

Repeated Games with Frequent Signals

Drew Fudenberg and David K. Levine¹

Departments of Economics,

Harvard University and Washington University in St. Louis

First Version: February 6, 2007

This version: December 13, 2007

Abstract: We study repeated games with imperfect public information when the public signal corresponds to the aggregate of many discrete events such as sales over a small time period. The set of equilibrium payoffs in the limit as the observation period goes to 0 depends on both the probability law governing the discrete events and on the level of aggregation. It in general differs from the equilibrium payoffs of the limiting continuous-time game, which correspond to the limit equilibria when the underlying signal process is binomial.

¹ We thank Mihai Manea for careful proofreading. NSF grants SES-03-14713 and SES-04-26199 provided financial assistance.

1. Introduction

In repeated games with imperfect public information, players do not necessarily observe one another's actions; they observe a public signal whose distribution depends on those actions. This is the case, for example, in the classic papers by Green and Porter [1983] (where players observe the market price but not the outputs of other players) and Radner, Myerson and Maskin [1986] (where players observe the partnership's output but not the effort of other players). As in repeated games with observed actions, the normalized equilibrium payoffs of these games depend on the discount factor: we expect a larger set of equilibrium payoffs when players are more patient. Equilibrium payoffs also depend on how much information about players' actions is revealed by the signal: we expect a larger set of equilibrium payoffs when players have more precise information (Kandori [1992]).²

One argument for the relevance of high discount factors is that play is frequent. However, as Abreu, Milgrom, and Pearce [1991] argue, rapid play not only means that players are more patient, but are exposed to noisier information. These forces push in opposite directions, so determining how the equilibrium payoffs vary with the period length requires some analysis. We focus on a class of games where the impact of information is relatively easy to assess, namely games between one long-run player and a sequence of short-run opponents, as in the classic Klein and Leffler [1981] model of a long-run firm facing a sequence of short-run consumers. One reason for interest in these games is that even for discount factors converging to one, the set of equilibrium payoffs is larger with perfectly observed actions than if actions are observed with noise, as shown in Fudenberg and Levine [1994]. A second reason is that here it is relatively simple to determine how the information determines the set of equilibrium payoffs, so that most of our work can be on characterizing the informativeness of different signal structures as opposed to calculations that are specific to repeated games. In more general games, the set of equilibrium payoffs will depend on the information structure in more complicated

² Both of these "expectations" are correct when the set of equilibrium payoffs is convex, as it will be if players have access to a public randomization device. For more on the role of this convexity condition see Mailath and Samuelson [2007] and Yamamoto [2007].

ways, but our calculation of the “limit informativeness” of various sequences of signal structures will still apply.

We examine various ways of sending the time period of the game to zero and passing to a continuous-time limit. Our main point is that these limits all correspond to the idea that players act “very frequently,” but can have very different limit equilibria. In addition, the paper highlights the role of “information aggregation” in determining the limit equilibrium payoffs. That is, when does observing the sum of many signals lead to a larger equilibrium set than observing the signals one at a time? Finally, one of the limits we consider corresponds to the case of players observing the state of a diffusion process, as in Sannikov [2007a], Sannikov and Skrypczak [2007], Faingold and Sannikov [2007] and Fudenberg and Levine [2007a]; our results here help provide intuition for how the set of equilibrium payoffs depends on the variances of the signals under different actions.

Briefly, we consider sequences of games indexed both by how frequently individual pieces of information are generated by Nature and by how frequently players observe the information that has been generated. Our focus is on situations where the information process as seen by an outside observer converges to a controlled diffusion. That is, for any fixed observation time, when players’ earlier actions are constant, the signal will have a normal distribution, with mean and variance that will depend on the actions chosen. However, our focus is on cases where the diffusion-process description of information is not exact, but arises from aggregating a large number of small frequent events – for example the public signal might be something like “sales” or “revenues.” Here players would observe at most a single transaction in each period if they monitor the process at a sufficiently high frequency, but if the period between observations is long compared to the arrival rate, then players may observe an aggregate of many events and not the individual events themselves. Consequently their observation structure, like that of an outside observer, can be approximated by a diffusion process. As we will see, though, even in this latter case it is not always true that the set of limit equilibrium payoffs can be determined by studying the equilibria of a game with continuous-time monitoring of a controlled diffusion.

We use an extended analysis of a trinomial signal structure to show that there are many ways of passing to a limit with the same limiting diffusions yet different limit equilibria. Notice that there are two different aspects of information arrival. One is how frequently events occur, which we refer to as the “event frequency.” The other is how frequently players get to publicly observe the aggregated data, which we call the “observation frequency.” The simplest reason that we may obtain different limiting equilibria even when information process converges to the same limiting diffusion is because the event and observation frequencies may approach the limit at different rates. For example, we show in Section 6.1 that if the event process is trinomial, and the “bad” action has a higher variance, the best equilibrium payoff is bounded away from efficiency if the event and observation frequencies are equal. However, the best limit equilibria is fully efficient if the event frequency goes to infinity faster than the observation frequency, so that asymptotically the public signal aggregates infinitely many events.

Less obviously, even two discrete-time processes that converge to the same diffusions and are observed at the same frequency can have sets of discrete-time equilibria that converge to different limits. This is the case for example with the “good news trinomials” considered in Section 6.2. In this case, the limit equilibria depends on the probability γ of a “null outcome,” but this parameter has no effect on the drifts and volatilities of the limit diffusions. This is also the case for the particular “bad news trinomials” we construct at the end of section 6.1. These trinomials converge to a pair of diffusion processes with equal volatility, and Faingold and Sannikov have shown that in the corresponding continuous-time game the only equilibria are trivial. However, in our construction the two signals have slightly different variances at each point before the limit is reached, and we show that this allows a limit equilibrium that is not only non-trivial but first-best for the long-run player.

To put this into perspective, recall that the best limit equilibrium is trivial when players observe diffusions that have the same volatility. Equal volatilities of the limit diffusion is a knife-edge case when the event process is trinomial, but the case of a binomial event process is very different. Here we show that if the event processes corresponding to the two actions are binomials with the same step sizes, and both processes aggregate to a diffusion limit, then the two diffusions must have the same

volatility. As a consequence we can show in many cases that the limit equilibrium must be trivial along any sequence of binomials that converge to a pair of diffusions. However, there are sequences of trinomials with unequal variances that converge to a pair of diffusions with the same volatility, but where there is a fully efficient limit equilibrium.

Our work builds on our earlier paper Fudenberg and Levine [2007a]. That paper, like this one, considers a game between one long-run player and a sequence of short-run opponents. It provides general conditions for a sequence of discrete-time games with period length going to zero to have a non-trivial limit equilibrium (that is, a limit of discrete-time equilibria where the long-run player's payoff is strictly above that of the static Nash equilibrium) when the public information varies with the period length in a fairly arbitrary way. These conditions relate to the "asymptotic informativeness" of the public information: it must be possible to use the public information to construct trigger strategies that in the limit have both a non-zero chance of triggering punishment and a sufficiently large "signal-to-noise" ratio. When these conditions are satisfied, the best equilibrium payoff is monotone in the signal to noise ratio, and the highest possible payoff is attained when this signal-to-noise ratio is infinite.

After developing the general result, the paper applies it to the case where the public information in the discrete-time games comes from observing the state of a diffusion. In the context of this paper, the diffusion arises from a limit on the event frequency, so the diffusion case corresponds to the case in which event frequency is much higher than the observation frequency.

Using the general result from the earlier paper, we can reduce the study of the limit equilibria to the analysis of the asymptotic informativeness of the signal structure. The per-event informativeness is all that matters if players observe each event separately, yet many processes with different per-event informativeness converge to the same diffusions; this is why the equilibria of the controlled-diffusion case can be different than the limit equilibria. In the "bad news" trinomial, where deviations increase the volatility of the observed signal, the per event informativeness stays constant as the time interval goes to 0. In addition aggregating many signals leads to a better limit value of the signal to noise ratio, because the strategies can be chosen "punish" only if every outcome was

extreme. We use this to explain our earlier finding that there are fully efficient limit equilibria in the case of bad-news diffusions. In contrast, in the “good news” case where the deviating actions lead to a lower volatility, the optimal strategy is to punish on “intermediate” signals; this is why the limit equilibrium payoffs are lower.

To understand the intuition, consider a trinomial distribution where the underlying signal can take on three values -1 , 0 , or $+1$. If an aggregate of two signals is observed, an “intermediate” aggregate outcome of, say, 0 could either come from two occurrences of 0 , a $+1$ followed by a -1 , or the reverse. Thus observing an aggregate of 0 is less informative than knowing that both outcomes were 0 ; whether it is more informative than a single 0 can depend on some of the fine detail of the model. The upshot is that the effect of aggregation in the good news case is ambiguous; this helps explain why the good news case has a relatively smaller set of limit equilibrium payoffs.

Like the earlier paper, this one is related to Abreu, Milgrom, and Pearce [1991], who studied the strongly symmetric pure-strategy equilibria of a repeated partnership game in discrete time when players observe the realization of a Poisson process. Our work is also similar to Sannikov and Skrypacz [2007]. They consider a game with two long-run players observing the infinite-dimensional sample path of a continuous-time process at discrete intervals. They provide an algorithm for computing an upper bound on the set of payoffs to pure-strategy equilibria as the time period shrinks, and an algorithm that computes the equilibrium payoffs in the iterated limit where first the time period shrinks and then the interest rate goes to 0 ; they also show that strategies that obtain extremal payoffs use the diffusion and Poisson parts of the public signal very differently.. Instead of considering games with two long-run players, we study a game with a single long-run player facing a sequence of short-run opponents, each of whom plays only once but knows about past outcomes. Unlike Sannikov and Skrypacz, we allow mixed as well as pure strategies.

Finally, our work is related to papers that construct a series of discrete-time games whose limit equilibria correspond to the equilibria of the continuous-time game with diffusion signals, as in Hellwig and Schmidt [2002] and Sannikov [2007b]. The main difference is in focus: The earlier papers are in the spirit of a lower-hemicontinuity argument, showing that there exists a sequence of discrete-time games that provide a

foundation for the limit game. Our paper is in the spirit of upper hemi-continuity, and asks for which sequences of discrete-time game is the usual continuous-time game the corresponding limit object? Our results show (modulo some technical restrictions) that this sort of upper hemi-continuity holds when the signals are an aggregate of binomials, but in general fails otherwise. This failure results not any of the relevant functions becoming discontinuous, but rather because information is lost in the limit, so that the strategy space in the standard continuous-time models is too small to incorporate all of the relevant limit objects. This suggests that upper hemi-continuity could be restored by enlarging the strategy space, as in Fudenberg and Tirole [1985] and Simon and Stinchcombe [1989], but that remains a topic for future work.

2. The Model

A long-run player 1 plays a stage game with a short-run player 2 who is completely impatient. To focus attention on the information-theoretic aspects of the problem we restrict attention to the following 2x2 stage game

		Player 2	
		Out	In
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 Out 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 In but he can only induce player 2 to play In by avoiding playing -1 . The highest feasible payoff for player 1 is $\bar{u} + g$. The Stackelberg payoff of $\bar{u} + g/2$ can be obtained by a publicly observed commitment to play the mixed strategy $(1/2, 1/2)$ but the highest repeated game payoff is \bar{u} when actions are observed (Fudenberg, Kreps and Maskin [1990]) and the highest payoff with imperfect public monitoring is strictly less than that (Fudenberg and Levine [1994]). Our focus will be on determining when the repeated game with vanishingly small time periods has equilibria with normalized discounted payoffs that exceed \underline{u} , and

when it has equilibria with payoff approaching \bar{u} , which we refer to as the “first-best” payoff.

When the game is repeated, the length of a period is τ , and the subjective continuous time interest rate for the long-run player is r , so that her rate of time discount is $\delta = e^{-r\tau}$. Each period, the stage game is played, and then the long-run player and subsequent short-run players observe a public signal $z \in \mathbb{R}$ that depends only on the action a_1 of the long-run player. The public signal has finite support; its distribution is described by the density function $f(z | a_1, \tau)$. In addition, we assume that the support of the signal is independent of the action played, so that every possible signal has positive probability under every action.

There is also a publicly observed public randomization device each period before actions are taken. The public history is the history of the signal and the public randomization device.³ Our solution concept is *perfect public equilibrium* 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.⁴

The characterization of perfect public equilibria in this setting is straightforward, using standard dynamic programming techniques in the spirit of Abreu, Pearce and Stachetti [1990]. Because we allow public randomization, the set of perfect public equilibrium payoffs to LR is a line segment between a best and worst equilibrium; because the static Nash equilibrium involves no entry and gives LR her minmax, the worst equilibrium is for LR to get \underline{u} . So the set of PPE payoffs to the LR player is completely described its upper bound, which we denote by v^* .⁵ Proposition 1 in the Appendix shows that v^* can be computed as the solution to a static linear programming problem, where the control variables are the “continuation payoffs” $w(z)$ that the player

³ 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.

⁴ 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.

⁵ The arguments of Fudenberg and Levine [1983] or Abreu, Pearce and Stachetti [1990] can be adapted to show that the set of PPE payoffs in this game is compact, so the best PPE payoff is well-defined

expects to receive following each signal z ; this result is used in the proof of our next proposition.

Now suppose that the continuation payoffs are restricted to the two values v^* (“reward”) and \underline{u} (“punishment”). Define p as the probability of punishment when the action chosen is **+1** (that is, p is the probability under action **+1** of signals such that continuation play is “punishment”) and define q as the probability of the punishment outcome when the action chosen is **-1**. We say that a pair (p, q) is *feasible* if it can be generated by some specification of the function w .

Proposition 2: (Fudenberg and Levine [2007a])

(a) *For a fixed discount factor δ , there is an equilibrium with the long-run player’s payoff above \underline{u} if and only if there are feasible $p, q \in [0, 1]$ that satisfy*

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

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

$$\bar{u} - \frac{pg}{q - p}. \quad (**)$$

When () is not satisfied, the highest PPE payoff is \underline{u} .*

(b) *There is a PPE that supports the highest PPE payoff that has the “cutoff likelihood property:” There is a cutoff λ^* such that if $f(z | a_1 = -1, \tau) / f(z | a_1 = +1, \tau) > \lambda^*$ then $w(z) = \underline{u}$, if $f(z | a_1 = -1, \tau) / f(z | a_1 = +1, \tau) < \lambda^*$ then $w(z) = v^*$.⁶*

Note that the best equilibrium v^* is close to the first best if there are feasible (p, q) with $p / (q - p)$ small.

3. Continuous-Time Limits

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. We consider then families of

⁶ Note that when the likelihood ratio is exactly λ^* the continuation value may lie anywhere in the interval $[\underline{u}, \bar{u}]$.

games indexed by the period length τ . We must now describe how the signal z varies with the period length τ . Our basic scenario is that z is an aggregate of discrete random variables representing, for example, sales, prices, or other transaction data.⁷ Specifically, we suppose that z is the sum of some number of “events,” by which we mean independent identically distributed random variables Z_j whose support is a fixed finite set, regardless of the action profile.

Recall that the length of a period, that is, the time between moves, is τ ; the “observation frequency” we mentioned in the introduction is thus $1/\tau$. We assume that the length of time between events (that is between realizations of the Z_j) is $\Delta \leq \tau$, so that the *event frequency* is $1/\Delta$. For simplicity we assume that τ is an integer multiple of Δ , so that an integer number $k = \tau/\Delta$ of events occur in a single period. Notice that the distribution of Z_j and z depends in general on the action taken by the long-run player.

We are interested in the case in which $\tau \rightarrow 0$ (implying that $\Delta \rightarrow 0$ as well). It is convenient to assume that τ is a specified continuous strictly increasing function of Δ with $\tau(0) = 0$. In general, we allow the distribution of Z_j and its support to depend on Δ ; recall that this is necessary for the distribution of the aggregate z to approach a diffusion. However we will assume that the cardinality of the support of Z_j is a constant, independent of Δ .

The information available at the end of the period beginning at t is the signal $z = \sum_{j=t/\Delta}^{(t+\tau)/\Delta} Z_j$. Our goal is to characterize the set of equilibrium payoffs in the limit. Specifically, if for a given interest rate r there are positive $\bar{\tau}$ and ε such that for all non-negative smaller values $0 < \tau < \bar{\tau}$ the game with period length τ and interest rate r has an equilibrium with payoff at least $\underline{u} + \varepsilon$, we say that there is a *non-trivial limit equilibrium for r* . If there is any positive interest rate r for which there is a non-trivial limit equilibrium, we say simply that there is a non-trivial limit equilibrium. If for all $r > 0$ and all sequences $\tau \rightarrow 0$ the equilibrium payoff converges to \underline{u} we say there is only a trivial limit. (In principle there can be cases where the limit depends on the sequence chosen, however, we do not provide names for these cases.) If there is an

⁷ This model does not capture the case where the signal involves an occasional catastrophic event, such as a failed surgery, bad reaction to a drug, or airplane crash. That type of signal is better modeled in continuous time as a Poisson process. See Celentani, Levine and Martinelli [2007].

$\bar{r} > 0$ such that for all $0 < r < \bar{r}$, all $\varepsilon > 0$, and all sequences $\tau \rightarrow 0$, there is a sequence of equilibria with payoff converging to $\bar{u} - \varepsilon$ we say there is an *efficient limit equilibrium*. If for all $(\tau, r) \rightarrow (0, 0)$ there are equilibria that have payoffs converging to \bar{u} , we say that there is an *efficient patient equilibrium*.

Note that the definition of a non-trivial limit equilibrium allows the interest rate r to be arbitrarily small, but it requires that the payoff in question to be supportable as an equilibrium when that interest rate is held fixed as the period length τ goes to 0. The definition of an efficient patient equilibrium allows 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 2 contains key information about when a sequence of equilibria has a limit that is non-trivial or efficient. The following corollaries all apply to sequences of equilibria for the games indexed by observation period τ . First, for each fixed τ we define $\rho(\tau) = (q(\tau) - p(\tau))/p(\tau)$ which we may view as the signal to noise ratio for the specified equilibrium. From (***) we see that if $v^*(\tau) > \underline{u}$, then it must be that

$$\rho(\tau) > \frac{g}{\bar{u} - \underline{u}};$$

We also see that in order for the payoffs to converge to \bar{u} it must be that $\lim_{\tau \rightarrow 0} \rho(\tau) \rightarrow \infty$; it will be helpful to remember that $\lim_{\tau \rightarrow 0} \rho(\tau) \rightarrow \infty$ implies $p(\tau) \rightarrow 0$.

Corollary 3:⁸

(a) *If for some $(r, \tau) \rightarrow (0, 0)$ there is no sequence of equilibria with $\rho(\tau) \rightarrow \infty$ then there is not an efficient patient equilibrium.*

(b) *If for all $r > 0$, all $\tau \rightarrow 0$, and all equilibria, $\rho(\tau) \rightarrow 0$, then there is only a trivial limit equilibrium.*

⁸ Our earlier paper states a result with the same conclusion under the additional hypothesis that the sequence of equilibria is “regular,” meaning that $\rho(\tau)$ and $(q(\tau) - p(\tau))/\tau$ both converge.

The condition

$$\rho(\tau) > \frac{g}{\bar{u} - \underline{u}}$$

implies that the LHS of (*) is positive, so the condition is sufficient for $\lim_{\tau \rightarrow 0} v^*(\tau) > \underline{u}$ in cases where the RHS of (*) converges to 0, which is true in particular when $p(\tau)$ is bounded away from 0.

Corollary 4: *If for all $r > 0$ and all $\tau \rightarrow 0$ there is a sequence of equilibria with*

$$\rho(\tau) > \frac{g}{\bar{u} - \underline{u}}$$

and $p(\tau)$ bounded away from 0, there is a non-trivial limit equilibrium for any r .

In many cases of interest, the best equilibria will have $p(\tau)$ converging to 0. To address this possibility, note that we can rewrite (*) as

$$\frac{(\bar{u} - \underline{u})}{g} \rho - 1 \geq \frac{r\nu(\tau)}{p/\tau}$$

where $\nu(\tau) = (e^{r\tau} - 1)/r\tau$ converges to 1 as $\tau \rightarrow 0$. Because $\rho = (q - p)/p$, $p = q/(\rho + 1)$, so (*) is equivalent to

$$\frac{q}{\tau} \geq \frac{\rho + 1}{\frac{(\bar{u} - \underline{u})}{g} \rho - 1} rg(\tau).$$

Note that the RHS of this inequality is bounded below by

$$\frac{rg}{\bar{u} - \underline{u}},$$

to which it converges as $\rho \rightarrow \infty$. This yields three further conclusions:

Corollary 5:

(a) *If for all $r > 0$ and all $\tau \rightarrow 0$, along any sequence of best equilibria $\rho(\tau) > g/(\bar{u} - \underline{u})$ implies $q/\tau \rightarrow 0$, then there is only be a trivial limit.⁹*

⁹ The best equilibrium payoff exists for each τ but there may be multiple equilibria with this payoff.

(b) if for every $\theta > 0$ and every sequence $(r, \tau) \rightarrow (0, 0)$ there is a sequence of equilibria with $q(\tau)/\tau \geq \theta$ and $\rho(\tau) \rightarrow \infty$, then there is an efficient patient equilibrium

(c) If for every $\lambda > 1$ and every sequence $\tau \rightarrow 0$ there is a sequence of equilibria with $p(\tau)$ constant and $\lim_{\tau \rightarrow 0} q(\tau)/p(\tau) = \lambda$ then there is an efficient limit equilibrium.

Proof: (a) and (b) are immediate. For (c), note that when λ is sufficiently large, the LHS of (*) is positive and bounded away from 0; the RHS converges to 0, so using the strategies that generate these probabilities yields a non trivial limit equilibrium, and the payoff to this equilibrium converges to \bar{u} as $\lambda \rightarrow \infty$. \square

4. Converging to Diffusions: General Results

We now restrict attention to information processes that converge to diffusion processes in the limit, because we want to relate this limit to the predictions of continuous-time controlled-diffusion models. The idea is that the diffusion process reflects the aggregation of information, with the limiting normal distribution arising from central limit theory.

Our basic diffusion hypothesis is that for each fixed action $i = +1, -1$ of the long-run player the sum $z = \sum_{j=1}^{\lfloor t/\Delta \rfloor} Z_j$ converges to a diffusion as $\Delta \rightarrow 0$. That is, in the limit $\Delta \rightarrow 0$, when the long-run player's action is held fixed from time 0 to any time t , the value of the observed signal at time t is a normally distributed random variable with mean $\mu_i t$ and standard deviation $\sigma_i^2 t$. We continue to assume that the support of the z 's is independent of the action chosen, so that when $\tau = \Delta$ and players observe each individual realization of Z , no outcome perfectly reveals which of the two actions was played.

As Δ and τ converge to 0, the distribution of the Z 's will change; we let f^Δ denote the density of the Z 's, and f^τ the density of the aggregate z .

The simplest case is where $\Delta = \tau$: Here for each Δ the maximum possible value of q/p is obtained by punishing only on the signal or signals that maximize $f^\Delta(Z | -1)/f^\Delta(Z | +1)$. Since this ratio is finite (because the support of Z is independent of the action played) there cannot be an efficient equilibrium for any fixed value of Δ , and we can use the asymptotic value of this ratio to determine when there are efficient or non-trivial equilibria. We can also use the asymptotic behavior of this ratio to

derive upper bounds on the equilibrium payoffs when $\tau = k\Delta$ for some fixed “level of aggregation” k . Define $M(k, \Delta) = \max_Z \left(f^\Delta(Z | -1) / f^\Delta(Z | +1) \right)^k$.

Proposition 6: *Suppose that $\lim_{\Delta \rightarrow 0} \tau(\Delta) / \Delta = k < \infty$. Then*

(a) *If $\limsup_{\Delta \rightarrow 0} M(k, \Delta) < \infty$ there is no efficient patient equilibrium.*

(b) *If $\limsup_{\Delta \rightarrow 0} M(k, \Delta) = 1$ there is only a trivial limit equilibrium.*

Proof: The best likelihood ratio for any test based on observing the sum of k realizations of Z is no higher than the best likelihood ratio when observing each draw independently. Because the draws in each period are i.i.d the highest feasible likelihood ratio when observing each of the draws is

$$M(k, \Delta) = \max_Z \left(f^\Delta(Z | -1) / f^\Delta(Z | +1) \right)^k.$$

In case (a), along any sequence of games and strategies we have $\limsup_{\Delta \rightarrow 0} \rho(\Delta) < \infty$, so by corollary 3a there is no efficient patient equilibrium; in case b we have $\limsup_{\Delta \rightarrow 0} \rho(\Delta) = 0$ and the conclusion follows from corollary 3b. ☑

It turns out that there are three cases of interest depending on the variances of the limit diffusions. Recall that σ_{+1}^2 is the variance of the limit diffusion when the long-run player is friendly, and that σ_{-1}^2 is the variance when he is unfriendly. If $\sigma_{-1} / \sigma_{+1} > 1$ then a large draw of z is “bad news” and means that the long-run player should be punished. So we refer to this as the “bad news” case. Conversely if $\sigma_{-1} / \sigma_{+1} < 1$ a large draw of z is “good news.” Finally there is the case of equal variances $\sigma_{+1} = \sigma_{-1}$.

Proposition 7: *In the bad news case ($\sigma_{-1} / \sigma_{+1} > 1$) if $\lim_{\Delta \rightarrow 0} \tau(\Delta) / \Delta = \infty$ there is an efficient limit equilibrium.*

Proof: Our previous paper showed that there is an efficient limit equilibrium when players observe the state of a bad-news diffusion process; the proof uses that fact and a continuity argument to construct a sequence of equilibria satisfying the conditions of Corollary 5c. See Appendix 2 for details. ☑

5. Binomial Arrays Converging to Diffusions

Just as the central limit theorem applies to the sum of any i.i.d random variables processes with the appropriate bounds on moments, many triangular arrays converge to diffusions. We begin with the simplest and most obvious case, in which the individual events Z_j follow a binomial distribution. As we will see, this case is quite special, as it appears to be the one case in which the limit equilibrium is necessarily trivial, so readers already familiar with the mechanics of the binomial arrays may want to skip this section on first reading and go on to Section 6.

Section 5.1 considers the usual binomial construction of a diffusion, where the “step size” (that is, the magnitude of the discrete-time individual events Z_j) is equal to the volatility times $\Delta^{1/2}$. With this construction, two binomial arrays with the same step sizes converge to diffusions with the same volatility. Since our earlier results on the case where Z_j is drawn from a diffusion shows that in the equal variance case there is only a trivial limit equilibrium, we might expect that this when the signals correspond to the standard textbook binomial array that converges to a diffusion. This is easiest to show when $\lim_{\Delta \rightarrow 0} k(\Delta)\tau(\Delta) = 0$, as assumed by Proposition 8, whose short proof is included to provide intuition for our more general results.

Section 5.2 considers general pairs of binomial arrays with the same support. We show first in Proposition 9 if these arrays converge to non-degenerate diffusions, the diffusions must have the same volatility. We use this fact in Proposition 10 to show that there is only a trivial limit equilibrium if $\lim_{\Delta \rightarrow 0} \tau(\Delta) \exp(k(\Delta)^{2/7}) \rightarrow \infty$. The proof of Proposition 10 is fairly involved, as it involves results from the theory of large deviations, so we have relegated it to the appendices.

5.1. The Standard Binomial Construction

Consider then a diffusion with drift μ_1 and volatility σ_1^2 . The standard binomial construction¹⁰ sets the “step size” of Z to $h = \sigma_1 \Delta^{1/2}$ and the support to $\pm h$, with $\Pr(Z = h) = .5 + .5\mu_1 \Delta^{1/2} \sigma_1^{-1}$. Then the mean of Z is $\mu_1 \Delta$, and the mean of

¹⁰ As in Stokey [2007], for example.

$\lim_{\Delta \rightarrow 0} E \sum_{j=1}^{\lfloor t/\Delta \rfloor} Z_j$ is $\mu_1 t$. Similarly, we can compute that the variance of Z is $\sigma_1^2 \Delta - (\mu_1 \Delta)^2$, and $\lim_{\Delta \rightarrow 0} \text{var} \left(\sum_{j=1}^{\lfloor t/\Delta \rfloor} Z_j \right) = \sigma_1^2 t$. Since this construction sets the step size h to be proportional to the volatility of the target diffusion, it forces any two binomials with the same step size to have the same variance. Thus if we consider a second binomial to correspond to action **-1** with the same support, we must have $\sigma_{-1} = \sigma_1 = h\Delta^{-1/2}$.

Now consider the implications of this for the limit equilibria of our two-player game. First, suppose there is no information aggregation at all, so that $\tau = \Delta$. We assume that the bad action has a higher mean, so the best equilibrium punishes with some positive probability π when outcome $+h$ occurs. In other words, $q = \pi \left(.5 + .5\mu_{-1}\Delta^{1/2}\sigma_1^{-1} \right)$ and $p = p \left(.5 + .5\mu_1\Delta^{1/2}\sigma_1^{-1} \right)$, so

$$\rho = \frac{q-p}{p} \rightarrow \frac{(\mu_{-1} - \mu_{+1})\Delta^{1/2}\sigma_1^{-1}}{1 + \mu_{+1}\Delta^{1/2}\sigma_1^{-1}} = 0.$$

Hence, while there may be a non-trivial equilibrium for some finite period lengths (depending on the interest rate and the informativeness of the signal as compared to the short-run gain from deviating), Corollary 3b shows that there can only be a trivial equilibrium in the limit, because the informativeness of the underlying signal converges to 0.

Next, suppose that the agents aggregate a fixed number of signals, that is, $k = \tau/\Delta$ is constant. Now the most informative test (the one with the highest value of $\rho = (q-p)/p$) is to punish only if the sum equals kh so that every realization was $+h$. With this test, the probabilities of punishment under **-1** and **+1** are $q = \phi \left(.5 + .5\mu_{-1}\Delta^{1/2}\sigma_{+1}^{-1} \right)^k$ and $p = \phi \left(.5 + .5\mu_{+1}\Delta^{1/2}\sigma_{+1}^{-1} \right)^k$. This aggregation may allow for better equilibria for some values of Δ , but it is still true that $\rho \rightarrow 0$ as $\Delta \rightarrow 0$, so Proposition 6 implies that the aggregation does not lead to a better limit equilibrium.

Of course, the case where agents aggregate a finite number of observations is not quite the same as when they observe the diffusion limit, as observing a diffusion corresponds to aggregating infinitely many signals, so that $k = \tau/\Delta$ tends to infinity. The next result allows for this to happen but requires that k not grow too quickly:

Proposition 8: *If $\lim_{\Delta \rightarrow 0} k(\Delta)\tau(\Delta) = 0$ then in the standard binomial construction of diffusions there are only trivial limit equilibria.*

Proof: We compute

$$\rho = \frac{q - p}{p} = \frac{q}{p} - 1 = \left(\frac{.5 + .5\mu_{-1}(\tau/k)^{1/2}\sigma_{+1}^{-1}}{.5 + .5\mu_{+1}(\tau/k)^{1/2}\sigma_{+1}^{-1}} \right)^k - 1.$$

This goes to zero if the log of the first term goes to zero. We calculate

$$\begin{aligned} & \lim_{\Delta \rightarrow 0} k(\Delta) \log \left(.5 + .5\mu_{-1} [\tau(\Delta)/k(\Delta)]^{1/2} \sigma_{+1}^{-1} \right) - k(\Delta) \log \left(.5 + .5\mu_{+1} [\tau(\Delta)/k(\Delta)]^{1/2} \sigma_{+1}^{-1} \right) = \\ & \lim_{\Delta \rightarrow 0} k(\Delta) \left(\mu_{-1} [\tau(\Delta)/k(\Delta)]^{1/2} \sigma_{+1}^{-1} - \mu_{+1} [\tau(\Delta)/k(\Delta)]^{1/2} \sigma_{+1}^{-1} \right) = \\ & \lim_{\Delta \rightarrow 0} (k(\Delta)\tau(\Delta))^{1/2} \sigma_{+1}^{-1} (\mu_{-1} - \mu_{+1}) = 0 \end{aligned}$$

where the last equality follows from $\lim_{\Delta \rightarrow 0} k(\Delta)\tau(\Delta) = 0$. Thus by corollary 4b, there is only a trivial limit equilibrium. \checkmark

Sequences of games with $\lim_{\Delta \rightarrow 0} k(\Delta)\tau(\Delta) > 0$ are more complicated, as we can obtain a limiting value of ρ that is non-zero by considering strategies of the form “only punish if every outcome was $+h$ ”. However along this sequence we have $q/\tau \rightarrow 0$, so as Corollary 5a shows this is no help. This is the limit-of-discrete signals version of the Mirrlees result that the tail of the normal has arbitrarily high likelihood. Of course this is only one possible sequence of strategies; to show that there are no non-trivial limit equilibria regardless of the values of the payoff parameters, one must show that every method of constructing strategies fails the test in Corollary 6a. We did this in Fudenberg and Levine [2007a] Proposition 2 for the case where observations in each period are normally distributed with the same variance; from the central limit theorem, this corresponds to passing to the limit $\Delta \rightarrow 0$ for fixed τ and only then sending τ to 0, so that τ/Δ is effectively infinite. The next subsection has a result that applies when $k = \tau/\Delta$ goes to infinity sufficiently quickly; its hypothesis overlaps with the condition $\lim_{\Delta \rightarrow 0} k(\Delta)\tau(\Delta) = 0$, so that combined the results cover all cases where k goes to infinity.¹¹

¹¹ The case where k is bounded is covered by Proposition 6.

5.2 General Binomial Arrays

So far we have only discussed the standard construction of a diffusion from the sum of binomials, but one can also consider more general constructions where the step sizes and probabilities of the binomials depend on Δ in other, more complicated ways.¹²

Proposition 9: *Suppose the period length is Δ , and that we have i.i.d. binomials $Z_i(\Delta)$ where the common outcomes are $x(\Delta) > y(\Delta)$ and the probability of $x(\Delta)$ under action i is $\alpha_i(\Delta)$, with $\lim_{\Delta \rightarrow 0} \alpha_i(\Delta) = \alpha_i, 0 < \alpha_i < 1$. If the sums $\sum_{j=1}^{\lfloor t/\Delta \rfloor} Z_j$ converge to a diffusion with drift μ_i and volatilities σ_i^2 then $\sigma_1 = \sigma_2$.*

Proof: In Appendix 3. ☑

The requirement that the binomials have the same step sizes and both converge to diffusions has a strong implication:

Proposition 10: *Suppose*

(i) $\lim_{\Delta \rightarrow 0} \tau(\Delta) \exp(k(\Delta)^{2/7}) \rightarrow \infty$

(ii) *the signals are i.i.d. binomials $Z(\Delta)$ where the common outcomes are $x(\Delta) > y(\Delta)$, and the probability of $x(\Delta)$ under action $i = +1, -1$ is $\alpha_i(\Delta)$ with $\lim_{\Delta \rightarrow 0} \alpha_i(\Delta) = \alpha_i, 0 < \alpha_i < 1$*

(iii) $\sum_{j=1}^{\lfloor t/\Delta \rfloor} Z_j$ *converge to non-degenerate diffusions with means μ_i and variances σ_i^2 .*

Then all limit equilibria are trivial.

Proof: in Appendix 4. ☑

Note that if $k(\Delta) \rightarrow \infty$ and $\lim_{\Delta \rightarrow 0} k(\Delta)\tau(\Delta) > 0$, then $\lim_{\Delta \rightarrow 0} \tau(\Delta) \exp(k(\Delta)^{2/7}) = \infty$, so condition (i) covers the cases ruled out by the hypotheses of Proposition 8.

¹² As an example of a “non standard binomial array,” the appendix uses the case of binomials with limit probabilities converging to (1/3, 2/3) instead of (1/2, 1/2).

The idea is to bootstrap off a result from our previous work Fudenberg and Levine [2007a] which showed this was true if the underlying process was a diffusion rather than a binomial. The rough idea is to use the central limit theorem to show that the binomial partial sums are “close enough” to a normal that only a trivial equilibrium can be sustained. If we could restrict the analysis to strategies where the cutoff for punishment was fixed relative to the standard error of the signal this strategy would work, but unfortunately we have to consider also punishment cutoffs that become large relative to the standard error. This requires us to use a “large deviations” argument that extends an argument from Feller [1972] from the sum of i.i.d random variable to the case of triangular arrays.

To understand why this result is needed, recall that the usual central limit theorem concludes that the probability $F^n(x)$ that the normalized sum of n draws is below any fixed x converges to the probability $\Phi(x)$ that a standard normal variable is below x . Feller’s large deviations argument extends this to give conditions under which

$$\lim_{n \rightarrow \infty} \frac{1 - F^n(x_n)}{1 - \Phi(x_n)} = 1$$

when x_n is not fixed but rather $\lim_{n \rightarrow \infty} x_n = \infty$. Feller’s result does not directly apply because in the limit we consider the distribution of the underlying variables changes with the period length. We report the necessary extension in Fudenberg and Levine [2007b].

6. Trinomial Informational Limits

While some data, such as sale or no sale, may have a binomial distribution, other data, such as the number of units sold, or their price, will generally have more than two values. The simplest case beyond the binomial is the trinomial: we shall see that the trinomial breaks the link between the volatilities under the two different actions, so the equal variance/degenerate limit case seems to be the exception rather than the rule. Moreover, in these more general limits, the equilibria of the game with the limiting diffusion does not capture well the limit of the equilibria when the signal is the aggregate of many small events.

Fix a pair of drifts μ_1, μ_{-1} and a pair of volatilities $\sigma_1^2, \sigma_{-1}^2$. We will construct a particular family of pairs of trinomials such that each trinomial converges to a diffusion with the corresponding drift and volatility.

We shall use this family to explore in greater detail various ways of passing to the continuous time limit. We focus on three simple cases: a “bad news” case where the drifts are equal and deviating increases the volatility; a “good news” case with equal drifts where deviating decreases the volatility, and the case of equal volatilities and unequal drifts.

The pairs of diffusions will be indexed by a free parameter γ that is not determined by the limit diffusions; this will let us illustrate our point that different processes that converge to the same pair of diffusions can have different limit equilibria. For any $\gamma \geq 1$, we set $\bar{\gamma} = \gamma \max(\sigma_1^2, \sigma_{-1}^2)$; we will now consider a pair of trinomial distributions with the same three possible outcomes, $x = -h(\Delta), 0, h(\Delta)$, where Δ is the period length and $h(\Delta) = \bar{\gamma}^{1/2} \Delta^{1/2}$. The probability distributions on the outcomes depends on action $i = +1, -1$, and γ as follows: The probability of outcome 0 is $\alpha_i = (\bar{\gamma} - \sigma_i^2) / \bar{\gamma}$, independent of Δ (note that this is always non negative and less than 1), and the probability of outcome $+h$ is

$$\beta_i(\Delta) = \frac{1 - \alpha_i}{2} + \frac{\mu_i \Delta^{1/2}}{2\bar{\gamma}^{1/2}}.$$

As we shall see, in both the bad news and good news cases, the per-event informativeness of the individual events is constant as $\Delta \rightarrow 0$. In the bad news case the informativeness of the best test, and thus the best limit equilibrium payoff, is independent of the parameter γ . However, in the good news case γ determines the informativeness of the best test and also the best limit equilibrium payoff.

The good and bad news cases also differ in their aggregation properties: In the bad news case, aggregating more signals leads to a more informative test; so that when $k = \tau / \Delta \rightarrow \infty$ the best equilibrium approaches full efficiency; which is the result when players observe a diffusion. In the good news case, aggregating more signals can lead to a less informative test, and the effect of aggregation is ambiguous, and depends on the “free” parameter γ .

To analyze the trinomial example, we begin by computing the means and variances. We let E_i denote the expectation conditional on a_i . Then the expected values of the process are

$$\begin{aligned} E_i Z &= \beta_i h + \alpha_i 0 + (1 - \alpha_i - \beta_i)(-h) \\ &= (2\beta_i - (1 - \alpha_i))h \\ &= \mu_i \frac{\Delta^{1/2}}{\bar{\gamma}^{1/2}} h = \mu_i \frac{\Delta^{1/2}}{\bar{\gamma}^{1/2}} \bar{\gamma}^{1/2} \Delta^{1/2} = \mu_i \Delta \end{aligned}$$

and the variances are

$$\begin{aligned} E_i(z^2) - (E_i Z)^2 &= (1 - \alpha_i)h^2 - \mu_i^2 \Delta^2 \\ &= (1 - \alpha_i)\bar{\gamma}\Delta - \mu_i^2 \Delta^2 = \\ &= \sigma_i^2 \Delta - \mu_i^2 \Delta^2 \end{aligned}$$

Thus if we hold fixed the actions up to a real time t , the sum of the process up to time t has mean $\mu_i t$ and variance $\sigma_i^2 t - \mu_i^2 t \Delta$, which converges to $\sigma_i^2 t$ as Δ goes to 0. Moreover, if we look at the sum up to time $\tau(\Delta)$, where $\tau(\Delta)/\Delta$ goes to infinity, we again have a triangular array, so the position at $\tau(\Delta)$ is again described approximately by a normal.

6.1: Bad news case $\sigma_{-1}^2 > \sigma_{+1}^2$ and Zero Means.

In the bad news case, we can show that if the ratio of volatilities is sufficiently favorable then the limit equilibrium is non-trivial, regardless of the amount of information aggregation. We also show that if the amount of aggregation, as measured by the ratio $k = \tau / \Delta$, goes to infinity, then the first best can be approximated arbitrarily closely, so there is an efficient patient equilibrium. Of course; full efficiency is not possible with a finite amount of information aggregation. This shows that the limit equilibria are not determined by the assumptions that the limit distribution of the signals is a fixed pair of diffusions and that the τ and Δ both go to zero. Finally, by allowing the variance of the trinomials to converge to a common limit as τ and Δ go to zero, we can construct a sequence of games with an efficient limit equilibrium even though the

limit information structure – a diffusion with common volatilities – has only a trivial equilibrium.

To begin consider $\tau(\Delta) = \Delta$. Since the bad action has a higher volatility, and the two actions both have zero means, the best equilibrium payoff for period length $\tau = \Delta$ can be attained with a strategy that punishes with some positive probability $\pi(\Delta)$ following the signals $+h$ and $-h$ and punishes with probability 0 when the signal is 0. (The likelihood ratio for punishing on 0 is less than 1, and the symmetry of the problem means that treating $+h$ and $-h$ symmetrically is one of the solutions to the linear programming problem that defines the optimum.). Such strategies have signal to noise ratio

$$\rho = \frac{q}{p} - 1 = \frac{1 - \alpha_{-1}}{1 - \alpha_{+1}} - 1 = \frac{\bar{\gamma}\sigma_{-1}^2}{\bar{\gamma}\sigma_{+1}^2} - 1 = \frac{\sigma_{-1}^2 - \sigma_{+1}^2}{\sigma_{+1}^2},$$

independent of $\bar{\gamma}$, γ , Δ , and $\pi(\Delta)$. If π is a constant independent of Δ then $p = \pi(1 - \alpha_{+1})$ is independent of Δ as well. Hence by Corollary 4 these strategies support a non-trivial limit equilibrium for interest rate r if the ratio, $\sigma_{-1}^2 / \sigma_{+1}^2$ is sufficiently large. Moreover the simple form of the observation structure here lets us compute the best limit equilibrium payoff: Since no choice of cutoff can yield a higher value of ρ than

$$\frac{\sigma_{-1}^2 - \sigma_{+1}^2}{\sigma_{+1}^2},$$

the best limit equilibrium payoff is

$$\bar{u} - \frac{\sigma_{+1}^2}{\sigma_{-1}^2 - \sigma_{+1}^2} g.$$

Now consider $\tau = 2\Delta$. Here the most informative test is to punish only if the outcome is $+2h$ or $-2h$, this has

$$\rho = \frac{((1 - \alpha_{-1})/2)^2}{((1 - \alpha_{+1})/2)^2} - 1 = \frac{(1 - \alpha_{-1})^2}{(1 - \alpha_{+1})^2} - 1 = \left(\frac{\sigma_{-1}^2}{\sigma_{+1}^2} \right)^2 - 1$$

independent of Δ , and since the punishment probability is also independent of Δ we again have a non-trivial equilibrium. Moreover, because the maximal value of ρ

(consistent with non-zero punishment probability) has increased, aggregating two signals allows a better limit equilibrium payoff.

Now consider the case $\tau = \Delta^{1/2}$ so that $k(\Delta) = \Delta^{-1/2} \rightarrow \infty$ as $\Delta \rightarrow 0$. Since $\sigma_{-1}^2 > \sigma_{+1}^2$ and the signals converge to a diffusion, the hypotheses of Proposition 7 apply, and we conclude that there is an efficient limit equilibrium. More strikingly consider a sequence $\{\sigma_{-1,n}^2 / \sigma_{+1,n}^2\} \downarrow_n 1$ and to each n associate a sequence of trinomial observation structures with the specified pair of variances and $\Delta \rightarrow 0$. Then for any sequence of strictly positive $\varepsilon_n \rightarrow 0$ we know that for each n , there is a sequence of equilibria with limit payoff $\bar{u} - \varepsilon_n$ as $\Delta \rightarrow 0$, that is for each n there is a sequence of games $G_{n,\Delta}$ and associated PPE $\sigma_{n,\Delta}$ with payoff $u_{n,\Delta}$ converging to $\bar{u} - \varepsilon_n$. We can now use a diagonalization argument to conclude that there are sequences of trinomials that converge to diffusions with equal volatility but nevertheless support efficient limits. This shows that conclusions based on the hypothesis that the variances are equal in the limit do not apply to the limit of the equilibria along the sequence without additional information, such as the rate at which the variances become equal.

6.2. Good News Case $\sigma_{-1}^2 < \sigma_{+1}^2$ and Zero Means

Once again, we begin with the case $\tau(\Delta) = \Delta$. The optimal equilibrium with this signal structure prescribes punishment with positive probability when $Z = 0$ and 0 probability of punishment on $+h, -h$, so the signal to noise ratio is

$$\rho = \frac{q - p}{p} = \frac{\gamma\sigma_1^2 - \sigma_{-1}^2}{(\gamma - 1)\sigma_1^2} - 1 = \frac{\sigma_1^2 - \sigma_{-1}^2}{(\gamma - 1)\sigma_1^2}.$$

This is independent of Δ , but not independent of γ ; which is striking because γ is not pinned down by the limit diffusion. As we will see, γ will matter not only for the limit equilibria in the case of no information aggregation, but also for whether the best limit equilibrium payoff is improved by increased aggregation.

From Proposition 2 a necessary condition for a non-trivial limit equilibrium is

$$\frac{\sigma_1^2 - \sigma_{-1}^2}{(\gamma - 1)\sigma_1^2} > g/(\bar{u} - \underline{u}).$$

Note that $\rho \rightarrow \infty$ as $\gamma \rightarrow 1$. This is because when $\gamma = 1$, outcome 0 has probability 0 under action 0, so incentives can be provided at no cost. Conversely, $\rho \rightarrow 0$ as $\gamma \rightarrow \infty$, because in this case outcome 0 occurs with probability near 1 regardless of the choice of action.

A similar argument to that of the previous subsection shows that when there is a non trivial limit equilibrium, the best limit equilibrium payoff is $\bar{u} - \frac{(\gamma - 1)\sigma_{-1}^2}{\sigma_1^2 - \sigma_{-1}^2}g$. With this case as a baseline we now investigate the effect of information aggregation on the limit equilibrium payoffs.

The simplest case of information aggregation is $\tau(\Delta) = 2\Delta$. Because agents only observe the sum of the two periods outcomes, the possible signals take the values $\{-2, -1, 0, 1, 2\}$. As before, the payoffs in the optimal limit equilibria will depend on the highest possible limiting value of $\rho = \frac{q}{p} - 1$, so we want to determine the maximal value of q/p .

Even without a thorough analysis, it is immediate that aggregation hurts when $\gamma = 1$: Here when $\tau = \Delta$, $p = 0$, $\rho = \infty$, and the equilibrium is fully efficient, while clearly $p > 0$ when $\tau(\Delta) = 2\Delta$, so that the highest attainable ρ is finite and thus the limit equilibrium payoff is bounded away from efficiency.

At the other extreme, where $\gamma \rightarrow \infty$, we have $p = q = 1$ when $\tau = \Delta$, so that $\rho = 0$ and there is only the trivial equilibrium. In this case aggregating two signals could in principle lead to a higher value of ρ and a better limit equilibrium payoff. Appendix 5 gives a detailed analysis of this case, and shows that for some parameter configurations aggregating two signals does indeed lead to a better limit equilibrium payoff, specifically the case where γ is very large and the short-run gain to deviating, g , is very small.

Now consider the case $\tau(\Delta) = \Delta^{1/2}$, so that the signals observed by the players in each period converge to a pair of diffusions. It is important to note that the properties of the limiting diffusion, and thus its limit equilibria, are independent of γ . Thus by

specifying $\frac{\sigma_{+1} - \sigma_{-1}}{\sigma_{-1}} > g/(\bar{u} - \underline{u})$ (so there is a non trivial limit equilibrium for the diffusion) and γ large we can construct examples where there is only the trivial equilibrium when $\tau = \Delta$ and a non trivial limit when players aggregate infinitely many signals, while by specifying γ near 1, and σ_{-1} near σ_{+1} , we have examples with a non trivial limit when $\tau = \Delta$ and a trivial limit when players observe the diffusion. Thus there is no necessary connection between the equilibrium sets in the two cases, and the parameters of the limit diffusion are not sufficient to determine the nature of the equilibrium set when players observe each realization of the underlying process.

We should also point out that when $0 < \frac{\sigma_{+1} - \sigma_{-1}}{\sigma_{-1}} < g/(\bar{u} - \underline{u})$, so that the volatilities are in the interior of the range where the diffusion case has only trivial limit equilibria, then necessarily any sequence $\{\sigma_{+1,n}^2 / \sigma_{-1,n}^2\} \rightarrow_n \sigma_{+1}^2 / \sigma_{-1}^2$ will eventually lie in the interior of this range as well. We conjecture that we could thus use the large-deviations arguments of Appendix 4 to show that any sequence of trinomials with variances converging to $\sigma_{+1}^2, \sigma_{-1}^2$ as $\Delta \rightarrow 0$ will have only trivial equilibria. This result would leave open the question of whether the same conclusion holds for all processes that converge to the specified pair of diffusions.

6.3. Equal variance unequal mean

Finally we turn to trinomials with equal variances and unequal means; this case will be very similar to the binomial case we discussed in Section 5. As there, we suppose that the bad action has a higher mean. With equal variances, $\alpha_{-1} = \alpha_1$; $(1 - \alpha)\gamma = 1$ so $\gamma = 1 + \gamma\alpha, \alpha = (\gamma - 1)/\gamma$; note that the standard binomial case corresponds to $\gamma = 1$ so that $\alpha = 0$.

We begin with the case $\tau = \Delta$. Here the best equilibria punish when the outcome is $+h$ which has probability $\beta_i(\Delta) = \frac{1 - \alpha_i}{2} + \frac{\mu_i \Delta^{1/2}}{2\bar{\gamma}^{1/2}}$, so

$$\frac{q}{p} = \frac{\bar{\gamma}^{1/2}(1 - \alpha) + \mu_{-1}\Delta^{1/2}}{\bar{\gamma}^{1/2}(1 - \alpha) + \mu_{+1}\Delta^{1/2}} = \frac{\sigma\gamma^{-1/2} + \mu_{-1}\Delta^{1/2}}{\sigma\gamma^{-1/2} + \mu_{+1}\Delta^{1/2}} = \frac{\sigma + \mu_{-1}\Delta^{1/2}\gamma^{1/2}}{\sigma + \mu_{+1}\Delta^{1/2}\gamma^{1/2}}$$

Note that as $\Delta \rightarrow 0$, $q/p \rightarrow 1$, just as in the binomial case, and as there the underlying per-event signal becomes completely uninformative in the limit. Thus from Corollary 4b, there is only the trivial equilibrium with any fixed level of aggregation, that is when $\tau = k\Delta$. By analogy with our other results, we believe that this is also true when $\lim_{\tau \rightarrow 0} \tau/\Delta = \infty$, but since the result does not apply to sequences where the variances are only equal in the limit, we have not tried to provide a formal proof.

7. Conclusion

Many different arrays converge to a given diffusion process, and the limit equilibria of these arrays is not in general determined by the parameters of the limiting diffusions, but binomial arrays are an exception to this result. Thus the equilibria of continuous-time games where players monitor the state of a diffusion process are perhaps best thought of as applying to cases where the diffusion specification is either exact or arises from aggregating binomial events.

We have assumed throughout that players observe the aggregate of the process at each period; this is consistent with the idea that the diffusion comes from aggregation. If instead players do not merely see the aggregate, but observe the entire empirical cdf, they then get the first best limit payoff when variances are different and $\tau/\Delta \rightarrow \infty$ regardless of good or bad news. This parallels the observation that observing the infinite-dimensional path of a diffusion for a finite time interval reveals its volatility, which is what underlies the folk wisdom in the continuous time literature that any difference in volatilities leads to full efficiency. However, this full-revelation argument requires that the entire path of the diffusion process is observed, and in many applications, only the aggregate is available as a public signal. For example, firms may have access to one another's revenues or sales data through annual reports, which may possibly disaggregate down to the quarterly level, but will firm not generally have access to the individual sales data of their competitors, which are highly proprietary. Similarly, government reports many aggregates, ranging from money supply figures, to GDP, do hours worked, but the disaggregated data is quite closely held.

Appendix 1: Proof of Proposition 1

Proposition 1: *The most favorable PPE payoff v^* is the largest value v for which there is a function: $w : Z \rightarrow R$ such that (v, w) satisfies the constraints*

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

or $v = \underline{u}$ if no solution exists.

This result was asserted but not proved in Fudenberg and Levine [2007a]. It was used to prove what is Proposition 2 in this paper, so a proof is needed to support our subsequent analysis. The reason a proof is needed is that the conclusion of the theorem applies to both pure and mixed equilibria, but only pure actions are considered in the program (C). This simplification 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.

Proof: We need to show that it is sufficient to consider pure actions. Suppose that the best PPE for the long run player gives more than the static Nash payoff, and fix an equilibrium that attains this payoff. In the first period of this equilibrium, the short-run player must play **In** with positive probability, so the long run player must play some friendly action with positive probability. Fix such an equilibrium, and suppose that the short-run player plays **Out** with positive probability in the first period. Since the short-run player's actions are observed, the strategy profile where LR plays as in the original equilibrium and SR plays **In** with probability 1 in the first period and follows the original strategies thereafter is also a PPE, which shows that SR does not randomize in the first period of the best equilibrium. Finally, if the long-run player randomizes in the first period, the conditions in (C) apply to every action in the support of the first-period distribution, so the maximized value can be attained with a pure strategy. Finally, observe that we require only $v \geq w(z) \geq \underline{u}$ since any payoff in between the best and worst can be attained with public randomization. ☑

Appendix 2: Proof of Proposition 7

Proposition 7: *In the bad news case ($\sigma_{-1}/\sigma_{+1} > 1$) if $\lim_{\Delta \rightarrow 0} \tau(\Delta)/\Delta = \infty$ then there is an efficient limit equilibrium.*

Proof: Consider the strategy of punishing whenever the absolute value of z exceeds a threshold z^* or equivalently when the absolute value of $\zeta(\tau) = \frac{z}{\sigma_{+1}\tau^{1/2}}$

exceeds $\zeta^* = \frac{z^*}{\sigma_{+1}\tau^{1/2}}$. The proof of Proposition 4 of Fudenberg and Levine [2007a]

shows that when the observed outcomes correspond to observing the limit diffusions, specifying a fixed and large value of ζ^* makes $p(\tau)$ and $\lim_{\tau \rightarrow 0} q(\tau)/p(\tau)$ as large as we like. Let $q_*(\tau), p_*(\tau)$ denote the values of q and p computed when observations come from the limit diffusions and a fixed the cutoff rule ζ^* , and let $q_{\Delta(\tau)}(\tau), p_{\Delta(\tau)}(\tau)$ denote the punishment probabilities when the outcomes correspond to observing the sum of $\tau/\Delta(\tau)$ draws of the Z^Δ and the same cut off rule is used. Since we have assumed $\lim_{\Delta \rightarrow 0} \tau(\Delta)/\Delta = \infty$, and the cutoff ζ^* is fixed relative to the standard errors, we can apply the central limit theorem to conclude these probability distributions converge to a normal, so we obtain the same limit values of $\rho = (q - p)/p$ along the triangular arrays corresponding to $\tau(\Delta)$ as we do in the diffusion limit. Consequently, the proof from the earlier paper's Proposition 4 shows that there is an efficient limit equilibrium. □

Appendix 3: Binomial Convergence to Diffusions

Here we prove some results about the convergence of binomials to diffusions needed in proving Proposition 10 in Appendix 4.

Proposition 9: *Suppose the period length is Δ , and that we have i.i.d. binomials $Z_i(\Delta)$ where the common outcomes are $x(\Delta) > y(\Delta)$ and the probability of $x(\Delta)$ under action i is $\alpha_i(\Delta)$, with $\lim_{\Delta \rightarrow 0} \alpha_i(\Delta) = \alpha_i, 0 < \alpha_i < 1$. If the sums $\sum_{j=1}^{\lfloor t/\Delta \rfloor} Z_j$*

converge to a diffusion with drift μ_i and volatilities σ_i^2 then $\sigma_1 = \sigma_2$.

Proof: First we examine what it means for the Z_1 sum to converge to a diffusion. It is convenient to replace the parameters x, y with the parameters $\mu_1^\Delta, \sigma_1^\Delta > 0$, where

$$\begin{aligned} x(\Delta) &= \mu_1^\Delta \Delta + \sigma_1^\Delta \Delta^{1/2} \left(\frac{1 - \alpha_1(\Delta)}{\alpha_1(\Delta)} \right)^{1/2} \\ y(\Delta) &= \mu_1^\Delta \Delta - \sigma_1^\Delta \Delta^{1/2} \left(\frac{\alpha_1(\Delta)}{1 - \alpha_1(\Delta)} \right)^{1/2} \end{aligned}$$

With this new parametrization, we can calculate that $E_1 Z = \mu_1^\Delta \Delta$ and $\text{var}_1 Z = (\sigma_1^\Delta)^2 \Delta$.

If the limit process is a diffusion, then its position at t has the normal distribution with mean μ_1 , variance σ_1^2 . With the reparameterization, this is equivalent to $\mu_1^\Delta \rightarrow \mu_1$ and $\sigma_1^\Delta \rightarrow \sigma_1$. As an illustration, consider the standard binomial limit discussed in section 4: Here we have

$$x = -y = \sigma_1 \Delta^{1/2}, \alpha_1 = (1 + \mu_1 \Delta^{1/2} / \sigma_1) / 2, 1 - \alpha_1 = (1 - \mu_1 \Delta^{1/2} / \sigma_1) / 2,$$

so in this case $\mu_1^\Delta = \mu_1$ and $(\sigma_1^\Delta)^2 = (\sigma_1)^2 - (\mu_1)^2 / [1 / \Delta]$

Now we examine a second sequence of binomial distributions that converges to a different diffusion process with mean μ_2 and variance σ_2^2 . As we discussed earlier, it is important that this second sequence has the same increments $x(\Delta), y(\Delta)$; otherwise, a single realization could be fully informative. So we now have

$$\begin{aligned} (A1) \quad x(\Delta) &= \mu_1^\Delta \Delta + \sigma_1^\Delta \Delta^{1/2} \sqrt{\frac{1 - \alpha_1(\Delta)}{\alpha_1(\Delta)}} = \mu_2^\Delta \Delta + \sigma_2^\Delta \Delta^{1/2} \sqrt{\frac{1 - \alpha_2(\Delta)}{\alpha_2(\Delta)}} \\ y(\Delta) &= \mu_1^\Delta \Delta - \sigma_1^\Delta \Delta^{1/2} \sqrt{\frac{\alpha_1(\Delta)}{1 - \alpha_1(\Delta)}} = \mu_2^\Delta \Delta - \sigma_2^\Delta \Delta^{1/2} \sqrt{\frac{\alpha_2(\Delta)}{1 - \alpha_2(\Delta)}} \end{aligned}$$

We now solve this system to see the possible values of $\sigma_2^\Delta, \alpha_2(\Delta)$ as a function of $\sigma_1^\Delta, \mu_1^\Delta, \mu_2^\Delta, \alpha_1(\Delta)$.

$$\begin{aligned} (\mu_1^\Delta - \mu_2^\Delta) \Delta^{1/2} + \sigma_1^\Delta \sqrt{\frac{1 - \alpha_1(\Delta)}{\alpha_1(\Delta)}} &= \sigma_2^\Delta \sqrt{\frac{1 - \alpha_2(\Delta)}{\alpha_2(\Delta)}} \\ (\mu_1^\Delta - \mu_2^\Delta) \Delta^{1/2} - \sigma_1^\Delta \sqrt{\frac{\alpha_1(\Delta)}{1 - \alpha_1(\Delta)}} &= -\sigma_2^\Delta \sqrt{\frac{\alpha_2(\Delta)}{1 - \alpha_2(\Delta)}} \end{aligned}$$

Divide the two equations to eliminate σ_2^Δ and plug into (A1) to find

$$(A2) \sigma_2^\Delta = \left((\mu_1^\Delta - \mu_2^\Delta) \Delta^{1/2} + \sigma_1^\Delta \sqrt{\frac{1 - \alpha_1(\Delta)}{\alpha_1(\Delta)}} \right) \sqrt{\frac{\sigma_1^\Delta \sqrt{\frac{\alpha_1(\Delta)}{1 - \alpha_1(\Delta)}} - (\mu_1^\Delta - \mu_2^\Delta) \Delta^{1/2}}{\sigma_1^\Delta \sqrt{\frac{1 - \alpha_1(\Delta)}{\alpha_1(\Delta)}} + (\mu_1^\Delta - \mu_2^\Delta) \Delta^{1/2}}}.$$

Since $\lim_{\Delta \rightarrow 0} \alpha_i(\Delta) = \alpha_i, 0 < \alpha_i < 1$, and it follows directly that $\sigma_2^\Delta \rightarrow \sigma_1$. \square

Lemma A3.1: *Suppose the period length is Δ , and that we have i.i.d. binomials $Z_i(\Delta)$ where the common outcomes are $x(\Delta) > y(\Delta)$ and the probability of $x(\Delta)$ is $\alpha_i(\Delta)$, with $\lim_{\Delta \rightarrow 0} \alpha_i(\Delta) = \alpha_i, 0 < \alpha_i < 1$. If the sums $\sum_{i=1}^{\lfloor t/\Delta \rfloor} Z_i$ converge to a non-degenerate diffusion with means μ_i and variances σ_i^2 and $\Delta/\alpha_1 \rightarrow 0$, then*

$$\frac{|(\sigma_1^\Delta)^2 - (\sigma_2^\Delta)^2|}{\Delta^{1/4}} \rightarrow 0.$$

Proof: We wish to show that for $\beta < 1/4$

$$\frac{|(\sigma_1^\Delta)^2 - (\sigma_2^\Delta)^2|}{\Delta^\beta} \rightarrow 0.$$

Since $\frac{|(\sigma_1^\Delta)^2 - (\sigma_2^\Delta)^2|}{\Delta^\beta} = \frac{|\sigma_1^\Delta - \sigma_2^\Delta| |\sigma_1^\Delta + \sigma_2^\Delta|}{\Delta^\beta}$ and $|\sigma_1^\Delta + \sigma_2^\Delta| \rightarrow 2\sigma_1$, this is the same as $\frac{|\sigma_1^\Delta - \sigma_2^\Delta|}{\Delta^\beta} \rightarrow 0$, so

$$\frac{|\sigma_1^\Delta - \sigma_2^\Delta|}{\Delta^\beta} = \left| \sigma_1^\Delta \Delta^{-\beta} - \left((\mu_1^\Delta - \mu_2^\Delta) \Delta^{1/2-\beta} + \Delta^{-\beta} \sigma_1^\Delta \sqrt{\frac{1 - \alpha_1(\Delta)}{\alpha_1(\Delta)}} \right) \sqrt{\frac{\sigma_1^\Delta \sqrt{\frac{\alpha_1(\Delta)}{1 - \alpha_1(\Delta)}} - (\mu_1^\Delta - \mu_2^\Delta) \Delta^{1/2}}{\sigma_1^\Delta \sqrt{\frac{1 - \alpha_1(\Delta)}{\alpha_1(\Delta)}} + (\mu_1^\Delta - \mu_2^\Delta) \Delta^{1/2}}} \right|$$

Since $\beta < 1/4$, the term involving the means goes away and a series of algebraic leads to

$$\begin{aligned} \lim_{\Delta \rightarrow 0} \frac{|\sigma_1^\Delta - \sigma_2^\Delta|}{\Delta^\beta} &= \lim_{\Delta \rightarrow 0} \sigma_1^\Delta \frac{|\mu_1^\Delta - \mu_2^\Delta| \Delta^{1/2-2\beta} |2\alpha_1(\Delta) - 1|}{\alpha_1 \sqrt{\sigma_1^\Delta} \sqrt{\frac{1-\alpha_1(\Delta)}{\alpha_1(\Delta)}} \left(2\sqrt{\sigma_1^\Delta} \sqrt{\frac{1-\alpha_1(\Delta)}{\alpha_1(\Delta)}} \right)} = \\ &= \frac{|\mu_1 - \mu_2|}{2\sigma_1} \lim_{\Delta \rightarrow 0} \left(\frac{\Delta}{\alpha_1} \right)^{1/2} \Delta^{-2\beta} \frac{|2\alpha_1(\Delta) - 1|}{\sqrt{1-\alpha_1(\Delta)}} \end{aligned}$$

Since $\lim_{\Delta \rightarrow 0} \alpha_i(\Delta) = \alpha_i, 0 < \alpha_i < 1$ and $1/2 - 2\beta > 0$, the result follows. \square

Using the central limit theorem, the conditions $\mu_1^\Delta \rightarrow \mu_1, \sigma_1^\Delta \rightarrow \sigma_1$ and $\alpha_1(\Delta) \rightarrow \alpha_1, 0 < \alpha_1 < 1$ can be shown to be sufficient for a triangular array to converge to a diffusion. To construct a non-standard binomial array, it is convenient just to take $\mu_1^\Delta = \mu_1, \sigma_1^\Delta = \sigma_1, \alpha_1(\Delta) = \alpha_1$. Using our alternative parametrization from above we find for example that if $\mu_1 = 0, \sigma_1 = 1, \alpha_1 = 1/3$, we have the binomial taking on the values $x(\Delta) = \sqrt{2}\Delta^{1/2}, y(\Delta) = -(\sqrt{2}/2)\Delta^{1/2}$ with probability of $x(\Delta)$ equal to $1/3$, which generate a triangular array that converges to a diffusion with drift μ_1 and volatility σ_1 .

Appendix 4: Proof of Proposition 10

As in the text, we consider a sequence of games with both the event period Δ and the observation period $\tau(\Delta)$ converging to 0, and define $k(\Delta) = \tau(\Delta)/\Delta$.

Proposition 10: *Suppose*

(i) $\lim_{\Delta \rightarrow 0} \tau(\Delta) \exp(k(\Delta)^{2/7}) = \infty$

(ii) *the signals are i.i.d. binomials $Z(\Delta)$ where the common outcomes are $x(\Delta) > y(\Delta)$, and the probability of $x(\Delta)$ under action $i = +1, -1$ is $\alpha_i(\Delta)$ with $\lim_{\Delta \rightarrow 0} \alpha_i(\Delta) = \alpha_i, 0 < \alpha_i < 1$*

(iii) $\sum_{j=1}^{\lceil t/\Delta \rceil} Z_j$ converge to non-degenerate diffusions with means μ_i and variances σ_i^2 .

Then all limit equilibria are trivial.

Note that condition (i) of this proposition requires that $\lim_{\Delta \rightarrow 0} k(\Delta) = \infty$; we will maintain that restriction throughout this appendix. We start by summarizing some

notation and key results from other places. Let Φ, ϕ respectively denote the c.d.f. and density of the standard normal distribution.

Fact 1 [Fudenberg and Levine [2007a] Proposition 2]: *Suppose the signals are normally distributed with means $-a_1\tau$ and variance $\sigma^2\tau$. Then for any $\rho_0 > 0$, $\rho > \rho_0$ implies $q/\tau \rightarrow 0$ and so there is no non-trivial limit equilibrium.*

For a fixed distribution F , let $\psi_F(\zeta) \equiv \log \int_{-\infty}^{\infty} e^{\zeta x} F(dx)$ be the logarithm of the generating function. We will be interested in the distributions corresponding to the binomial distributions referred to in Proposition 10: In this case we have

$$\psi_{\Delta}(\zeta) = \log(\alpha_{\Delta} e^{\zeta x(\Delta)} + (1 - \alpha_{\Delta}) e^{-\zeta y(\Delta)})$$

Fact 2 [Large Deviations Theorem for Triangular Arrays, from Fudenberg and Levine [2007b]] : *Suppose that for each n there is a sequence Z_j^n $j = 1, \dots, n$ of i.i.d. random variables with zero mean, variance σ_n^2 and distribution F_n , and that $z_n = \sum_{j=1}^n Z_j^n$ has distribution F^{n*} , while the normalized sum $z_n / \sigma_n \sqrt{n}$ has distribution F^n . If*

1. *For some $\bar{s} > 0$ and all $0 \leq \zeta \leq \bar{s}$ there is a continuous function $\psi^2(\zeta) > 0$ and constant $B > 0$ such that $\lim_{n \rightarrow \infty} \sup_{0 \leq \zeta \leq \bar{s}} |\psi_n''(\zeta) - \bar{\psi}''(\zeta)| \rightarrow 0$ and $\sup_{0 \leq \zeta \leq \bar{s}} |\psi_n''''(\zeta)|, |\psi_n''''(\zeta)\zeta|, |\psi_n''''(\zeta)\zeta^2| < B$*

2. $\sigma_n \rightarrow \sigma$, $M_{3n} \equiv E|Z_n^i|^3 \rightarrow M_3 < \infty$

3. $n^{-1/6} x_n \rightarrow 0$

4. $x_n \rightarrow \infty$

Then

$$\frac{1 - F^n(x_n)}{1 - \Phi(x_n)} \rightarrow 1$$

In what follows, we will take the limit on $k \rightarrow \infty$ rather than $\Delta \rightarrow 0$, implicitly considering a sequence $\tau^k \rightarrow 0$ and $\Delta^k = \tau^k / k$.¹³ As in the proof of Proposition 9 we define new parameters $\mu_i^k = E_i Z, (\sigma_i^k)^2 = \text{var}_i Z / \Delta^k$.

We are interested in applying the Large Deviation Theorem to $\hat{z} \equiv \sum_{j=1}^k Z_j$

This leads us to define

$$\tilde{Z}_i^k = \frac{Z - \mu_i^k \Delta_k}{\sqrt{\Delta_k}}$$

so that $E \tilde{Z}_i^k = 0, \text{var } \tilde{Z}_i^k = (\sigma_i^k)^2 \rightarrow \sigma_i^2$ and the values taken on by the reparameterized binomial are

$$\begin{aligned} \tilde{x}_i^k &= \sigma_i^k \left(\frac{1 - \alpha_i(\Delta_k)}{\alpha_i(\Delta_k)} \right)^{1/2} \\ \tilde{y}_i^k &= -\sigma_i^k \left(\frac{\alpha_i(\Delta_k)}{1 - \alpha_i(\Delta_k)} \right)^{1/2}. \end{aligned}$$

Note that $\lim_{k \rightarrow \infty} x_i^k = \sigma_i$: the reparameterized binomial has step size tending to a non-zero constant.

Lemma A.4.1: *Suppose that the signals are i.i.d. binomials $Z(\Delta)$ where the common outcomes are $x(\Delta) > y(\Delta)$, and the probability of $x(\Delta)$ under action $i = +1, -1$ is $\alpha_i(\Delta)$ with $\lim_{\Delta \rightarrow 0} \alpha_i(\Delta) = \alpha_i, 0 < \alpha_i < 1$, that $k \rightarrow \infty$, and that $\sum_{j=1}^{\lfloor t/\Delta \rfloor} Z_j$ converge to non-degenerate diffusions with means μ_i and variances σ_i^2 .*

Then the reparameterized binomials satisfy conditions 1 and 2 of the large deviations theorem for both $i = -1, +1$.

Proof: We consider the case $i = +1$, the case $i = -1$ is identical save for notation. The cumulant generating function for +1 is

$$\psi_k(\zeta) = \log \left(\alpha_{+1}(\Delta_k) \exp \left(\zeta \sigma_i^k \left(\frac{1 - \alpha_{+1}(\Delta_k)}{\alpha_{+1}(\Delta_k)} \right)^{1/2} \right) + (1 - \alpha_{+1}(\Delta_k)) \exp \left(-\zeta \sigma_i^k \left(\frac{\alpha_{+1}(\Delta_k)}{1 - \alpha_{+1}(\Delta_k)} \right) \right) \right)$$

$$\text{Let } \hat{\psi}(\zeta) = \log \left(\alpha_{+1} \exp \left(\zeta \sigma_i \left(\frac{1 - \alpha_{+1}}{\alpha_{+1}} \right)^{1/2} \right) + (1 - \alpha_{+1}) \exp \left(-\zeta \sigma_i \left(\frac{\alpha_{+1}}{1 - \alpha_{+1}} \right) \right) \right).$$

¹³ We write k as a superscript as the subscript denotes the action $i = +1, -1$.

Because $\alpha_{+1}(\Delta_k) \rightarrow \alpha_{+1}$ and $\sigma_i^k \rightarrow \sigma_i$, we know that $\lim_{k \rightarrow \infty} \sup_{0 \leq \zeta \leq \bar{s}} |\psi_k''(\zeta) - \hat{\psi}''(\zeta)| \rightarrow 0$ so the first part of condition 1 is satisfied, and it is clear by inspection that the other necessary conditions hold as well. \square

We turn now to the main proof. The idea is to show that if there were strategies that led to a non-trivial limit equilibrium in the binomial case, we could construct a non-trivial limit equilibrium in the normal case. There are several details that need to be attended to in order for this argument to work. First, the approximating normals corresponding to the two different actions will have different variances, while Fact 1 supposes that the variances are equal before the limit is reached. Lemma A.3.1 adds a condition on the rate of convergence that enables us to extend Fact 1 to the case where the variances are different before the limit is reached. Moreover, while we know that within each period z is converging to a normal, the cutoff for punishment might be going to infinity, so the standard central limit theorem does not apply. Hence we use the large deviation theorem described above. The idea is to show that if the cutoff grows faster than $k^{1/6}$ the probability of punishment is so low that it cannot sustain a non-trivial equilibrium, while if it grows at $k^{1/6}$ the normal approximation is so good that we can make use of Fact 1.

First we give a Lemma needed to deal with variances that are unequal before the limit is reached.

Lemma A.4.2: *Suppose $\sigma_n \rightarrow \sigma$, and that $\zeta_n^2 / |\sigma_n^2 - \tilde{\sigma}_n^2| \rightarrow 0$. Then*

$$\frac{1 - \Phi\left(\frac{\zeta_n - \mu}{\sigma_n}\right)}{1 - \Phi\left(\frac{\zeta_n - \mu}{\tilde{\sigma}_n}\right)} \rightarrow 1$$

proof: Observe from L'Hopital's rule that if $x \rightarrow \infty$ then $(1 - \Phi(x)) / \phi(x) \rightarrow 0$. Using that fact, we may again apply L'Hopital's rule to see that $\lim_{x \rightarrow \infty} x(1 - \Phi(x)) / \phi(x) = 1$. Define

$$x_n = \frac{\zeta_n - \mu}{\sigma_n} \text{ and } \tilde{x}_n = \frac{\zeta_n - \mu}{\tilde{\sigma}_n}.$$

Then

$$\begin{aligned}
& \lim_{n \rightarrow \infty} \frac{1 - \Phi\left(\frac{\zeta_n - \mu}{\sigma_n}\right)}{1 - \Phi\left(\frac{\zeta_n - \mu}{\tilde{\sigma}_n}\right)} = \lim_{n \rightarrow \infty} \frac{1 - \Phi(x_n)}{1 - \Phi(\tilde{x}_n)} \\
&= \lim_{n \rightarrow \infty} \frac{(1 - \Phi(x_n)) \frac{\phi(x_n)}{x_n (1 - \Phi(x_n))}}{(1 - \Phi(\tilde{x}_n)) \frac{\phi(\tilde{x}_n)}{\tilde{x}_n (1 - \Phi(\tilde{x}_n))}} \\
&= \lim_{n \rightarrow \infty} \frac{\tilde{x}_n \phi(x_n)}{x_n \phi(\tilde{x}_n)} = \lim_{n \rightarrow \infty} \frac{\sigma_n}{\tilde{\sigma}_n} \exp\left(\frac{\left(\frac{\zeta_n - \mu}{\tilde{\sigma}_n}\right)^2 - \left(\frac{\zeta_n - \mu}{\sigma_n}\right)^2}{2}\right) \\
&= \lim_{n \rightarrow \infty} \exp\left(\frac{(\sigma_n^2 - \tilde{\sigma}_n^2)(\zeta_n - \mu)^2}{2\sigma_n^2 \tilde{\sigma}_n^2}\right) = 1
\end{aligned}$$

□

Lemma A.4.3: *The MLRP is satisfied by z .*

Proof: This is well known and can be verified by directly calculating the likelihood ratio for the multinomial sum of binomials. □

Define the random variable normalized for the agent taking action +1 as

$$\tilde{z}_{+1}^k \equiv \frac{\hat{z} - \mu_{+1}^k \Delta k}{\sigma_{+1}^k \sqrt{k} \Delta} = \frac{\sum_{j=1}^k \tilde{Z}_j}{\sigma_{+1}^k \sqrt{k}}.$$

Since the MLRP is satisfied, we may assume a strategy of the form “punish” if $\tilde{z}_{+1}^k > \bar{\zeta}^k$.

Lemma A.4.4: *If $\lim_{k \rightarrow \infty} \bar{\zeta}_k k^{-1/6} > 0$ and $\tau^k \exp(k^{2/7}) \rightarrow 0$ then $p/\tau \rightarrow 0$, $q/\tau \rightarrow 0$.*

Proof: It suffices to prove $q/\tau \rightarrow 0$ since $q \geq p$. To compute q , we need to consider the distribution of \tilde{z}_{+1}^k when the action taken is -1 . This does not have zero mean or unit

variance, so we renormalize, defining $\tilde{z}_{-1}^k \frac{\sigma_{+1}^k}{\sigma_{-1}^k} \tilde{z}_{+1}^k - (\mu_{-1}^k - \mu_{+1}^k) \Delta_k k$, which has zero

mean and unit variance when the action taken is -1 . Denote the cdf of this random variable when the action is -1 by G^k ; we may potentially apply the large deviations theorem to this distribution. In the new normalization, the cutoff is

$$x^k = \frac{\sigma_{+1}^k}{\sigma_{-1}^k} \bar{\zeta}_k - (\mu_{-1}^k - \mu_{+1}^k) \Delta_k k$$

so that $q^k = 1 - G^k(x^k)$. However, since $\lim \bar{\zeta}_k k^{-1/6} > 0$, $|\sigma_{+1}^k - \sigma_{-1}^k| \rightarrow 0$ and μ_i^k is bounded, we see that $\lim_{k \rightarrow \infty} x_k k^{-1/6} > 0$. Hence we can choose a sequence $\bar{x}_k \leq x_k$ so that $k^{-1/6} \bar{x}_k \rightarrow 0$, $\lim_{k \rightarrow \infty} k^{-1/7} \bar{x}_k = \sqrt{2\pi}$, so that the conditions 3,4 of the large deviations theorem above are satisfied for \bar{x}_k . It follows that for k sufficiently large $q \leq 2(1 - \Phi(\bar{x}_k))$. Next use L'Hopital's rule to see that $\frac{1 - \Phi(\bar{x}_k)}{\phi(\bar{x}_k)} \rightarrow 0$ so for large

enough k we have

$$q \leq 3\phi(\bar{x}_k) = 3 \exp\left(-\frac{(\bar{x}_k)^2}{2\pi}\right) / \sqrt{2\pi}.$$

Since $\lim k^{-1/7} \bar{x}_k = \sqrt{2\pi}$, we have that $p/\tau \leq C(\tau \exp(k^{2/7})) \rightarrow 0$ □

Proposition 10: *Suppose*

(i) $\lim_{\Delta \rightarrow 0} \tau(\Delta) \exp(k(\Delta)^{2/7}) = \infty$

(ii) *the signals are i.i.d. binomials $Z(\Delta)$ where the common outcomes are $x(\Delta) > y(\Delta)$, and the probability of $x(\Delta)$ under action $i = +1, -1$ is $\alpha_i(\Delta)$ with $\lim_{\Delta \rightarrow 0} \alpha_i(\Delta) = \alpha_i, 0 < \alpha_i < 1$*

(iii) $\sum_{j=1}^{\lfloor t/\Delta \rfloor} Z_j$ *converge to non-degenerate diffusions with means μ_i and variances σ_i^2 .*

Then all limit equilibria are trivial.

Proof: By Lemma A.4.3 the signals satisfy the MLRP, so we can restrict attention to strategies that punish when the observed signal exceeds some cutoff. By Lemma A.4.4 if there is a non-trivial limit, we may assume that the cutoff satisfies $\lim \bar{\zeta}_k k^{-1/6} = 0$. By Lemma A.4.1 this means that we may compute $p/\tau, q/\tau, \rho$ asymptotically using normal distributions with variances $\text{var}(Z_j | a_1)$. From Lemma A.3.1

$$\frac{k^{1/4} |(\sigma_{+1}^k)^2 - (\sigma_{-1}^k)^2|}{\tau^{1/4}} = \frac{|(\sigma_{+1}^k)^2 - (\sigma_{-1}^k)^2|}{(\tau/k)^{1/4}} \rightarrow 0,$$

so that $k^{1/4} | (\sigma_{+1}^k)^2 - (\sigma_{-1}^k)^2 | \rightarrow 0$. Since $\lim \zeta_k k^{-1/6} = 0$ also $\lim \zeta_k k^{-1/4} = 0$, implying $\zeta_k | (\sigma_{+1}^k)^2 - (\sigma_{-1}^k)^2 | \rightarrow 0$. Consequently Lemma A.4.2 applies, so that we may assume that the normals have the same variance, implying a non-trivial limit in that case. This contradicts Fact 1. \square

Appendix 5: Aggregating Two Good- News Signals

We want to show that aggregating two trinomial good-news signals leads to a better limit equilibrium payoff when γ is very large and the short-run gain to deviating, g , is very small. To do this we determine the best limit equilibrium payoff when aggregating two signals.

Punishing when the sum of the signals -2 and +2 will minimize and not maximize the target ratio, and with a 0 mean the signals -1 and +1 are symmetric. Thus it will be enough to determine q/p for the signals 0 and +1. To do this we first calculate q and p for these two signals. Note that for $i = 1, -1$

$$\begin{aligned} \Pr(\sum_{j=1,2} Z_j = 0) &= \Pr\{(0,0), (h, -h), (-h, h)\} = \\ \alpha_i^2 + 2\left(\frac{1 - \alpha_i}{2}\right)^2 &= \frac{3}{2}\alpha_i^2 + \frac{1}{2} - \alpha_i \end{aligned}$$

This is minimized at $\alpha_i = 1/3$, where it has value $1/3$, and has values $1/2$ at 0 and 1 at 1.

Next $\Pr(\sum_{j=1,2} Z_j = 1) = \alpha_i(1 - \alpha_i) = \alpha_i - \alpha_i^2$. This is maximized at $1/2$.

Thus if the strategies punish when the sum is 1, $q/p = \frac{\alpha_{-1} - \alpha_{-1}^2}{\alpha_{+1} - \alpha_{+1}^2} = B$, and if

the strategies punish when the sum is 0 we have $q/p = \frac{3\alpha_{-1}^2 + 1 - 2\alpha_{-1}}{3\alpha_{+1}^2 + 1 - 2\alpha_{+1}} = A$.

To compare $\max(A, B)$ to the likelihood ratio $C = \frac{\alpha_{-1}}{\alpha_{+1}}$ for a single observation, we first compare B and C:

$$\frac{B}{C} = \frac{\frac{\alpha_{-1} - \alpha_{-1}^2}{\alpha_{+1} - \alpha_{+1}^2}}{\frac{\alpha_{-1}}{\alpha_{+1}}} = \frac{1 - \alpha_{-1}}{1 - \alpha_{+1}} < 1$$

so unsurprisingly (0,1) and (1,0) are less informative than (0).

Next, we ask when $A < C$. This is true when $\frac{3\alpha_{-1}^2 + 1 - 2\alpha_{-1}}{3\alpha_{+1}^2 + 1 - 2\alpha_{+1}} < \frac{\alpha_{-1}}{\alpha_{+1}}$.

Note that $\alpha_{-1} > \alpha_{+1}$ because we are in the good news case. Observing that all the expressions are non-negative we can write this as $\frac{3\alpha_{-1}^2 + 1 - 2\alpha_{-1}}{\alpha_{-1}} < \frac{3\alpha_{+1}^2 + 1 - 2\alpha_{+1}}{\alpha_{+1}}$. The same function $f(\alpha) = \frac{3\alpha^2 + 1 - 2\alpha}{\alpha}$

appears on both the left and right hand side of this inequality. Its derivative is

$f'(\alpha) = \frac{6\alpha^2 - 2\alpha - 3\alpha^2 - 1 + 2\alpha}{\alpha^2}$, so $f'(1) > 0$, and thus when $\alpha_1 < \alpha_{-1}$ and both are sufficiently close to 1, we have $f(\alpha_{-1}) > f(\alpha_{+1})$; since $\alpha_1, \alpha_{-1} \rightarrow 1$ as $\gamma \rightarrow \infty$, aggregating two signals together improves the best likelihood ratio as $\gamma \rightarrow \infty$. On the other hand, the maximized value of this likelihood decreases to 1 as $\gamma \rightarrow \infty$. Thus for some payoff functions, the values of γ for which aggregating two signals improves the likelihood ratio may be so large that even with two signals there is only a trivial limit equilibrium. On the other hand, aggregation can allow a switch from non trivial to trivial limits if both γ is very large and g is very small, so that $\frac{\sigma_1^2 - \sigma_{-1}^2}{(\gamma - 1)\sigma_1^2} = g/(\bar{u} - \underline{u})$, and the one period likelihood ratio is just on the edge of the region that supports a non trivial limit.

References

- Abreu, D., P. Milgrom and D. Pearce [1991] "Information and Timing in Repeated Partnerships." *Econometrica*, 59, 1713–33.
- Abreu, D., D. Pearce, and E. Stachetti [1990] "Towards a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica* 58: 1041-1064.
- Celentani, M., D. Levine and C. Martinelli [2007], "Punishing, Investigating or Auditing?" mimeo.
- Faingold, E. and Y. Sannikov [2007] "Equilibrium Degeneracy and Reputation Effects," mimeo.
- Feller, W. [1971] *An Introduction to Probability Theory and Its Applications*, Volume II, John Wiley & Sons, New York.
- Fudenberg, D., D M. Kreps, and E. Maskin [1990] "Repeated Games with Long-run and Short-run Players" *Review of Economic Studies*, 57: 555-573.

- Fudenberg, D. and D.K. Levine [1993] "Limit Games and Limit Equilibria" *Journal of Economic Theory*, 38{261-279.
- 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 [2007a] "Continuous-Time Limits of Discrete-Time Games," mimeo, forthcoming in the *Review of Economic Dynamics*.
- Fudenberg, D. and D.K. Levine [2007b] "A Large Deviation Theorem for Triangular Arrays," mimeo.
- Fudenberg, D. and J. Tirole [1991] *Game Theory*, MIT Press, Cambridge, MA.
- Fudenberg, D. and J. Tirole [1985] "Preemption and Rent Equalization in the Adoption of New Technology," *Review of Economic Studies*, 52 (1985), 383-402;
- Green, E. and R. Porter [1984] "Non-Cooperative Collusion Under Imperfect Price Information," *Econometrica* 52: 87-100.
- Hellwig, M. and K. Schmidt [2002]. "Discrete-Time Approximations of the Holmström-Milgrom Brownian-Motion Model of Intertemporal Incentive Provision," *Econometrica*, 70: 2225-2264.
- Kandori, M. [1992] "The Use of Information in Repeated Games with Imperfect Monitoring," *Review of Economic Studies*, 59, 581-594.
- Klein, B. and K. Leffler [1981] "The Role of Market Forces in Assuring Contractual Performance," *Journal of Political Economy*, 89: 615-41.
- Levin, J. [2003], "Relational incentive contracts," *American Economic Review* 93, 835-57.
- Mailath, G. and L. Samuelson [2007] *Repeated Games and Reputations: Long-Run Relationships*, MIT Press, Cambridge, MA.
- 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.
- Radner, R., R. Myerson, and E. Maskin [1986] "An Example of a Repeated Partnership Game with Discounting and with Uniformly Inefficient Equilibria," *Review of Economic Studies*, 53: 59–69.
- Sannikov, Y. [2007a] "Games with Imperfectly Observed Actions in Continuous Time," *Econometrica*, 75:1285-1329.
- Sannikov, Y. [2007b] "Repeated Games with Imperfect Monitoring: Connection between Discrete and Continuous Time," mimeo.
- Sannikov, Y., and A. Skrzypacz [2007] "The Role of Information in Repeated Games with Frequent Actions," mimeo.
- Sannikov, Y., and A. Skrzypacz, "The Impossibility of Collusion under Imperfect Monitoring with Flexible Production," mimeo.
- Simon, L. and M. Stinchcombe [1989] "Extensive Form Games in Continuous Time: Pure Strategies," *Econometrica* 57:1171-1214.
- Stokey, N. [2007] *Brownian Models in Economics*, mimeo.
- Yamamoto, Y [2007] "The Use of Public Randomization in Discounted Repeated Games," mimeo.