



Munich Personal RePEc Archive

Frequent Monitoring in Repeated Games under Brownian Uncertainty

António M. Osório-Costa

Universidad Carlos III de Madrid

4. January 2009

Online at <http://mpa.ub.uni-muenchen.de/13104/>

MPRA Paper No. 13104, posted 2. February 2009 00:02 UTC

Frequent Monitoring in Repeated Games under Brownian Uncertainty*

António M. Osório-Costa[†]

Universidad Carlos III de Madrid

(Dated: First Draft 28th September 2007 - This Draft 4th January 2009)

Abstract

This paper studies frequent monitoring in a simple infinitely repeated game with imperfect public information and discounting, where players observe the state of a continuous time Brownian process at moments in time of length Δ . It shows that efficient strongly symmetric perfect public equilibrium payoffs can be achieved with imperfect public monitoring when players monitor each other at the highest frequency, i.e. $\Delta \rightarrow 0$. The approach proposed places distinct initial conditions on the process, which depend on the unknown action profile simultaneously and privately decided by the players at the beginning of each period of the game. The strong decreasing effect on the expected immediate gains from deviation when the interval between actions shrinks, and the associated increase precision of the public signals, make the result possible in the limit. The existence of a positive monotonic relation between payoffs and monitoring intensity is also found.

JEL: C73, D82.

KEYWORDS: Repeated Games, Frequent Monitoring, Imperfect Public Monitoring, Brownian Motion, Moral Hazard.

*Financial support from Fundação para Ciência e Tecnologia grant SFRH/BD/17157/2004 is gratefully acknowledged. I would like to thank Luis Corchón, Ramiro De Elejalde, Daniel García-González, Ángel Hernando-Veciana, Kenneth Judd, David Levine, David Rahman, Juan Pablo Rincón-Zapatero, Andrzej Skrzypacz and Galina Zudenkova, as well as the seminar participants at University Carlos III Madrid, 2008 University of Konstanz Workshop in Game Theory, and The 2008 Meetings of the International Society in Dynamic Games, for helpful comments and discussions. All remaining errors are mine.

[†]Department of Economics, Universidad Carlos III de Madrid, C\ Madrid, 126, 28903-Getafe, Madrid, Spain; E-mail: aosorio@eco.uc3m.es; Cellphone (+34) 666 71 40 69.

I. INTRODUCTION

In the general repeated games theory it is common to assume that the period length between each repetition of the stage game has length one. When the monitoring is perfect, letting the discount factor $\delta \rightarrow 1$ either by making the players more patient (a decrease in r) or by shrinking the period length between actions (a decrease in Δ) are equivalent exercises. The former approach has been preferred to prove many folk theorems and to show the existence of efficient equilibria.

When monitoring is not perfect but public, by making players increasingly patient, Fudenberg, Levine and Maskin (1994) were able to prove a folk theorem under some informational assumptions¹. Incentives are sustained through transfers of value between the players. When we turn our attention to strongly symmetric equilibria in problems with two-sided imperfect public monitoring, the pairwise identifiability assumptions typically fail, limiting to a great extent the provision of incentives. Payoffs are then bounded away from efficiency, even when $r \rightarrow 0$. Destruction of value through punishments is the only way to provide incentives.

The question was then; if by letting $\Delta \rightarrow 0$ we could obtain efficient strongly symmetric payoffs. In their pioneer work in frequent monitoring, Abreu, Milgrom and Pearce (1991), in a setting where the public information arrival is modelled with a Poisson process, have shown that different but inefficient results arise depending on whether $\delta \rightarrow 1$ is due to $r \rightarrow 0$ or to $\Delta \rightarrow 0$. In the latter case, they found that equilibrium payoffs above the static Nash, but not efficient, can be sustained when the jumps in the process represent "bad news", which are more likely to occur when some player has deviated².

Recent work by Sannikov (2007)³ and Faingold and Sannikov (2007) on repeated games modeled directly in continuous time, has renewed the interest in frequent monitoring. The latter work, in a part which is relevant here, reports a degeneracy of the set of *strongly symmetric perfect public equilibrium* (SSPPE henceforth) payoffs in a game where a known

¹ Their stronger assumption is called pairwise identifiability, which means that a deviation from a given player impacts on the distribution of the public signals differently than any deviation from any other player.

² In an infinitesimal time interval, the absence of realizations of the Poisson process is infinitely more likely than the occurrence of a realization. For that reason, the same result does not extend when the information arrivals represent "good news".

³ This paper presents a characterization of the set of PPE payoffs using continuous time methods.

normal type long-run player faces a sequence of short run players. By degeneracy they mean that payoffs outside the convex hull of the Nash equilibria payoffs set cannot be sustained in continuous time, where the noisy public information is modeled through a Brownian process. Also, more in the spirit of the present paper, i.e. by studying the limit of the discrete time games, Fudenberg and Levine (2007) and Sannikov and Skrzypacz (2007a) report the same degeneracy result. These results came as surprising, since Brownian motion is an infinitesimal variation process and we would expect payoffs at least above the static Nash payoffs.

This paper explores frequent monitoring in a partnership game with imperfect public monitoring and discounting, along the lines of the work of Radner, Myerson and Maskin (1986)⁴. It analyses the limit of the sequence of the discrete time games indexed by Δ . The public signal is the observed state of an arithmetic Brownian motion (ABM henceforth) process, in intervals of length Δ . The observation is compared against a previously chosen threshold, based on this decision rule; players adjust their actions for the following period.

It focus on the value of the best SSPPE and shows that in the limit⁵ efficient payoffs can be achieved, independently on how players discount the future and on the level of uncertainty. Moreover, the value of the best SSPPE payoff improves monotonically when the monitoring intensity increases. The characterization of the optimal decision rule for different values of Δ is also novel.

The approach proposed places distinct initial conditions on the process, which depend on the unknown action profile privately and simultaneously decided by the players at the beginning of each period of the game. This modelling approach changes the results drastically.

This paper is the first to show efficient results when the time interval between observations is taken to the limit without placing assumptions on the uncertainty parameter. The result is possible because the information extracted from the public signals becomes increasingly

⁴ Other classical situations involving imperfect public monitoring about players actions are Green and Porter (1984) and Porter (1983) where the market price is an imperfect signal of the quantities supplied by firms. See Fudenberg and Tirole (1991) and Mailath and Samuelson (2006) for complete surveys of the problems and methods used to solve games with imperfect public information.

⁵ By "in the limit" we mean the length of time interval $\Delta \rightarrow 0$, sometimes also referred to as the "highest monitoring intensity" or "continuous monitoring". During the paper we frequently mention "an increase in the monitoring intensity" or "an increase in the monitoring frequency"; they refer to a decrease in Δ .

precise as to players actions, increasing the payoffs monotonically⁶. As a consequence, the expected immediate gains associated with a deviation from the equilibrium path become less attractive, not only because the periods between the actions becomes shorter, but also because the expected number of periods during which a deviator can enjoy these gains decreases.

Before going through the methodology presented in this paper in more detail, and in order to better integrate it within the existent contributions, we will review the modelling approach employed in the papers that are more closely related, with the goal to access the source of their results.

Sannikov and Skrzypacz (2007a)⁷ studied extensively how monitoring intensity affects the equilibrium payoffs of the repeated Cournot duopoly game. They report the impossibility of achieving payoffs higher than the static Nash payoffs when the public information arrives continuously and disturbed by a Brownian motion. Crucial for their results is the assumption that the public signal observed by the players, at moments in time $t = \Delta, 2\Delta, \dots$, is the state of an ABM price process divided by the length of the time interval Δ . Such modelling of the observed public signal becomes extremely noisy when observed at high frequency, creating a degeneracy effect on the payoffs. The root of the problem lies in the fact that the accumulated Brownian increments in a given time interval Δ are of a higher order than the underlying time interval, making any inference about the drift of the process inefficient in the limit, but not exclusively so.

Fudenberg and Levine (2007)⁸ study a repeated game with a finite action space where a long-run player of a known type faces a sequence of myopic short-run players. The public

⁶ These results share similarities with the ones obtained in Abreu, Milgrom and Pearce (1991). For the "bad news" case, the most efficient equilibrium (however not fully efficient) is obtained in the limit, and there is also a monotonic improvement on the payoff with the monitoring intensity. Here the most efficient result is also obtained in the limit; it is however fully efficient. In the "good news" case, the degeneracy of the best SSPPE is similar in shape to the results obtained by Fudenberg and Levine (2007) and Sannikov and Skrzypacz (2007a), although clearly distinct in the mechanics behind it. This issue will be discussed in more detail below.

⁷ See also, Sannikov and Skrzypacz (2007b) where they bound the set of equilibria when $\Delta \rightarrow 0$ for Levy processes by placing restrictions on how information from Brownian and Poisson components are used to provide incentives in the most efficient way.

⁸ See also, Fudenberg and Levine (2008) where they consider different ways of passing to the continuous time limit, i.e. binomial and trinomial approximations of the Brownian paths, linking monitoring intensity with event frequency, by taking the limit of the former.

signal is the observed state of an ABM process, which can be influenced by the actions of the long-run player. The long-run player would like to sustain the equilibrium associated with a particular action, in the case where the drift of the process decreases with Δ , but she may also deviate to get a larger expected short run payoff, in which case the drift of the process increase with Δ . The results are driven by the assumption that the initial conditions on the process are the same independently of the actions of the long-run player. When Δ becomes small the distribution of the public signals cannot provide reliable information about the long-run player actions, creating the degenerating effect, which may also occur even before the limit. However, they show that if a deviation by the long-run player increases the uncertainty parameter; equilibria can be achieved that are arbitrary close to efficiency. This is so far the most positive limit result obtained with frequent monitoring under Brownian uncertainty.

Before proceeding, it is important to understand what information we can obtain from statistical inference for Brownian type processes⁹. While the variance parameter can be consistently estimated from the path of the process ($y_s, s \in (t, t + \varepsilon]$), for a small but measurable ε , the same does not happen with the drift of the process. Even if players are able to observe the full path of the process realized from time t to time $t + \Delta$, a relatively large Δ is needed for the actions of the players to be statistically distinguishable. The first observation is on the basis of Fudenberg and Levine's (2007) almost efficient result. The second explains why the set of equilibria payoffs deteriorates when Δ decreases and equilibria above the static Nash are not possible as reported in Fudenberg and Levine (2007), and Sannikov and Skrzypacz (2007a).

The present paper models the public signal observed by the players at moments in time $t = \Delta, 2\Delta, \dots$, in a different way. Since players focus on the observed public signal, the monitoring technology is in that sense similar to Fudenberg and Levine (2007). However, when each player privately selects her action, the initial condition on the process will reflect the aggregate of these individual decisions. Two different actions of player i have associated different initial conditions that can be statistically distinguished from each other for high monitoring frequencies. There is a measurable distance between an initial condition associated with mutual cooperation and an initial condition associated with a profile of actions

⁹ See Prakasa Rao (1999) for a formal treatment of the statistical methods for diffusion processes.

where a unilateral deviation has occurred. When switching from cooperation to defection, player i causes a movement in the process similar to a jump. When away from the limit, such jump might be hard to separate from the aggregate of infinitesimal realizations, but in the limit such movement is almost surely caused by deviating behaviour. This is even more evident when the action space is discrete.

Even though they present different results for the same problem, it is important to stress that the modelling proposed in this paper, the Fudenberg and Levine (2007) and Sannikov and Skrzypacz (2007a) are not competing approaches, but rather they complement each other. The degeneracy effect reported in the latter approaches seems suitable to model situations where reliable information needs time to build-in, in such cases an early access to this "not ready" information generates perverse effects. These papers have in common the fact that the informativeness of the public signals increases with the lack of monitoring. A player that is actually cooperating needs some time for the associated output to be separated in statistical terms from the situation from which she has deviated. Examples of situations that fit into such a framework are a judge's verdict or an investment report. Asking a judge to read her sentence five minutes after the start of judgement might be premature. Probably one week later she will be in better position to deliver a more qualified decision. The same happens with the investment report.

However, in problems where the information contained in the public signals increases with the monitoring intensity, the approach of the present paper is more appropriate; for example when the variable that is the object of monitoring is the physical activity of a worker, for which information is continuously available. When monitored with high frequency a deviating worker will find it hard to justify a low observation of the process. If such a low value of the process is observed, it is very likely that it has been caused by deviating behaviour. On the other hand, when the time interval between monitoring activities is large, a low observation can be justified both as a result of nature as well as a result of deviating behaviour. The monitor is then more likely to incur in either type I or type II errors¹⁰. In problems of this type frequent monitoring may produce a disciplinary effect among the players¹¹, but everything depends on what variable players are monitoring. In general terms,

¹⁰ We refer to a type I error as the event of punishing a non-deviator, and a type II error as the event of not punishing a deviator.

¹¹ The disciplinary effect of monitoring seems to be easier to support. See, for example, the classical work

the monitoring approach presented in this paper fits in situations where effort is modelled as an action that has a more physical meaning as opposed to an intellectual one attached to it.

In this sense, the present paper fills a gap in the existing literature in frequent monitoring in repeated games, by enlarging the spectrum of potential real life situations that can be studied with this framework.

The paper is organized as follows. Section II presents the repeated game model and the public information producing process. Section III explains in detail the approach of this paper, in particular the connection between the initial conditions of the process and the associated distributions. Section IV computes the bounds on the set of SSPPE payoffs and characterizes the optimal decision rule for varying Δ . Section V focuses on the limit case and presents the main results of this paper. Section VI discusses extensions to the continuous time case and to games with a continuous action space. Section VII concludes.

II. THE REPEATED GAME MODEL

To present the main arguments outlined in the previous section, we explore frequent monitoring in a simple partnership game played between two long-run players. The history of the game is the following. At moments in time $t = 0, \Delta, 2\Delta, \dots$, two players can choose from two different effort levels $a_i^t = 1$ or $a_i^t = 0$. In the former case, player i is providing effort, in the latter case she is shirking. More formally, let $A_i^t = \{0, 1\}$ denote player i 's $N = \{1, 2\}$ non-empty and compact action space with generic element a_i^t representing an action, and denote $A^t = A_1^t \times A_2^t$ as the set of action profiles endowed with the product topology of the individual action spaces, with generic element $a^t = (a_1^t, a_2^t)$ denoting a profile of actions.

Independently of their private effort decisions, players at moments in time $t = \Delta, 2\Delta, \dots$, observe and divide equally between themselves the total output generated during these intervals of time of length Δ . Given a profile of actions a^t at time t , the total output and, at same time, the publicly information observed at time $t + \Delta$ ¹², is driven by the following

of Alchian and Demsetz (1972) for an early defense of this perspective.

¹² There is also the possibility that in the end of each period of length Δ , players observe the full path of the process $(y_s, s \in (t, t + \Delta])$ realized from t to $t + \Delta$. This case provides more information to the monitor.

ABM¹³,

$$y_{t+\Delta} = y_t + \sigma \int_t^{t+\Delta} dZ_s, \text{ with } Z_t = 0 \text{ and } t = 0, \Delta, 2\Delta, \dots, \quad (2.1)$$

where $y_t \equiv 2d(a_1^t + a_2^t)$ is the initial condition of the process at a given time t , a function of the unknown profile of actions.

All the uncertainty arises from the standard Brownian motion $\{Z_s; s \geq 0\}$. Observe that information about the evolution of the total output is produced continuously every infinitesimal time interval a new realization of the process is available. The monitoring frequency Δ and the frequency of signals are different; in some sense our goal is to equate the monitoring frequency to the signal frequency.

All the relevant information relevant about players' actions is contained in the initial condition of the process at each moment in time $t = 0, \Delta, 2\Delta, \dots$. Since players cannot revise their actions during the interval of length Δ , the process (2.1) has no drift. In this way, we also eliminate any trend in the process that is irrelevant for the study of the problem in our setting¹⁴. Implicitly, we are imposing that the public information producing process satisfies the martingale property,

$$E(y_{t+\Delta} | y_t, \Delta \geq 0) = y_t,$$

with respect to some filtration. As a consequence, the transition density of the process places equal mass above and below the initial condition, i.e. above and below its mean¹⁵.

Player $i \in \{1, 2\}$ realized payoff (ex-post) from the partnership is given by,

$$r_i(a_i^t, y_{t+\Delta}) \equiv y_{t+\Delta}/2 - (d+c)a_i^t,$$

where $d > c > 0$ and $y_{t+\Delta}$ is the realized output from the partnership¹⁶, which is divided

It can be shown, that when compared with the case where only the state of the process is observed, a lower threshold is always required to sustain a particular equilibrium and as a consequence it always has larger associated payoffs. Note, however, that in the limit both cases are equivalent.

¹³ The ABM assumption simplifies the analyses and the proofs of many of the results which will be presented later. However, all results are valid for the geometric Brownian motion and Ornstein-Uhlenbeck process.

¹⁴ Using a public information producing process with drift leads to the same results. In that case, the threshold would be some function preserving the distance to the initial condition. For example, for the ABM process with drift, the threshold would be the function $b + \mu(y_t, t)t$, where b is the threshold associated with an ABM process without drift and $\mu(y_t, t)t$ is the drift of the process. Both approaches lead to the same distribution of the public signals.

¹⁵ Examples of other processes with the proposed specifications are $\mu(y_t, t) = 0$ for the geometric Brownian motion, and $\mu(y_t, t) = \rho(y_0 - y_t)$ for the Ornstein-Uhlenbeck process.

¹⁶ We assume that public signal is not only an action dependent function but also represents the evolution

equally between the players independently of the effort that they have supplied. The second term on the right-hand side (RHS) denotes the cost of providing effort for player i in utility terms. By the martingale property, the expected payoff (ex-ante) of player $i \in \{1, 2\}$ from the partnership is then,

$$\pi_i(a^t) \equiv E(r_i(a_i^t, y_{t+\Delta}) | y_t) = d(a_1^t + a_2^t) - (d + c)a_i^t.$$

Under the expected utility hypothesis, this is the expression that is relevant for studying our problem.

In summary, at each moment in time $t = 0, \Delta, 2\Delta, \dots$, players repeatedly play the stage game,

	1	0
1	$d - c, d - c$	$-c, d$
0	$d, -c$	$0, 0$

Providing no effort is a dominate strategy for both players. The minimax value of the game coincides with the stage game Nash payoffs and equals 0 for both players. This game has the same structure as a prisoners' dilemma, and can be treated as such.

I will focus on strongly symmetric public strategies, where after every public history¹⁷ the same action is chosen by both players. A strategy is public if at any moment t it depends only on the public histories and not on player i 's private history. Given a public history, a profile of public strategies that induces a Nash equilibrium on the continuation game from time t on, is called a *perfect public equilibrium* (PPE). Moreover, if the other player $-i$ is playing a public strategy, player i 's best reply can only be a public strategy.

Players discount the futures according to a common discount factor, assuming exponential discounting $\delta \equiv e^{-r\Delta}$, where r is the discount rate.

Player i 's $\in \{1, 2\}$ payoffs in the Δ -indexed infinitely repeated game is the discounted normalized sum of the stage-game expected payoffs, $(1 - \delta) \sum_{t=0, \Delta, 2\Delta, \dots}^{\infty} \delta^t \pi_i(a^t)$, induced by the profile of actions a^t for every $t = 0, \Delta, 2\Delta, \dots$

of the aggregate output of the partnership. Other formulations of the public signal could have also been considered, provided that they would depend on both players' actions. It is also important that $r_i(\cdot)$ does not depend on a_{-i} explicitly.

¹⁷ A public history, a time t , is a sequence of realizations of the observed state of the process, denoted by $h^t = (y_{t-\Delta}, y_{t-2\Delta}, \dots, y_0) \in Y^t$ with $h^0 = Y^0 \equiv \emptyset$. The sequence of player i 's private actions is player i 's private history.

III. THE INITIAL CONDITIONS AND THE DISTRIBUTION OF PUBLIC SIGNALS

As discussed in the introductory section, the proposed approach places distinct initial conditions on the process, which depend on the unknown action profile simultaneously and privately decided by the players in the beginning of each period of the game. Since this is a critical issue, in this section we examine in more detail the monitoring technology employed in this paper.

To keep the notation simple, from now on we will drop the t index, and simply denote with a Δ index an end of period object, and without any index when it refers to the beginning of the period.

Since the focus is on SSPPE, we are interested in sustaining the strongly symmetric profile $(1, 1)$ against potential unilateral deviations. The space of public signals is continuous; players use a threshold decision rule¹⁸ to distinguish realizations suggesting cooperation from realizations suggesting defection. The threshold value creates a partition of the signal space; in signals suggesting cooperative behaviour $\{y_\Delta > b\}$, which we call "good" signals, and signals suggesting defective behaviour $\{y_\Delta \leq b\}$, called "bad" signals.

In general, for a given initial condition $y \equiv 2d(a_1 + a_2)$, the probability that the state of the public process (2.1) appears below b in the end of the period of length Δ is

$$\Pr(y_\Delta \leq b) = \Phi\left(\frac{b - 2d(a_1 + a_2)}{\sigma\sqrt{\Delta}}\right),$$

where $\Phi(\cdot)$ is the standard zero mean and unit variance Gaussian distribution.

We assume that the type of uncertainty they are facing is common knowledge among the players. In particular, they know the value of parameter σ ¹⁹. This way they can compute the above probability and the impact of a deviation on the distribution of the public signals.

Depending on the unknown profile of actions that arise from each player's private effort

¹⁸ Sannikov and Skrzypacz (2007a) show that a threshold decision rule is the best test to detect unilateral deviations. Later, we will focus on the optimal threshold value for varying parameterizations of Δ . For now we contend with an arbitrary threshold value, denoted as b , and we will abstain from referring its dependence on Δ as well as other parameters of the model.

¹⁹ The assumption that players know the variance of the process is not as strong as it might seem. According to the discussion in the introductory section, since the focus is on the limit case and this is equivalent to full path observation, if players do not know the value of this parameter they can estimate it consistently in a very small but measurable time interval ε .

decision $a_{i \in N}$, we have different initial conditions for the public process. In the partnership game there are four possible profiles of actions: the strongly symmetric profile $a \equiv (1, 1)$, the unilateral deviation profiles $a' \equiv (0, 1) \equiv (1, 0)$ ²⁰ and the Nash or punishment profile $a^N \equiv (0, 0)$. The latter profile is a Nash equilibrium, and for that reason it is trivially self-enforcing. We will focus on the first two possibilities:

i) The strongly symmetric profile a has a initial condition $4d$ and a probability of punishment denoted as

$$F(b, \Delta) \equiv \Phi \left(\frac{b - 4d}{\sigma \sqrt{\Delta}} \right).$$

This is type I error probability; even though no player has deviated, punishment will be exerted when a low realization of the process is observed.

ii) The unilateral deviation profile a' has an initial condition $2d$ and a probability of punishment denoted as

$$F'(b, \Delta) \equiv \Phi \left(\frac{b - 2d}{\sigma \sqrt{\Delta}} \right).$$

Analogously a type II error is given by $1 - F'(b, \Delta)$, see footnote 10 above.

Observe that the actions decided independently at the beginning of the period by each of the players are clearly printed in the initial conditions of the process, as they should be. At the end of the period of length Δ , the profile of actions chosen by the players appears disturbed by some noise. Lowering Δ , we are decreasing the uncountable number of infinitesimal contributions to the noise, leading to a more likely observation of the process around its initial condition. An observation of the process far away from the mutual effort initial condition is then a sign of a deviation.

If we want to generalize this methodology to other more complex environments, this correspondence between actions taken and the associated expected public signal observed in the end of the period always have to be present. More elaborated information structures such as the pairwise full-rank condition of Fudenberg, Levine and Maskin (1994) are not required.

Active Coordination

²⁰ Fudenberg, Levine and Maskin (1994) pairwise full-rank condition typically fails for strongly symmetric equilibria. There is no loss in generality when placing no distinction between the profiles $(1, 0)$ and $(0, 1)$, and denoting them as a' . Also note that to keep the notation standard, until now, a has denoted a general action profile. With a slight abuse of notation, a now denotes the strongly symmetric effort profile.

When the game is a repeated partnership, it is natural to assume that the process is reset at the end of each period, after players have observed and split the realized output, starting again in the point associated with the new action profile decided simultaneously by both players. However, the same assumption does not fit in some other contexts. For example in the repeated Cournot game, it is not reasonable to assume that players can reset the process every period, since it represents the market price. In this case, the new equilibrium actions have to be adjusted to take into account the observed state of the process²¹. Although different, these two possibilities can be studied within the same methodology. To illustrate this point consider the following example.

In the Cournot duopoly game the public process observed at time Δ is the market price, $P_\Delta = P + \sigma \int_0^\Delta dZ_t$, where $P = (\alpha - q_1 - q_2)$. Suppose that at time 0 the game starts with both players supplying the monopoly quantities, i.e. $q_i = \alpha/4$ ²². The expected market price (and initial condition) is then $P = \alpha/2$ and players select some threshold value $b < \alpha/2$. In period Δ , let the realization of the process be $P_\Delta = P + \sigma z_\Delta$. If $P_\Delta \leq b$, players play their punishment actions, with respect to the observed state of the process. For simplicity, consider the Nash punishment; where in the following period each player will supply $q_{i\Delta} = (\alpha + \sigma z_\Delta)/3$.

On the other hand, if $P_\Delta > b$, the next period 50/50 monopoly quantities will be chosen also taking into account the state of the process, hence, requiring each player to supply the quantities $q_{i\Delta} = (\alpha + \sigma z_\Delta)/4$ with a necessary adjustment in the decision rule, i.e. $b_\Delta = b + \sigma z_\Delta$.

We will call this type of coordinated behaviour between players in a dynamic setting *Active Coordination*, since it requires players to adjust their equilibrium path actions taking into account the observed state of the process. Players do not reset the value of the process, but rather they reset the uncertainty.

²¹ The martingale condition guarantees that such a procedure does not change the properties of the distribution of the public signals.

²² In the Cournot duopoly game under imperfect public information, it is never efficient to produce exactly the monopoly quantities, but rather an amount slightly larger (except in the limit). This issue and others are discussed in more detail in section VI B. Take this example as an illustration.

IV. THE BEST STRONGLY SYMMETRIC EQUILIBRIA

This section furnishes the reader with a set of general results that are independent of the monitoring intensity, which will be particularly useful for the following section when we focus on the limit case. It also presents a characterization of the optimal decision rule associated with the value of the best SSPPE of the infinitely repeated partnership game. In some occasions, the results presented are standard in repeated games; a brief description will be presented, sufficient to keep the exposition self-contained.

Since the game is symmetric there is no loss in generality when studying a single player incentives, hence we remove players indexes.

Recall from the previous section that the equilibrium profile we want to sustain is $a \equiv (1, 1)$. A profile where a single player deviates is denoted as a' , and the Nash profile denoted as a^N . These profiles have associated with them the stage game payoffs, $\pi \equiv \pi(a) = d - c$, $\pi' \equiv \pi(a') = d$ and $\pi^N \equiv \pi(a^N) = 0$, respectively. See the expected payoffs table in section II. The public information is produced continuously and is generated by the simple ABM process given by (2.1), with initial conditions depending on the unknown action profile.

Since the public information process can virtually return any value in \mathbb{R} , we can use the Abreu, Pearce and Staccetti (1986, 1990) bang-bang result to compute the best SSPPE payoff.

Lemma 1 (Sannikov and Skrzypacz (2007a Lemma 3)) *In a strongly symmetric equilibria each player solves the problem,*

$$\max_{v(y) \in [\underline{v}, \bar{v}]} (1 - \delta) \pi + \delta \int_{-\infty}^{\infty} v(y) f(y) dy$$

s.t. for all $y \in \mathbb{R}$,

$$(1 - \delta) (\pi' - \pi) \leq \delta \int_{-\infty}^{\infty} v(y) (f(y) - f'(y)) dy,$$

where $v(\cdot) \equiv v$ denotes the continuation value of the game, and $f(\cdot)$ and $f'(\cdot)$ are the Gaussian densities associated with the public process (2.1) for the initial conditions $4d$ and $2d$ respectively and common variance $\sigma^2 \Delta$. When a solution to this problem exists, it takes the form

$$v(y) = \begin{cases} \bar{v} & \text{if } y > b \\ \underline{v} & \text{if } y \leq b \end{cases}$$

for $b \in (-\infty, 3d]$.

Here \bar{v} and \underline{v} will denote respectively the upper and the lower bounds on the set of SSPPE payoffs²³. The Lemma tell us that the continuation value v of the game is some non-trivial combination between the expected value \bar{v} when play starts with the observation of a good signal and the expected value \underline{v} when play starts with an observation of a bad signal. Such continuation value must enforce the profile a .

Here punishments are not executed in terms of a deterministic number of time periods as suggested by Porter (1983) and Green and Porter (1984) but rather in a probabilistic sense as in Abreu, Pearce and Staccetti. In our partnership problem, as we shall see, since the distribution of the public signals is not convex, optimality requires an infinite punishment length and both approaches do not differ. In some other problems, in particular when the minmax value differs from the Nash value, the differences might be substantial.

Lemma 1 allows us to rewrite players' problem in a more tractable way, and to use simple dynamic programming methods to search for expressions for the values of \bar{v} and \underline{v} that are exclusively represented as functions of the parameters of the model. The value of the best SSPPE is then written,

$$\bar{v} = (1 - \delta) \pi + \delta [(1 - F(b, \Delta)) \bar{v} + F(b, \Delta) \underline{v}]. \quad (4.1)$$

It can be interpreted as follows. While both players provide effort, each receives the immediate discounted normalized expected payoff associated with mutual effort, as well as a discounted expectation over the expected values associated with the two potential signals that might be observed.

The value in expression (4.1) has to satisfy the usual set of constraints,

$$\begin{aligned} \bar{v} &\geq (1 - \delta) \pi' + \delta [(1 - F'(b, \Delta)) \bar{v} + F'(b, \Delta) \underline{v}], \\ \bar{v} &\geq \underline{v} \geq \pi^N, \\ b &\in (-\infty, 3d]. \end{aligned}$$

The first constraint is the enforceability condition of Lemma 1; we will frequently use the terminology incentive compatibility (IC henceforth) instead. It has a simple interpretation:

²³ Since later we will allow players to correlate their actions on some public signal we are calling \bar{v} and \underline{v} extreme points of the set of SSPPE payoffs. This set is the collection of all payoffs that can be achieved with strongly symmetric public strategies, i.e. when both players choose the same action after every public history. See also the discussion after Proposition 2.

the expected value of the game associated with mutual effort has to be at least as good as the expected value of the game associated with a potential unilateral deviation, even if this deviation just last one period. The second constraint requires both \underline{v} and \bar{v} to be feasible and individually rational. The third condition is technical and comes from Lemma 1. It restricts b to the interval where $\partial F'(b, \Delta) / \partial b \geq \partial F(b, \Delta) / \partial b$, i.e. the likelihood ratio is larger than or equal to one. An alert of deviation is more likely to occur when somebody has in fact deviated. The likelihood difference $F'(b, \Delta) - F(b, \Delta)$ is non-negative and is an increasing function of b in the interval $(-\infty, 3d]$. In Appendix A we show that this condition does not impose any particular restriction since it is always satisfied.

When any of these conditions fails, no equilibria either the infinite repetition of the static Nash equilibrium can be sustained, i.e. $\bar{v} = \pi^N$. When this turns out to be the case, in particular in the limit, we say that the set of SSPPE payoffs degenerates²⁴.

We can solve the system composed by expression (4.1) and the IC constraint (assuming it holds with equality), for \bar{v} and \underline{v} , to obtain the solutions

$$\bar{v} = \pi - \frac{F(b, \Delta)}{F'(b, \Delta) - F(b, \Delta)} (\pi' - \pi), \quad (4.2)$$

and

$$\underline{v} = \bar{v} - \frac{(1 - \delta)}{\delta} \frac{\pi' - \pi}{F'(b, \Delta) - F(b, \Delta)}. \quad (4.3)$$

Expression (4.2) and (4.3) characterize the value of the upper and lower bounds on a set of SSPPE payoffs respectively.

Appendix A shows that the constraint $b \in (-\infty, 3d]$ is always satisfied. Additionally as a necessary condition for optimality, we have imposed that the IC condition would have to hold with equality. Now we are just left with the feasibility constraint to be concerned about. The following result establishes conditions on the threshold b in order for the value \bar{v} to be optimal among all the SSPPE values that are feasible.

Proposition 2 *In pure strategies and with public correlation, under (2.1) and given Δ , any punishment strategy with associated expected value $\underline{v} \in [\pi^N, \bar{v}]$ is feasible. Moreover, the strategy profile that achieves the largest upper bound \bar{v}^* on the set of SSPPE payoffs*

²⁴ The statement is not exactly correct. It might happen that condition $b \in (-\infty, 3d]$ is not satisfied and we are still able to achieve payoffs that satisfy all the other conditions. However in such a case the value b has not been chosen efficiently. The concept of efficient threshold is defined in Appendix A.

requires perpetual punishment the first time the process is observed below $b^*(\Delta)$. Where $b^*(\Delta) \equiv b^*$ is called the optimal threshold and is the solution to $\underline{v} = \pi^N$, i.e. b^* solves

$$F'(b, \Delta) (\pi - \pi^N) - F(b, \Delta) (\pi' - \pi^N) - (1 - \delta) (\pi' - \pi) / \delta = 0. \quad (4.4)$$

Proof. See Appendix B. ■

The result tells us that among all the feasible punishment schemes, perpetual punishment is the optimal one. Expression (4.4) gives an implicit function to compute b^* .

For a given Δ , when a solution b^* to (4.4) exists, the value in expression (4.2) returns the largest SSPPE payoff \bar{v}^* . The optimal threshold b^* is the lowest value of b that is feasible and IC. It establishes the right balance between gains and losses associated with right and wrong inference about players' actions.

Independently of Δ , the largest strongly symmetric payoff is attained using the most severe punishment available²⁵, which happens because the ABM is a Gaussian process and the distribution of the public signals is not convex in all of its domain²⁶.

Notice that the two point set $\{\pi^N, \bar{v}^*\}$ associated with (4.4) in Proposition 2 is self-generating using only pure strategies, since the continuation values π^N and \bar{v}^* are elements of the set. However we have added a public correlated signal and we have called the set of SSPPE payoffs $[\pi^N, \bar{v}^*]$. When we consider the set $[\underline{v}, \bar{v}]$ with $\underline{v} > \pi^N$, a public correlated signal is required in order for the set to be self-generating. The reason is that the continuations associated with the value of the game that starts with cooperation, and the value of the game that starts with punishment must be elements in the interval $[\underline{v}, \bar{v}]$. The important point here is that such a generalization allows us to call the set of SSPPE an interval even when $\underline{v} = \pi^N$ and also because the result in Proposition 5 of Section V will be shown to hold not only under optimal behaviour but also for more general sets $[\underline{v}, \bar{v}] \subseteq [\pi^N, \bar{v}^*]$.

However, not explicitly mentioned in order to keep the notation simple, the solution b^* depends on all the parameters in the model, i.e. d , c , r and σ . Note also that since the values of $F(\cdot)$ and $F'(\cdot)$ are endogenously determined by b^* , expression (4.2) necessarily depends on players impatience level.

²⁵ Nonetheless, we will keep working with expression (4.2), which is more general and can harbour the perpetual punishment case as a particular case.

²⁶ For a discussion of this issue in the context of repeated games, see Porter (1983).

By proposition 2, if a solution to (4.4) exists, it is an optimal threshold b^* . The following result establishes that such a solution, indeed exists, and is unique and differentiable on some interval $(0, \bar{\Delta})$. The value $\bar{\Delta}$ depends on all the parameters of the model, and denotes the maximum length of time between observations of the process in the infinitely repeated partnership that can support some nontrivial equilibria.

Proposition 3 *There is a $\bar{\Delta} > 0$ such that for all $\Delta \in (0, \bar{\Delta})$, a solution $b^*(\Delta)$ to (4.4) exists, is unique and differentiable.*

Proof. See Appendix B. ■

Outside the interval $(0, \bar{\Delta})$ there is no value of b^* that satisfies (4.4), and shirk becomes a dominate strategy for both players.

The Optimal Threshold - Numerical

Proposition 2 shows that the largest value of \bar{v} is obtained using the most severe punishment. Such punishment has associated with it an optimal decision rule whose existence is guaranteed by Proposition (3). Now we will attempt to provide the intuition behind some properties of the resulting decision rule.

Figure 1 shows the optimal threshold value as a function of Δ , for the partnership game, and for different parameterizations of σ and r , when $d = 3$ and $c = 1$.

It is clear from the figure that independently of the values that σ and r can take, in the limit as $\Delta \rightarrow 0$, the value of b^* converges to $2d$, the expected signal that would arise if there was a deviation. Such a result is formally shown in Proposition 6 in Section V²⁷.

Notice that the value $\bar{\Delta}$ increases either when players get more patient or the public signal becomes less noisy, see Figure 1.

The more impatient the players are (larger r) the tighter the monitoring²⁸ must be in order to create incentives for cooperative behaviour. This can be seen in Figure 1 where for the same parameterizations as above, we increase r from 0.1 to 0.2.

²⁷ It can also be shown that $\partial b^*/\partial \Delta \rightarrow -\infty$ when $\Delta \rightarrow 0$ and $\partial b^*/\partial \Delta \rightarrow \infty$ when $\Delta \rightarrow \bar{\Delta}^-$, however such results are not particularly relevant.

²⁸ By "tighter monitoring" as opposed to "relaxed monitoring", I mean a higher threshold value. The larger the threshold the more likely players are to detect deviations, but they are also more likely to commit type I errors.




FIG. 1: The optimal threshold as a function of Δ .

When we consider the noise parameter of the public process σ , it is not always true that large uncertainty leads to a lower threshold. For high levels of monitoring intensity this would be so, but when the monitoring frequency is low a tighter threshold might be required even if the uncertainty level is higher, since the incentives for deviation increase with Δ . This can be seen in Figure 1, for example when $\sigma = 6$; for Δ around 2.3 we observe that the associated threshold tight above the threshold associated with $\sigma = 3$.

The convex shape of the threshold function in $\Delta \in (0, \bar{\Delta})$ is caused by two effects that operate in the same direction. They are the increasing informativeness of the public signal when the monitoring activity intensifies, and the expected immediate gains from deviating behaviour that also decrease when Δ becomes small. For infinitesimal Δ , the public signal is extremely informative and the expected immediate gains from deviation become negligible; for that reason the optimal threshold approaches $2d$. When the monitoring decreases in intensity, the sum of infinitesimal variations of the process is more likely to generate a "bad signal". Wrong punishments in the equilibrium path become more likely, forcing

the monitoring to relax, but in a decreasing way, because at the same time the expected immediate gains from deviating behaviour are becoming more attractive. At a certain point, when it reaches its minimum value, the optimal threshold value starts increasing with Δ , creating the convex shape. This happens because the expected gains from deviation become increasingly important and the monitoring has to become tighter in order to keep players with incentives. Finally, for large values of $\Delta \geq \bar{\Delta}$, there is no threshold value that can sustain mutual effort; shirk is a dominant strategy for both players.

V. MONITORING FREQUENCY AND LIMIT EFFICIENCY

In this section we look at the value \bar{v} of the upper bound on the set of SSPPE of the partnership game when the monitoring reaches its highest intensity, i.e. when $\Delta \rightarrow 0$. We will see that when such a value is feasible then is also efficient in the limit.

We will also look at the limit value of b^* and we will show that under Brownian uncertainty, monitoring intensity always has a positive effect on the payoffs. These results are independent of how much players discount the future and of the uncertainty level.

A. Monotonicity of the Best SSPPE Payoff

We start to present a result that establishes monotonicity between the value of the optimal upper bound on the set of SSPPE payoff and monitoring intensity. The result holds for any monitoring intensity $\Delta \in (0, \bar{\Delta})$.

Proposition 4 *For all $\Delta \in (0, \bar{\Delta})$ and $b \in (-\infty, 3d]$, the best SSPPE payoffs \bar{v}^* increase monotonically with the monitoring intensity, i.e. with a decrease in Δ .*

Proof. See Appendix C. ■

The monotonicity of the best SSPPE payoff to the monitoring intensity, in our setting, suggests that more monitoring always improves the payoffs. More monitoring is then preferred to less monitoring. The result is a consequence of the increased precision of the public signal when Δ becomes small. Kandori (1992) shows a similar result, where an exogenous improvement of the public signals causes a necessarily expansion on the set of PPE.

B. Limit Efficiency of the SSPPE Payoffs

We now turn our attention to the limit case, i.e. when $\Delta \rightarrow 0$. Observe that a threshold that is slightly larger (more tight) than b^* still satisfies the feasibility conditions but at the expense of a lower SSPPE payoff, while a lower value (more relaxed) than b^* cannot satisfy the feasibility constraint. The threshold b^* gives us exactly the point that maximizes \bar{v} according to Proposition 2. As we will see, limit efficiency does not require an optimal value for b , as in the limit, the same results hold with a larger threshold inside some region.

To obtain an efficient result in the limit we must see $F(b, \Delta) \rightarrow 0$ when $\Delta \rightarrow 0$, suggesting that in the limit, the probability of a type I error converges to 0. It confirms the comments previously made, that under Brownian uncertainty when the monitoring intensity is taken to the limit, the public signals become almost perfectly informative of players' actions. On the other hand, in the case of a deviation, the probability of detection $F'(b, \Delta)$ must not vanish in the limit, in order for an efficient result to be sustainable. We should not concern ourselves too much with the exact limit value of this probability since it depends on the exact form of $b^*(\Delta)$ which is unknown to us. As well will see, these simple conditions are sufficient to keep players with incentives to provide effort in the limit. Fudenberg and Levine (2007) discuss the necessity of similar conditions for the existence of an efficient limit equilibrium.

Under the optimal punishment scheme of the previous section, we had the relation $\bar{v}^* \geq \underline{v}^* = \pi^N$ for some $\Delta \in (0, \bar{\Delta})$, with $\bar{\Delta}$ depending on the parameters of the model, and $\bar{v}^* = \underline{v}^* = \pi^N$ in some interval $[\bar{\Delta}, \infty)$, i.e. the infinite repetition of the static Nash payoff. These are in accordance with Proposition 2 and Proposition 3 of Section IV. We can, however, set a lower bound on the set of SSPPE payoffs, such that $\underline{v} > \pi^N$. The punishment stage is not absorbing anymore; returning to the cooperative path is then possible. Such relaxation on the expected value (a larger \underline{v}) of a game that starts with an observation of a bad signal impacts negatively on the value of the upper bound \bar{v} , since it associates a lower continuation value. Consequently, the following relations must hold $\bar{v}^* \geq \bar{v} \geq \underline{v} \geq \pi^N = \underline{v}^*$ and $b^* \leq b \leq 3d$.

We will now present the main result of the paper. The idea is to take a feasible interval $[\underline{v}, \bar{v}] \subseteq [\pi^N, \bar{v}^*]$ and show that the associated best SSPPE payoff is efficient in the limit. Intuitively, if $\bar{v} \rightarrow \pi$ and since $\bar{v} \leq \bar{v}^*$, then \bar{v}^* must converge to π as well.

Proposition 5 *When the public signal follows (2.1) with distinct initial conditions for dif-*

ferent effort profiles, the upper bound on the set of SSPPE payoffs of the infinitely repeated partnership game $\bar{v} \rightarrow \pi$ as $\Delta \rightarrow 0$ providing that $r > 0$ and $\sigma < \infty$, for any $b \in [b^*, 3d]$.

Proof. See Appendix C. ■

Proposition 5 tell us that when the initial conditions of the process reflect different effort profiles and the public signals are Brownian and continuously available, we can achieve efficient SSPPE payoff by monitoring the process at the highest frequencies.

There is a large set of thresholds that make the result possible in the limit, which was shown for any $b \in [b^*, 3d]$. Since the set of feasible and admissible decision rules is just a subset $[b^*, b_*]$ of $[b^*, 3d]$, we can restrict our result to this subset, see Appendix A.

The result is independently of how players discount the future and independently of the magnitude of the noise parameter in the process. The lack of importance of the players' patience in the limit is easy to reconcile. Even if players have a large discount rate, in the limit δ is a continuous function and potential gains from deviating behaviour are negligible. The independence from σ is more surprising. No matter how large this parameter is, any realization of the process is still multiplied by $\sqrt{\Delta}$, which goes to zero with Δ , explaining its limit irrelevance.

It is important to understand that even though $\underline{v} \rightarrow \pi$, punishment it still present. What is happening is that once in the punishment state the probability of returning to the reward state converges to one, implying that the continuation payoff associated with \underline{v} converges to the value \bar{v} , becoming the dominant component. Punishments are expected to last no longer than a few numbers of instantaneous periods, sufficient to keep players with incentives.

Figure 2 shows how the value of the best SSPPE payoff \bar{v}^* associated with the thresholds of Figure 1 converge to the most efficient outcome $\pi = 2$, for varying monitoring intensities. Figure 2 clearly shows the strict monotonic improvement in the best SSPPE payoffs towards efficiency as the monitoring increases, in parallel with the statement of Proposition 4 and Proposition 5.

The increased informativeness of the public signals for high levels of monitoring intensity is the key aspect. It is due to the measurable distance between different initial conditions of the process²⁹. Such a distance in a process of infinitesimal variation plays a crucial role and makes destruction of value along the equilibrium path converging to zero with Δ . Another

²⁹ In the context of the partnership game, by measurable distance, we are referring to $2d = |2d - 4d|$.




FIG. 2: The best strongly symmetric payoff as a function of Δ .

important effect that occurs for small Δ is the decrease in the expected immediate gains from a deviation. At the end of Section IV above, these two effects were discussed in some detail.

Wrong inference about players equilibrium actions tends to vanish in the limit. Relevant uncertainty arises only if players cannot observe the public process during some measurable time interval. Then the accumulated sum of infinitesimal normal events may be misleading, which is even more likely the larger the time interval is, during which the process was left unattended.

The mechanics of the repeated games plays a role in the result. That is, actions are decided by players at the beginning of each period, at the end of the period the state of the process is observed, uncertainty is reset³⁰, new actions are taken for the following period and so on. Under Brownian uncertainty, when we shrink the time interval of this cyclical

³⁰ Uncertainty reset is equivalent to saying that either the value of the process is reset or players *actively coordinate*. See the discussion on *active coordination* in Section III.

process to the limit, what Fudenberg and Levine (2007) call "fast play", we get signals that are "almost" perfectly informative of players actions. Almost in the sense that we still have some infinitesimal variation and also because we are still not able to determine the identity of the deviator³¹. In Section VI A we discuss in more detail the meaning of the terminology "almost" perfect monitoring.

We now present a last result, concerning the limit value of the optimal threshold b^* .

Corollary 6 *When $\Delta \rightarrow 0$ the optimal threshold $b^*(\Delta)$ converges to $2d$, i.e. the expected signal associated with the deviation with less impact on the distribution of the public signals.*

Proof. See Appendix C. ■

For our particular game, the result simply states that $b^* \rightarrow 2d$ in the limit. The result is, however, more general; b^* must converge to the expected signal associated with the deviation with less impact on the distribution of the public signals. This is true even if the game has a continuous action space.

The focus here is on the limit of a sequence of discrete time games; however, the result also suggests that the continuous time analogue to b^* should be $2d$. We will discuss these issues in more detail in Section VI A.

A few comments before complete this section. As mentioned before, the true functional form of b^* is unknown, but clearly must depend on all the parameters of the model. It is however nontrivial to verify the rate at which b^* converges to $2d$. A wide number of numerical simulations suggest that this convergence must occur at a rate lower than or equal to $\sqrt{\Delta}$. Suppose that the optimal threshold function has the following structure: $b^*(\Delta) = 2d - \Delta^\alpha k(\cdot)$ where $k(\cdot)$ is some function of all the parameters in the model that converges to some bounded value and $\alpha > 0$. While choosing $\alpha \geq 1/2$ we can show that efficient and feasible results hold in the limit. On the other hand, when choosing a smaller α than $1/2$ we can bound b^* from below in some interval, violating, however feasibility. These results suggest that $b^* \rightarrow 2d$ at a rate of $\Delta^{0.499999(9)}$. We will not develop this idea further, since it is not required for any of the results. Nonetheless, the remark is left here.

³¹ Information about the identity of the deviator is on the basis of strong folk theorems under perfect monitoring, see for example Fudenberg and Maskin (1986), as well as Fudenberg, Levine and Maskin (1994) for a folk theorem under imperfect public monitoring.

The results presented in this paper also hold if we consider other well known Gaussian processes of infinitesimal variation, as for example, the geometric Brownian motion or the Ornstein-Uhlenbeck process. Finally, as discussed in section III, without any extra conditions on the information structure, these results can be generalized to others settings with a finite number of players and many actions.

VI. POSSIBLE EXTENSIONS - SOME COMMENTS

In this section we briefly discuss two important extensions associated with the methodology presented in this paper. The first, is how the results from a continuous time analogue partnership game would differ from the discrete time version. Another important extension would be to see how the approach of the present paper would perform if players have a continuum of actions available. This is important since the methodology exploits the jumping effect caused by a discrete deviation.

A. The Continuous Time

In this paper the emphasis has been placed on the limit of a discrete time repeated game. It is also possible to employ the methods of this paper directly in continuous time.

It is important to interpret the limit of the time interval Δ as an approximation or an abstraction for what would be an infinitesimal time transition. In this way, the tails of the distribution would not vanish. The limit degeneracy effect observed on the distribution of the public signals (see the proof of Proposition 5) would not occur directly in continuous time.

Brownian paths are continuous but not differentiable. The lack of smoothness of the paths allows us to infer the existence of infinitesimal variation. It is not clear that the limit of a sequence of discrete time games as presented in this paper and the analogous model directly in continuous times would exactly agree³². However, they should not differ significantly.

³² There are some reasons for this: in particular Ito's integral is defined in mean square sense and Brownian motion is a construction based on infinitesimal approximations. Because of that, it is difficult to translate the meaning of infinitesimal transition to discrete time. The distribution associated with a single infinitesimal transition is then not well defined, degenerating in the limit.

The approach was based on building a sequence of thresholds for each monitoring frequency that was shown to converge to the point $2d$. When we take the limit of a discrete time game, the rate of convergence to the limit point plays a role and we have to take it into consideration. Directly in continuous time we do not have such a problem. This paper suggests that the optimal threshold in continuous time is a value in the low neighbourhood of $2d$.

In the partnership game, if the equilibrium expected signal is $4d$ in continuous time, the probability that in an instantaneous transition would touch the threshold $2d$, does not vanish, although it is a very unlikely event. Note, however, that in an infinitely repeated game played in continuous time, these events of almost zero probability are likely to happen infinitely often. Their destruction of value in the game is negligible however. The reason is because the expected number of infinitesimal periods of punishment needed to sustain the equilibrium path will have zero measure. Thus we have an uncountable number of moments where the path of the game is in the static Nash, but with total measure zero. In an intuitive way, we can say that the path of the game is always in the cooperative equilibrium, with an expected equilibrium payoff of π , except for an uncountable number of periods with total measure zero.

This discussion explains in some way the use of the terminology "almost" perfect information to describe the limit situation.

B. A Game with a Continuous Action Space

When the action space is discrete in the limit, deviations from the equilibrium path are equivalent to jumps in the process. Since Brownian paths are continuous but not smooth, such defective behaviour is almost surely detected. In this section we briefly discuss the infinitely repeated Cournot game. This game is of interest since it has a continuum of actions available for each player and so deviations can be infinitesimal.

We apply the bang-bang result for the strongly symmetric equilibrium in the spirit of Abreu, Pearce and Stacchetti (1986 and 1990)³³.

In brief, the stage game expected payoffs (ex-ante) are given by $\pi_i(q_1, q_2) = q_i P(Q)$ where

³³ I thank Andrzej Skrzypacz for providing me with the material needed to compute and understand the mechanics of the best SSPPE payoffs in the Cournot game.




FIG. 3: The best strongly symmetric equilibrium in the Cournot game.

$P(Q)$ is the inverse demand function and $Q = q_1 + q_2$ is the aggregate supply. Firms, always have the possibility of staying out of the market and producing nothing. For simplicity and without loss generality, production costs are zero and firms face no capacity constraints, i.e. $q_i \in [0, \infty)$.

The stage game best strongly symmetric equilibrium is achieved when each firm supplies half of the monopoly quantity. Denote this quantity as q_i^M .

The two firms decide their supply quantities simultaneously and independently at moments in time $t = 0, \Delta, 2\Delta, \dots$ and observe the market price (the public signal) at times $t = \Delta, 2\Delta, \dots$. The public signal is given by (2.1)³⁴. Where now $y_{t+\Delta}$ is the observed market price at a given moment in time $t + \Delta$, and $y_t \equiv P(Q_t)$ is the inverse demand and the initial condition of the process, it reflects the individual private supply decisions taken by each player at the beginning of the period t .

³⁴ The observed state of the ABM price process may present a negative value. This undesirable feature of the model is irrelevant for the issue we wish to study and there is no loss in generality when considering such a process. A geometric Brownian motion process would be more adequate, but it also generates other technical problems.




FIG. 4: The expected public signal and the optimal threshold value.

For the case where $P(Q) = 1 - q_1 - q_2$ and $\sigma = r = 0.1$, Figure 3 shows the value of the best SSPPE for varying Δ . In particular when Δ becomes small the best SSPPE payoff converges to the value $1/8$, the perfect monitoring 50/50 monopoly split payoff. The numerical approximation suggests that the same efficient limit result, shown for the finite action space case (Proposition 5) also holds in the continuous action case. A monotonic improvement in the payoffs is clear when Δ becomes small, similar to the statement of Proposition 4.

It is also interesting to contrast the evolution of the optimal threshold against the expected signal of the process $P(Q^*)$ for varying Δ . In Figure 4 can be seen that as Δ gets small, the threshold value becomes tighter, in a similar fashion to the discrete actions case. In the limit converging to the expected signal associated with the most collusive equilibrium $P(Q^M) = 1/2$, as can be seen in Figure 4, reflecting the decreasing uncertainty when $\Delta \rightarrow 0$. This result is just an extension of Corollary 6 for the continuous action space case. In the limit the optimal threshold $b^*(\Delta)$ must converge to the expected signal associated with the deviation with less impact on the distribution of the public signals. In the continuous action case, such a deviation is an infinitesimal deviation.

The discussion suggests that all the results presented in Section V also hold when we consider games with a continuous action space. This is the main point highlighted by this section.

VII. FINAL COMMENTS

In the simple setting of a repeated partnership game, this paper shows that efficient SSPPE payoffs can be achieved in the limit when the public signal observed by the players is the state of an ABM and the initial conditions associated with different effort profiles are distinct.

Another conclusion of this paper is that payoffs improve monotonically with monitoring intensity.

Under Brownian uncertainty a degeneracy of the equilibrium in the limit or the opposed efficient scenario as shown here, depends crucially on the modelling adopted. Whereas the former approach tends to fit better in problems where the precision of the signals increases with the time interval between observation of the process, as shown in the work of Fudenberg and Levine (2007) and Sannikov and Skrzypacz (2007a), the present approach is more appropriate in situations where the informativeness of the signals improve with the monitoring intensity.

Reality evolves continuously, however, in economics, and contrary to other sciences, it is hard to think of situations where information is continuously available. An example could be the listed price of very liquid stocks and certain commodities that are available at high frequencies, but not continuously. Nevertheless, if such a possibility were available the most efficient outcomes would be achieved by continuously monitoring the state of process.

The reason why a continuous monitoring of the available information is not seen in real world economic problems is because it is often extremely costly and does not compensate the potential benefits. In other situations the nature of the problem might recommend monitoring at lower frequencies as suggested in the previously mentioned contributions.

Quoting Alchian and Demsetz (1972, p. 780), *"If detecting such behavior were costless, neither party would have an incentive to shirk, because neither could impose the cost of his shirking on the other."*

In practical applications when a partner continuously monitors the other, she probably

cannot devote her time to other activities, such as providing effort to the partnership. In practice what we observe are agents monitoring at discrete moments in time. In many occasions these monitoring events might be random, in the sense that an agent does not know the exact moment in time when the monitor is going to observe the public signal. Osório-Costa (2008) studies repeated games problems of this kind.

Even though the economic reality apparently finds no place for continuous monitoring, it is important to stress the importance of the results presented here. Although its focus was in the limit, this paper connects monitoring and actions frequency with their associated payoffs and decision rules. Repeated games theory can now mix actions frequency and players patience, creating a wide spectrum of potential applications and developments for the theory itself. While the impatience level of the players has typically been presented as an exogenous element in the repeated games theory, monitoring frequency has an enormous appeal to be endogenously determined by the problem at hand. This allows repeated games theory in discrete time to study problems in a richer fashion.

APPENDIX A: EFFICIENT DECISION RULES AND THE SMALLEST SET OF SSPPE.

Lemma 1 has constrained b to be in $(-\infty, 3d]$. In what follows we will show that such a constraint is always satisfied when the decision rule is admissible. Now, we define what is meant by an admissible threshold in our setting, which is similar to the meaning given in decision theory.

Definition 7 *Given a lower bound \underline{v} on the set of SSPPE payoffs, we say that a threshold value b is admissible when it has associated with it the largest feasible upper bound \bar{v} .*

The function \underline{v} is strictly concave in b , taking the maximum value $\underline{v}_* \leq \pi$ for some $b_* \in (-\infty, 3d]$. For that reason, the upper bound \bar{v} might not be unique. For example, for the same $\underline{v} \in (\pi^N, \underline{v}_*)$, we might be able to obtain two different upper bounds \bar{v}_1 and \bar{v}_2 with $\bar{v}_1 \geq \bar{v}_2$ and associated thresholds $b_1 \leq b_2$ respectively. In this case, it is not admissible to choose an decision rule other than the one associated with the largest upper bound, i.e. b_1 dominates b_2 .

Proposition 8 Given a $\Delta \in (0, \bar{\Delta})$, the set of feasible and admissible thresholds for the partnership game is the interval $[b^*, b_*] \subseteq (-\infty, 3d]$, where

$$b_* = \left\{ \arg \max_{b \in \mathbb{R}} \underline{v} \text{ s.t. } \underline{v} \geq \pi^N \right\},$$

and b^* is given by (4.4).

Given a $\Delta \in (0, \bar{\Delta})$, a threshold b is guaranteed to be admissible if it satisfies $\partial \underline{v} / \partial b \geq 0$ and is said to be feasible if satisfies $\underline{v} \geq \pi^N$.

When a decision rule belongs to the interval $(-\infty, 3d]$ the condition on b of Lemma 1 is necessarily satisfied. Moreover, associated with each bound b_* and b^* with $b_* \geq b^*$ on the set of SSPPE payoffs we have respectively, the smallest $[\underline{v}_*, \bar{v}_*]$ and the largest $[\underline{v}^*, \bar{v}^*]$ sets of SSPPE payoffs.

Proof. Note that $\underline{v} \geq \pi^N$ for all $\Delta \in (0, \bar{\Delta})$. Then to find b_* we just need to look at the derivative of \underline{v} . The $\partial \underline{v} / \partial b \geq 0$ is equivalent to

$$\frac{1 - \delta}{\delta} \geq \frac{F'(b, \Delta) F_b(b, \Delta) - F(b, \Delta) F'_b(b, \Delta)}{F'_b(b, \Delta) - F_b(b, \Delta)}, \quad (\text{A1})$$

where $F'_b(b, \Delta)$ and $F_b(b, \Delta)$ denote respectively the partial derivatives of $F'(b, \Delta)$ and $F(b, \Delta)$ with respect to b . The LHS does not depend on b . Suppose that the RHS is monotonic for $b \in (-\infty, 3d]$. When $b \rightarrow -\infty$ the ratio on the RHS converges to 0, while if $b \rightarrow 3d$ the ratio on the RHS converges to ∞ . So a value b_* must exist above which the preceding inequality is reversed. So $\partial \underline{v} / \partial b \geq 0$ for $b \in (-\infty, b_*]$ and $\partial \underline{v} / \partial b < 0$ for $b \in (b_*, 3d]$. Then, by monotonicity, \underline{v} is strictly concave, reaching the maximum value \underline{v}_* when $b = b_* \leq 3d$.

Since $\partial \underline{v} / \partial b \geq 0$ for $b \in (-\infty, b_*]$, following the same arguments for the optimality of \bar{v} as in Proposition 2, we lower the value of \underline{v} until $\underline{v} = \pi^N$ by decreasing b as well. Implying that $b_* \geq b^*$, and we must also have $\partial \underline{v} / \partial b > 0$ at $b = b^*$. In summary we have $\underline{v} \leq \underline{v}_*$ for $b \in (-\infty, b_*] \subseteq (-\infty, 3d]$ and $\underline{v} \geq \underline{v}^* = \pi^N$ for $b \in [b^*, 3d] \subseteq (-\infty, 3d]$. Putting all of this together we obtain that an feasible and admissible threshold b must belong to $(-\infty, b_*] \cap [b^*, 3d] = [b^*, b_*]$. ■

Note that the optimal decision rule b^* of Proposition 2 is a boundary value for the set of feasible and admissible decision rules.

Finally, it could have also been shown that: $b_* \rightarrow 3d$ when $\Delta \rightarrow 0$, that b_* is a strictly convex function in Δ and that b_* exists, and is unique and differentiable inside the interval

$(0, \bar{\Delta})$.

APPENDIX B:

Proof of Proposition 2. The best strongly symmetric equilibria is obtained by the value b that maximizes (4.2) subject to be feasible, i.e.

$$\pi^N \leq \underline{v} = \bar{v} - \frac{(1-\delta)}{\delta} \frac{\pi' - \pi}{F'(b, \Delta) - F(b, \Delta)} \leq \bar{v}.$$

The equality establishes IC and the first and second inequalities establish feasibility.

The payoff we want to maximize \bar{v} , increases monotonically when b decreases, i.e. $\partial \bar{v} / \partial b < 0$ for any $b \in \mathbb{R}$ since

$$F(b, \Delta) F'_b(b, \Delta) - F'(b, \Delta) F_b(b, \Delta) < 0,$$

for the Gaussian distribution. Where $F'_b(b, \Delta)$ and $F_b(b, \Delta)$ denote respectively the partial derivatives of $F'(b, \Delta)$ and $F(b, \Delta)$ with respect to b . On the other hand, $\partial \underline{v} / \partial b \geq 0$ only for $b \in (-\infty, b_*] \subseteq (-\infty, 3d]$, where b_* is the threshold value that solves (A1) with equality. See Proposition 8 and the respective proof in Appendix A. Then, clearly for $b \in (-\infty, b_*] \subseteq (-\infty, 3d]$ the derivatives of \bar{v} and \underline{v} have opposite signs.

The maximum \bar{v} is given by the lowest value of b that makes IC hold with equality and pushing the value of \underline{v} to the lower boundary of the feasible set, i.e. to π^N . Then, the optimal threshold value b^* must satisfy,

$$\pi(a^N) = \underline{v}^* = \bar{v}^* - \frac{(1-\delta)}{\delta} \frac{\pi' - \pi}{F'(b^*, \Delta) - F(b^*, \Delta)}. \quad (\text{B1})$$

Note that while $\underline{v}_* \equiv \max_{b \in \mathbb{R}} \underline{v} \geq \underline{v}^* = \pi^N$ we must have $b^* \leq b_*$, otherwise the set of SSPPE equilibria degenerates. See the proof of Proposition 8 in Appendix A for a more precise meaning of \underline{v}_* .

Plug equality (B1) into expression (4.2) and rearrange it, to obtain, after some manipulations

$$\bar{v}^* = \frac{(1-\delta)\pi + \delta F(b^*, \Delta)\pi^N}{1-\delta + \delta F(b^*, \Delta)},$$

which is the payoff associated with perpetual punishment. Plug this payoff back into the binding IC constraint (B1), which after some rearrangement gives expression (4.4). ■

Proof of Proposition 3. We extend the usual local implicit function theorem to hold in some convex interval $(0, \bar{\Delta})$. For now take as given the existence of an upper bound $\bar{\Delta}$ on the interval $(0, \bar{\Delta})$. Denote the LHS of equality (4.4) by $I(b, \Delta)$, which I rewrite here,

$$I(b, \Delta) \equiv F'(\pi - \pi^N) - F'(\pi' - \pi^N) - (1 - \delta)(\pi' - \pi) / \delta. \quad (\text{B2})$$

Where, to save on notation, we have denoted as $F \equiv F(b, \Delta)$ and $F' \equiv F'(b, \Delta)$, and the partial derivatives with respect to b will be denoted respectively as $F_b \equiv \partial F(b, \Delta) / \partial b$ and $F'_b \equiv \partial F'(b, \Delta) / \partial b$. Since F , F' and δ are continuous and differentiable w.r.t. to its own variables $\Delta \in (0, \bar{\Delta})$ and $b \in (-\infty, 3d]$, so thus the mapping $I(b, \Delta)$ is continuous and differentiable.

To simplify assume that the point (i) of Sandberg's (1981, p. 146) global implicit function theorem holds, that is, for some $\Delta_0 \in (0, \bar{\Delta})$ there is exactly one $b_0 \in (-\infty, 3d]$ such that $I(b_0, \Delta_0) = 0$. Now we need to verify that $I(b, \Delta)$ is locally solvable in the neighborhood of the point (b_0, Δ_0) and then by invoking continuity of $b(\Delta)$, Sandberg's theorem must hold for all $\Delta \in (0, \bar{\Delta})$. The condition for local solvability is $\partial I(b_0, \Delta_0) / \partial b \neq 0$ implying that $I(b_0, \Delta_0) = 0$. Since $\partial v / \partial b > 0$ for $b_0 \in (-\infty, b_*) \subseteq (-\infty, 3d]$, where b_* is the threshold value that solves (A1) with equality. (See Proposition 8 and the respective proof in Appendix A) Then $I(b_0, \Delta_0) = 0$ only when $\partial I(b_0, \Delta_0) / \partial b > 0$, or equivalently

$$\frac{\pi - \pi^N}{\pi' - \pi^N} > \frac{F_b}{F'_b} = e^{\frac{4b_0d - 12d^2}{2\Delta_0\sigma^2}}.$$

Since $\pi' > \pi$, the ratio on the LHS is constant, with a value strictly larger than zero but less than one. For any $b_0 \in (-\infty, b_*)$, $F_b/F'_b \rightarrow 0$ when $\Delta_0 \rightarrow 0$ which is the lowest value it can take. When $\Delta_0 \rightarrow \infty$ we have $F_b/F'_b \rightarrow 1$, the largest value it can take for $\Delta_0 > 0$. The ratio F_b/F'_b is clearly continuous, then exist always some nonempty interval of the form $(0, \bar{\Delta})$ where the above inequality hold, where $\bar{\Delta}$ is the lowest value in the pair (b_0, Δ_0) such that the above inequality stops to hold. Then for each point inside the interval $(0, \bar{\Delta})$ there is a unique continuous differentiable function $b(\Delta)$ on an open ball around (b_0, Δ_0) that locally satisfies $I(b, \Delta) = 0$.

Now apply Sandberg's implicit function theorem; by continuity of $b(\Delta)$, for each $S \in A$ there is a $T \in B$, where A and B are families of compact subsets of $(0, \bar{\Delta})$ and $(-\infty, 3d)$ respectively, $b(S)$ is compact as well and belongs to T . Then, there is a unique and continuous function $b^*(\Delta)$ such that $I(b^*, \Delta) = 0$ for all $\Delta \in (0, \bar{\Delta})$.

Intuitively, depending on the parameters of the model $\bar{\Delta}$ might be larger or smaller. The threshold $b^*(\Delta)$ adjusts to keep (B2) holding with equality for a given Δ . For $\Delta \in [\bar{\Delta}, \infty)$ there is no $b(\Delta)$ that makes (B2) hold, in this case $\partial I(b^*, \Delta) / \partial b = 0$. While $\Delta \in (0, \bar{\Delta})$ we can find $b^*(\Delta)$ that makes (B2) hold. ■

APPENDIX C:

Proof of Proposition 4. Start by defining $\bar{v}^*(\Delta) = \max_b \bar{v}(b, \Delta)$ subject to (4.4), i.e. $I(b, \Delta) = 0$ and write the Lagrangian $\mathcal{L}(b, \Delta) = \bar{v}(b, \Delta) - \lambda I(b, \Delta)$. By Proposition (3) the solution $b^*(\Delta)$ is a continuous and differentiable function of Δ ; assume the same holds for the Lagrangian multiplier λ . Then by the envelope theorem for constrained maximization problems we can write $\partial \bar{v}^*(\Delta) / \partial \Delta = \partial \mathcal{L}(b^*, \Delta) / \partial \Delta$. Our goal is then to show that

$$\frac{\partial \bar{v}^*(\Delta)}{\partial \Delta} = \frac{\partial \bar{v}(b^*, \Delta)}{\partial \Delta} - \lambda \frac{\partial I(b^*, \Delta)}{\partial \Delta} < 0 \quad (\text{C1})$$

where λ is obtained from solving $\partial \mathcal{L}(b, \Delta) / \partial b = 0$. Expression (C1) has the following three components, which we develop to

$$\begin{aligned} \frac{\partial \bar{v}(b^*, \Delta)}{\partial \Delta} &= -\frac{F'F_{\Delta} - FF'_{\Delta}}{(F' - F)^2} (\pi' - \pi), \\ \frac{\partial I(b^*, \Delta)}{\partial \Delta} &= F'_{\Delta} (\pi - \pi^N) - F_{\Delta} (\pi' - \pi^N) - \frac{r}{\delta} (\pi' - \pi), \\ \text{and } \lambda &= -\frac{(F'F_b - FF'_b) (\pi' - \pi) / (F' - F)^2}{F'_b (\pi - \pi^N) - F_b (\pi' - \pi^N)}. \end{aligned}$$

Where, to save on notation, we have denoted $F \equiv F(b, \Delta)$, $F' \equiv F'(b, \Delta)$, $F_{\Delta} \equiv \partial F(b, \Delta) / \partial \Delta$, $F'_{\Delta} \equiv \partial F'(b, \Delta) / \partial \Delta$, $F_b \equiv \partial F(b, \Delta) / \partial b$ and $F'_b \equiv \partial F'(b, \Delta) / \partial b$ when evaluated at $b = b^*$. Replacing these expressions in (C1), and after some algebra we obtain

$$[F' (\pi - \pi^N) - F (\pi' - \pi^N)] (F'_b F_{\Delta} - F_b F'_{\Delta}) > -\frac{r}{\delta} (\pi' - \pi) (F' F_b - F F'_b),$$

using the expression (4.4), we can simplify further to get

$$(1 - \delta) (F'_b F_{\Delta} - F_b F'_{\Delta}) > -r (F' F_b - F F'_b).$$

Notice that $F_{\Delta} = -\frac{(b^* - 4d)}{2\Delta} F_b$ and $F'_{\Delta} = -\frac{(b^* - 2d)}{2\Delta} F'_b$, then we are left with

$$(1 - \delta) F_b F'_b d / \Delta > -r (F' F_b - F F'_b).$$

The LHS is clearly always positive. Since for any $\Delta > 0$ we have $F'F_b - FF'_b > 0$ for the Gaussian distribution, the RHS is always negative. ■

Proof of Proposition 5. Our goal is to show that it is possible to obtain $\lim_{\Delta \rightarrow 0} \bar{v} = \pi$ without violating the feasibility constraint. Denote $b_0 \equiv \lim_{\Delta \rightarrow 0} b(\Delta)$ and focus on the values of $b_0 \in (-\infty, 3d]$. We will look at limit values of $b \in (-\infty, 3d]$ as required by Lemma (1). It insures that we cover all the relevant limit cases, including some potentially not feasible and inadmissible values of b , see Appendix A. Then, we obtained

$$F'(b, \Delta) \rightarrow \begin{cases} 0 & b_0 \in (-\infty, 2d) \\ 0 & b(\Delta) \rightarrow 2d \text{ slower than } \sqrt{\Delta} \\ \Phi(.) & b(\Delta) \rightarrow 2d \text{ at } \sqrt{\Delta} \text{ rate} \\ 1 & b(\Delta) \rightarrow 2d \text{ faster than } \sqrt{\Delta} \\ 1 & b_0 \in (2d, \infty) \end{cases},$$

and $F(b, \Delta) \rightarrow 0$ when $b \in (-\infty, 4d)$. Notice that when $b(\Delta) \rightarrow 2d$ we observe different limit values for $F'(b, \Delta)$ depending on the rate of convergence. Note also that when $b(\Delta) \rightarrow 2d$ at $\sqrt{\Delta}$ rate $F'(b, \Delta) \rightarrow \Phi(.)$, where the value of $0 < \Phi(.) < 1$ is left unspecified (this is without loss in generality), it depends on the true functional form of b^* which is unknown for us. In the limit we have a degeneracy on the distribution of the public signals; it can only take one of three possible values, 0, $\Phi(.)$ and 1 (keep in mind that taking the limit $\Delta \rightarrow 0$ is just an abstraction for an infinitesimal transition).

Since the IC condition has been imposed and $b \in (-\infty, 3d]$, we just need to verify when feasibility condition $\bar{v} \geq \underline{v} \geq \pi^N$ holds in the limit. Consider the following relevant possibilities about the limit value of b .

i) $b_0 \in (-\infty, 2d)$, we have $F'(b, \Delta) - F(b, \Delta) \rightarrow 0$ implying that $\underline{v} \rightarrow -\infty$, which is not feasible in the limit.

ii) $b_0 = 2d$ with $b(\Delta) \rightarrow 2d$ slower than $\sqrt{\Delta}$. In this case suppose $b(\Delta) = 2d - \Delta^\alpha k(.)$ with $0 < \alpha < 1/2$. We have $F'(b, \Delta) \rightarrow 0$ implying that $\underline{v} \rightarrow -\infty$, which is not feasible in the limit.

iii) $b_0 = 2d$ with $b(\Delta) \rightarrow 2d$ at least at the $\sqrt{\Delta}$ rate. Consider the extreme case $b(\Delta) = 2d - \sqrt{\Delta}k(.) \geq b^*$ for all $\Delta \in (0, \bar{\Delta})$. Then $F'(b, \Delta) = \Phi(-k(.)/\sigma) > 0$, implying that $\underline{v} \rightarrow \bar{v}$ and $\bar{v} \rightarrow \pi$, which is feasible and approaches efficiency in the limit.

iv) $b_0 \in (2d, 3d]$, we also obtain feasibility and convergence to the efficient outcome in the limit, i.e. $\underline{v} \rightarrow \bar{v}$ and $\bar{v} \rightarrow \pi$.

In both cases iii) and iv) we observe \bar{v} and \underline{v} converging to π , but with $\bar{v} > \underline{v}$ for all infinitesimal Δ . To see this observe that

$$\lim_{\Delta \rightarrow 0} (\bar{v} - \underline{v}) = \lim_{\Delta \rightarrow 0} \frac{(1 - \delta)}{\delta} \frac{\pi' - \pi}{F'(b, \Delta) - F(b, \Delta)} > 0,$$

which is equivalent to looking at $\lim_{\Delta \rightarrow 0} (1 - \delta)(\pi' - \pi) > 0$. Differentiating with respect to Δ , even though both \bar{v} and \underline{v} converge to π , we obtain that if $r > 0$ we must have $\underline{v} < \bar{v}$ for small infinitesimal Δ . The irrelevance of σ is always present in the proof. ■

Proof of Corollary 6. We know that the lower the value of b the larger the value of \bar{v} , since $\partial \bar{v} / \partial b < 0$ is always true for the Gaussian distribution, see Section IV. Our problem is then to find the lowest feasible value of b that maximizes \bar{v} . Using the results obtained in the proof of Proposition 5, the lowest value of b such that $F'(b, \Delta) \rightarrow \Phi(\cdot) > 0$ and at same time $F(b, \Delta) \rightarrow 0$, is the value to which b must converge in the limit. So $\lim_{\Delta \rightarrow 0} b^*(\Delta) = 2d$ is the lowest value which is feasible, and which achieves efficiency. Other limit values either are not lower or do not satisfy the feasibility constraint. ■

REFERENCES

1. Abreu, D., P. Milgrom and D. Pearce (1991). "Information and Timing in Repeated Partnerships." *Econometrica*, 59, 1713-1733.
2. Abreu, D., D. Pearce and E. Stacchetti (1986). "Optimal Cartel Equilibria with Imperfect Monitoring," *Journal of Economic Theory*, 39, 251-269.
3. Abreu, D., D. Pearce and E. Stacchetti (1990). "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." *Econometrica*, 58, 1041-1063.
4. Alchian, A. and H. Demsetz (1972). "Production, Information Costs, and Economic Organization." *American Economic Review*, 62, 777-795.
5. Faingold, E., Y. Sannikov (2007). "Reputation Effects and Equilibrium Degeneracy in Continuous-Time Games," mimeo.
6. Fudenberg, D., D. Levine and E. Maskin (1994). "The Folk Theorem with Imperfect Public Information." *Econometrica*, 62, 997-1040.
7. Fudenberg, D. and D. Levine (2007) "Continuous Time Models of Repeated Games with Imperfect Public Monitoring." *Review of Economic Dynamics*, 10(2), 173-192.
8. Fudenberg, D. and D. Levine (2008) "Repeated Games with Frequent Signals." *Quarterly Journal of Economics*, forthcoming.
9. Fudenberg, D. and E. Maskin (1986). "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, 54, 533-554.
10. Fudenberg, D. and J. Tirole (1991) *Game Theory*, MIT Press, Cambridge, MA.
11. Green, E. and R. Porter (1984). "Noncooperative Collusion under Imperfect Price Information." *Econometrica*, 52, 87-100.
12. Kandori, M. (1992) "The Use of Information in Repeated Games with Imperfect Monitoring." *Review of Economic Studies*, 59, 581-593.
13. Mailath, G. and L. Samuelson (2006) *Repeated Games and Reputations: Long-run Relationships*. Oxford University Press, New York.

14. Osório-Costa, A. (2008), "Efficiency Gains in Repeated Games at Random Moments in Time," mimeo.
15. Porter, R. (1983). "Optimal Cartel Trigger Price Strategies." *Journal of Economic Theory*, 29, 313-338.
16. Prakasa Rao, B.L.S. (1999) *Statistical Inference for Diffusion Type Processes*. Arnold Publishers and Oxford University Press.
17. Radner, R., R. Myerson, and E. Maskin (1996) "An Example of a Repeated Partnership Game with Discounting and with Uniformly Inefficient Equilibria," *Review of Economic Studies*, 53, 59-69.
18. Sandberg, I. W. (1981) "Global Implicit Function Theorems," *IEEE Transactions on Circuits and Systems*, 28, 145–149.
19. Sannikov, Y. (2007). "Games with Imperfectly Observable Actions in Continuous Time," *Econometrica*, 75, 1285–1329.
20. Sannikov, Y. and A. Skrzypacz (2007a) "Impossibility of Collusion under Imperfect Monitoring with Flexible Production," *American Economic Review*, 97, 1794-1823.
21. Sannikov, Y. and A. Skrzypacz (2007b) "The Role of Information in Repeated Games with Frequent Actions," mimeo.