Games and Economic Behavior 72 (2011) 86–99

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Repeated games with asynchronous monitoring of an imperfect signal ☆

Drew Fudenberg, Wojciech Olszewski *

Northwestern University, Department of Economics, 2003 Sheridan Road, Evanston, IL, United States

ARTICLE INFO

Article history:
Received 11 June 2009
Available online 16 September 2010

JEL classification:
C02
C73

Keywords:
Repeated games
Poisson and diffusion signal processes
Asynchronous monitoring
Continuous-time limits

ABSTRACT

We consider a long-run player facing a sequence of short-run opponents who receive noisy signals of the long-run player's past actions. We modify the standard, synchronous-action, model by supposing that players observe an underlying public signal of the opponent's actions at random and privately known times. In one modification, the public signals are Poisson events and either the observations occur within a small epsilon time interval or the observations have exponential waiting times. In the second modification, the underlying signal is the position of a diffusion process. We show that in the Poisson cases the high-frequency limit is the same as in the Fudenberg and Levine (2007, 2009) study of limits of high-frequency public signals, but that the limits can differ when the signals correspond to a diffusion.

1. Introduction

This paper studies repeated games where players have imperfect information about one another's past actions and where their signals and moves occur asynchronously. There is a substantial literature on repeated games with imperfect monitoring and synchronous moves, and a smaller one on repeated games with perfectly observed actions and asynchronous moves; this paper extends both of them in the context of a particular example of a game between a long-run player and a sequence of short-run opponents. When the time between consecutive signals or moves is long, there is no reason to expect the same equilibria under the synchronous and asynchronous scenarios. Our focus is on the comparison between the two scenarios in the high-frequency limit when signals and moves occur rapidly.

Our starting point is the observation that standard models of repeated games with imperfect monitoring assume that both monitoring and play are *exactly* synchronous. This simplifying assumption is hard to motivate, even as an approximation, when signals and play occur at a very high frequency, which raises the question of whether results about high-frequency synchronous signals extend to settings with asynchronous observation of a common signal. By this we have in mind cases where the players' signals correspond to the current state of a signal process such as units sold, revenue, or customer satisfaction, and where any player who observes the state of the signal process at a given date observes the same thing. The case of synchronous signals corresponds to all players observing this process at the same sequence of dates t, $2t$, etc.; this paper studies cases where each player may observe the process at slightly different times.

The paper can also be seen as an extension of the literature on asynchronous play in repeated games. Until very recently, work on asynchronous play only considered the case of observed actions and all long run players. This is a subclass of

☆ We thank David K. Levine, Satoru Takahashi, and a referee for helpful comments, and NSF grants SES 0646816 and CAREER SES 0644930 for financial support.

* Corresponding author.
E-mail address: wo@northwestern.edu (W. Olszewski).

0899-8256/$ – see front matter
doi:10.1016/j.geb.2010.08.005

stochastic games (Shapley, 1953) for which a folk theorem holds (Dutta, 1995).[1] Subsequent to this paper, Fudenberg and Yamamoto (2009) and Hörner et al. (2009) provide folk theorems for irreducible stochastic games with imperfect public monitoring. Their models do not include ours, as in our model asynchronous imperfect observations lead to a particular form of private monitoring, as neither player knows the signals that the other player observed. They also do not study the high-frequency limit of their games, which requires a specification of the way that the monitoring structure and state transitions vary with the period length.

To study high-frequency limits, we parameterize the effect of monitoring frequency on information by supposing there is a family of continuous-time signal processes whose distribution depends on the players' actions but not directly on the monitoring frequency itself. As in most work on high-frequency repeated games, we suppose for analytic simplicity that this continuous-time process is either a diffusion or a Poisson process.[2] As in Fudenberg and Levine (2007) (FL) and Fudenberg and Levine (2009) and Faingold and Sannikov (2007), we restrict attention to a game between a long player 1 and a sequence of short-run player 2's (or equivalently a single player 2 who is myopic). In such games, the best limit equilibrium payoffs with synchronous public signals can be obtained with trigger strategies that "punish" with reversion to a static Nash equilibrium whenever a sufficiently bad signal is observed, so there is a relatively straightforward link between the informativeness of the signals and the set of payoffs to perfect public equilibria. Moreover, this is perhaps the simplest case where the accuracy of information has an impact on equilibrium payoffs in the limit of discount factors tending to 1, so it provides a natural test bed for studying high-frequency limits.[3]

The particular game that FL considered was based on Klein and Leffler's (1981) model of a long-run firm facing a sequence of short-run consumers. In that model, the firm's action (choice of effort or quality) is irrelevant if consumers play "Out," i.e. choose not to purchase, so in the static equilibrium (Low quality, Out) the firm is indifferent between its choices. This indifference did not matter in the synchronous-move case considered by FL – relaxing it has no effect on their results – but we find that it can matter in some specifications of asynchronous play.

We begin with the case of players observing the state of a Poisson process. Here we consider two formulations of asynchronous monitoring. In the first, each player observes the state and revises their action at times $t - \varepsilon_{i,1}, 2t - \varepsilon_{i,2}, \ldots, kt - \varepsilon_{i,k}, \ldots$, where $(\varepsilon_{i,k})_{k=1}^{\infty}$ is a sequence of i.i.d. random variables distributed uniformly on an interval $[0, \varepsilon]$ with $\varepsilon \geqslant 0$. In this case, for any fixed t, the observation times are highly correlated as $\varepsilon \to 0$, and so are the observations themselves. In particular, for sufficiently small values of ε, the players are very likely to observe the same number of Poisson events, which makes it relatively straightforward to show that the limit payoffs (as $t \to 0$ and $\varepsilon/t \to 0$) are the same as in FL, both for the payoff matrix they considered and for a variant where the long-run player's has an incentive to "shirk" regardless of the short-run player's action.

In our second version of asynchronous Poisson signals, we suppose that players observe the Poisson process at exponential times with intensity γ. To separate the effect of stochastic period length from that of asynchronous observations, we first suppose that the observation times of the players are perfectly correlated, so the players observe and act simultaneously at exponentially distributed waiting times. With observed actions, this sort of stochastic period length is equivalent to the usual discrete-time repeated game with a lower discount factor (Kamada and Kandori, 2009), so the folk theorem holds as usual when both the discount factor and the interaction probability go to 1. Our results show that the limit equilibrium payoffs when the observation frequencies $\gamma \to \infty$ coincide with those of the discrete-time (synchronous-move) model when the time between consecutive periods $t \to 0$.

We then consider the case where the players' observation times are independent. Of course, synchronous and asynchronous play can differ for a fixed value of γ, but we show that the effect of asynchronous signals vanishes $\gamma \to \infty$, so the limit conclusions for the asynchronous model are the same as for the synchronous one. Intuitively, this is because when γ is high the players observe the process much more frequently than events occur, so that the incentive constraints for equilibrium are governed by the wait for an occurrence of a Poisson event, and the wait to observe an event once it occurs is negligible.

The Poisson case is relatively simple to analyze because the trigger strategies that support the best limit equilibrium payoff give the players a strict incentive to conform to equilibrium play. The case where the underlying signal process is a diffusion is more complicated, as here there is a continuum of signals, and the optimal equilibrium strategies with simultaneous moves take the form of "cutoffs": continuation play depends on whether the observed signal is inside or outside of a given interval. For signals near the cutoff, the players are nearly indifferent, yet cutoff strategies implicitly rely on a coordinated response that is harder to implement with asynchronous moves. For this reason, asynchronous signals have a greater impact with diffusion signals, and in particular the best limit payoffs in trigger strategies can be strictly lower with asynchronous signals than with synchronous ones. The exact argument for this is too complicated to summarize here, but some intuition can be gained from the simpler case where players observe the diffusion simultaneously but with a small

[1] Note that Dutta uses the full-dimension condition of Fudenberg and Maskin (1986), which Abreu et al. (1994) show is stronger than necessary. The "Anti-folk-theorem" of Lagunoff and Matsui (1997) shows that stronger full-dimension conditions than Abreu, Dutta and Smith's are needed to obtain the folk theorem with asynchronous moves.

[2] Sannikov and Skrzypacz (2009) allow for both sorts of signals. More generally, it is well-known that any continuous-time process with independent identical increments can be decomposed into the sum of a k-dimensional diffusion and a jump process.

[3] Alternatively one can study symmetric games and restrict attention to strongly symmetric equilibria as in Abreu et al. (1991); this restricts the set of equilibrium payoffs to be one-dimensional, just as with a single long-run player.

Table 1

		Player 2	
		Out	In
Player 1	+1	$\underline{u}, 0$	$\bar{u}, 1$
	−1	$\underline{u}, 0$	$\bar{u}+g, -1$

Table 2

		Player 2	
		Out	In
Player 1	+1	$\underline{u}, 0$	$\bar{u}, 1$
	−1	$\underline{u}+g, 0$	$\bar{u}+g, -1$

amount of noise, so that conditional on her own signal each player believes the opponent's signal in normally distributed with a small positive variance. Then if players use cutoff strategies, and a player's observation is very close to a cutoff level, she believes that the opponent will play each of the two actions with probability close to 1/2, as opposed to the perfect coordination that is possible with completely public signals. This intuition is quite similar to that behind the anti-robustness result of Sugaya and Takahashi (2009) for the effect of a small amount of private noise.

2. The discrete-time model

A long-run player 1 (who will also be called LR) plays a stage game with a short-run player 2 (also be called SR) who is completely impatient in a 2×2 stage game. FL consider the case of the payoff matrix in Table 1 where neither player's payoff depends on player 1's action when player 2 plays Out. As we explained in the introduction, this aspect of the payoffs comes from the Klein–Leffler example used to motivate the game, and was not important for FL's results. However, some of the asynchronous-monitoring results for the game from Table 1 seem to rely on player 1 being indifferent when 2 plays Out. For this reason, we will consider both this payoff matrix and the following small modification where $\underline{u}+g<\bar{u}$. In both games, player 2 plays Out in the unique static Nash equilibrium, while player 1 would prefer that player 2 plays In. The only difference is that in Table 2, player 1's payoff from action profile (−1, Out) is higher by g, so that (−1, Out) is a strict Nash equilibrium, and the "punishment payoff" for player 1 is now $\underline{u}+g$.

Turning to the repeated version of the game, we study a continuous-time model, in which players act at discrete times kt, $k=0,1$, etc., so that the length of a period between two consecutive actions is t. Thus, the long-run player's discount factor per period is $\delta=\exp(-rt)$ where r is her (subjective) interest rate. We assume that at times kt, before taking actions, players observe some public signals. These public signals include the action taken by the short-run player at time $(k-1)t$, an imperfect signal y_k about the action taken by the long-run player at time $(k-1)t$, and a public randomization device (that is, a draw from the uniform distribution on the interval $[0,1]$).[4] To tie the distribution of y_k to the period length, we suppose that for each action of LR there is a continuous-time process Z_t with independent and identically distributed increments, and that the signal y_k corresponds to the increment $Z_{kt}-Z_{(k-1)t}$. In our applications, we set Z_t to be either a Poisson process or a diffusion, and let f_t^+ and f_t^- denote the probability density of y_k under actions +1 and −1, respectively, when the time period has length t. The public history at time kt consists of the public signals $y_t, y_{2t}, \ldots, y_{kt}$ observed at times $t, 2t, \ldots, kt$. The private history of the long-run player at time kt includes the public history and also her actions taken at times $0, t, \ldots, (k-1)t$; the private history for the short-run players is simply the public history.

Our solution concept here is perfect public equilibrium or PPE (Fudenberg et al., 1994): these are strategy profiles for the repeated game in which (a) each player's strategy depends only on the public history, and (b) no player wants to deviate at any private history. In both stage games, the Stackelberg payoff of $\bar{u}+g/2$ can be obtained by a publicly observed commitment to play the mixed strategy $(1/2, 1/2)$, but when actions are observed the highest subgame-perfect equilibrium payoff of the repeated game is $\bar{u}$ (Fudenberg et al., 1990), and the highest PPE payoff with imperfect public monitoring is strictly less than that (Fudenberg and Levine, 1994).

FL show it is without loss of generality to look at PPE with continuation payoffs restricted to the two values ("reward") and ("punishment"), where punishment corresponds to play of the static Nash equilibrium, and "reward" means that the players continue to play (+1, In) until a punishment occurs. Define p to be the probability of punishment when the action chosen is +1 (that is, p is the probability under action +1 of signals such that continuation play is "punishment"), and define q to be the probability of punishment when the action chosen is −1. We say that a pair (p,q) is feasible if it can be generated by some specification of the function w that maps public signals to the two continuation payoffs.

Proposition 1. *(See Fudenberg and Levine, 2007.) Consider the payoff matrix of Table* 1. *For a fixed discount factor* δ, *there is an equilibrium with the long-run player's payoff above* $\underline{u}$ *if and only if there are feasible* p *and* q *that satisfy*

[4] The order in which players observe the components of the public signal at time kt does not matter.

$$\frac{(\bar{u}-\underline{u})}{g}\frac{(q-p)}{p} \geqslant \frac{(1-\delta)}{\delta p}.$$

In this case the highest PPE payoff to the long-run player is

$$\max_{feasible\ p,q} \bar{u} - \frac{pg}{q-p}.$$

It is easy to see from their proofs that similar results hold for the payoffs of Table 2; the only difference is that the punishment payoff $\underline{u}$ is replaced by $\underline{u}+g$ in the incentive constraint for a non-trivial equilibrium. FL use this result to give various partial characterizations of the limit equilibria as $t \to 0$. Rather than restate their general results, we will state the specific versions relevant for our cases in the relevant cases.

3. Poisson events

In this section, we assume that Z_t is a Poisson process, so that with synchronous moves the signal y_k corresponds to the number of Poisson events over the time interval from $(k-1)t$ to kt. Let the expected waiting time for a Poisson event be $1/\lambda_p$ if LR plays $+1$, and $1/\lambda_q$ if LR plays -1. Assume that $1/\lambda_q < 1/\lambda_p$, so that the events are more likely to occur under action -1; in the Klein–Leffler game this corresponds roughly to events being a product failure or episode of unsatisfactory customer support. (This is the "bad news" case; FL show that there is no non-trivial limit equilibrium in the case $1/\lambda_q > 1/\lambda_p$, where events are "good news" such as a drug curing a disease.)

FL show that with synchronous signals and moves the highest limit PPE payoff can be achieved by "cutoff strategies," by which we mean strategies of the following form: Start out playing $(+1, \text{In})$; once SR is observed to play Out, switch to playing $(-1, \text{Out})$ from then on; when players observe a Poisson event, use the public randomizing device to switch to playing $(-1, \text{Out})$ from then on with some probability; and otherwise keep playing $(+1, \text{In})$.[5]

In general, whether there is a non-trivial equilibrium when time periods are very short can depend on the interest rate r as well as on the stage-game payoffs and the parameters of the monitoring technology. Following FL, we will focus on whether there are non-trivial equilibria for interest rates near 0.[6] FL say that there is a *"non-trivial limit equilibrium"* if there is some interest rate $\bar{r} > 0$ such that for all $r < \bar{r}$ there is a sequence of equilibria for the games with period lengths $t \to 0$ whose payoffs are uniformly higher than in the static equilibrium. FL show that with the payoffs of Table 1, a necessary and sufficient condition for the existence of non-trivial limit equilibrium is

$$g/(\bar{u}-\underline{u}) < (\lambda_q - \lambda_p)/\lambda_p,$$

and that when this inequality is satisfied, the best limit PPE payoff is

$$v^* = \bar{u} - g\lambda_p/(\lambda_q - \lambda_p).$$

It is easy to see from their arguments that the condition for a non-trivial limit PPE with the payoffs of Table 2 is

$$g/(\bar{u}-\underline{u}-g) < (\lambda_q - \lambda_p)/\lambda_p;$$

the best limit PPE payoff is the same as with the payoffs of Table 1.

3.1. Poisson events observed at almost the same time

Now suppose that players observe signals and act at times $t-\varepsilon_{i,1}, 2t-\varepsilon_{i,2}, \ldots, kt-\varepsilon_{i,k}, \ldots$, where $(\varepsilon_{i,k})_{k=1}^{\infty}$ is a sequence of i.i.d. random variables distributed uniformly on an interval $[0,\varepsilon]$ with $\varepsilon \geqslant 0$ and $\varepsilon < t$. More precisely, player i observes at time $kt-\varepsilon_{i,k}$ the current state $Z_{kt-\varepsilon_{i,k}}$ of the Poisson process, the value of the accompanying public randomization, and SR's action taken at time $(k-1)t-\varepsilon_{2,k-1}$. Then, also at time $kt-\varepsilon_{i,k}$, player i can change her action. Both here and in all subsequent sections of the paper, LR's payoff is the discounted flow of the stage game payoffs,

$$\int \exp(-r\tau)U_1(a_{1\tau}, a_{2\tau})\,d\tau$$

where U_1 is LR's stage-game payoff function where $a_{i\tau}$ is the action that was chosen by player i at his most recent move. We suppose that SR is completely myopic, and so chooses a best response to her beliefs about the current actions of LR.[7]

[5] Because the density of y_t satisfies the monotone likelihood ratio property (i.e. $f_t^-(z)/f_t^+(z)$ is strictly increasing in z), the optimal equilibrium with synchronous moves must take the form of a cutoff: Switch to $(-1, \text{Out})$ if more than c events are observed in the current period; switch with some probability α if exactly c events are observed; and do not switch if fewer than c events are observed. As $t \to 0$ the probability of more than one event in a time period is $O(1/t^2)$; FL show that this implies that if there is a non-trivial limit equilibrium, the optimal strategies have a cutoff of $c = 1$ for small t.

[6] Although our analysis also has implications for the nature of equilibrium payoffs for larger interest rates, the main points can be seen from the case of interest near zero, and stating the results this way is simpler and facilitates the comparison with past work.

[7] This is the simplest way of extending the LR/SR model to asynchronous moves. One could also study the case where each SR tries to maximize expected payoff until the next move; we conjecture that this will not change our results. In particular, we have checked that it does not change the conclusions of Proposition 2.

This monitoring structure is motivated by the objective of making only a minimal departure from the synchronous case studied in FL.

To facilitate the analysis of optimal PPE, we assume in the asynchronous case that when player i observes $Z_{kt-\varepsilon_{i,k}}$, if there have been any Poisson events between $kt-\varepsilon_{i,k}$ and $(k-1)t-\varepsilon_{i,k-1}$, the player also observes the time at which each of these events occurred. Moreover, each of these events is accompanied by a realization of the public randomizing device, which occurs exactly at the same time as the event occurs. (We make the latter assumptions to make the analysis easier; we believe that the same results would obtain here without it, but more complicated strategies would be needed.)

Proposition 2. *Fix either the payoffs of Table* 1 *or those of Table* 2*. Suppose that there is a non-trivial limit equilibrium of the simultaneous-move games as $t \to 0$. Then every limit payoff that can be attained with synchronous moves is the limit of sequential equilibrium for positive small values of ε as $t \to 0$, $\varepsilon/t \to 0$.*

Proof. Consider first the payoffs of Table 1, and suppose $g/(\bar{u}-\underline{u}) < (\lambda_q-\lambda_p)/\lambda_p$. We claim that a modification of the strategy analyzed by FL is an equilibrium for r close to 0 and sufficiently small positive values of t and ε/t. More precisely, take an $r<\bar{r}$ as in the definition of non-trivial limit equilibrium. The modification is necessary because if SR observes that an event has arrived later than at time $kt-\frac{\varepsilon}{2}$, the probability that LR has not observed the event yet is higher than $\frac{1}{2}$. Then SR prefers playing In, even if action Out is prescribed.

To facilitate the description of the modified strategies, fix a probability α that LR switches to -1 when observing a Poisson event, and say that a "punishment event" occurs at time t when a Poisson event is realized at t and the value of the accompanying public randomization is less than or equal to α.[8] We consider strategies of the following form: player 1 starts out playing $+1$ and sticks with it until the first time she sees a punishment event or she sees player 2 play Out; once player 1 starts playing -1, she plays -1 from then on. Player 2 starts out playing In, and plays In up to the time she sees a punishment event. If she observes a punishment event arrived in interval $[kt-\varepsilon, kt-\frac{\varepsilon}{2}]$, she switches to playing Out. If she observes a punishment event arrived in interval $(kt-\frac{\varepsilon}{2}, kt]$ she switches to playing Out at time at her next move, which will be at $(k+1)t-\varepsilon_{2,k+1}$.

Consider first the case when LR observed no punishment event till time kt, and SR played In at all times less than or equal to $(k-1)t$. Then, when $\varepsilon/t \to 0$, the probability assigned by LR to the event that SR observed no punishment event herself tends to 1. Thus, the difference between LR's payoff to playing -1 and $+1$ at time $kt-\varepsilon_{1,k}$ received between $kt-\varepsilon_{1,k}$ and $(k+1)t-\varepsilon_{1,k+1}$ differs compared to the case of $\varepsilon=0$ only by $O(\varepsilon/t)$.

The expected difference between the continuation payoffs from time $(k+1)t-\varepsilon_{1,k+1}$ on, to playing $+1$ and -1 between times $kt-\varepsilon_{1,k}$ and $(k+1)t-\varepsilon_{1,k+1}$ also changes only by $O(\varepsilon/t)$. Indeed, LR's action between times $kt-\varepsilon_{1,k}$ and $(k+1)t-\varepsilon_{1,k+1}$ affects SR's signal observed at time $(k+2)t-\varepsilon_{2,k+2}$, because the probability of SR observing an event at time $(k+2)t-\varepsilon_{2,k+2}$ depends on LR's action between time $(k+1)t-\varepsilon_{2,k+1}$ and time $(k+1)t$. It also affects, for the same reason, LR's signal observed at time $(k+2)t-\varepsilon_{1,k+2}$, and so LR's action at time $(k+2)t-\varepsilon_{1,k+2}$. The consequences of LR's action between times $kt-\varepsilon_{1,k}$ and $(k+1)t-\varepsilon_{1,k+1}$ propagate throughout the entire future. However, the consequences for LR's payoff are of order $O(\varepsilon/t)$.

Consider now the case when LR observed a punishment event at time $kt-\varepsilon_{1,k}$, and SR played In at all times less than or equal to $(k-1)t$. Then, LR knows that SR will observe the event at time $(k+1)t-\varepsilon_{2,k+1}$ and will then switch to playing Out, so LR will play -1.

The incentives of SR are straightforward, and similar arguments apply to the payoffs of Table 2. □

Past work on games with all long-run players by Mailath and Morris (2002), Hörner and Olszewski (2009) and Mailath and Olszewski (2008) suggests that results on public monitoring would be robust to perturbations that make the monitoring almost public. (Indeed, Hörner and Olszewski, 2009, prove folk theorems for repeated games with finite sets of actions and signals, a finite number of long-run players, and private monitoring structures converging to public monitoring.) And, as we remarked in the introduction, LR who observes no punishment event is almost certain that SR has observed no punishment event. Thus the event that no punishment event was observed is "almost common knowledge" in the sense of Monderer and Samet (1989), so that in the case of observing no Poisson event, the monitoring structure is "almost public." However, this monitoring structure is not "almost public." When a player observes a punishment event, she need not be almost certain that her opponent also observed the event, because the probability depends on when the event has arrived. Thus the techniques of the papers on almost public monitoring are not directly applicable here.[9] The departure from almost-public monitoring turns out to be inessential in the case of Poisson process, but as we will see when we study diffusion signals, the distinction between almost-public monitoring and other sorts of arguably "small" perturbations remains important.

[8] The assumption that each Poisson event is accompanied by a realization of the public randomizing device is important here. Without it, players could disagree whether an event is a punishment event.

[9] In addition, some of the papers use indifference constructions that are not feasible with short-run players, and others use a restriction on the memory of the strategies.

3.2. Poisson events and exponentially distributed observation times

In the model of the previous subsection, the observation times of the players become highly correlated as $\varepsilon/t \to 0$. We now study the less artificial case where players observe the number of Poisson events that have occurred at exponentially distributed times with arrival rate γ independent of the actions played, and that players can reevaluate their action each time they make an observation, so that the time interval between moves is stochastic. We will consider two versions of this model, one with synchronous (perfectly correlated) observation times, and one in which the two players' observation time are independent. In each version, the wait between an observation of player 1 and player 1's next observation has the same distribution as the wait between an observation by player 1 and the next observation by player 2, so in a sense the "correlation coefficient" of the two players' observation times remains constant (at either 1 or 0, respectively) regardless of γ; in contrast the observation times of the previous subsection become increasingly correlated as $\varepsilon/t \to 0$.

We maintain the assumption that when players observe an event, they observe the time at which the event occurred and the accompanying realization of the public randomizing device; as in the previous section, we believe that the same results would obtain without it. We will show that in this model too the synchronous and asynchronous scenarios have the same high-frequency limit, where "high-frequency" here corresponds to $\gamma \to \infty$. Intuitively, when γ is large the Poisson event is observed almost as soon as it occurs; the key lag in the model, and the reason that the equilibrium will not in general be fully efficient, is the waiting time for the event to occur, which requires that an event be followed by a non-negligible increase in the probability of punishment.

3.2.1. Synchronous exponentially distributed moves

We begin with the synchronous case where players observe signals and move at the same time, and the observations have exponential waiting time with parameter γ.[10] Thus, players move at random times $t_1, t_2, \ldots$, etc.; at each such time, they observe the corresponding value of Z_t, the times at which all events occurred, and also the corresponding realizations of the public randomizing device. This defines the public histories at times t_k. As before, players also have private information about their own past actions, but we will restrict attention to PPE so that this private information is irrelevant.

Note that the probability of two or more events occurring before an observation is $O(1/\gamma^2)$, so as with deterministic period lengths, the strategy of punishing on two or more events cannot support a non-trivial limit equilibrium as $\gamma \to \infty$. To begin the analysis, we consider the grim strategies "Play (+1, In) until an event occurs; when it does play (−1, Out) forever."

With the payoffs in Table 1, this strategy profile has payoff

$$v^* = \int_{y=0}^{\infty}\int_{s=y}^{\infty} \lambda_p \exp(-\lambda_p y)\gamma \exp\bigl(-\gamma(s-y)\bigr)\bigl(\bar{u}\bigl(1-\exp(-rs)\bigr) + \underline{u}\exp(-rs)\bigr)\,ds\,dy$$
$$= \frac{\bar{u}r(\gamma+\lambda_p+r) + \underline{u}\gamma\lambda_p}{(\lambda_p+r)(\gamma+r)}.$$

Note that as $\gamma \to \infty$, v^* converges to

$$\bar{v} = \frac{\bar{u}r + \underline{u}\lambda_p}{\lambda_p + r}.$$

In general, even when γ is very large the optimal equilibrium will not use this grim strategy, but by adapting the arguments of FL, one can see that the highest PPE payoffs as $\gamma \to \infty$ can be obtained with a strategy that is "α-grim," meaning that it starts in the cooperative phase but makes use of the public randomizing device to only punish with probability α if an event occurs. (That is, each time an event occurs, play switches to "(−1, Out) forever" with probability α, and with probability $1-\alpha$ ignores the event and stays in the cooperative phase.)

This strategy has payoff $v^{*\alpha}$, where $v^{*\alpha}$ solves

$$v^{\alpha *} = \int_{y=0}^{\infty}\int_{s=y}^{\infty} \lambda_p \exp(-\lambda_p y)\gamma \exp\bigl(-\gamma(s-y)\bigr)\bigl(\bar{u}\bigl(1-\exp(-rs)\bigr) + \bigl(\alpha\underline{u} + (1-\alpha)v^*\bigr)\exp(-rs)\bigr)\,ds\,dy$$
$$= \frac{\bar{u}r(\gamma+\lambda_p+r) + \alpha\underline{u}\gamma\lambda_p}{(\lambda_p+r)(\gamma+r)} + (1-\alpha)\frac{\gamma\lambda_p v^{\alpha *}}{(\lambda_p+r)(\gamma+r)},$$

so that

$$v^{*\alpha} = \frac{\bar{u}r(\gamma+\lambda_p+r) + \alpha\underline{u}\gamma\lambda_p}{\alpha\gamma\lambda_p + \lambda_p r + \gamma r + r^2}.$$

[10] As noted in the introduction, Kamada and Kandori (2009) study this sort of timing in a repeated game with observed actions; it is also used by Ambrus and Lu (2009) to study non-cooperative bargaining.

Note that as $\gamma \to \infty$, $v^{*\alpha}$ converges to

$$\bar{v}^{\alpha} = \frac{\bar{u}r + \alpha \underline{u}\lambda_p}{\alpha\lambda_p + r}.$$

Following FL, we say that there is a "non-trivial limit equilibrium" if there is some interest rate $\bar{r} > 0$ such that for all $r < \bar{r}$ there is a sequence of equilibria for the games with $\gamma \to \infty$ whose payoffs are uniformly higher than in the static equilibrium.

Proposition 3. (a) *For the payoffs of Table* 1*, with synchronous exponential observations and Poisson signals, there is a non-trivial limit equilibrium as* $\gamma \to 0$ *iff*

$$\frac{g}{\bar{u} - \underline{u}} < \frac{(\lambda_q - \lambda_p)}{\lambda_p}. \tag{3.1}$$

The best limit equilibrium payoff is

$$\bar{u} - \frac{g\lambda_p}{(\lambda_q - \lambda_p)}$$

just as in FL.

(b) *With the payoffs of Table* 2*, the condition for a non-trivial limit is*

$$\frac{g}{\bar{u} - \underline{u} - g} < \frac{(\lambda_q - \lambda_p)}{\lambda_p},$$

and the best limit equilibrium payoff is again

$$\bar{u} - \frac{g\lambda_p}{(\lambda_q - \lambda_p)}.$$

Proof. Suppose players use α-grim strategies. The one-period expected cost of playing $+1$ (compared to playing -1) is $gr/(\gamma + r)$; this decreases the probability of a regime change by $(\lambda_q - \lambda_p)\gamma\alpha/(\lambda_p + \gamma)(\lambda_q + \gamma)$, which yields a benefit of $v^{*\alpha} - \underline{u}$.

Thus the strategies satisfy the one-stage deviation principle if

$$\frac{g}{\bar{u} - \underline{u}} \leqslant \frac{(\lambda_q - \lambda_p)\gamma\alpha}{(\lambda_p + \gamma)(\lambda_q + \gamma)} \frac{(\gamma + \lambda_p + r)(\gamma + r)}{[\alpha\gamma\lambda_p + \lambda_p r + \gamma r + r^2]}. \tag{3.2}$$

For large gamma, this requires

$$\frac{g}{\bar{u} - \underline{u}} \leqslant \frac{(\lambda_q - \lambda_p)\alpha}{\alpha\lambda_p + r} + O(1/\gamma). \tag{3.3}$$

As expected, punishing with probability $\alpha < 1$ when an event occurs, instead of probability 1, makes the incentive constraint harder to satisfy, because it weakens punishment, so if there is a non-trivial limit equilibrium, there is a non-trivial limit equilibrium with $\alpha = 1$. When

$$\frac{g}{\bar{u} - \underline{u}} < \frac{(\lambda_q - \lambda_p)}{\lambda_p + r}, \tag{3.4}$$

incentives for playing $+1$ can be provided by setting $\alpha < 1$, and since the equilibrium payoff is monotone in α, the highest limit equilibrium payoff can be found by setting (3.3) to hold with equality (ignoring the $O(1/\gamma)$ term) to solve for the optimal α^*, and then substituting α^* into the formula for $\bar{v}^{\alpha}$. Because $\alpha^* = rg/((\bar{u} - \underline{u})(\lambda_q - \lambda_p) - \lambda_p g)$, and $\bar{v}^{\alpha} = \bar{u}r + \alpha\underline{u}\lambda_p/(\alpha\lambda_p + r)$, the interest rate r factors out of the resulting formula, and we have

$$\bar{v}^{\alpha} = \frac{\bar{u}r + \alpha\underline{u}\lambda_p}{\alpha\lambda_p + r} = \bar{u} - \frac{g\lambda_p(\bar{u} - \underline{u})}{g\lambda_p + (\bar{u} - \underline{u})(\lambda_q - \lambda_p) - \lambda_p g} = \bar{u} - \frac{g\lambda_p}{(\lambda_q - \lambda_p)}. \tag{3.5}$$

Thus when (3.4) holds for a fixed r, there is a sequence of equilibria with payoffs uniformly above the static Nash level. Consequently there is a non-trivial limit equilibrium whenever (3.4) holds with $r = 0$, that is when

$$\frac{g}{\bar{u} - \underline{u}} < \frac{(\lambda_q - \lambda_p)}{\lambda_p};$$

and conversely regardless of r, there are no non-trivial limit equilibria when

$$\frac{g}{\bar{u}-\underline{u}} \geqslant \frac{(\lambda_q-\lambda_p)}{\lambda_p}.$$

So far we have only discussed the payoffs in Table 1, but the extension to the payoffs in Table 2 is immediate: we simply need to replace $\underline{u}$ with $\underline{u}+g$ in both the definition of $\bar{v}^{\alpha}$ – which becomes

$$\frac{\bar{u}r+\alpha(\underline{u}+g)\lambda_p}{\alpha\lambda_p+r}, \tag{3.6}$$

and in the incentive constraint for large γ and fixed r,

$$\frac{g}{\bar{u}-\underline{u}} < \frac{(\lambda_q-\lambda_p)\alpha}{\lambda_p+r},$$

which becomes

$$\frac{g}{\bar{u}-\underline{u}-g} < \frac{(\lambda_q-\lambda_p)\alpha}{\lambda_p+r}. \tag{3.7}$$

Once again, solving the program defining the highest equilibrium payoff for α^* shows that the optimal punishment probability is proportional to r, so that r factors out of the formula for the highest limit equilibrium payoff, which is

$$\bar{u}-\frac{g\lambda_p}{(\lambda_q-\lambda_p)},$$

just as it would be with the deterministic periods of FL, and the conclusion of the proposition follows from considering when (3.7) can be satisfied for small values of r. □

Inspection of the proof shows that the high-frequency limit of synchronous exponential moves coincides with the high-frequency limit of the usual discrete time model for any interest rate, and not just for interest rates near 0. However, the main points of the present paper can be seen from the case of interest rates near zero, and stating the results for this case is simpler and facilitates the comparison with past work.

3.3. Asynchronous exponentially distributed moves

We maintain the assumption that Z_t is a Poisson process, and now suppose the observation and action times of the two players are i.i.d. exponential waiting-time processes, each with intensity γ. That is, at times $t_{i,1}, t_{i,2}, \ldots,$ player i observes the value of Z_t, the times at which all events occurred, and then chooses an action. We continue to suppose that each time an event occurs, it is accompanied by a realization of a public randomizing device, and these realizations are observed by the players together with the corresponding events.

An adaptation of the α-grim strategy to this private monitoring setting turns out to do as well in the limit as it did with synchronous signals and moves. (Note that it is still the case that the probability of two or more events between observations is $O(1/\gamma^2)$, so there is no point in having a cutoff at 2 events or more; the only relevant strategies for large γ are those that punish with some probability when a single event is observed and punish with probability one contingent on two or more events.)

We will establish the following result:

Proposition 4. *Suppose that there is a non-trivial limit equilibrium of the games with simultaneous exponentially distributed moves as $\gamma \to \infty$. With either the payoffs of Table* 1 *or Table* 2*, the limit equilibrium payoffs with asynchronous exponentially distributed moves* (*and exponentially distributed observation times*) *are at least as high.*

Proof. Consider first the payoffs of Table 1. We will show that a small modification of the grim strategies outlined at the beginning of the previous section is an equilibrium under the same parameter restrictions as with synchronous signals, and supports the same limit equilibrium payoffs. The modification is necessary because if player 2 observes the state at t, and sees that an event occurred at $t-\rho$, where ρ is small compared to $1/\gamma$, then player 2 knows it is unlikely that player 1 has yet observed an event. Because by assumption player 2 is completely myopic, player 2 will not play "Out" in this situation; so, we modify player 2's strategy to be incentive compatible, while maintaining the same player 1 strategy as in the previous subsection. Given player 1's strategy, player 2 is indifferent between In and Out when $1-\exp(-\gamma\rho)=.5$ (equivalently $\rho=\ln(2)/\gamma$), so there is probability $1/2$ that player 1 has already seen the event and is playing -1. For all smaller ρ player 2 prefers playing In.[11]

[11] This is the point where we use the assumption that the time of the event is observed. Otherwise, when player 2 sees an event, instead of conditioning on whether the event was more or less than $\ln(2)/\gamma$ time units ago, the player needs to form beliefs about the time of the event. These beliefs would depend on the time that player 2 observed the process previously, and the setting would be much less tractable.

As in Section 3.1, fix a probability α such that LR switches to -1 with this probability when observing a Poisson event, and say that a "punishment event" occurs at time t when a Poisson event is realized at t and the value of the accompanying public randomization is less than or equal to α. We consider strategies of the following form: player 1 starts out playing $+1$ and sticks with it until the first time she sees a punishment event, or she sees player 2 play Out; once player 1 starts playing -1, she plays -1 from then on. Player 2 starts out playing In, and plays In up to the time $\ln(2)/\gamma$ after the occurrence of a punishment event that she has observed – that is, if the event occurs at t, 2 switches to playing Out at the first date $t' > t + \ln(2)/\gamma$ at which she takes an action.

For this profile to be an equilibrium, it is necessary that player 1 prefers her equilibrium strategy to the alternative of always playing -1. Moreover, due to the stationarity of player 2's strategy, if deviating to "-1 forever" is not profitable, then no other strategy for player 1 can yield an improvement.

Suppose player 1 conforms to this strategy, then her payoff is $\bar{u}$ initially, and stays at $\bar{u}$ at least until the first time t an event occurs. If player 1 observes the event before t', where t' is the first time after $t + \ln(2)/\gamma$ that player 2 observes the event, then player 1's flow payoff is $\bar{u} + g$ from the time she observes the event until t', and $\underline{u}$ thereafter. Otherwise, player 1 gets flow payoff $\bar{u}$ until t' and $\underline{u}$ thereafter. The resulting formula for player 1's equilibrium payoff v^{asynch} is complicated, but it has the same limit as in the synchronous case as $\gamma \to \infty$; in fact,

$$v^{asynch} = \frac{\bar{u}r + \alpha\underline{u}\lambda_p}{\alpha\lambda_p + r} + O(1/\gamma).$$

Similarly, the waiting time for observations is also asymptotically unimportant for the payoff to deviating; the delay that is important is the waiting time $1/\lambda_p$ or $1/\lambda_q$ for the Poisson event. That is, playing -1 gives an immediate gain of g, with no cost until the Poisson event occurs and is observed. Playing -1 increases the arrival rate of the bad event, and so reduces the expected value of the time T when flow payoffs fall from $\bar{u}$ to $\underline{u}$. In equilibrium, the expected discounted value of reduction in T must be large enough to offset the gain of g until T occurs. The time T is the first time when (a) an event has occurred and (b) has been observed by player 2; since the waiting time for an observation is $O(1/\gamma)$, this second waiting time is negligible when γ is large. Thus, we can compute that the payoff to the deviation "always play -1" is

$$\frac{r(\bar{u} + g) + \alpha\lambda_q\underline{u}}{\alpha r + \lambda_q} + O(1/\gamma),$$

and algebra shows that the incentive constraint for the grim strategy is the same as before up to $O(1/\gamma)$:

$$\frac{g}{\bar{u} - \underline{u}} < \frac{\alpha(\lambda_q - \lambda_p)}{\lambda_p + r}.$$

It is therefore sufficient for the existence of non-trivial equilibrium that

$$\frac{g}{\bar{u} - \underline{u}} < \frac{(\lambda_q - \lambda_p)}{\lambda_p + r},$$

which is the same condition as with the limit of synchronous exponentially distributed Poisson signals; and for $r < \bar{r}$ as in the definition of non-trivial limit equilibrium, it is sufficient that

$$\frac{g}{\bar{u} - \underline{u}} < \frac{(\lambda_q - \lambda_p)}{\lambda_p}.$$

Next, we consider what happens with the payoffs of Table 2. As before, the main change this makes is that the punishment payoff is now $\underline{u} + g$, so that the incentive constraint for the grim strategy is

$$\frac{g}{\bar{u} - \underline{u} - g} < \frac{(\lambda_q - \lambda_p)}{\lambda_p + r}.$$

We conclude that the best equilibrium payoff with asynchronous signals and either the payoffs of Table 1 or Table 2 is at least as high as in FL. □

4. Diffusion signals

Now, we consider the case where the underlying signal process Z_t is a diffusion process, and when players observe the process, they see its current level. We have in mind that the underlying signal is the aggregate of many small frequent events, and that the diffusion process is an approximation of the distribution of this aggregate.[12] Because Z_t is a diffusion, its increments $Z_{kt} - Z_{(k-1)t}$ are normally distributed with variance proportional to t; the variance is $\sigma_{+1}^2 t$ and $\sigma_{-1}^2 t$ as the

[12] If players observe the infinite-dimensional time path of the process, instead of its current position, they can learn its volatility from an arbitrarily short sample, but this seems impractical, and in many cases only the aggregate is publicly available.

action chosen is $a_1 = +1$ or -1. To avoid unimportant details, we simplify by setting the drift of the diffusion to be 0 regardless of the actions played.[13] We consider the "good news case" where $1 < \sigma_{+1}^2/\sigma_{-1}^2$. For the payoffs in Table 1, FL show (see Propositions 5 and 6) that there exist $\bar{\lambda}, \underline{\lambda} > 1$ such that if $1 < \sigma_{+1}^2/\sigma_{-1}^2 < \underline{\lambda}$, there is no non-trivial PPE in the limit of high-frequency synchronous play; and such that if $\sigma_{+1}^2/\sigma_{-1}^2 > \bar{\lambda}$, there is a non-trivial limit PPE. Moreover, the best limit PPE payoffs can be obtained with two-sided cut-off strategies of the form: Start out playing $+1$, In; once SR is observed to play Out, or player $i = 1, 2$ observes at any time an increment from $[\underline{y}, \bar{y}]$, she plays -1 or Out, respectively, from then on.

However for any fixed $\sigma_{+1}^2/\sigma_{-1}^2 > 1$, all limit PPE payoffs are bounded away from efficiency.[14] It is clear from inspection of the proof that these results extend to the payoff of Table 2; but as we will see, the two payoff functions have different limit equilibria once monitoring is asynchronous.

4.1. Diffusion process monitored at almost identical times

Now, suppose that each player i observes the state of the diffusion process at times $t - \varepsilon_{i,1}, 2t - \varepsilon_{i,2}, \ldots, kt - \varepsilon_{i,k}, \ldots,$ where $(\varepsilon_{i,k})_{k=1}^{\infty}$ is a sequence of i.i.d. random variables distributed uniformly on an interval $[0, \varepsilon]$ with $\varepsilon \geqslant 0$. Thus, at time $kt - \varepsilon_{i,k}$ player i learns the increment $y_{kt,i} = Z_{kt-\varepsilon_{i,k}} - Z_{(k-1)t-\varepsilon_{i,k-1}}$; she also observes the action taken at time $(k-1)t - \varepsilon_{2,k-1}$ by player 2. Player 1's private history is the sequence of increments $y_{kt,1}$ and actions taken by player 2, along with her own past actions; player 2's private history is simply the sequence of increments $y_{kt,2}$ and her own past actions. Because the normal distribution has a continuous density, in this section we do not need to allow for public randomizations.

We consider the stage-game payoffs of Table 2, and continue to focus on the "good news case" where $1 < \sigma_{+1}^2/\sigma_{-1}^2$. Moreover, we will restrict attention to a particular class of equilibria, namely equilibria in *history-dependent cutoff strategies*. These are strategies of the form: Start out playing $+1$, In; once SR is observed to play Out, or player $i = 1, 2$ observes at time $kt - \varepsilon_{i,k}$ an increment from $[\underline{y}_i(\tilde{h}_i^{k-1}), \bar{y}_i(\tilde{h}_i^{k-1})]$, she plays -1 or Out, respectively, from then on. Notice that the cutoffs are allowed to depend on player i's *extended* private history $\tilde{h}_i^{k-1}$, which comprises the history h_i^{k-1}, i.e., the increments observed at all previous times, the exact time of observing these increments, and in the case of LR, also on the actions played at all previous times; and the extended history also includes $kt - \varepsilon_{i,k}$, the exact time at which the player observes the current signal and chooses the current action.

It is known that the restriction to cutoff strategies is without loss of generality in the case of synchronous moves, but with asynchronous moves the status of this restriction is open. Note that in the literature on repeated games with private monitoring, there are no characterizations of the full equilibrium set except when the folk theorem holds – in other cases the existing results are all for specific classes of equilibria (e.g. Ely et al., 2005 and Yamamoto, 2009).

Proposition 5. *Consider the payoff of Table 2, and suppose $1 < \sigma_{+1}^2/\sigma_{-1}^2$. Consider the limit of LR's payoffs in equilibria in history-dependent cutoff strategies in the asynchronous case, in the iterated limit where first $\varepsilon \to 0$ and then $t \to 0$.*

(i) *If*

$$\frac{\sigma_{+1}}{\sigma_{-1}} < \underline{\lambda}, \tag{4.1}$$

then there is no non-trivial equilibrium in trigger and cutoff strategies in the limit.

(ii) *If*

$$\frac{\sigma_{+1}}{\sigma_{-1}} > \bar{\lambda},$$

then in the synchronous case, there are non-trivial limit payoffs. The limit of the highest LR's equilibrium payoffs in history-dependent cutoff strategies in the asynchronous case is strictly lower than the limit of the highest LR's PPE payoffs when signals are synchronous.[15]

Note that here we restrict attention to the iterated limit, while our result on this sort of asynchronicity and Poisson signals applies to any double limit where $\varepsilon/t \to 0$ and $t \to 0$; this restricted scope still permits us to show that a discontinuity can arise. Note also that part (ii) leaves open the question of whether there are any non-trivial equilibria. When signals are asynchronous, the monitoring is private and neither almost perfect nor almost public. We are unaware of any characterizations of the entire set of equilibrium payoff vectors for such information structures.

---

[13] In applications, it will often seem natural for the players' actions to influence the drift of the diffusion as well as its volatility, but in the high-frequency limit the effect of actions on drift does not matter. Intuitively, this is because the "signal to noise ratio" for strategies based on the difference in drift goes to 0.

[14] If $\sigma_{+1}^2/\sigma_{-1}^2 < 1$, there are limit equilibria with first-best payoffs.

[15] Recall that the highest limit of LR's PPE payoffs in the synchronous case can be attained in cutoff strategies.

Proof. Part (i) follows from FL and upper-hemi continuity of the set of equilibrium payoffs with respect to ε.

We now show part (ii). Suppose that in the model with asynchronous signals, the limit of the highest equilibrium payoffs in cutoff strategies exceeds $\underline{u}+g$. (The static equilibrium can be seen as a degenerate cutoff equilibrium, so the set of history-dependent cutoff strategy equilibria is not empty. If the limit is equal to $\underline{u}+g$, part (ii) follows directly from the fact there are non-trivial equilibria with synchronous signals.)

Given a sequence of non-trivial equilibria for any sequence of t's, which converge to zero, and any sequence of (sufficiently small) ε's, we shall first construct a sequence of equilibria in the synchronous case (for the same sequence of t's) with the property that the limit LR's payoff to the equilibria in the synchronous case is at least as high as the limit LR's payoff to the equilibria in the asynchronous case; and next, we will construct other equilibria in the synchronous case with a strictly higher LR's payoff in the limit.

Let for $k=1,2,\dots,$

$$\underline{y}^k(t)=E^{\tilde{h}_1^{k-1}}\underline{y}_1\big(\tilde{h}_1^{k-1}\big)\quad\text{and}\quad\bar{y}^k(t)=E^{\tilde{h}_1^{k-1}}\bar{y}_1\big(\tilde{h}_1^{k-1}\big),$$

where the expectation over extended histories $\tilde{h}_1^{k-1}$ is taken with respect to the equilibrium strategies.[16] Since SR is indifferent between playing In and Out exactly when she assigns equal probabilities to LR playing $+1$ and -1,

$$E^{\tilde{h}_2^{k-1}}\underline{y}_2\big(\tilde{h}_2^{k-1}\big)-E^{\tilde{h}_1^{k-1}}\underline{y}_1\big(\tilde{h}_1^{k-1}\big)\underset{\varepsilon\to 0}{\to}0$$

and

$$E^{\tilde{h}_2^{k-1}}\overline{y_2}\big(\tilde{h}_2^{k-1}\big)-E^{\tilde{h}_1^{k-1}}\overline{y_1}\big(\tilde{h}_1^{k-1}\big)\underset{\varepsilon\to 0}{\to}0.$$

Notice that the cutoffs $\underline{y}^k(t)$ and $\bar{y}^k(t)$, $k=1,2,\dots,$ define cutoff strategies in the synchronous case. The cutoffs depend on date kt, but are independent of the (public) histories up to date kt. (The cutoffs that maximize LR's payoff in the synchronous case are also independent of the histories, and even independent of the date.)

To verify incentives to play these strategies, we only need to show that LR prefers action $+1$ to action -1 when she observes an increment from the complement of $[\underline{y}^k(t),\bar{y}^k(t)]$, given that SR has played In at all previous times; other incentive conditions are satisfied in an obvious manner. This will in fact establish the upper-hemi continuity of the set of LR's equilibrium payoffs in trigger and cutoff strategies.

Let $v^k(t)$ denote LR's expected continuation payoffs from time kt, $k=0,1,\dots,$ contingent on SR playing In at date kt, and all previous dates, and LR playing $+1$ at date kt. Let $\bar{p}^k(t)$ and $\bar{q}^k(t)$ denote the probability of switching to the stage-game Nash equilibrium at date $k+1$, contingent on playing $+1$ and -1, respectively, at time kt.

The difference between LR's instant payoff to playing -1 and $+1$ between time kt and time $(k+1)t$ is equal to $(1-\delta)g$, and the difference in continuation payoffs is equal to

$$\delta\big[\bar{q}^k(t)-\bar{p}^k(t)\big]\big[v^{k+1}(t)-\underline{u}-g\big].$$

As $\varepsilon\to 0$, all cutoffs $\underline{y}_1(\tilde{h}_1^{k-1})$ and $\underline{y}_2(\tilde{h}_2^{k-1})$ converge to $\underline{y}^k(t)$, and all cutoffs $\bar{y}_1(\tilde{h}_1^{k-1})$ and $\bar{y}_2(\tilde{h}_2^{k-1})$ converge to $\bar{y}^k(t)$. At any cutoff $\underline{y}_1(\tilde{h}_1^{k-1})$ and $\bar{y}_1(\tilde{h}_1^{k-1})$, LR is indifferent between playing $+1$ and -1. And, as $\varepsilon\to 0$, the difference between LR's instant payoff to playing -1 and $+1$ between time $kt-\varepsilon_{1,k}$ and time $(k+1)t-\varepsilon_{1,k+1}$ is equal to $(1-\delta)g$, and the difference in continuation payoffs (at the cutoffs) converges to

$$\delta\frac{1}{2}\big[\bar{q}^k(t)-\bar{p}^k(t)\big]\big[v^{k+1}(t)-\underline{u}-g\big].$$

The factor $1/2$ comes from the fact that at any cutoff, LR assigns a probability converging to $1/2$ to the event that SR plays Out at $kt-\varepsilon_{2,k}$, in which case the action of LR does not affect the continuation payoff. This yields that

$$(1-\delta)g=\delta\frac{1}{2}\big[\bar{q}^k(t)-\bar{p}^k(t)\big]\big[v^{k+1}(t)-\underline{u}-g\big]$$

at $\underline{y}^k(t)$ and $\bar{y}^k(t)$. Therefore LR prefers action $+1$ to action -1 when she observes an increment from the complement of $[\underline{y}^k(t),\bar{y}^k(t)]$, given that SR has played In at all previous dates.

The LR payoff in this synchronous-case equilibrium does not exceed the payoff in the cutoff equilibrium such that players play ($+1$, In) when they observe an increment from the complement of $[-y(t),y(t)]$, where $y(t)$ is such that

$$(1-\delta)g=\delta\frac{1}{2}\big[q(t)-p(t)\big]\big[v(t)-\underline{u}-g\big]\tag{4.2}$$

and

[16] The cutoffs could equally well have been defined in terms of the expected values of $\underline{y}_2(\tilde{h}_2^{k-1})$ and $\bar{y}_2(\tilde{h}_2^{k-1})$.

$$v(t) = (1-\delta)\bar{u} + \delta\big[p(t)(\underline{u}+g) + \big(1-p(t)\big)v(t)\big], \tag{4.3}$$

where

$$\zeta(t) = \frac{y(t)}{\sigma_{+1}t^{1/2}}$$

and

$$p(t) = 1 - 2\Phi\big(-\zeta(t)\big) \quad \text{and} \quad q(t) = 1 - 2\Phi\left(-\frac{\sigma_{+1}}{\sigma_{-1}}\zeta(t)\right),$$

i.e., $p(t)$ and $q(t)$ denote the probabilities that the increment observed at $k+1$ will belong to $[-y(t), y(t)]$ contingent on playing $+1$ and -1, respectively.

Indeed, it follows from (4.2) and (4.3) that

$$v(t) = \bar{u} - \frac{2p(t)g}{q(t)-p(t)}. \tag{4.4}$$

Therefore, the highest $v(t)$ among all equilibria in cutoff strategies that satisfy condition (4.2) is obtained by choosing the interval of signals that trigger actions $(-1, \text{Out})$ to maximize the "signal to noise" ratio

$$\frac{q(t)-p(t)}{p(t)}.$$

This highest "signal to noise" ratio is attained for the interval $[-y(t), y(t)]$, because given $q(t)-p(t)$, the lowest $p(t)$ is uniquely attained at such an interval.[17]

The optimal synchronous-case cutoff $y^*(t)$ is, in contrast, determined by

$$(1-\delta)g = \delta\big(q^*(t)-p^*(t)\big)\big(v^*(t)-\underline{u}-g\big), \tag{4.5}$$

and

$$v^*(t) = (1-\delta)\bar{u} + \delta\big[p^*(t)(\underline{u}+g) + \big(1-p^*(t)\big)v^*(t)\big],$$

i.e.,

$$v^*(t) = \frac{(1-\delta)\bar{u} + \delta p^*(t)(\underline{u}+g)}{1-\delta(1-p^*(t))},$$

where

$$\zeta^*(t) = \frac{y^*(t)}{\sigma_{+1}t^{1/2}}$$

and

$$p^*(t) = 1 - 2\Phi\big(-\zeta^*(t)\big) \quad \text{and} \quad q^*(t) = 1 - 2\Phi\left(-\frac{\sigma_{+1}}{\sigma_{-1}}\zeta^*(t)\right).$$

Denote by v^* the limit of $v^*(t)$, and by v the limit of $v(t)$. It follows from the equation for $v^*(t)$ that

$$\frac{p^*(t)}{1-\delta} \underset{t\to 0}{\to} \frac{\bar{u}-v^*}{v^*-\underline{u}-g}.$$

Since

$$\zeta^*(t) \underset{t\to 0}{\to} 0,$$

and

$$q^*(t) = 1 - 2\Phi\left(\frac{\sigma_{+1}}{\sigma_{-1}}\Phi^{-1}\left(\frac{1-p^*(t)}{2}\right)\right),$$

applying l'Hopital's rule,

[17] More precisely, suppose that highest "signal to noise" ratio is attained for a different interval. Take the interval $[-y, y]$ with the same $q(t)-p(t)$, and thus a lower $p(t)$. Thus the value of $v(t)$ computed by (4.4) increases. This implies that the $q(t)-p(t)$ must be even lower, and $v(t)$ even higher, in the cutoff strategies that satisfy condition (4.2), and punishment is triggered by increments from interval $[-y(t), y(t)]$.

$$\frac{q^*(t)}{1-\delta} \underset{t\to 0}{\to} \frac{\sigma_{+1}}{\sigma_{-1}} \cdot \frac{\bar{u}-v^*}{v^*-\underline{u}-g}.$$

Thus, by (4.5),

$$g = \left(\frac{\sigma_{+1}-\sigma_{-1}}{\sigma_{-1}}\right)(\bar{u}-v^*).$$

Similarly, referring to (4.2),

$$g = \frac{1}{2}\left(\frac{\sigma_{+1}-\sigma_{-1}}{\sigma_{-1}}\right)(\bar{u}-v).$$

(Remember that $v^* > \underline{u}+g$ whenever $\sigma_{+1}/\sigma_{-1} > \bar{\lambda}$.) Therefore, $v < v^*$. □

This proposition shows that asynchronous monitoring lowers the equilibrium payoffs to cutoff strategies for the payoffs of Table 2. With the payoffs of Table 1, Eq. (4.2) in the proof becomes

$$(1-\delta)\frac{1}{2}g = \delta\frac{1}{2}\big[q(t)-p(t)\big]\big[v(t)-\underline{u}-g\big], \tag{4.6}$$

so the $\frac{1}{2}$ factors out, and the proof unravels. This appears, however, to be a specific feature of the payoffs of Table 1. Specifically, the best limit payoff is lower with asynchronous moves whenever (−1, Out) is a strict stage-game Nash equilibrium, i.e., the LR payoff of this action profile exceeds $\underline{u}$. In this case the left-hand side of (4.6) would then be strictly higher than $(1-\delta)\frac{1}{2}g$.

4.2. Almost-public synchronous monitoring

Some intuition for the result on diffusions with asynchronous monitoring can be gained from a simpler model of private monitoring, that of synchronous observation of a noisy private signal of the diffusion process. That is, suppose players move and act simultaneously, but that each player i observes the state of the diffusion with noise. More specifically, suppose that at each time t, player i observes $Z_t+\varepsilon_i$, where the ε_i are normal random variables with mean 0 and variance ε^2.

Now suppose that play corresponds to a cutoff equilibrium with cutoffs $[\underline{y},\bar{y}]$ and consider play in period 1. Since the strategies specify that player 1 plays +1 at time 0, each player believes that Z_1 is normally distributed with mean 0 and variance σ_{+1}^2. Thus conditional on $y_{1,i}=Z_1+\varepsilon_i$, player i believes Z_1 is normal with mean $y_{1,i}/(1+\varepsilon^2\sigma_{+1}^2)$ and variance $\varepsilon^2\sigma_{+1}^2/(\varepsilon^2+\sigma_{+1}^2)$, and so i believes that the opponent's signal $y_{1,-i}$ is normal with mean $y_{1,i}/(1+\varepsilon^2\sigma_{+1}^2)$ and variance $\varepsilon^2+\varepsilon^2\sigma_{+1}^2/(\varepsilon^2+\sigma_{+1}^2)=\varepsilon^2(\varepsilon^2+2\sigma_{+1}^2)/(\varepsilon^2+\sigma_{+1}^2)$. Consequently, if player i observes a signal just equal to $\underline{y}$, she thinks the opponent's signal is normal with mean $\underline{y}/(1+\varepsilon^2\sigma_{+1}^2)$ and variance $\varepsilon^2(\varepsilon^2+2\sigma_{+1}^2)/(\varepsilon^2+\sigma_{+1}^2)$, so she assigns probability tending to $\frac{1}{2}$ as $\varepsilon\to 0$ to the opponent's signal being below $\underline{y}$.[18]

This is quite different than the case of $\varepsilon=0$, in which the player knows which action will be played by the opponent, and in fact the actions that will be played following such signals are not almost common knowledge. For this reason, the costs of "miscoordination" enter into the player's incentive constraints, and so when player 1 decides whether to play −1 or +1, she must consider the probabilities of both actions of player 2. This intuition is quite similar to that behind the anti-robustness result of Sugaya and Takahashi (2009) for the effect of a small amount of private noise on the equilibria of the repeated prisoner's dilemma with imperfect public information.

One may wonder if the discontinuity in this noisy synchronous-move model arises from the fact that the diffusion process has an uncountable number of possible signals. Suppose that the diffusion process is replaced with a stochastic process that takes a finite number of values, and that the probability assigned by a player i to the event that player $j\neq i$ receives signal s, contingent on player i receiving signal s herself, tends to 1 as ε tends to 0. Then, by analogy with results on strict, bounded memory equilibria for games with two long run players, we conjecture that for any fixed time period t the set of LR's equilibrium payoffs converges as ε converges to 0 to the set of LR's equilibrium payoffs for $\varepsilon=0$. Our intuition is that the perfectly observable actions of one player play a role somewhat similar to bounded memory.

To relate this observation about finite-support signals to the monitoring of diffusion signals, note that we can approximate the diffusion processes for each action by arrays that converge to the diffusions as the time period goes to 0, as in Fudenberg and Levine (2009). That is, we divide each interval $[kt,(k+1)t]$ into m subintervals. In each subinterval, a single, positive or negative event may arrive. Positive events count as +1 and negative events count as −1, and players observe at time $(k+1)t$ only the total number of events that arrived in interval $[kt,(k+1)t]$. Now modify their model by assuming that LR observes the sum of the events at times $(k+1)t$ with a little noise, which is multinomial on $\{-c/m,\dots,-1/m,0,1/m,\dots,c/m\}$. In this case, the signals are almost public along the limit if the probability that the noise is equal to 0 converges to 1, so players are almost certain to observe the same thing. Here too we would expect that

[18] The probability is $\Phi(\underline{y}\varepsilon\sigma_{+1}^2(\varepsilon^2+\sigma_{+1}^2)^{1/2}/(1+\varepsilon^2\sigma_{+1}^2)(\varepsilon^2+2\sigma_{+1}^2)^{1/2})$.

the limit equilibria converge to the same limit as in the game without the noise. Conversely, if the signals do not become almost public along the limit, we would expect cutoff equilibria to do poorly. More speculatively, we conjecture that the same "asymptotically almost public information" condition is the key for the difference between the limits of synchronous and asynchronous monitoring of arrays that converge to diffusions.

5. Conclusion

This paper is only a first look at the implications of asynchronous imperfect monitoring. We found that high-frequency limits of Poisson monitoring are relatively robust, but that the high-frequency limit of monitoring a diffusion is more sensitive to the details of the monitoring structure. Our sense is that this difference stems from the fact that in the Poisson case punishment events become common knowledge with a short expected lag, while this is not the case with diffusion signals, in which case punishment increments observed by one player may never be observed by the other. It would of course be interesting to extend our results on high-frequency limits to more general games, but this seems difficult and indeed the problem of characterizing high-frequency limits even with synchronous moves is still open.[19] It would also be interesting to study the low-interest-rate limit of games with asynchronous imperfect monitoring; if we allow for monitoring to occur at any time t, this amounts to an extension of the stochastic games literature to a continuum of states and a particular kind of private monitoring.

References

Abreu, D., Milgrom, P., Pearce, D., 1991. Information and timing in repeated partnerships. Econometrica 59, 1713–1733.
Abreu, D., Dutta, P., Smith, L., 1994. Folk theorems for repeated games: A NEU condition. Econometrica 62, 939–948.
Ambrus, A., Lu, S., 2009. A continuous-time model of multilateral bargaining. Mimeo.
Dutta, P., 1995. A folk theorem for stochastic games. J. Econ. Theory 66, 1–32.
Ely, J.C., Hörner, J., Olszewski, W., 2005. Belief-free equilibria in repeated games. Econometrica 73, 377–416.
Faingold, E., Sannikov, Y., 2007. Equilibrium degeneracy and reputation effects. Mimeo.
Fudenberg, D., Levine, D.K., 1994. Efficiency and observability with long-run and short-run players. J. Econ. Theory 62, 103–135.
Fudenberg, D., Levine, D.K., 2007. Continuous time limits of repeated games with imperfect public monitoring. Rev. Econ. Dynam. 10, 173–192.
Fudenberg, D., Levine, D.K., 2009. Repeated games with frequent actions. Quart. J. Econ. 124, 149–195.
Fudenberg, D., Maskin, E., 1986. The folk theorem in repeated games with discounting or with incomplete information. Econometrica 54, 533–556.
Fudenberg, D., Yamamoto, Y., 2009. The folk theorem for irreducible stochastic games with imperfect public monitoring. Mimeo.
Fudenberg, D., Kreps, D.M., Maskin, E., 1990. Repeated games with long-run and short-run players. Rev. Econ. Stud. 57, 555–573.
Fudenberg, D., Levine, D.K., Maskin, E., 1994. The folk theorem with imperfect public information. Econometrica 62, 997–1040.
Hörner, J., Olszewski, W., 2009. How robust is the folk theorem with public monitoring. Quart. J. Econ. 124, 1773–1814.
Hörner, J., Sugaya, T., Takahashi, S., Vieille, N., 2009. Recursive methods in discounted stochastic games: An algorithm for $\delta \to 1$ and a folk theorem. Mimeo.
Kamada, Y., Kandori, M., 2009. Revision games. Mimeo.
Klein, B., Leffler, K., 1981. The role of market forces in assuring contractual performance. J. Polit. Econ. 89, 615–641.
Lagunoff, R., Matsui, A., 1997. Asynchronous choice in repeated coordination games. Econometrica 65, 1467–1477.
Mailath, G., Morris, S., 2002. Repeated games with almost-public monitoring. J. Econ. Theory 102, 189–228.
Mailath, G., Olszewski, W., 2008. Folk theorems with bounded recall under (almost) perfect monitoring. Mimeo.
Monderer, D., Samet, D., 1989. Approximating common knowledge with common beliefs. Games Econ. Behav. 1, 170–190.
Sannikov, Y., Skrzypacz, A., 2009. The role of information in repeated games with frequent actions. Mimeo.
Shapley, L., 1953. Stochastic games. Proc. Natl. Acad. Sci. USA 39, 1095–1100.
Sugaya, T., Takahashi, S., 2009. Coordination failure in repeated games with private monitoring. Mimeo.
Yamamoto, Y., 2009. A limit characterization of belief-free payoffs in repeated games. J. Econ. Theory 144, 802–824.

[19] Sannikov and Skrzypacz (2009) give an upper bound on the high-frequency limit of pure-strategy equilibria in games with all long-run players, but their bound is only tight when the interest rate goes to 0.