Efficiency Gains in Repeated Games at Random Moments in Time*

António M. Osório-Costa†

Universidad Carlos III de Madrid

(Dated: First Draft August 8, 2008 - This Draft January 30, 2009)

Abstract

This paper studies repeated games where the time of repetitions of the stage game is not known or controlled by the players. Many economic situations of interest where players repeatedly interact share this feature, players do not know exactly when is the next time they will be called to play again. We call this feature random monitoring. We show that perfect random monitoring is always superior to perfect deterministic monitoring when players discount function is convex in time domain. Surprisingly when the monitoring is imperfect but public the result does not extend in the same absolute sense. The positive effect in the players discounting is not sufficient to compensate for a larger probability of punishment for all frequencies of play. However, we establish conditions under which random monitoring allows efficiency gains on the value of the best strongly symmetric equilibrium payoffs, when compared with the classic deterministic approach.

JEL: C73, D82.

KEYWORDS: Repeated Games, Frequent Monitoring, Random Monitoring, Perfect and Imperfect Monitoring, Moral Hazard, Stochastic Processes.

---

*Financial support from Fundação para Ciência e Tecnologia grant SFRH/BD/17157/2004 is gratefully acknowledged. I would like to thank Juan Pablo Rincón-Zapatero and Galina Zudenkova, as well as the seminar participants at ..., ..., 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

Many economic situations where agents repeatedly interact resemble and are studied as repeated games. While the research in repeated games as focus exclusively on the case where the repeated game is played at predetermined and equally spaced moments in time, this paper consider the possibility that the stage game is played at random and not equally spaced moments in time. In other words, we step further to meet the economic reality that suggests that indeed many economic interactions of interest are in fact similar to repeated games but not necessarily repeated with the same known frequency. Players do not know and do not control the moments in time in which they are called to play the stage game.

Under this possibility, form a technical point of view we then ask, what changes have to be introduce in the usual repeated games methodology in order to accommodate such a possibility? And in particular from an efficiency perspective, we are concerned on how the repeated games payoffs change when there is uncertainty about the time of repetitions of the stage game, compared with the usual case. This paper addresses the answers to these questions, not only when the monitoring is perfect but also when it is public and imperfect[1][2].

We show that under perfect monitoring, random monitoring is always superior to deterministic monitoring when players discount function is convex in time domain. Their decisions are then based on a larger discount factor, we call it the *expected discount factor effect*. Surprisingly when the monitoring is imperfect but public the result does not extend in the same absolute sense. The positive effect on the players discounting is not sufficient to compensate potential adverse effects on the distribution of the public signals for all frequencies of play. However, we establish conditions under which random monitoring allows efficiency gains on the value of the best strongly symmetric equilibrium payoffs, when

[1] When monitoring is imperfect but public player commonly observe a noisy public signals about other players actions. See Radner, Myerson and Maskin (1986) where the output of a partnership is a noisy signal of players' actions, or Green and Porter (1984) and Porter (1983) where the market price is an imperfect signal of the quantities supplied by firms. Fudenberg and Tirole (1991) and Mailath and Samuelson (2006) for complete surveys of the problems and methods used to solve repeated games.

[2] The case where the stage game is repeated at unknown and not equally spaced moments in time, we call it *random monitoring*. When the stage game is repeated at known and equally spaced moments in time we call it *deterministic monitoring*. These concepts should not be confused with perfect or imperfect monitoring. Perfect monitoring can be either random or deterministic, and the same with imperfect monitoring.

compared with the classic deterministic approach.

The study of random monitoring, as done in the present paper, would not be possible without the recent advances in the theory of frequent monitoring. After the seminal work of Abreu, Milgrom and Pearce (1991), renewed interest in frequent monitoring has reemerged in particular due to Sannikov (2007). In the spirit of the latter work, that is, studying repeated games directly in continuous time, see Faingold (2006) and Faingold and Sannikov (2007). More similar to the former work, that is, studying the limit of the discrete time games, see Fudenberg and Levine (2007 and 2008), Osório-Costa (2008) and Sannikov and Skrzypacz (2007a and b).

In this paper we do not focus in a particular monitoring intensity but rather we study discrete time games for the all spectrum of monitoring frequencies where repeated play can improve over the static Nash payoffs. For that reason the approach followed by the latter contributions is elected. Even though that these papers where mainly concern with the answer to the limit case, by studying the associated sequence of discrete time it is possible to study repeated games for arbitrary monitoring intensities.

Abreu, Milgrom and Pearce (1991) show that the value of the best strongly symmetric equilibrium degenerate in the limit when the realizations of the public process represent bad news. The lack of observed public signals becomes infinitely likely in the limit. Fudenberg and Levine (2007) and Sannikov and Skrzypacz (2007a) present a similar result when the public signal is Brownian rather than Poisson. Their degeneracy effect is due to a degradation of the information content of the public signals for high monitoring intensities.

When the realizations of the public process are interpreted as bad news, Abreu, Milgrom and Pearce (1991) have shown, under some conditions, that equilibrium payoffs above the static Nash, but not efficient can be sustained in the limit. It is also in the limit that the most efficient equilibrium is achieved. Under Brownian uncertainty, Osório-Costa (2008) presents a similar result but efficient in the limit, where payoffs improve monotonically with the monitoring intensity.

Monitoring is in general a costly activity. For that reason there is an enormous spectrum of potential applications for the results of the present paper, not only in repeated games and dynamic game theory in general, but also in the theory of contracts and mechanism design. Once we identify when random monitoring is superior to deterministic monitoring (and consequently, the other way around), we can select the monitoring technology that achieves larger payoffs for a same monitoring frequency[3] (The same expected costs) or from another perspective we fix the payoff and search for the less costly monitoring.

The paper is organized as follows. Section II presents the repeated game model and notation that is common for both the perfect and imperfect monitoring cases. Special attention is given to the expected discount factor. Section III focus on the perfect monitoring case, the first important result of this paper is presented. Section IV studies the imperfect random monitoring case. First we look at the distribution of the public signals, then we characterize the value of the best SSE payoff, and finally we present our results. Finally, in Section V we illustrate our findings, in the framework of some important discrete time approaches to frequent monitoring presented in the literature.

II. THE MODEL AND THE EXPECTED DISCOUNT FACTOR

Crucial in our problem is how players discount the future. Usually it is assumed that the stage game is repeated at predetermined moments in time and that players discount the future according to a common discount factor; we denote it as δ^{τ}. It is also commonly assumed that players either discount the future according to an exponential or a hyperbolic discount factor. In either case discounting is convex in time. For that reason and for convenience, without loss in generality we implicitly assume exponential discounting $\delta^{\tau} \equiv e^{-r\tau}$. Where r is the discount rate and τ is the interval of time between repetitions of the stage game, usually assumed equal to the unit. Here we will not make assumptions about the value of τ and we consider also the possibility of this parameter to be a random variable. Then the discount factor is no more a deterministic but rather a stochastic function of time.

At moments in time $\tau_0, \tau_1, \tau_2, ...$, each player $i \in N$ chooses its individual action $a_i^{\tau_k}$, with $k \in \mathbb{N} = \{0, 1, 2, ...\}$, from some finite action space $A_i^{\tau_k}$. Where $\tau_0 = 0$ is the known moment in time where players play for the first time, τ_1 is the moment in time where players play for the second time, and so on. Denote $A^{\tau_k} = \times_{i=1}^{n} A_i^{\tau_k}$ as the set of action profiles endowed with the product topology of the individual action spaces, with generic element $a^{\tau_k} = (a_1^{\tau_k}, ..., a_2^{\tau_k})$

[3] Under random monitoring, the monitoring events are random variables with support in some interval, when in expected terms these events match the deterministic and equally spaced monitoring moments; we use the terminology "same monitoring frequency (or intensity)". In expected terms the monitoring events happen the same number of times, see expression (2.2) below.

denoting a profile of actions.

At moments in time $\tau_1, \tau_2, ...$, players simultaneously observe a continuously available signal which provides noisy information about the action profile that has resulted from each player individual and private decisions. Actions are taken in the beginning of each period but payoffs are only collected in the end of the period[4].

Under the expected utility hypothesis, the player i's ex-ante expected payoff at time τ_k, denoted as $\pi_i(a^{\tau_k})$, is the relevant element for studying the game.

Player i's$\in N = \{1, ..., n\}$ infinite sum of payoffs[5] from the repeated game is then given by

$$\delta^{\tau_1}\pi_i(a^{\tau_0}) + \delta^{\tau_2}\pi_i(a^{\tau_1}) + \delta^{\tau_3}\pi_i(a^{\tau_2}) + ...$$

As mentioned before, the sequence of times $\tau_1, \tau_2, ...$ might be predetermined and known at the beginning of the game with each consecutive time being equally spaced in time $\tau_k - \tau_{k-1} = \tau$, or it might be unknown with each value τ_k, $k \in \mathbb{N}$, being repeatedly draw from some known distribution[6]. The former case we will call it deterministic monitoring while the second we will call it random monitoring.

When the repetitions of the stage game are random, since the stage game payoffs of the repeated game are independent of time, the expected discounted stream of payoffs is given by

$$\begin{aligned} & E(\delta^{\tau_1})\pi_i(a^{\tau_0}) + E(\delta^{\tau_2})\pi_i(a^{\tau_1}) + E(\delta^{\tau_3})\pi_i(a^{\tau_2}) + ... \\ & = E(\delta^x)\pi_i(a^{\tau_0}) + E(\delta^x)^2\pi_i(a^{\tau_1}) + E(\delta^x)^3\pi_i(a^{\tau_2}) + ... \end{aligned}$$

Where we have made use of the fact that $\tau_k = \tau_{k-1} + x_k$, where x_k is some random variables, and $x_0 = \tau_0 = 0$ is known. Recursively we find that $\tau_k = \sum_{j=0}^{k} x_j$. The expected discount factor for each moment in time is then,

$$E(\delta^{\tau_k}) = E\left(\delta^{\sum_{j=0}^{k} x_j}\right) = \prod_{j=0}^{k} E(\delta^{x_j}) = E(\delta^x)^k,$$

[4] This approach is more appealling. As usual, an appropriated normalization will turn the payoffs of the infinitely repeated game in the same units as the stage game.

[5] In a context of imperfect monitoring we call it the infinite sum of the ex-ante expected payoffs of the repeated game. In such a setting we also assume that this ex-ante expected payoffs are independent of the monitoring intensity, i.e. the stage game payoffs do not depend on τ.

[6] We abstain here to discuss on how players are informed about the time to play the stage game. Rather we focus in studying the associated expected payoffs. However, we can think of a public correlated signal.

where in last equality we have assumed that each x_k is i.i.d. across time[7].

Let the random variable x follows some continuous differentiable distribution $G(x)$ with support on the interval $(0, \overline{\tau})$ where $\overline{\tau} > 0$. It means that during the time interval $(0, \overline{\tau})$ a monitoring event occurs, however, not exactly known when. Then the expected discount factor is denoted and equals

$$E(\delta^x) \equiv \int_0^{\overline{\tau}} e^{-rx} dG(x). \tag{2.1}$$

In order to make comparisons between deterministic and random monitoring meaningful, the value of $\overline{\tau} > 0$ is chosen, such that it solves

$$E(x) = \int_0^{\overline{\tau}} x dG(x) = \tau. \tag{2.2}$$

Now depending on whether we consider perfect or imperfect public monitoring we develop each specific aspects in the respective section.

III. PERFECT MONITORING

When studying perfect monitoring with random repetitions of the stage game, the only relevant aspect that changes is the discount factor due to the time uncertainty brought to the problem by the random monitoring. When the stage game is repeated at predetermined and equal spaced moments in time players discount the future according to δ^τ on the other hand when the repetitions are not predetermined the discount factor to apply is $E(\delta^x)$. In the previous section we have seen the difference between these discount factors.

Before go further, we develop some more notation while we review some important concepts in repeated games. Following Abreu, Pearce and Stacchetti (1990) a payoff profile is a pure-strategy subgame-perfect equilibrium if after a given profile of actions there are associated continuation payoffs that are themselves subgame-perfect equilibrium of the continuation game.

More formally, denote $V^p \subset \mathbb{R}^n$ the set of pure-strategy subgame-perfect equilibria payoffs. If for each action profile a^* there are credible continuation promises $w(a^*) \in V^p$ such

[7] We can obtained the same result if instead, we had computed the recursive interacted conditional expectation $E(E(...E(E(\delta^{\tau_k}|\tau_{k-1})|\tau_{k-2})...|\tau_1)|\tau_0) = E(\delta^x)^k$, for $\tau_k = \tau_{k-1} + x$, with the expectation taken w.r.t. the i.i.d. random variables x.

that for each player $i \in N$, and all $a_i' \in A_i$,

$$v_i = (1-\delta^\tau)\,\pi_i(a^*) + \delta^\tau w_i(a^*) \geq (1-\delta^\tau)\,\pi_i\left(a_i', a_{-i}^*\right) + \delta^\tau w_i\left(a_i', a_{-i}^*\right),$$

then (by the inequality) the profile a^* is enforceable (i.e. incentive compatible) on V^p and the payoff profile $v \in V^p$ is a pure-strategy subgame-perfect equilibria[8].

With just a slight increase of notation the same concepts generalize with random monitoring, by simply substituting δ^τ for $E(\delta^x)$. For our proposes is enough to distinguish $\widetilde{V}^p \subset \mathbb{R}^n$ as the set of pure-strategy subgame-perfect equilibria payoffs under random monitoring.

We can now present our first result. Fix a monitoring intensity $\tau = E(x)$, in deterministic or random terms. Suppose we have a particular pure-strategy subgame-perfect equilibrium payoff profile v that can be sustained both by monitoring randomly, i.e. $v \in \widetilde{V}^p$ or by monitoring in a deterministic way, i.e. $v \in V^p$. Then we may ask, under which conditions one type of monitoring requires a lower discount factor than the other? The following result establishes such a condition for perfect monitoring in a general setting.

Theorem 1 *Given an monitoring intensity $\tau = E(x)$, to sustain a particular pure-strategy subgame-perfect equilibrium payoff profile v, random monitoring requires a lower discount factor than deterministic monitoring if $\delta^x \in (0,1)$ is strictly convex in x.*

Proof. Assuming all the conditions for pure-strategy subgame-perfect equilibria are satisfied. In the deterministic monitoring case, enforceability of the profile a^*, with respect to continuations in V^p, requires that

$$\delta^\tau \geq \sup_{i \in N \text{ and } a_i' \in A_i} \frac{\pi_i\left(a_i', a_{-i}^*\right) - \pi_i(a^*)}{\pi_i\left(a_i', a_{-i}^*\right) - \pi_i(a^*) + w_i(a^*) - w_i\left(a_i', a_{-i}^*\right)} \equiv \overline{\delta}_v,$$

and $E(\delta^x) \geq \overline{\delta}_v$ in the random monitoring case with continuations in $\widetilde{V}^p$. Where $\pi_i\left(a_i', a_{-i}^*\right) \geq \pi_i(a^*)$ and $w_i(a^*) \geq w_i\left(a_i', a_{-i}^*\right) \geq \underline{v}_i^p$ where $\underline{v}_i^p \equiv \min_{a_{-i} \in A_{-i}} \max_{a_i \in A_i} \pi_i(a_i, a_{-i})$. Given a same equilibrium payoff $v \in V^p$ and $v \in \widetilde{V}^p$, random monitoring requires a lower discount factor than deterministic monitoring if $E(\delta^x) > \delta^\tau$, which is always true by Jensen's inequality when δ^x is convex, since $\tau = E(x) > 0$. ■

[8] We directly assume that any value v and $w(a)$ belongs to V^p, for that reason we do not mention the "largest self-generating property" of the set of pure-strategy subgame-perfect equilibria payoffs.

Since the monitoring is perfect and we assume a general payoff profile v that is subgame-perfect equilibrium under deterministic and random monitoring, the associated lower bound on discounting does not depend on r and τ. For that reason the result is just centered around the comparison between discount factors.

We stress that it is commonly assumed that players discount the future in a convex way, according to either an exponential or a hyperbolic discount factor. For that reason it is very unlikely that conditions of the theorem stops to hold. Under perfect monitoring, only when δ^x is not strictly convex or when the distribution of the random time is degenerate the theorem does not apply.

Numerical Example: As an illustration, suppose that $x \sim U(0, \overline{\tau})$ with $\overline{\tau} = 2\tau$, implying that $E(x) = \tau$ as required in Theorem 1. A folk theorem[9] in pure-strategies for the prisoners' dilemma of Table I can be obtained providing that players discount factor δ^τ is greater than 2/3. The same folk theorem can be obtained with random monitoring with a lower discount factor of 0.645901. Similarly if we allow player to use public correlation we can obtain a folk theorem with a discount factor of 0.243968, rather than 1/3[10].□

The numerical example demonstrate that a lower discount factor is needed to obtain a particular folk theorem under random monitoring. We can also reinterpret Theorem 1 by saying that there are payoffs which cannot be sustained under deterministic monitoring, but can be sustained as pure-strategy subgame-perfect equilibria payoffs if the repetitions of the stage game are unknown to the players. The following result formalizes this intuition.

Corollary 2 *When the conditions of Theorem ... hold, then for a given discount rate r and an expected monitoring intensity τ, there are pure-strategy subgame-perfect equilibria payoffs that can be sustained only under random monitoring, the converse is not true.*

Proof. When $\delta^x \in (0, 1)$ is strictly convex in x and $E(x) = \tau$ by Jensen's inequality we have $E(\delta^x) > \delta^\tau$. Then exist equilibria where $\overline{\delta} > \delta^\tau$ and on same time it happens that $E(\delta^x) > \overline{\delta}$, for some combinations of r and τ. The converse requires $\delta^\tau > E(\delta^x)$ which clearly cannot be true if the conditions of Theorem 1 hold. ■

[9] We say that a folk theorem exist if any feasible and strictly individual rational payoff is contained in the set of subgame-perfect equilibria payoffs. For folk theorems with discounting in infinitely repeated games, see the classical papers of Friedman (1971) and Fudenberg and Maskin (1986).

[10] Mailath and Samuelson (2006, sections 2.5.3 and 2.5.6) show how to obtain these values for the perfect deterministic monitoring case.

The result tell us that depending on r and τ, the set of subgame-perfect equilibria obtained with deterministic monitoring is a subset of the set of subgame-perfect equilibria obtained with random monitoring, that is $V^p \subseteq \widetilde{V}^p$.

Numerical Example (revisited): In face of Corollary 2 we can reinterpret the above numerical example in different ways. Suppose $\delta^\tau = 0.65$ then with deterministic monitoring we must be able to sustain a great number of equilibria but not a folk theorem since $\delta^\tau < 2/3$. However, with random monitoring we get $E(\delta^x) = 0.670291 > 2/3$, allowing us to sustain any feasible and weakly individual rational payoff profile. Similarly, when public correlation is allowed, if $\delta^\tau = 0.3$ we can sustain simply the infinite repetition of the stage game Nash equilibria, however with random monitoring we can obtain a folk theorem since $E(\delta^x) = 0.377916 > 1/3$[11].□

Notice that when the monitoring is random, players' true discount factor does not change. However, under the expected utility hypothesis, their decisions are based on the expected discount factor, which has become larger. This reflects the intuition behind these results. The uncertainty about the moments where the stage game is repeated brings uncertainty about the value of gains that a potential deviator may contemplate. In expected terms they get smaller, the continuation value of the game becomes more important. In some sense it is equivalent as to say that under perfect random monitoring players become more patient.

One aspect that we did not explore in detail in this section, was the expected discount factor effect that tends to be stronger for lower frequencies of monitoring (equivalently - lower discount factor). Such can be seen in the numerical examples above. This fact will shows extremely useful to interpret the results of Section V.

IV. IMPERFECT PUBLIC MONITORING

As show in the previous section, the superiority of random monitoring in the context of perfect monitoring is absolute[12] when the discount factor is convex. The same it is not true in games with imperfect monitoring. The difference is that together with the discount factor,

[11] Stahl, II (1991) discuss in detail the discountinuity on the set of subgame-perfect equilibrium payoffs with public correlation.

[12] We say that random monitoring is superior in absolute terms to deterministic monitoring when for a given parameterization of the problem, it achieves larger payoffs for all monitoring intensities.

the distribution of the public signals is also affected by the uncertainty on the repetitions of the stage game. We can however establish conditions on when one monitoring technology is superior to the other.

For simplicity, we will consider general symmetric finite stage games which have a least one symmetric Nash equilibria (potentially mixed), that players use to coordinate the mutual punishments[13]. We focus on *strongly symmetric equilibria* (SSE henceforth), where after every public history the same action is chosen by both players. A strategy is public if in any moment in time where players are required to play the stage game, it depends only on the public histories and not on player i private history. Given a public history, a profile of public strategies that induces a Nash equilibrium on the continuation game from that time on, is called a *perfect public equilibrium* (PPE). Moreover if the other players $-i$ are playing public strategies, player i's best reply can only be a public strategy.

Player i's$\in N$ payoffs in the infinitely repeated game is the discounted normalized sum of the stage-game expected payoffs, $(1-\delta^{\tau})\sum_{t=\tau_0,\tau_1,\tau_2,...}^{\infty}\delta^{t}\pi_i\left(a^t\right)$, induced by the profile of actions a^t and associated public history, for every $t=\tau_0,\tau_1,\tau_2,...$ and $\tau_0=0$.

A. The Distributions of the Public Signals

Denote by $Y_{\tau}\left(a\right)$ the random variable associated with the state of the stochastic process when observed at the moments in time τ and denote by $\widetilde{Y}_{\tau}\left(a\right)$ the random variable associated with the state of the stochastic process when the observation time is also a random variable[14]. Notice the dependence of these random variables on the action profile a that results from the private choices made by each of the player. The equal τ in the notation emphasizes that the monitoring intensity is the same.

In general terms for $\tau=E\left(x\right)$ an observation of the public signal either suggests cooperative behavior which we call "a good signals" and denote it as $\overline{y}_{\tau}$ and $\widetilde{\overline{y}}_{\tau}$ depending on whether we are considering deterministic or random monitoring respectively, or suggests de-

[13] We have restricted the set of games in a considerable way, we do this in order to make the problem more tractable. Such restriction does not diminuish in any way the point this paper wants to clearly address.

[14] As before assume that time is a random variable with some continuous diferentiable distribution $G\left(x\right)$ with support on the interval $(0,\overline{\tau})$ where $\overline{\tau}>0$. In order to make comparations between deterministic and random monitoring meaningfull, $\overline{\tau}$ is chosen such that $E\left(x\right)=\tau$. Our concern is that monitoring intensity will be the same.

fective behavior, which we call "a bad signals" and analogously denote as $\underline{y}_{\tau}$ and $\underline{\widetilde{y}}_{\tau}$. Then, depending on whether we are considering deterministic or random monitoring we observe $y_{\tau} \in \left\{ \underline{y}_{\tau}, \overline{y}_{\tau} \right\}$ or $\widetilde{y}_{\tau} \in \left\{ \underline{\widetilde{y}}_{\tau}, \widetilde{\overline{y}}_{\tau} \right\}$ respectively.

The partition is convenient, what matters for the players is the information content of the public signal. For example a realized public signal y_{τ} or $\widetilde{y}_{\tau}$ might take any value in $\mathbb{R}$, the partition of the signal space has integrated the interpretation given by the players to each public signal, reducing this infinite space to just two signals. The public history of the game is just a binary sequence of "good" and "bad" signals[15].

The probability of observing a "bad public signal" differs depending on whether we consider deterministic or random monitoring. In the deterministic case, given a profile of actions a^* where each of the players privately plays the equilibrium action, we denote the probability of a "bad signal" as

$$F \equiv \Pr\left(Y_{\tau}\left(a^*\right) \in \underline{y}_{\tau} \right).$$

When the repetitions of the stage game are random variables, we denote this probability as

$$\widetilde{F} \equiv \Pr\left(\widetilde{Y}_{\tau}\left(a^*\right) = \underline{\widetilde{y}}_{\tau} \right).$$

When the resulting profile (denoted as a') has implicit a unilateral deviation by a single player, these probabilities are denoted respectively as F' and $\widetilde{F}'$.

In order for the action profile a^* to be enforceable clearly we must have $\left\{ Y_{\tau}\left(a^*\right) \in \underline{y}_{\tau} \right\} \subset \left\{ Y_{\tau}\left(a'\right) \in \underline{y}_{\tau} \right\}$ and $\left\{ \widetilde{Y}_{\tau}\left(a^*\right) = \underline{\widetilde{y}}_{\tau} \right\} \subset \left\{ \widetilde{Y}_{\tau}\left(a'\right) = \underline{\widetilde{y}}_{\tau} \right\}$ implying that $F' > F$ and $\widetilde{F}' > \widetilde{F}$ respectively[16].

Note also that the discounting process is independent from the decision process. In other words, the expected discount factor and the probabilities of the different signals are computed separately. Nonetheless, for a same monitoring frequency, the probability of a given signal is sensitive on whether the observation time is or not a random variable.

[15] To give some intuition on what the sets $\underline{y}_{\tau}$ or $\underline{\widetilde{y}}_{\tau}$ might be. In the Brownian motion case a "bad signals" is just an observation of the public signal below a given threshold, in Poisson "bad news" case it is just a single or more sudden movements in the public signal path, in the Poisson "good news" case it is the total or parcial absence of movements in the public path. See Section V.

[16] For more elaboarated informative conditions on the public signals see Fudenberg, Levine and Maskin (1994), and also Fudenberg and Levine (1994).

B. The Best Strong Symmetric Equilibrium

In this section we will first derive the expression for the best SSE payoffs under deterministic monitoring, and then we simply generalize it to the random monitoring case. Finally we present the conditions under which random monitoring is superior to deterministic monitoring.

To shorten on notation, since the setting is symmetric and the recursive structure of the problem is preserved, we respectively drop players' and time indices. Where by symmetry, $\pi(a')$ is equal to all players, i.e. the most likely deviation is the same for all the players and for any value of τ. The same applies to $\pi(a^*)$ and $\pi\left(a^N\right)$.

In the deterministic case we denote $\overline{v}$ and $\underline{v}$ respectively as the upper and the lower bounds on the set of SSE payoffs. In other words $\overline{v}$ is the expected value when the play starts with the observation of a good signal and $\underline{v}$ is the expected value of the game when play starts with an observation of a bad signal.

We apply Abreu, Pearce and Staccetti (1986 and 1990) bang-bang result for strongly symmetric equilibria. The restrictions we imposed on the structure of the stage game under consideration allow us to write players' problem in a tractable way and solve it using simple dynamic programming methods.

The expected normalized payoff $\overline{v}$ is the sum of two components; the immediate normalized expected payoff π associated with the equilibrium profile, plus, the sum of the discounted value associated with a good signal $\overline{v}$, which happens with probability $1-F$, and the lower value $\underline{v}$ if a bad signal is observed, which happen with probability F. The best SSE payoff, with correlation on some public signal, is just the largest value $\overline{v}$,

$$\overline{v} = (1-\delta^\tau)\,\pi(a^*) + \delta^\tau\left[(1-F)\,\overline{v} + F\left(\phi\overline{v} + (1-\phi)\,\underline{v}\right)\right], \tag{4.1}$$

that satisfies

$$\overline{v} \geq (1-\delta^\tau)\,\pi(a') + \delta^\tau\left[(1-F')\,\overline{v} + F'\left(\phi\overline{v} + (1-\phi)\,\underline{v}\right)\right], \tag{4.2}$$

and

$$\overline{v} \geq \underline{v} \geq \pi\left(a^N\right). \tag{4.3}$$

The first constraint guarantee that the equilibrium profile a^* is enforceable by some symmetric continuation value. The expected value of the game associated with the strongly symmetric profile, has to be at least as good as the expected value of the game associated with a potential unilateral deviation. The second set of constraint requires that the value $\overline{v}$ is individual rational and larger than $\underline{v}$, and that the value associated with the punishment stage $\underline{v}$ is credible. When this last set of constraint is meet we say that $\overline{v}$ and $\underline{v}$ are feasible.

A necessary condition for optimality requires condition (4.1) to hold with equality. Providing that δ^{τ} is large enough, for it to be the case, when required we allow players to correlate their actions on some public signal[17]. In this case a "bad" observation of the process is interpreted as a "bad signal" with probability $1-\phi$. Then ϕ is the probability with which players ignore a public signal that suggest deviating behavior. We can return to the general case without public correlation by setting $\phi = 0$.

Solving equations (4.1) and (4.2) for $\overline{v}$ and $\underline{v}$, in the general case with public correlation we obtain

$$\overline{v} = \pi\left(a^{*}\right) - \frac{F}{F' - F}\left(\pi\left(a'\right) - \pi\left(a^{*}\right)\right), \tag{4.4}$$

and

$$\underline{v} = \overline{v} - \frac{1}{1-\phi}\frac{1-\delta^{\tau}}{\delta^{\tau}}\frac{\pi\left(a'\right) - \pi\left(a^{*}\right)}{F' - F}, \tag{4.5}$$

which are expressed exclusively as functions of the parameters of the model.

Since $\pi\left(a^{*}\right) > \pi\left(a'\right)$ and $F' > F$ the second term on the RHS of (4.5) is always non-negative, the first inequality in (4.3) is then clearly satisfied. Then we simply say that a particular pair of values $\underline{v}$ and $\overline{v}$ are feasible if they satisfy the constraint $\underline{v} \geq \pi\left(a^{N}\right)$. When this conditions fails, no equilibria either than the infinite repetition of the static Nash equilibrium can be sustained, i.e. $\overline{v} = \pi\left(a^{N}\right)$.

Similarly when the length between repetitions of the stage game is uncertain, $\widetilde{\overline{v}}$ and $\widetilde{\underline{v}}$ have equivalent interpretations. They are respectively the extreme upper and the lower SSE payoffs when the repeated game is played at random moments in time. Following arguments analogous to the ones used to derive expressions (4.4) and (4.5) for the deterministic case, $\widetilde{\overline{v}}$ and $\widetilde{\underline{v}}$ can be expressed exclusively as functions of the parameters of the model, that is

$$\widetilde{\overline{v}} = \pi\left(a^{*}\right) - \frac{\widetilde{F}}{\widetilde{F}' - \widetilde{F}}\left(\pi\left(a'\right) - \pi\left(a^{*}\right)\right), \tag{4.6}$$

[17] Public correlation on some public signals is needed in the Poisson process case for condition (4.2) to bind. Instead of correlate on the public signals we could have employed strategies that ignore with certainty one or more "bad" observation of the public process, however public correlation does not present the same integer problems.

and

$$\underline{\widetilde{v}} = \widetilde{\overline{v}} - \frac{1}{1-\widetilde{\phi}} \frac{1-E(\delta^x)}{E(\delta^x)} \frac{\pi(a') - \pi(a^*)}{\widetilde{F}' - \widetilde{F}}, \tag{4.7}$$

where $\widetilde{\overline{v}} > \pi\left(a^N\right)$ is the value of the best SSE payoff with random monitoring when $\underline{\widetilde{v}} \geq \pi\left(a^N\right)$[18].

Among all the feasible punishment schemes, the largest SSE payoff is attained using the most severe punishment available[19], i.e. perpetual play of a stage game Nash equilibrium. Then the decision rule (when possible) is chosen to return the largest SSE payoff, conditioned to $\underline{v} = \pi\left(a^N\right)$ ($\underline{\widetilde{v}} = \pi\left(a^N\right)$ in the random monitoring case). Clearly when we consider public correlation on public signals, a bad observation of the public signal does not necessarily result in the immediate execution of the punishment rule.

As a consequence the two point sets $\left\{\pi\left(a^N\right), \overline{v}\right\}$ and $\left\{\pi\left(a^N\right), \widetilde{\overline{v}}\right\}$ associated with deterministic and random monitoring respectively are self-generating. The continuation values are themselves elements on these sets.

C. Efficiency Gains with Random Monitoring

In Section III it was shown that for any equilibrium and all monitoring intensities, perfect random monitoring always allows efficiency gains with respect to perfect deterministic monitoring. Even though no changes occurred on players true discount factor, their decisions where based on a larger discount factor. We call it, the expected discount factor effect, and it was the key for the obtained results.

Under imperfect public random monitoring the expected discount factor effect still present and with the same intensity. The difference is that now the uncertainty on the repetitions

---

[18] Where for the case of repeated games played at random moments in time, the probability $\widetilde{\phi}$ is denoted and given by

$$\widetilde{\phi} = 1 - \frac{1-E(\delta^x)}{E(\delta^x)} \frac{\pi(a') - \pi(a^*)}{\widetilde{F}'\left(\pi(a^*) - \pi\left(a^N\right)\right) - \widetilde{F}\left(\pi(a') - \pi\left(a^N\right)\right)}.$$

A simmilarly expression with F' and F instead of $\widetilde{F}'$ and $\widetilde{F}$ respectively and without expectation on the discount factor gives the case when time is deterministic.

[19] Since the our concern is mainly on Brownian and Poisson type processes, the assumptions is without loss in generality. The Brownian motion is a Gaussian process and the distribution of the public signals is not convex in all of its domain, perpectual punishment is then the most efficient way to provide incentives, see Porter (1983). Similarly result can be show for a Poisson process with public correlation over the public signals.

of the stage game also affect the distribution of the public signals, whether or not there was a deviation by any of the players. It is clear that $F' > F$ and $\widetilde{F}' > \widetilde{F}$ for all monitoring intensities, however we cannot establish such precise relations when comparing the likelihood ratios $\widetilde{F}'/\widetilde{F}$ and F'/F. As pointed by Abreu, Milgrom and Pearce (1991) the value of the likelihood ratio is crucial, so thus it is here the comparison between these two likelihood ratios.

The following result establish conditions under which imperfect public random monitoring allows gains in efficiency.

Theorem 3 *For a given monitoring intensity $\tau = E(x)$, the upper bound on the set of SSE payoff of an infinitely repeated game at random moments in time $\widetilde{v}$, is larger than the upper bound associated with the repeated game at deterministic moments in time $\overline{v}$;*

(i) if

$$F\widetilde{F}' > F'\widetilde{F},$$

when $\underline{v} = \pi\left(a^N\right)$ and $\underline{\widetilde{v}} = \pi\left(a^N\right)$, or

(ii) if

$$\widetilde{F}'\left(\pi\left(a^*\right) - \pi\left(a^N\right)\right) > \widetilde{F}\left(\pi\left(a'\right) - \pi\left(a^N\right)\right),$$

when $\underline{v} < \pi\left(a^N\right)$ and $\underline{\widetilde{v}} = \pi\left(a^N\right)$.

Proof. (i) When for a given monitoring intensity $\tau = E(x)$ and associated discount factors we can find decision rules such that both $\underline{v} = \pi\left(a^N\right)$ and $\underline{\widetilde{v}} = \pi\left(a^N\right)$ hold, then both $\overline{\widetilde{v}}$ and $\overline{v}$ attain values larger than $\pi\left(a^N\right)$. We need to compare the value of the best SSE payoff in each case, i.e. $\overline{\widetilde{v}} > \overline{v}$. Then the comparison between expressions (4.6) and (4.4) is equivalent to, after some arrangements,

$$\frac{F}{F' - F} > \frac{\widetilde{F}}{\widetilde{F}' - \widetilde{F}},$$

rearranging again, we obtain the desired condition

$$F\widetilde{F}' > F'\widetilde{F}.$$

(ii) When for a given monitoring intensity τ and associated discount factors we can find decision rules such that $\underline{\widetilde{v}} = \pi\left(a^N\right)$ holds, but not $\underline{v} = \pi\left(a^N\right)$. Then $\overline{\widetilde{v}} > \pi\left(a^N\right)$ while $\overline{v} = \pi\left(a^N\right)$. With imperfect public random monitoring we can sustain equilibria payoffs above the static Nash but not with perfect public deterministic monitoring. Then the condition is that $\widetilde{\widetilde{v}} > \pi^N$, or equivalently

$$\pi(a^*) - \frac{\widetilde{F}}{\widetilde{F}' - \widetilde{F}} (\pi(a') - \pi(a^*)) > \pi^N.$$

After some rearrangements we obtain the desired inequality

$$\widetilde{F}' \left(\pi(a^*) - \pi\left(a^N\right)\right) > \widetilde{F} \left(\pi(a') - \pi\left(a^N\right)\right).$$

■

The superiority of random monitoring relative to deterministic monitoring shown in Section III does not generalize in such an absolute way with imperfect public monitoring[20]. Theorem 3 tell us when such superior is possible. While in point (ii) only random monitoring achieve equilibria payoffs above the static Nash, in (i) we can sustain payoffs above the static Nash with both monitoring technologies, the issue is then to find conditions where random monitoring is superior to deterministic monitoring. Outside the cases stated in Theorem 3, either we say that random monitoring is equivalent to deterministic monitoring or that the former is less efficient.

The idea is that trough the expected discount factor effect the future has become more important, but it might also has become more uncertain due to the effect on the distribution of the public signals. A clear answer depends on the relation between different likelihood ratios.

V. NUMERICAL EXAMPLES - THE PRISONERS' DILEMMA

In this section we provide a numerical example of the value of the best SSE payoff when monitoring is random and when it is deterministic. For the case where monitoring is imperfect but public we study when efficiency gains are possible under random monitoring, on the context of the work of Abreu, Milgrom and Pearce (1991), Fudenberg and Levine (2007), Sannikov and Skrzypacz (2007) and Osório-Costa (2008). These papers main concern

[20] Recall that we say that random monitoring (deterministic monitoring) is superior in absolute terms to deterministic monitoring (random monitoring) when for a given parameterization of the problem, it achieves larger payoffs for all monitoring intensities.

	1	0
1	$2, 2$	$-1, 3$
0	$3, -1$	$0, 0$

TABLE I: A Prisioners' Dilemma - Stage Game Payoffs

is on the limit case ($\tau \to 0$), however, we extend their approach for more general monitoring frequencies.

To illustrate it, we consider that at each moment in time $\tau_0, \tau_1, \tau_2, \ldots$, players repeatedly play the prisoners' dilemma stage game of Table I.

Provide 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.

A. Abreu, Milgrom and Pearce (1991) Approach

In their paper, Abreu, Milgrom and Pearce (1991) tailor their approach to answer the limit case questions ($\tau \to 0$). Here we are also interested in other monitoring intensities. We simplify their setting by assuming that there are only two possible public signals at any time τ; the case where there is no observed events, and the case where one or more events are observed[21]. This allows us to work with the exponential distribution rather than with a sum of Poisson distributions, equal to the number of relevant events chosen in some optimal way[22].

Players, after a bad observation of the public signal, condition their actions (in particular the perpetual punishment decision) on some publicly observed random variable. There is a probability ϕ (or $\widetilde{\phi}$) with which players ignore a bad signal, and makes condition (4.2) hold with equality.

[21] The interpretation given to each of these two cases depend on whether we are considering the "bad news" or the "good news" case. More about it below.

[22] Even in the latter case we were not able to solve the integer problem refered in Footnote 17. The resourse to public correlation is still required in order for condition (4.2) to hold with equality.

1. The bad news case $\mu > \lambda > 0$

In this case one or more observed movements in the public process are interpreted as bad news with probability $1 - \phi$ (or $1 - \widetilde{\phi}$). Abreu, Milgrom and Pearce (1991) have shown that in the limit equilibrium payoffs above the static Nash but not efficient can be sustained depending on the value of the parameters μ and λ with respect to r[23]. In this case the value of the best SSE payoffs improve monotonically with the monitoring intensity[24].

Independently of the monitoring frequency, smaller the ratio λ/μ larger the payoffs that can the achieved, see Figure 1. For given parameterization of the problem, i.e. λ/μ sufficiently small and τ sufficiently large we observe random monitoring being superior to deterministic monitoring in the sense of part (i) of Theorem 3 while $\tau \in (\tau_-, \tau_+]$ and superior in the sense of part (ii) of Theorem 3 while $\tau \in (\tau_+, \tau_*]$, see the upper pair of curves in Figure 1. Deterministic monitoring is then superior for small values of τ. This result is explained by the expected discount factor effect that tends to be stronger for lower frequencies of monitoring.

Random monitoring can only be superior in absolute terms when we let the ratio $\lambda/\mu \to 0$.

When the ratio λ/μ is sufficiently large deterministic monitoring may be superior in a absolute terms[25]. The low pair of curves in Figure 1 illustrate this possibility.

Clearly the value of (4.4) and (4.6) do not vary with the discount rate r. However, a larger value of r (lower discounting) reduces the set of monitoring frequencies that can sustain payoffs above the static Nash. This is true for random and deterministic monitoring.

2. The good news case $\lambda > \mu > 0$

In this case, it is the lack of observed movements in the public process that are interpreted as "bad signals"[26]. Abreu, Milgrom and Pearce (1991) have shown that when τ gets small "bad signals" become very likely (infinitely likely in the limit) and the only equilibria that

[23] Where μ and λ denote the bad news arrival rates, for the case where there is a deviation and the case without any deviation respectively.

[24] Osório-Costa (2008) reports a similar monotonic result when the public signals are Brownian. However, his result is efficient in the limit.

[25] This happen when for any monitoring intensity, ϕ and $\widetilde{\phi}$ never cross and $\phi > \widetilde{\phi}$.

[26] Where μ and λ denote the intensity of good news arrival, for the case where there is a deviation and the case without any deviation respectively.

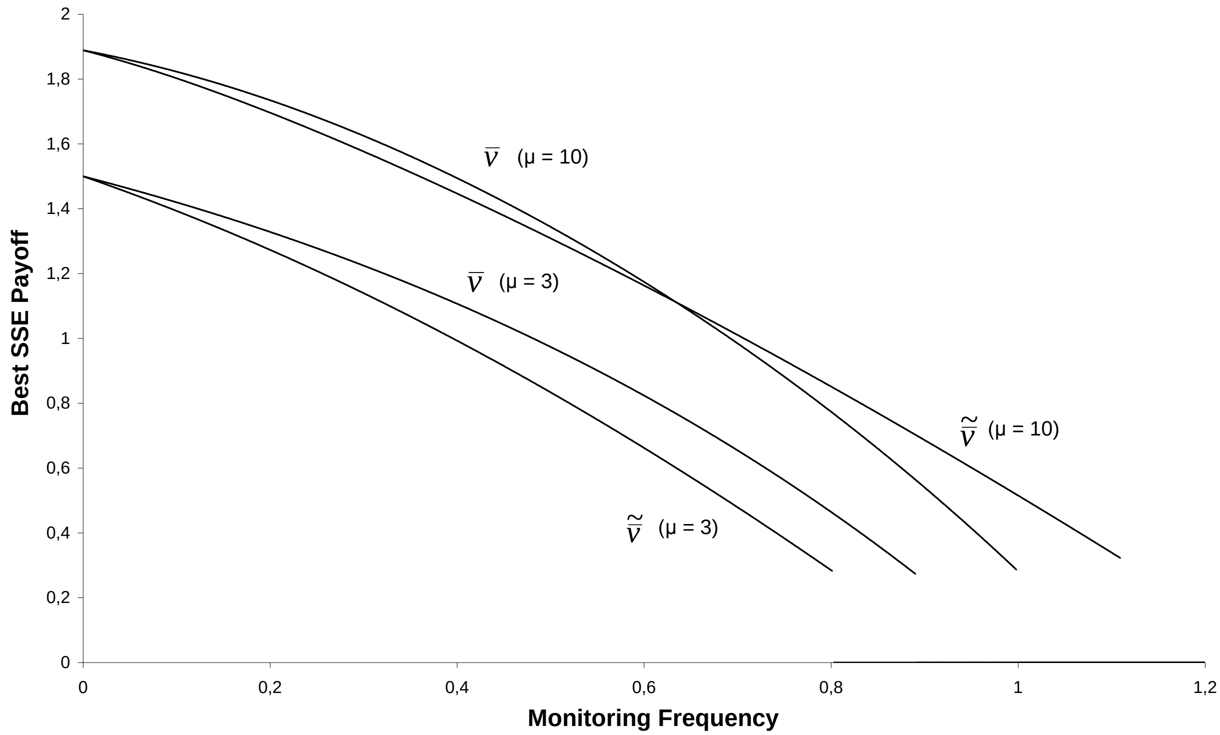

FIG. 1: The Bad News Case, $\lambda = 1$ and $r = 0.1$.

can be sustained is the infinite repetition of the stage game static Nash equilibria. Figure 2 illustrate the degeneracy of the best SSE payoffs for high monitoring intensities.

In the "good news" case random monitoring does not improving on the best SSE payoff in the sense of Theorem 3 part (i). Deterministic monitoring is in general superior in absolute terms, with exception for the cases where the ratio λ/μ is sufficiently large, where we can observe random monitoring being superior in accordance with Theorem 3 part (ii)[27]. It can be seen from Figure 2 when $\lambda = 5$. In this way random repeated monitoring can at most enlarge the spectrum of monitoring intensities in which payoffs above the static Nash can be sustained.

[27] This happens when ϕ and $\widetilde{\phi}$ cross once. Notice also that, here in the "good new case", a larger ratio λ/μ has positive effects on the value of the best SSE.

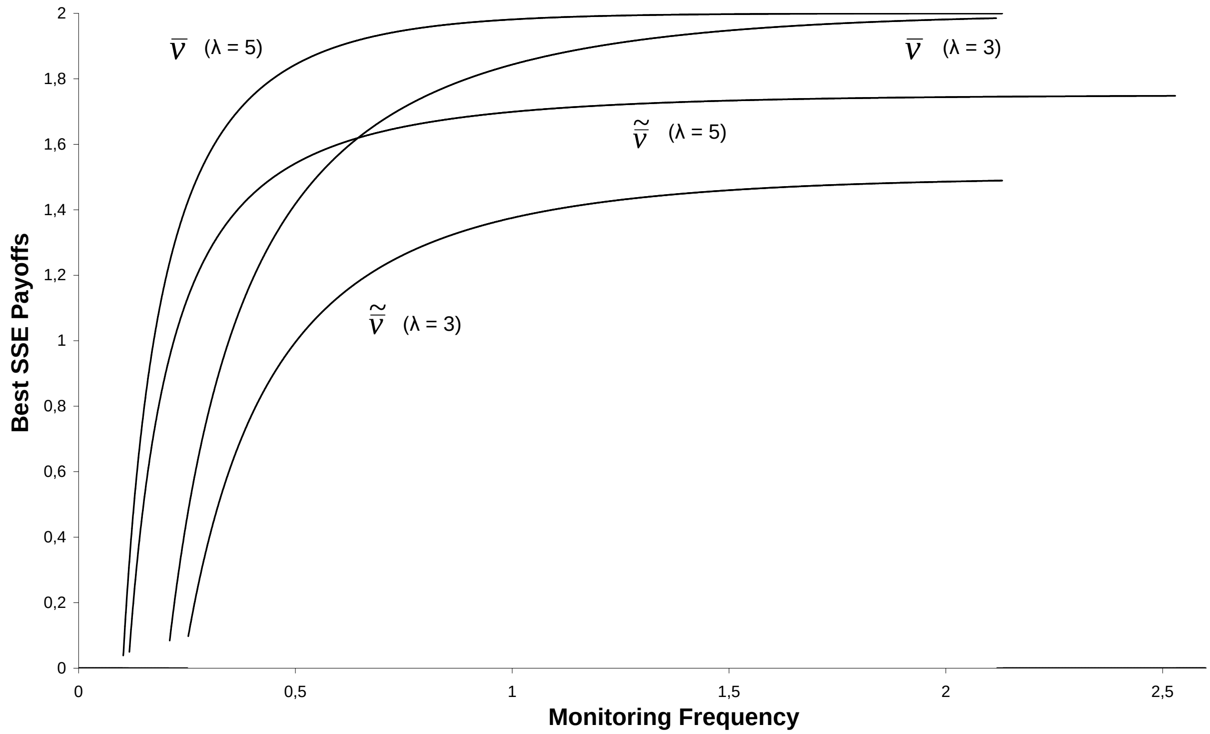

FIG. 2: The Good News Case, $\mu = 1$ and $r = 0.1$.

B. Fudenberg and Levine (2007)/Sannikov and Skrzypacz (2007) Approach

We explore some equivalence between Fudenberg and Levine (2007) and Sannikov and Skrzypacz (2007) frequent monitoring approaches, for that reason we treat them togheter. As reported in these papers, we observe a degeneracy of the set of SSE for high monitoring intensities. A result parallel to the one obtained by Abreu, Milgrom and Pearce (1991) in the good news case.

Away from the limit, a large value of σ or r, have negative effects on the payoffs[28], see Figure 3.

While in other approaches, random monitoring when superior, it is for monitoring intensities $\tau > \tau_{-}$, here it happens the opposite. For sufficiently low σ, we may find an interval of high monitoring intensities (low τ) where random monitoring is superior in the sense of Theorem 3 part (ii)[29], then as τ increases there is a subsequent interval where random mon-

[28] For sufficiently large values of σ no equilibria can be sustained for any monitoring intensity.

[29] Before that interval, i.e. for even higher monitoring intensities, both random and deterministic monitoring

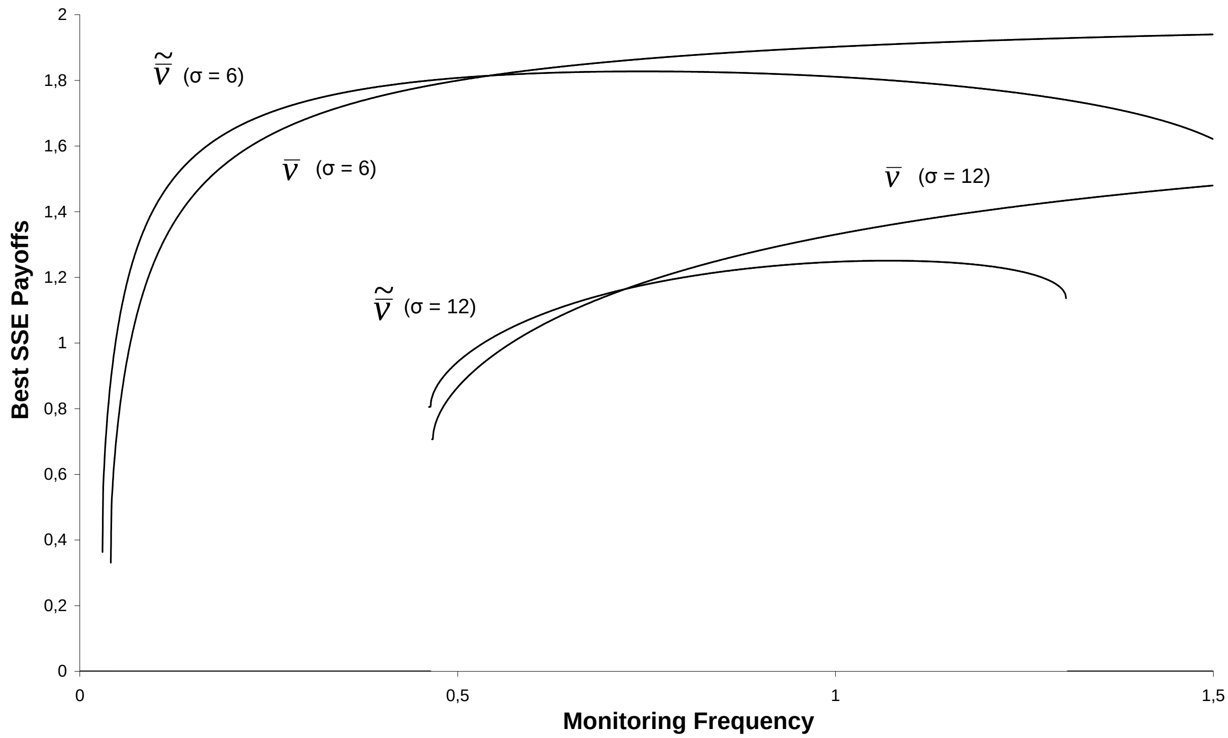

FIG. 3: Fudenberg and Levine/Sannikov and Skrzypacz Frequent Monitoring Approach, $r = 0.1$.

itoring dominates in the sense of Theorem 3 part (i). After that deterministic monitoring becomes superior. This point can clearly be seen for the upper pair of curves in Figure 3.

C. Osório-Costa (2008) Approach

In this setting a low observations of the public signal is interpreted as a "bad signals". The decision rule that separates "good" from "bad" signals is explained in great detail in Osório-Costa (2008). In this paper efficient results are obtained in the limit and payoffs improve monotonically with the monitoring intensity. The monotonicity result is similar to the one described above for the Abreu, Milgrom and Pearce (1991) bad news case. (See Figures 1 and **??**)

(...figure missing...)

Independently of the monitoring frequency, smaller the uncertainty parameter σ larger the

return degenerated payoffs.

payoffs that can the achieved. We observe random monitoring being superior to deterministic monitoring according to Theorem 3 part (i) while $\tau \in (\tau_-, \tau_+]$ and superior in accordance with Theorem 3 part (ii) while $\tau \in (\tau_+, \tau_*]$[30]. As we increase σ, random monitoring stops to dominate in the sense of part (i) of Theorem 3 and dominates exclusively in the sense of part (ii). The upper and midle pairs of curves in Figure **??** illustrate this idea. When σ gets sufficiently large deterministic monitoring dominates in absolute terms. (See the low pair of curves in Figure **??**)

If there is an interval of monitoring intensities for which random monitoring is superior, that interval does not vanish with varying r, however, its measure does. Low values of r have positive effects on the payoffs.

APPENDIX A:

[30] In this case the optimal decision rule for the deterministic and the random monitoring case cross twice. Also, independently of σ, deterministic monitoring tends to be superior for small values of τ.

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. Faingold, E. (2006). "Building a Reputation under Frequent Decisions," mimeo.

5. Faingold, E., Y. Sannikov (2007). "Reputation Effects and Equilibrium Degeneracy in Continuous-Time Games," mimeo.

6. Friedman, J. W. (1971). "A Noncooperative Equilibrium for Supergames," Review of Economic Studies, 38, 1-12.

7. Fudenberg, D. and D. Levine (1994) "Efficiency and Observability with Long-Run and Short-Run Players," Journal of Economic Theory, 62(1), 103-135.

8. Fudenberg, D. and D. Levine (2007) "Continuous Time Models of Repeated Games with Imperfect Public Monitoring." Review of Economic Dynamics, 10(2), 173-192.

9. Fudenberg, D. and D. Levine (2008) "Repeated Games with Frequent Signals." Quarterly Journal of Economics, forthcoming.

10. Fudenberg, D., D. Levine and E. Maskin (1994). "The Folk Theorem with Imperfect Public Information." Econometrica, 62, 997-1040.

11. Fudenberg, D. and E. Maskin (1986). "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," Econometrica, 54, 533-554.

12. Fudenberg, D. and J. Tirole (1991) Game Theory, MIT Press, Cambridge, MA.

13. Green, E. and R. Porter (1984). "Noncooperative Collusion under Imperfect Price Information." Econometrica, 52, 87-100.

14. Mailath, G. and L. Samuelson (2006) Repeated Games and Reputations: Long-run Relationships. Oxford University Press, New York.

15. Osório-Costa, A. M. (2008) "Frequent Monitoring in Repeated Games under Brownian Uncertainty," mimeo.

16. Porter, R. (1983). "Optimal Cartel Trigger Price Strategies." Journal of Economic Theory, 29, 313-338.

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. Sannikov, Y. (2007). "Games with Imperfectly Observable Actions in Continuous Time," Econometrica, 75, 1285–1329.

19. Sannikov, Y. and A. Skrzypacz (2007a) "Impossibility of Collusion under Imperfect Monitoring with Flexible Production," American Economic Review, 97, 1794-1823.

20. Sannikov, Y. and A. Skrzypacz (2007b) "The Role of Information in Repeated Games with Frequent Actions," mimeo.

21. Stahl, II, D. O. (1991) "The Graph of Prisoners' Dilemma Supergame Payoffs as a Function of the Discount Factor," Games and Economic Behavior, 3, 368-384.