Repeated Games with Frequent Signals

Citation

Fudenberg, Drew, and David K. Levine. 2009. Repeated games with frequent signals. Quarterly Journal of Economics 124, no. 1: 233-265.

Published Version

<http://dx.doi.org/10.1162/qjec.2009.124.1.233>

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:3160491>

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](#)

REPEATED GAMES WITH FREQUENT SIGNALS*

DREW FUDENBERG AND DAVID K. LEVINE

We study repeated games with frequent actions and frequent imperfect public signals, where the signals are aggregates of many discrete events, such as sales or tasks. The high-frequency limit of the equilibrium set depends both on the probability law governing the discrete events and on how many events are aggregated into a single signal. When the underlying events have a binomial distribution, the limit equilibria correspond to the equilibria of the associated continuous-time game with diffusion signals, but other event processes that aggregate to a diffusion limit can have a different set of limit equilibria. Thus the continuous-time game need not be a good approximation of the high-frequency limit when the underlying events have three or more possible values.

I. INTRODUCTION

We study the limits of equilibria of repeated games with imperfect public information as the frequency of observations and actions grows to infinity. To highlight the role of the information structure, we focus 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 consumers, each of whom purchases only once. In the Klein-Leffler model, in each period the firm chooses an intended quality level, but the production process is stochastic, so that the realized quality may differ from the intended one.¹ Consumers will only purchase if the firm is expected to try to produce high quality in the current period; the firm has a short-run incentive to cut costs and produce low quality, but there can be equilibria in which the firm tries for high quality to avoid losing future sales.² As a second illustration, the nonstrategic players could be shareholders, and the long-run player the manager of the firm. The manager chooses an effort level, but this is not observed by the shareholders, who do observe the realized sales in each period. Because the long-run player's action is observed with noise, the set of equilibria depends on the

*We thank Ed Glaeser and two referees for useful comments, and Mihai Manea for careful proofreading. NSF grants SES0426199 and SES0646816 provided financial assistance.

1. In the original Klein-Leffler model, the production technology was deterministic.

2. The key aspect of these "short-run players" is that they are strategically myopic and do not try to influence the future play of the long run player. The same analysis applies when the "short-run players" are replaced by a continuum of infinitesimal long-lived players, with the large player observing only the aggregate play of the small ones.

information structure, and typically efficient payoffs cannot be approximated by equilibria, even in the limit as the discount factor tends to 1 (Fudenberg and Levine 1994).

We show how the best equilibrium payoff for the long-run player depends on the information structure, and in turn on how the relevant characteristics of the information structure change as the observation period shrinks.³ Our work builds on our earlier paper Fudenberg and Levine (2007a), which provides general conditions for a sequence of discrete-time games with period length going to zero to have a nontrivial limit equilibrium. 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 some cases, frequent interactions permit fully efficient outcomes, for example, if consumers receive such accurate information that the firm can be induced to almost always produce high quality. In other cases, the equilibrium set collapses in the limit, and only the static equilibrium can be supported, so that the firm produces low quality forever.

We focus on cases where the public signal is an aggregate of several or many discrete events, such as sales, price changes, or components of quality, and in particular on the case where the distribution of this aggregate converges to a diffusion process. We feel that this is of relevance for interpreting results about continuous time games where players observe the state of a diffusion process, as in Faingold and Sannikov (2007), Fudenberg and Levine (2007a), Sannikov (2007a), and Sannikov and Skrypcz (2008), because in most settings of interest the diffusion assumption is an approximation for a sum of discrete events.⁴

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 the same limiting signal distribution

3. In more general games, the set of equilibrium payoffs will depend on the information structure in more complicated ways, but our calculation of the “limit informativeness” of various sequences of signal structures will still apply.

4. Diffusion processes are continuous, yet processes such as sales or price are inherently discrete, and so players would observe at most a single transaction in each period if they monitored the process at a sufficiently high frequency.

may correspond to ways of passing to the limit that have very different limit equilibria. We also highlight the role of “information aggregation” in determining the limit equilibrium payoffs. That is, we ask when observing the sum of many signals leads to a larger limit equilibrium set than observing the signals one at a time. Our previous paper showed that there are efficient limit equilibria if deviations increase the volatility of the diffusion but not when deviations decrease the volatility. We relate the differing conclusions in these two cases to their differing aggregation properties: when deviating leads to increased volatility, the signal structure is more informative when players observe the aggregate of the discrete events instead of observing each event as it occurs, and the informativeness becomes infinite as players aggregate more and more observations; this is not the case when deviating lowers volatility. Finally, our results show that the usual continuous-time games with controlled diffusions correspond to some but not all of the ways that the discrete-time observation structures can converge to a diffusion, so the standard continuous-time model is “too small” to incorporate all of the relevant limit objects.⁵

Like the earlier paper, this one is related to that of 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 that of Sannikov and Skrzypacz (2008). They consider a game with two long-run players observing the infinite-dimensional sample path of a continuous-time Lévy process at discrete intervals and provide an upper bound on the set of pure-strategy equilibrium payoffs as the time interval shrinks to zero. 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. This is also the case in Faingold and Sannikov (2007) and Sannikov and Skrzypacz (2007). Unlike those papers, we allow mixed as well as pure strategies and explicitly consider the limit of equilibria when signals are an aggregate of underlying underlying discrete-time events.

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

5. This suggests that one might want to construct a larger space of continuous-time games, as done by Fudenberg and Tirole (1985) and Simon and Stinchcombe (1989) in a related context, but that remains a topic for future work.

TABLE I
STAGE-GAME PAYOFFS

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

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 work points to, loosely speaking, a failure of upper hemicontinuity.

On a more practical level, the equilibria of games played at high but finite frequency depends on the informativeness of the available signals. Even when these signals can be well approximated by a diffusion, the equilibria of the standard continuous time models may not be a good approximation of the finite-frequency equilibria, unless the underlying signal process is binomial. Otherwise, whether or not the continuous time results are relevant is an empirical issue and is not a necessary consequence of the periods being short.

II. 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 2×2 stage game shown in Table I, 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 payoffs 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.⁶ 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.⁷

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 by its upper bound, which we denote by v^* .⁸ Proposition 1 in Appendix I shows that v^* can be computed as the solution to a static linear programming problem, where the control variables

6. Technically speaking, the public information also includes the short-run player’s action, but because 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 equilibria. 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. Because that expected value lies between the best and worst equilibria, there is an equivalent equilibrium in which the continuation value is constant and equal to the conditional expected value.

7. See Fudenberg and Tirole (1991) for a definition of this concept and an example of a nonpublic equilibrium in a game with public monitoring.

8. 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.

are the “continuation payoffs” $w(z)$ that the player expects to receive following each signal z ; this result is used in the proof of our first 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 1 (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

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

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

$$(2) \quad \max_{p, q \text{ feasible}} \bar{u} - \frac{pg}{q - p}.$$

When (1) 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

$$\begin{aligned} f(z \mid a_1 = -1, \tau) / f(z \mid a_1 = +1, \tau) &> \lambda^* \text{ then } w(z) = \underline{u}, \text{ if} \\ f(z \mid a_1 = -1, \tau) / f(z \mid a_1 = +1, \tau) &< \lambda^* \text{ then } w(z) = v^*. \end{aligned} \quad 9$$

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

III. 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 thus consider families of games indexed by the period length τ . We must now describe how the signal z varies with

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

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.¹⁰ 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 $1/\Delta$. 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 $\tau(0) = 0$. We then define $k(\Delta) = \tau(\Delta)/\Delta$; players observe the integer number $\lfloor k(\Delta) \rfloor$ of signals when the time between moves is $\tau(\Delta)$. In general, we allow the distribution of Z_j and its support to depend on Δ , and to emphasize this dependence we will write $Z_j(\Delta)$. (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 $Z_j(\Delta)$ is a constant, independent of Δ .

The information available at the end of the period beginning at t is the signal $z = \sum_{j=\lfloor t/\Delta \rfloor}^{\lfloor (t+\tau)/\Delta \rfloor} Z_j(\Delta)$. 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 nonnegative 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 *nontrivial limit equilibrium* for r . If there is any positive interest rate r for which there is a nontrivial limit equilibrium, we say simply that there is a nontrivial 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 $\bar{r} > 0$ 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

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

payoffs converging to $\bar{u}$, we say that there is an *efficient patient equilibrium*.¹¹

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 (2) 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 1.¹²

- (a) If for some sequence $(r, \tau) \rightarrow (0, 0)$ 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 sequences $\tau \rightarrow 0$, and all equilibria, $\rho(\tau) \rightarrow 0$, then there is only a trivial limit equilibrium.
- (c) If for all $r > 0$ and all sequences $\tau \rightarrow 0$ there are an $\varepsilon > 0$ and a sequence of equilibria with

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

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

The first two parts of this result are immediate. Part (c) follows from the observation that $\rho(\tau) > \frac{g}{\bar{u} - \underline{u}}$ implies that the LHS of (1) is positive, so $\lim_{\tau \rightarrow 0} v^*(\tau) > \underline{u}$ in cases where the RHS of (1) converges to 0, which is true in particular when $p(\tau)$ is bounded away from 0.

11. Note that the definition of a nontrivial limit equilibrium allows the interest rate r to be arbitrarily small, but it requires 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.

12. 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.

In many cases of interest, the best equilibria will have $p(\tau)$ converging to 0.

COROLLARY 2.

- (a) If for all $r > 0$ and all sequences $\tau \rightarrow 0$, along any sequence of best equilibria $\rho(\tau) > g/(\bar{u} - \underline{u})$ implies $q/\tau \rightarrow 0$, then there is only a trivial limit.¹³
- (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 there is an $\bar{r} > 0$ such that for all $0 < r < \bar{r}$, 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. We can rewrite (1) as

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

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

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

The RHS of this inequality is bounded below by $\frac{rg}{\bar{u} - \underline{u}}$, to which it converges as $\rho \rightarrow \infty$.

This immediately yields parts (a) and (b). For (c), note that when λ is sufficiently large, the LHS of (1) is positive and bounded away from 0; the RHS converges to 0, so using the strategies that generate these probabilities yields a nontrivial limit equilibrium, and the payoff to this equilibrium converges to $\bar{u}$ as $\lambda \rightarrow \infty$ ■

IV. CONVERGING TO DIFFUSIONS

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

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

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 $z = \sum_{j=1}^{\lfloor t/\Delta \rfloor} Z_j(\Delta)$ 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 $Z_j(\Delta)$, no outcome perfectly reveals which of the two actions was played. As Δ and τ converge to 0, the distribution of the $Z_j(\Delta)$ will change; we let f^Δ denote the $Z_j(\Delta)$, and f^τ the density of the aggregate z .

The $Z_j(\Delta)$ represent underlying economic events that are being aggregated. As long as their distribution is well-behaved, the central limit theorem applies and each of the triangular arrays converges to a diffusion.¹⁴ In fact, many different distributions on the $Z_j(\Delta)$ may generate the same diffusion. Our goal is to understand whether the limit diffusion is sufficient to characterize the set of limit equilibria or whether the details of the particular triangular array matter. In practice, the distribution of the underlying events depends on the situation being modeled. In some settings it is natural to think of the data as having a binomial distribution. For example, if the data being observed are sales data, and the items being sold are cars or refrigerators or other large durable goods, then it is reasonable to think that a consumer either buys the item or not, but does not buy several at once. On the other hand, for goods sold by volume or weight, each individual sale can take on many different values, so the underlying data being aggregated have a nontrivial distribution of their own.

We take up first the simplest case, where the underlying $Z_j(\Delta)$ do in fact follow a binomial distribution. This distribution has the special feature that its mean and its variance are linked to each other: the variance is equal to the product of the mean and one minus the mean. As the next result shows, this link between the mean and the variance implies that two binomial arrays that

14. The Lindberg-Feller condition for the central limit theorem is that the $Z_j(\Delta)$'s have finite mean and variance; to apply this to arrays, where the probability law changes with Δ , it is sufficient that these bounds hold uniformly in Δ .

converge to diffusions and have common outcomes must have the same volatility.

PROPOSITION 2. Suppose that the signals are sums of i.i.d. binomials $Z_j(\Delta)$ where the common outcomes are $x(\Delta) > y(\Delta)$, and that the probability of $x(\Delta)$ under action $i = +1, -1$ is $\alpha_i(\Delta)$ with $\lim_{\Delta \rightarrow 0} \alpha_i(\Delta) = \alpha_i$, $0 < \alpha_i < 1$. If under each action i the sums $\sum_{j=1}^{\lfloor t/\Delta \rfloor} Z_j(\Delta)$ converge to a diffusion with drifts μ_i and volatilities σ_i^2 as $\Delta \rightarrow 0$, then $\sigma_{+1} = \sigma_{-1}$.

Proof. In Appendix II. ■

The equal-volatility case is important because in this case the equilibrium with respect to the limiting diffusions must be trivial. This does not necessarily imply that if triangular arrays converge to limit diffusions with the same volatility the limit equilibria are trivial. Indeed, the next section gives an example where the signals converge to equal-volatility diffusions and yet there is an efficient limit equilibrium. However, the limit equilibrium must be trivial whenever the variances of the aggregate signals converge to the common limit at a sufficiently high rate, and if there is enough aggregation of signals so that we can apply an appropriate variant of the central limit theorem. The next result shows that the assumption of binomial signals plus some technical assumptions does lead to this result, where the “enough aggregation of signals” condition is that $k(\Delta) = \tau(\Delta)/\Delta$ grows quickly enough so that $\lim_{\Delta \rightarrow 0} \tau(\Delta) \exp(k(\Delta)^{2/7}) \rightarrow \infty$.

PROPOSITION 3. Suppose that

- (i) $\lim_{\Delta \rightarrow 0} \tau(\Delta) \exp(k(\Delta)^{2/7}) \rightarrow \infty$
- (ii) the signals are sums of i.i.d. binomials $Z_j(\Delta)$, 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) under each action i , $\sum_{j=1}^{\lfloor t/\Delta \rfloor} Z_j(\Delta)$ converge to nondegenerate diffusions with drifts μ_i and volatilities σ_i^2 .

Then all limit equilibria are trivial.

Proof Sketch. The proof is in Appendix III; here is a rough outline: The idea is to use the central limit theorem and a continuity argument to extend our earlier result that the limit equilibrium is trivial when the underlying signal structure is a pair of diffusions with equal volatility. If we could restrict the analysis to strategies

where the cutoff for punishment was fixed relative to the standard error of the signal the proof would be straightforward, but we have to also consider 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 (1971) 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 $\lim_{n \rightarrow \infty} x_n = \infty$ at a rate slowly enough so that $n^{-1/6}x_n \rightarrow 0$. Feller’s result does not directly apply to our setting, because the distribution of the underlying variables changes with the period length; we report the extension of his result to our case of triangular arrays in Fudenberg and Levine (2007b).

The proof also uses a sharpening of Proposition 2, reported in Lemma A.1: not only do the two binomial arrays converge to diffusions with a common volatility, but also the variances of the two signal processes converge to equality sufficiently quickly for our argument to be valid. ■

Proposition 3 assumes that $\lim_{\Delta \rightarrow 0} \tau(\Delta) \exp(k(\Delta)^{2/7}) = 0$. Without this condition, we cannot use the normal approximation, so we do not have a general result. However, one important special case is the binomial construction of diffusions found in many textbooks, such as Stokey (2008). Here $x(\Delta) = -y(\Delta) = \sigma \Delta^{1/2}$, $\alpha(\Delta) = .5 + .5\mu \Delta^{1/2}/\sigma$, and the triangular array converges to a diffusion with drift μ and volatility σ^2 . In this case it is clear how the assumption that the support of the binomials is the same under both actions forces the two diffusions to have the same volatilities, and moreover we can determine what happens when k is small.

PROPOSITION 4. If the player’s signals are as in the standard binomial construction of diffusions and $\lim_{\Delta \rightarrow 0} k(\Delta)\tau(\Delta) = 0$, then there are only trivial limit equilibria.

Proof. Assume that the drift μ_{-1} under action -1 exceeds the drift μ_{+1} under action $+1$. (The case $\mu_{-1} < \mu_{+1}$ is symmetric and it is obvious that there is only a trivial limit equilibrium when the drifts are equal.) For any finite number k of signals, the largest possible value of $\rho = (q - p)/p$ is obtained with strategies that punish only if the sum equals kh , so that every realization was $+h$. Recall that we must have the case of equal volatilities $\sigma_{+1} = \sigma_{-1}$. We compute

$$\rho = \frac{q - p}{p} = \frac{q}{p} - 1 = \left(\frac{0.5 + 0.5\mu_{-1}(\tau/k)^{1/2}/\sigma_{+1}}{0.5 + 0.5\mu_{+1}(\tau/k)^{1/2}/\sigma_{+1}} \right)^k - 1.$$

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

$$\begin{aligned} & \lim_{\Delta \rightarrow 0} k(\Delta) \log(0.5 + 0.5\mu_{-1}[\tau(\Delta)/k(\Delta)]^{1/2}/\sigma_{+1}) \\ & \quad - k(\Delta) \log(0.5 + 0.5\mu_{+1}[\tau(\Delta)/k(\Delta)]^{1/2}/\sigma_{+1}) \\ &= \lim_{\Delta \rightarrow 0} k(\Delta)(\mu_{-1}[\tau(\Delta)/k(\Delta)]^{1/2}/\sigma_{+1} - \mu_{+1}[\tau(\Delta)/k(\Delta)]^{1/2}/\sigma_{+1}) \\ &= \lim_{\Delta \rightarrow 0} (k(\Delta)\tau(\Delta))^{1/2}(\mu_{-1} - \mu_{+1})/\sigma_{+1} = 0, \end{aligned}$$

where the last equality follows from $\lim_{\Delta \rightarrow 0} k(\Delta)\tau(\Delta) = 0$. Thus by Corollary 1b, there is only a trivial limit equilibrium. ■

Proposition 4's hypothesis that $\lim_{\Delta \rightarrow 0} k(\Delta)\tau(\Delta) = 0$ overlaps with Proposition 3's hypothesis that $\lim_{\Delta \rightarrow 0} \tau(\Delta) \exp(k(\Delta)^{2/7}) \rightarrow \infty$, so combining the two results gives a complete characterization of the limit of standard binomials:

COROLLARY 3. If the player's signals are as in the standard binomial construction of diffusions, then there are only trivial limit equilibria.

To relate this result to the previous one, and to our earlier general analysis, note that when $\lim_{\Delta \rightarrow 0} k(\Delta)\tau(\Delta) > 0$, the sequence of strategies "only punish if every outcome was $+h$ " has a limiting value of ρ that is nonzero. However, along this sequence we have $q/\tau \rightarrow 0$, so as Corollary 2a shows, this is no help.

V. TRINOMIAL INFORMATIONAL LIMITS

Although 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 do not correspond to the limit of the equilibria when the signal is the aggregate of many small events. This suggests that the continuous time game is “too small” to capture all of the more general ways that signal processes can be approximated by diffusions.

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, and use this family to explore 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 signal processes will be indexed by a free parameter γ that is not determined by the limit diffusions. For any $\gamma \geq 1$, we set $\bar{\gamma} = \gamma \max(\sigma_{+1}^2, \sigma_{-1}^2)$. 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 depend 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 nonnegative and less than 1), and the probability of outcome $+h$ is

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

A simple example may help put this in perspective. The sign of x and the size of the step are simply normalizations, so that the normalized signals converge to a diffusion, so we may think of the underlying data as “0 sales” corresponding to $x = -h(\Delta)$, “a single sale” corresponding to $x = 0$, and “a double sale” corresponding to $x = +h(\Delta)$. Let us focus on the bad-news case, where $\sigma_{-1} > \sigma_{+1}$, and take $\sigma_{-1} = 2, \sigma_{+1} = 1$ and $\gamma = 2$. Then $\bar{\gamma} = 4, \alpha_{+1} = 3/4, \alpha_{-1} = 1/2$, so that a sale is more likely if action $+1$ is taken. Ignoring the “small noise term” of order $\Delta^{1/2}$, the probability of no sale or a double sale when $a = +1$ is $1/8$, and

when $a = -1$ the probability of no sale or a double sale is $1/4$. That is, action $+1$ increases the likelihood of a single sale at the expense of both no sales and double sales.

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 action i . Then the expected values of the trinomial distribution described above are

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

and the variances are

$$\begin{aligned} E_i[Z_j(\Delta)^2] - (E_i[Z_j(\Delta)])^2 &= (1 - \alpha_i)h^2 - \mu_i^2 \Delta^2 \\ &= (1 - \alpha_i)\bar{\gamma} \Delta - \mu_i^2 \Delta^2 = \sigma_i^2 \Delta - \mu_i^2 \Delta^2. \end{aligned}$$

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.

V.A. *Bad-News Case*: $\sigma_{-1}^2 > \sigma_{+1}^2$

In the bad-news case, we can show that if the ratio of volatilities is sufficiently large then the limit equilibrium is nontrivial, 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 simplify the presentation, we restrict attention to the case where both diffusions have zero drift, but this is not important for the results.

To begin, consider $\tau(\Delta) = \Delta$. Because 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 zero when the signal is 0. (The likelihood ratio for punishing on 0 is less than one, 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 a 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 1c these strategies support a nontrivial 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: Because 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 because the punishment probability is also independent of Δ we again have a nontrivial equilibrium. Moreover, because the maximal value of ρ (consistent with nonzero punishment probability) has increased, aggregating two signals allows a better limit equilibrium payoff.

Now consider the case $\tau = \Delta^{1/2}$ such that $k(\Delta) = \Delta^{-1/2} \rightarrow \infty$ as $\Delta \rightarrow 0$. Here we make use of the following more general result:

PROPOSITION 5. 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 2c. See Appendix V for details. ■

COROLLARY 4. There are sequences of information structures with efficient limit equilibrium where players observe the sum of discrete events, and these sums converge to a pair of diffusions with the same volatilities.

Proof. The idea is to use a diagonalization argument to obtain a sequence of trinomials where the ratio of variances goes to 1 sufficiently slowly so that there is an efficient limit. To do this, consider a sequence $\{\sigma_{-1n}^2/\sigma_{+1n}^2\} \downarrow 1$ and to each $(\sigma_{-1n}^2, \sigma_{+1n}^2)$ associate a trinomial signal structure distribution $\{S_{n\Delta}\}_\Delta$ as defined above, so that the sum of the public signal under information structure $S_{n\Delta}$ converges to a pair of diffusions with drift 0 and volatilities $\sigma_{-1}^2, \sigma_{+1}^2$. Let $G_{n\Delta}$ be the game with event frequency $1/\Delta$, information structure $S_{n\Delta}$, and period length $\tau(\Delta) = \Delta^{1/2}$. From Proposition 5, for any sequence of strictly positive $\varepsilon_n \rightarrow 0$,

there is a sequence of PPE $P_{n\Delta}$ for the $G_{n\Delta}$ with limit payoff $\bar{u} - \varepsilon_n$ as $\Delta \rightarrow 0$.

Now we diagonalize: For each j , pick Δ_j so that $P_{j\Delta_j}$ has payoff at least $\bar{u} - 2\varepsilon_j$; let $G_j \equiv G_{j\Delta_j}$ be the corresponding game; then the sequence of games $\{G_j\}$ has a sequence of PPE $P_{j\Delta_j}$ with payoffs converging to $\bar{u}$. ■

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.

V.B. Good-News Case $\sigma_{-1}^2 < \sigma_{+1}^2$

Once again, we simplify by setting the drifts equal to 0, and begin with the case $\tau(\Delta) = \Delta$. The optimal equilibrium with this signal structure prescribes punishment with positive probability when $Z_j = 0$ and zero probability 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 , even though 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 1 a necessary condition for a nontrivial 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 zero under action $+1$, 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 one regardless of the choice of action.

An argument similar to that of the previous section shows that when there is a nontrivial 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 $\rho = (q/p) - 1$, 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, whereas 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 V 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 in the case where γ is very large and the short-run gain to deviating, g , is very small.

Now consider the case $\tau(\Delta) = \Delta^{1/2}$, where 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 nontrivial limit equilibrium for the diffusion) and γ large we can construct examples where there is only the trivial equilibrium when $\tau = \Delta$ and a nontrivial limit when players aggregate infinitely many signals, whereas by specifying γ near 1, and σ_{-1} near σ_{+1} , we have examples with a nontrivial 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, necessarily any sequence $\{(\sigma_{+1n}^2, \sigma_{-1n}^2)\} \rightarrow_n (\sigma_{+1}^2, \sigma_{-1}^2)$ will eventually lie in the interior of this range as well. We conjecture that we could thus use the large-deviations arguments of Appendix III 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.

V.C. 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 IV. As there, we suppose that the bad action has a higher mean. With equal variances, $\alpha_{-1} = \alpha_{+1}$, so $\alpha = (\gamma - 1)/\gamma$; the standard binomial case corresponds to $\gamma = 1$ and $\alpha = 0$.

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

$$\begin{aligned} \frac{q}{p} &= \frac{\gamma^{-1/2}(1-\alpha) + \mu_{-1}\Delta^{1/2}}{\gamma^{-1/2}(1-\alpha) + \mu_{+1}\Delta^{1/2}} = \frac{\sigma\gamma^{-1/2} + \mu_{-1}\Delta^{1/2}}{\sigma\gamma^{-1/2} + \mu_{+1}\Delta^{1/2}} \\ &= \frac{\sigma + \mu_{-1}\Delta^{1/2}\gamma^{1/2}}{\sigma + \mu_{+1}\Delta^{1/2}\gamma^{1/2}}. \end{aligned}$$

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. As in the proof of Proposition 4, this implies that 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 because 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.

VI. CONCLUSIONS

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 in 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 cumulative distribution, they get the first-best limit payoff when volatilities are different and $\tau/\Delta \rightarrow \infty$ regardless of the ratio of the volatilities. 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 be 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 firms do 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 to hours worked, but the disaggregated data are quite closely held.

APPENDIX I: PROOF OF PROPOSITION 1

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

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

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 1 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 +1 with positive probability. Fix such an equilibrium, and suppose that the short-run player plays Out with positive probability in the first period. Because 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 a PPE in which LR has a higher payoff, 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 $v \geq w(z) \geq \underline{u}$ because any payoff between the best and worst can be attained with public randomization. ■

APPENDIX II: BINOMIAL CONVERGENCE TO DIFFUSIONS

Here we prove some results about the convergence of binomials to diffusions needed in proving Proposition 3 in Appendix III.

PROPOSITION A.2. Suppose that the signals are sums of i.i.d. binomials $Z_j(\Delta)$, where the common outcomes are $x(\Delta) > y(\Delta)$, and that the probability of $x(\Delta)$ under action $i = +1, -1$ is $\alpha_i(\Delta)$ with $\lim_{\Delta \rightarrow 0} \alpha_i(\Delta) = \alpha_i$, $0 < \alpha_i < 1$. If under each action i the sums $\sum_{j=1}^{\lfloor t/\Delta \rfloor} Z_j(\Delta)$ converge to a diffusion with drift μ_i and volatilities σ_i^2 as $\Delta \rightarrow 0$, then $\sigma_{+1} = \sigma_{-1}$.

Proof. First we examine what it means for the sum of the $Z_j(\Delta)$ to converge to a diffusion under action +1. 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 parameterization, we can calculate that $E_{+1}Z_j(\Delta) = \mu_{+1}^\Delta \Delta$ and $\text{var}_{+1}Z_j(\Delta) = (\sigma_{+1}^\Delta)^2 \Delta$.

If the limit process is a diffusion, then its position at t has the normal distribution with mean μ_{+1} and 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 IV: Here we have $x = -y = \sigma_{+1} \Delta^{1/2}$ and $\alpha_{+1} = (1 + \mu_{+1} \Delta^{1/2} / \sigma_{+1}) / 2$, so $\mu_{+1}^\Delta = \mu_{+1}$ 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_{-1} and variance σ_{-1}^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

$$\begin{aligned} x(\Delta) &= \mu_{+1}^\Delta \Delta + \sigma_{+1}^\Delta \Delta^{1/2} \sqrt{\frac{1 - \alpha_{+1}(\Delta)}{\alpha_{+1}(\Delta)}} = \mu_{-1}^\Delta \Delta + \sigma_{-1}^\Delta \Delta^{1/2} \sqrt{\frac{1 - \alpha_{-1}(\Delta)}{\alpha_{-1}(\Delta)}} \\ y(\Delta) &= \mu_{+1}^\Delta \Delta - \sigma_{+1}^\Delta \Delta^{1/2} \sqrt{\frac{\alpha_{+1}(\Delta)}{1 - \alpha_{+1}(\Delta)}} = \mu_{-1}^\Delta \Delta - \sigma_{-1}^\Delta \Delta^{1/2} \sqrt{\frac{\alpha_{-1}(\Delta)}{1 - \alpha_{-1}(\Delta)}}. \end{aligned} \quad (3)$$

We now solve this system to see the possible values of $\sigma_{-1}^\Delta, \alpha_{-1}(\Delta)$ as a function of $\sigma_{+1}^\Delta, \mu_{+1}^\Delta, \mu_{-1}^\Delta, \alpha_{+1}(\Delta)$:

$$\begin{aligned} (\mu_{+1}^\Delta - \mu_{-1}^\Delta) \Delta + \sigma_{+1}^\Delta \sqrt{\frac{1 - \alpha_{+1}(\Delta)}{\alpha_{+1}(\Delta)}} &= \sigma_{-1}^\Delta \sqrt{\frac{1 - \alpha_{-1}(\Delta)}{\alpha_{-1}(\Delta)}} \\ (\mu_{+1}^\Delta - \mu_{-1}^\Delta) \Delta - \sigma_{+1}^\Delta \sqrt{\frac{\alpha_{+1}(\Delta)}{1 - \alpha_{+1}(\Delta)}} &= -\sigma_{-1}^\Delta \sqrt{\frac{\alpha_{-1}(\Delta)}{1 - \alpha_{-1}(\Delta)}}. \end{aligned}$$

Divide the two equations to eliminate σ_{-1}^Δ , solve for $\alpha_{-1}(\Delta)$, and plug back into the second equation to find

$$\begin{aligned} \sigma_{-1}^\Delta &= \left((\mu_{+1}^\Delta - \mu_{-1}^\Delta) \Delta + \sigma_{+1}^\Delta \sqrt{\frac{1 - \alpha_{+1}(\Delta)}{\alpha_{+1}(\Delta)}} \right) \\ &\quad \times \frac{\sigma_{+1}^\Delta \sqrt{\frac{\alpha_{+1}(\Delta)}{1 - \alpha_{+1}(\Delta)}} - (\mu_{+1}^\Delta - \mu_{-1}^\Delta) \Delta}{\sigma_{+1}^\Delta \sqrt{\frac{1 - \alpha_{+1}(\Delta)}{\alpha_{+1}(\Delta)}} + (\mu_{+1}^\Delta - \mu_{-1}^\Delta) \Delta}. \end{aligned}$$

Because $\lim_{\Delta \rightarrow 0} \alpha_i(\Delta) = \alpha_i$, $0 < \alpha_i < 1$, and it follows that $\sigma_{-1}^\Delta \rightarrow \sigma_{+1}$. ■

LEMMA A.1. Suppose that the signals are sums of i.i.d. binomials $Z_j(\Delta)$ where the common outcomes are $x(\Delta) > y(\Delta)$, and that the probability of $x(\Delta)$ under action $i = +1, -1$ is $\alpha_i(\Delta)$ with $\lim_{\Delta \rightarrow 0} \alpha_i(\Delta) = \alpha_i$, $0 < \alpha_i < 1$. If under each action i the sums $\sum_{j=1}^{\lfloor t/\Delta \rfloor} Z_j(\Delta)$ converge to a diffusion with drift μ_i and volatilities σ_i^2 as $\Delta \rightarrow 0$, then

$$\lim_{\Delta \rightarrow 0} \frac{|(\sigma_{+1}^\Delta)^2 - (\sigma_{-1}^\Delta)^2|}{\Delta^{1/5}} = 0.$$

Proof. Because

$$\frac{|(\sigma_{+1}^\Delta)^2 - (\sigma_{-1}^\Delta)^2|}{\Delta^{1/5}} = \frac{|\sigma_{+1}^\Delta - \sigma_{-1}^\Delta| |\sigma_{+1}^\Delta + \sigma_{-1}^\Delta|}{\Delta^{1/5}}$$

and $|\sigma_{+1}^\Delta + \sigma_{-1}^\Delta| \rightarrow 2\sigma_{+1}$, this is the same as $|\sigma_{+1}^\Delta - \sigma_{-1}^\Delta|/\Delta^{1/5} \rightarrow 0$, so

$$\begin{aligned} & \frac{|\sigma_{+1}^\Delta - \sigma_{-1}^\Delta|}{\Delta^{1/5}} \\ &= \left| \sigma_{+1}^\Delta \Delta^{-1/5} - \left((\mu_{+1}^\Delta - \mu_{-1}^\Delta) \Delta^{3/10} + \Delta^{-1/5} \sigma_{+1}^\Delta \sqrt{\frac{1 - \alpha_{+1}(\Delta)}{\alpha_{+1}(\Delta)}} \right) \right| \\ & \quad \times \sqrt{\frac{\sigma_{+1}^\Delta \sqrt{\frac{\alpha_{+1}(\Delta)}{1 - \alpha_{+1}(\Delta)}} - (\mu_{+1}^\Delta - \mu_{-1}^\Delta) \Delta^{1/2}}{\sigma_{+1}^\Delta \sqrt{\frac{1 - \alpha_{+1}(\Delta)}{\alpha_{+1}(\Delta)}} + (\mu_{+1}^\Delta - \mu_{-1}^\Delta) \Delta^{1/2}}}. \end{aligned}$$

Algebraic manipulation leads to

$$\begin{aligned} \lim_{\Delta \rightarrow 0} \frac{|\sigma_{+1}^\Delta - \sigma_{-1}^\Delta|}{\Delta^{1/5}} &= \lim_{\Delta \rightarrow 0} \sigma_{+1}^\Delta \frac{|\mu_{+1}^\Delta - \mu_{-1}^\Delta| \Delta^{1/10} |2\alpha_{+1}(\Delta) - 1|}{\alpha_{+1} \sqrt{\sigma_{+1}^\Delta \sqrt{\frac{1 - \alpha_{+1}(\Delta)}{\alpha_{+1}(\Delta)}}} \left(2\sqrt{\sigma_{+1}^\Delta \sqrt{\frac{1 - \alpha_{+1}(\Delta)}{\alpha_{+1}(\Delta)}}} \right)} \\ &= \frac{|\mu_{+1} - \mu_{-1}|}{2\sigma_{+1}} \lim_{\Delta \rightarrow 0} \Delta^{1/10} \frac{|2\alpha_{+1}(\Delta) - 1|}{\sqrt{\alpha_{+1}(1 - \alpha_{+1}(\Delta))}}. \end{aligned}$$

Because $\lim_{\Delta \rightarrow 0} \alpha_i(\Delta) = \alpha_i$, $0 < \alpha_i < 1$, the result follows. ■

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 nonstandard binomial array with the probabilities of the two steps not converging to $1/2$, it is convenient to set $\mu_{+1}^\Delta = \mu_{+1}$, $\sigma_{+1}^\Delta = \sigma_{+1}$, $\alpha_{+1}(\Delta) = \alpha_{+1}$. Using our alternative parameterization 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 0 and volatility 1.

APPENDIX III: PROOF OF PROPOSITION 3

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 $k(\Delta) = \tau(\Delta)/\Delta$. Note that condition (i) of Proposition 3 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. Recall that when players observe the state of a diffusion process at discrete intervals, the signals are normally distributed; let Φ , ϕ respectively denote the c.d.f. and density of the standard normal distribution.

Fact A.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 nontrivial 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 3: In this case we have

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

Fact A.2 [Large Deviations Theorem for Triangular Arrays, from Fudenberg and Levine (2007b)]. Suppose that for each n there is a sequence Z_i^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*} , whereas 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 $\sup_{0 \leq \zeta \leq \bar{s}} |\psi_n'''(\zeta)|, |\psi_n''''(\zeta)| |\psi_n''''(\zeta) \zeta^2| < B$
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 $\tau^k \rightarrow 0$, with $\Delta^k = \tau^k/k$ and $Z_j^k = Z_j(\Delta^k)$.¹⁵ As in the proof of Proposition 2, we define new parameters $\mu_i^k = E_i Z_j^k, (\sigma_i^k)^2 = \text{var}_i Z_j^k / \Delta^k$.

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

$$\tilde{Z}_{ij}^k = \frac{Z_j^k - \mu_i^k \Delta^k}{\sqrt{\Delta^k}}$$

so that $E \tilde{Z}_{ij}^k = 0, \text{var} \tilde{Z}_{ij}^k = (\sigma_i^k)^2 \rightarrow \sigma_i^2$ 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 nonzero constant.

LEMMA A.2. Consider two i.i.d. binomials $Z_j(\Delta)$ with common outcomes $x(\Delta) > y(\Delta)$, where the probability of $x(\Delta)$ under action $i = +1, -1$ is $\alpha_i(\Delta)$ with $\lim_{\Delta \rightarrow 0} \alpha_i(\Delta) = \alpha_i, 0 < \alpha_i < 1$. If under each action i the sums of i.i.d. binomials $\sum_{j=1}^{\lfloor t/\Delta \rfloor} Z_j(\Delta)$ converges to a diffusion with drift μ_i and volatilities σ_i^2 as $\Delta \rightarrow 0$, then the reparameterized binomials satisfy conditions 1 and 2 of the large deviations theorem for both $i = -1, +1$.

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

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

$$\begin{aligned} \psi_k(\zeta) = & \log \left(\alpha_{+1}(\Delta^k) \exp \left(\zeta \sigma_i^k \left(\frac{1 - \alpha_{+1}(\Delta^k)}{\alpha_{+1}(\Delta^k)} \right)^{1/2} \right) \right. \\ & \left. + (1 - \alpha_{+1}(\Delta^k)) \exp \left(-\zeta \sigma_i^k \left(\frac{\alpha_{+1}(\Delta^k)}{1 - \alpha_{+1}(\Delta^k)} \right) \right) \right). \end{aligned}$$

Let

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

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. ■

We turn now to the main proof. The idea is to show that if there were strategies that led to a nontrivial limit equilibrium in the binomial case, we could construct a nontrivial limit equilibrium when players observe the position of a diffusion. 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, whereas Fact A.1 supposes that the variances are equal before the limit is reached. Lemma A.1 adds a condition on the rate of convergence that enables us to extend Fact A.1 to the case where the variances are different before the limit is reached. Moreover, although 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 nontrivial equilibrium, whereas if it grows at $k^{1/6}$ the normal approximation is so good that we can make use of Fact A.1.

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

LEMMA A.3. 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'Hôpital's rule that if $x \rightarrow \infty$ then $(1 - \Phi(x))/\phi(x) \rightarrow 0$. Using that fact, we may again apply L'Hôpital's rule to see that $\lim_{x \rightarrow \infty} x(1 - \Phi(x))/\phi(x) = 1$. Define

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

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. When the signal is the sum of binomials with common support, and action -1 has a higher mean, the monotone likelihood ratio property is satisfied for the pair of signals.

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

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

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

LEMMA A.5. If $\liminf_{k \rightarrow \infty} \bar{\zeta}^k k^{-1/6} > 0$ and $\tau^k \exp(k^{2/7}) \rightarrow 0$ then $p/\tau \rightarrow 0, q/\tau \rightarrow 0$.

Proof. It suffices to prove $q/\tau \rightarrow 0$ because $q \geq p$. To compute q , we need to consider the distribution of $\bar{z}_{+1}^k$ when the action taken is -1 . This does not have zero mean or unit variance, so we renormalize, defining $\bar{z}_{+1}^k \frac{\sigma_{+1}^k}{\sigma_{-1}^k} \bar{z}_{+1}^k - (\mu_{-1}^k - \mu_{+1}^k) \Delta^k k$, which has zero mean and unit variance when the action taken is -1 . Denote the c.d.f. 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, because $\lim \bar{\zeta}^k k^{-1/6} > 0$, $|\sigma_{+1}^k - \sigma_{-1}^k| \rightarrow 0$, and μ_i^k is bounded, we see that $\lim_{k \rightarrow \infty} x^k k^{-1/6} > 0$. Hence we set $\bar{x}^k = k^{1/7} \sqrt{2\pi}$, and because $k^{-1/6} \bar{x}^k \rightarrow 0$, we know that $\bar{x}^k \leq x^k$ for large enough k , and that the conditions 3 and 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'Hôpital's rule to see that $\frac{1 - \Phi(\bar{x}^k)}{\phi(\bar{x}^k)} \rightarrow 0$, so for large enough k we have

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

Because $\bar{x}^k = k^{1/7} \sqrt{2\pi}$, we have that $p/\tau \leq C(\tau \exp(k^{2/7})) \rightarrow 0$. ■

Now we can prove Proposition 3, which we restate for convenience.

PROPOSITION 3. Suppose

- (i) $\lim_{\Delta \rightarrow 0} \tau(\Delta) \exp(k(\Delta)^{2/7}) \rightarrow \infty$
- (ii) the signals are sums of i.i.d. binomials $Z_j(\Delta)$ 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) under each action i , $\sum_{j=1}^{\lfloor t/\Delta \rfloor} Z_j(\Delta)$ converges to a nondegenerate diffusion with drifts μ_i and volatilities σ_i^2 .

Then all limit equilibria are trivial.

Proof. By Lemma A.4, the signals satisfy the MLRP, so we can restrict attention to strategies that punish when the observed signal exceeds some cutoff. By Lemma A.5, if there is a nontrivial limit, we may assume that the cutoff satisfies $\lim_{k \rightarrow \infty} \bar{\xi}^k k^{-1/6} = 0$. By Lemma A.2 and Fact A.2, this means that we may compute $p/\tau, q/\tau, \rho$ asymptotically using normal distributions. From Lemma A.1,

$$\lim_{k \rightarrow \infty} \frac{|(\sigma_{+1}^k)^2 - (\sigma_{-1}^k)^2|}{(\tau/k)^{1/5}} = 0,$$

so that $\lim_{k \rightarrow \infty} k^{1/5}|(\sigma_{+1}^k)^2 - (\sigma_{-1}^k)^2| = 0$ because $\lim_{k \rightarrow \infty} \bar{\xi}^k k^{-1/6} = 0$, $\lim_{k \rightarrow \infty} \bar{\xi}^k k^{-1/5} = 0$ and so $\lim_{k \rightarrow \infty} \bar{\xi}^k |(\sigma_{+1}^k)^2 - (\sigma_{-1}^k)^2| = 0$. Consequently Lemma A.3 applies, so that we may assume that the normals have the same variance, implying a nontrivial limit in that case. This contradicts Fact A.1. ■

APPENDIX IV: 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 two signals are aggregated.

Punishing when the sum of the signals is -2 or $+2$ 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 $i = 1, -1$,

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

This is minimized at $\alpha_i = 1/3$, where it has value $1/3$. Next $\Pr_i(\sum_{j=1,2} Z_j = 1) = \alpha_i(1 - \alpha_i) = \alpha_i - \alpha_i^2$. 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 $C = \alpha_{-1}/\alpha_{+1}$ 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 $\frac{3\alpha_{+1}^2 + 1 - 2\alpha_{-1}}{3\alpha_{+1}^2 + 1 - 2\alpha_{+1}} < \alpha_{-1}/\alpha_{+1}$.

Note that $\alpha_{-1} > \alpha_{+1}$ because we are in the good-news case. Observing that all the expressions are nonnegative, 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 sides 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 $f(\alpha_{-1}) > f(\alpha_{+1})$; because $\alpha_{+1}, \alpha_{-1} \rightarrow 1$ 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 nontrivial 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 nontrivial limit.

APPENDIX V: PROOF OF PROPOSITION 5

PROPOSITION 5. 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 $z \equiv \sum_{j=1}^{\lfloor t/\Delta \rfloor} Z_j(\Delta)$ exceeds a threshold z^* or equivalently when the absolute value of $\zeta(\tau) = z/\sigma_{+1}\tau^{1/2}$ exceeds $\zeta^* = z^*/\sigma_{+1}\tau^{1/2}$. 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 $p(\tau)$ a constant independent of τ 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 players observe the position of the limit diffusions and use strategies with a fixed normalized cutoff 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. Because we have assumed that $\lim_{\Delta \rightarrow 0} \tau(\Delta)/\Delta = \infty$, 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. ■

DEPARTMENT OF ECONOMICS, HARVARD UNIVERSITY

DEPARTMENT OF ECONOMICS, WASHINGTON UNIVERSITY IN ST. LOUIS

REFERENCES

- Abreu, D., P. Milgrom, and D. Pearce, "Information and Timing in Repeated Partnerships," *Econometrica*, 59 (1991), 1713–1733.
- Abreu, D., D. Pearce, and E. Stachetti, "Towards a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58 (1990), 1041–1064.
- Celentani, M., D. Levine, and C. Martinelli, "Punishing, Investigating or Auditing?" mimeo, Washington University of St. Louis, 2007.
- Faingold, E., and Y. Sannikov, "Equilibrium Degeneracy and Reputation Effects," mimeo, Yale University, 2007.
- Feller, W., *An Introduction to Probability Theory and Its Applications*, Volume II (New York, NY: Wiley, 1971).
- Fudenberg, D., D. M. Kreps, and E. Maskin, "Repeated Games with Long-Run and Short-Run Players," *Review of Economic Studies*, 57 (1990), 555–573.
- Fudenberg, D., and D. K. Levine, "Subgame-Perfect Equilibria of Finite and Infinite Horizon Games," *Journal of Economic Theory*, 31 (1983), 251–268.
- , "Efficiency and Observability with Long-Run and Short-Run Players," *Journal of Economic Theory*, 62 (1994), 103–135.
- , "Continuous Time Limits of Repeated Games with Imperfect Public Monitoring," *Review of Economic Dynamics*, 10 (2007a), 173–192.
- , "A Large Deviation Theorem for Triangular Arrays," mimeo, Harvard University, 2007b.
- Fudenberg, D., and J. Tirole, "Preemption and Rent Equalization in the Adoption of New Technology," *Review of Economic Studies*, 52 (1985), 383–402.
- , *Game Theory* (Cambridge, MA: MIT Press, 1991).
- Hellwig, M., and K. Schmidt, "Discrete-Time Approximations of the Holmström-Milgrom Brownian-Motion Model of Intertemporal Incentive Provision," *Econometrica*, 70 (2002), 2225–2264.
- Klein, B., and K. Leffler, "The Role of Market Forces in Assuring Contractual Performance," *Journal of Political Economy*, 89 (1981), 615–641.
- Sannikov, Y., "Games with Imperfectly Observed Actions in Continuous Time," *Econometrica*, 75 (2007a), 1285–1329.

- , “Repeated Games with Imperfect Monitoring: Connection between Discrete and Continuous Time,” mimeo, 2007b.
- Sannikov, Y., and A. Skrzypacz, “The Impossibility of Collusion under Imperfect Monitoring with Flexible Production,” *American Economic Review*, 95 (2007), 314–331.
- , “The Role of Information in Repeated Games with Frequent Actions,” mimeo, Stanford University, 2008.
- Simon, L., and M. Stinchcombe, “Extensive Form Games in Continuous Time: Pure Strategies,” *Econometrica*, 57 (1989), 1171–1214.
- Stokey, N., “Brownian Models in Economics,” mimeo, 2008.