

Leniency Programs and Cartel Prosecution

Massimo Motta
European University Institute, Florence
Universitat Pompeu Fabra, Barcelona

Michele Polo
Bocconi University and IGIER, Milan ^a

May 10, 1999

Abstract

We study the enforcement of competition policy against collusion under Leniency Programs, which give reduced fines to firms revealing information to the Antitrust Authority. Such programs give firms an incentive to break collusion, but may also have a pro-collusive effect, since they decrease the expected cost of misbehaviour. We analyze the optimal policy under alternative rules and with homogeneous and heterogeneous cartels, obtaining a ranking of the different schemes and showing when the use of reduced fines may improve antitrust enforcement.

1 Introduction

The enforcement of competition policy against collusion and price fixing agreements is one of the main fields of antitrust intervention. Recent developments show that the attention devoted by antitrust authorities to collusive agreements has not diminished over time. Recently, the DGIV, the Directorate-General of the European Union in charge of competition policy, has established a special investigation unit against cartels¹, and a similar pattern can be found in the US, where the Antitrust Division of the Department of Justice has reallocated and improved the resources of the Criminal section in charge for cartel prosecution².

In the design of the policy we find today richer and more complex mechanisms than those based simply on an increase in fines. Since 1978 the US Antitrust Division of the Department of Justice has allowed for the possibility of avoiding criminal sanctions if

^aWe benefited from discussions with Luis Cabral, Federico Ghezzi, Patrick Rey, Thomas Von Ungern-Sternberg and seminar participants at the University of Salerno and the European University Institute (Fiesole). Their comments are gratefully acknowledged.

¹See Venit (1996), p.92, and European Union (1999, p.22).

²See Bingman and Spratling (1995).

specific conditions occurred. In 1993 this policy has been redesigned in the Corporate Leniency Policy, which establishes that criminal sanctions can be avoided in two cases: either if a colluding firm reveals information before an investigation is opened, as it was in the previous regime, or if the Division has not yet been able to prove collusion when a firm decides to cooperate³. The new Leniency Policy has shown in the first years of application a significant success in terms of the number of cases that the Division has been able to open and successfully conclude.

The current EU system draws from the US experience. In order to reach a more effective deterrence of collusive practices, the DGIV has initially focused its enforcement policy on a sharp increase in fines: the average fine given to firms involved in collusion cases, up to the mid Eighties has remained below 500.000 ECU while in the last decade it has reached an average of around 1.500.000 ECU⁴. However, if the firms anticipate a low probability of having collusive practices discovered (and proved), fines alone will be insufficient to prevent firms from establishing cartels. Although it is hard to quantify such expected probabilities, there seems to be a wide perception that the deterrence effects of the fines has been relatively poor, and that various types of collusive practices are still widespread. This has pushed the European Union to introduce⁵ a new regime in which reduced fines can be given to firms which cooperate with the antitrust authority by providing evidence of a collusive agreement in which they have been involved. A 75-100% reduction in fines⁶ can be given if firms reveal information before an inquiry is opened, while a lower reduction (50-75%) can be granted if cooperation occurs after an investigation has started, but that investigation has failed to provide sufficient grounds for initiating a procedure leading to a decision. A 10-50% reduction in fines can be given for partial cooperation, such as providing additional evidence or not contesting the facts on which the Commission bases its allegations. Moreover, only the first firm which cooperates can obtain a reduction, provided that it is not the promoter and major partner of the cartel. It is too early to evaluate the effects of this new policy, although the US experience suggests that enforcement against cartels might become more effective.

In this paper we want to investigate the different effects that the introduction of a Leniency Program⁷ can have on both firms' behaviour and deterrence. Our work is related to other papers on the optimal enforcement of law, specifically those on pre-trial

³Some additional restrictions on the firms entitled to benefit from this regime are introduced, as the fact that only the first can be given a fine reduction, and that it must be a junior partner in the cartel.

⁴See Furse (1995), p.114.

⁵European Union (1996).

⁶Notice that while in the US the regime applies to criminal sanctions (which include both fines and incarceration), in the EU reductions are referred only to monetary fines. Criminal sanctions do not exist under EU competition law.

⁷The US Leniency Program involves both reductions of fines and the elimination of the threat of incarceration. In this paper we focus on reduced monetary fines. Hence, we use the term Leniency Programs in a broad sense.

negotiation and settlement⁸ and on plea-bargaining⁹, in which these alternative judicial procedures have been studied with a general reference to the US judicial system, although not explicitly to antitrust law.

There are however several important differences between our work and the existing literature. The papers on pre-trial negotiation have considered mainly the properties of these procedures in saving trial costs preventing wasteful litigation. In the plea bargaining literature the enforcer acts more explicitly on behalf of taxpayers, balancing the goal of condemning the guilty agents and not condemning the innocent ones with the minimization of resources devoted to enforcement. In both cases, the issue of deterrence is not really addressed: agents have (possibly) already committed a crime, and in most papers, whether the agent is innocent or guilty and how strong is the evidence against him (agent's type) is exogenous in the model, and it is not explained in terms of incentives to commit a crime. The effects of the legal procedures on preventing the crime or making it to cease are instead at the center of our analysis.

In our paper we are mainly concerned with the deterrence and desistance properties of negotiations between the Antitrust Authority and private firms. The enforcer is motivated by the maximization of social welfare and aims at minimizing the occurrence of collusion among firms by committing on a certain set of policy parameters¹⁰. In order to focus on deterrence, in our setting we exclude other ingredients already studied in the literature. First of all we do not consider (variable) litigation costs on either party, a central issue in the pre-trial negotiation literature. The enforcer's budget is set at the beginning of the game and enforcement costs are sunk, i.e. they are already allocated among the different tasks of the organization, as general monitoring or prosecution. Secondly, we do not consider the possibility of wrong sentences, analyzed in the plea-bargaining papers: at the end of an investigation either a guilty firm is condemned or no evidence is reached.

In this setting we consider several issues. First of all we analyze the reaction of firms to different policy regimes, i.e. on the incentive to collude and on the decision to reveal or not information to the Antitrust Authority. A perverse effect can arise under this respect: since a Leniency Program allows firms to pay reduced fines, it may have ex-ante a pro-collusive effect, decreasing the expected cost of anticompetitive behaviour. But we show that, if the Antitrust Authority has limited resources, and is therefore unable to prevent collusion ex-ante, the use of Leniency Programs can improve the effectiveness of the policy, by sharply increasing the probability of interrupting collusive practices. Hence, in a second best perspective, fine reductions may be desirable because they allow to better implement ex-post desistance from collusion.

There is however a third component that operates in equilibrium: in order to induce firms to reveal, a Leniency Program has to commit resources to guarantee a sufficiently

⁸Bebchuk (1984), Nalebu (1987), Schweizer (1989), Shavell (1989).

⁹Grossman and Katz (1983), Reinganum (1988).

¹⁰Other papers that are related to our own are Kobayashi (1992) and Marshall, Muerer and Richard (1994).

high probability of independent prosecution. This is the implicit cost of a reduced fines regime, since those resources are subtracted from the general monitoring activity, which determines the frequency of "revelations" and successful inquiries. As the resources committed to prosecution become too costly, a Leniency Program loses its appeal, and a full fines regime may become more convenient again. The conditions under which these results hold will be identified in both a homogeneous and a heterogeneous cartel setting. The effects and desirability of alternative leniency rules will also be studied.

The paper is organized as follows. In section 2 we set up the basic model, in which every firm which decides to cooperate with the Antitrust Authority is given a fine reduction. In section 3 we consider alternative Leniency Programs, in which fine reductions can be granted only if cooperation occurs before an investigation is opened, or in which only the first comer, or a specific firm, is entitled to a reduced fine. Finally, in section 4 we extend the basic model to the heterogeneous cartels case. Section 5 concludes the paper.

2 The Model

Throughout the paper, we assume that the Antitrust Authority (AA from now on) aims at maximizing a utilitarian social welfare function and is able to commit to a certain set of policy parameters¹¹, which consist of full and reduced monetary fines and probabilities of enforcement. In the basic model of this section we consider a regime in which all firms which cooperate in the investigation even after this has been opened, and which simultaneously provide useful evidence to prove collusion¹², can benefit from a reduction in fines. In the following sections we shall consider alternative rules and compare them with this benchmark case.

The AA is (exogenously) endowed with a per-period budget B : in line with the literature, we assume that setting the fines at any level is not costly, while increasing the probability of enforcement requires resources. More precisely, we assume that the maximum fine that firms can receive if found guilty of collusion is exogenously given by law and equal to F , a fixed amount of money: then, being costless, it is always optimal to set the full fine at this maximum level. However, the AA can commit to a Leniency Program which allows for reduced fines $R < F$ to firms which reveal information useful to prove the existence of collusion. Indirectly, that is via the allocation of its given resources among different tasks, the AA determines the probability θ of opening an investigation and the probability p of proving firms guilty. The former refers to the

¹¹This is in line with actual experience, in which little discretion is left by the law to the Authority as to the conditions under which reductions can be given, and their amount.

¹²Throughout the paper, we assume that information given by a single firm is enough to prove that all the firms which have taken part in the collusion are guilty. This might be interpreted as the case where each firm has access to the minutes of the meetings which take place among all the colluding firms, or has copies of letters, faxes or e-mail messages which all the firms have used to coordinate on the collusive outcome. Since an important component in the working of cartels is the coordination of moves among participants, the access of each partner to some information regarding the others seems quite realistic.

preliminary activities (general monitoring) necessary to open an investigation such as collecting information about the firms in the industry, interviewing firms, suppliers and customers, collecting data from the different sources; the latter (prosecution) involves collecting further more focused information on the case, ordering surprise "raids" in the firms' headquarters, processing the information collected and preparing the case against the firms according to the existing laws. The AA, allocating resources to these two groups of tasks can obtain a combinations of these probabilities according to their specific production functions¹³. The budget constraint is then:

$$B = w_m \theta + w_p p \quad (1)$$

where w_m and w_p are the (constant) unit cost of monitoring and prosecution. We assume that firms know the probabilities θ and p chosen by the AA and its budget constraint.

The AA objective function is a standard utilitarian welfare function, i.e. the sum of producers and consumers surplus. Fines, whether full or reduced, are pure transfers, i.e. they go to the general government budget and are redistributed to consumers without distortions, and cannot be used by the AA to increase its budget. The agency problem can therefore be described as choosing the incentive scheme $(R; \theta; p)$ in order to influence firms' behaviour and maximize social welfare. The incentive compatibility constraints will be derived from the analysis of the subgame perfect equilibria in the supergame played by firms once the policy parameters are set. After observing the policy parameters chosen by the AA, n identical firms decide whether to collude or not, by correctly taking into account the probabilities $(\theta; p)$ and by knowing whether a Leniency Program R is in place or not.

We follow the usual supergame literature and consider the incentive of each firm to play an action which leads to the collusive outcome given that all other firms take the collusive action. If a firm deviates it earns a profit π_D in the current period but it triggers the punishment of the other firms, which will play the one-shot non-cooperative equilibrium action forever afterwards, by giving the deviating firm a total discounted payoff of $\pi_D + \delta \pi_N = (1 - \delta) \pi_D$. If instead the firm decides to take the collusive action, then it earns a payoff of π_M (with $\pi_N < \pi_M < \pi_D$) in the current period.

We assume that the existence of a collusive outcome in the industry is perfectly observed by the antitrust agency, but this is not enough for collusion to be proved in courts. To be able to build a case against the firms (which would otherwise win the appeal in a Court), the AA needs to find some "hard" information about collusion. Such information might consist of any document proving that firms have agreed on prices or

¹³More precisely, let the AA budget constraint be $B = w(l_m + l_p)$; where B is the total budget available to the Authority; l_m the number of hours allocated to general monitoring and l_p those devoted to prosecution, w the wage rate. In turn, the probabilities are determined given the resources according to the production functions $\theta = k_m l_m$, and $p = k_p l_p$, with θ and $p \in [0, 1]$, characterized for simplicity by positive and constant marginal productivity. Then the labor requirement to obtain θ and p are $l_m(\theta) = \theta/k_m$ and $l_p(p) = p/k_p$ respectively and the total cost of implementing θ and p are $w l_m(\theta) = w \theta/k_m = w_m \theta$ and $w l_p(p) = w p/k_p = w_p p$. It will be clear in the analysis that assuming decreasing marginal productivity, which would imply a concave budget line and a convex budget set, would not alter all our conclusions.

have met to coordinate on the prices to be charged¹⁴. Perfect observability of collusive prices also implies that the antitrust agency will never open an investigation on firms which do not collude at equilibrium.

For simplicity we consider the case where firms decide once and for all at the initial period whether to collude or to deviate from the projected cartel¹⁵. From our discussion so far, the timing of the game, represented in Figure 1, is as follows:

- t = 0 The Antitrust Authority determines the policy parameters $R; \theta; p$, which are observed by all firms. The reduced fine R is granted to any firm cooperating even after the investigation is opened.
- t = 1 Firms $i = 1; \dots; n$ decide whether to collude or deviate and realize the per-period associated payoff[®].
- t = 2 The AA opens an investigation with probability $\theta \in [0; 1]$. If the inquiry is not opened, each firm realizes the per-period profits associated to the previous choice. If the investigation starts, firms simultaneously decide whether to reveal information that the AA will find useful to prove collusion; if at least one firm reveals, the AA is able to prove them guilty. The firm(s) which cooperated with the AA pays $R - F$ while the others pay the full fine F . If no firm reveals, the AA is able to prove them guilty with probability $p \in [0; 1]$. If the AA has not been able to prove the firms guilty of collusion at the end of this inquiry, the firms will never be investigated again in the future. If proved guilty, they will behave non-cooperatively forever in the future.
- t > 2 If up to the previous period the AA has not started an investigation, with probability θ it opens an inquiry in t, firms decide whether to reveal, and so on.

Figure 1 about here

We can now solve for the equilibrium of this game. Our first step is to identify the incentive compatibility constraints, which requires to work out, for given policy

¹⁴To this purpose, note that in our model, like any repeated game with an infinite horizon, there exists a continuum of possible equilibria, and firms need some coordination to select the fully collusive outcome giving them the per-period profit π_M .

¹⁵In our setting, this is not a completely innocent assumption since the game becomes stationary only after the initial period, once firms have started colluding: considering the choice of deviating for any $t > 1$ is equivalent, since in this case firms, having participated for some periods to a cartel, pay an expected fine even if they deviate later on. When deviating initially, on the contrary, a firm can avoid the fine, since it never participated to the illegal agreement. However, notice that for this reason a deviation at the beginning is more attractive than breaking down the cartel later on, and the associated constraints are more stringent. Since the alternative case makes the analysis more complex but gives the same qualitative results, we have preferred to keep the simplest version where firms decide only at $t = 1$ whether to deviate or collude.

parameters, the subgame perfect equilibria of the game starting at $t = 1$, characterized by firms colluding or deviating and by the choice of revealing or not information to the AA. We first consider the "revelation game" which is played from $t = 2$ on if an investigation is opened by the AA. The following Lemma identifies the conditions for the existence of Nash equilibria in which firms cooperate or not with the AA.

Lemma 1 Let

$$1 - \beta > \frac{(1 - \beta)(\beta M - \beta N)}{\beta F - R} \quad \beta \in (p; F; R) \quad (2)$$

Provided that an investigation has been opened, in the "revelation" game an equilibrium always exists in which all firms reveal information.

If

- 1) $\beta F < R$ or
- 2) $\beta F \geq R$ and $\beta \geq \beta(p; F; R)$

an equilibrium exists in which no firm reveals. If this latter exists, it Pareto dominates the equilibrium outcome in which the firms reveal.

Proof: See Appendix.

2

Figure 2.a below illustrates the critical locus of points β : To the right of this curve, firms reveal if an investigation has been opened by the AA. To the left of it, they do not. This curve, which always passes through the upper right corner of the picture, rotates to the left as the reward from revealing information increases (that is, the lower R) and the larger becomes the fine F to be paid if found guilty: in other words, revelation occurs for a wider set of parameters as the incentive to cooperate with the AA is sharpened.

We can now consider the initial decision to join the proposed agreement or deviate at $t=1$. Three possible outcomes can occur: firms might prefer not to collude (NC), since they expect an immediate deviation. Alternatively, collusion may start, followed by the decision not to reveal (CNR) or to reveal (CR) if an investigation is opened by the AA. To simplify the statement of the results, it is convenient to introduce the following expressions. $\beta_{NC}(\beta; p; F)$ is the value which solves:

$$\frac{\beta M(1 + \frac{\beta(1-\beta)}{1-\beta}) + \beta p(\frac{\beta N}{1-\beta} - F)}{1 - \beta(1 - \beta)} = \beta D + \frac{\beta N}{1 - \beta} \quad (3)$$

while $\beta_{CR}(\beta; R)$ is:

$$\frac{\beta D - \beta M}{(1 - \beta)(\beta D - \beta N) - \beta R} = \beta_{CR}(\beta; R) \quad (4)$$

The following proposition identifies the conditions on the discount factor β for the three outcomes to occur.

Proposition 2 For given policy parameters (F; R; θ ; p):

- ² if $\pm_{CR}(\theta; R) \leq \pm(p; F; R)$, firms collude and reveal if monitored (CR).
- ² if $\pm \geq \max\{\pm_{NC}(\theta; p; F); \pm(p; F; R)\}$, firms collude and do not reveal if monitored (CNR).
- ² if $\pm < \min\{\pm_{NC}(\theta; p; F); \pm_{CR}(\theta; R)\}$ firms do not collude (NC).

Proof: See Appendix.

2

Figure 2.a below illustrates the line corresponding to \pm_{CR} , for given values of θ and R: this locus does not depend on p (it is flat) since in the region to the right of \pm firms cooperate with the AA once an investigation is opened and p becomes irrelevant. Above the line, firms prefer to collude even though they anticipate that, if an investigation is opened, collusion would collapse because firms would reveal information to the AA. Below the line, firms, anticipating revelation, prefer to deviate, and the collusive outcome never occurs.

Consider now \pm_{NC} , which identifies the regions where firms start colluding (above) or not (below). For $\theta = 0$ or $p = 0$, we have $\pm_{NC} = \pm_{CR} = \frac{D_1 + M}{D_1 + N}$, and the condition for collusion amounts to the "textbook" critical discount factor, which is in fact derived under the condition of no antitrust enforcement. Positive values of θ and p (and higher values of the full fine F) increase \pm_{CR} and make the cartel harder to sustain, since the expected collusive profits are reduced.

Note also that the more generous the Leniency Program (the lower the reduced fine R) the lower \pm_{CR} : if firms expect that in case an investigation is opened they have the possibility to reveal information and get away with a small fine, this will give an incentive to choose the collusive strategy. In other words, a generous Leniency Policy might stimulate ex-ante collusion. (We shall come back to this issue below.)

Figures 2.a and 2.b about here

The curves represented in Figure 2.a define, for a given θ , the conditions that must hold for a collusive agreement to emerge, and those which induce revelation or not if an inquiry is opened by the AA. More precisely, if no Leniency Program is introduced ($R = F$) firms have no reason to reveal information to the Authority if an investigation is opened, and the equilibrium outcomes would be defined uniquely by the line \pm_{NC} : Above the line, firms would collude (CNR); below, they would not (NC), because any proposed agreement would break down immediately. Reduced fines modify the situation: in the region to the left of \pm firms don't reveal if monitored, and the same argument above still applies. To the right of that curve, firms anticipate that they reveal information if monitored: below \pm_{CR} they prefer not to collude and above they initially collude and then reveal if monitored.

We can notice that the conditions for collusion are more demanding with respect to the standard case when no AA operates: the critical discount factor needed for a collusive outcome is always higher than $(\delta_{i|j}^M) = (\delta_{i|j}^N)$ when θ and p are positive. When a firm considers whether to join a cartel or deviate, in fact, it evaluates the collusive profits taking into account that with a certain probability collusion will be detected, inducing a double loss: the fine to be paid and the lost collusive profits from there on. The higher the probability of these losses, the lower the collusive profits. Hence, we need a higher and higher discount factor to balance the temptation to deviate.

To understand the role of Leniency Programs on the sustainability of collusion, consider what happens when, starting with a situation in which no Leniency Program is used, we introduce reduced fines. This has two effects which are shown in Figure 2.a. On the one hand, the Leniency Program might have an adverse, pro-collusive effect. By reducing the expected value of the fine to be paid if an investigation is opened, the Leniency Program might give an incentive to collusion. This occurs in the area (1) included between the dotted part of the curve δ_{NC} and the line δ_{CR} . In this region, no collusion can be sustained in the industry if full fines are given (NC), but under a Leniency Program firms would engage in collusion and, if monitored, they would reveal (CR) and pay the reduced fine $R < F$.

On the other hand, there exists an area (2) where collusion will break down (because the firms reveal information) if the AA starts monitoring the industry (CR), whereas in the absence of a Leniency Program collusion could stop only after a successful complete investigation (CNR). This is the area comprised between the dotted part of the curve δ_{NC} and the curve δ .¹⁶

We can now move to the analysis of the optimal policy, having identified the implementable allocations. So far we have expressed the conditions for the different equilibrium outcomes in the space $(p; \delta)$: this was useful because we obtained the conditions of cartel stability in terms of critical discount factors, thereby allowing a comparison with the modern theory of collusion. To proceed with the analysis of the optimal policy design, it is convenient to rewrite the critical loci found above in the space $(p; \theta)$ of policy parameters.

Firms would reveal if monitored if:

$$p > \frac{\delta_{i|j}^M + \delta_{i|j}^N + R(1 - \delta)}{\delta_{i|j}^M + \delta_{i|j}^N + F(1 - \delta)} = p(\delta; R; F): \quad (5)$$

Firms would prefer to collude rather than deviate, when they anticipate that the opening of an investigation would result in collusion broken down by revelations, if:

¹⁶If the Leniency Program were unanticipated, firms would decide whether to collude or not on the basis of an expected fine $R = F$ and therefore would not cooperate unless $\delta \geq \delta_{NC}$. When the leniency program is introduced unexpectedly, collusion would break down in all the area below the curve δ (that is, (1) plus (2)), without any adverse effect arising.

$$\textcircled{R} \frac{\delta M \delta D + \delta(\delta D \delta N)}{\delta(\delta D \delta N + R)} = \textcircled{R}_{CR}(\delta; R): \quad (6)$$

Finally, collusion arises in the case where firms anticipate that no revelation would occur after the opening of an investigation, if:

$$\textcircled{R} \frac{(1 - \delta)[\delta M \delta D + \delta(\delta D \delta N)]}{\delta[pF(1 - \delta) + p(\delta M \delta N) + \delta(1 - \delta)\delta M + \delta \delta N]} = \textcircled{R}_{NC}(\delta; p; F): \quad (7)$$

The three loci above allow to define, in the space of policy parameters, three regions associated with different implementable allocations, in which firms do not collude (NC), collude and reveal if monitored (CR) and collude and do not reveal (CNR) :

$$A_{NC} = f(\textcircled{R}; p) \in [0; 1]^2 \text{ j } \textcircled{R} \geq \max\{f_{NC}(\textcircled{R}); \textcircled{R}_{CR}\} \quad (8)$$

$$A_{CNR} = f(\textcircled{R}; p) \in [0; \beta] \times [0; 1] \text{ j } \textcircled{R} < \textcircled{R}_{NC}(p) \quad (9)$$

$$A_{CR} = f(\textcircled{R}; p) \in [0; \textcircled{R}_{CR}] \times [\beta; 1] \quad (10)$$

When no Leniency Program is introduced the only outcomes are NC, if $(\textcircled{R}; p)$ are above the \textcircled{R}_{NC} curve, or CNR otherwise. If $R < F$ the threshold β becomes lower than 1 and CR is an outcome if $\textcircled{R} < \textcircled{R}_{CR}$ and $p > \beta$. Notice that $\textcircled{R}_{NC}(\beta) = \textcircled{R}_{CR}$, that is the upper left corner of the region associated to CR shifts up along the \textcircled{R}_{NC} curve as R is reduced. When $R = 0$ we obtain the widest CR region.

We find also in the $(\textcircled{R}; p)$ space the same adverse effect of Leniency Programs already discussed: the intersection of A_{NC} when $R = F$ and A_{CR} when $R < F$ is non empty. That means that there are policy combinations which prevent collusion when full fines are given and that induce firms to collude and reveal once a Leniency Program is introduced.

Moreover, if $\delta < \frac{\delta D \delta M}{\delta D \delta N}$, where the latter term is the standard critical discount factor for collusion when no antitrust prosecution is considered, $\textcircled{R}_{NC} < 0$ and $\textcircled{R}_{CR} < 0$, i.e. the only admissible outcome for any value of the policy parameters is NC. Figure 2.b illustrates the equilibrium outcomes when $\delta > \frac{\delta D \delta M}{\delta D \delta N}$ and $R < F$, and it is the dual of Figure 2.a - see above.

We summarize the subgame perfect equilibrium outcomes of the supergame played by firms for given policy parameters and discount factor δ in the following proposition, which is the dual of Proposition 2.

Proposition 3 Given the gains δM and δD specified above,

- ² If the policy combination $(\textcircled{R}; p) \in A_{NC}$ there is a unique subgame perfect equilibrium in which firms will abstain ex ante from collusion (NC).
- ² If $(\textcircled{R}; p) \in A_{CNR}$ there is a unique subgame perfect equilibrium in which firms collude and don't reveal if monitored (CNR).

² If $(\theta; p) \in A_{CR}$ there is a unique subgame perfect equilibrium in which firms collude and reveal if monitored (CR). If $R = F$, A_{CR} is an empty set.

The AA chooses $(\theta; p; R)$ given the incentive compatibility constraints, summarized in Figure 2.b and Proposition 3, in order to maximize a utilitarian welfare function in which fines are pure transfers. Let $K = DWL = (1 - \beta) \sum_{t=0}^{\infty} \beta^t DWL_t$ be the discounted sum of the deadweight loss DWL, which can be thought of as the net social benefit from preventing collusion. Moreover, let W_j be the present value of the welfare gain if the policy induces the equilibrium outcome $j = NC; CR; CNR$. Then we have, for given $(\theta; p)$, $W_{NC} = K > W_{CR} = \theta K = (1 - \beta) \sum_{t=0}^{\infty} \beta^t K > W_{CNR} = \theta p K = (1 - \beta) \sum_{t=0}^{\infty} \beta^t p K$.

It is useful to identify the (welfare) indifference curves for the policy problem in the $(\theta; p)$ space: if we do not introduce fine reductions, in all the region A_{NC} we have full deterrence ex-ante and the associated welfare gains are K for all the policy parameters in the A_{NC} region. In the region A_{CNR} the indifference curve for a level of welfare gains \bar{W}_{CNR} is $\theta = \bar{W}_{CNR} / (1 - \beta) = p K / (1 - \beta)$, i.e. it is a decreasing and convex curve in the $(\theta; p)$ space: ex-post desistance in this case depends on both θ and p according to the trade-off described by the curve.

Figure 3 about here

Moreover, it is easy to show that the indifference curves in the A_{CNR} region have a shape similar to the θ_{NC} curve as defined in (7), which is the upper boundary of that region, and in the limit they overlap with that curve. In fact, if we consider the indifference curves for given \bar{W}_{CNR} and the $\theta_{NC}(p)$ curve which is the upper boundary of the A_{CNR} region and equate them we obtain after rearranging:

$$\frac{\bar{W}_{CNR}}{K} = \frac{\beta(D_i - N_i) + (D_i - M_i)}{\beta(F(1 - \beta) + M_i - N_i)} \cdot \beta$$

The right hand side expression corresponds to the upper intercept of the $\theta_{NC}(p)$ curve at $\theta = 1$, i.e. $\theta_{NC}(\beta) = 1$. Hence, looking at the expression above, if $\bar{W}_{CNR} = \beta K$ the indifference curve overlaps with the upper boundary of the A_{CNR} region, that is with the $\theta_{NC}(p)$. For $\bar{W}_{CNR} < \beta K$ the indifference curve in the A_{CNR} region shifts toward the origin.

When a Leniency Program is introduced, below the A_{NC} region we have the A_{CR} and A_{CNR} regions. The indifference curves across the region A_{CR} are $\theta = (1 - \beta) \bar{W}_{CR} = (K - \beta \bar{W}_{CR})$: those curves are horizontal, since in the CR case ex-post desistance depends only on θ . The same level of welfare in the A_{CNR} region can be obtained only if θ is higher; that means that the indifference curve is discontinuous at β and jumps up as we move from the A_{CR} to the A_{CNR} region ¹⁷.

¹⁷Notice that $W_{CNR} = W_{CR}$ for the same θ when $p = 1$. Hence, if we extend the W_{CNR} indifference curve in the A_{CR} region up to $p = 1$ we find the level of θ such that $W_{CR} = W_{CNR}$ and we are able to identify the level of the indifference curve in the A_{CR} region, as shown in Figure 3.

The iso-welfare curves in the A_{CNR} and A_{CR} regions do not identify a convex set of policy parameters. We proceed therefore convexifying the indifference curves in the following way. Consider an indifference curve in the A_{CNR} and A_{CR} region; draw a line which passes through the point of discontinuity ($\theta = (1 - \alpha)W_{CR} = (K - \alpha W_{CR})$; $p = \bar{p}$) and which is tangent to the indifference curve in the A_{CNR} region. Let the tangency point be $e(W_{CNR})$; repeating this procedure for different values of W_{CNR} an entire locus $e(W_{CNR})$ is obtained. Define E_{CNR} the subset of A_{CNR} to the left of that locus, which is represented in Figure 3. Notice that, constructing E_{CNR} , we have excluded those points on the indifference curves in the A_{CNR} region which are dominated by a combination of policy parameters in (at the boundary of) the A_{CR} region, obtaining a convex set of policy parameters.

We can now analyze the optimal policies. According to the values of B ; w_θ and w_p , i.e. the position of the budget constraint $B = w_\theta \theta + w_p p$ in the $(\theta; p)$ space, we can have different solutions to the optimal policy problem.

Proposition 4 Consider the optimal policies given the budget constraint.

- 2 If the budget constraint is above or on the $\theta_{NC}(p)$ curve, the optimal policy implements NC at a tangency point between the budget constraint and the $\theta_{NC}(p)$ curve, and the set of possible equilibrium outcomes includes all the curve, i.e. $E_{NC} = f(\theta; p) \mid p \in [0; 1]; \theta = \theta_{NC}(p)$.
- 2 If the budget constraint is below the $\theta_{NC}(p)$ curve the optimal policy implements either CR or CNR.
 - { In a CR equilibrium the optimal policy sets $R = 0$, $p = \bar{p}$ and θ along the budget constraint, and the policy combinations lie along the vertical line \bar{p} , i.e. in the set $E_{CR} = f(\theta; p) \mid \theta \in [0; \theta_{CR}]; p = \bar{p}$.
 - { In a CNR equilibrium the optimal policy combinations are at the tangency point between the budget constraint and the indifference curve.
- 2 If the budget constraint is tangent to an indifference curve in the E_{CNR} region defined above, the optimal policy implements a CNR outcome; otherwise CR is the equilibrium outcome.

Proof: See Appendix.

2

Proposition 4 gives the conditions which in general allow to identify the optimal policies for given budget constraint and it defines three sets of policy parameters which correspond to the different equilibrium outcomes, as represented in Figure 3. It is useful to consider the sequence of policy regimes that are associated with lower and lower budget constraints. Notice that two possible sequences can arise, according to the way in which the budget constraint shrinks: either we move from a NC to a CNR regime, if the budget constraint is initially very steep and the tangency point on the $\theta_{NC}(p)$ curve which implements the NC outcome lies in the neighborhood of the E_{CNR} region, or we

have, for a different budget constraints, a NC-CR-CNR sequence if the tangency point with the $A_{NC}(p)$ curve is in its lower part. This latter case seems quite interesting and allows to get the intuition of the pros and cons of the Leniency Programs.

Consider the optimal policies for parallel shifts of the budget constraint; for a relatively high total endowment a NC outcome can be implemented at a tangency point with the $A_{NC}(p)$ curve: in this case reduced fines would be harmful, inducing collusion (and revelation) when otherwise the AA would be able to prevent collusion. When the budget constraint shifts downwards and lies below the A_{NC} curve, it is no longer possible to obtain ex-ante deterrence of collusion. In this case it is optimal to implement a CR outcome by granting maximum fine discounts and setting the policy parameters along the p vertical locus: intuitively, when the AA is only able to implement ex-post desistance, reduced fines become appealing as a less costly way of proving and interrupting collusion. The implicit cost of such a policy is the need to sink resources in order to make independent prosecution a credible threat which induces revelation. As a consequence, when the total endowment is further reduced (the budget constraint shifts further downwards), fewer and fewer resources are left for general monitoring, which in the end determines the likelihood of interrupting collusion and the desirability of such a policy. At some point, we find that the (low) budget constraint becomes tangent to the iso-welfare curve in the E_{CNR} region: it means that we obtain a higher expected welfare moving to the region where firms collude and do not reveal, abandoning the Leniency Program and changing the mix of policy parameters in a more favourable way.

3 Alternative Leniency Rules

In this section we adapt the model to alternative Leniency Rules that have been adopted in the recent experience in the US and in the European Union. The first extension refers to the possibility of giving reduced fines only if firms reveal information before an inquiry is opened by the AA. Another regime assigns the reduction in fines only to the first firm which offers cooperation with the agency. Next, we suggest that if only one specific firm is entitled to benefit from a Leniency Program, this policy would be even more successful.

3.1 Fine reductions only before the inquiry is opened

As mentioned in the introduction, the initial Leniency Program introduced in the US in 1978 entitled firms with a reduction in fines only if the cooperation started before an inquiry was opened. On the same line, the actual regime chosen in the EU with the July 1996 Notice is more favourable for firms who reveal information before the AA has opened an official investigation. It is therefore interesting to analyze whether this rule can be justified in terms of enforcement effectiveness. We show that this is not the case.

Let us consider a "fine reductions only before an inquiry is opened" regime. The cor-

responding game structure is described for the general case of n firms in the following¹⁸:

- $t = 0$: The AA sets the policy parameters $\theta; p; R$ which are observed by the firms.
- $t = 1$: Firms $i = 1; \dots; n$ decide whether to collude or deviate and realize the associated payoffs.
- $t = 2$: At the beginning of the period, firms simultaneously choose whether to reveal the existence of the cartel to the AA, benefiting of reduced fines, or not; if no firm reveals, the AA opens an investigation with probability $\theta \in [0; 1]$, proving them guilty with probability $p \in [0; 1]$. Then, payoffs are realized.
- $t > 2$: if up to the previous period the AA has not started an investigation, the game restarts as from $t = 2$, etc.¹⁹

Consider first the subgame starting at $t = 1$ after a decision to collude. To find the conditions under which not revealing is an equilibrium, we have to compare the payoff from revealing when the other firms do not reveal, namely $\frac{1}{1-\theta} \pi_i^R$, with the payoff from not revealing when the other firms do not reveal. The latter is given by:

$$\pi_{nr} = \theta [p(\frac{1}{1-\theta} \pi_i^F) + (1-p)(\frac{1}{1-\theta} \pi_i^M)] + (1-\theta)(\pi_i^M + \pi_{nr});$$

whence:

$$\pi_{nr} = \frac{\theta [p(\frac{1}{1-\theta} \pi_i^F) + [(1-\theta) + \theta(1-p)] \frac{1}{1-\theta} \pi_i^M]}{1-\theta + \theta(1-\theta)}$$

It is simple algebra to check that this payoff is higher than $(p(\frac{1}{1-\theta} \pi_i^F) + (1-p)\frac{1}{1-\theta} \pi_i^M)$, the expected payoff from not revealing after the investigation has been opened, which was the relevant one under the rule analyzed in the previous section. Since the payoff from revealing is the same in both cases, it follows that the equilibrium where firms do not reveal is more likely to occur when the Leniency Program is applied only for revelations before the inquiry is opened. In other words, the curve θ moves to the right and collusion is less likely to be broken by revelations in this regime. This is hardly surprising, because the probability of the event "being found guilty and thus fined" is lower before seeing if the industry will be monitored than after an investigation is actually opened.

We have now to consider if the Leniency Program might change the ex-ante incentives of firms to collude. It turns out that there would never be collusion in the industry when firms expect that there would be revelation of information to the AA in the following period: this implies that an equilibrium in which firms choose to collude and reveal does not exist. In fact, by colluding when expecting the cartel to be broken by information

¹⁸The payoffs in the different outcomes are similar to the model analysed above, and will be omitted here in the description of the game.

¹⁹Allowing firms to choose whether to reveal or not before an investigation is opened at any period would not change the results.

given to the AA, a firm would get $V_c = \frac{1}{M} + \frac{1}{N}(1 - \frac{1}{M})R$. By deviating, it would get $V_d = \frac{1}{D} + \frac{1}{N}(1 - \frac{1}{D})$. Since $\frac{1}{D} > \frac{1}{M}$ and $R > 0$, it follows that $V_c < V_d$.

In the case, considered in the previous section, where firms were entitled to fine discounts after the opening of an investigation, the expected profit from collusion decreases when the event "opening of an investigation" realizes, leading firms to reveal information to the agency. In the case we are considering here, instead, nothing new happens between the moment when the firms decide on collusion and the moment when they are asked to cooperate with the authorities to break down the cartel.

Our analysis reveals that if Leniency Programs are to be effective in breaking down cartels, they should be extended to benefit firms which reveal after the industry is put under monitoring. Since proving firms guilty of collusion is a very lengthy and complex issue, which does not always end up with the firms being condemned, a great amount of resources can be saved and a final positive outcome guaranteed by ensuring that firms have the proper incentives to collaborate with the AA even after an investigation has been started.

This result is consistent with the US experience, where initially the Leniency Program was used only for firms which spontaneously offered evidence before the inquiry was opened by the AA. In this initial regime the program was quite ineffective while, once allowed in 1993 for reduced fines even after the inquiry was opened, the number of cases in which firms cooperate with the judges increased significantly. In the 1994 Annual Report of the Antitrust Division it is stressed that in the first year of the new regime "an average of one corporation per month come forward with information on unilateral conspiracies, compared to an average of one per year under the previous policy. The policy thus allowed the Division to extend the reach of its criminal enforcement activities with relatively little expenditure of resources"²⁰.

According to our results, the new regulation on Leniency Programs²¹ adopted by the EU should be widened. The regulation states that firms which denounce a cartel before the Commission has opened an investigation are entitled to a reduction of 75-100% of the fines. Firms which denounce a cartel after that a "verification" has been opened are entitled to a 50-75% reduction of the fines, but only if those verifications had not been fruitful and had not led to the opening of a procedure. Basically, this means that Leniency Programs are opened only for firms operating in industries which are not under the scrutiny of the AA. This narrows too much the scope of the application of the regulation, and fails to provide the firms with enough incentives for revealing information which can be useful to break the cartel.²²

²⁰Antitrust Division (1994), p.6-7.

²¹See Official Journal of the European Communities, Series C, 207, 18-7-1998.

²²Taken literally, our analysis would also suggest that when firms reveal they should always receive a zero fine ($R = 0$), since this would give them the greatest incentive to denounce the cartel. However, we are assuming that collaborating with the AA is a binary variable. Either one does not collaborate, or if it does it can give all the information necessary to prove the participation in the cartel of all the firms. In reality, the type of information that firms can provide would be more of a continuous variable, and tuning the fine reductions to the quality of the information revealed makes sense.

3.2 Only the first comer obtains a fine reduction

The criteria that determine which firms can receive the benefits of a reduced fine have been restricted in different ways both in the US and in the EU experience. An interesting case is where only the first firm which offers evidence is given a fine reduction, as it is the case in the EU regulation.

In this case the game structure is the same as in our initial model. The only difference is that if all firms decide to reveal information to the judges, as it happens in a subgame perfect equilibrium in which firms reveal if monitored, the expected payoff becomes:

$$\frac{1-N}{1-\beta} \left(\frac{R + (n-1)F}{n} \right) \quad (11)$$

where n is the number of firms in the cartel: every firm is ex-ante the first one to disclose information to the AA with probability $1/n$. Notice, however, that when we check for the existence of an equilibrium in which no firm reveals, a deviating firm obtains the reduced fine R for sure, being the only one which cooperates with the judges. Hence, the condition for an equilibrium in which no firm reveals is $\beta \geq \beta^*$, exactly as in the case treated above.

Moreover, it is easy to see that if an equilibrium exists in which no firm reveals if monitored, it also Pareto dominates the equilibrium in which all firms cooperate with the AA.

Consider now the decision of firms on collusion at $t = 1$: if $\beta \geq \beta^*$ firms will not reveal if monitored and everything is as in the basic model. However, if $\beta < \beta^*$, revelation will follow the opening of an investigation, but firms' incentives to collude are modified in the present regime, since the expected payoff if monitored is lower than in the previous case where all firms could benefit from the Leniency Program.

One can check that firms will abstain from collusion if

$$\beta \geq \frac{(1-\beta)(1-N)\beta + (1-\beta)(1-N)\beta + (1-\beta)(1-N)\beta}{(1-\beta)(1-N)\beta + (1-\beta)(1-N)\beta + (1-\beta)(1-N)\beta} \geq \beta_{CR}^1 \quad (12)$$

It is immediate to notice that $\beta_{CR} < \beta_{CR}^1$, that is, the region of parameters that induce firms to abstain from collusion is larger than in the previous "all firms get the reduction" regime - see Figure 4.

Figure 4 about here

The intuition of this result is as follows: in the more restrictive set of rules analyzed in this section, the expected reduction in fines is smaller when all firms choose to cooperate with the judges, although it is equivalent when we consider the incentive for a firm to cheat the partners when they do not reveal. Hence, when firms anticipate that they all will confess if monitored, they expect higher sanctions. Consequently, in some cases they find it less attractive to collude and reveal as an alternative to deviating from the

beginning and avoiding the fine. The regime therefore is able to partially reduce the ex-ante incentive to collusion without reducing the power of the program in making firms denounce a cartel after an inquiry is opened.

This case suggests an alternative rule which might increase the effectiveness of a Leniency Program, by further reducing the ex-ante incentive of engaging in collusion induced by the expected reduction in fines.

3.3 Only a specific firm receives a fine reduction

23

As we have repeatedly emphasized, a Leniency Program influences firms in two ways. The first is that it stimulates ex-post breaking of the cartel via revelation of information to the AA; the second (adverse effect) is to increase the incentive of collusion via a reduction in the punishment in case of being found guilty. We have also seen that granting a reduction in fines only to the first firm which reveals works because it leaves unchanged the first effect but reduces the second. The effectiveness of the Leniency Program could be increased even further by increasing asymmetries in the industry and specifying ex-ante that only a specific firm could be entitled to the LP, no matter the way in which such a firm is selected. The way of interpreting this rule is that of defining ex-ante a set of parameters which allow all the participants in each specific situation to identify a single firm entitled to a reduced fine²⁴: the firms involved in the cartel, applying the rule, are able to work out which one will be the firm selected. Denote this firm with a number, say 1. The conditions under which revelation occurs are the same as usual: If $\alpha < \beta$, the cartel would break because firm 1 denounces it. On the other hand, the conditions under which ex-ante collusion occurs will change. For the $n - 1$ firms which are not entitled to the Leniency Program, the condition for taking part in the collusion will be:

$$\alpha > \frac{\beta D_i + M}{(1 - \beta)(\beta D_i + N) + \beta F} - \alpha_{CR}^F \quad (13)$$

For firm 1, the condition is laxer:

$$\alpha > \frac{\beta D_i + M}{(1 - \beta)(\beta D_i + N) + \beta R} \quad (14)$$

However, since all the firms must find it convenient to take part in collusion, the latter condition does not play any role, while the former is binding and determines the existence of collusion. Also, notice that $\alpha_{CR}^F > \alpha_{CR}^I > \alpha_{CR}$ - see Figure 4. In other words, if the authority targets the Leniency Program to a specific firm, it will be able to reduce the ex-ante adverse effect of it without decreasing the ex-post incentive to reveal information. Hence, collusion becomes less likely because the firms excluded from the program find it less appealing to engage in a cartel which includes a likely cheater²⁵.

²³We thank P.Rey for suggesting this extension.

²⁴For instance, it might be the firm located in the smallest city, or the last one in alphabetical order, etc.

²⁵Of course, leniency rules which limit the applicability of the fine reduction to only one firm will result

4 Heterogeneous Cartels

So far we have considered homogeneous cartels, in which the payoffs in each possible outcome were the same across partners. Notice however that, in all our arguments, if the participants have heterogeneous payoffs and they know the payoffs of the partners in each possible outcome, the equilibria are governed by the conditions of one of the firms, the one whose constraints bind. This decisive agent is the point of reference for the others, whether they expect such firm to deviate or to reveal information after colluding, and drives the equilibrium conditions of the entire cartel. Hence, in a sense, our previous analysis allows to consider heterogeneous firms within a cartel, but it assumes that in each cartel in the economy such a decisive agent is always the same. It is therefore interesting to consider the case in which the cartels are truly heterogeneous, in the sense that the participants may differ in payoffs and the decisive partner may be different across cartels.

We consider in this setting the design of an optimal enforcement policy which cannot be made conditional on cartel's type, due to informational and/or institutional restrictions. Hence, the AA has to design a single, general policy facing many different industries, characterized by heterogeneous market conditions and potential gains from collusion. In this case, the policy implemented will induce different effects in the various industries, reaching a more or less effective deterrence of collusion and inducing different types to choose different reactions: hence we might have some cartel types colluding and not revealing while others will prefer not to join the cartel; or we might have all types colluding, but only a subset of them revealing information when monitored, etc. Hence, the different effects that we have identified in the previous sections will be combined in a richer way once the AA faces heterogeneous types.

From the previous analysis we already know that the incentive compatibility constraints for given policy parameters depend on two variables of cartel's type: π_M and π_D . Hence, multiple types would require to deal with a bivariate distribution, related to the gains from collusion and from deviation. To maintain the analysis simple, we assume in this section Bertrand competition (with constant marginal costs) in the non-cooperative equilibrium: hence $\pi_N = 0$ and $\pi_D = n\pi_M$: the gains from collusion are now perfectly correlated to those from deviation, and we can consider a univariate distribution of types. Cartel types refer to the gains from collusion, due for example to different marginal cost levels, with $\pi_M \in [\underline{\pi}_M; \overline{\pi}_M]$; the AA does not observe cartel types but knows their distribution $g(\pi_M)$, and is not able to condition the policy chosen to some observable that can make it contingent on a message. In other words, the AA sets a single combination of policy parameters taking into account that there exist many cartel types in the economy.

Under the assumption of Bertrand competition the standard critical discount factor when antitrust is absent, $\alpha^* = (\pi_D / \pi_M) = (\pi_D / \pi_N)$, is $(n - 1)/n$. We can rewrite the

in a larger amount of money collected through fines. In a world where non-distortionary transfers are not available, this would be an additional advantage of such rules.

relevant loci as:

$$\mathbb{R}_{CR} = \frac{\lambda_M(1 - \lambda - n + n\lambda)}{\lambda_M n\lambda} = (\lambda - \lambda^2) = \lambda$$

which does not depend on cartel's type,

$$\mathbb{R}_{NC}(p) = \frac{(1 - \lambda)n(\lambda - \lambda^2)\lambda_M}{\lambda[pF(1 - \lambda) + \lambda_M(p - n(\lambda - \lambda^2))]}$$

and

$$p = \frac{\lambda_M}{(\lambda_M + F(1 - \lambda))}$$

which are both increasing in λ_M . Moreover, \mathbb{R}_{CR} is always above \mathbb{R}_{NC} at $p = 1$. Hence, when $R < F$ we can distinguish 5 regions which are represented in Figure 5.

Figure 5 about here

In region A all types choose CNR and the corresponding welfare is $W_A = \frac{\mathbb{R}_p}{1 - \lambda + \mathbb{R}_\lambda} E(K)$ where $E(K)$ is the expected value of the gains from deterrence given the distribution of types $g(\lambda_M)$. In region B all types choose CR with $W_B = \frac{\mathbb{R}_\lambda}{1 - \lambda + \mathbb{R}_\lambda} E(K)$ while in E all types abstain from collusion and welfare is $W_E = E(K)$. In region C some types choose CNR and others CR: let λ_M^C be the type whose p equals the p chosen by the AA in region C: all types lower than λ_M^C collude and reveal while the cartels more profitable collude and don't reveal. The expected welfare is therefore

$$W_C = \frac{\mathbb{R}_\lambda}{1 - \lambda + \mathbb{R}_\lambda} \int_{\lambda_M^C}^{\lambda_M} K(\lambda_M)g(\lambda_M)d\lambda_M + \frac{\mathbb{R}_p}{1 - \lambda + \mathbb{R}_\lambda} \int_{\lambda_M^C}^{\lambda_M} K(\lambda_M)g(\lambda_M)d\lambda_M$$

Analogously, in region D lower types abstain from collusion and higher types collude and don't reveal, with the threshold type λ_M^D such that the actual policy combination in D lies on that type's $\mathbb{R}_{NC}(p)$ curve. The expected welfare is then

$$W_D = \int_{\lambda_M^D}^{\lambda_M} K(\lambda_M)g(\lambda_M)d\lambda_M + \frac{\mathbb{R}_p}{1 - \lambda + \mathbb{R}_\lambda} \int_{\lambda_M^D}^{\lambda_M} K(\lambda_M)g(\lambda_M)d\lambda_M$$

When $R = F$ only regions A,D and E exist, defined by the set of $\mathbb{R}_{NC}(p)$ curves which extend up to $p = 1$.

The analysis of the optimal policy proceeds in three steps, which are developed analytically in the Appendix. First, the iso-welfare curves in each of the five regions A-E are derived; then, we check how the same welfare level is obtained passing (eventually) from one region to the neighbouring one, distinguishing whether these reductions are given or not; finally, comparing the two cases, it is selected whether reduced fines R allow to save enforcement costs, defining a set of iso-welfare curves along which Leniency Programs are optimally used.

The result of this analysis is shown in Figure 6.a: the lower bold curve is the iso-welfare (cost minimizing) curve setting $R = 0$, which passes through the regions A-C-B. The upper bold curve passing in the D-C-B regions entails the use of reduced fines only in a subset of the B and C regions. The policy combinations $(R; \rho; p)$ which minimize the cost of reaching the same expected welfare are summarized in a map of iso-welfare curves which are not convex: as before, we have to convexify them excluding from the set of possible equilibrium outcomes those policy combinations which belong to the non-convex portions of the indifference curves.

Given the map of indifference curves that minimize the cost of a given expected welfare, we exclude those portions which can never be selected given our linear budget constraint²⁶. For the indifference curves in the A region we obtain a subset of points E_{CNR} analogous to the one obtained in the single type case already discussed. In region C we find a subset of points E_{CNR}^{CR} in which some types choose CR and higher types choose CNR. In region B we select only the boundary to the left, which corresponds to the E_{CNR}^{CR} case when all types opt for CR. A subset of D, E_{CNR}^{NC} is obtained where low type select NC and high types choose CNR, and finally from region E we select the lower bound. Once excluded the non-convex portions of the iso-welfare curves, the optimal policies for given budget constraint can be established along the same lines of Proposition 4's proof. We summarize the results in the following Proposition, which is represented in Figure 6.b.

Proposition 5 Consider the optimal policy under asymmetric information given the budget constraint and the distribution of cartel types.

- ² If the budget constraint passes through region E, the optimal policy implements NC for all types at a tangency point between the budget constraint and the $\rho_{NC}(p)$ curve of the highest type.
- ² If the budget constraint passes through region D and is tangent to an indifference curve in E_{CNR}^{NC} , the optimal policy is at the tangency point with no fine reduction, and implements a CNR-NC outcome according to the different types.
- ² If the budget constraint passes through region A and is tangent to an indifference curve in E_{CNR} , the optimal policy is at the tangency point and implements CNR for all types.
- ² If the budget constraint passes through C and is tangent to an indifference curve in E_{CNR}^{CR} , that is the optimal policy and implements a CNR-CR outcome.
- ² In all the other cases the optimal policy implements CR for all types setting p equal to the ρ of the highest type along the budget constraint.

²⁶Any convex budget set, as that obtained under the assumption of decreasing returns to enforcement, would allow a similar exercise.

Figure 6.a and 6.b about here

We can give an explanation of the result above considering the sequence of optimal policies when the budget constraint becomes steeper and steeper as a result of an increase in the cost of independent prosecution (higher w_p). For low values of w_p the policy implements full deterrence ex-ante for all types. As the budget constraint rotates toward the origin we initially move to a CNR-NC mixed outcome with no τ -ne reduction, in which the more profitable cartels are not deterred. Granting τ -ne discounts in this case would shift low types from NC to CR: the pro-collusive effect of Leniency Programs would dominate reducing welfare. However, when the fraction of low types which choose NC shrinks further, reduced τ -nes are introduced, inducing all types to collude and reveal. In this case the improvement in prosecution allowed by reduced τ -nes becomes predominant. A further increase in w_p moves the equilibrium outcome in the E_{CNR}^{CR} region with an increasing portion of high types that choose CNR while low types collude and reveal. The implicit cost of the Leniency Programs, which forces the AA to commit resources to independent prosecution to make it a credible threat, becomes heavier and heavier as the resources left to open inquiries decrease and as the fraction of types which are induced to reveal shrinks. In the end we move to the E_{CNR} region, abandoning the Leniency Program. Hence, the optimal policy is determined, in a sense, by the relative importance of the pro-welfare effect of Leniency Program, that allows to obtain more effective ex-post desistance, and the welfare decreasing effects of reduced τ -nes: the incentive to collude (and reveal) instead of abstaining from collusion, and the need to sink resources to make independent prosecution credible, which reduces the probability of opening an inquiry and of obtaining ex-post desistance.

5 Conclusions

In this paper we have analyzed the effects of Leniency Programs on the incentives of firms to collude and to reveal information that helps the Antitrust Authority to prove illegal behaviour. The benchmark regime gives to any firm a reduction in τ -nes even if revelation occurs after an investigation is opened. We show that reducing the expected τ -nes may induce a pro-collusive reaction: combinations of policy parameters which, without Leniency Programs, would prevent collusion, may induce firms to collude (and reveal if monitored) when τ -ne reductions are given. Hence, if the resources available to the AA are sufficient to prevent collusion using full τ -nes, Leniency Programs should not be used.

However, when the AA has limited resources, Leniency Programs may be optimal in a second best perspective. Fine reductions, inducing firms to reveal information once an investigation is opened, increase the probability of ex-post desistance and the expected welfare gains. The optimal scheme requires maximum τ -ne reductions and a shift of resources from prosecution to monitoring.

A fixed amount of resources, however, must be committed in any case to make

independent prosecution a credible threat, since no firm would reveal if it expects that the AA is unable to prove them guilty. When independent prosecution is very costly, too few resources are left to general monitoring, which in the end determines the effectiveness of Leniency Programs. In this case it may become more convenient to shift back to a full fines regime with a more favourable mix of policy parameters.

We have then compared our benchmark regime with alternative sets of rules: the first allows to give fine reductions only to firms which cooperate with the Antitrust Authority before an inquiry is opened, as initially established in the US policy in 1978, and similar to the approach followed by the EC Notice on the non-imposition of fines, and we proved this regime to be inferior with respect to our benchmark case. We have then considered other rules which restrict the set of the firms that can benefit from a Leniency Program. We showed that by granting a fine reduction only to the first firm which cooperates with the AA the perverse pro-collusive effect of the Leniency Program would be reduced without softening the incentives to reveal information. Better still, the AA might target a specific firm and allow only this one to benefit from the reduction. The intuition for this result, which makes the Leniency Program even more effective, lies in the asymmetry that the policy introduces among otherwise identical firms, between the entitled firm and the excluded ones: the latter would more often prefer to abstain from collusion rather than join a cartel together with a likely cheater.

Finally, the case of multiple cartel types has been considered: the AA is assumed to be unable, for informational or institutional reasons, to implement Leniency Programs contingent on cartel's type, and therefore has to set general rules. For instance, the AA cannot shape the policy to the conditions of each specific industry, but has to choose a general rule of behaviour, obtaining different effects in different industries. Then, according to the position of the budget constraint in the set of policy combinations, we characterized the optimal policy: it turns out that the policy parameters and the regime of full or reduced fines are chosen according to the relative weight of the three effects described above, where the weights depend on the share of types which choose the different equilibrium outcomes (no collusion; collusion; collusion and revelation).

We believe that, despite the simple setting, our paper sheds some light on the desirable features of leniency programs, and suggests some changes in the EC leniency policy. First of all, if it is optimal to use a leniency program (as in the realistic case where the antitrust agency has limited resources), then the program should be as generous and certain as possible with the firms which provide fresh evidence that establishes the existence of a cartel. In contrast, the EC policy of keeping some degree of discretionality instead of granting automatic and total reduction of fines even to those firms which fulfil all the (strict) conditions laid down in the EC Notice undermines the success of the leniency program, as it does not give certainty to the prospective cooperating firm and reduces the incentive to break the cartel.

Likewise, some of the conditions required by the EC policy are too strict. For instance, a firm must "maintain continuous and complete co-operation throughout the investigation" to be entitled to have a very substantial reduction (more than 75%) of

the fine. This has led the Commission to give only a 50% reduction to a firm, Tate & Lyle, which had spontaneously brought conclusive evidence of a cartel to the attention of the Commission (at a time when the Commission did not even suspect the existence of an agreement), but had later (partially) contested some of the allegations made by the Commission²⁷. The strict wording and application of the Notice will reduce the incentive of the firms to reveal information²⁸.

Furthermore, our analysis indicates that a leniency program should be equally applicable to information disclosed before and after an investigation has started, whereas the EC policy does not create enough incentives for post-investigation disclosure of information. It gives only 50-75% reduction of the fines for cooperation after an investigation has been undertaken already but only if such an investigation has failed to provide sufficient material for initiating a procedure leading to a decision.

The US experience (where after the 1993 policy revision a corporation is granted leniency after an investigation has begun provided that "the Division, at the time the corporation comes in, does not yet have evidence against the company that is likely to result in a sustainable conviction" - point B.2.) clearly shows that extension of the leniency program to post-investigation amnesty (along with the automatic granting of the amnesty) is a crucial ingredient for success: "...under the old policy on average only one corporation per year applied for amnesty," (Spratling (1998, page 2) whereas under the revised policy, "Amnesty applications over the past year have been coming in at the rate of approximately two per month" (Spratling (1999, page 2).

So far, the leniency program of the EC has been applied to a very reduced number of cases, since its introduction in the end of 1996. There was no case in 1997 and only four in 1998²⁹. We believe that granting higher and automatic reductions of fines and extending the leniency program to after-investigation cooperation would greatly increase the success of this policy.

²⁷This is the case "British Sugar", EC Decision of 14 October 1998, published in the Official Journal of the EC, L76, 22 March 1999.

²⁸Hornsby and Hunter (1997) also point out that the firms do not have enough incentives to cooperate under the EC policy. Part of the problem is also due to the fact that the Notice cannot provide immunity from civil proceedings. Admission of an infringement leads to a formal Commission Decision on which an action for damages can be built, without the plaintiff[®] having to prove the infringement again. This problem does not exist in the US, where the cooperating firms can resort to a consent decree.

²⁹Information provided by an EC official, Guerrin (1999). Of these, three regarded instances of minor cooperation (firms were given discounts for not having contested the Commission's allegations or for providing additional evidence which helped establishing the facts). These cases are "Alloy surcharge", EC Decision of 21 January 1998; "Pre-insulated pipes", EC Decision of 21 October 1998; "Greek Ferries", EC Decision of 9 December 1998. The fourth case was the "British Sugar" case reported in the previous footnote, which might be a discouraging precedent for firms considering cooperation with the Competition Commission. According to Guerrin, in some further half a dozen current cases the Leniency Notice has been invoked. No further details were given for reasons of confidentiality.

References

- [1] Antitrust Division (1994), Annual Report for Fiscal Year 1994.
- [2] Bebchuk L. (1984), Litigation and Settlement under Imperfect Information, *Rand Journal of Economics*, 15, 404-413.
- [3] Bingman A., Spratling G., (1995) Criminal Antitrust Enforcement, mimeo.
- [4] European Union (1996), Notice on the Non-Imposition or Reduction of Fines in Cartel Cases, *Official Journal*, v.207, p.4.
- [5] European Union (1999), Competition Policy Newsletter, No.1, February.
- [6] Furse M., (1995), Article 15(2) of Regulation 17: Fines and the Commission's Discretion, *European Competition Law Review*, p.110.
- [7] Grossman G., Katz M. (1983), Plea-Bargaining and Social Welfare, *American Economic Review*, 7, 749-757.
- [8] Guerrin M.(1999), Private communication. Electronic message of 5 March 1999.
- [9] Hornsby S., Hunter, J. (1997), New incentives for "whistle-blowing": will the E.C. Commission's Notice bear fruit?, *European Competition Law Review*, 1, 38-41.
- [10] Kobayashi B. (1992), Deterrents with Multiple Defendants: an Explanation for "Unfair Plea Bargains", *Rand Journal of Economics*, 22, 507-517.
- [11] Marshall R., Muerer M., Richard J. (1994), Litigation, Settlement and Collusion, *Quarterly Journal of Economics*, 1, 211-239.
- [12] Nalebu® B. (1987), Credible Pre-Trial Negotiation, *Rand Journal of Economics*, 18, 198-210.
- [13] Reinganum J. (1988) Plea-Bargaining and Prosecutorial Discretion, *American Economic Review*, 78, 713-728.
- [14] Reinganum J., Wilde L. (1987), Settlement, Litigation and the Allocation of Litigation Costs, *Rand Journal of Economics*, 17, 557-566.
- [15] Schweizer U. (1989), Litigation and Settlement under Two-Sided Incomplete Information, *Review of Economic Studies*, 56, 163-178.
- [16] Shavell S. (1989), Sharing Information Prior to Settlement or Litigation, *Rand Journal of Economics*, 20, 183-195.
- [17] Spratling G.R. (1998), The corporate leniency policy: answers to recurring questions, speech of the Deputy Assistant Attorney General, presented at the Bar Association of the District of Columbia, February 16, 1999.

- [18] Spratling G.R. (1999), Making companies an offer they shouldn't refuse, speech of the Deputy Assistant Attorney General, presented at the ABA Antitrust Section 1999 Spring Meeting, April 1, 1998.
- [19] Venit J. (1996), EU Competition Law - Enforcement and Compliance: An Overview, Antitrust Law Journal, p.87.

Appendix

Proof of Lemma 1

If a firm reveals, it gets a payoff of $\frac{1}{1-i} \pi$ independently of the action chosen by the other firms. If a firm does not reveal any information but at least one firm does, then the former firm receives a payoff of $\frac{1}{1-i} \pi$. Hence, it is always (weakly) better to reveal if the other firms are expected to reveal, which establishes the existence of the "revelation" equilibrium. Finally, if no firm reveals any information, each firm receives an expected payoff

$$p \left(\frac{1}{1-i} \pi \right) + (1-p) \frac{1}{1-i} \pi \quad (15)$$

If a firm expects the others not to reveal, the best reply is trivially not to reveal if $pF < R$. If $pF \geq R$, when the other firms don't reveal, a firm prefers not to reveal as well if the payoff above is higher than $\frac{1}{1-i} \pi$, which simplifies to $\pi \geq \pi(p; F; R)$. Hence, in this case a "no revelation" equilibrium exists. Moreover, the same inequality implies that the "no revelation" equilibrium gives higher payoffs to all firms than the "revelation" equilibrium.

Proof of Proposition 2

We consider the decision to collude or deviate in both cases, when a firm will decide to reveal if investigated, and when they will prefer not to cooperate with the AA.

- ² Case 1: $\pi < \pi$. In this case firms reveal if an investigation is opened by the AA. Define π_R as the expected profit immediately before an investigation is opened. It is easy to see that:

$$\pi_R = \pi \left(\frac{1}{1-i} \pi \right) + (1-\pi) (\pi_M + \pi \pi_R)$$

from which we obtain:

$$\pi_R = \frac{(1-\pi) \pi_M + \pi \left(\frac{1}{1-i} \pi \right)}{1-i + \pi}$$

If a firm decides to set the collusive price, then its expected discounted payoff will be:

$$V_{CR} = \frac{1}{i} M + \frac{1}{i} R = \frac{\frac{1}{i} M + \frac{1}{i} (1 - i) R}{1 - i + i(1 - i)}$$

If instead a firm decides to deviate from the collusive strategy, then its payoff is given by:

$$V_D = \frac{1}{i} D + \frac{1}{i} N$$

Collusion can arise if $V_{CR} \geq V_D$, that is if the following condition is satisfied:

$$i \geq \frac{\frac{1}{i} D + \frac{1}{i} M}{(1 - i)(\frac{1}{i} D + \frac{1}{i} N) + i R} \equiv i_{CR}(R); \quad (16)$$

² Case 2: $i \geq i$. In this case, firms anticipate that even if an investigation is started, no firm will reveal any information. Collusive outcome will be obtained unless the AA can prove the firms guilty of collusion.

Write the expected profit immediately before knowing if an investigation is opened as:

$$i_{NR} = \frac{1}{i} [p(\frac{1}{i} N + i F) + (1 - p)(\frac{1}{i} M)] + (1 - i) (\frac{1}{i} M + \frac{1}{i} NR)$$

We can then obtain:

$$i_{NR} = \frac{\frac{1}{i} [p(\frac{1}{i} N + i F) + (1 - p)(\frac{1}{i} M)] + (1 - i) \frac{1}{i} M}{1 - i + i(1 - i)}$$

If a firm follows the collusive strategy its expected discounted payoff is given by:

$$V_{CNR} = \frac{1}{i} M + \frac{1}{i} i_{NR} = \frac{\frac{1}{i} M (1 + \frac{1}{i} (1 - p)) + \frac{1}{i} p (\frac{1}{i} N + i F)}{1 - i + i(1 - i)}$$

As before, a firm which deviates obtains a payoff:

$$V_D = \frac{1}{i} D + \frac{1}{i} N$$

The inequality $V_{CNR} \geq V_D$ implicitly defines the locus of points $i_{NC} = i_{NC}(R; p; F)$.

Proof of Proposition 4

We proceed in two steps: first we show, for each given outcome NC, CR, CNR, which is the associated optimal policy; second, we show the conditions under which a particular outcome is better than the others. If a NC outcome is implemented, we can have two cases: either the budget constraint is above the lower boundary of the A_{NC} region or it is

tangent: in the latter case the tangency point with the \mathbb{W}_{NC} curve is trivially the optimal solution; if the budget constraint is above that curve, the tangency point can still be suggested under a cost saving argument. In this case we set $R = F$ since granting τ discounts would shrink the A_{NC} region. If a CR outcome is implemented, the welfare gain depends only on \mathbb{W} , which therefore must be maximized: we can therefore set $R = 0$, shifting to the left the \mathbb{p} threshold, and setting $\mathbb{p} = \mathbb{p}$; with \mathbb{p} at its lowest level in the A_{CR} region we can set $\mathbb{W} = B = W_{\mathbb{W}} | (w_{\mathbb{p}} = w_{\mathbb{W}}) \mathbb{p}$ along the budget constraint. Finally if a CNR outcome is chosen, a tangency point between the budget constraint and the indifference curve in the A_{CNR} region must be chosen.

Consider now the choice among the three outcomes: since W_{NC} is always dominant for any set of policy parameters, if the budget constraint is not below the $\mathbb{W}_{NC}(\mathbb{p})$ curve, NC is the optimal outcome implemented. The choice between a CR and a CNR outcome is more complex, since both W_{CNR} and W_{CR} depend on the associated policy parameters, which, in turn, are different at the optimal points in the two regimes. Suppose the tangency point in the A_{CNR} region belongs to the subset E_{CNR} : from the definition of E_{CNR} , even if the budget constraint in its lower portion reaches the A_{CR} region, it passes through indifference curves lower than the initial one: hence, picking the tangency point in the E_{CNR} region and implementing a CNR outcome is the optimal policy. On the contrary, if the tangency point is in A_{CNR} but not in E_{CNR} , the budget constraint reaches the A_{CR} at a higher indifference curve, and a CR outcome is optimal.

The iso-welfare curves with heterogeneous cartel types

In the following three Lemmas we identify the iso-welfare curves when the AA faces heterogeneous cartels.

Lemma 6 The iso-welfare curves in each of the τ -ve regions have the following pattern:

- 2 in E all the policy combination give the same welfare;
- 2 in A and D they replicate the shape of the $\mathbb{W}_{NC}(\mathbb{p})$ curves;
- 2 in B they are horizontal;
- 2 in C they are decreasing;

Proof: Since W_E does not depend on the policy parameters, all the region correspond to the same expected welfare. In region A all types choose CNR. From our analysis of the single type case we already know that the iso-welfare curves when no type colludes have a shape similar to the \mathbb{W}_{NC} curves (one for each type) and in the limit overlap with those. Hence, in region A the indifference curves replicate the \mathbb{W}_{NC} curves shape. In region D high types choose CNR and lower types choose NC: as long as we move along the \mathbb{W}_{NC} curve of type $i \in D$, the threshold type does not change and the first term in W_D is unchanged as well; moreover, we know that moving along a \mathbb{W}_{NC} curve

we keep the expected welfare for types choosing CNR constant. We conclude that the indifference curves in region D correspond to the θ_{NC} curves through it. In region B all types choose CR and the expected welfare depends only on θ , i.e. we have θ at iso-welfare curves. Finally, in the C region high types select CNR and low types CR: since W_C increases when θ is higher (more frequent revelation) as well as when p increases (more effective prosecution and more types induced to reveal), the iso-welfare curve in the C region must be decreasing. 2

Notice that when no Leniency Program is used, only the regions A, D and E exist, and the result above states that all the curves in A (or D) never pass through another region. Hence, the three relevant sets of indifference curves are completely defined. When $R < F$ all the five regions A-E exist; the Lemma above defines the iso-welfare curves in each region, but now the iso-welfare curves in A (D) eventually continue through region C and B. Hence, we have to carefully check how the iso-welfare curves behave moving from one region to the other.

Lemma 7 Consider the case $R < F$. The iso-welfare curves passing through the regions A-C-B are continuous and kinked at the boundaries of the A and B regions. The iso-welfare curves passing through the region D-C-B discontinuously shift to the right passing from D to C.

Proof: We start by identifying the indifference curves that pass through the A-C-B regions. We already know, borrowing from the analysis of the single type case, that the iso-welfare curve jumps down from A to B, when all types choose CNR and then CR: however, from the definitions of the expected welfare we can notice that W_C tends to W_A or W_B as the threshold type θ_M^C tends to θ_M or $\bar{\theta}_M$. Hence, the indifference curves are now continuous; it is easy to check also that they are kinked at the boundaries of the A and B regions, with the indifference curve steeper in C than in the other two regions³⁰.

Consider next the indifference curves passing through the D-C-B regions: we already established that in D the curves replicate the θ_{NC} curves shape, with some types choosing CNR and others NC; once moving into the C regions, some types still select CNR while others CR. Since the welfare associated to NC is higher than that when CR occurs, it must be that, moving from region D to region C along a iso-welfare curve, less types choose CNR. That requires a discontinuous jump to the right of the iso-welfare curve once entering in the C region. 2

Figure 6.a shows the two cases of indifference curves, one through A-C-B and the other through D-C-B. In the two Lemmas above we have completely characterized the iso-welfare curves when a Leniency Program is introduced and when it is not. Our next step is to verify when it is convenient to offer reduced fines in order to reach a certain expected welfare. This exercise corresponds to comparing the iso-welfare curves in the two cases, selecting in the different regions the lower one.

³⁰The concavity or convexity of the indifference curve in C cannot be stated in general, since it depends on the distribution of types $g(\theta_M)$. In what follows we consider the case of concave indifference curves, while the extension to the convex case is left to the reader.

Lemma 8 When the iso-welfare curve with $R = 0$ passes through the A-C-B regions, it is always optimal to use the leniency Programs. When the iso-welfare curve with $R = 0$ lies in the D-C-B region, the Leniency Programs are optimal only in a subset of the C and B regions.

Proof: Since $R = 0$ is the more effective way of inducing CR, we compare the case $R = F$ and $R = 0$, selecting the lower of the two iso-welfare curves. Since in A and D the iso-welfare curves are the same in both regimes, our problem amounts to selecting the lower curve in regions C and B. This can be done by distinguishing the case in which the indifference curve with τ ne reduction passes through the A-C-B region and that in which it lies in the D-C-B areas. Comparing the full and reduced τ nes indifference curves is immediate for the A-C-B case: in the A region they overlap while in the C and B region the iso-welfare curve is lower when $R = 0$, as shown in figure 6.a. Hence, the iso-welfare curve through A-C-B is that identified when the Leniency Program is used.

More complex is the comparison of the indifference curves with and without τ ne reductions when the former passes through the D-C-B region. In this case, in fact, the iso-welfare curve with τ ne reductions jumps to the right entering the C region, while with no Leniency Program the indifference curve, which is the same as before in the D area, goes on smoothly in the C region³¹. Hence, entering the C region from the top, the lower indifference curve is initially the one associated with no τ ne reduction. It may happen, as shown in figure 6.a, that continuing along it, the indifference curve with τ ne reductions becomes the lower one for a while. Finally, moving further to the right, the indifference curve with full τ nes lies again below that with τ ne reductions. For higher levels of the expected welfare, the indifference curve with full τ nes always dominates that with reduced τ nes. Hence, when we consider the indifference curves for increasing value of the expected welfare, as long as we are in the A-C-B region we use Leniency Programs, while entering the D-C-B region we adopt reduced τ nes only with a subset of policy parameters $(\tau; p)$, as shown in figure 6.a. 2

³¹Strictly speaking, with no Leniency Program no C region exists; hence we refer to the C region as those policy combinations defined in case of τ ne reductions.

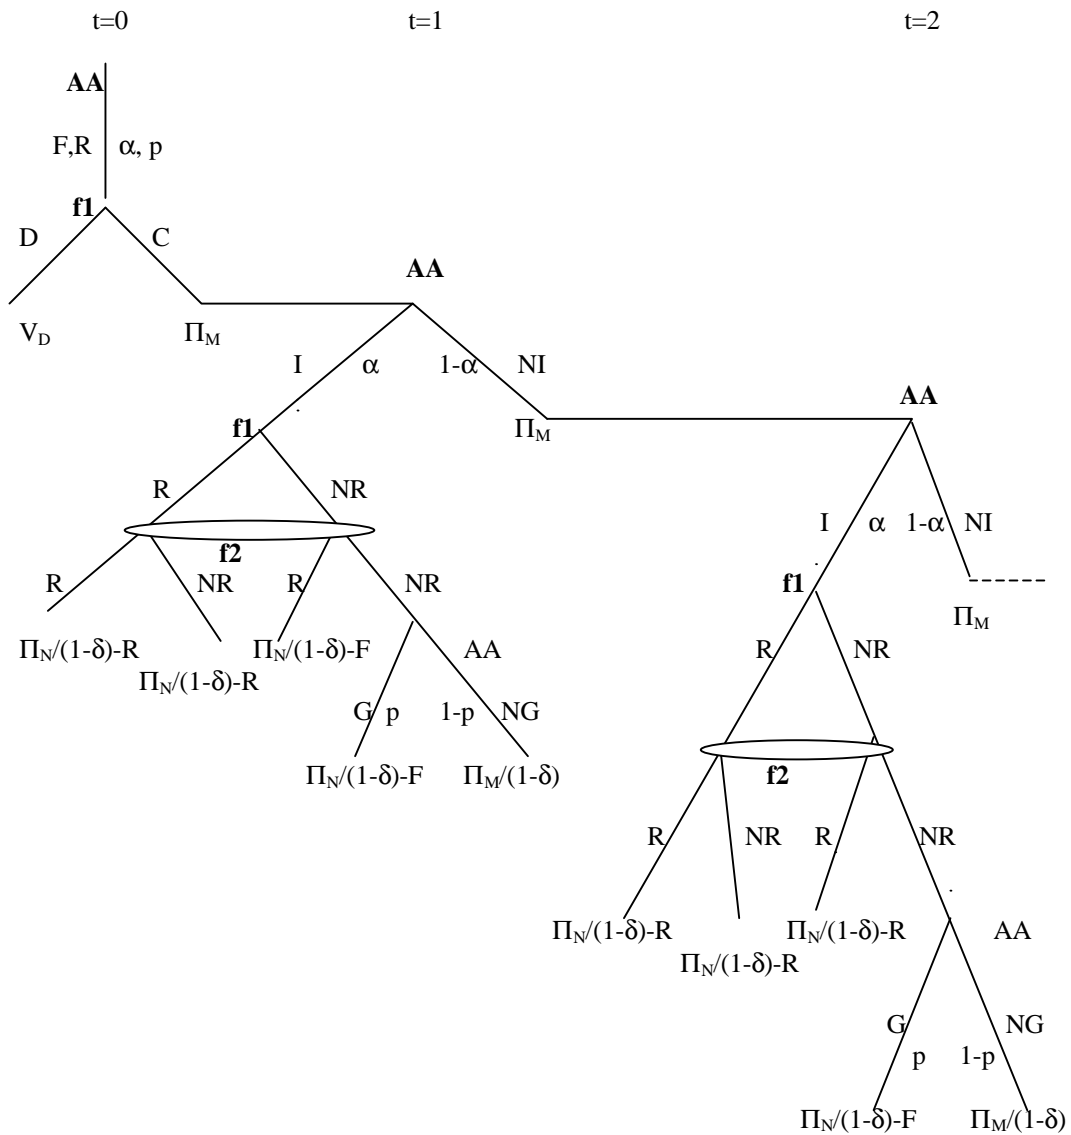


Figure 1: The game tree (D: deviation; C: collusion; I: investigation; NI: no investigation; R: reveal; NR: not reveal; G: guilty; NG: not guilty)

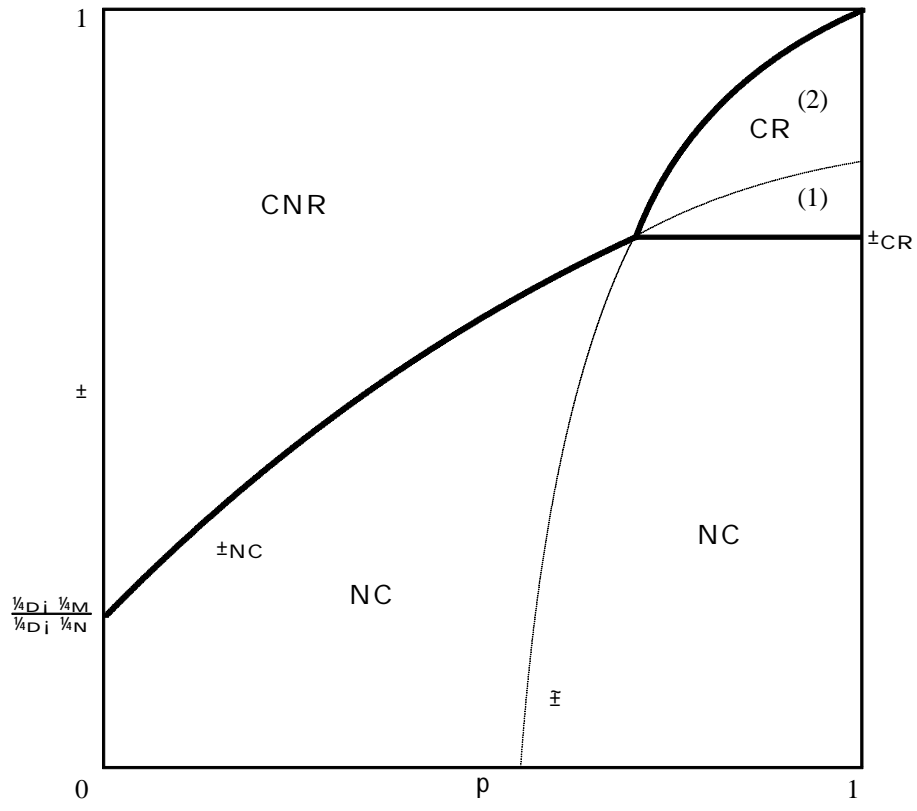


Figure 2.a - Implementable allocations




Figure 2.b - Implementable allocations

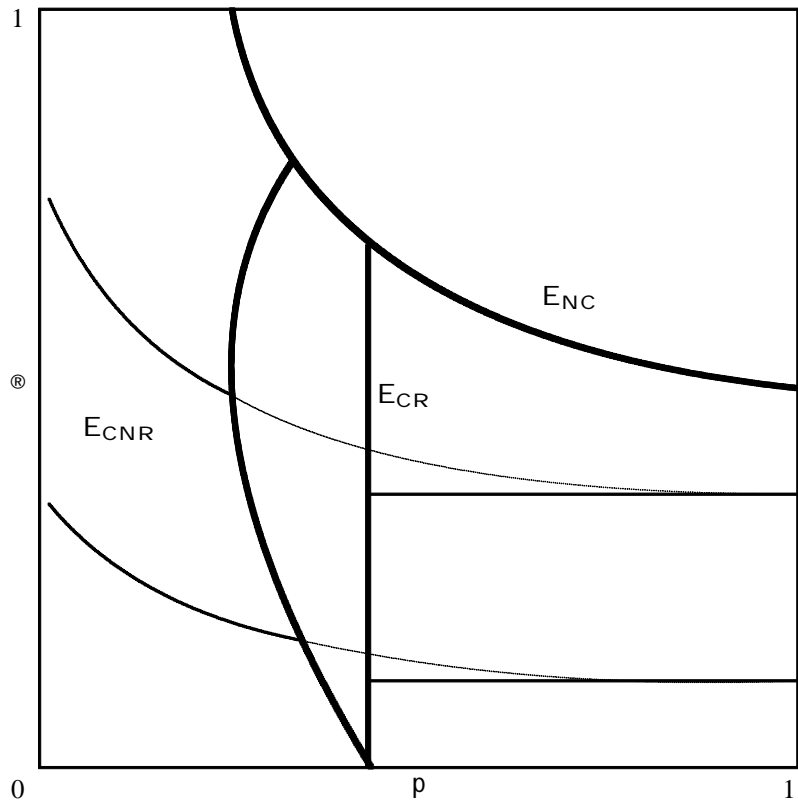


Figure 3 - Equilibrium policy combinations




Figure 4 - Alternative Leniency rules and implementable allocations




Figure 5 - Multiple types and implementable allocations




Figure 6.a - Iso-welfare curves with multiple types




Figure 6.b - Equilibrium policy combinations with multiple types