Repeated Games with Asynchronous Monitoring

Drew Fudenberg and Wojciech Olszewski

January 16, 2009

Abstract

We study two modifications of the standard model of repeated games with public signals. In one modification, players osberve an underlying public signal at random and privately known times, where the public signals are Poisson events and either the observations occur within a small epsilon time interval or the obsevations have an exponential waiting times. In the second modification, the players observe the position of a diffusion process with a small amount of noise. We show that in the Poisson cases the high-frequency limist are the same as in the Fudenberg and Levine [2007, 2008] study of high-frequency public signals, but that the limits can differ when the signals correspond to a diffusion.

1. Introduction

Standard models of repeated games with imperfect monitoring assume that monitoring and play are both exactly synchronous. This sort of simplifying assumption is particularly hard to motivate, even as an approximation, when signals and play occur at a very high frequency, so it seems important to investigate the extent to which past results about 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 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.

Past work has shown that asynchronous and synchronous play can differ with perfectly observed actions, even in the limit as the discount factor goes to 1.[1] Asynchronous imperfect observations have a different effect: they lead to a particular form of private monitoring, as neither player knows the signals that the other player observed. When the period lengths are long, there is no reason to expect the same equilibria with synchronous and asynchronous signals, but it is still possible that the asynchronicity makes no difference when signals are observed with sufficiently high frequency; this is the subject of our study.

As in Fudenberg-Levine (FL) [2007,2008] and Faingold and Sannikov [2007], we study the high-frequency limit of 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 this game, the best limit equilibrium payoffs with 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, which makes it easier to determine the way that the equilibrium set varies with the information structure.

We begin with the case where the signal process is the number of Poisson events that have occurred. Here we consider two formulations of almost synchronous high frequency signals. In the first, each player observes the state, and revises their action, at times $t-\varepsilon_{i,1}, 2t-\varepsilon_{i,2}, , ..., kt-\varepsilon_{i,k}, ...$, where $(\varepsilon_{i,k})_{k=1}^{\infty}$ is a sequence of i.i.d. random variables distributed uniformly on an interval $[0, \varepsilon]$ with $\varepsilon \geq 0$. In this case, as $\varepsilon \to 0$ for any fixed t the observation times are highly correlated, 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 small generalization that makes the static equilibrium strict instead of weak.

Next we consider the case where 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 study the case where the observation times of the players are perfectly correlated, so that they jointly observe the Poisson process at exponentially distributed waiting times. In this case the condition for the existence of a non-trivial equilibrium differs slightly from that with deterministic periods, even in the "high-frequency" limit of $\gamma \to \infty$, but when the condition is satisfied the highest limit equilibrium payoff is exactly as in FL 2007, both for the

---

[1] Specifically, the "Anti-folk-theorem" of Lagunoff and Matsui [1997] combined with Dutta's [1995] folk theorem for stochastic games show that stronger full dimenionality conditions are needed in games with alternating moves than in the usual simultaneous-move case studied by Fudenberg and Maskin [1986] and Abreu, Dutta, and Smith [1994].

payoff matrix they considered and for our generalization. For a fixed value of γ, there is no reason to expect the equilibria of the synchronous case to resemble those of the "asynchronous" case where the observation times of the two players are independent. However, 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 the trigger strategies used involve a cutoff, and for signals near the cutoff the players are nearly indifferent. Because of these complications, which we discuss in more detail later on, we have only been able to analyze a somewhat artificial model, where both players make observations at fixed times $t, 2t, ...$, and the LR player observes the sum of the diffusion process and a noise term $\eta \sim N(0, \varepsilon^2\sigma^2)$, while the short-run players observe the state perfectly. In this model it turns out that the limit conclusions of FL apply to the payoff matrix considered by FL in the limit $t \to 0$, $\varepsilon/t \to 0$, but the fact that the LR player assigns equal probabilities to the true state being either less than or greater than his signal results in a lower limit equilibrium payoff than in FL for the generalized payoff matrix. As we explain, this is because with a continuously distributed signal (such as the position of a diffusion process) even with arbitrarily little observation noise, the signals are not ”almost public,” that is, neither player thinks it is likely that the other player observed exactly the same signal that he did.

2. Review of the Standard Discrete-Time Model

A long-run player 1 plays a stage game with a short-run player 2 who is completely impatient in a 2x2 stage game. FL consider the case of the payoff matrix in Table 1

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

Table 1

where neither player's payoff depends on player 1's action when player 2 plays Out. We will consider both this payoff matrix and the following small modification

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

Table 2.

where $\underline{u} + g < \overline{u}$. In both games, player 2 plays Out in every 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, when the length of a period is t, and the subjective continuous time interest rate for the long-run player is r, her rate of time discount per period is $\delta = \exp(-rt)$. There is also a publicly observed public randomization device each period before actions are taken. The public history is the history of the signal and the public randomization device. Our solution concept here is perfect public equilibrium or PPE (Fudenberg, Levine and Maskin [1994]): these are strategy profiles for the repeated game in which (a) each player's strategy depends only on the public information, and (b) no player wants to deviate at any public history. In both 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 the highest repeated game payoff is $\bar{u}$ when actions are observed (Fudenberg, Kreps, and Maskin [1990]) and the highest payoff with imperfect public monitoring is strictly less than that (Fudenberg and Levine [1994]).

FL [2007] 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 outcome 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 signals and the public randomizing device to continuation payoff.

Proposition: (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

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

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

As in FL (2007), assume that the expected waiting time for a Poisson event is λ_p if LR plays $+1$, and $\lambda_q > \lambda_p$ if LR plays -1. (This is the case when the event is "bad news;" FL show that there is no non-trivial limit equilibrium in the case of "good news.") FL show that highest limit equilibrium payoff can be achieved by strategies of the following form: Start out playing (+1, In); once SR is observed to play Out, or players observe a Poisson event, use the public randomizing device to switch to playing (-1, Out) from then on with some probability, and otherwise keep playing (+1, In.)[2] With the payoffs of Table 1, these strategies implement a non-trivial limit equilibrium if $\frac{g}{\bar{u}-\underline{u}} < \frac{(\lambda_q-\lambda_p)}{\lambda_p}$, when this is satisfied, the best limit equilibrium payoff is $v^* = \bar{u} - \frac{g\lambda_p}{\lambda_q-\lambda_p}$. With the payoffs of Table 2, the condition for a non-trivial limit equilibrium is that $\frac{g}{\bar{u}-\underline{u}-g} < \frac{(\lambda_q-\lambda_p)}{\lambda_p}$; the best limit equilibrium payoff is the same as with the payoffs of Table 1.

3.1. Poisson events observed at almost the same time

Now assume that players observe the state of the Poisson process at time $t-\varepsilon_{i,1}, 2t-\varepsilon_{i,2}, , ..., kt-\varepsilon_{i,k}, ...$, where $(\varepsilon_{i,k})_{k=1}^{\infty}$ is a sequence of i.i.d. random variables distributed uniformly on an interval $[0, \varepsilon]$ with $\varepsilon \geq 0$. To facilitate the analysis of optimal equilibria we will assume here that when players observe an event and accompanying

---

[2] Considering the likelihood ratios shows that the optimal equilibrium must take the form of a cut-off:: switch to (-1, Out) if more than c events are observed, switch with some probability α if exactly c events are observed, and don't 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.

realization of the public randomizing device, they observe the time at which the event occurred. (We make the latter assumption to make the analysis easier; we believe that the same results would obtain here without it, but more complicated strategies would be needed.

Proposition 1. *Consider the limit of payoffs of trigger strategy equilibria in the asynchronous case) as* $t \to 0$, $\varepsilon/t \to 0$. *Then every limit of equilibrium payoffs that can be attained in equilibrium by LR for* $\varepsilon = 0$ *can be approximated by LR's equilibrium payoffs for positive small values of* ε *for both the game from Table 1 of FL and the game from Table 2.*

Proof: Consider first the payoffs of Table 1, and suppose $\frac{g}{\bar{u}-\underline{u}} < \frac{(\lambda_q-\lambda_p)}{\lambda_p}$. We claim that a modification of the strategy profile from FL is an equilibrium for sufficiently small positive values of t and ε/t. 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 grim strategies, suppose that the public randomizing device is uniformly distributed on [0,1], 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. We consider strategies of the following form: player 1 starts out playing +1 and sticks with it until the first time he sees a punishment event or he sees player 2 play Out; once player 1 starts playing -1 he 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 at time kt. If she observes a punishment event arrived in interval $(kt-\frac{\varepsilon}{2}, kt]$ she switches to playing Out at time $(k+1)t$.

Consider first the case when LR observed no punishment event till time kt, and SR played In in periods $t, 2t, ..., (k-1)t$. 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 received between kt and $(k+1)t$ differs compared to the case of $\varepsilon = 0$ only by $O(\varepsilon/t)$.

The expected difference between the continuation payoffs from time $(k+1)t$ on, to playing $+1$ and -1 between times kt and $(k+1)t$ also changes only by $O(\varepsilon/t)$. Indeed, LR's action between times kt and $(k+1)t$ 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 $(k+2)t$. The consequences of LR's action between times kt and $(k+1)t$ 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 in periods $t, 2t, ..., (k-1)t$. Then, LR knows that SR will observe the event at time $(k+1)t-\varepsilon_{2,k+1}$. Even if the event arrived late and it is unlikely that SR observed it at time $kt-\varepsilon_{i,k}$, or SR observed at $kt-\varepsilon_{i,k}$ that it arrived late and she will play In at time kt, she will switch to playing Out no later than at time $(k+1)t$. Therefore, LR has no incentive to play $+1$.

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

As we remarked in the introduction, an important aspect of this model is that for sufficiently small values of ε, LR who observes no punishment event is almost certain that SR has observed no punishment event, and will play In at time $(k+1)t$. 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." Past work on games with all long-run players by Mailath and Morris (2002, 2006), Hörner and Olszewski (2008) 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 (2008) 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.)

However, our 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. It depend on when the event has arrived. The techniques of the papers on almost public monitoring are not directly applicable here.[3] 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.

3.2. Poisson events and Exponentially Distributed Observation times.

In the τ, ε model, the observation times of the players become highly correlated as $\varepsilon/\tau \to 0$. We will now suppose that 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

---

[3] 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.

make an observation. In this model, unlike the previous one, 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 instead of converging to 1.

We maintain the assumption that when players observe an event and accompanying realization of the public randomizing device, they observe the time at which the event occurred; as in the previous case, we believe that the same results would obtain without it. We will focus on the high-frequency limit as $\gamma \to \infty$. Even in the synchronous case, the high-frequency limit is slightly different than the short deterministic periods considered by FL, but the synchronous and asynchronous actions have the same high-frequency limit. 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 that the waiting time for the event to occur requires that an event be followed by a non-negligible increase in the probability of punishment.

3.2.1. Synchronous case

We begin with the synchronous case where players move and observe simultaneously and observations have exponential waiting time with parameter γ. 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 $\gamma \to \infty$. So consider the grim strategies "Play (+1, In) until an event occurs; when it does play (-1, Out) forever."

With the payoffs in Table 1, if this strategy is followed the payoff in the cooperative state is

$$
\begin{aligned}
v^* &= \int_{y=0}^{\infty}\int_{s=y}^{\infty} \lambda_p \exp(-\lambda_p y)\gamma \exp(-\gamma(s-y))(\bar{u}(1-\exp(-rs)) + \underline{u}\exp(-rs))dsdy \\
&= \frac{\bar{u}r(\gamma+\lambda_p+r)+\underline{u}\gamma\lambda_p}{(\lambda_p+r)(\gamma+r)}.
\end{aligned}
$$

Note that as $\gamma \to \infty$ this payoff 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 will instead make use of the public randomizing device to only punish with probability α if the event occurs (and otherwise stay in the cooperative phase.) This strategy has payoff $v^{*\alpha}$, where $v^{*\alpha}$ solves

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

so that

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

As $\gamma \to \infty$, this payoff converges to

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

The one-period cost of playing +1 is $\frac{gr}{\gamma+r}$; this decreases the probability of a regime change by $\frac{(\lambda_q-\lambda_p)\gamma\alpha}{(\lambda_p+\gamma)(\lambda_q+\gamma)}$, which yields a benefit of $v^{*\alpha} - \underline{u}$.

Thus the strategies satisfy the no-one-stage-deviations test if

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

For large gamma this requires

$$\frac{g}{\bar{u}-\underline{u}} <= \frac{(\lambda_q-\lambda_p)\alpha}{\alpha\lambda_p + r} + O(1/\gamma) \tag{3.1}$$

As expected using public randomization makes the IC 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

$$\frac{g}{\bar{u}-\underline{u}} < \frac{(\lambda_q-\lambda_p)}{\lambda_p + r} + O(1/\gamma) + \zeta \tag{3.2}$$

holds for some fixed $\zeta > 0$, then equilibrium is consistent with smaller values of α, and since the equilibrium payoff is monotone in α, the highest limit equilibrium payoff can be found by setting (3.1) to hold with equality (ignoring the $O(1/\gamma)$ term) to solve for the optimal value public randomizing probability and then substituting α^* into the formula for $\bar{v}^\alpha$. Because $\alpha^* = \frac{rg}{(\bar{u}-\underline{u})(\lambda_q-\lambda_p)-\lambda_p g}$, and $\bar{v}^\alpha = \frac{\bar{u}r+\alpha\underline{u}\lambda_p}{\alpha\lambda_p+r}$, the interest rate r factors out of the resulting formula, and we have

$$\bar{v}^\alpha = \frac{\bar{u}r+\alpha\underline{u}\lambda_p}{\alpha\lambda_p+r} = \bar{u} - \frac{g\lambda_p(\bar{u}-\underline{u})}{g\lambda_p+(\bar{u}-\underline{u})(\lambda_q-\lambda_p)-\lambda_p g} = \bar{u} - \frac{g\lambda_p}{(\lambda_q-\lambda_p)} \tag{3.3}$$

Note that the limit equilibrium payoff in (3.3) is exactly the same is in FL, but that the condition for a non-trivial equilibrium, (3.2) only reduces to the corresponding condition in FL's limit of frequent periods of deterministic length as $r \to 0$, while for non-zero interest rates, the condition (3.2) is more demanding. The interest rate appears in the incentive constraint (3.2) due to the dispersion of observation times, but when the the incentive constraint can be strictly satisfied, public randomization is used to reduce the frequency of punishment to the minimum required in a way that removes the dependence on r.

To gain some intuition into the relationship between this model and FL, it may be instructive to compare the model with observation rate γ to the deterministic-periods model with the same expected period length, namely $t = 1/\gamma$. The one-period cost of playing +1 here is $\frac{gr}{\gamma+r}$, while the corresponding cost in FL is $g(1-\exp(-r/\gamma))$; the ratio of these costs tends to one, so the same amount of payoff must be "sacrificed" in the two settings to give player 1 the right incentives. The fact that the necessary conditions for a non-trivial limit nevertheless differ in the two settings shows that the dispersion of observation times has a first-order effect on the effectiveness of punishment as $\gamma \to \infty$ even though it has only a second-order effect on the one-period cost/benefit ratio.

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

$$\frac{\bar{u}r+\alpha(\underline{u}+g)\lambda_p}{\alpha\lambda_p+r} \tag{3.4}$$

and in the constraint $\frac{g}{\bar{u}-\underline{u}} < \frac{(\lambda_q-\lambda_p)\alpha}{\lambda_p+r}$, which becomes

$$\frac{g}{\bar{u}-\underline{u}-g} < \frac{(\lambda_q-\lambda_p)\alpha}{\lambda_p+r} \tag{3.5}$$

Once again, solving this incentive constraint 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.

We summarize this discussion in the following result:

Proposition *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 $\frac{g}{\bar{u} - \underline{u}} < \frac{(\lambda_q - \lambda_p)}{\lambda_p + r}$. *The best limit equilibrium payoff is* $\bar{u} - \frac{g\lambda_p}{(\lambda_q - \lambda_p)}$, *just as in FL. With the payoffs of Table 2, the condition for a non-trivial limit is* $\frac{g}{\bar{u} - \underline{u} - g} < \frac{(\lambda_q - \lambda_p)}{\lambda_p + r}$, *and the best limit equilibrium payoff is again* $\bar{u} - \frac{g\lambda_p}{(\lambda_q - \lambda_p)}$.

3.2.2. Asynchronous signals

Now suppose the observation times of the two players are i.i.d. waiting time processes, each with intensity γ. It is still the case that the probability of two or more events between observations is $O(1/\gamma^2)$, so the only relevant strategies for large γ are those that punish with some probability when a single event is observed.

We will establish the following result:

Proposition 2. *With either the payoffs of Table 1 or Table 2, the limit equilibrium payoffs with asynchronous Poisson signals and exponentially distributed observation times are at least as high as with synchronous signals.*

To prove this, we will argue that a small variant on the grim strategies outlined above can support the same limit equilibrium payoffs as with synchronous signals. The modification is necessary because if player 2 observes the state at t, and sees that an event occurred at $t - \varepsilon$, where ε is small compared to $1/\gamma$, then 2 knows it is unlikely that 1 has yet observed an event. Player 2 is indifferent between In and Out when there is probability 1/2 that 1 has already seen the event and so is playing -1; this occurs when $1 - \exp(-\gamma\varepsilon) = .5$ or $\varepsilon = \ln(2)/\gamma$; for all smaller ε, 2 prefers playing In.[4]

As in Section 3.1, suppose that the public randomizing device is uniformly distributed on [0,1], 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. We consider strategies of the following form: player 1 starts out playing +1 and sticks with it until the first time he sees a punishment event or he sees player 2

---

[4] 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, and these beliefs will depend on the time that player 2 last observed the process.

play Out; once player 1 starts playing -1 he 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 prefer his 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 his 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 player 2 does, then player 1's flow payoff is $\bar{u} + g$ until t', and $\underline{u}$ thereafter, where t' is the first time after $t + \ln(2)/\gamma$ that player 2 observes the event. If player 2 is the first to observe the event, 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 $\gamma \to \infty$, and 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. Intuitively, 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 IC constraint for the grim strategy is the same as before up to $O(1/\gamma)$: It is sufficient that $\frac{g}{\bar{u}-\underline{u}} < \frac{\alpha(\lambda_q - \lambda_p)}{\lambda_p + r}$.

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 with synchronous signals 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 payoffs and either the payoffs of Table 1 or Table 2 is at least as high as in FL. We conjecture that no higher payoffs are possible, but we have not attempted to characterize all of the equilibria in this game with private monitoring - indeed, we are unaware of any complete characterization of the limit equilibria in any game of private monitoring in a case where the folk theorem fails.

4. Diffusion Signals

4.1. Public Monitoring of a Diffusion with different variances

Consider the case of synchronous signals and actions, in which players revise their actions at the instant they receive a new signal, i.e., at dates t, $2t$, etc.

As in FL, we assume that over any interval of length t, the signals have variance $\sigma_{+1}^2 t$ and $\sigma_{-1}^2 t$ as the action chosen is $a_1 = +1$ or -1. (The means in Fudenberg and Levine are $\mu_{+1}t$ and $\mu_{-1}t$, but we assume for simplicity that $\mu_{+1} = \mu_{-1} = 0$.) We will focus on the "good news case" where $1 < \sigma_{+1}^2/\sigma_{-1}^2$. Here FL show (see Propositions 5 and 6) that there exist $\overline{\lambda}, \underline{\lambda} > 1$ such that if $1 < \sigma_{+1}^2/\sigma_{-1}^2 < \underline{\lambda}$, there is no non-trivial equilibrium; and if $\sigma_{+1}^2/\sigma_{-1}^2 > \overline{\lambda}$, there is a non-trivial equilibrium. However for any fixed $\sigma_{+1}^2/\sigma_{-1}^2$, all equilibrium payoffs are bounded away from efficiency. Fudenberg and Levine study only the game from their Table 1, but their proof works also for the game from Table 2. Of course, taking the lower of the two values of $\underline{\lambda}$ and the higher of the two values of $\overline{\lambda}$, we can assume that there exist $\overline{\lambda}$ and $\underline{\lambda}$ common for both games.

4.2. Noisy Observation of Diffusion signals

Now assume that LR's signal is noisy. More precisely, assume that SR observes perfectly the state of the diffusion process at dates $t, 2t, ...$, while LR observes this state perturbed by an i.i.d. random variable $\eta \sim N(0, \varepsilon^2\sigma^2)$. Later, i.e., at date $(k+1)t$, LR already observes perfectly the state of the diffusion process at date kt. This is not the same as assuming that the players observe the process at slightly different times, but it does seem like a very small perturbation of the model, so it is interesting to note that even this change is enough to lead to different limit results than in the case of deterministic signals. After explaining why this is the case, we will explain why we expect similar results to hold with truly asynchronous signals.

We shall first keep t fixed, and show that every payoff that can be attained in equilibrium for $\varepsilon = 0$ can be approximated by equilibrium in the limit of positive ε tending to zero for the game from Table 1 of FL. However, as we argue later, the result does not generalize to other (similar) two-person, two-action stage games, and in particular it does not generalize to the game from Table 2.

Recall that the equilibrium constructed in Fudenberg and Levine has the following form: Start out playing +1, In; if SR has played Out in the previous period, or the observed increment of the stochastic process lies beyond the interval $[-z^*(t), z^*(t)]$,[5]

---

[5] Recall that we assume that $\mu_{+1} = \mu_{-1} = 0$. This is the reason for the positive and negative thresholds being equal.

play -1 or Out. Otherwise, play $+1$, In. Fudenberg and Levine also show that the optimum LR's payoff can be achieved by strategies of this form.

The optimal cutoff $z^*(t)$ is determined by the condition that LR is indifferent between playing $+1$ and -1. Let

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

let also

$$\delta = \exp(-rt).$$

The probability of the signal taking a value $z < -z^*(t)$ or $z > z^*(t)$ under action $+1$ and -1, respectively, is

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

And LR is indifferent if

$$(1-\delta)g = \delta(q^*(t) - p^*(t))(v^*(t) - \underline{u}), \tag{4.1}$$

where

$$v^*(t) = (1-\delta)\overline{u} + \delta\left[p^*(t)\underline{u} + (1 - p^*(t))v^*(t)\right],$$

i.e.,

$$v^*(t) = \frac{(1-\delta)\overline{u} + \delta p^*(t)\underline{u}}{1 - \delta(1 - p^*(t))}.$$

The left-hand side of (4.1) is the difference between LR's payoff to playing -1 and $+1$ received between time kt and time $(k+1)t$ conditional on SR playing In; the right-hand side represents the expected difference between the continuation payoffs from time $(k+1)t$ on, contingent on playing ($+1$,In) and (-1,In) respectively, between time kt and time $(k+1)t$. All payoffs are computed at time kt.

If it happens that the left-hand side of (4.1) exceeds the right-hand side for any value of $z^*(t)$, which is for example the case when t is sufficiently large (or equivalently, δ is sufficiently small), then no non-trivial equilibrium exists. This "corner" case is, however, straightforward, since trivial equilibria exist for any value of ε. So, we disregard the corner case in what follows.

Suppose now that $\varepsilon > 0$. We claim that if ε is positive but sufficiently small, then similar strategies are still an equilibrium. We only need to modify the equations for cutoffs. Let now LR play -1 once she observes a signal from the complement of $[-z_1^*(t), z_1^*(t)]$, and SR play Out once she observes a signal from the complement of $[-z_2^*(t), z_2^*(t)]$. In the case of LR, the signal means the difference between the state observed at time kt (which is a noisy signal of the actual state) and the actual state at time $(k-1)t$ (which LR already knows at time kt). (Of course, players play -1, Out also when SR was observed to play Out in the past.) Let $v^*(t)$ denote LR's expected continuation payoffs from time kt on, contingent on SR playing In at date kt, and all previous dates, and LR playing $+1$ at date kt. Notice that $v^*(t)$ does not depend on k, or LR's signals observed at dates $t, 2t, , ..., kt$. Let, finally,

$$p^*(t) = 1 - 2\Phi(-\zeta_2^*(t)) \text{ and } q^*(t) = 1 - 2\Phi(-\frac{\sigma_{+1}}{\sigma_{-1}}\zeta_2^*(t));$$

that is, $p^*(t)$ and $q^*(t)$ are determined by SR's cutoff.

We will describe how to modify condition (4.1) in order to guarantee that the prescribed strategies are an equilibrium. We will first show how to pick $z_2^*(t)$ (or equivalently, $\zeta_2^*(t)$) to make LR, independently of the signal she observed, indifferent between playing -1 and $+1$. Consider the difference between LR's instant payoff to playing -1 and $+1$ between time kt and time $(k+1)t$. This difference is equal to πg, where π denotes the probability that SR plays In at date kt. Of course, this probability π depends on LR's signal at time kt.

Consider in turn the difference between the continuation payoffs from time $(k+1)t$ on, conditional on playing $+1$ and -1 at time kt. The continuation payoff contingent on playing $+1$ is

$$(1-\pi)\underline{u} + \pi\left[\pi_{+1,In}v^*(t) + \pi_{-1,In}v(t) + \pi_{Out}\underline{u}\right],$$

where $\pi_{+1,In}$, $\pi_{-1,In}$, and π_{Out} stand for the probability that players will play (+1,In), (-1,In), and SR will play Out, respectively, at date $(k+1)t$; and $v(t)$ denotes the expected continuation payoffs from time $(k+1)t$ on, contingent on SR playing In at date $(k+1)t$, and all previous dates, and LR playing -1 at date $(k+1)t$. It is important to note that $\pi_{+1,In}$, $\pi_{-1,In}$, π_{Out}, and $v(t)$ are independent of LR's signal at time kt. Similarly, the continuation payoff contingent on playing -1 is

$$(1-\pi)\underline{u} + \pi\left[\pi'_{+1,In}v^*(t) + \pi'_{-1,In}v(t) + \pi'_{Out}\underline{u}\right];$$

since $\pi_{+1,In} \to 1 - p^*(t)$, $\pi'_{+1,In} \to 1 - q^*(t)$, when $\varepsilon \to 0$, and $\pi_{-1,In}$, $\pi'_{-1,In} \sim O(\varepsilon)$, the difference in continuation payoffs can be represented as

$$\pi\left[(q^*(t) - p^*(t))(v^*(t) - \underline{u}) + O(\varepsilon)\right].$$

The term $O(\varepsilon)$ is independent of LR's signal at time kt.

Therefore, the analogue of equation (4.1) has the form

$$(1-\delta)g = \delta(q^*(t) - p^*(t))(v^*(t) - \underline{u}) + O(\varepsilon). \tag{4.2}$$

Similarly, the formula for the payoff $v^*(t)$ changes only by $O(\varepsilon)$, and is independent of LR's signal at time kt.

Finally, SR's cutoff $z_2^*(t)$ must have the property that she is indifferent between playing In and Out, i.e., she must assign probability 1/2 to LR observing an increment from interval $[-z_1^*(t), z_1^*(t)]$. This yields an additional equation that determines $z_2^*(t)$ as a function of $z_1^*(t)$. If SR observes signal $z_2^*(t)$, she believes that LR's signal is $N(z_2^*(t), \varepsilon^2\sigma^2)$. Again, it is important that SR's beliefs depend only on the increment she observed at time kt.[6] So we see that the $z_2^*(t)$ that makes SR indifferent converges to $z_1^*(t)$ as $\varepsilon \to 0$.

Thus, if thresholds $z_1^*(t)$ and $z_2^*(t)$ are chosen in this manner, the FL trigger strategies are an equilibrium also for positive, but small, values of ε; moreover, if $\varepsilon \to 0$, the equations that determine LR's payoff converge to the equations that determine the highest equilibrium payoff she obtains when $\varepsilon = 0$.

Consider now the stage-game payoffs of Table 2. We restrict attention to *trigger strategies*, i.e., the strategies of the form: Start out playing +1, In; once SR is observed to play Out, or player $i = 1, 2$ observes an increment the stochastic process from the complement of $[-z_i^*(t), z_i^*(t)]$, she plays −1 or Out, respectively, from then on. Otherwise, they keep playing +1, In. We emphasize that in the case of LR, the increment means the difference between the state observed at time kt (which is a noisy signal of the actual state) and the actual state at time $(k-1)t$ (which LR already knows at time kt).

The equations that determine the optimum LR's equilibrium payoff in trigger strategies are the same as in the case of the game from Table 2, except equation (4.2). The difference between LR's instant payoff to playing −1 and +1 between time kt and time $(k+1)t$ becomes equal to g, because playing −1, compared to +1, yields an instant payoff higher by g, independently of SR's action.

The difference between the continuation payoffs from time $(k+1)t$ on, conditional on playing +1 and −1 at time kt, on the other hand, remains unaltered, except that $\underline{u}$ must be replaced with $\underline{u}+g$, because LR's stage-game equilibrium payoff has changed. Thus, equation (4.2) in the determination of LR's optimum equilibrium payoff must be replaced with

$$(1-\delta)g = \pi\delta(q^*(t) - p^*(t))(v^*(t) - \underline{u} - g) + O(\varepsilon). \tag{4.3}$$

Recall that π depends on LR's signal at time kt. We take π to be the probability that SR plays In at date kt, assigned by LR with signal $z_1^*(t)$; the probability is equal to the probability assigned by LR with signal $-z_1^*(t)$. If LR observes $z \in (-z_1^*(t), z_1^*(t))$, then π is larger, and LR has an incentive to play +1; and if LR

---

[6] Otherwise, we would have distinct equations for distinct types of SR player.

observes $z \notin (-z_1^*(t), z_1^*(t))$, then π is smaller, and LR has an incentive to play -1, exactly as she is prescribed.

Notice that the characterization of optimal equilibria can be used to obtain limit results for $t \to 0$, $\varepsilon/t \to 0$, similar to ones obtained in Fudenberg and Levine for $\varepsilon = 0$. [7]

Proposition 3. *Suppose* $1 < \sigma_{+1}^2/\sigma_{-1}^2$. *Consider the limit of LR's payoffs of trigger-strategy equilibria in the noisy case as* $t \to 0$, $\varepsilon/t \to 0$ *and the limit of equilibrium payoffs, as* $t \to 0$, *in the synchronous case studied by FL.*

(i) If

$$\frac{\sigma_{+1}}{\sigma_{-1}} < \underline{\lambda}, \tag{4.4}$$

then with either synchronous or asynchronous signals, there is no non-trivial trigger strategy equilibrium in the limit for both the game from Table 1 of FL and the game from Table 2.

(ii) If

$$\frac{\sigma_{+1}}{\sigma_{-1}} > \overline{\lambda}$$

then, for the payoffs of Table 1 of FL, there are non-trivial limit payoffs to trigger-strategy equilibria, and the limit of the sets of LR's equilibrium payoffs with asynchronous signals coincides with the limit of the set of equilibrium payoffs when signals are synchronous.

(iii) IF

$$\frac{\sigma_{+1}}{\sigma_{-1}} > \overline{\lambda}$$

then, for the payoffs of Table 2, there are non-trivial limit payoffs to trigger-strategy equilibria, and the limit of the highest LR's equilibrium payoff with asynchronous signals is strictly lower than the limit of the highest LR's equilibrium payoff when signals are synchronous.

Proof: Part (i) follows from Fudenberg and Levine and upper-hemi continuity of the set of equilibrium payoffs with respect to ε. Part (ii) follows from the fact that equation (4.2) converges to (4.1) as $\varepsilon \to 0$.

---

[7] We focus on the good-news case as here the equilibrium set is more sensitive to the exact nature of the information received. In the "bad news" case $\sigma_{+1}^2 < \sigma_{-1}^2$ for any interest rate r, and both the game from Table 1 of FL and the game from Table 2, it is easy to show that in the limit $t \to 0$, $\varepsilon/t \to 0$ there are equilibria whose payoffs converge to efficiency, just as with synchronous signals.. To see this, recall that FL Proposition 4 shows that any payoff strictly smaller than efficiency can be attained in the limit as $t \to 0$ by taking a $\zeta^*(t) \equiv \zeta^*$ independent of t. For this value of ζ^*, the left-hand sides of (4.1) and (4.3), respectively, must be strictly smaller than the right-hand sides for sufficiently small values of t.

Finally, consider the game from Table 2, and suppose to the contrary that the limits of the highest equilibrium payoffs coincide for synchronous and asynchronous signals. Denote this limit by v^*. It now follows from the equation for $v^*(t)$ that

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

In the case of good news,

$$p^*(t) = 1 - 2\Phi(-\zeta_2^*(t)) \text{ and } q^*(t) = 1 - 2\Phi(-\frac{\sigma_{+1}}{\sigma_{-1}}\zeta_2^*(t)),$$

where

$$\zeta_2^*(t) \to_{t\to 0} 0.$$

Thus,

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

Applying l'Hopital's rule

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

The limits $p^*(t)/(1-\delta)$ and $q^*(t)/(1-\delta)$ have been computed, and are the same, for both synchronous and asynchronous signals. Thus, by (4.1),

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

and by (4.3) and the assumption that $\varepsilon/t \to 0$,

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

(Notice that (4.1) and (4.3) must hold for small enough values of t, since $v^* > \underline{u} + g$ whenever $\sigma_{+1}/\sigma_{-1} > \overline{\lambda}$.) Therefore, we obtain a contradiction, since $v^* < \overline{u}$ for any $\sigma_{+1}/\sigma_{-1} > 1$.

The discontinuity of the set of LR's payoffs in the case of the game from Table 2 may seem surprising, when we compare the result to the existing papers on private but almost public monitoring, since when ε is small we might think that the monitoring structure is in some sense close to its limit, public, form.[8] However, the monitoring

---

[8] See for example Mailath and Morris (2002, 2006), Hörner and Olszewski (2008) and Mailath and Olszewski (2008). Indeed, Hörner and Olszewski (2008) prove the "limit" 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.

structure studied in this paper is not almost public, in the sense defined by Mailath and Morris (2002). Indeed, an important implication of their definitions is that if a player receives a signal s, she is almost certain that other players also received signal s, so that "everyone observed s" is common p-belief for a p close to 1.

When LR player observes a diffusion process with a non-zero ε amount of noise, each player believes that the opponent's signal in normally distributed with a positive, although small, variance. Most importantly, if players play cutoff strategies, and a player observes the threshold (or close to the threshold) increment, she believes that the opponent will play each of the two actions with probability close to 1/2; this is quite different than the case of $\varepsilon = 0$, in which the player knows which action will be played by the opponent.

The previous work on repeated games with almost public signals assumed that the support of the signal distribution is finite, which raises the question of whether the set of LR's equilibrium payoffs would be continuous in ε if signals had a finite range of values. So suppose that the diffusion process we have analyzed so far is replaced with a stochastic process that takes a finite number of values. Suppose in addition that the noise in the LR's signal has property 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, 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 $\varepsilon = 0$.[9]

To see how this conjecture relates to observing diffusion processes at high frequency, recall that Fudenberg and Levine (2008), studying synchronous signals, approximate the diffusion processes for each action by arrays that converge to the diffusions as the time period goes to 0. To do this, they 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 to asynchronous signals, by assuming that LR observes the total number of events at times $(k+1)t$ with a little noise, which is multinomial on $\{-c/m, ..., -1/m, 0, 1/m, ..., c/m\}$. 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. Otherwise, we would expect similar issues

---

[9] We are more confident about this conjecture if (as FL assumed) players have access to a public randomizing device. Note that the conjecture will show that a form of tigger strategy is robust to almost public montoring; yet, Mailath and Morris (2002) found that trigger strategies are not equilibrium strategies under almost public monitoring in repeated games with all long-run players and finite sets of actions and signals. The key difference between the two models, which enabled us to use trigger strategies, is that LR observes SR's actions, and with a one-period delay, she also observes SR's signals.

to arise.

We would also expect similar issues to arise in the model where state is observed without noise but the observation times are slightly different. Intuitively, if players' strategies involve cutoffs, then as with noisy observations a player whose signal is exactly at the cutoff assigns probability .5 to the other player's signal being on either side of it. However, when the observation times differ, the construction of any equilibrium strategies encounters difficulties typical for repeated games with imperfect private monitoring: complex statistical inference and the lack of recursive structure (see Kandori (2002). These difficulties seems to be magnified by the fact that the support of the signal distribution is a continuum.

To see the problem suppose that players observe, as in Section 3.1, the state of diffusion at time $t - \varepsilon_{i,1}, 2t - \varepsilon_{i,2}, , ..., kt - \varepsilon_{i,k}, ...$, where $(\varepsilon_{i,k})_{k=1}^{\infty}$ is a sequence of i.i.d. random variables distributed uniformly on an interval $[0, \varepsilon]$ with $\varepsilon \geq 0$, and they take actions at time t, 2t, etc. Further, suppose that players' choice of actions at date kt involves cutoffs in the increments observed at time kt. Then LR's cutoff must depend not only on the increment observed at time $kt - \varepsilon_{1,k}$, but also the increment observed at $(k-1)t - \varepsilon_{1,k-1}$. Indeed, since it happens with a positive probability that $(k-1)t - \varepsilon_{2,k} < (k-1)t - \varepsilon_{1,k}$, the increment observed at $(k-1)t - \varepsilon_{1,k-1}$ contains information about the likelihood of SR's signal at $kt - \varepsilon_{2,k+1}$ being on either side of her cutoff.[10] This yields a continuum of LR's cutoffs. Further, the increment observed by SR at time $(k-1)t - \varepsilon_{2,k}$ contains information about LR's signal at $(k-1)t - \varepsilon_{1,k}$, and therefore about her cutoff. This yields a continuum of SR's cutoffs at kt, and implies that the increment observed by LR at time $(k-2)t - \varepsilon_{2,k}$ contains information about SR's cutoff at kt. Thus, the cutoffs must depend on the entire history up to date kt, and the relation between them seems to be very complex.

---

[10] Similarly, LR's cutoff must depend on LR's action at time $(k-1)t$.

REFERENCES

Abreu, D. , P. Milgrom and D. Pearce [1991] “Information and Timing in Repeated Partnerships.” *Econometrica*, 59, 1713–33.

Abreu, D., D. Pearce, and E. Stachetti [1990] “Towards a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica*, 58: 1041-1064.

Dutta, P. [1995] “A Folk Theorem for Stochastic Games,” *Journal of Economic Theory,* 66, 1-32.

Faingold, E. and Y. Sannikov [2007] “Equilibrium Degeneracy and Reputation Effects,” mimeo.

Fudenberg, D. and D. K. Levine [2007] “Continuous Time Limits of Repeated Games with Imperfect Public Monitoring” *Review of Economic Dynamics*, 10, 173-192

Fudenberg, D., D M. Kreps, and E. Maskin [1990] “Repeated Games with Long-run and Short-run Players” *Review of Economic Studies*, 57, 555-573.

Fudenberg, D. and D. K. Levine [1994]: “Efficiency and Observability with Long-Run and Short-Run Players,” *Journal of Economic Theory*, 62, 103-135.

Fudenberg, D. and D. K. Levine [2008] “Repeated Games with Frequent Actions,” forthcoming in the *Quarterly Journal of Economics.*

Fudenberg, D., D.K. Levine, and E. Maskin [1994]: “The Folk Theorem with imperfect public information,” *Econometrica*, 62, 997-1040.

Fudenberg, D. and E. Maskin [1986] “The Folk Theorem in Repeated Games with Discounting or with Incomplete Information,” *Econometrica,* 54, 533-556.

Horner, J. and W. Olszewski [2006] “The Folk Theorem for Games with Private Almost-Perfect Monitoring,” *Econometrica,* 74, 1499-1545.

Hörner, J. and W. Olszewski [2008] “How Robust is the Folk Theorem with Public Monitoring,” mimeo.

Kandori M. [2002] “Introduction to Repeated Games with Private Monitoring,” *Journal of Economic Theory*, 102: 1-15.

Lagunoff, R. and A. Matsui [1997] "Asynchronous Choice in Repeated Coordination Games," *Econometrica,* 65, 1467-1477.

Mailath, G. and S. Morris [2002] “Repeated Games with Almost-Public Monitoring,” *Journal of Economic Theory,* 102, 189-228.

Mailath, G. and W. Olszewski [2008] “Folk Theorems with Bounded Recall under (Almost) Perfect Monitoring,” mimeo.

Monderer, D. and D. Samet [1989] “Approximating Common Knowledge with Common Beliefs,” *Games and Economic Behavior,* 1, 170-190.