Continuous Time Limits of Repeated Games with Imperfect Public Monitoring

Citation

Fudenberg, Drew, and David K. Levine. 2007. Continuous time limits of repeated games with imperfect public monitoring. *Review of Economic Dynamics* 10(2): 173-192.

Published Version

<http://dx.doi.org/10.1016/j.red.2007.02.002>

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:3196334>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Continuous Time Limits of Repeated Games with Imperfect Public Monitoring¹

Drew Fudenberg and David K. Levine²

This version: 5/25/2007

First version: 12/15/2005

Abstract: In a repeated game with imperfect public information, the set of equilibria depends on the way that the distribution of public signals varies with the players' actions. Recent research has focused on the case of "frequent monitoring," where the time interval between periods becomes small. Here we study a simple example of a commitment game with a long-run and short-run player in order to examine different specifications of how the signal distribution depends upon period length. We give a simple criterion for the existence of efficient equilibrium, and show that the efficiency of the equilibria that can be supported depends in an important way on the effect of the player's actions on the variance of the signals, and whether extreme values of the signals are "bad news" of "cheating" behavior, or "good news" of "cooperative" behavior.

¹ Financial support from NSF grants SES-03-1471 and SES-0426199 is gratefully acknowledged. We would like to thank Antonio Miguel Osorio da Costa, Eduardo Faingold, Yuliy Sannikov, and Andrzej Skrzypacz for helpful conversations, and Satoru Takahashi for careful proofreading.

² Departments of Economics, Harvard University, and Washington University in St. Louis.

1. Introduction

In a repeated game with imperfect public information, the set of equilibria depends on the way that the distribution of public signals varies with the players' actions. When considering the case of "frequent monitoring," where the time interval between periods becomes small, it seems natural to suppose that the distribution of signals changes in some way as the time period shrinks. In this paper, we model the dependency of the information structure on the period length by supposing that there the players observe the state of a fixed continuous-time process at the start of each period, and that this process is either Poisson or a diffusion.

Intuitively, if the public signal is "sales" or "revenues," it corresponds to the aggregate of a number of individual transactions, so that over small enough time intervals we would observe at most a single transaction. Even for a monetary aggregate that measures all transactions in an economy, in any given picosecond we are unlikely to observe more than a single trade, so the discrete Poisson process, then, is one natural way to model the frequent observation of revenues." In practice, however, it is often not practical or possible to observe at a high enough frequency to track every discrete event. Instead, what is observed over the relevant time period is an aggregate of many events, and under standard conditions this aggregate converges to a diffusion as the period between events and their size both become small at a particular relative rate. The continuous-time limit we compute here thus corresponds to the iterated limit where the observation period of the players, though short, is much longer than the period between events.³

Our goal is to illuminate some conceptual points about the relationship between discrete and continuous time repeated games, and not to present a general theory, so we specialize throughout the paper to a specific example of a repeated game between a single long-run player and a sequence of short-run opponents. In this setting, the best equilibrium payoff can be attained by a "grim" strategy that prescribes the efficient outcome so long as the public signal above a critical threshold. Our first main result,

³ We examine more general ways of passing to the continuous time limit in a companion paper, Fudenberg and Levine [2007]; this lets us explore the sensitivity of results about the diffusion case to the amount of "information aggregation" within each period.

Proposition 1, shows how the existence of efficient or non-trivial equilibria in the limit of time periods shrinking to zero can be determined by two properties of the limits of the probabilities p and q that punishment is triggered under the equilibrium action and defection, respectively. Specifically, the key variables are the limit of the signal-to-noise ratio $(q - p)/p$, which we denote by ρ , and the limit μ of the rate at which deviation increases transitions to the punishment regime $\mu = (q - p)/\tau$ where τ is the length of the period. We show that there is a non-trivial limit equilibrium if ρ is sufficiently large and $\mu > 0$, and that there is an efficient equilibrium in the iterated limit where first τ and then r go to 0 if $\rho = \infty$ and $\mu > 0$.

Proposition 1 applies for arbitrary specifications of how the signal structure depends on the period length; the remainder of this paper considers the case where the signals comes from observing an underlying Poisson or diffusion process. We find that the equilibrium set is larger (and so efficient outcomes are more likely to be supportable by equilibria) when the public signals correspond to the aggregation of a great many signals, that is, in the diffusion case, and that efficiency is less likely with Poisson signals. In addition, we find that when the signal is based on a diffusion what matters is the effect of the players' actions on the variance, of the aggregate, as opposed to its mean: Efficiency is more likely when the "tempting" or "cheating" actions generate a higher variance. (Note that in a Poisson process (aggregated or not), the mean and variance are linked, so actions that increase the variance must increase the mean.) Our results show that the case where player's actions control the drift but not the variance of a diffusion process, is a knife-edge, at least when the long-run player has only two actions, as the conclusions about the frequent-monitoring limit can change discontinuously if actions have even a small effect on the variance.⁴

Finally, we extend the result of Abreu, Milgrom and Pearce [1991] (AMP) who show that Poisson events that correspond to bad news, meaning increased likelihood of "cheating," lead to more efficient outcomes than Poisson events that correspond to good news. These results about the most efficient limit equilibria are summarized in the following table:

⁴ If the long-run player controls more than two actions, there may be mixed deviations that generate the same limit variance as the efficient action does. The implications of this for the limits of discrete-time equilibria has not yet been worked out.

	Poisson	Diffusion	Diffusion, constant variance
Bad News	Non-trivial	Efficient	Trivial
Good News	Trivial	Non-trivial	Trivial

Table 1: Most Efficient Limit Equilibrium Under Various Signal Structures

Because discrete-time games are simpler and more familiar than games in continuous time, our analysis helps provide intuition for existing results on continuous-time repeated games. In particular, we can use elementary calculus (l’Hopital’s rule) to show why diffusion signals with constant variance are relatively ineffective in supporting repeated play. Our methods also facilitate the analysis of diffusions where actions do change the variance of the signals.

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 [1994] (FLM) 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 $E(1) = \lim_{\delta \rightarrow 1} E(\delta)$; on the conditions of the FLM theorem, a payoff vector v is feasible and individually rational if and only if it is in $E(1)$. It is important to recall that the identification conditions used for this theorem are purely qualitative; when they are satisfied, the set $E(1)$ is independent of the exact nature of the distribution of signals and in particular of any quantitative measure of their “informativeness.” FLM also explain why the highest equilibrium payoff in symmetric strategies can be bounded away from efficiency when there are equilibrium payoffs that are symmetric and almost efficient.⁵

Sannikov [2005] characterizes the equilibrium payoffs in a repeated game with two long-run players in continuous time, where players control the mean of a vector-valued diffusion process; he shows that this set is not degenerate but that for a fixed

⁵ The FLM and FL results both use a “full-dimension” condition; see Fudenberg, Levine and Takahashi [2006] for a characterization of $E(1)$ without this condition.

interest rate r it can be bounded away from full efficiency, in contrast to the FLM result. Under a somewhat stronger identification decision (what FLM called a “product structure”), he proves a folk theorem for the limit $r \rightarrow 0$.

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 use 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.⁶ 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 Fudenberg-Levine characterization 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 can explain with our discrete-time methodology.

A second existing result that we explain 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. AMP 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 informativeness of the signals. Our setting of a repeated game between a long-run player and a sequence of short-run opponents is essentially equivalent to their model, as in each case there is no way to “efficiently punish” one player by simultaneously

⁶ The reason for this was first noted by Fudenberg, Kreps, and Maskin [1990], which covers the case of perfectly observed actions.

rewarding his opponent, and the only way to provide incentives is to lower the equilibrium payoff after some of the signals, and the size of $E(1)$ thus depends on the probability that punishment is triggered.⁷

This probability of punishment 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 the simple game we study, this minimum probability depends on a particular likelihood ratio that we identify. In the case of sampling from a fixed-intensity Poisson process, 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 AMP. In contrast, the key likelihood ratio converges to 0 when the signals correspond to sampling the diffusion process studied by Faingold and Sannikov, which provides a discrete-time explanation of their equilibrium degeneracy result.

In games with all long-run players, the identification conditions imply that there are equilibria where incentives can be provided at negligible efficiency cost by efficient punishments; this is what FLM call “enforcement on tangent hyperplanes.” Because the punishments can be efficient (i.e. tangential) their probability does not influence $E(1)$. This is related to Sannikov’s [2005] result that diffusion signals that satisfy an identification condition do allow non-trivial equilibria in games with all long-run players. 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.⁸

In contemporaneous work, Sannikov and Skrzypacz [2006] provide a linear-programming characterization of the limit of the equilibria of repeated games with two long-run players in discrete time as the period length shrinks and the interest rate goes to 0, where the public signal is derived by sampling an underlying continuous-time Levy process (a combination of a diffusion process and a Poisson process) whose parameters are independent of the sampling length. They show, loosely speaking, that near the

⁷ The FLM and FL results both use a “full-dimension” condition. See Fudenberg, Levine and Takahashi [2006] for a characterization of $E(1)$ without this condition and Mailath and Samuelson [2006] for a review of much of the related literature.

⁸ The diffusion case is more complicated because second-order terms are not negligible, so that the variance of the diffusion process does have an impact on the set $E(r)$.

boundary of the equilibrium set only the Poisson process can be used to provide non-tangential incentives, while both sorts of processes can be used to provide incentives on tangent hyperplanes. Our analysis differs in allowing the underlying process to vary with the sampling length, and in considering diffusions whose variance is influenced by the players' actions. Of course our analysis also differs in considering an example of games with a short-run player (so enforcement on tangent hyperplanes is not possible) as opposed to their treatment of games with two long-run players.

We should also acknowledge Hellwig and Schmidt's [2002] study of the limits of discrete-time principal-agent games as the time period shrinks. Instead of assuming that the discrete-time games correspond to sampling a diffusion process at discrete intervals, Hellwig and Schmidt suppose that the discrete-time games have a multinomial signal structure that converges to a diffusion as the time period shrinks, and compare the resulting limits to Holmstrom and Milgrom's analysis of the corresponding continuous-time game,. Thus their work resembles our companion paper more than it does this one.

2. The Repeated Commitment Game

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

		Player 2	
		L	R
Player 1	+1	$\underline{u}, 0$	$\bar{u}, 1$
	-1	$\underline{u}, 0$	$\bar{u} + g, -1$

Table 2: Stage-Game Payoffs

where $\underline{u} < \bar{u}$ and $g > 0$. In the stage 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 2 to play **R** by avoiding playing **-1**.

At the end of each play of the stage game, players observe a public signal $z \in \mathbb{R}$ that depends only on the action taken by player 1; player 2's action is publicly observed,

as is the outcome of a public randomizing device.⁹ The probability distribution of the public signal is $F(z | a_1)$. We assume that F is either differentiable and strictly increasing, or that it corresponds to a discrete random variable. (When F is strictly increasing we endow the real line with the Lebesgue sigma-algebra and suppose that strategies are Lebesgue measurable.) In either case, let $f(z | a_1)$ denote the density function. We assume the monotone likelihood ratio condition that $f(z | a_1 = -1)/f(z | a_1 = +1)$ 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.¹⁰ 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 $\delta = \exp(-r\tau)$ facing an infinite series of short-run opponents. We restrict attention to the set of *perfect public equilibria*, or *PPE*: 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.¹¹ The most favorable perfect public equilibrium for LR is characterized¹² by the largest value v that satisfies the constraints

$$\begin{aligned}
 v &= (1 - \delta)\bar{u} + \delta \int w(z)f(z | a_1 = +1)dz \\
 (C) \quad v &\geq (1 - \delta)(\bar{u} + g) + \delta \int w(z)f(z | a_1 = -1)dz \\
 v &\geq w(z) \geq \underline{u}
 \end{aligned}$$

⁹ Technically speaking the public information also includes the short-run player’s action, but since public randomizations are available we can restrict attention to strategies that ignore the past actions of the short-run player, and obtain the same set of outcomes of perfect public equilibria. To see this, observe that continuation payoffs can always be arranged by a public randomization between the best and worst equilibrium. If continuation payoffs depend on the play of the short-run player, the long-run player cares only about the expected value conditional on the signal of his own play. Since that expected value lies between the best and worst equilibrium, there is an equivalent equilibrium in which the continuation value is constant and equal to the conditional expected value.

¹⁰ Because player 1 has only two actions, this assumption is without loss of generality, as we can always re-order the signals so that it is satisfied.

¹¹ See Fudenberg and Tirole [1991] for a definition of this concept and an example of a non-public equilibrium in a game with public monitoring.

¹² The arguments of Fudenberg and Levine [1983] or Abreu, Pearce and Stachetti [1990] can be adapted to show that the set of PPE payoffs here is compact, so the best equilibrium payoff is well-defined.

or $v = \underline{u}$ 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 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 variable z with a positive density this condition is quite straightforward; Levin [2003] and Sannikov-Skrzypacz [2005] prove the analogous result for games with two long-run players. When the distribution of z has atoms, the argument 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 $\underline{z}^* < \bar{z}^*$ just below and above z^* , with probability $(\bar{z}^* - z^*)/(\bar{z}^* - \underline{z}^*)$ 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 cutoff $\underline{z}^*$ or cutoff $\bar{z}^*$.

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

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

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

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

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

$$\frac{f(z \mid a_1 = +1)}{f(z \mid a_1 = -1)} > \text{ or } < \nu$$

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 the case $z = z^*$ when z is discrete. Since in that case $\underline{u} \leq w(z^*) \leq v$ 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 simplified constraints

$$\begin{aligned}
(C') \quad & v = (1 - \delta)\bar{u} + \delta(v - p(v - w)) \\
& v = (1 - \delta)(\bar{u} + g) + \delta(v - q(v - w)) \\
& \underline{u} \leq w \leq v \\
& 0 \leq p, q \leq 1
\end{aligned}$$

Let the solution to this be v^* or $v^* = \underline{u}$ if no solution exists.¹³ 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 $w = \underline{u}$. Manipulating the first two lines of system (C') shows that $(1 - \delta)g = \delta(q - p)(v - w)$, and plugging this into the first line of (C') shows that if a solution exists, its value is

$$v = \bar{u} - \frac{pg}{q - p}.$$

So we conclude that the highest equilibrium payoff is

$$(1) \quad v^* = \max\left\{\bar{u} - \frac{pg}{q - p}, \underline{u}\right\}.$$

Note that this converges to the first best as $p/(q - p) \rightarrow 0$. It remains to determine when a solution to C' exists. Substitution into the equation for w shows that

$$w = \bar{u} - \frac{pg}{q - p} - \frac{(1 - \delta)}{\delta} \frac{g}{q - p}.$$

This payoff is feasible if it is at least $\underline{u}$, which is equivalent to

$$(2) \quad \frac{(\bar{u} - \underline{u})(q - p)}{g} - 1 \geq \frac{(1 - \delta)}{\delta p}.$$

Moreover, because $\delta < 1$, when (2) is satisfied, we have

$$\frac{(\bar{u} - \underline{u})(q - p)}{g} > 1,$$

which implies that

¹³ Here we use the fact that the incentive constraint is binding at the optimum, this is why the second line is an equality and not an inequality.

$$v^* = \bar{u} - \frac{pg}{q-p} > \underline{u}.$$

This proves the following result:

Corollary 2: *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 $p, q \in [0, 1]$ that satisfy (2). If such (p, q) exist for a given δ , they exist for all $\delta' > \delta$.*

Inspecting (1) and (2) shows that the highest equilibrium payoff is obtained by choosing z^* to maximize the “signal to noise” ratio

$$\frac{q-p}{p}$$

subject to the constraint that (2) is satisfied. In games with a finite set of signals, the likelihood ratio is obviously finite for any cut-off such that $p > 0$, so the best equilibrium payoff is bounded away from the first best irrespective of δ . This need not be the case when the set of signals is infinite. Indeed, as noted by Mirrlees [1974], this likelihood ratio can become infinite when the signals are normally distributed with a fixed variance and mean that depends on action. In the static principal-agent problem Mirrlees considered, the set of transfers was unbounded, so the fact that the signal to noise ratio can be made arbitrarily large implied that the first-best outcome can be approximated arbitrarily closely. In our setting, in contrast, because of the bound on the continuation payoffs, the first best can not be approximated arbitrarily closely for any fixed $\delta < 1$, but it can be approximated in the limit as $\delta \rightarrow 1$. Intuitively, as $\delta \rightarrow 1$, the bounds on continuation payoffs become unimportant, because even a small change in continuation payoff outweighs any one period gain. We say more about the case of unbounded signal to noise ratios and the normal distribution in section 4.

3: Sending the Time Interval to Zero

Our interest is in how the set of PPE payoffs varies with the period length, and in particular its behavior as the time period shrinks to zero, because we want to relate this limit to the predictions of various continuous-time models. To facilitate taking the continuous-time limit, we substitute $e^{-r\tau}$ into (2) and rearrange terms, to obtain

$$(3) \quad \frac{e^{r\tau} - 1}{\tau} \leq \frac{(q - p)(\bar{u} - \underline{u})}{\tau g} - \frac{p}{\tau}.$$

Let p and q be functions of τ such that $p(\tau)$ and $q(\tau)$ satisfy (3) for each τ ; we say that p and q are regular if the limits $\rho = \lim_{\tau \rightarrow 0} (q(\tau) - p(\tau)) / p(\tau)$ and $\mu = \lim_{\tau \rightarrow 0} (q(\tau) - p(\tau)) / \tau$ exist. The first limit ρ can be thought of as the limit of 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 difference between the bad news signal arrival rate over the good news rate. When p and q are regular, the limit of the right-hand-side of (3) exists, resulting in the limit inequality

$$(4) \quad r \leq (\mu / \rho) (((\bar{u} - \underline{u}) / g) \rho - 1)$$

and moreover

$$(5) \quad \lim_{\tau \rightarrow 0} v^* = \lim_{\tau \rightarrow 0} \bar{u} - \frac{gp(\tau)}{q(\tau) - p(\tau)} = \bar{u} - \frac{g}{\rho}.$$

Now fix regular functions (p, q) . If there exists positive $\bar{\tau}, \bar{r}$ and ε such that for all non-negative smaller values $0 < \tau < \bar{\tau}, 0 < r < \bar{r}$ the game with period length τ and interest rate r has an equilibrium with punishment probabilities $p(\tau)$ and $q(\tau)$ with payoff at least $\underline{u} + \varepsilon$, we say (p, q) supports a non-trivial limit equilibrium. If for all $(\tau, r) \rightarrow (0, 0)$ there are equilibria with punishment probabilities $p(\tau)$ and $q(\tau)$ that have payoffs converging to $\bar{u}$, we say that (p, q) supports an efficient patient equilibrium. We say that there is a non-trivial or efficient limit if there is a regular (p, q) that supports it.

Note that the definition of a non-trivial limit equilibrium requires the payoff in question to be supportable as an equilibrium when the interest rate r is held fixed as the period length τ goes to 0. The definition of an efficient patient equilibrium requires the interest rate to go to 0 as well. However the efficient payoff must be attained in the limit regardless of the relative rates at which τ and r converge, so that in particular efficiency must be obtained if we first send τ to 0 with r fixed and only then decrease r . The other

order of limits, with r becoming small for fixed τ , corresponds to the usual folk-theorem analysis in discrete-time games.

Proposition 1: *Regular (p, q) support a non-trivial limit equilibrium if $\rho > g/(\bar{u} - \underline{u})$ and $\mu > 0$; it supports an efficient limit equilibrium if $\rho = \infty$ and $\mu > 0$. Conversely, there is a non-trivial limit equilibrium only if there is a (p, q) with $\rho > g/(\bar{u} - \underline{u})$ and $\mu > 0$, and there is an efficient patient equilibrium only if there is a regular (p, q) with $\rho = \infty$ and $\mu > 0$.*

Proof: If $\rho > g/(\bar{u} - \underline{u})$ and $\mu > 0$, then the right hand side of (4) is strictly positive, so we can find $\bar{r} > 0$ such that (4) is satisfied for all sufficiently small τ and all $r < \bar{r}$. If $\rho = \infty$ and $\mu > 0$, then (4) is positive for small r , and moreover from (5) the corresponding limit payoff is efficient.

Conversely, if (p, q) is regular and either $\rho \leq g/(\bar{u} - \underline{u})$ or $\mu = 0$, the right-hand side of (4) is non-positive, and so for any fixed positive r (4) must be violated for τ sufficiently small, so there cannot be an equilibrium with payoffs above $\underline{u}$. Finally, if $\mu > 0$ and $\rho < \infty$ then from (5) the limit payoff cannot be efficient, $v^* > \bar{u}$ and (3) is satisfied for τ sufficiently large. From (1) $v^* > \underline{u}$, and then (3) is positive if $\mu > 0$. Moreover, from (2) $v^* = \underline{u}$ if and only if $\rho = \infty$.

□

The proof of Proposition 1 does not use the fact that the optimal equilibrium has continuation payoff $\underline{u}$ after bad signals: For the existence of non-trivial limit equilibria, it is necessary and sufficient that $\rho > g/(\bar{u} - \underline{u})$ and $\mu > 0$ for some family of cut-point equilibria. Of course the conditions are also necessary and sufficient for limits of families of optimal equilibria. Note that Proposition 1's sufficient condition for a non-trivial limit equilibrium is an extension of Proposition 2 of AMP, which applies only to the case of sampling from a fixed Poisson distribution that we study in the next section.¹⁴

Restricting to these equilibria gives a useful lemma that makes it easier to check the conditions of Proposition 1.

¹⁴ Their result covers only pure strategy equilibria, and has no condition on μ , which is implicitly assumed to be positive.

Lemma 3: Suppose that the interest rate $r(\tau)$ depends on the period length. If $(p(\tau), q(\tau))$ are optimal non-trivial equilibria for $(\tau, r(\tau))$ and $\rho = \lim_{\tau \rightarrow 0} (q(\tau) - p(\tau)) / p(\tau)$, $\mu = \lim_{\tau \rightarrow 0} (q(\tau) - p(\tau)) / \tau$, $r = \lim_{\tau \rightarrow 0} r(\tau)$ and $v^* = \lim v^*(\tau)$ exist, then $\mu / \rho = r(\bar{u} - v^*) / (v^* - \underline{u})$.

Proof: Substituting $w = \underline{u}$ into the equation for the equilibrium payoff in the first line of C' we have

$$v^*(\tau) = (1 - \delta)\bar{u} + \delta[v^*(\tau) - p(\tau)(v^*(\tau) - \underline{u})].$$

This may be rearranged as

$$\frac{p(\tau)}{\tau} = \frac{(1 - e^{-r\tau})}{e^{-r\tau}\tau} \frac{\bar{u} - v^*(\tau)}{v^*(\tau) - \underline{u}}.$$

We can rewrite

$$p(\tau) / \tau = \frac{(q(\tau) - p(\tau)) / \tau}{(q(\tau) - p(\tau)) / p(\tau)},$$

so

$$\frac{(q(\tau) - p(\tau)) / \tau}{(q(\tau) - p(\tau)) / p(\tau)} = \frac{(1 - e^{-r\tau})}{e^{-r\tau}\tau} \frac{\bar{u} - v^*(\tau)}{v^*(\tau) - \underline{u}}$$

and taking the limit $\tau \rightarrow 0$ gives the desired result.

□

Because a non-trivial limit equilibrium requires that the limit payoff exceeds $\underline{u}$, lemma 1 implies that $\mu / \rho = \lim_{\tau \rightarrow 0} p(\tau) / \tau$ must be finite, a fact we use below. Intuitively, if the optimal equilibria were to have $\lim_{\tau \rightarrow 0} p(\tau) / \tau = \infty$, when τ is small the probability of quickly reaching the punishment phase with a continuation payoff of $\underline{u}$ is close to one. That implies that the equilibrium payoff is close to $\underline{u}$.

In applying Proposition 1 to a family of optimal equilibria $(p(\tau), q(\tau))$ corresponding to the game with period length and interest rate $(\tau, r(\tau))$, it is possible that this family is not regular. In this case, we can restrict attention to a sequence of optimal equilibria that are regular, that is, along a subsequence τ^n for which $\lim_{\tau^n \rightarrow 0} (q(\tau^n) - p(\tau^n)) / p(\tau^n)$ and $\lim_{\tau^n \rightarrow 0} (q(\tau^n) - p(\tau^n)) / \tau^n$ do converge. For

notational simplicity we have stated our results for the case in which the optimal equilibria are regular, but they apply equally well to subsequences that are regular.

4. Fixed-Intensity Poisson Signals

In this section we suppose that the public signal of the long-run player's action is generated by observing a Poisson process in continuous time. The players simultaneously observe the result of this process every τ units of time, and the process is held fixed as the time period becomes small. This means that the probability of two or more events in a time interval is second order in τ .

The arrival rate of the Poisson process is λ_p if the action taken by LR is $+1$ and λ_q if the action taken by LR is -1 . As in Abreu, Milgrom and Pearce's [1991] analysis of a partnership game with two long-run players, a critical role is played by whether the Poisson event 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 event is “bad news.” In this case we take random variable z to be the non-negative discrete random variable representing the number of events that occurred during the previous interval of length τ . If $\lambda_q < \lambda_p$ the event 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 events received during the previous interval of length τ . In this way we preserve the convention that high z is bad news. As in Abreu, Milgrom and Pearce [1991], 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 case of bad-news events $\lambda_q > \lambda_p$. The cutoff point is how many events must occur before the punishment $v - w$ is triggered. If we take the cutoff to be two or more events, then the probability of triggering punishment is of order τ^2 ; this implies that $\mu = \lim_{\tau \rightarrow 0} (q(\tau) - p(\tau)) / \tau$ is equal to 0, which from lemma 1 rules out a non-trivial limit equilibrium.

So the only interesting cutoff is to punish whenever any event is received. Consider the suboptimal rule in which punishment occurs with probability α when one or more events occur, independent of the period length. The probability of one or more events is $p(\tau) = 1 - e^{-\lambda_p \tau}$, or $q(\tau) = 1 - e^{-\lambda_q \tau}$, as the long-run player plays -1 or $+1$. We may then directly compute

$\rho = \alpha(\lambda_q - \lambda_p) / \alpha\lambda_p = (\lambda_q - \lambda_p) / \lambda_p$ and $\mu = \lambda_q - \lambda_p$. Hence the condition for a non-trivial limit equilibrium is $\rho = (\lambda_q - \lambda_p) / \lambda_p > g / (\bar{u} - \underline{u})$, and the best limit equilibrium payoff is $v^* = \bar{u} - g\lambda_p / (\lambda_q - \lambda_p)$, as shown by Proposition 2 of APM.¹⁵ The significant feature of this solution is that it is independent of the payoff $\underline{u}$.

Now we analyze the case of “good news” events $\lambda_q < \lambda_p$. Here the punishment is triggered by a small number of events, rather than a large number. If there is to be any punishment at all, then punishment must certainly occur when no event occurs. Suppose the probability of punishment when there is no event 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 case of “good news” events, 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 case of bad-news events) or if there are several long run players so that punishments can take the form of efficient transfers.

5. Diffusion Signals With Common Variance

Faingold and Sannikov [2005] study the case where 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. In this section, we suppose that the players observe the underlying process at intervals of length τ , as in Sannikov and Skrzypacz [2005], who study repeated games with two long run players.¹⁶ To do this, we apply Proposition 3, and consider a slight generalization of the diffusion assumption: we

¹⁵ Because the probability of having more than a single event vanishes at rate τ^2 , the limit equilibrium payoff $v^* = \bar{u} - g\lambda_p / (\lambda_q - \lambda_p)$ computed from this sub-optimal rule is the same as the payoff computed assuming certain punishment whenever there are two or more events. We could also allow α to depend on the period length without changing our conclusions.

¹⁶ Thus our approach differs from that of Hellwig and Schmidt [2002], which constructs the limit diffusion process and limit continuous time game as the limit of discrete-time games with multinomial signals, and the signal process and payoff functions are rescaled in the appropriate way as $\tau \rightarrow 0$.

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 show below that there is no sequence of cutoffs that supports a non-trivial limit equilibrium. Before doing so, we examine a simpler result that provides intuition for the general one – it is not possible to support a non-trivial limit when the probability p of punishment along the equilibrium path is held fixed as $\tau \rightarrow 0$.

Let $z^*(\tau)$ denote the cutoff when the period length is τ , and define the normalized cutoff

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

Then

$$p = \Phi(-\zeta(\tau))$$

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

If p is held fixed, then $\zeta = -\Phi^{-1}(p)$ is fixed as well, and

$$\rho = \lim_{\tau \rightarrow 0} (q - p) / p = \lim_{\tau \rightarrow 0} \left(\Phi\left(\frac{2\tau^{1-\alpha}}{\sigma} - \zeta\right) - \Phi(-\zeta) \right) / \Phi(-\zeta) = 0.$$

From Proposition 1, this sequence cannot support a non-trivial limit equilibrium.

Of course this does not show that there are not non-trivial limits for all sequences of cut-offs. In particular, it is always possible to find a series of cut-offs such that $\rho = \infty$, by sending the normalized cutoff to ∞ . However, in contrast to the principal-agent problem considered by Mirrlees, rewards and punishments in our model are bounded, so $\rho = \infty$ is not sufficient for a non-trivial limit equilibrium; we also need to ensure that $\mu > 0$, so that the increase in punishment caused by cheating vanishes no faster than the period length. We will show that the combination of these two conditions

is only possible if α is (at least) one; a similar result for the case $\alpha = 1/2$ can be found in Sannikov and Skrzypacz [2005].

Proposition 2: *For any $\alpha < 1$ there is no non-trivial limit equilibrium.*

Proof: We will show that for τ sufficiently small, (3) is necessarily violated.¹⁷

Let $z^*(\tau)$ denote the cutoff when the period length is τ . It is convenient to work with the normalized cutoff $\zeta(\tau)$ defined above. From Proposition 2, if there is to be a non-trivial limit equilibrium, we must have $\rho = \lim_{\tau^n \rightarrow 0} (q(\tau^n) - p(\tau^n)) / p(\tau^n) > 0$ and $\mu > 0$, and this second inequality requires that $(q - p) / \tau$ is bounded away from zero. We will show that this implies that $\mu / \rho = \lim p(\tau) / \tau = \infty$, which contradicts Lemma 1.

Let $\phi = \Phi'$ denotes the density of the standard normal. First we compute $(q - p) / \tau$ using the mean value theorem to observe that for each τ there is a number $f(\tau)$, $0 \leq f(\tau) \leq 1$, such that

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

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

$$\begin{aligned} \phi \left(-\zeta(\tau) + f(\tau) \frac{2\tau^{1-\alpha}}{\sigma} \right) &= c(\tau) \sigma \tau^\alpha / 2 \\ \frac{1}{\sqrt{2\pi}} \exp \left(-\frac{1}{2} \left(-\zeta(\tau) + f(\tau) \frac{2\tau^{1-\alpha}}{\sigma} \right)^2 \right) &= c(\tau) \sigma \tau^\alpha / 2 \\ \left(-\zeta(\tau) + f(\tau) \frac{2\tau^{1-\alpha}}{\sigma} \right)^2 &= -2 \log(\sqrt{2\pi} c(\tau) \sigma / 2) - 2\alpha \log(\tau) \\ \zeta(\tau) &= \sqrt{-2 \log(\sqrt{2\pi} c(\tau) \sigma / 2) - 2\alpha \log(\tau)} + f(\tau) \frac{2\tau^{1-\alpha}}{\sigma} \end{aligned}$$

We now want to use this to show that $p / \tau \rightarrow \infty$. Since we have assumed that $c(\tau)$ is bounded away from zero, $\zeta(\tau) \leq \sqrt{-\log b - 2\alpha \log \tau} + a\tau^{1-\alpha}$. This gives

¹⁷ Note the difference with the result of Muller [2000], who observes that the principal can implement the first best in a discrete-time repeated moral hazard problem with diffusion signals. The difference stems from the fact that Muller allows unbounded rewards, while the feasible payoffs here are bounded above and below.

$p/\tau \geq \Phi(-\sqrt{-\log b - 2\alpha \log \tau} - a\tau^{1-\alpha})/\tau$. We now apply the mean value theorem to the function $g(\tau) = \Phi(-\sqrt{-\log b - 2\alpha \log \tau} - a\tau^{1-\alpha})/\tau$, where we define $g(0) = 0$.¹⁸

Again using the mean value theorem we know that for some $\tau' \in [0, \tau]$, we have

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

Denote the expression in the last step of this series of equalities by $J(\tau')$; this has the form $k(x(\tau') - \gamma) \exp\left(\frac{-1}{x(\tau')}\right) y(\tau')$, where $x(\tau') = \frac{\tau'^{\alpha-1}}{\sqrt{-\log b - 2\alpha \log \tau'}}$, $\gamma = a(1-\alpha)$, and $y(\tau') = \exp(-1/2) (a^2(\tau')^{2-2\alpha}) \rightarrow 1$. As $\tau' \rightarrow 0$, $-\log(\tau')$ diverges to infinity more slowly than any fixed negative power of τ' so $x(\tau') \rightarrow \infty$, so $J(\tau') \rightarrow \infty$ and thus $p/\tau \rightarrow \infty$

☑

Notice that in the limit for $\alpha < 1$ the equilibrium collapses to the static Nash; in particular this is true even when $\alpha > 1/2$, so that 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?

¹⁸ We can apply the mean value theorem here because g is continuous on the closed interval $[0, \tau]$ and differentiable on its interior.

It is useful to contrast the diffusion case $\alpha = 1/2$ to the case $\alpha = 1$, where the mean and the standard error are both of order τ . When we take the limit of such a sequence of processes, the limit is a deterministic process without noise.

Proposition 3: *If $\alpha = 1$ there is an efficient patient equilibrium.*

Proof: From Proposition 1, it is enough to find functions p, q that satisfy (3) and such that $\rho = \infty$ and $\mu > 0$. As in the proof of Proposition 2, it is convenient to work with the normalized cutoff

$$\zeta(\tau) = \frac{1 + z^*(\tau)/\tau}{\sigma},$$

Then

$$p = \Phi(-\zeta(\tau))$$

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

Observe then that if $\zeta(\tau) \rightarrow \infty$ then as in Mirrlees [1974] the likelihood ratio $q(\tau)/p(\tau) \rightarrow \infty$ implying that $\rho = \infty$. Moreover if we take

$$\zeta(\tau) = \sqrt{\log(1/\tau)}$$

then

$$\begin{aligned} \mu &= \lim_{\tau \rightarrow 0} \frac{q(\tau) - p(\tau)}{\tau} \\ &\geq \lim_{\tau \rightarrow 0} \frac{2\phi(-\zeta(\tau))}{\sigma\tau} = \lim_{\tau \rightarrow 0} \frac{2}{\sigma\sqrt{2\pi}} \frac{e^{-\frac{[-\zeta(\tau)]^2}{2}}}{\tau} \\ &= \lim_{\tau \rightarrow 0} \frac{2}{\sigma\sqrt{2\pi}} \frac{e^{-\log(1/\tau)}}{\tau} = \frac{2}{\sigma\sqrt{2\pi}} > 0 \end{aligned}$$

as required in Proposition 1.

☑

6. Diffusion Signals with Unequal Variances

In the previous section, we followed Faingold and Sannikov [2005] by assuming that the instantaneous variance of the diffusion process was independent of the action

taken. We now look at the opposite case. We assume that over an interval of length τ the signals have variance $\sigma_{+1}^2\tau$ or $\sigma_{-1}^2\tau$ as the action chosen is either $a_1 = +1$ or -1 ; the means are still $\mu_{+1}\tau$ and $\mu_{-1}\tau$. With unequal variances, the monotone likelihood ratio condition necessarily fails. However, we will show below, that the optimum is to have two cutoff-points, rather than one. In other words, the relevant cutoff rules are of the form $\underline{z}(\tau), \bar{z}(\tau)$, together with strategies that treat signals that are extreme $z < \underline{z}(\tau), z > \bar{z}(\tau)$ in one way, and signals that are intermediate $\underline{z}(\tau) \leq z \leq \bar{z}(\tau)$ the other way.

In what follows, it will generally be useful to define the normalized cutoffs

$$\bar{\zeta}(\tau) = \frac{\bar{z}(\tau) - \mu_{+1}\tau}{\sigma_{+1}\tau^{1/2}} \text{ and } \underline{\zeta}(\tau) = \frac{\underline{z}(\tau) - \mu_{+1}\tau}{\sigma_{+1}\tau^{1/2}}$$

together with the regions $\zeta < \underline{\zeta}(\tau), z > \bar{\zeta}(\tau)$ and $\underline{\zeta}(\tau) \leq \zeta \leq \bar{\zeta}(\tau)$. This gives rise also to the inverse relations

$$\sigma_{+1}\tau^{1/2}\underline{\zeta}(\tau) + \mu_{+1}\tau = \underline{z}(\tau)$$

$$\sigma_{+1}\tau^{1/2}\bar{\zeta}(\tau) + \mu_{+1}\tau = \bar{z}(\tau).$$

6a) **Extreme values are bad news:** $\sigma_{-1} > \sigma_{+1}$.

We will show that in this case there is an efficient patient equilibrium. Here punishment should occur whenever the realized value of the signal z is extreme. To prove that there is an efficient patient equilibrium, we do not need to compute the optimal cutoff rule; we simply need to find a cutoff rule that supports equilibrium payoffs that are arbitrarily close to efficiency. To do this we fix a positive number ζ^* and consider the cutoff rule $\underline{\zeta}(\tau) = -\zeta^*, \bar{\zeta}(\tau) = \zeta^*$, which implies that $\underline{z}(\tau) = -\sigma_{+1}\tau^{1/2}\zeta^* + \mu_{+1}\tau$ and $\bar{z}(\tau) = \sigma_{+1}\tau^{1/2}\zeta^* + \mu_{+1}\tau$. Note that this cutoff rule is symmetric around the mean of the $+1$ action, and that the normalized cut-offs are independent of τ . Then the probability of punishment under action $+1$ is $p^* = 2\Phi(-\zeta^*)$, and the probability of punishment under action -1 is

$$\begin{aligned}
q(\tau) &= \Phi\left(\frac{z(\tau) - \mu_{-1}\tau}{\sigma_{-1}\tau^{1/2}}\right) + 1 - \Phi\left(\frac{\bar{z}(\tau) - \mu_{-1}\tau}{\sigma_{-1}\tau^{1/2}}\right) \\
&= \Phi\left(\frac{z(\tau) - \mu_{-1}\tau}{\sigma_{-1}\tau^{1/2}}\right) + \Phi\left(\frac{-\bar{z}(\tau) + \mu_{-1}\tau}{\sigma_{-1}\tau^{1/2}}\right) \\
&= \Phi\left(\frac{-\sigma_{+1}\tau^{1/2}\zeta^* + \mu_{+1}\tau - \mu_{-1}\tau}{\sigma_{-1}\tau^{1/2}}\right) + \Phi\left(\frac{-\sigma_{+1}\tau^{1/2}\zeta^* + \mu_{-1}\tau - \mu_{+1}\tau}{\sigma_{-1}\tau^{1/2}}\right) \\
&\rightarrow 2\Phi\left(-\frac{\sigma_{+1}}{\sigma_{-1}}\zeta^*\right) \equiv q^*
\end{aligned}$$

This shows that the effect of the actions on the means has vanishing impact on the punishment probability, essentially because the means under the two actions are converging at rate τ while the standard deviations only converge at rate $\tau^{1/2}$.

Proposition 4: *If $\sigma_{-1}/\sigma_{+1} > 1$ then for any interest rate there are equilibria whose payoffs are arbitrarily close to efficiency.*

Proof: Set

$$\kappa(\zeta^*) = \frac{\Phi(-\zeta^*)}{\Phi\left(-\left(\frac{\sigma_{+1}}{\sigma_{-1}}\right)\zeta^*\right)}.$$

Then $\kappa(0) = 1$, and using L'Hopital's rule we have that

$$\begin{aligned}
\lim_{\zeta \rightarrow \infty} \kappa(\zeta) &= \left(\frac{\sigma_{-1}}{\sigma_{+1}}\right) \lim_{\zeta \rightarrow \infty} \frac{\phi(-\zeta^*)}{\phi\left(-\left(\frac{\sigma_{+1}}{\sigma_{-1}}\right)\zeta^*\right)} \\
&= \left(\frac{\sigma_{-1}}{\sigma_{+1}}\right) \lim_{\zeta \rightarrow \infty} \frac{\exp(-\zeta^{*2}/2)}{\exp\left(-\left(\frac{\sigma_{+1}}{\sigma_{-1}}\right)^2 \zeta^{*2}/2\right)} \\
&= \left(\frac{\sigma_{-1}}{\sigma_{+1}}\right) \lim_{\zeta \rightarrow \infty} \exp\left(-\left(1 - \left(\frac{\sigma_{+1}}{\sigma_{-1}}\right)^2\right)\zeta^{*2}\right) = 0
\end{aligned}$$

Hence by taking ζ^* sufficiently large we can make $\kappa(\zeta^*)$ arbitrarily close to 0.

With these trigger strategies, the payoffs converge to

$$v = \bar{u} - \frac{gp^*/q^*}{1 - p^*/q^*}.$$

This is more than $\underline{u}$ whenever $\kappa(\zeta^*)$ is sufficiently small, so (1) is satisfied; moreover it converges to $\bar{u}$ as $\kappa \rightarrow 0$.

From Corollary 1, the other condition for equilibrium is

$$p^* \left(\frac{(\bar{u} - \underline{u})}{g} (q(\tau)/p^* - 1) - 1 \right) > e^{r\tau} - 1.$$

Observe that $q(\tau)/p^*$ converges to $1/\kappa$ so for small enough κ this inequality will hold for all sufficiently small τ .

☑

6b) Extreme values are good news: $\sigma_{+1} > \sigma_{-1}$.

In this case we have two results, one positive and one negative. We begin with the positive result, which is simpler, because we can prove it using the same cutoff rule we used above $\zeta(\tau) = -\zeta^*, \bar{\zeta}(\tau) = \zeta^*$.

Proposition 5: *For any fixed r there is a $\lambda > 1$ such that if $\sigma_{+1}/\sigma_{-1} > \lambda$ there is a non-trivial limit equilibrium. Moreover, as $\lambda \rightarrow \infty$, the best limit payoff converges to the first best.*

Proof: Set $v(\kappa) = \bar{u} - \frac{\kappa g}{1 - \kappa}$, and set $\bar{\kappa} = \frac{\bar{u} - \underline{u}}{\bar{u} - \underline{u} + g}$; note that $v(\kappa) > \underline{u}$ exactly when $\kappa > \bar{\kappa}$ and that $\lim_{\kappa \rightarrow 0} v(\kappa) = \underline{u}$. Set

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

Then from l'Hopital's rule

$$\kappa(0) = \frac{\sigma_{-1}}{\sigma_{+1}} < 1$$

and κ is monotone increasing, with $\lim_{\zeta \rightarrow \infty} \kappa(\zeta) = 1$. Thus we can set $\lambda = 1/\bar{\kappa}$, and whenever $\sigma_{+1}/\sigma_{-1} > \lambda$, there will be a ζ such that $\kappa(\zeta) < \bar{\kappa}$. In this case we claim

that there is an equilibrium with cutoff ζ and payoff $v(\kappa(\zeta))$ for all τ less than some $\bar{\tau} > 0$.

To see this, note that with these strategies the payoff is

$$v = \bar{u} - \frac{p^G(\zeta)g}{q^G(\zeta) - p(\zeta)} > \underline{u},$$

so (1) is satisfied by construction. From Corollary 1, the other condition for equilibrium is

$$\frac{(\bar{u} - \underline{u})}{g}(q - p) - p > e^{r\tau} - 1.$$

Because p and q are held fixed independent of τ , and the left-hand side of this inequality is positive, the inequality will hold for all sufficiently small τ .

Finally, note that as $\lambda \rightarrow \infty$, $\kappa(0) \rightarrow 0$, so sending ζ to 0 as τ goes to 0 results in payoffs that converge to the first best.

☑

For the next result we need to consider the optimal equilibria.

Lemma 4: *The optimal policy is determined by a number $B(\tau)$ such that*

$$(\sigma_{+1}^2 - \sigma_{-1}^2)\zeta^2 + 2(\mu_{-1} - \mu_{+1})\sigma_{+1}\tau^{1/2}\zeta + B(\tau) = 0$$

has two real roots $\underline{\zeta}(\tau) < \bar{\zeta}(\tau)$ and the policy of punishing when $\underline{\zeta}(\tau) \leq \zeta \leq \bar{\zeta}(\tau)$.

Proof: Recall from the proof of Lemma 1, prior to the use of the MLRP condition, that the optimal strategy must maximize the Lagrangean

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

In other words, $w(z)$ must be minimized (punishment) when $f(z | a_1 = +1) - \nu f(z | a_1 = -1)$ is negative, and take on the maximum value of v when this expression is positive. So the punishment region is bounded by critical values of z where

$$\frac{1}{\sqrt{2\pi\tau}\sigma_{+1}} e^{-\frac{1}{2}\frac{(-z-\mu_{+1}\tau)^2}{\sigma_{+1}^2\tau}} = \nu \frac{1}{\sqrt{2\pi\tau}\sigma_{-1}} e^{-\frac{1}{2}\frac{(-z-\mu_{-1}\tau)^2}{\sigma_{-1}^2\tau}}$$

Taking logs, this becomes the quadratic

$$\sigma_{+1}^2(-z - \mu_{-1}\tau)^2 - \sigma_{-1}^2(-z - \mu_{+1}\tau)^2 = 2\tau\sigma_{+1}^2\sigma_{-1}^2 \log\left(\frac{\sigma_{+1}}{\sigma_{-1}}\nu\right).$$

Substituting the normalization

$$-z = -\sigma_{+1}\tau^{1/2}\zeta(\tau) + \mu_{+1}\tau$$

then gives

$$(-\zeta\sigma_{+1} - (\mu_{-1} - \mu_{+1})\tau^{1/2})^2 - \sigma_{-1}^2\zeta^2 = 2\sigma_{-1}^2 \log\left(\frac{\sigma_{+1}}{\sigma_{-1}}\nu\right) / \tau.$$

✓

Since large observations are good news, punishment occurs when observations take on intermediate values so

$$p(\tau) = \Phi(\bar{\zeta}(\tau)) - \Phi(\underline{\zeta}(\tau))$$

and

$$\begin{aligned} q(\tau) &= \Phi\left(\frac{\bar{z}(\tau) - \mu_{-1}\tau}{\sigma_{-1}\tau^{1/2}}\right) - \Phi\left(\frac{\underline{z}(\tau) - \mu_{-1}\tau}{\sigma_{-1}\tau^{1/2}}\right) \\ &= \Phi\left(\frac{\sigma_{+1}\tau^{1/2}\bar{\zeta}(\tau) + \mu_{+1}\tau - \mu_{-1}\tau}{\sigma_{-1}\tau^{1/2}}\right) - \Phi\left(\frac{\sigma_{+1}\tau^{1/2}\underline{\zeta}(\tau) + \mu_{+1}\tau - \mu_{-1}\tau}{\sigma_{-1}\tau^{1/2}}\right) \end{aligned}$$

Proposition 6: *a) There is a $\lambda > 1$ such that if $1 < \sigma_{+1} / \sigma_{-1} < \lambda$ there is no non-trivial limit equilibrium.*

b) If $\sigma_{+1} / \sigma_{-1} > 1$ there is no efficient patient equilibrium.

Proof: Suppose there is a non-trivial equilibrium; then as we indicated in the discussion after Lemma 1, we know that $p(\tau) \rightarrow 0$ and since $p(\tau) = \Phi(-\bar{\zeta}(\tau)) - \Phi(-\underline{\zeta}(\tau)) \rightarrow 0$, it follows that $\bar{\zeta}(\tau) - \underline{\zeta}(\tau) \rightarrow 0$. From the quadratic formula and Lemma 8, we can compute

$$\underline{\zeta}(\tau) = -\frac{2(\mu_{-1} - \mu_{+1})\sigma_{+1}}{\sigma_{+1}^2 - \sigma_{-1}^2}\tau^{1/2} - \bar{\zeta}(\tau),$$

from which it follows that $\bar{\zeta}(\tau) + \underline{\zeta}(\tau) \rightarrow 0$, so we conclude that $\underline{\zeta}(\tau), \bar{\zeta}(\tau) \rightarrow 0$.

Now we calculate ρ

$$\rho = \lim_{\tau \rightarrow 0} \frac{\Phi\left(\frac{\sigma_{+1}\tau^{1/2}\bar{\zeta}(\tau) + \mu_{+1}\tau - \mu_{-1}\tau}{\sigma_{-1}\tau^{1/2}}\right) - \Phi\left(\frac{\sigma_{+1}\tau^{1/2}\underline{\zeta}(\tau) + \mu_{+1}\tau - \mu_{-1}\tau}{\sigma_{-1}\tau^{1/2}}\right) - (\Phi(\bar{\zeta}(\tau)) - \Phi(\underline{\zeta}(\tau)))}{\Phi(\bar{\zeta}(\tau)) - \Phi(\underline{\zeta}(\tau))}$$

By the mean value theorem, we can find $\underline{\zeta}(\tau) \leq \zeta^*(\tau), \zeta^{**}(\tau) \leq \bar{\zeta}(\tau)$ such that

$$\begin{aligned} \rho &= \lim_{\tau \rightarrow 0} \frac{\phi\left(\frac{\sigma_{+1}\tau^{1/2}\zeta^{**}(\tau) + \mu_{+1}\tau - \mu_{-1}\tau}{\sigma_{-1}\tau^{1/2}}\right) \frac{\sigma_{+1}}{\sigma_{-1}} (\bar{\zeta}(\tau) - \underline{\zeta}(\tau))}{\phi(\zeta^*(\tau)) (\bar{\zeta}(\tau) - \underline{\zeta}(\tau))} - 1, \\ &= \frac{\sigma_{+1} - \sigma_{-1}}{\sigma_{-1}} \end{aligned}$$

where the final step follows from $\underline{\zeta}(\tau), \bar{\zeta}(\tau) \rightarrow 0$, implying $\zeta^*(\tau), \zeta^{**}(\tau) \rightarrow 0$. So we conclude that there is no efficient patient equilibrium, and that there is no non-trivial limit if

$$\rho = \frac{\sigma_{+1} - \sigma_{-1}}{\sigma_{-1}} \leq g/(\bar{u} - \underline{u}).$$

☑

7. Conclusion

To recapitulate our conclusions, we repeat Table 1's description of the most efficient limit equilibria:

	Poisson	Diffusion	Diffusion, constant variance
Bad News	Non-trivial	Efficient	Trivial
Good News	Trivial	Non-trivial	Trivial

Notice that with the exception of the constant-variance knife-edge, the most efficient equilibrium is more efficient if there are more events in a given time period (comparing Poisson vs. diffusion on a given row) and also more extreme signals are bad news. The first comparison is suggestive of a more general result about aggregation of signals being beneficial; this is a focus of the companion paper. Notice also that our approach here relies on the assumption that both agents simultaneously observe the state and revise their chosen action. The simultaneity allows us to use the techniques of discrete-time repeated games with publicly observed signals. It would be interesting to study the harder case where agents sample and revise at random times, and do not observe the times of the other agent's revision opportunities; this would result in a game of imperfect and private monitoring.

References

- Abreu, D., P. Milgrom and D. Pearce [1991]: “Information and Timing in Repeated Partnerships,” *Econometrica*, **59**: 1713-1733.
- Abreu, D., D. Pearce, and E. Stachetti [1990] “Towards a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica* 58:1041-1064.
- Faingold, E. and Y. Sannikov [2005] “Equilibrium Degeneracy and Reputation Effects in Continuous Time Games,” mimeo.
- Fudenberg, D., D. Kreps, and E. Maskin [1990]. “Repeated Games with Long-run and Short-run Players,” *Review of Economic Studies*, 57: 555-573.
- Fudenberg, D., D. K. Levine and E. Maskin [1994]: “The Folk Theorem with Imperfect Public Information,” *Econometrica*, **62**: 997-1039.
- Fudenberg, D. and D. K. Levine [1983]: “Subgame-Perfect Equilibria of Finite- and Infinite-Horizon Games,” *Journal of Economic Theory*, **31**: 251-258
- 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 [2007] “The Aggregation of Information in High-Frequency Repeated Games,” in preparation.
- Fudenberg, D., D. K. Levine, and S. Takahashi [2006] “Perfect Public Equilibrium when Players are Patient,” forthcoming in *Games and Economic Behavior*.
- Fudenberg, D. and J. Tirole [1991] *Game Theory*, MIT Press, Cambridge, MA.
- Hellwig, M. and K. Schmidt [2002]. “Discrete-Time Approximations of the Holmström-Milgrom Brownian-Motion Model of Intertemporal Incentive Provision,” *Econometrica*, **70**: 2225-2264.
- Holmström, B. and P. Milgrom [1987] “Aggregation and Linearity in the Provision of Intertemporal Incentives.” *Econometrica* 55: 303-328.
- Levin, J. [2003] “Relational Incentive Contracts.” *American Economic Review* 93:835-857.
- Mailath, G. and L. Samuelson [2006], *Repeated Games and Reputations: Long-Run Relationships*, Oxford University Press.

- Mirrlees, J. A. [1974] "Notes on welfare economics, information and uncertainty," in M. Balch, D. McFadden and S.-Y. Wu, editors, *Essays on Economic Behavior under Uncertainty*, North-Holland, Amsterdam.
- Muller, H. [2000] "Asymptotic Efficiency in Dynamic Principal-Agent Problems," *Journal of Economic Theory* 91, 292-301.
- Sannikov, Y. [2005] "Games with Imperfectly Observed Actions in Continuous Time," mimeo.
- Sannikov, Y. and A. Skrzypacz [2005] "Impossibility of Collusion under Imperfect Monitoring with Flexible Production," mimeo.
- Sannikov, Y., and A. Skrzypacz [2006] "The Role of Information in Repeated Games with Frequent Actions," mimeo.