Brownian Signals: Information Quality, Quantity and Timing in Repeated Games.*

António Osório
Universität Rovira i Virgili and CREIP†

March 14, 2017

Abstract

This paper examines different Brownian information structures over varying time intervals. We focus on the non-limit case, and on the trade-offs between information quality and quantity when making a decision whether to cooperate or defect in a prisoners' dilemma game. In the best-case scenario, the information quality gains are strong enough so that agents can substitute information quantity with information quality. In the second best-case scenario, the information quality gains are weak and must be compensated for with additional information quantity. In this case, information quality improves but not quickly enough to dispense with the use of information quantity. For sufficiently large time intervals, information degrades and monitoring becomes mostly based on information quantity. The results depend crucially on the particular information structure and on the rate at which information quality improves or decays with respect to the discounting incentives.

JEL: C73, D82, D86.

KEYWORDS: Repeated Games; Frequent Monitoring; Information Quantity; Information Quality.

1 Introduction

Coordination and cooperation depend crucially on the frequency of interaction between the involved parties. This fact affects the agents' behavior and the stability of these relationships.

*I would like to thank David Rahman, Juan Pablo Rincón-Zapatero and Ricardo Ribeiro, as well as several seminars and congresses participants for helpful comments and discussions. I am also grateful to the editor and two anonymous referees for their constructive comments. All remaining errors are mine.

†Department of Economics; Reus (Barcelona); E-mail: superacosta@hotmail.com.

In the simplest setting, with perfect information in which actions and monitoring occur simultaneously at the same frequency, smaller time intervals facilitate cooperation by making agents more patient, and as a consequence the deviation incentives decrease. Conversely, large time intervals render cooperation more difficult because agents are less patient and the incentives to deviate increase.

However, with imperfect monitoring, varying the time interval leads to different results, because the length of the time interval between actions affects information quality (Abreu et al., 1991; Fudenberg and Levine, 2007, 2009; Osório, 2012; Sannikov and Skrzypacz, 2007). For instance, in small time intervals the deviations incentives are weak, but information quality might be so low that coordination and cooperation are impossible. However, we may also have situations in which information quality is high enough to make coordination and cooperation possible.

In this context, information quality is critical to provide incentives and to determine the equilibrium payoffs. Information quality depends crucially on how actions feedback into signals - the *information structure.* Therefore, it matters whether actions affect the *fundamental value* of a given variable or the *variance* (Fudenberg and Levine, 2007, 2009). For instance, at moments of market stability, when an OPEC (Organization of the Petroleum Exporting Countries) member country unilaterally decides to increase its oil supply, the oil price (i.e. the fundamental value) is likely to decrease. However, at moments of market instability and conflict, the same action may have an insignificant effect on the fundamental value, but may instead induce additional market volatility (i.e. the variance). Therefore, in the former case, the parties involved should monitor the fundamental value, while in the latter case the parties involved should monitor the variance.

Is it more difficult for the agents to coordinate and cooperate when actions affect the fundamental value or when they affect the variance? How do coordination and cooperation depend on the length of the time interval between actions and monitoring activities?

This paper attempts to answer these questions. We examine how different information structures depend on the time interval between actions and monitoring activities, which are assumed to have the same frequency. We examine how information **quality** and **quantity** vary with the length of the time interval between actions, and how they are used to provide incentives (Mirrlees, 1974). *Information quality* refers to the signals' precision, while *information quantity* refers to the number of signals used to enforce cooperation. In this context, the starting point is to acknowledge that the provision of incentives depends crucially on how information quality improves or decays in relation to the deviation incentives (due to discounting) and its actual level.

The aim of this paper is not to present a general theory, which would be very complex, but rather to illustrate some relevant aspects related to the use of information quantity and quality in the provision of incentives. Contrary to the existing literature (Abreu et al., 1991; Fudenberg and Levine, 2007; Sannikov and Skrzypacz, 2007), which has focused exclusively on the limit case (i.e. when the length of the time interval between actions tends towards zero), we consider the non-limit case (i.e. any length of the time interval other than zero). While the limit case is interesting for theoretical and tractable reasons, the non-limit case is more realistic, although it requires the use of numerical and computational methods. Nonetheless, gaining an understanding of the existing information trade-offs justifies the use of these tools.

Abreu et al. (1991) were the first to initiate the literature on frequent monitoring. In a setting with Poisson signals, they have shown that cooperation can be sustained at the limit when the observations of the process represent bad news, but not when they represent good news. Nonetheless, in both cases, the best payoffs are not obtained at the limit (see also Osório, 2015). More recently, Sannikov (2007) and Faingold and Sannikov (2007) have renewed the interest in frequent monitoring by considering continuous time methods for studying repeated games. Simultaneously, by studying the limit of the discrete time games with Brownian signals, Fudenberg and Levine (2007) have considered several information structures (see also Fudenberg and Levine 2009; Osório 2012; Sannikov and Skrzypacz 2007, 2010).[1] They found that full efficient results are possible in the limit case if deviations increase the noise of the process. In this respect, the present paper can be seen as a non-limit extension to the Fudenberg and Levine's (2007) limit analysis, and we hope it will help to provide a better understanding of the strengths and the weaknesses of each of the information structures discussed in the literature.

The provision of incentives depends crucially on the rate to which information quality improves or decays relative to the deviation incentives (due to discounting) and the level of information quality. In the ideal scenario, we would like to have simultaneously the highest information quality and the lowest information quantity. However, with imperfect information this objective is not attainable because in general information quality is limited.

In this context, it is natural to think that information quantity and quality are **substitutes** because *ceteris paribus*, higher information quality, at the margin, reduces the quantity

[1] A notable variation of the original model is Fudenberg and Olszewski (2011), who study repeated games with stochastic asynchronous monitoring. Outside the limit case, Fudenberg et al. (2014) show that if players wait long enough, then it is likely that everybody will have observed the same signal and a folk theorem may be possible. Kamada and Kominers (2010) also consider a time-varying information structure. However, their argument and information notions are different; see also Kandori (1992).

of information needed, and vice versa. However, we found that substitution occurs only if the information quality gains or losses are strong enough to dispense (positive substitution) or require (negative substitution), respectively, the use of additional information quantity. The reason is that variations in the length of the time interval between actions affect not only the quality of information but also the deviation incentives. Consequently, there are no *ceteris paribus* situations. We also found that when the information quality gains or losses are weak, information quantity must be used to **compensate** for the weaknesses of information quality in providing incentives. In this sense, compensation is a weaker form of substitution in which information quantity and quality are jointly used in the same direction.

Since information quality depends crucially on how actions affect the distribution of the Brownian public signals, we have considered three different information structures:

In section 4.1 we consider the case where **actions affect the drift** of the process (i.e. the fundamental value). This information structure captures the most commonly observed situations in real life. For instance, when a worker reduces its effort, the output is likely to decrease, or when a firm increases its supply, the market price is likely to decrease. With this information structure, information quality tends to improve with the time interval, except for large time intervals. The intuitive argument is that reliable inference about the drift requires a sufficiently large time interval.[2] This observation is especially true for small time intervals because increasing the length of the time interval leads to strong improvements in information quality. Consequently, information quality substitutes information quantity. However, as the time interval increases, the marginal improvement in information quality diminishes to the point where it is not enough to decrease the number of signals used to provide incentives. In this case, information quality needs to be complemented with information quantity. Lastly, for large time intervals, and before the equilibrium degenerates, signals become extremely noisy. In order to provide players with incentives, the falling information quality is substituted by information quantity. This pattern, observed for large time intervals and before the equilibrium degenerates, is common to all information structures.

However, if **deviations affect the noise of the process** (i.e. the variance), information quality always decays with the length of the time interval between actions. Contrary to the case in which actions affect the drift, the best inference about the noise parameter is obtained in the smallest time intervals (Prakasa Rao, 1999).

In section 4.2 we consider the case where **deviations increase the noise** of the process. This information structure captures situations in which the observation of extreme events,

[2]With this information structure Sannikov and Skrzypacz (2007) have shown that at the limit the equilibrium degenerates. A different approach and implications are shown in Osório (2012).

	$e_{2t} = 1$	$e_{2t} = 0$
$e_{1t} = 1$	g, g	$-u, g + u$
$e_{1t} = 0$	$g + u, -u$	$0, 0$

Table 1: The Prisoners' Dilemma Stage Game Payoffs.

such as high market volatility or high sales variation, are associated with misbehavior or lack of effort. With this information structure, for small time intervals, information quality is high but decays with the length of the time interval between actions, at a rate that compensates for the increase in deviation incentives.[3] Therefore, while the information quality is sufficiently high, it is possible to reduce the quantity of signals. For sufficiently large time intervals and before the equilibrium degenerates, the deviation incentives increase and the signals become extremely noisy (information quality is low); the falling information quality needs to be offset by increasing the quantity of signals.

In section 4.3 we consider the case where **deviations decrease the noise** of the process. This information structure captures situations in which the observation of stable events, such as unchanging profits, are associated with a lack of effort or commitment. With this information structure, we found that information quality is low and decays with the length of the time interval.[4] Efficient monitoring substitutes the falling information quality with information quantity.

Section 5 concludes with additional comments, identifies possible information patterns and discusses avenues for further research.

2 The model

We explore frequent monitoring in a simple prisoners' dilemma-type game with two long-run players $i = 1, 2$. At moments in time $t = 0, \tau, 2\tau, ...$ players simultaneously choose whether or not to provide effort, i.e. $e_{it} = 1$ or $e_{it} = 0$, respectively, where τ denotes the length of the time interval between actions and monitoring, which have the same frequency. We consider the stage game payoffs in Table 1, with $g > u > 0$ implying that no effort is a dominant strategy for both players.[5]

[3] With this information structure Fudenberg and Levine (2007) have shown that full efficiency is possible at the limit.

[4] With this information structure Fudenberg and Levine (2007) have shown that cooperation is possible at the limit, but not full efficient payoffs.

[5] We have restricted our analysis to the simplest setting. The results obtained in the following sections generalize to other discount factors, payoff structures and games.

In the subsequent period $t+\tau$, an imperfect public signal $s_{t+\tau}$ about the players' actions, taken in period t, is publicly observed by both players. Consequently players use the public signal as a coordinating device. The public signal is generated by the *arithmetic Brownian motion* (ABM) process:

$$s_{t+\tau} - s_t = \mu_t \tau + \sigma_t \int_t^{t+\tau} dW_x, \; with \; W_t = 0 \; and \; t = 0, \tau, 2\tau, ..., \tag{1}$$

where $\{W_x; x \geq 0\}$ is the standard Brownian motion, and $s_t = 0$ is the initial value. The parameters μ_t and σ_t depend on the effort profile $e_t = (e_{1t}, e_{2t})$ and control, respectively, the drift (i.e. the fundamental value) and the noise components (i.e. the variance) of the process at time t. In sections 4.1, 4.2 and 4.3 we consider different assumptions regarding how actions affect the drift and the noise components of the process (1).

Let $S_{t+\tau}$ denote the set of signals $s_{t+\tau}$ that suggest deviation or no effort. These are observations of the process that fall inside a certain region bounded by one or more thresholds (i.e. $\underline{s}_{t+\tau}$ and/or $\overline{s}_{t+\tau}$). In equilibrium these thresholds will depend not only on how actions affect the drift and the noise of the process (1), but also on the length of the time interval between actions. Consequently, the number of signals that are inside the set $S_{t+\tau}$ varies with the time interval.

Given the effort profile $e_t = (e_{1t}, e_{2t})$, the probability of punishment - the probability that the state of the public process (1) falls inside the set $S_{t+\tau}$ - is Gaussian distributed and given by:

$$p(e_{1t}, e_{2t}) = \Pr(s_{t+\tau} \in S_{t+\tau} | e_t) = \int_{S_{t+\tau}} \exp(-(x-\mu_t\tau)^2/(2\sigma_t^2\tau))/(2\pi\sigma_t^2\tau)^{1/2} dx. \tag{2}$$

Similarly, the no punishment probability - the probability that the state of the public process (1) falls outside the set $S_{t+\tau}$ - is given by: $1 - p(e_{1t}, e_{2t})$. In this case, signals are interpreted as suggesting cooperation or mutual effort. Clearly, the higher the effort, the higher is this probability, and vice versa.

The common discount factor is $\delta \equiv e^{-r(t+\tau)}$, where $r \in (0, \infty)$ denotes the discount rate.

We look at strategy profiles that form a strong symmetric *perfect public equilibrium* (PPE).[6] Since the equilibrium is stationary, the particular moment in time t is irrelevant. For that reason, we will set $t = 0$ and refer only to the length of the time interval τ.

[6] A strategy is public if it depends only on the public history (of signals) and not on the private history (of signals and of individual efforts). 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 PPE.

3 Equilibrium and Information Properties

In order to understand how information quantity and quality are simultaneously used to enforce cooperation, we compute the value of the best strong symmetric equilibrium and the associated incentive compatible constraint. In addition, in this section, we define the meaning of information quantity and quality.

Equilibrium payoffs and equilibrium conditions

Our aim is to sustain the infinite repetition of the mutual effort profile $e_0 = (1,1)$. In this setting, the best equilibrium payoff is attained by a "grim" strategy that prescribes effort so long as the public signal is outside the punishment region.

The normalized discounted value of the relationship when both players provide effort, denoted as $\overline{v}$, is given by (which, by symmetry, is the same for both players):

$$\overline{v} = (1-\delta)\, g + \delta \left[(1 - p(1,1))\overline{v} + p(1,1)\underline{v} \right], \tag{3}$$

where $\underline{v}$ is the normalized discounted value of the relationship in case of punishment, and $p(1,1)$ is the probability of [mistaken] punishment. In expression (3), players receive the mutual effort payoff g, plus a discounted expectation over the expected values $\overline{v}$ and $\underline{v}$, that are associated with the two types of signals that might be observed (i.e. signals suggesting cooperation and defection, respectively).

Simultaneously, the equilibrium must be self-enforceable with respect to unilateral deviations, i.e. the profiles $e_0 = (1,0)$ and $e_0 = (0,1)$, which are the same by symmetry. Therefore, the normalized discounted value of the relationship, when both players provide effort, has to be at least as good as the normalized discounted value of the relationship in case of deviation, i.e.:

$$\overline{v} \geq (1-\delta)\,(g+u) + \delta \left[(1 - p(1,0))\, \overline{v} + p(1,0)\underline{v} \right], \tag{4}$$

where $p(1,0)$ is the probability of punishment in case of deviation. On the right-hand side of inequality (4), the deviator receives the instantaneous payoff $g + u$, plus a discounted expectation over the expected values $\overline{v}$ and $\underline{v}$, that are associated with the two types of signals that might be observed (i.e. signals suggesting cooperation and defection, respectively). In order to minimize the probability of mistaken punishment, in equilibrium inequality (4) must hold with equality.

Lastly, since the distribution of public signals is not convex, optimally requires an infinite punishment length (Mirrlees, 1974; Porter, 1983), therefore, $\underline{v} = 0$.

To find the expression that characterizes the highest equilibrium payoff, we solve the dynamic programming problem composed by expressions (3), (4) and $\underline{v} = 0$. We apply the bang-bang result of Abreu et al. (1986, 1990) to compute the best strongly symmetric equilibrium payoff, which is given by:

$$\overline{v} = g - u \times p(1,1)/(p(1,0) - p(1,1)) = g - u/(l_\tau - 1), \tag{5}$$

where $l_\tau = p(1,0)/p(1,1)$ is the likelihood ratio of correct punishment (with respect to mistaken punishment). In addition, in order for mutual effort to be played in equilibrium, expression (5) must be non-negative, i.e. $\overline{v} \geq 0$, and the following incentive compatible condition must be satisfied:

$$g/u = (1 - \delta(1 - p\,(1,1)))/(\delta(p\,(1,0) - p\,(1,1))). \tag{6}$$

Otherwise, the mutual effort equilibrium fails to exist, i.e. the *equilibrium degenerates* - players cannot sustain an equilibrium other than the infinite repetition of the stage game Nash equilibrium $e_t = (0,0)$ in Table 1.[7]

In equilibrium, condition (6) will establish how information quality and quantity are used to enforce cooperation while maximizing the expected payoff (5).

Remark (equilibrium computation) *The correct and mistaken punishment probabilities,* $p\,(1,0)$ *and* $p\,(1,1)$, *respectively, depend on the same set of signals* S_τ, *which is bounded by one or more thresholds, i.e.* $\underline{s}_\tau$ *and/or* $\overline{s}_\tau$ *(depending on the information structure). Consequently, the equilibrium obtained is in "threshold strategies". In other words, we search for the threshold root/s that make condition (6) holds with equality, and simultaneously maximize the value in Expression (5).*[8] *For this reason, since the punishment probabilities* $p\,(1,0)$ *and* $p\,(1,1)$ *are cumulative Gaussian distribution functions (see Expression (2)), the computation of the equilibrium thresholds requires the use of numerical and computational methods. The threshold root or roots establish the interval of integration of the cumulative Gaussian distribution.*[9]

[7] A more detailed derivation of Expressions (5) and (6) can be found in Abreu et al. (1991) or Fudenberg and Levine (2007), among others.

[8] Sannikov and Skrzypacz (2007) and Fudenberg and Levine (2007) have shown that the best equilibrium is obtained with a "grim" strategy that prescribes mutual effort as long as the public signal falls within the region bounded by some critical threshold/s. Once these thresholds are crossed, i.e. the public signal falls in the region defined by the set S_τ, the punishment stage is initiated.

[9] The results obtained are particularly robust. The predictions of the model do not depend significantly on the value of the parameters. Nonetheless, clearly, the greater the difference between these parameters (i.e. greater the impact of deviations in the drift or in the noise components of the process), the easier it is to see the reported effects. Figures 1a, 1b, 2a, 2b, 3a and 3b are sufficiently representative of the information

Information Quantity and Quality

We shall now define the meaning of information quantity and quality in our setting.

Definition 1 (information quantity) *Information quantity is determined by the size of the set S_τ. We say that information quantity increases if the size of the set S_τ increases, and decreases if the size of the set S_τ decreases.*

Information quantity is measured by the number of signals suggesting defection that are used to monitor the players' actions, and consequently to enforce cooperation. The optimal choice of $\underline{s}_\tau$ and/or $\overline{s}_\tau$, that satisfies $\overline{v} \geq 0$ and condition (6), determines the size of the set S_τ.

Definition 2 (information quality) *Information quality is determined by the likelihood ratio $l_\tau = p(1,0)/p(1,1)$. We say that information quality increases if the likelihood ratio l_τ increases, and decreases if the likelihood ratio l_τ decreases.*

Intuitively, information quality depends positively on the probability of correct punishment $p(1,0)$, where deviations are correctly punished, and depends negatively on the probability of mistaken punishment $p(1,1)$, where no deviations are incorrectly punished (Mirrlees, 1974). These probabilities are positively related by their common dependence on the set S_τ. Consequently, information quantity and quality are necessarily linked.

Note also the positive and monotonic relationships between information quality and payoffs in Expression (5).

In the ideal scenario, we would like to have the highest possible information quality, but using the lowest possible number of signals (information quantity). However, under imperfect information, this goal is not possible because information quality is always limited and information quantity is always needed to provide incentives. Consequently, information quality and quantity must always be used simultaneously to provide incentives. In some cases information quality and information quantity are substitutes, while in other cases information quality and information quantity need to compensate for each other.

4 Information Structures

We now consider three different information structures that have been discussed in the literature. The intuition and discussion is presented for an increasing time interval. The

structures considered in the present paper.

reader is free to consider the opposite exercise. For instance, if for an increasing time interval we have a situation of positive substitution, in the opposite direction, i.e. for a decreasing time interval, we have a situation of negative substitution, and vice versa.

The crucial point here is to understand how much information quality improves or decays relatively to the deviation incentives due to discounting, and the actual levels of information quality. Strong/weak information quality gains/losses are therefore always relatively to the deviation incentives and the level of information quality.

4.1 Information structure: deviations affect the drift

With this information structure, a deviation affects the drift μ_0 of the process (1). This is probably the most common structure in real life situations. For example, when a worker reduces his or her effort, output is likely to decrease, or when a firm increases the supply, the market price is likely to decrease.

In our context, we can capture these types of situations in several ways. For instance, we can have $\mu_0 = c(e_{10} + e_{20})$ with $\sigma_0 = \sigma > 0$ constant, where $c > 0$ is a parameter. In this case, a unilateral deviation reduces the drift of the process from $c(2)$ to $c(1)$, and increases both the mistaken and the correct probabilities of punishment, $p(1, 1)$ and $p(1, 0)$, respectively.

Under this information structure, Sannikov and Skrzypacz (2007) and Fudenberg and Levine (2007) have shown that one-sided threshold strategies are optimal for detecting deviations. In the case where a deviation reduces the drift of the process, players employ a common one-sided threshold $\underline{s}_\tau$ to distinguish observations suggesting cooperation, i.e. $\{s_\tau > \underline{s}_\tau\}$, from observations suggesting defection, i.e. $S_\tau = \{s_\tau \leq \underline{s}_\tau\}$.[10] Since a deviation reduces the drift of the process, with lower value signals becoming more likely, optimal monitoring establishes a region composed of lower value signals as the one that suggests deviation.

Figure 1a provides an illustration. The dark green line denotes the optimal threshold strategy for varying time interval τ. The light green shaded region corresponds to the set S_τ for varying time interval τ, i.e. the measure of information quantity (Definition 1). In this context, as the length of the time interval between actions varies, the optimal threshold strategy, and consequently, the number of signals used to sustain cooperation rises.

[10] In the case where a deviation increases the drift of the process, the common one-sided threshold $\overline{s}_\tau$, which is employed to distinguish the observations that suggest mutual effort $\{s_\tau < \overline{s}_\tau\}$ from the observations that suggest deviation $S_\tau = \{s_\tau \geq \overline{s}_\tau\}$, has a symmetric intuition.

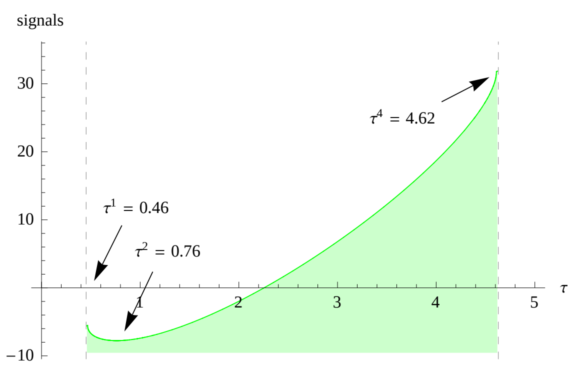

(a) Information quantity.

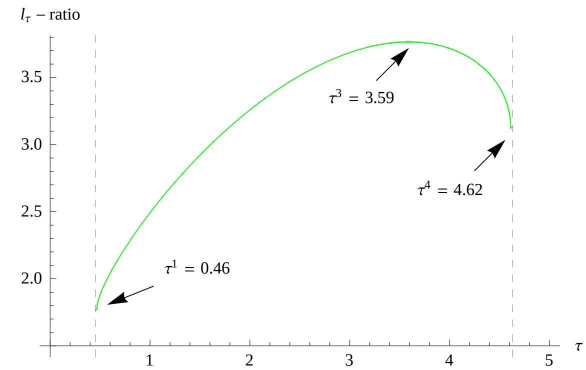

(b) Information quality.

Figure 1: **Deviation decreases the drift -** (a) Information quantity, measured by the set S_τ (light green shaded area), and the optimal one-sided threshold strategy for varying time interval τ (dark green lines). (b) Information quality, measured by the likelihood ratio. The cooperation equilibrium ceases to exist for τ below $\tau^1 = 0.46$ and above $\tau^4 = 4.62$.

Figure 1b provides a numerical illustration of how information quality measured by the likelihood ratio (Definition 2) varies with the length of the time interval τ.

The following result shows the efficient use of information quantity and quality in providing incentives under the information structure in this subsection.

Proposition 3 *When deviations affect the drift, as the length of the time interval between actions increases:*

(∅) $0 < \tau \le \tau^1$: *the equilibrium degenerates,*

(i) $\tau^1 < \tau \le \tau^2$: *the strong gains in information quality substitute information quantity (positive substitution),*

(ii) $\tau^2 < \tau \le \tau^3$: *the weak gains in information quality are compensated with information quantity (positive compensation),*

(iii) $\tau^3 < \tau \le \tau^4$: *the strong losses in information quality are substituted by information quantity (negative substitution),*

(∅) $\tau^4 < \tau < \infty$: *the equilibrium degenerates,*

where $0 < \tau^1 \le \tau^2 \le \tau^3 \le \tau^4 < \infty$ *are time interval cutoffs.*

The result is valid regardless of whether deviations increase or decrease the drift of the process. Figures 1a and 1b provide numerical illustrations of Proposition 3 for $\mu_0 =$

$6\,(e_{10} + e_{20})$, $\sigma_0 = 12$, $r = 0.1$, $g = 3$ and $u = 1$. In this case, the time interval cutoffs are $\tau^1 = 0.46$, $\tau^2 = 0.76$, $\tau^3 = 3.59$ and $\tau^4 = 4.62$.

As in Sannikov and Skrzypacz (2007) and Fudenberg and Levine (2007), we found an equilibrium degeneracy for small time intervals (this situation occurs for $0 < \tau \leq \tau^1$). In other words, it is not possible to sustain mutual effort and cooperation for small time intervals. Despite the low deviation incentives (due to discounting), information quality is so bad that monitoring is impossible (see the likelihood ratio in Figures 1b).

This observation is related to the fact that reliable inference about the drift of the process requires a sufficiently large time interval (Prakasa Rao, 1999). Note that the drift of the process (1) is multiplied by the length of the time interval (see also Expression (2)). Therefore, for small time intervals, it becomes difficult to distinguish whether the drift carries mutual effort or not. In order to illustrate this argument, note that for small τ it is difficult to separate $c(2)\tau$ from $c\,(1)\,\tau$, i.e., we have $c(2)\tau \approx c\,(1)\,\tau$. However, for sufficiently large τ, it becomes clear that $c(2)\tau > c\,(1)\,\tau$.

However, there is a limit on how much information quality improves with the length of the time interval between actions (this situation occurs for $\tau^1 < \tau \leq \tau^3$), because for sufficiently large time intervals the process becomes extremely noisy. In other words, the noise component in Expression (1) dominates the drift component, and it becomes more difficult to distinguish the players' actions. Consequently, for $\tau^3 < \tau \leq \tau^4$, information quality, measured by the likelihood ratio decays.

Regarding the efficient use of information quantity and quality in the provision of incentives, for time intervals in the region $\tau^1 < \tau \leq \tau^4$, we found three different situations:

(i) As we move away from the limit case, information quality improves and mutual effort is possible if the information quality improvements are enough to compensate for the deviation incentives (due to the decreasing discount factor). In this case, as the time interval increases, incentives are sustained with fewer, but more precise signals (this situation occurs for $\tau^1 < \tau \leq \tau^2$). Information quantity is smoothly replaced by information quality. This is a situation of positive substitution because the increasing information quality improves the payoffs (see Expression (5)).

(ii) At a certain point, the informational quality improvements are not sufficiently strong to compensate for the increasing deviation incentives. Consequently, the weak information quality improvements must be compensated for by an increasing number of signals (this situation occurs for $\tau^2 < \tau \leq \tau^3$). This is a situation of positive compensation because the increasing information quality improves the payoffs (see Expression (5)).

(iii) At a certain point, the informational quality decays. In order to provide incentives for cooperation and mutual effort, the monitoring technology substitutes the decreasing quality with an even stronger increase in the quantity of signals (this situation occurs for $\tau^3 < \tau \leq \tau^4$). This is a situation of negative substitution, because the decreasing information quality reduces the payoffs (see Expression (5)).

Finally, at a certain point, for larger time intervals (i.e., for $\tau > \tau^4$), the deviation incentives are so strong (and the information quality so weak) that the mutual effort equilibrium degenerates.

4.2 Information structure: deviations increase the noise

With this information structure, a deviation causes an instantaneous increase in the variance of the process (1). This information structure fits with situations in which the observation of extreme events is interpreted as bad news. For example, events such as high market volatility, high instability or high sales variation are associated with misbehavior.

In our context, we can capture these types of situations in several ways. For instance, we can have an action dependent noise component of the form $\sigma_0 = \sigma\left(k - e_{10} - e_{20}\right)$, where $\sigma > 0$, $k > 2$ and the drift $\mu_0 = c \geq 0$ are constants. In this case, a unilateral deviation increases the noise parameter σ_0 from $\sigma\left(k-2\right)$ to $\sigma\left(k-1\right)$, and increases both the mistaken and the correct probabilities of punishment, $p(1,1)$ and $p(1,0)$, respectively.

With this information structure Fudenberg and Levine (2007) have shown that the optimal provision of incentives is achieved by means of a two-sided threshold strategy that distinguishes observations suggesting equilibrium play, i.e. $\left\{\underline{s}_\tau < s_\tau < \overline{s}_\tau\right\}$, from observations suggesting defection, i.e. $S_\tau = \left\{s_\tau \leq \underline{s}_\tau \cup s_\tau \geq \overline{s}_\tau\right\}$. Since deviations increase the volatility of the process, extreme observations become more likely; optimal monitoring considers the signals that are farther away from the mean as the ones that suggest deviation. On the contrary, the observations that are closer to the mean suggest mutual effort.

Figure 2a provides an illustration. The dark green lines denote the two-sided optimal threshold strategy $\underline{s}_\tau$ and $\overline{s}_\tau$ for varying time interval τ. The light green shaded regions (above and below the zero-mean) correspond to the set S_τ for varying time interval τ, i.e., the measure of information quantity (Definition 1). In this case, as the length of the time interval between actions varies, thus the optimal threshold strategy, and consequently, the number of signals used to sustain cooperation, thus rise.

Similarly, Figure 2b illustrates how information quality, measured by the likelihood ratio (Definition 2), varies with the time interval τ.

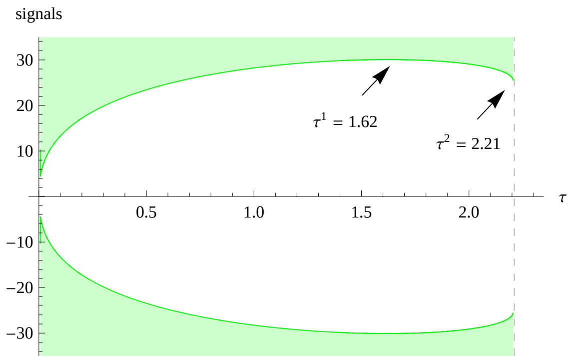

(a) Information quantity.

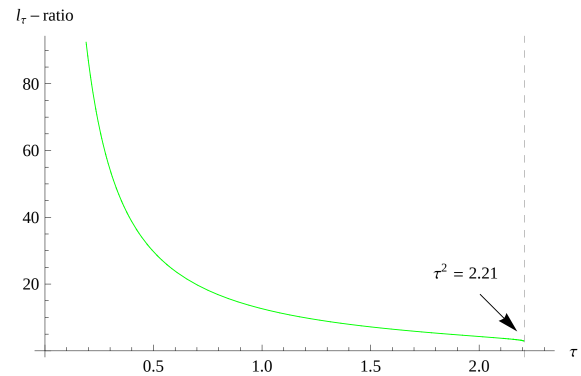

(b) Information quality.

Figure 2: **Deviation increases the noise -** (a) Information quantity, measured by the set S_τ (light green shaded area), and the optimal two-sided threshold strategy for varying time interval τ (dark green lines). (b) Information quality, measured by the likelihood ratio. The cooperation equilibrium ceases to exist for τ above the dashed line.

The following result shows the efficient use of information quantity and quality in providing incentives under the information structure in this subsection.

Proposition 4 *When deviations increase the noise, as the length of the time interval between actions increases:*

(i) $0 < \tau \leq \tau^1$: *the weak losses in information quality compensate information quantity (negative compensation),*

(ii) $\tau^1 < \tau \leq \tau^2$: *the strong losses in information quality are substituted by information quantity (negative substitution),*

(∅) $\tau^2 < \tau < \infty$: *the equilibrium degenerates,*

where $0 < \tau^1 < \tau^2 < \infty$ *are time interval cutoffs.*

Figures 2a and 2b provide numerical illustrations of Proposition 4 for $\mu_0 = 0$, $\sigma_0 = 5\,(4 - e_{10} - e_{20})$, $r = 0.1$, $g = 2$ and $u = 1$. In this case, the time interval cutoffs are $\tau^1 = 1.62$ and $\tau^2 = 2.21$.

Note that while we assume that a zero drift is not necessary, that drift can be non-zero and the results hold as long as it does not depend on the actions. A drift other than zero would simply result in a geometric rotation around the fixed point $(0, 0)$, in Figure 2a, in the direction determined by the drift. Similarly, in Figure 2b, a drift other than zero would result in a geometric transformation without implications for the results. These observations are also true in the next section.

This is the information structure proposed by Fudenberg and Levine (2007). In this case, information quality is maximum at the limit (when τ approaches zero), but decays with time. The likelihood ratio l_τ converges to infinite at the limit, and then is strictly decreasing in τ (see Figure 2b). Consequently, the full efficient equilibrium, in which expected payoffs converge to g, is possible at the limit.

In intuitive terms, reliable inference about the noise parameter can be obtained in small time intervals (Prakasa Rao, 1999). This is the case because Brownian motion is (almost surely) an infinitesimal variation process with continuous paths. Monitoring extreme events at small time intervals is easier because these events are unlikely. Consequently, efficient monitoring has sufficient freedom to establish a cutoff threshold that is infinitely more likely to be reached if there was a deviation than if there was no deviation.[11] However, this freedom is restricted as the time interval between actions increases. Inference becomes more difficult because the noise component of the process becomes increasingly important. Consequently, extreme observations of the process become increasingly likely, regardless of whether there was a deviation or not.

Considering time intervals in the region $0 < \tau \leq \tau^2$, we found two different situations:

(i) For small time intervals, incentives are provided with a large quantity of high quality signals. However, as the length of the time interval increases, information quality deteriorates and deviation incentives due to low discounting increase (this situation occurs for $0 < \tau \leq \tau^1$). However, since information quality decays slowly - at a rate that is still enough to compensate for the increasing deviation incentives - efficient monitoring reduces the number of signals needed to sustain mutual effort, and consequently, reduces the probability of punishment. Therefore, we observe a simultaneous decrease in information quality and quantity, which is possible because, in this region, the level of information quality is still high, and the deviation incentives are still low. This is a situation of negative compensation because information quality and quantity move in the same direction and payoffs decrease monotonically (see Expression (5)).

(ii) Since information quality degrades and the deviation incentives increase with the length of the time intervals, there is a point beyond which the decreasing information quality needs to be reinforced with an increasing number of signals (this situation occurs for $\tau^1 < \tau \leq \tau^2$). This is a situation of negative substitution because information quantity substitutes the falling information quality and payoffs decrease monotonically (see Expression (5)).

Finally, for larger time intervals (i.e. for $\tau > \tau^2$), the deviation incentives are so strong

[11] In technical terms, the difference in results between Section 4.1 and this section is related to the fact that the noise component converges to zero slower (at rate $\tau^{1/2}$) than the drift component, which converges to zero faster (at a rate τ). This aspect is crucial at the limit.

(and the information quality so weak), that the mutual effort equilibrium degenerates.

4.3 Information structure: deviations decrease the noise

With this information structure, a deviation causes an instantaneous decrease in the variance of the process (1). This information structure represents situations in which the observation of stable events, such as, no change in profits, could be associated with a lack of effort or commitment.

We can capture this type of situations in several ways. For instance, we can have an action dependent noise component of the form $\sigma_0 = \sigma(k + e_{10} + e_{20})$, where $\sigma > 0$, $k \geq 0$ and the drift $\mu_0 = c \geq 0$ are constants. Then, a unilateral deviation would decrease the noise parameter σ_0 from $\sigma(k+2)$ to $\sigma(k+1)$, and increases both the mistaken and the correct probabilities of punishment, $p(1,1)$ and $p(1,0)$, respectively.

With this information structure Fudenberg and Levine (2007) have shown that the optimal provision of incentives is achieved with a two-sided threshold strategy that distinguishes observations suggesting mutual effort, i.e. $\{s_\tau \leq \underline{s}_\tau \cup s_\tau \geq \overline{s}_\tau\}$, from observations suggesting defection, i.e. $S_\tau = \{\underline{s}_\tau < s_\tau < \overline{s}_\tau\}$. In the case of deviation, less extreme observations become more likely because a deviation decreases the volatility. Consequently, optimal monitoring considers the signals that are closer to the mean as the ones that suggest deviation. On the contrary, extreme realizations of the process suggest mutual effort.

Figure 3a provides an illustration. The dark green lines denote the symmetric two-sided optimal threshold strategy $\underline{s}_\tau$ and $\overline{s}_\tau$ for varying time interval τ. The light green shaded region (around the zero-mean) corresponds to the set S_τ for varying time interval τ, i.e., the measure of information quantity (Definition 1).

Similarly, Figure 3b illustrates how information quality, measured by the likelihood ratio (Definition 2), varies with the time interval τ.

The following result shows the efficient use of information quantity and quality in providing incentives under the information structure in this subsection.

Proposition 5 *When deviations decrease noise, as the length of the time interval between actions increases:*

(i) $0 < \tau \leq \tau^1$: *the strong losses in information quality are substituted by information quantity (negative substitution),*

(∅) $\tau^1 < \tau < \infty$: *the equilibrium degenerates,*

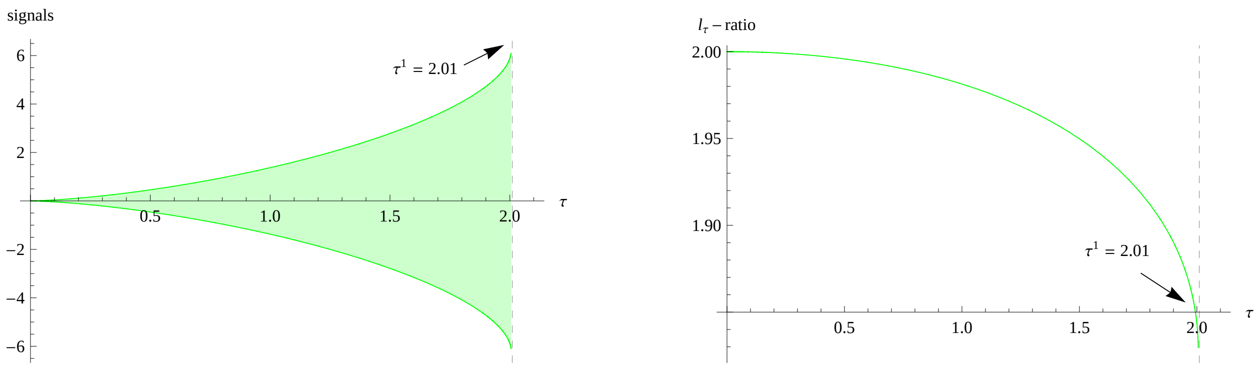

(a) Information quantity. (b) Information quality.

Figure 3: **Deviation decreases the noise** - (a) Information quantity, measured by the set S_τ (light green shaded area), and the optimal two-sided threshold strategy for varying time interval τ (dark green lines). (b) Information quality, measured by the likelihood ratio. The cooperation equilibrium ceases to exist for τ above the dashed line.

where $0 \leq \tau^1 < \infty$ *are time interval cutoffs.*

Figures 3a and 3b provide numerical illustrations of Proposition 5 for $\mu_0 = 0$, $\sigma_0 = 5\left(e_{10} + e_{20}\right)$, $r = 0.1$, $g = 2$ and $u = 1$. In this case, the time interval cutoff is $\tau^1 = 2.01$.

Note that in Section 4.1 equal size decreases and increases in the drift were equivalent. However, proportional decreases and increases in the noise parameter (in this section and Section 4.2, respectively) are not equivalent exercises.

Fudenberg and Levine (2007) have shown that the highest payoffs are obtained at the limit. However, the payoff values depend on the relative impact of deviations in the noise parameter. If the impact is sufficiently strong, we may get arbitrarily close to efficiency. However, if the impact is smooth, we may be unable to establish incentives for any frequency of play. In this case, we would have $\tau^1 = 0$ and the interval $0 < \tau \leq \tau^1$ would be empty. The following remark discusses this issue in more detail.

Remark (no equilibrium) *In the limit we have that* v approaches $g - u\sigma_0\left(1,0\right)/(\sigma_0\left(1,1\right) - \sigma_0\left(1,0\right))$, *where* $\sigma_0\left(1,0\right)$ *and* $\sigma_0\left(1,1\right)$ *are the noise parameters under defection and cooperation, respectively. Note that if* $v < 0$ *the mutual effort equilibrium fails to exist (not only in the limit, but for any length of time* τ*), which occurs for* $\sigma_0\left(1,1\right)/\sigma_0\left(1,0\right) < (g+u)/g$*, i.e., when deviations have small impact in the noise parameter. In the numerical example of Figures 3a and 3b, we do not have this situation because* $\sigma_0\left(1,1\right)/\sigma_0\left(1,0\right) = 2 > (g+u)/g = 3/2$*, consequently,* $\tau^1 = 2.01 > 0$.

In the interesting case where the interval $0 < \tau \leq \tau^1$ is non-empty, information quality decreases with the length of the time interval, i.e. l_τ is strictly decreasing with τ (Figure 3b).

Since the level of information quality is already low, information quality losses are strong in relative terms. This is the main difference compared to Section 4.2. The intuitive argument is similar, in the sense that the best inference about the value of the noise parameter is obtained in the smallest time intervals (Prakasa Rao, 1999). However, in the present section there are additional inference difficulties because deviations create an inward movement in the noise of the process. Since Brownian motion is (almost surely) an infinitesimal variation process with continuous paths, monitoring non-extreme events is more difficult than monitoring extreme events, because the former are likely to occur, regardless of whether there was a deviation or not. For that reason, the level of information quality is much lower than the one in Section 4.2. Nonetheless, as we have seen before, information quality might be enough to sustain the mutual effort equilibrium; it all depends on the impact of deviations in the noise parameter.

Regarding the efficient use of information quantity and quality in the provision of incentives, we found one single scenario (i.e. for $0 < \tau \leq \tau^1$): (i) In small time intervals, incentives for cooperation are provided by a small number of signals. However, as the length of the time interval between actions increases, mutual effort requires an increasing quantity of signals (Figure 3a). This pattern remains until the equilibrium collapses. There is a negative substitution of decreasing information quality for information quantity, and payoffs decrease monotonically for all frequencies of play (see Expression (5)).

Finally, for larger time intervals (i.e. for $\tau > \tau^1$), the deviation incentives are so strong (and the information quality so weak), that the mutual effort equilibrium degenerates.

5 Conclusion and further comments

Conclusions regarding the efficient use of information depend crucially on how information quality improves or decays relatively to: (i) the deviation incentives (due to discounting), and (ii) the actual level of information quality.

Nonetheless, payoffs only increase if information quality improves. Otherwise, even if there is a reduction in the quantity of signals used to provide incentives, payoffs do not improve. In this sense, our results are in line with Kandori (1992) and Mirrlees (1974).

Detailed conclusions depend on how each information structure uses information quality and quantity to provide incentives to mutual effort and cooperation. We have considered three information structures discussed in the literature.

When actions affect the drift (Section 4.1), information quality improves with the length of the time interval between actions (except for large time intervals, in which case the noise component of the process dominates). The intuitive argument is that reliable inference about the direction of the drift requires a sufficiently large time interval. This is the case because the drift component converges to zero relatively fast (at rate τ). Therefore, in small time intervals, it is difficult to infer whether the drift carries mutual effort or not.

On the other hand, if actions affect the noise parameter, information quality decays with the length of the time between actions (Sections 4.2 and 4.3). This decay occurs because the noise component of the process becomes increasingly important as the time interval grows, which makes the distinction between actions more difficult, and increases the probability of punishment.

Contrary to the drift, the best inference about the noise parameter is obtained in the smallest time intervals. Nonetheless, we must distinguish between the cases in which deviations increase and decrease the noise parameter. Since Brownian motion is an infinitesimal variation process with continuous paths, in small time intervals, monitoring extreme events (i.e. the case in which deviations increase the noise) is easier than monitoring non-extreme events (i.e. the case in which deviations decrease the noise). The former are less likely to occur than the later, regardless of whether there was a deviation or not. For this reason, when deviations increase the noise, efficient monitoring has more freedom to establish a cut-off threshold that is infinitely more likely to have been reached if there was a deviation than if there was no deviation. On the contrary, when deviations decrease the noise, the inward movement in the process makes it more difficult to detect if there was a deviation or not.

We also found that information quality and quantity could be combined in different ways to provide incentives to mutual effort and cooperation. These combinations depend crucially on the relative strength of the information quality gains and losses for varying time interval. We found that information quality and quantity substitute for each other when the information quality gains or losses are sufficiently strong in relative terms. However, when the information quality gains or losses are weak in relative terms, information quantity must be used to compensate for the poor information quality.

For instance, as the length of the time interval between actions increases (for large time intervals), we observe - in all information structures - a negative substitution of information quality by information quantity. The reason is twofold: the incentives for deviations (due to discounting) are high and the signals become extremely noisy for large time intervals. This pattern remains until the cooperation equilibrium ceases to exist.

The study of the trade-offs between information quality and quantity seems to have been consistently ignored in the literature. We call for a research agenda on these issues. Despite the complexity of the topic, the present paper is an attempt in this direction. We hope that our findings will help researchers and practitioners to better understand the role played by actions and monitoring frequency on the individuals' incentives to cooperate when there is imperfect monitoring. In particular, our results may guide researchers and practitioners choosing the modelling approach that better fits each concrete situation.

References

Abreu, D., P. Milgrom and D. Pearce (1991). "Information and Timing in Repeated Partnerships." Econometrica, 59, 1713-1733.

Abreu, D., D. Pearce and E. Stacchetti (1986). "Optimal Cartel Equilibria with Imperfect Monitoring." Journal of Economic Theory, 39, 251-269.

Abreu, D., D. Pearce and E. Stacchetti (1990). "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." Econometrica, 58, 1041-1063.

Faingold, E., and Y. Sannikov (2011). "Reputation in Continuous-Time Games," Econometrica, 79(3), 773-876.

Fudenberg, D., Y. Ishii and S. Kominers (2014). "Delayed-Response Strategies in Repeated Games with Observation Lags." Journal of Economic Theory, 150, 487-514.

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

Fudenberg, D. and D. Levine (2009). "Repeated Games with Frequent Signals." Quarterly Journal of Economics, 124, 233-265.

Fudenberg, D. and W. Olszewski (2011). "Repeated Games with Asynchronous Monitoring of an Imperfect Signal," Games and Economic Behavior, 72, 86-99.

Kamada, Y. and S. Kominers (2010). "Information can wreck cooperation: A counterpoint to Kandori (1992)," Economics Letters, 107, 112-114.

Kandori, M., (1992). "The use of information in repeated games with imperfect monitoring." Review of Economic Studies, 59, 581-594.

Mirrlees, J., (1974). "Notes on welfare economics, information and uncertainty." In: Balch, M., McFadden, D., Wu, S. (Eds.), Essays on Economic Behavior under Uncertainty. North-Holland, Amsterdam.

Osório, A. (2012). "A Folk Theorem for Games when Frequent Monitoring Decreases Noise," The B.E. Journal of Theoretical Economics, 12(1), Article 11.

Osório, A. (2015). "Some Notes and Comments on the Efficient use of Information in Repeated Games with Poisson Signals." Operations Research Letters, 43(3), 288-292.

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

Prakasa Rao, B. (1999). Statistical Inference for Diffusion Type Processes. Arnold Publishers and Oxford University Press.

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

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

Sannikov, Y. and A. Skrzypacz (2010). "The Role of Information in Repeated Games with Frequent Actions," Econometrica, 78, 847-882.