Continuous Time Models of Repeated Games with Imperfect Public Monitoring

Drew Fudenberg and David K. Levine

This version: 9/1/2005

First version: 12/15/2005

1. Introduction

In this paper we consider a number of different ways that a sequence of discrete-time repeated games can approach a continuous-time limit. Our purpose is to clarify the effects of three different factors: 1) The distribution of signals in a fixed discrete-time game, 2) How the distribution (and notably its variance) changes with the period length, and 3) whether some of the player roles are filled by sequences of short-run players. Our interest in these questions was sparked by recent papers by Sannikov [], Sannikov and Skrypcaz [], Faingold and Sannikov [2005] and Faingold [2005] on various repeated games in continuous time, and one of our goals is to simply to make some of their results more accessible and intuitive by providing discrete-time analogs. A second goal is to generalize their results, by considering sequences of discrete-time games that do not converge to the class of stochastic processes these papers considered. A third, and related, goal, is to illuminate the situations which are and are not well described by diffusion processes. Finally, as in Fudenberg, Levine, and Takahashi [2005], we would like to make explicit the parallel between the assumption that some of the players are short run and restrictions (either on payoff functions or on the set of equilibria) that reduce the dimensionality of the space of possible payoff vectors. Such restrictions have previously been shown to bound the set of equilibrium payoffs away from efficiency; we explain how this bounding effect can lead to a collapse of the equilibrium set to the static equilibria in some (but not all) continuous-time limits.

To set the stage for the issues we will address in this paper, a brief review will be useful. Under some identification conditions, Fudenberg, Levine and Maskin [1995] provide a folk theorem for the case of all long-run players, showing that any individually rational payoff vector can be approximated by an equilibrium payoff if the common discount factor of the players is sufficiently close to 1. More precisely, let $E(\delta)$ be the set of perfect public equilibrium payoffs for a fixed δ, and let $E(1) = \lim_{\delta \to 1} E(\delta)$; on the conditions of the FLM theorem, if payoff vector v is feasible and individually rational then it is in $E(1)$. FLM also explain the highest equilibrium payoff in symmetric strategies can be bounded away from efficiency when there are equilibrium payoffs that are (almost?) symmetric and almost efficient. Sannikov [2005] characterizes the equilibrium payoffs in a repeated game in continuous time, where players control the mean of a vector-valued diffusion process; under a somewhat stronger identification decision ("product structure") he also proves a folk theorem for the limit of interest rates $r \to 0$.

For the case of games with both long-run and short-run players, Fudenberg and Levine [1994] provide a linear programming algorithm for computing the limit of the equilibrium payoffs as the discount factor of the long-run players converges to 1, and used this to prove a characterization of the limit payoffs in games with a product structure. This limit set is typically smaller than if all players were long run, and in particular the highest equilibrium payoff of a long-run player is bounded away from what it would be if all players were long run.[1] However, the limit set typically does include payoff vectors that cannot be generated by static equilibria. For this reason it is striking to note that Faingold and Sannikov [2005] show that the set of equilibria in a repeated game with one LR player facing SR opponents in continuous time when the public information is a diffusion process is simply the static equilibria, irregardless of the interest rate, so that the folk theorem fails. Thus changing the standard model by assuming both short run players and a diffusion process makes a more significant qualitative difference than either change on its own; this is one of the findings we would like to explain with discrete-time methodology.

A second existing result that we would like to better understand is that the effect just described is specific to the diffusion process, and does not in general extend to the case of continuous time repeated games with Poisson signals. Abreu, Pearce and Milgrom [1991] (APM) investigate how the set of equilibrium payoffs varies with period length in a two-action partnership game with two long-run players, where what is observed in each time period is the number of Poisson-distributed "signals" that have arrived in the period. They restrict attention to symmetric equilibria, and determine the limit of the highest symmetric equilibrium payoff as the time period shrinks to 0; whether this limit is degenerate (that is, includes only the static equilibrium payoff) or not depends on the relationship between the parameters of the payoff matrix and the appropriate likelihood function. Like APM, we restrict our analysis of continuous-time limits to some simple

---

[1] The reason for this was first noted by Fudenberg, Kreps, and Maskin [], which coves the case of perfectly observed actions.

games, as opposed to providing a general classification. Still we believe that our analysis illustrates some more general propositions.

One of the points we would like to emphasize is the similarity between a repeated game between a long-run player and a sequence of short-run opponents, and a repeated game between two long-run players when a symmetry condition is imposed. In both cases, unlike with general equilibria of games with two long-run players, there is no way to "efficiently punish" one player by simultaneously rewarding his opponent. The key to the FLM folk theorem in games with two long-run players is that the identification conditions imply that there are equilibria where incentives can be provided at negligible efficiency cost by such efficient punishments; this is what FLM call "enforcement on tangent hyperplanes."[2] The identification conditions are purely qualitative; when they are satisfied, the set $E(1)$ is independent of the exact nature of the distribution of signals.[3] In contrast, with a symmetry restriction or short-run players, so that efficient punishments are impossible, the size of $E(1)$ depends on the probability that punishment is triggered. This probability is endogenously determined as part of the equilibrium, but to characterize the most efficient equilibrium what matters is how small the probability can be made without giving a player an incentive to deviate; in our examples, this depends on a particular likelihood ratio that we identify. In the Poisson case, this likelihood ratio is constant as the time period shrinks, when it is sufficiently large, the equilibrium set is non-degenerate in the continuous-time limit, just as in APM. In contrast, the key likelihood ratio converges to 0 in the diffusion case, which provides a discrete-time explanation of Faingold and Sannikov's equilibrium degeneracy result.

2. Long-Run versus Short-Run

We consider the two-person two-action stage game with payoff matrix

	Player 2

[2] Enforcement on tangent hyperplanes is also a key part of Sannikov's [2005] characterization of the equilibrium set in continuous time. In each case (both discount factors going to 1 and time periods shrinking to 0) the equilibrium continuation payoffs vary only slightly with each observation, and moving along a tangent hyperplane means that the efficiency loss is second order.

[3] The FLM and FL results both use a "full-dimension" condition; see Fudenberg, Levine and Takahashi [2005] for a characterization of $E(1)$ without this condition.

		L	R
Player 1	**+1**	$\underline{u}$,0	u,1
	-1	$\underline{u}$,0	$u+g$,-1

where $\underline{u} < u, g > 0$. In this game, player 2 plays **L** in every Nash equilibrium, so player 1's static Nash equilibrium payoff is $\underline{u}$, which is also the minmax payoff for player 1. Naturally player 1 would prefer that player 2 play **R**, but he can only induce player to play **R** by avoiding playing **–1**.

At the end of each stage game, a public signal $z \in \mathbb{R}$ is depends only on the action taken by player 1; player 2's action is publicly observed. The probability distribution of the public signal is $F(z \mid a_1)$. We assume that F is either differentiable and strictly increasing, or that it corresponds to a discrete random variable. In either case, let $f(z \mid a_1)$ denote the density function. We assume the monotone likelihood ratio condition that $f(z \mid a_1 = -1) / f(z \mid a_1 = +1)$ is strictly increasing in z. This says that z is "bad news" about player 1's behavior in the sense that large z means that player was probably playing **–1**, a reputation player 1 would like to avoid if he is to keep player 2 in the game. We assume also the availability of a public randomization device; the outcome of this device is observed at the start of each period, before actions are taken.

Let τ denote the length of the period. We suppose that player 1 is a long-run player with discount factor $\delta = 1 - r\tau$ facing an infinite series of short-run opponents. The most favorable perfect public equilibrium for LR is characterized[4] by the largest value v that satisfies the incentive constraints

$$v = (1-\delta)u + \delta \int w(z) f(z \mid a_1 = +1) dz$$
$$v \geq (1-\delta)(u+g) + \delta \int w(z) f(z \mid a_1 = -1) dz$$
$$v \geq w(z) \geq \underline{u}$$

or $v = 0$ if no solution exists. Notice that this formulation is possible only because the existence of a public randomizing device implies that any payoff $w(z)$ between v and $\underline{u}$ can be attained by randomizing between the two equilibria. Note that the second incentive constraint must hold with equality, since otherwise it would be possible to

[4] The set of PPE payoffs is compact so the best equilibrium payoff is well-defined.

increase the punishment payoff w while maintaining incentive compatibility, and by doing so increase utility on the equilibrium path. This is a simple extension to the case of a continuous signal of the result proven in Fudenberg and Levine [1994].

Because of the monotone likelihood ratio condition, equilibria that give the long-run player the maximum utility have a cut-point property, with fixed punishment occurring if the signal exceeds a threshold z^*. In the case of a continuous z this condition is quite straightforward; in the discrete case it is complicated by the fact that the threshold itself will typically be realized with positive probability. For this reason it is useful for a given threshold $z^* \in \Re$ to use public randomization to define a random variable $\tilde{z}^*$ that in the continuous case is equal to z^* and in the discrete case picks the two grid points $\underline{z}^* < \overline{z}^*$ just below and above z^*, with probability $(\overline{z}^* - z^*)/(\overline{z}^* - \underline{z}^*)$ of picking $\underline{z}^*$. After the signal z is observed, and before play in the next period, the public randomizing device is used to determine whether z is compared to $\underline{z}^*$ or $\overline{z}^*$ In this case

Proposition 1: *There is a solution to the LP problem characterizing the most favorable perfect public equilibrium for the long-run player with the continuation payoffs* $w(z)$ *given by*

$$w(z) = \begin{cases} w & z \geq \tilde{z}^* \\ v & z < \tilde{z}^* \end{cases}$$

and indeed, $w = \underline{u}$.

Proof: Let $w(z)$ be a solution to the LP problem, and let

$$W = \int w(z) f(z \mid a_1 = -1) dz$$

Clearly $w(z)$ must also solve the problem of maximizing

$$\int w(z) f(z \mid a_1 = +1) dz$$

subject to

$$\int w(z) f(z \mid a_1 = -1) dz \leq W$$
$$v \geq w(z) \geq \underline{u}$$

Ignoring for a moment the second set of constraints, and letting ν be the Lagrange multiplier on the first constraint, the derivative of the Lagrangean is

$$\int w(z)[f(z \mid a_1 = +1) - \nu f(z \mid a_1 = -1)]dz\,.$$

By the monotone likelihood ratio condition, there is a z^* such that

$$\frac{f(z^* \mid a_1 = +1)}{f(z^* \mid a_1 = -1)} \stackrel{>}{<} \nu$$

as $z < z^*, z > z^*$, and in the continuous case there is a unique z^* for which the condition holds with equality.

This now shows that for $z < z^*$ we must have $w(z) = v$ and for $z > z^*$ we must have $w(z) = \underline{u}$. That leaves when z is discrete the case $z = z^*$. Since in that case $\underline{u} \le w(z^*) \le v$ can be realized by a public randomization between $\underline{u}, v$, we may use the $\tilde{z}^*$ construction for some appropriately chosen z^*.

☑

In the continuous case, we can now define

$$p = \int_{z^*}^{\infty} f(z \mid a_1 = +1)dz, q = \int_{z^*}^{\infty} f(z \mid a_1 = -1)dz$$

to be the probability of punishment conditional on each of the two actions. In the discrete case, we can make a similar definition, taking account of the public randomization implicit in $\tilde{z}^*$. Consider, then, the LP problem of maximizing v subject to the incentive constraints

$$v = (1-\delta)u + \delta(v - p(v-w))$$
$$v = (1-\delta)(u+g) + \delta(v - q(v-w))$$
$$v \ge w \ge \underline{u}$$

or $v = \underline{u}$ if no solution exists. Choosing the cutoff point z^* which leads to the largest solution of this optimization problem then characterizes the most favorable perfect public equilibrium for the long-run player; we know also that in this optimal solution $w = \underline{u}$.

In order to accommodate our subsequent analysis of games with two long-run players, it is useful to slightly generalize the constraints to

$$v - w \ge 0, v - w \le \beta(v - \underline{u})$$

which corresponds to

$$v \geq w \geq \underline{u}$$

when $\beta = 1$.

The linear programming problem characterizing the most favorable equilibrium is easily found to have a solution if and only if

(*) $$r \leq \frac{1}{\tau + \beta^{-1}(((u-\underline{u})/g)[(q-p)/\tau]-[p/\tau])^{-1}}$$

in which case the solution is given by

(**) $$v^* = u - gp/(q-p).$$

We are interested in the case in which τ is small; in particular we are interested in the case in which $q(\tau), p(\tau)$ are functions of τ. Call $\rho, \mu \in \Re \cup \{\infty\}$ *regular values* of $q(\tau), p(\tau)$ if along some sequence $\tau^n \to 0$ we have $\rho = \lim_{\tau^n \to 0}(q(\tau^n) - p(\tau^n))/p(\tau^n)$ and $\mu = \lim_{\tau^n \to 0}(q(\tau^n) - p(\tau^n))/\tau^n$. The first limit ρ can be thought of as the signal to noise ratio, since $q - p$ is a measure of how different the distribution of outcomes is under the two different actions, and p is a measure of how often the "punishment" signal arrives when in fact the long-run player has been well-behaved. The second limit μ is a measure of the signal arrival rate.

Notice that along a sequence $\tau^n \to 0$ defining regular values ρ, μ the limit $\overline{v} = \lim_{\tau^n \to 0} \max\{\underline{u}, v^*\}$ of most favorable equilibria exists, as does the limit of the right-hand-side of (*)

(***) $$\frac{1}{(\mu/\beta\rho)(((u-\underline{u})/g)\rho - 1)^{-1}}$$

If this expression is positive and $\overline{v} > \underline{u}$ we say that there is a *non-trivial* limit equilibrium:[5] in this case from (*) and (**), there exists positive τ, r such that for all smaller values[6] there is an equilibrium giving the long-run player more than $\underline{u}$. Conversely, if either $\overline{v} \leq \underline{u}$ or (***) is non-positive then for any fixed $r > 0$ along the

[5] Notice that we require (***) to be strictly positive. If we allowed (***) to be zero, then at best we could assure that there is some sequence of r^n's giving rise to a non-trivial equilibrium for τ^n. This allows the interest rate to go to zero faster than the period length, and corresponds to the usual type of folk-theorem experiment, where the discount factor goes to 1 for fixed period length. Here we hold fixed r as we let the period length get shorter.

[6] The values of τ must be restricted to lie in the sequence τ^n.

sequence τ^n the best equilibrium for the long-run player must converge to $\underline{u}$. In case $\overline{v} = u$ we say that the limit equilibrium is efficient.

Proposition 2: *Suppose that* ρ, μ *are regular. Then there is a non-trivial limit equilibrium if and only if* $\mu > g/(u - \underline{u})$ *and* $\rho > 0$. *There is an efficient limit equilibrium if and only if* $\mu > g/(u - \underline{u})$ *and* $\rho = \infty$.

Proof: From (**) $\overline{v} > \underline{u}$ if and only if $\mu > g/(u - \underline{u})$. Given that condition, (***) is positive if and only if $\rho > 0$. Moreover, from (**) $\overline{v} = u$ if and only if $\rho = \infty$.

☑

Finally, we remark that in checking $\mu > g/(u - \underline{u})$ and $\rho > 0$ we need not limit ourselves to optimal cutoff points: if the condition holds for some sequence of cutoff points, it certainly holds for the optimal sequence; while if it holds for no sequence it cannot hold for the optimal sequence.

The Poisson Case

Suppose that the public signal of the long-run player's action is generated by observing an underlying Poisson process in continuous time. The Poisson arrival rate is λ_p if the action taken by LR is **+1** and λ_q if the action taken by LR is **–1**. As in Abreau, Pearce and Milgrom's [????] analysis of a partnership game with two long-run players, a critical role is played by whether the Poisson signal is good news – meaning that the long-run player probably played the commitment action **+1**, or bad news, meaning he probably deviated to **–1**. If $\lambda_q > \lambda_p$ the signal is "bad news." In this case we take random variable z to be the non-negative discrete random variable representing the number of signals received during the previous interval of length τ. If $\lambda_q < \lambda_p$ the signal is "good news." In this case we take the random variable z to be the non-positive discrete random variable representing the negative of the number of signals received during the previous interval of length τ. In this way we preserve the convention the high z is bad news. As in Abreu, Pearce and Milgrom [????], we will show that we get a non-trivial limit equilibrium in the bad news case, but not in the good news case.

To begin we analyze the bad-news case $\lambda_q > \lambda_p$. The cutoff point is how many signals must be received before the punishment $v - w$ is triggered. If we take the cutoff to be two or more signals, then the probability of triggering punishment is of order τ^2

meaning that $\mu = 0$. So the only interesting cutoff is to punish whenever any signal is received. We do not need to consider the optimal cutoff; if suffices to consider the cutoff in which punishment always occurs when any signal is received. The probability this occurs is $p(\tau) = 1 - e^{-\lambda_p \tau}, q(\tau) = 1 - e^{-\lambda_q \tau}$, as the long-run player plays **–1** or **+1**. We may then directly compute $\rho = (\lambda_q - \lambda_p)/\lambda_p, \mu = \lambda_q - \lambda_p$. Hence the condition for a non-trivial limit equilibrium is $\lambda_q - \lambda_p > g/(u - \underline{u})$. Notice that because the probability of getting more than a single signal vanishes at the rate τ^2 the limit equilibrium payoff $v^* = u - g\lambda_p/(\lambda_q - \lambda_p)$ computed from the sub-optimal cutoff is the same as the computed from the optimal cutoff. The significant feature of this solution is that it is independent of the payoff $\underline{u}$.

Now we analyze the "good news" case $\lambda_q < \lambda_p$. Here the punishment is triggered by a small number of signals, rather than a large number. If there is to be any punishment at all, then punishment must certainly occur when no signals arrive. Suppose the probability of punishment when there is no signal is $\gamma(\tau)$. Then $p(\tau) = \gamma(\tau)e^{-\lambda_p \tau}, q(\tau) = \gamma(\tau)e^{-\lambda_q \tau}$. Regardless of $\gamma(\tau)$ this implies $\rho = 0$.

As we observed in the introduction, the fact that one player is short run means that providing incentives to the long-run player has a non-trivial efficiency cost. In the "good news" case, providing incentives requires frequent punishment, but if there are many independent and non-trivial chances of a non-trivial punishment in a small interval of real time, the long run player's present value must be so low that it is impossible to improve on the static equilibrium. In contrast, there can be non-trivial equilibrium even in the limit when the signal used for punishment has negligible probability (as in the bad-news case) or_ if there are several long run players so that punishments can take the form of efficient transfers.

The Diffusion Case [Faingold and Sanikov [2005]

Faingold and Sannikov study the case where the signals about the action of the long-run player are generated by a diffusion process in continuous time, with the drift in the process controlled by the long-run player's action. If we supposed that players sample this process at intervals of length τ, the signals would have variance $\sigma^2\tau$. We will consider a slight generalization of the diffusion assumption: we allow the variance of the signal z to be given by $\sigma^2\tau^{2\alpha}$ where $\alpha < 1$, so that the diffusion case corresponds to $\alpha = 1/2$. The mean of the process is $-a_1\tau$ (recall that a_1 = **+1** or **–1**).

Observe that if Φ is the standard normal cumulative distribution, then

$$p = \Phi\left(\frac{-z*-\tau}{\sigma\tau^{\alpha}}\right)$$
$$q = \Phi\left(\frac{-z*+\tau}{\sigma\tau^{\alpha}}\right)$$

We fix x and consider the behavior of the best equilibrium for player 1 as $\tau \to 0$. We will show that for τ sufficiently small, (*) is necessarily violated.

Proposition 3*: For any $\alpha < 1$ there exists $\underline{\tau} > 0$ such that for $0 < \tau < \underline{\tau}$ there is no non-trivial limit equilibrium.*

Proof: Let $z*(\tau)$ denote the cutoff when the period length is τ. It is convenient to work with the normalized cutoff

$$\zeta(\tau) = \frac{\tau + z*(\tau)}{\sigma\tau^{\alpha}}$$

Then

$$p = \Phi(-\zeta(\tau))$$
$$q = \Phi\left(\frac{2\tau^{1-\alpha}}{\sigma} - \zeta(\tau)\right).$$

From Proposition 1, if there is to be a non-trivial limit equilibrium, we must have $\mu > 0$. We will show that this implies that $p(\tau^n)/\tau^n \to \infty$, forcing $\rho = 0$, leading to the conclusion that there can be no non-trivial limit equilibrium.

First we compute $(q-p)/\tau$ using the mean value theorem to observe that for each τ there is a number $f(\tau)$, $0 \le f(\tau) \le 1$, such that

$$\frac{q-p}{\tau} = \frac{\Phi\left(\frac{2\tau^{1-\alpha}}{\sigma} - \zeta(\tau)\right) - \Phi(-\zeta(\tau))}{\tau}$$
$$= \left(\frac{1}{\sigma\tau^{\alpha}}\right)\phi\left(-\zeta(\tau) + f(\tau)\frac{2\tau^{1-\alpha}}{\sigma}\right)$$

Let $c(\tau) = (q-p)/\tau$, we can invert this relationship to find

$$\phi\left(-\zeta(\tau)+f(\tau)\frac{2\tau^{1-\alpha}}{\sigma}\right)=c(\tau)\sigma\tau^{\alpha}$$

$$\frac{1}{\sqrt{2\pi}}\exp(-\frac{1}{2}\left(-\zeta(\tau)+f(\tau)\frac{2\tau^{1-\alpha}}{\sigma}\right)^2)=c(\tau)\sigma\tau^{\alpha}$$

$$\left(-\zeta(\tau)+f(\tau)\frac{2\tau^{1-\alpha}}{\sigma}\right)^2=-2\log(\sqrt{2\pi}c(\tau)\sigma)-2\alpha\log(\tau)$$

$$\zeta(\tau)=-\sqrt{-2\log(\sqrt{2\pi}c(\tau)\sigma)-2\alpha\log(\tau)}+f(\tau)\frac{2\tau^{1-\alpha}}{\sigma}$$

We now want to use this to show that $p/\tau\to\infty$. Since we have assumed that $c(\tau)$ is bounded away from zero, $\zeta(\tau)\le\sqrt{-\log b-2\alpha\log\tau}+a\tau^{1-\alpha}$. This gives $p/\tau\ge\Phi(-\sqrt{-\log b-2\alpha\log\tau}-a\tau^{1-\alpha})/\tau$. Again using the mean value theorem we know that for some $\tau'\in[0,\tau]$, we have

$$\begin{aligned}
&\Phi(-\sqrt{-\log b-2\alpha\log\tau}-a\tau^{1-\alpha})/\tau\\
&=\left(\frac{\alpha}{\tau'\sqrt{-\log b-2\alpha\log\tau'})}-a(1-\alpha)(\tau')^{-\alpha}\right)\phi(-\sqrt{-\log b-2\alpha\log\tau'}-a(\tau')^{1-\alpha})\\
&=\frac{1}{\sqrt{2\pi}}\left(\frac{\alpha}{\tau'\sqrt{-\log b-2\alpha\log\tau'})}-a(1-\alpha)(\tau')^{-\alpha}\right)\\
&\times\exp(-1/2)\left(-\log b-2\alpha\log\tau'+2a(\tau')^{1-\alpha}\sqrt{-\log b-2\alpha\log\tau'}+a^2(\tau')^{2-2\alpha}\right)\\
&=\frac{b^{1/2}}{\sqrt{2\pi}}\left(\frac{\alpha}{\tau'\sqrt{-\log b-2\alpha\log\tau'})}-a(1-\alpha)(\tau')^{-\alpha}\right)\\
&\times(\tau')^{\alpha}\exp(-a(\tau')^{1-\alpha}\sqrt{-\log b-2\alpha\log\tau'})\exp(-1/2)\left(a^2(\tau')^{2-2\alpha}\right)\\
&\to\frac{\alpha b^{1/2}}{\sqrt{2\pi}}\left(\frac{(\tau')^{\alpha-1}}{\sqrt{-2\alpha\log\tau'})}\right)\exp(-a\frac{\sqrt{-2\alpha\log\tau'}}{(\tau')^{\alpha-1}})
\end{aligned}$$

It suffices then to apply L'Hopital's rule to show that

$$\frac{\tau^{\alpha-1}}{\sqrt{-\log\tau}}\to(1-\alpha)\frac{\tau^{\alpha-2}}{1/\tau}=1-\alpha)\tau^{\alpha-1}\to\infty.$$

☑

Notice that in the limit for $\alpha<1$ the equilibrium collapses to the static Nash; in particular this is true even when $\alpha>1/2$, in which case the process converges to a deterministic one. By way of contrast, the "bad news" Poisson case, which like the diffusion case corresponds to $\alpha=1/2$, does not collapse in the limit. This shows that the exact form of the noise matters: is it a series of unlikely negative events, as in the "bad news" Poisson case, or a sum of small increments as in the normal case?

It is useful to contrast the diffusion case $\alpha = 1/2$ with a sum of small increments where the scale of the increment is proportional to the length of the interval, so that the standard error of the signal is of order τ. This corresponds to the case $\alpha = 1$. When we take the limit of such a sequence of processes, the limit is a deterministic process without noise.

Proposition 4: *If* $\alpha = 1$ *there exists* $\underline{\tau}$ *such that for all* $0 < \tau < \underline{\tau}$ *(*) is satisfied, and* $\lim_{\tau \to 0} v^* = u$.

Proof: Later.

☑

To understand this result, recall that $v^* = u - gp/(q-p)$, and consider for fixed τ taking a very large cutoff $z^* \to \infty$. As observed by Mirlees [????], this causes the likelihood ratio $q/p \to \infty$, so that $p/(q-p) \to 0$. Consequently, $v^* \to 1$, the pure commitment value, and the best PPE payoff for the most the long-run player can get when there is no noise in the signal. Notice, however, that this solution has $p, q \to 0$, which means that for fixed $\underline{u}$ and τ and z^* sufficiently large, (*) must be violated. On the other hand, for any choice of z^*, r, τ, there is always a $\underline{u}$ sufficiently small that (*) holds. In other words the worst punishment determines the best equilibrium: by going far enough into the tail of the normal, arbitrarily reliable information can be found about whether a deviation occurred, and this information can be used to create incentives, provided a sufficiently harsh punishment is available. The key point here is that when $\alpha = 1$, as the signal to noise ratio improves sufficiently quickly that we can exploit the shorter intervals to choose a more negative cutoff value of ζ.

Limits of Poisson Processes

Whether there is a non-trivial limit equilibrium as we approach continuous time depends on the nature of the underlying process: is it a Poisson "good" or "bad" news process? Is it a diffusion process? In many economic settings, it is natural to think of the signal that is observed is some kind of aggregate of underlying Poisson events. Consider for examples "sales" or "revenues." Each of these is made up of the sum of a number of individual transactions, and if we consider small enough time intervals we will observe at most a single transaction. Even for a monetary aggregate that measures all transactions in an economy, in any given nanosecond we are unlikely to observe more than a single trade take place. Thus Poisson processes seem natural for studying economic signals.[7] However, in many circumstances, individual transactions are not visible, so the first model we considered is the relevant one. That is, we suppose that the public signal of the long-run player's action is generated by observing an underlying Poisson process in continuous time. The Poisson arrival rate is λ_p if the action taken by LR is **+1** and λ_q if the action taken by LR is **–1**. What is observed is the number of Poisson events (or in the "good news" case its negative).

If we fix the Poisson process generating "sale" we conclude as we did above that if the time interval is sufficiently small, we get a non-degenerate limit if and only if the events correspond to "bad news." On the other hand, fixing the Poisson process and allowing the time interval to grow small, only individual transactions are ever observed, and as we indicated, this is not plausible in many settings. So it is important to consider allowing the Poisson parameters to vary with the time interval; that is consider general functions $\lambda_p(\tau), \lambda_q(\tau)$.

The expected number of events over an interval of length τ is then $\lambda_p(\tau)\tau, \lambda_q(\tau)\tau$. Besides the case where these are fixed constants, the most interesting case is the opposite case in which the expected number of events over a period is very large – that is, the limit $\lambda_q(\tau)\tau, \lambda_p(\tau)\tau \to \infty$. In this case, from the Central Limit Theorem, the Poisson random variable y has $(y - Ey)/(\operatorname{var} y)^{1/2}$ approach a standard normal. Moreover, when the action is **+1** we have $Ey = \lambda_p(\tau)\tau, \operatorname{var} y = \lambda_p(\tau)\tau$ and when it is **-1** we have $Ey = \lambda_q(\tau)\tau, \operatorname{var} y = \lambda_q(\tau)\tau$. In the bad news case $\lambda_q > \lambda_p$ define $z = (y - \lambda_p(\tau)\tau)/(\lambda_p(\tau)\tau)^{1/2}$, and in the good news case define $z = -(y - \lambda_p(\tau)\tau)/(\lambda_p(\tau)\tau)^{1/2}$. Then $\mu_q =\mid \lambda_q(\tau) - \lambda_p(\tau) \mid \tau^{1/2} / (\lambda_p(\tau))^{1/2}$. So the diffusion case is going to correspond to λ_p constant and $\lambda_q - \lambda_p = \tau^{1/2}$, which should be the Hellwig [????] case as well. $\tau^{1/2}$.

[7] We may want to consider the broader class of processes where jumps arrive according to a Poisson process, but the size of the jump itself is random – corresponding, for example, to the size of a transaction or a price at which it takes place. Since this adds little to the understanding of the limit, we stick with jumps of fixed size.

Poisson and Diffusion

Show that if we have two signals, one Poisson and one a diffusion in the limit of short intervals, only the Poisson matters.

Two Long-Run Players

We consider a discrete variation of a Bertrand competition game. This example is designed to be satisfy the requirements of the Folk Theorem, to be simple to analyze, and to have a linear Pareto frontier – the significance of which will become clear below. In this game, each firm can cooperate **C** or undercut **U**. If both cooperate, price in the market is the monopoly price and they split the market; if both undercut, price in the market is the competitive price and they split the market; if one cooperates and one undercuts, price in the market is the monopoly price and the undercutter gets $(1-\alpha)$ of the market, where $\alpha < 1/2$. In the market, sales opportunities arrive according to a Poisson process, with arrival rate λ_H if price is the monopoly price and λ_L if price is the competitive price. Firm strategies and prices are not observed, but if a sale is made, it is observed which firm made the sale. We will later get the diffusion process by making the arrival rate of sales very fast relative to the length of the interval, so that each firm makes many sales in a period.

Firms that sell nothing get nothing. A sale at the competitive price is worth L. A sale at the monopoly price is worth H. So the normal form of this game is

		Player 2	
		C	**U**
Player 1	**C**	$\lambda_H H/2, \lambda_H H/2$	$\alpha\lambda_H H, (1-\alpha)\lambda_H H$
	U	$(1-\alpha)\lambda_H H, \alpha\lambda_H H$	$\lambda_L L/2, \lambda_L L/2$

We assume that $\lambda_L L/2 > \alpha\lambda_H H$ so that it is better to get half the market at a low price than an α of the market at a high price – if this is not the case, then it is an equilibrium for one firm to undercut, and the game becomes a game of chicken.

There are three possible values of the signal: no sale; player 1 makes a sale; player 2 makes a sale. The instantaneous arrival rates for the latter two events is given by

	Player 1 sale (A)	Player 2 sale (B)
CC	$\lambda_H/2$	$\lambda_H/2$
CU	$\alpha\lambda_H$	$(1-\alpha)\lambda_H$
UC	$(1-\alpha)\lambda_H$	$\alpha\lambda_H$
UU	$\lambda_L/2$	$\lambda_L/2$

We assume that $(1-\alpha)\lambda_H < \lambda_L/2$ as would be the case with linear demand if *H* and *L* correspond to the monopoly and competitive price respectively.

Analysis of the discrete time game

What is the set of exactly efficient equilibria?

Suppose that $v = v_1, v_2$ is an equilibrium that maximizes firm 1's payoff on the efficiency frontier. Notice that by symmetry v_2, v_1 maximizes player 2's payoff on the efficiency frontier. If $v_1, v_2 = \lambda_H H/2, \lambda_H H/2$ this would be the only equilibrium on the efficiency frontier. By standard repeated game arguments, this can be shown to be impossible. So suppose $v_1 > \lambda_H H/2$. This implies that (U,C) must be played in the first period.

Recall that A corresponds to player 1 making a sale, B to player 2 making a sale and let 0 indicate that neither has made a sale. Player 1's utility at this equilibrium is

$$v_1 = (1-\delta)(1-\alpha)\lambda_H H + \delta((1-\alpha)\lambda_H w_1(A) + \alpha\lambda_H w_1(B) + (1-\lambda_H)w_1(0))$$

If he deviates to C he gets instead

$$(1-\delta)\lambda_H H/2 + \delta(\lambda_H w_1(A)/2 + \lambda_H w_1(B)/2 + (1-\lambda_H)w_1(0)) \leq v_1 .$$

Since $w_1(A), w_1(B), w_1(0) \leq v_1$ and we are assuming $v_1 > \lambda_H H/2$, this constraint always holds strictly, and so can be ignored in light of the other constraints.

For player 2 we have

$$v_2 = (1-\delta)\alpha\lambda_H H + \delta[((1-\alpha)\lambda_H w_2(A) + \alpha\lambda_H w_2(B) + (1-\lambda_H)w_2(0)$$

and for a deviation to U

$$(1-\delta)\lambda_L L/2 + \delta[w_2(A)\lambda_L/2 + w_2(B)\lambda_L/2 + (1-\lambda_L)w_2(0)] \leq v_2$$

We also know that payoffs are constrained to lie on the efficiency frontier, so that $w_2 = \lambda_H H - w_1$. When we substitute this into the equilibrium utility for player 2, we get the same equation as that for player 1, so this equation is redundant. Hence equilibrium is characterized by an equality and an inequality

$$v_1 = (1-\delta)(1-\alpha)\lambda_H H + \delta((1-\alpha)\lambda_H w_1(A) + \alpha\lambda_H w_1(B) + (1-\lambda_H)w_1(0))$$
$$v_1 \le (1-\delta)(\lambda_H H - \lambda_L L/2) + \delta[w_1(A)\lambda_L/2 + w_1(B)\lambda_L/2 + (1-\lambda_L)w_1(0)]$$
$$v_1 \ge w_1(A), w_1(A), w_1(B)$$

We may rewrite this substituting for v_1 in the incentive constraint

$$v_1 = (1-\delta)(1-\alpha)\lambda_H H + \delta((1-\alpha)\lambda_H w_1(A) + \alpha\lambda_H w_1(B) + (1-\lambda_H)w_1(0))$$
$$0 \le (1-\delta)(\lambda_H H - \lambda_L L/2 - (1-\alpha)\lambda_H H) +$$
$$\delta[(\lambda_L/2 - (1-\alpha)\lambda_H)w_1(A) + (\lambda_L/2 - \alpha\lambda_H)w_1(B) + (\lambda_H - \lambda_L)w_1(0)]$$
$$v_1 \ge w_1(A), w_1(A), w_1(B)$$

Observe that the coefficients of $w_1(A)$ (by the assumption that $(1-\alpha)\lambda_H < \lambda_L/2$) and $w_1(B)$ are positive in the incentive constraint. Increasing either improves the constraint and increases v_1, so solving the LP problem requires $w_1(A) = w_1(B) = v_1$. In addition, since $w_1(0)$ has a negative coefficient in the constraint, the constraint must hold with equality. In other words

$$0 = (1-\delta)(\alpha\lambda_H H - \lambda_L L/2) + \delta[(\lambda_L - \lambda_H)(v_1(A) - w_1(0))], \text{ so}$$

$$w_1(0) = \frac{(1-\delta)}{\delta(\lambda_L - \lambda_H)}(\alpha\lambda_H H - \lambda_L L/2) + v_1$$

We can then substitute this into the expression for v_1

$$v_1 = (1-\delta)(1-\alpha)\lambda_H H + \delta(\lambda_H v_1 + (1-\lambda_H)w_1(0))$$

to get

$$v_1 = (1-\delta)(1-\alpha)\lambda_H H + \delta(\lambda_H v_1 + (1-\lambda_H)\left(\frac{(1-\delta)}{\delta(\lambda_L - \lambda_H)}(\alpha\lambda_H H - \lambda_L L/2) + v_1\right)$$
$$v_1 = \left[(1-\alpha)\lambda_H H - \frac{(1-\lambda_H)}{(\lambda_L - \lambda_H)}(\lambda_L L/2 - \alpha\lambda_H H)\right]$$

The Poisson case is obvious in this setting. I claim we can also get the diffusion case by letting the number of events (sales) per period grow in the limit, so that what is observed are the aggregate of a great many sales per period. Depending on how we do this we should be able to get either the diffusion or the deterministic limit.

References

Abreu, D., P. Milgrom and D. Pearce(1991): .Information and Timing in Repeated Partnerships, *Econometrica*, 59(6), 1713-1733.

Faingold, E. and Y. Sannikov [2005] "Equilibrium Degeneracy and Reputation Effects in Continuous Time Games," mimeo

Fudenberg, D., D. K. Levine and E. Maskin

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

Mirlees

Sannikov, Y. (2004): .Games with Imperfectly Observed Actions in Continuous Time," mimeo.