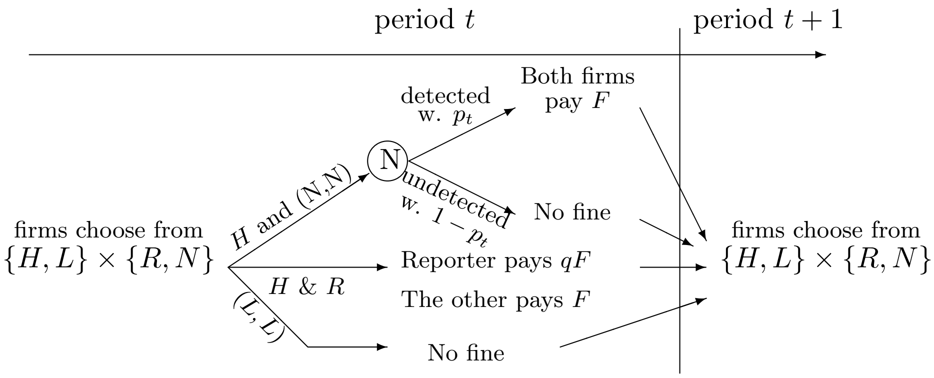

Figure 2: Timeline of the game with a leniency program

3.1 Constant policy with leniency

To sustain (full) collusion under a constant detection probability $p_t = p$ for all $t = 1, 2, \ldots$, it is necessary and sufficient that the following trigger strategy played by both firms is a subgame

[13] If (L, L) is chosen and a firm chooses to report the evidence of collusion, then nothing is reported since there is no evidence of collusion.

perfect equilibrium: Choose H and do not report to the AA as long as no firm deviates, but choose L forever after if a firm deviates from this path.

A deviating firm can choose between not reporting to the AA and using the leniency program, whichever gives a lower (expected) fine. Hence, the following incentive condition must be satisfied to sustain full collusion.

$$V = \frac{B - pF}{1-\delta} \geqq 2B - \min\{pF, qF\} \tag{10}$$

Leniency programs are relevant only if it is used in the optimal deviation, i.e.,

$$p > q. \tag{11}$$

In what follows, we assume (11). Then, condition (10) can be rewritten as follows.

$$\frac{B - pF}{1-\delta} \geqq 2B - qF \iff B \geqq \frac{\{pF - (1-\delta)qF\}}{2\delta - 1} =: \underline{B}^C(p;q), \tag{10'}$$

where the superscript C stands for a "Constant" policy. Note that

$$\underline{B}^C(p;q) = \frac{\delta pF}{2\delta - 1} + \frac{(1-\delta)(p-q)F}{2\delta - 1} = \underline{B}(p) + \frac{(1-\delta)}{2\delta - 1}(p-q)F,$$

which implies that $\underline{B}^C(p;q) > \underline{B}(p)$ for any (p,q) as long as (11) holds. As Chen and Rey (2013) pointed out, it is always desirable to offer some leniency, since that would tighten the incentive condition and make collusion harder to sustain. The extra part $\frac{(1-\delta)}{2\delta-1}(p-q)F$ can be interpreted as the advantage of using the leniency program for a deviating firm. This raises the minimum benefit from collusion to sustain it. It is increasing in the amnesty rate $1-q$.

3.2 Binary stochastic policy under leniency

Consider full collusion deterrence when the AA uses a binary stochastic policy with the support $\{p_1, p_2\}$ such that $p_1 < p_2$. By the mean-preservation (7) and the relevance assumption (11), we must have $p_2 > p > q$. That is, a deviating firm in state p_2 always uses the leniency program to pay qF instead of taking chances to pay the expected fine of p_2F. Thus full collusion is sustainable if and only if the following two conditions are simultaneously satisfied:

$$V_1 = B - p_1F + \delta\frac{B - pF}{1-\delta} \geqq 2B - \min\{p_1F, qF\} \tag{12}$$

$$V_2 = B - p_2F + \delta\frac{B - pF}{1-\delta} \geqq 2B - qF. \tag{13}$$

We now show that, if (13) holds, then the condition (12) also holds. Hence the necessary and sufficient condition for full collusion is (13).

Lemma 3 *Under a binary stochastic policy, full collusion is sustainable if and only if the incentive condition (13) of the more risky state* p_2 *is satisfied.*

Proof. By rearrangements, (13) is equivalent to

$$\delta\frac{B - pF}{1-\delta} \geqq B + (p_2 - q)F, \tag{13'}$$

while (12) is equivalent to

$$\delta\frac{B - pF}{1-\delta} \geqq B + \max\{0, (p_1 - q)F\}. \tag{12'}$$

Since $p_2 > p_1$, the RHS of (13') is strictly greater than the RHS of (12'). ■

From (13) or (13') and $p_2 > p$, we have the following conclusion.

Corollary 1 *The bound to* B *that sustains full collusion under a binary stochastic policy is*

$$B \geqq \frac{\delta pF + (1-\delta)(p_2 - q)F}{2\delta - 1} = \underline{B}(p) + \frac{(1-\delta)}{2\delta - 1}(p_2 - q)F =: \underline{B}^F(p_2; p, q),$$

where the superscript F *stands for full collusion deterrence. For any* (p, q) *such that* $q < p$ *and any mean-preserving binary stochastic policy* G*,*

$$\underline{B}^F(p_2; p, q) > \underline{B}^C(p; q).$$

Hence, with a leniency program installed, the class of markets where full collusion is sustainable becomes strictly smaller under binary stochastic policies than under a constant policy with the same mean detection probability. This is because the advantage of the reduced fine is larger in the "risky state" of p_2 than the "average" state p, which makes full collusion more difficult. Since $\underline{B}^F(p_2; p, q)$ is increasing in p_2, by increasing only the higher detection probability makes full collusion more difficult.

Next, consider partial collusion deterrence. By a similar logic to the derivation of (8), partial collusion only in state p_k is sustained if and only if

$$W_k^k = B - p_k F + \delta \frac{x_k(B - p_k F)}{1 - \delta} \geqq 2B - \min\{p_k F, qF\} \tag{14}$$

$$\iff \delta \frac{x_k(B - p_k F)}{1 - \delta} \geqq B + \max\{0, (p_k - q)F\} \tag{15}$$

We have an analogous result to Lemma 1.

Lemma 4 *Assume that there is a leniency program with $q < p$. If full collusion is deterred under a binary stochastic policy, then partial collusion in which the firms collude only in state* 2 *is deterred.*

Proof. The incentive condition (14) for $k = 2$ simplifies to

$$W_2^2 = B - p_2 F + \delta \frac{x_2(B - p_2 F)}{1 - \delta} \geqq 2B - qF,$$

while the necessary and sufficient condition for full collusion under leniency was

$$V_2 = B - p_2 F + \delta \frac{B - pF}{1 - \delta} \geqq 2B - qF. \tag{13}$$

Since $x_2 < 1$ and $p < p_2$, the continuation value of collusion satisfies

$$\delta \frac{B - pF}{1 - \delta} > \delta \frac{x_2(B - p_2 F)}{1 - \delta}.$$

That is, $V_2 > W_2^2$ and the deviation value is the same. Therefore, if full collusion is deterred, the partial collusion only in state 2 is deterred. ■

Lemma 5 *Assume that there is a leniency program with $q < p$. The partial collusion only in state* 1 *is deterred for any B if and only if*

$$x_1 < \frac{1-\delta}{\delta}. \tag{9}$$

Proof. By rearrangements, the incentive condition (15) for $k = 1$ becomes

$$(x_1 - \frac{1-\delta}{\delta})B \geqq x_1 p_1 F + \frac{1-\delta}{\delta} \max\{0, (p_1 - q)F\}.$$

Since the RHS is non-negative, if $x_1 < \frac{1-\delta}{\delta}$, then the above inequality is violated for any $B(> 0)$, i.e., partial collusion in state 1 is not sustainable. ■

We summarize the above analysis as follows.

Proposition 3 [Superiority of Stochastic Policies with Leniency] *Fix a constant policy with p and a leniency program with $q < p$.*

1. *Any mean-preserving binary stochastic policy makes full collusion* ***strictly more*** *difficult than the constant policy. Moreover, sustaining full collusion becomes* ***strictly more*** *difficult as the highest detection probability p_2 is increased.*

2. *Whenever the constant policy with some amnesty deters a cartel, there always exists its mean-preserving stochastic policy that can also deter the same cartel with* ***strictly smaller*** *amnesty rates.*

3. *Any mean-preserving binary stochastic policy such that*

$$x_1 < \frac{1-\delta}{\delta} \tag{9}$$

makes ***any type of collusion strictly more difficult*** *than the constant policy.*

Proof. 1 and 2 follow from Lemma 3 and Corollary 1. 3 follows from Lemmas 4 and 5. ■

Therefore, by implementing stochastic policies, the AA can reduce the amnesty rate without compromising the effectiveness of cartel deterrence. In this way, leniency programs and non-constant cartel investigations complement each other. By an analogous argument, we can extend the superiority result to arbitrary finite-support stochastic policies.

4 Conclusion

In this paper, we study how cartel behaviors are affected by dynamic antitrust enforcement by the regulator. Our focus is to compare the (often-assumed) constant investigation policy with the mean-preserving stochastic policies that randomize cartel-detecting probabilities for each period. We illustrate that, in the simple Bertrand-type competition model and without a leniency program, the two types of policies are identically effective. Whereas, in the presence of leniency programs, some stochastic policies can outperform the constant policy in deterring collusion and reducing the amnesty rate. These findings suggest that leniency programs and fluctuations in cartel investigation complement each other. We expect this would provide a new scope for the competition policy.

To derive the above results in the simplest possible setting, the current model assumes that firms have only two actions and there are only two firms. Allowing continuous prices instead of binary actions would be straightforward. We also expect that considering more than two firms does not give qualitatively new insight as long as only the first informant gets the amnesty. However, a model with more than two firms involves more policy choice variables: how many firms can get the amnesty and how the amnesty rates should differ in the order of report. Landeo and Spier (2020) have already investigated this type of design problem of the optimal ordered leniency program. It would be an important future research topic to analyze how non-constant detection probabilities affect the optimal ordered-leniency policy.

Another simplifying assumption is that the cartel-detection probability and the level of fine are the same for (H, H) and (L, H) and is independent of B. In Appendix, we generalize these assumptions a little. In Section A.1 we show that as long as the difference of cartel-detecting probabilities between (H, H) and (L, H) is constant, the qualitative results are the same. In Section A.2., we allow that the fine level F also depends on whether the action combination was symmetric ((H, H)) or asymmetric ($(L, H)/(H, L)$), and is increasing in the profits earned. There, we show that in this case, it is optimal for the AA to charge the same fine level across action combinations and the profits earned. While a much more general structure of the cartel detection process is important, we leave it for future research.[14]

References

[1] Abreu, D. (1988). On the Theory of Infinitely Repeated Games with Discounting. *Econometrica*, 56(2), 383-396.

[2] Chen, Z. and Rey, P. (2013). On The Design of Leniency Programs. *Journal of Law and Economics*, 56(4): 917-957.

[3] Dal Bó, P. (2007). Tacit Collusion under Interest Rate Fluctuations. *RAND Journal of Economics*, 38(2): 1-8.

[4] Frezal, S. (2006). On Optimal Cartel Deterrence Policies. *International Journal of Industrial Organization*, 24(6): 1231-1240.

[5] Harrington, J. (2008a). Optimal Corporate Leniency Programs. *Journal of Industrial Economics*, 56(2): 215-246.

[6] Harrington, J. (2008b). Detecting Cartels. *Handbook of Antitrust Economics*, P. Buccirossi (ed.), MIT Press, Cambridge and London, 213-258.

[7] Harrington, J. and Chang, M. (2015) When Can We Expect a Corporate Leniency Program to Result in Fewer Cartels? *Journal of Law and Economics*, 58: 417-449.

[8] Gärtner, D. (2022) Corporate Leniency in a Dynamic World: The Preemptive Push of an Uncertain Future. *Journal of Industrial Economics*, 70: 119-146.

[9] Ivaldi, M. Jullien, B, Rey, P., Seabright, P. and Tirole, J. (2003). The Economics of Tacit Collusion. Final Report for DG Competition, European Commission.

[10] Landeo, C. and Spier, K. (2020). Optimal Law Enforcement with Ordered Leniency. *Journal of Law and Economics*, 63: 71-111.

[14] A related model is considered by Harrington and Chang (2015), where the probability of the success of investigation by the AA depends on the mass of the leniency cases and the non-leniency cases.

[11] Marshall, R. and Marx, L. (2012). *The Economics of Collusion: Cartels and Bidding Rings.* The MIT Press, Cambridge MA and London.

[12] Marvão, C. and Spagnolo, G. (2018) Cartels and Leninecy: Taking Stock of What We Learnt. *Handbook of Game Theory and Industrial Organization*, Corchón and Marini (eds.), Edward Elgar Publishing, Northampton, MA.

[13] Rotemberg, J. and Saloner, G. (1986). A Supergame-theoretic Model of Price Wars during Booms. *American Economic Review*, 76(3): 390-407.

A Appendix

Our qualitative result is robust to some changes in the model, in particular, even if the probability of getting fined at (L, H) is different from the one at (H, H), which can be more natural. Moreover, there is a foundation that the AA wants to impose the same fine level for (L, H) and (H, H), or different profits earned by a collusive firm.

A.1 Lower detection probability for a deviant

We continue to denote by p the probability of detection at (H, H) under a constant policy. Take a constant γ such that $0 \leqq \gamma < p$ and let $p - \gamma$ be the probability of detection at (L, H), that is, the firm that slightly undercuts the collusive price is charged a fine with probability $p - \gamma$ under the constant policy, probably due to its lower price than the rival firm. Then the incentive condition of (1) becomes

$$V = \frac{B - pF}{1 - \delta} \geqq 2B - (p - \gamma)F. \tag{16}$$

Take any binary stochastic policy that also differs in the probability of detection by the same γ at (L, H), given a state. That is, if the AA's realized cartel detection probability of (H, H) is p_2, the one at (L, H) is $p_2 - \gamma (> 0)$. If the detection probability at (H, H) is p_1, then the one at (L, H) is $\max\{0, p_1 - \gamma\}$. For simplicity, we focus on the binary stochastic policies such that $p_1 > \gamma$.

Without leniency, the modified incentive conditions (12) and (13) for full collusion become as follows.

$$V_1 = B - p_1 F + \delta\{x_1 V_1 + x_2 V_2\} \geqq 2B - (p_1 - \gamma)F \tag{17}$$

$$V_2 = B - p_2 F + \delta\{x_1 V_1 + x_2 V_2\} \geqq 2B - (p_2 - \gamma)F \tag{18}$$

As in our analysis in Section 2.2, (16) is equivalent to the conditions (17) and (18). Therefore, the two types of policies are identically effective in the absence of a leniency program.

Lemma 1 continues to hold in this model because the modified incentive condition of partial collusion in only state 2 is

$$W_2^2 = B - p_2F + \delta\frac{x_2(B-p_2F)}{1-\delta} \geqq 2B - (p_2-\gamma)F,$$

and this is violated if $V_2 < 2B - (p_2 - \gamma)F$, i.e., if full collusion is deterred by the binary stochastic policy.

Lemma 2 also holds. To see this, the modified incentive condition of partial collusion in only state 1 is

$$\begin{aligned} W_1^1 = B - p_1F + \delta\frac{x_1(B-p_1F)}{1-\delta} &\geqq 2B - (p_1-\gamma)F \\ \iff (x_1 - \frac{1-\delta}{\delta})B &\geqq x_1p_1F + \frac{1-\delta}{\delta}\gamma F. \end{aligned}$$

Hence the same class of binary stochastic policies prevent all kinds of collusion for the same range of B as the constant policy with p.

Next, assume that there is a leniency program. The incentive condition (10) under the constant policy becomes

$$V = \frac{B-pF}{1-\delta} \geqq 2B - \min\{(p-\gamma)F, qF\}. \tag{19}$$

The RHS of (19) implies that a leniency program is relevant if and only if $p-\gamma > q$. If we replace the relevance condition (11) ($q < p$) to $q < p - \gamma$, all analyses in Section 3.2 go through. Hence, in the class of leniency programs such that $q < p - \gamma$, the superiority of (binary) stochastic policies continues to hold.

A.2 Profit-dependent fine

In this section we show that even if the AA can set different F depending on the earned profit of a collusive firm, it is better to choose the same level. Suppose that both the detection probability and the fine level depend on whether collusion is symmetric (H,H) or asymmetric $(L,H)/(H,L)$. We keep using the same notations, p and F, for the values in the symmetric action case. Let $\tilde{p}$ and $\tilde{F}$ be the corresponding values in the asymmetric case. Then, the incentive condition to sustain collusion under a non-stochastic but action-dependent policy becomes

$$V = \frac{B - pF}{1-\delta} \geqq 2B - \tilde{p}\tilde{F} \iff B \geqq \frac{pF - (1-\delta)\tilde{p}\tilde{F}}{2\delta - 1} \tag{20}$$

Let us denote the difference of the two fine levels by ΔF, that is,

$$\tilde{F} = F + \Delta F. \tag{21}$$

Since the amount of the fine is usually non-decreasing in the amount of the (excess) profits of the colluding firms, we assume $\Delta F \geqq 0$. Substituting (21) into (20), the incentive condition is now expressed as

$$B \geqq \frac{pF - (1-\delta)\tilde{p}(F + \Delta F)}{2\delta - 1} = \underline{B} - \frac{1-\delta}{2\delta - 1}\{F(\tilde{p} - p) + \delta\tilde{p}\Delta F\} \tag{22}$$

Note the RHS of (22) is decreasing in ΔF. This implies that collusion becomes *easier* as the gap ΔF gets larger. Given that the amount of fine is non-decreasing, it is optimal to set $\Delta F = 0$. That is, for constant policies (given $\tilde{p}$), the AA should charge the constant fine $F = \tilde{F}$ independent of whether collusion is symmetric or asymmetric.