Abstract

We study a coordination game with randomly changing payoffs and small frictions in changing actions. Using only backwards induction, we find that players must coordinate on the risk dominant equilibrium. More precisely, a continuum of fully rational players are randomly matched to play a symmetric 2×2 game. The payoff matrix changes according to a random walk. Players observe these payoffs and the population distribution of actions as they evolve. The game has frictions: opportunities to change strategies arrive from independent random processes, so that the players are locked into their actions for some time. As the frictions disappear, each player ignores what the others are doing and switches at her first opportunity to the risk dominant equilibrium. History dependence emerges in some cases when frictions remain positive.

J.E.L. No.: C73. **Field:** Game Theory.

Fast Equilibrium Selection
by Rational Players
Living in a Changing World

Krzysztof Burdzy,

David M. Frankel,

and Ady Pauzner[1]

First version: July 1996.

This version: June 27, 1999

1 Introduction

Games with multiple strict Nash equilibria present a major challenge for game theory. Most equilibrium refinements do not select a unique equilibrium in such games, and those that do are not unanimous in their predictions. These difficulties have motivated an interest in explicitly modelling dynamic processes by which play might evolve. By tracing a gradual process in which players adjust their actions in response to past play, one could hope to identify conditions that govern whether players will eventually coordinate on an equilibrium and, if so, on which one.

Most dynamic models have focused on the evolution of play in a fixed world. We assume instead that the world changes over time. In many of the concrete examples of multiple equilibria, such as the choice between technological standards or the economy's coordination on high or low activity, the assumption that the world is changing is more realistic. The state of technological knowledge, oil prices, and weather conditions are only a few of many factors that change over time and affect the relative payoffs from different choices.

The model has a continuum of fully rational players who are randomly matched in pairs to play a symmetric 2x2 game with strategic complementarities. There are frictions: in each period, each player has only a small probability of being able to change her action. The payoff matrix of the game change every period, according to a random walk. We assume that, in the (perhaps very distant) future, the cumulative changes in payoffs have the potential to make either action strictly dominant.

We focus on the limit case in which the frictions become small, so that players' chances to change actions come more and more often.[2] The model becomes dominance-solvable. *The equilibrium that is risk-dominant at any given time must be played.* Moreover, convergence to that equilibrium occurs as fast as the frictions allow.

This contrasts with the predictions of most dynamic random matching models with players who are not fully rational (such as Kandori, Mailath and Rob [15] and Foster

---

[2]In taking this limit as frictions become small, we also shrink the time between periods, in order to retain the property that each player has only a small chance to change actions in a given period.

and Young [8]). In such models, when there is selection (e.g., as mutations go to zero), it takes a relatively weak form: the selected equilibrium is played almost all of the time in the ultra-long run. Our result also contrasts with selection in other matching models with rational players, where equilibrium refinements have typically been needed (e.g., Matsui and Matsuyama [18]).

The result is driven by a contagion argument. Because of strategic complementarities in the static game, players have an incentive to pick the same action that their opponents are likely to choose. Hence, if a player picks her action when the game's parameters are near a region where an action (call it a) is strictly dominant, she has an incentive to pick that action as well. This is because while she is locked into her action, random payoff changes are likely to move the game into the region of strict dominance, leading some of the player's opponents to pick a. So a must be played when sufficiently close to the dominance region. But then a must be played when sufficiently close to this *adjacent* region, and so on.

In the limit as frictions become small, this iterative process covers the entire parameter space, giving unique predictions of what players will do in any situation. This result depends crucially on the assumption that only a small proportion of players change actions in each period, so that two players matched at random will nearly always have chosen their actions seeing different payoffs. Without this, there might be multiple equilibria for some parameters, since a player's optimal action could depend on what others plan to do at the same payoffs.

We also consider the case in which frictions remain large but payoff shocks shrink to zero. If players are completely patient, the same result obtains: the risk-dominant equilibrium is selected. With impatience, players' strategies are still uniquely determined, but they now take into account the distribution of actions in the population. In particular, there can be *history dependence*: for a range of payoffs, if enough players are initially playing a given action, the rest of the population must follow.

This paper is related to papers by Carlsson and van Damme [6] and Matsui and Matsuyama [18]. The former analyzes a static game in which two players each receive a slightly noisy signal of the game's payoffs. Iterative strict dominance leads to the

selection of the risk-dominant equilibrium through a contagion argument. Matsui and Matsuyama study a dynamic model in which a large population of players are randomly matched to play a static game. There are multiple rational expectations equilibria; however, only the stationary state in which the risk-dominant equilibrium is played possesses certain stability properties. Our model is like that of Matsui and Matsuyama, but the payoffs in the static game change randomly over time. This gives rise to a unique outcome. Like Carlsson and van Damme, we prove this using a contagion argument. However, our contagion mechanism is fundamentally different since our model is dynamic and, as such, has state variables that have the potential to generate history dependence. The key to our argument is showing that the effect of the state variables is offset by players' future behavior. The relations among these three papers are discussed in detail in section 4.

Another related paper is Frankel and Pauzner [9]. They apply some of the techniques of this paper to a continuous-time macroeconomic model in which workers choose between two sectors, one of which has external increasing returns and is subject to productivity shocks. The paper also studies extensions to cases not analyzed in this paper. In particular, it identifies conditions under which there is a unique equilibrium of the dynamic game with nonvanishing frictions and shocks.

The rest of this paper is organized as follows. The model is presented in section 2. Section 3 presents the results and their intuitions, with proofs deferred to the appendix. In section 4 we review related literature. Section 5 concludes with a discussion of how the results depend on the various features of the model.

2   The Model

2.1   The Static Game

We consider a symmetric static game with two actions, R and L. Payoffs depend on a random parameter B_t that changes over time: if a player playing a meets a player playing a' at time t, her payoff in the static game is $u(a, a', B_t)$. Higher values of B_t raise the

3

relative payoff to playing R while lower values make L more desirable. More precisely, the *relative payoff to playing R* against the action a, $\Delta(a, B_t) = u(R, a, B_t) - u(L, a, B_t)$, is strictly increasing and continuously differentiable in B_t.

The game has strategic complementarities. That is, the relative payoff to playing R is higher when one's opponent is playing R: $\Delta(R, B_t) > \Delta(L, B_t)$.[3] The following table gives an example of the time-t payoff matrix.

	R	**L**
R	$3 + B_t, 3 + B_t$	$2 + B_t, 0$
L	$0, 2 + B_t$	$4, 4$

An action is *p-dominant* (Morris, Rob and Shin [19]) if it is a best response whenever the opponent is expected to play that action with probability at least p. We say that an action is *exactly p-dominant* if a player is *indifferent* when her opponent puts a weight of exactly p on that action.[4] (Equivalently, p is the smallest number for which the action is *p-dominant*.) Clearly, R is exactly p-dominant if and only if L is exactly $(1 - p)$-dominant.

This terminology permits a convenient rescaling. We denote by B^p the value of B_t at which R is exactly p-dominant in the static game:

$$pu(R, R, B^p) + (1 - p)u(R, L, B^p) = pu(L, R, B^p) + (1 - p)u(L, L, B^p) \qquad (1)$$

In the above game, for instance, when $B_t = 0$, R is exactly 0.4-dominant, so $0 = B^{0.4}$. Note that B^p is decreasing in p: if a player's opponent plays R with higher probability, the player will be willing to play R at *lower* values of B_t. An action is *risk-dominant* (Harsanyi and Selten [13]) if it is a best response when one's opponent is expected to play both actions with equal probabilities. In our terminology, R is risk-dominant whenever $B_t \geq B^{1/2}$ and L is whenever $B_t \leq B^{1/2}$.

---

[3]This assumption also ensures that only (R, R) and (L, L) can ever be pure Nash equilibria.

[4]For the purpose of the definition, we allow p to take values also outside the interval $[0, 1]$. For example, if R is exactly -0.2-dominant, it is strictly dominant.

4

2.2 The Dynamic Context

The game takes place in a sequence of periods at times $t = \tau, 2\tau, 3\tau, \ldots$. There is a continuum of players of measure 1. In each period, a subset of the players is randomly selected to revise their actions. Each player has an independent chance $k\tau$ of receiving an action revision opportunity.[5] We assume no aggregate uncertainty;[6] hence, in each period a measure $k\tau$ of players can change actions. When k is large, we say that frictions are small since a player's expected lock-in time $1/k$ is short.

After actions are revised, the payoff parameter changes randomly (from $B_{t-\tau}$ to B_t). We assume that it follows a random walk: it jumps by either $\mu\tau + \sigma\sqrt{\tau}$ or $\mu\tau - \sigma\sqrt{\tau}$, with equal probabilities. We use this formulation so that the real-time mean and variance do not depend on the parameter τ: for any times $t > t'$ that are both integer multiples of τ, the random variable $B_t - B_{t'}$ has mean $\mu(t - \hat{t})$ and variance $\sigma^2(t - \hat{t})$. We refer to μ as the trend of B and to σ^2 as the variance.

Following the change in the payoff parameter, another subset of the players is selected and randomly matched in pairs. Each pair collects the static game payoffs corresponding to the actions they are currently locked into and to the new value of the random parameter. Each player has an independent chance $m\tau$ of being matched; the measure of matched players is thus $m\tau$.

Let X_t be the proportion of players locked into R after actions are revised in period t. We refer to X_t as the "state of play," B_t as the "state of the world," and the pair (B_t, X_t) as the "state of the environment." The *public history* at time t is the evolution of the environment until period $t - \tau$, $(B_v, X_v)_{v=0,\tau,\ldots,t-\tau}$. (The initial values (B_0, X_0) are given.) A player's *private history* at time t consists of her actions and the details of her matches through period $t - \tau$. A player's information set at time $t = \tau, 2\tau, \ldots$ is given by the public history, together with her private history. *Strategies* are functions

---

[5]The assumption that the action revision rate is independent of a player's current action is for expositional simplicity only; see the discussion in section 5.2.

[6]Judd [14] discusses some technical problems that arise with a continuum of i.i.d. variables. Boylan [3] and Gilboa and Matsui [12] offer possible solutions in the context of random matching.

from the set of all information sets to the set of mixtures over $\{R, L\}$ that indicate what a player will do should she have an action revision opportunity.[7]

When a player has a revision opportunity, she chooses the best action given the probability distribution over paths $(B_v)_{v=t,t+\tau,...}$ and her beliefs about the path of play $(X_v)_{v=t,t+\tau,...}$ that will result from any given realization of $(B_v)_{v=t,t+\tau,...}$. In any period v, a player is matched with probability $m\tau$ to an opponent who plays R with probability X_v and L with probability $1 - X_v$. If she is locked into action $a \in \{R, L\}$, her expected payoff in the period is thus $m\tau(X_v u(a, R, B_v) + (1 - X_v)u(a, L, B_v))$. This payoff is relevant to her time t decision only if she has no revision opportunities between times t and v. This event has the probability $(1 - k\tau)^{\frac{v-t}{\tau}}$. Hence, her relative payoff to choosing R in period t is:[8]

$$E\left[\sum_{v=t,t+\tau,...} m\tau(1 - k\tau)^{\frac{v-t}{\tau}} e^{-r(v-t)} \left(X_v\Delta(R, B_v) + (1 - X_v)\Delta(L, B_v)\right)\right] \qquad (2)$$

where $r \geq 0$ is the common discount rate. A player chooses R if this relative payoff is positive and L if it is negative. Note that the relative weight on period v, $(1 - k\tau)^{\frac{v-t}{\tau}} e^{-r(v-t)}$, depends on k. When revision opportunities are frequent (k large), players put most of the weight on payoffs that they receive very soon. Nevertheless, we will see that more distant events remain important because of backwards induction.

Finally, to give our iterative dominance argument a place to start, we assume the existence of dominance regions. For $B_{t-\tau}$ large enough,

$$E\left[\sum_{v=t,t+\tau,...} m\tau(1 - k\tau)^{\frac{v-t}{\tau}} e^{-r(v-t)}\Delta(L, B_v)dv \mid B_{t-\tau}\right] > 0 \qquad (3)$$

---

[7]A given strategy profile gives rise to a particular dynamical system. Let h_t denote the vector of private histories at time t. Let $R(h_t)$ denote the probability that a randomly selected player chooses R after the history h_t. (If players use only pure strategies, $R(h_t)$ is just the proportion who choose R.) Then $X_t = X_{t-\tau}(1 - k\tau) + k\tau R(h_t)$, while $B_t = B_{t-\tau} + \mu\tau \pm \sigma\sqrt{\tau}$ with equal probabilities.

[8]This holds because the player's action choice can have no influence over which path $(X_v)_{v=t,t+\tau,...}$ will occur, since any player's action will be observed by only a countable number of other players.

which implies that R is strictly dominant. For $B_{t-\tau}$ small enough,

$$E\left[\sum_{v=t,t+\tau,\dots} m\tau(1-k\tau)^{\frac{v-t}{\tau}} e^{-r(v-t)}\Delta(R, B_v)dv \mid B_{t-\tau}\right] < 0 \qquad (4)$$

so that L is strictly dominant.

3   Solving the Model

Rather than looking for equilibria, we analyze the game using a more primitive solution concept: the *iterative elimination of conditionally dominated strategies* (see Fudenberg & Tirole [10, pp. 128 ff.]). This is essentially the extension of backwards induction to infinite horizon games. It is important to note that, in our model, iterative conditional dominance is not a refinement of Nash equilibrium: every Nash equilibrium outcome of the dynamic game survives the iterative process.[9]

When carrying out the iterative procedure, we use a simple method to keep track of the strategies that survive each successive round of elimination. We consider the space $\Re \times [0, 1]$ of all conceivable values of the current state (B_t, X_t). When we discover that R must be played in all surviving strategies after any public history h_t that ends in the

---

[9]This is because players are small and anonymous, so that no unilateral deviation can alter the probability distribution of reached information sets. Therefore, given a Nash equilibrium, one can alter the strategies in any way at unreached information sets and the resulting strategy profile will remain a Nash equilibrium. In particular, one can adjust the strategies at unreached information sets so that the overall equilibrium is subgame perfect and thus survives iterated conditional dominance. To make this argument precise, let s be a Nash equilibrium strategy profile, and let $s(b, x)$ denote the play prescribed by s if the initial state is $(B_0, X_0) = (b, x)$. We construct a subgame perfect equilibrium $\hat{s}$ that has the same distribution of equilibrium paths as s as follows. If the public history at time t is consistent with s, players continue to play according to s. (In particular, they ignore deviations by individual opponents.) Otherwise, let $v \leq t$ be the earliest period such that the public history at time t is consistent with all players having 'reset their clocks to zero' at time v and having played according to $s(B_v, X_v)$ thereafter. Under the new profile $\hat{s}$, players continue to conform to $s(B_v, X_v)$ after seeing the history h_t. Clearly, $\hat{s}$ induces the same equilibrium play as s and is subgame perfect. (We thank Philip Reny for suggesting this argument.)

state (B_t, X_t), we tag the state with 'R'.[10] We do the same for L.

Suppose a player receives an action revision opportunity at time t. If B_t is large enough, R is strictly dominant, so the player will choose R regardless of her beliefs over which strategies are used by the other players. Let f^0 be the boundary of the region where R is strictly dominant (i.e., where (3) holds). This is depicted in Figure 1. To the right of f^0, we know that the player must choose R; on or to the left of f^0 we cannot yet say what the player does. We tag the region to the right of f^0 with 'R'. This corresponds to eliminating all the strategies in which any player ever chooses L in states (B_t, X_t) that are to the right of f^0.
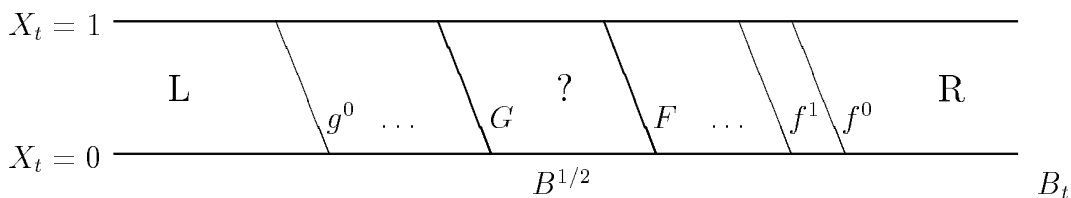

Figure 1: The iterative elimination procedure.

In the second step we assume that a player believes that other players will always choose R when they are strictly to the right of f^0. With this belief, there is a new boundary, f^1, such that a player must choose R when she is strictly to the right of f^1. f^1 must lie weakly to the left of f^0, since knowing that other players will sometimes choose R can only make R a more appealing action. In the next step we find f^2 and so on. Let F be the limit of the sequence $f^0, f^1, \ldots$. Whenever (B_t, X_t) is strictly to the right of F, any player who is called to act must choose R. In a similar way, starting an iterative process from the left side of the environment space where the action L is dominant, we construct a sequence $g^0, g^1, \ldots$, with limit G.

In all strategies that survive the iterative procedure, R is played to the right of F and L is played to the left of G. We cannot preclude the possibility of a region (denoted by '?' in Figure 1) between the two curves. Because it is not tagged with 'R' or 'L', we

---

[10]Since the random walk is a discrete process, $\Re \times [0, 1]$ may include some states that cannot be reached from a given (B_0, X_0). This does not present a problem; by tagging unreachable states, we do not alter the set of surviving strategies.

cannot say how players will act in this region (if it exists). Different strategies might partition this region into R and L in different ways, and their prescriptions might even depend on aspects of the history that are not reflected in the time-t environment space.

3.1 Results

We have results for two limiting cases: small frictions and small noise. In both cases, we take the period length τ to zero first. This ensures that the impact of the τk players who revise actions simultaneously in each period is negligible relative to the effects of the random changes in B, which are of order $\sigma \sqrt{\tau}$.[11]

Theorem 1 pertains to the case of small frictions: when k tends to infinity. In the limit, F and G coincide with the vertical line at $B^{1/2}$. This is depicted in Figure 2. Apart from a vanishing range of values of B_t around $B^{1/2}$, a player's choice between L and R is uniquely determined by the state of the world B_t. The player simply plays the action that would be a best response in the static game against an opponent who puts equal weight on R and L.
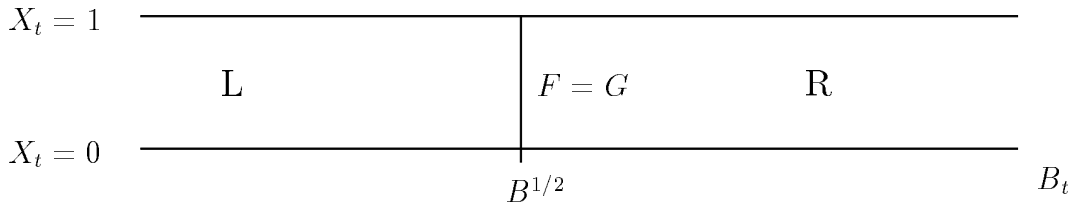

Figure 2: Case of vanishing frictions ($k \to \infty$).

Theorem 1 *Holding σ, μ, and r fixed, $\lim_{k \to \infty} \lim_{\tau \to 0} F = \lim_{k \to \infty} \lim_{\tau \to 0} G = B^{1/2}$. Equivalently, for any $\varepsilon > 0$, there is a $\underline{k}$ and a function $\underline{\tau}(\cdot) > 0$ such that if $k > \underline{k}$ and $\tau < \underline{\tau}(k)$, then R must be played whenever $B_t > B^{1/2} + \varepsilon$ and L whenever $B_t < B^{1/2} - \varepsilon$.*

This result has strong implications for the evolution of aggregate play in the population. For almost any initial state, the population quickly coordinates on the action that

---

[11]The fact that the random changes in B are proportional to the square root of the time interval is not a special feature of a random walk; it must hold for any stochastic process with independent increments. See section 5.1.

is risk-dominant. A dynamic interpretation is that the population follows the evolution of the world. Whenever B_t is above $B^{1/2}$, the whole population plays R, while when B_t is below $B^{1/2}$, all players coordinate on L. Switching between equilibria occurs very quickly whenever B_t crosses $B^{1/2}$.

Our second result (Theorem 2) concerns the case in which lock-in remains positive and instead the trend and variance of the payoff parameter tend to zero. Once again, the environment space is divided uniquely into R and L regions. If the players are perfectly patient, the risk-dominant equilibrium is again selected. But if players are impatient, the curve that separates the R and L regions is downward sloping rather than vertical. An example is depicted in Figure 3. (In general the curve need not be straight.)

The fact that the indifference line is downwards sloping means that for some intermediate range of values of B_0 there is *history dependence*: the action players choose depends on the current state of play. In this region, if enough players are playing R, the rest of the population will follow; otherwise, play converges to L.
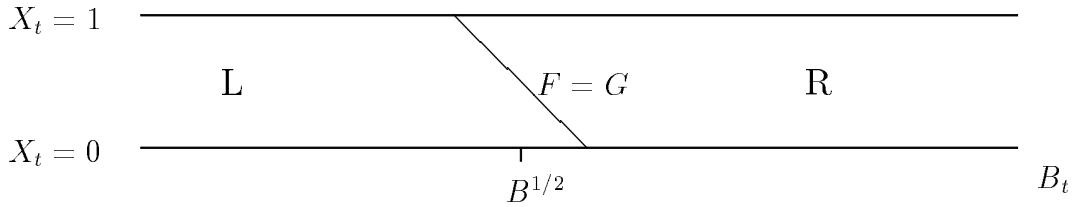

Figure 3: Case of slowly changing world and fixed frictions.

Theorem 2 *Let $h(x) = \frac{rx+k}{r+2k}$ and let $B^{h()}$ denote the curve $B_t = B^{h(X_t)}$. Holding k and r fixed, $\lim_{\sigma,\mu \to 0} \lim_{\tau \to 0} F = \lim_{\sigma,\mu \to 0} \lim_{\tau \to 0} G = B^{h()}$.*

Equivalently, for any $\varepsilon > 0$, there is a $\overline{\sigma} > 0$, $\overline{\mu} > 0$, and a function $\underline{\tau}(\cdot) > 0$, such that if $\sigma < \overline{\sigma}$, $|\mu| < \overline{\mu}$, and $\tau < \underline{\tau}(\sigma)$, then R must be played whenever $B_t > B^{h(X_t)} + \varepsilon$ and L whenever $B_t < B^{h(X_t)} - \varepsilon$.

Remark: *as players become patient ($r \to 0$), $h(x) \to \frac{1}{2}$, so the result converges to that of Theorem 1.*

The proofs of the two theorems are somewhat technical and appear in the appendix. Intuitions are given below.

3.2 Intuition for Theorem 1

We will explain why R must be played to the right of $B^{1/2}$ in the limit as $\tau \to 0$ and then $k \to \infty$. The argument that L must be played to the left of $B^{1/2}$ is analogous. To go through the iterative elimination, we need to build a sequence of curves such that a player must play R to their right. The functions f^n, defined on p. 8, are difficult to work with since we know little about their shapes. Instead, we work with sequences of functions whose form we do know: downwards sloping, almost vertical straight lines.

Fix a slope ρ that is negative and large in absolute value. Let q_0 be the leftmost line with slope ρ such that R is strictly dominant to the right of q_0. Then inductively define q_n for $n \geq 1$ as the leftmost translation of q_0 such that R is a strict best response at any state to the right of q_n even for a player who expects others to choose L when to the left of q_{n-1} and R to the right. Since this is the strategy profile that makes R the least appealing among not yet eliminated profiles, R must be played to the right of q_n.

Let Q be the limit of the q_n's as n goes to infinity. R must be played to the right of Q. Moreover, there must be at least one point on Q where a player would be indifferent between R and L if she believed that others would choose R when to the right of Q and L to the left. Otherwise, the iterative process would have gone beyond Q.[12,13] Consider such a player; let us call her 'A'. We will show that under A's beliefs, in the limit as $\tau \to 0$ and then $k \to \infty$, the proportion of R players among the opponents A expects to meet during her lock-in period converges to one half. Moreover, A meets all of her opponents very quickly in the limit, so her discount rate r plays essentially no role. Thus, since she is willing to play L, it must be that B_t converges to a number no greater than $B^{1/2}$.[14] Since Q can be taken as close to vertical as we like, R must be

---

[12] There might be only one point of indifference since we restrict the shape of Q. If the shape of Q were unrestricted (as in the construction of F), a player would be indifferent at all points.

[13] We are implicitly assuming that the relative payoff from choosing R is a continuous function of the state; the proof does not make this assumption.

[14] Recall that if a player expects her opponent to play R with probability p, the state at which she is indifferent is B^p. Thus, if she is willing to play L, the state must be no greater than B^p.

played at any state to the right of $B^{1/2}$ in the limit.

It remains to count how many of A's opponents she expects to play R. Suppose A revises her action at time $t + \tau$, after observing the state (B_t, X_t). While locked into her new action, she will meet a sequence of opponents. We divide them into three groups: "new" players, "old" players, and "same-time" players. A player is old if, at the time of his match with A, he will be locked into an action that he will have chosen at or before time t. A new player is one who will have chosen his action after time $t + \tau$. A same-time player will have chosen his action simultaneously with A, at time $t + \tau$.

The expected number of old players equals the expected number of new players. Why? If we select one of A's matches at random, knowing nothing about when she or her opponent chose her action, the probability that A chose first equals the probability that her opponent chose first. This is because all players have the same rate of revision opportunities. Thus, over A's lifetime, she is expected to meet the same number of new and old opponents. Since the current period of commitment is, *ex ante*, the same as any other, the expected numbers of new and old players in the period must also be equal.

We now count the proportion of R players A expects to meet while locked into her action when we first take τ to zero and then take k to infinity. Since this means that $k\tau$ goes to zero, the set of same-time opponents becomes negligible. Hence, half of A's opponents will be old players in the limit; of these, X_t will play R. Half of A's opponents will be new players; we will show that under A's beliefs, the probability that a new player plays R converges to $1 - X_t$. Therefore, the total proportion of R players indeed converges to

$$\frac{1}{2} \cdot X_t + \frac{1}{2} \cdot (1 - X_t) = \frac{1}{2} \tag{5}$$

How does A forecast the behavior of new players? Suppose that, as A believes, all players do choose R when to the right of Q and L when to the left. While the state is to the right of Q, all players currently playing L switch to R when they get the chance. In each period v, a measure $k\tau$ of players have chances to change actions. The proportion of L players is $1 - X_v$, so X_v increases by $k\tau(1 - X_v)$. Similarly, when $B_v < Q(X_v)$, players switch from R to L, and the proportion of R players is X_v, so X_v decreases by $k\tau X_v$.
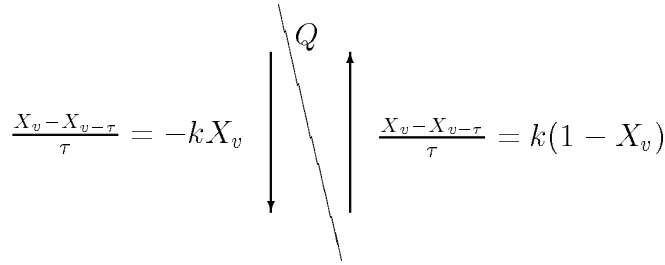
Figure 4: Local dynamics around Q.

Figure 4 illustrates these dynamics. Recall that A chooses her action when the state is on Q; consider what happens afterwards. The dynamics are unstable since the changes in X always pull us further away from Q. These dynamics will eventually lead to a *bifurcation*: a time beyond which the state remains on one side of the line until X (almost) reaches either 0 or 1.[15]

The basic properties of bifurcations are presented in Lemma 1 (see section 3.4 below). By the lemma, in the limit as $\tau \to 0$ and then $k \to \infty$, two things happen. First, the bifurcation time goes to zero relative to A's expected lock-in time, $1/k$. This means that essentially all of A's new opponents will have chosen their actions *after* the bifurcation. Hence, either all new opponents choose R (if there is an upward bifurcation) or all choose L (if the bifurcation is downward). Second, the probability of an upward bifurcation converges to $1 - X_t$. Thus, the probability that a "new" player will play R converges to $1 - X_t$, as claimed.

3.3 Intuition for Theorem 2

The intuition is essentially the same as for Theorem 1, with one exception. With nonvanishing frictions, a player's expected lock-in time does not go to zero. If she is impatient, she puts more weight on old players than on new players. This is because old players are typically encountered earlier than new players during the period of commitment. This means that in equation (5), she puts a weight of more than one half on the X_t and

---

[15]Note that X never actually reaches 0 or 1, since for every given length of time there is a positive fraction of players who have not yet received a chance to change their actions.

less than one half on the $1 - X_t$. Thus, a higher X_t makes R a more appealing choice. As a result, the division line is not vertical but rather slopes downward.[16] The rest of the argument is still valid. The relevant properties of bifurcations still hold since they depend only on k becoming large *relative* to σ and μ. The intuition for this is explained below.

3.4 Bifurcation properties

Both theorems rely on two crucial properties of the dynamical system depicted in Figure 4: when $B_t = Q(X_t)$, (a) the system bifurcates very quickly and (b) the ratio of probabilities of an upward and downward bifurcation converges to the ratio of the speeds at which the system moves up and down, $(1 - X_t)/ X_t$. Lemmas 1 and 2 state these properties for the cases corresponding to Theorems 1 and 2, respectively. An intuition follows; rigorous proofs appear in the appendix.

We formally define a bifurcation as follows. Let us say that a state (B_t, X_t) is *close to Q* if it is within one horizontal jump of Q.[17] Fix a small $\delta > 0$ that is less than the smaller of the distances from X_t to 0 and 1.[18] If (B_t, X_t) is close to Q, the (random) *time of a δ-bifurcation* is the last time $v \geq t$ at which the state is close to Q before X reaches a δ-neighborhood of either zero or one.[19] We say that the bifurcation is *upward* if, following the bifurcation, the state remains to the right of Q (so that X grows towards 1). Otherwise it is a *downward* bifurcation.

Lemma 1 *Fix σ, μ, and δ. Let $Q : [0, 1] \to \Re$ be any strictly decreasing, continuously differentiable function. Consider the dynamical system depicted in Figure 4.[20] For any*

---

[16]Because of this, we cannot do the iterations with straight lines as in Theorem 1. Instead, we use translations of the curve $B^h()$. For details, see the proof.

[17]This holds if $-\sigma\sqrt{\tau} + \mu\tau < B_t - Q(X_t) < \sigma\sqrt{\tau} + \mu\tau$.

[18]I.e., $\delta < \min\{X_t, 1 - X_t\}$.

[19]If X_t equals 0 or 1, δ can be any small positive number. The bifurcation time is then the last time at which the state is close to Q before X comes within δ of 1 or 0, respectively.

[20]In this system, $B_v - B_{v-\tau}$ equals $\mu\tau \pm \sigma\sqrt{\tau}$ with equal probabilities. $X_v - X_{v-\tau}$ equals $-k\tau X_{v-\tau}$

$\varepsilon > 0$, there is a $\underline{k}$ and a function $\underline{\tau}(\cdot) > 0$ such that if $k > \underline{k}$ and $\tau < \underline{\tau}(k)$, and if (B_t, X_t) is 'close to Q', then

1. with probability at least $1 - \varepsilon$, a δ-bifurcation will happen within time ε/k;
2. the probability that the bifurcation will be upward is within ε of $1 - X_t$.

Lemma 2 *Fix k. Let $Q : [0,1] \to \Re$ be any strictly decreasing, continuously differentiable function. Consider the dynamical system depicted in Figure 4. For any $\varepsilon > 0$, there are $\underline{\sigma} > 0$, $\underline{\mu} > 0$ and a function $\underline{\tau}(\cdot) > 0$ such that if $\sigma < \underline{\sigma}$, $|\mu| < \underline{\mu}$, and $\tau < \underline{\tau}(\sigma)$, and if (B_t, X_t) is 'close to Q', then*

1. with probability at least $1 - \varepsilon$, a δ-bifurcation will happen within time ε/k;
2. the probability that the bifurcation will be upward is within ε of $1 - X_t$.

3.5   Intuitions for Lemmas 1 and 2

We begin with the intuition for Lemma 2. Lemma 1 will follow by a change of variables. It is convenient to linearize the dynamical system: let $D_t = B_t - Q(X_t)$, so that $|D_t|$ is the horizontal distance between the state and Q. We can think of the process D as the sum of a random component (mean zero jumps of $\pm\sigma\sqrt{\tau}$) and two deterministic trends: positive jumps of approximately $\lambda_R \tau = k(1 - X)|Q'(X)|\tau$ when D is positive and negative jumps of about $-\lambda_L \tau = -kX|Q'(X)|\tau$ when D is negative. (We can ignore the $\mu\tau$ term since μ goes to zero while k is fixed.) The trends λ_R and λ_L are depicted in Figure 5.

Figure 5: Deterministic component of D.
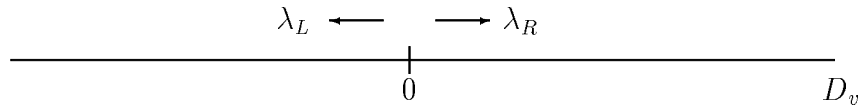
An upward (downward) bifurcation of (B, X) corresponds to a positive (negative) bifurcation of D: a time after which D remains positive (negative) for some long period.

if $B_{v-\tau} < Q(X_{v-\tau})$ and $k\tau(1 - X_{v-\tau})$ if $B_{v-\tau} > Q(X_{v-\tau})$. If $B_{v-\tau} = Q(X_{v-\tau})$, $X_v - X_{v-\tau}$ can take any value in $[-k\tau X_{v-\tau}, k\tau(1 - X_{v-\tau})]$.

The smaller is the δ-neighborhood that X must reach, the longer D must remain positive (or negative) for a bifurcation to occur. The reasoning behind Theorems 1 and 2 is that by taking δ to be very small, we can ignore what happens in the distant future after X reaches a δ-neighborhood of zero or one. We will approximate this case of small δ by considering the probabilities that D will stay positive or negative *forever*.

We first need to explain why the bifurcation occurs almost instantly. Intuitively, after spending any time on one side of 0 the deterministic trend will take D a positive distance from 0. But as the random jumps shrink, so does the chance that a reverse movement in D will take us back to 0. This means that a increasingly short stay on one side of 0 suffices to make bifurcation very likely. Hence, the bifurcation must happen quickly.

To compute the odds of a positive vs. a negative bifurcation, we use several approximations. First, since X changes little before the bifurcation, the trends λ_R and λ_L can be treated as a fixed constant $(k|Q'(X_t)|)$ times the constants $1 - X_t$ and X_t, respectively. Moreover, since τ is taken to zero before the other limits, we can treat D as moving in a continuous way. With this last approximation, D can be said to bifurcate at time v if $D_v = 0$ and $D_s \neq 0$ for all $s > v$.

Suppose that $D_t = 0$. Let $P^\uparrow$ be the probability of a positive bifurcation of D occurring at any time after t. Let $P^\downarrow$ be the probability of a negative bifurcation. For any constant $\gamma > 0$, let $P^\uparrow_\gamma$ and $P^\downarrow_\gamma$, respectively, be the probabilities of a positive and negative bifurcation occurring at some time $v \in [t, t + \gamma]$. We claim that, the ratio $P^\uparrow/P^\downarrow$ of bifurcation probabilities equals the ratio $P^\uparrow_\gamma/P^\downarrow_\gamma$. Why? Let $P_\gamma = P^\uparrow_\gamma + P^\downarrow_\gamma$. If there is no bifurcation in the interval $[t, t + \gamma]$, then D must equal 0 at some time $v > t + \gamma$. As of time v, the probability of an upwards bifurcation is once again $P^\uparrow$. Thus, $P^\uparrow = P^\uparrow_\gamma + (1 - P_\gamma)P^\uparrow$. This shows that $P^\uparrow = P^\uparrow_\gamma/P_\gamma$. Likewise, $P^\downarrow = P^\downarrow_\gamma/P_\gamma$. Therefore, $P^\uparrow/P^\downarrow = P^\uparrow_\gamma/P^\downarrow_\gamma$.

Since this equality holds for all γ, it also holds as γ goes to zero. This limit turns out to be easy to compute. For an upwards bifurcation to occur in $[t, t+\gamma]$, two things must happen. First, D must be positive at time $t + \gamma$. Second, D must remain positive forever after. Where is D at $t + \gamma$? Since $D_t = 0$, the value of D at time $t + \gamma$ is dominated by

the noise in B. This is because the standard deviation of $B_{t+\gamma} - B_t$ is proportional to $\sqrt{\gamma}$, while the linear trends λ_R and λ_L (coming from changes in X) produce a change of order only γ. For small γ, $\sqrt{\gamma}/\gamma$ is arbitrarily large. So for small γ we can treat the distribution Ψ_γ of $D_{t+\gamma}$ as approximately symmetric around zero. (In particular, it is approximately normal with mean 0 and standard deviation $\sigma\sqrt{\gamma}$).[21]

Given that D is positive at $t + \gamma$, what is the probability that it remains so forever after? If $D_{t+\gamma} = z$ and z is small, this probability is approximately proportional to the distance z times the trend λ_R. This relies only on the fact that a random walk has i.i.d. increments. To see why, let p be this probability, and consider what happens if $D_{t+\gamma}$ starts twice as far away, at $2z$. What is the probability that D never hits zero? It is the probability p that D never hits z, plus the probability $1 - p$ that D hits z times the probability p that, from z, D never reaches zero. This is $p + (1 - p)p$, which is approximately $2p$ since, for small z, $1 - p$ is close to one. This shows that the probability that D never hits zero if it starts at z is approximately proportional to z for small z.

Why is this probability also approximately proportional to λ_R? Let us multiply the time scale by 4 and the space scale by 2, giving a new process $\widehat{D}_v = 2D_{v/4}$. As long as D and $\widehat{D}$ remain positive, both have constant trends, so they have the same variance:

$$\mathrm{Var}(\widehat{D}_w - \widehat{D}_v) = \mathrm{Var}(2D_{w/4} - 2D_{v/4}) = 4\sigma^2\left(w/4 - v/4\right) = \sigma^2(w - v) = \mathrm{Var}(D_w - D_v)$$

The new process $\widehat{D}$ begins at $2z$. Since the time scale is stretched by twice the space scale, the trend of $\widehat{D}$ is $\lambda_R/2$, half the trend of D. But a change in the scaling cannot affect the probability of never hitting zero, which must still be p after doubling the initial distance z and halving the trend λ_R. Since the probability that D never hits zero is linear in z, it must also be linear in λ_R.

Hence, if $D_{t+\gamma} = z$ and z is small, the probability that D_v remains positive for all $v > t + \gamma$ is proportional to $z\lambda_R$. Since most of the weight of the c.d.f. Ψ_γ of z is close to zero, the probability $P_\gamma^\uparrow$ is approximately proportional to

$$\int_{z=0}^\infty z\lambda_R d\Psi_\gamma(z) \;=\; \lambda_R \int_{z=0}^\infty z d\Psi_\gamma(z)$$

---

[21] This property of Brownian motions, that the noise swamps any linear trend over short intervals, is an implication of independent increments. See note 26 in section 5.1.

The probability $P_\gamma^\downarrow$ that D has a negative bifurcation in $[t, t + \gamma]$ is approximately proportional to

$$\int_{z=-\infty}^{0} |z| \lambda_L d\Psi_\gamma(z) \approx \int_{z=0}^{\infty} z \lambda_L d\Psi_\gamma(z) = \lambda_L \int_{z=0}^{\infty} z d\Psi_\gamma(z)$$

because Ψ_γ is approximately symmetric. Therefore, the ratio $P_\gamma^\uparrow / P_\gamma^\downarrow$ equals λ_R / λ_L as γ goes to zero. This ratio is approximately $(1 - X_t)/X_t$. This explains why the relative probability of bifurcating up vs. down converges to $(1 - X_t)/X_t$ as τ and then σ and μ shrink to zero.

Lemma 1 follows from Lemma 2 by a change of variables. We stretch the time scale by replacing B_t with $\hat{B}_t = B_{t/k}$ and X_t with $\hat{X}_t = X_{t/k}$. In the new time coordinates, the rate of change of $\hat{X}$ is fixed while the variance σ/k and trend μ/k of $\hat{B}$ shrink to zero, so Lemma 2 can be applied.

4   Relation to the Literature

This paper is related to two research programs. The first studies how connections among "nearby" games can determine how rational players will behave in a given game. The second is the literature on dynamic population models.

Connections Among Nearby Games

Our framework is closely related to that of Carlsson and van Damme [6] (henceforth, CvD).[22] That paper studies a one-shot 2×2 game whose payoffs are not common knowledge. Rather, each player receives a slightly noisy signal of the true payoffs. The space of possible payoffs includes regions where each action is strictly dominant. Iterative strict dominance gives rise to a contagion effect that starts from these regions and determines how players will play throughout the space of possible payoffs. For small enough noise in the signals, the players must play the risk-dominant equilibrium of the true game.

There is an analogy between CvD's game and our own. Since CvD's game has incomplete information, a player must consider what opponents of different types will do. (In CvD, a player's type is her signal.) The correlation between the two players'

---

[22]See also Morris, Rob and Shin [19].

signals leads players to focus on opponents whose types are close to their own. This (together with the dominance regions) is what causes the contagion effect. In our model, payoffs change randomly. Thus, a player who wishes to predict what players will choose after her must consider a distribution of possible "types", each seeing different payoffs. The persistence of the random walk plays a role similar to that of correlation in CvD: most of these players will choose actions at states that are close to that at which the given player chooses her action.

However, there is a fundamental difference. In our model, when a player chooses her action, half of her future opponents have already chosen theirs. Any reasoning she does about what different types of players must do can relate only to the half who *haven't* chosen their actions yet. This contrasts with CvD, where a player's conclusions from tracing other players' decision problems apply to *all* of her potential opponents.

Indeed, trying to apply CvD's technique directly to our model (i.e., doing iterative dominance on the payoff parameter alone) establishes only a weaker result: that an action must be played if it is better than 1/4-dominant. No prediction can be reached if B_t is greater than $B^{3/4}$ but less than $B^{1/4}$. Formally, doing iterative dominance on B_t alone is equivalent to iterating with vertical lines. Suppose that such an iterative process starting from the region where R is strictly dominant ends at $B_t = b$. Consider a player who picks actions at $X_t = 0$ with B_t slightly greater than b. We know that she is willing to play R. With small frictions, she cares only about the very near future, when the random walk is equally likely to be above or below b. Therefore, she expects that half of the new players will play R and half L. Since all old players are playing L, she expects just quarter of her opponents to play R. Hence, B_t must be $B^{1/4}$. Similarly, iterating from the left also stops short of $B^{1/2}$, at $B^{3/4}$.

The key to our stronger result is observing that the new and old players offset each other. The reason is that once the population starts shifting towards a given action, the speed at which this occurs is proportional to the size of the population currently playing the *other* action. On the curve F, these speeds govern the bifurcation probabilities.[23]

---

[23]This follows from our assumption that payoffs change frequently (i.e., that τ tends to zero faster than anything else). This point is discussed further in section 5.1.

Hence, when the proportion X_t of old players playing R is low, the chance $1 - X_t$ that the new players will play R is high. When the payoff weight on the two groups is equal (i.e., when frictions are small or players are patient), this leads to a strong selection result: X_t plays no role at all in a player's decision. When old players receive more weight (i.e., when frictions are large and players are impatient), there is history dependence since the new players only partially offset the old.

To see more clearly the difference between our model and a static model of incomplete information game, it is instructive to look at the model of Matsui [17]. His also has a payoff parameter that changes randomly in each period. There is a sequence of players, each of whom chooses one action. Unlike in our model, each player's payoff comes only from her interaction with the player who chooses *after* her. Matsui shows that his model is isomorphic to a static incomplete information game like that of CvD. Such an isomorphism cannot be found with our model, since the actions of past players (captured by the state variable X_t) are payoff-relevant. The lack of such an isomorphism is most obvious in the case of small noise and large frictions, where we find history dependence.

Dynamic Population Models

There is an extensive literature that studies models in which players in a large population are randomly matched, from time to time, to play some normal-form game. With less than fully rational players and mutations, these models can yield equilibrium selection in the form of a long-run ergodic distribution with most of its weight on a given equilibrium.[24]

Our model, which assumes fully rational players, is most related to that of Matsui and Matsuyama [18] (henceforth, MM). The essential difference between the two models is that payoffs in MM are fixed. In our model this corresponds to taking B_t as a constant parameter. Figure 6 illustrates what happens with constant B_t, in the case of small frictions (which is the case on which MM focus).

---

[24]Examples include Foster and Young [8], Fudenberg and Harris [11], Kandori, Mailath and Rob [15], and Young [21]. Models of bounded rationality with *local interaction* yield faster convergence; see, for example, Ellison [7] and Blume [2].
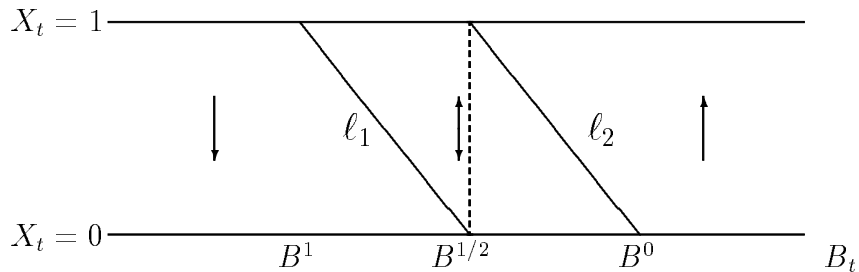
Figure 6: Case of constant B_t, corresponding to MM.

To the left of ℓ_1, only L can be selected in any rational expectations equilibrium. R must be chosen to the right of ℓ_2.[25] Between the curves, there are multiple equilibria, including one in which players always select R and another in which they always choose L. However, MM find support for the prediction that the risk-dominant equilibrium of the static game is more likely to be played, since it is 'globally accessible' and 'uniquely absorbing.' Suppose, e.g., that $B_0 > B^{1/2}$. Since the initial state must be to the right of ℓ_1, it is an equilibrium for all players to choose R for any initial value X_0: R is 'globally accessible.' Moreover, if X_0 is large enough, the initial state lies above ℓ_2, so all players *must* choose R. Hence, R is also 'uniquely absorbing.'

Our model shows that with stochastic payoff changes, the area between ℓ_1 and ℓ_2 splits into two regions. The part to the left of $B^{1/2}$ joins the area where players must choose L, while the part to the right of $B^{1/2}$ joins the area where they must choose R. Moreover, the changes in payoffs make the model dominance-solvable.

5   Concluding Remarks

5.1   The Assumption of Frequent Jumps in the Random Walk

For all of our results we take the time τ between jumps of the random walk to be small given the other key parameters of the model, k and σ. Without this, there could be regions with multiple equilibria. To see why, consider the group of players who choose actions simultaneously in a given period. This group constitutes a proportion $k\tau$ of the

---

[25]These curves correspond roughly to g^0 and f^0 in Figure 1. In the limit as frictions vanish, this correspondence is exact.

21

players. If they all choose the action a, their relative payoff to playing a is increased, and this might justify their choice.

The group's choice affects their relative payoff in two ways. First, members of the group may play each other. To make this effect negligible, we take τ small given k, so that $k\tau$ is small. An analogous issue emerges in the static incomplete information game of CvD [6]. There, multiple equilibria might arise if the probability that the two players see the same signal of the true payoff parameter is large. CvD make this probability zero by assuming a continuous signal distribution.

A second consideration, which has no analogue in CvD's model, comes from the dynamic structure of our model: if the $k\tau$ others pick R, X will rise, possibly moving the state into a region where yet more players will pick R. Taking τ small given k and σ eliminates this possibility by ensuring that the effect of the players' simultaneous action choices, which is of order $k\tau$, is swamped by the random jumps in B, which are of size $\sigma\sqrt{\tau}$.[26]

5.2 Extensions

The Stochastic Process

We assume that payoffs change according to a random walk. However, our results remain valid if payoffs follow any discrete process with independent, stationary increments.[27] For example, the jump in each period might be normally distributed. Or the waiting times between jumps might themselves be random. As τ shrinks, all such processes behave in essentially the same way and share the properties of a random walk that we use. Our results also hold when τ *equals* zero, so that B is a Brownian motion and time is continuous (see Burdzy, Frankel, and Pauzner [4]). Hence, there is

---

[26]The property that the noise $\sigma\sqrt{\tau}$ is proportional to the square root of the time τ between jumps is a consequence of the random walk having independent increments. The change in B over a given interval of length, say, 1, is the sum of $1/\tau$ independent jumps. The only way this sum can retain a nontrivial variance as τ shrinks is if the size of each jump is of order $\sqrt{\tau}$.

[27]Stationarity means that the distribution of increments over a given time interval can depend only on its length.

continuity at the limit.

The assumption of independent, stationary increments can also be weakened. B may be any strictly increasing, continuously differentiable function of a random walk, as long as it can reach values at which either action is strictly dominant.[28] This lets B represent bounded processes such as prices or temperatures. Our results also hold if the trend parameter, rather than being constant, is any bounded Lipschitz function of t and B_t. For instance, the trend may have a seasonal component or be mean-reverting. Our results may be of greater interest with mean-reverting processes. This is because a random walk tends to wander away from $B^{1/2}$ to regions where one action is strictly dominant; a mean reverting process with mean close to $B^{1/2}$ spends a positive fraction of its time in the area where the static game has multiple equilibria.

In the case we study (small τ), the stochastic process changes almost continuously over time. For some applications, it may be more natural to assume that in addition to its small frequent jumps, B also has large infrequent jumps at random times. Theorem 1 holds for any such process, since during a player's shrinking lock-in period the chance of a large jump becomes very small. For Theorem 2, the noise can become small in two ways: the large jumps may become less and less frequent, or they may retain their frequency but become smaller and smaller. Theorem 2 holds in both cases. The reason is that the bifurcation occurs earlier and earlier, so that the probability that a random jump will precede the bifurcation goes to zero; hence, infrequent jumps have almost no effect on the dynamics of X.

Dominance Regions and Monotonicity

We assume that extreme values of B_t make either action strictly dominant and that the relative payoff to R against either action is strictly increasing in B_t. The first assumption gives the contagion argument a place to start. The second guarantees that the iterations starting from the two dominance regions will meet.

If these assumptions do not hold, the following weaker version of Theorem 1 can still

---

[28] This is because any such transformation of B is equivalent to a change in the utility function: instead of $u(a, a', \tilde{B}_t)$ where $\tilde{B}_t = g(B_t)$, we use the utility function $\tilde{u}(a, a', B_t) = u(a, a', g(B_t))$.

be proved. Assume only that the game has strategic complementarities ($\Delta(R, B_t) > \Delta(L, B_t)$) and that $\Delta(a, B_t)$ is continuously differentiable in B_t. Let b be any value of B_t at which R is strictly dominant. Let $(\underline{b}, \overline{b})$ be the largest (potentially infinite) interval that includes b, such that R is risk-dominant (and L is not) throughout the interval. In the limit as $\tau \to 0$ and then $k \to \infty$, R must be played at all points in this interval. An analogous result can be proved for L.

Different Revision Rates

Our model assumes that a player's rate of action revision opportunities is independent of the action she is currently playing. With differing revision rates, players still choose the risk-dominant action in the case of small frictions (Theorem 1). To see why, let k_a be the revision rate of a player currently locked into action a. We compute the relative payoff of player 'A' (p. 11) to playing R vs. L by first conditioning on bifurcations in each direction. If A knew that the bifurcation would be upward, she would play R from now on. By picking L rather than R, she loses utility only until her next revision opportunity, when she will switch to R. These opportunities arrive at rate k_L, which equals the rate at which other players change actions (since only L players switch). Hence, her expected numbers of new and old opponents while she is locked into L are still equal.[29] X_t of the old players and *all* of the new players play R, so the expected number of R players is approximately $1/2 + X_t/2$. The expected lock-in time is $1/k_L$. Hence, her relative payoff from playing R is $\frac{1}{k_L} \left[\frac{1+X_t}{2} \Delta(R, B_t) + \frac{1-X_t}{2} \Delta(L, B_t) \right]$. Similarly, conditional on a downward bifurcation, the relative payoff from R is $\frac{1}{k_R} \left[\frac{X_t}{2} \Delta(R, B_t) + \left(\frac{2-X_t}{2} \right) \Delta(L, B_t) \right]$. The probabilities of upward and downward bifurcations are proportional to the speeds at which X moves up and down, $k_L(1 - X_t)$ and $k_R X_t$. Therefore, the (unconditional) relative payoff to choosing R is proportional to

$$k_L(1 - X_t) \frac{1}{k_L} \left[\frac{1 + X_t}{2} \Delta(R, B_t) + \frac{1 - X_t}{2} \Delta(L, B_t) \right]$$
$$+ k_R X_t \frac{1}{k_R} \left[\frac{X_t}{2} \Delta(R, B_t) + \left(\frac{2 - X_t}{2} \right) \Delta(L, B_t) \right]$$

---

[29]More precisely, since R players never change their actions in an upward bifurcation, their revision rate is irrelevant. So for the purposes of computation, we can assume that the R players also have the revision rate k_L and hence that the expected numbers of new and old opponents are equal.

$$= \frac{1}{2} \left[\Delta(R, B_t) + \Delta(L, B_t) \right]$$

Thus, the division line is again vertical at $B^{1/2}$.

In the case of small noise, the weight on old players depends on the revision rate corresponding to each direction of bifurcation. Hence, the revision rates do not cancel out in the computation. The result is qualitatively the same as in Theorem 2, but the division line now depends on the two revision rates, k_R and k_L.

Homogeneity of Players

We assume all players have the same payoff function. Suppose instead that the payoff of player $i \in [0, 1]$ is $u(a, a', B_t + \theta_i)$. Assume that most of the players have taste parameters θ_i in the range $[\theta, \theta']$: no more than ε are below θ, while no more than ε' are above θ'. A modification of our argument shows that as frictions vanish, at least $1 - \varepsilon$ of the players choose R when $B_t > B^{1/2 - \varepsilon/2} - \theta$ and at least $1 - \varepsilon'$ choose L when $B_t < B^{1/2 + \varepsilon'/2} - \theta'$. This shows that there is continuity: when most of the players have very similar tastes, our results hold approximately.

A Proofs

Proof of Theorem 1

We show that R must be played for $B_t > B^{1/2} + \varepsilon$; the proof that L must be played
for $B_t < B^{1/2} - \varepsilon$ is analogous. Let $q_{-1}(x) = \infty$. For $n \geq 0$, let $q_n(x) = -\varepsilon x/3 + \lambda_n$
where λ_n is the smallest constant such that at all states (B_t, X_t) to the right of q_n (i.e.,
satisfying $B_t > q_n(X_t)$), R is a strict best response if all other players are expected
to choose R whenever the state is to the right of q_{n-1} and L otherwise. Since R is
strictly dominant for large enough B_t, $q_n(x)$ must be finite for all $n \geq 0$. By strategic
complementarities, $q_n(x) \leq q_{n-1}(x)$. Moreover, by the existence of a dominance region
for L, the sequence $(q_n(x))_{n \geq -1}$ is bounded below. Hence $Q(x) = \lim_{n \to \infty} q_n(x)$ exists
and is finite.

All players must choose R when the state is to the right of Q. Since the iterative
process stops at Q, there must also be states arbitrarily close to (or on) Q at which
a player is willing to choose L if she thinks that all others will choose R when to the
right of Q and L otherwise. Let 'A' denote a player with these beliefs who observes
such a state (b, x) and then chooses an action at time t. (Note that b, x, and Q are all
functions of k and τ.) We can assume that (b, x) is within one horizontal jump of Q.
Let $\underline{\tau}()$ be small enough that this implies $b - Q(x) < \varepsilon/3$.[30] By (2),

$$0 \geq E\left[\sum_{v=t,t+\tau,\dots} m\tau(1 - k\tau)^{\frac{v-t}{\tau}} e^{-r(v-t)}[X_v\Delta(R, B_v) + (1 - X_v)\Delta(L, B_v)]\right]$$

Taking τ to zero for any given k, the right hand side converges to

$$E\left[\int_{v=t}^{\infty} me^{-(k+r)(v-t)}[X_v\Delta(R, B_v) + (1 - X_v)\Delta(L, B_v)]\,dv\right]$$

since $\lim_{\tau \to 0}(1 - k\tau)^{\frac{v-t}{\tau}} = e^{-k(v-t)}$. Since Δ is a continuous function, as $k \to \infty$ this
converges to

$$\int_{v=t}^{\infty} me^{-(k+r)(v-t)}[E(X_v)\Delta(R, b) + (1 - E(X_v))\Delta(L, b)]\,dv \tag{6}$$

---

[30]Since $b - Q(x) < \sigma\sqrt{\tau} + |\mu|\,\tau$, we must have $b - Q(x) < \varepsilon/3$ if $\tau < \min\left\{\varepsilon/3\sigma, [\varepsilon/3\sigma]^2\right\}$, which we
can guarantee by choosing $\underline{\tau}(\sigma)$ to be smaller than $\min\left\{\varepsilon/3\sigma, [\varepsilon/3\sigma]^2\right\}$.

By Lemma 1, for any $\delta \in (0, \min(x, 1 - x))$, in the limit as first $\tau \to 0$ and then $k \to \infty$ player A expects a δ-bifurcation to happen almost instantaneously relative to her expected lock-in time, $1/k$. The probability that the δ-bifurcation will be upwards is $1 - x$. Since δ can be arbitrarily small, $E(X_v)$ converges to $(1 - x)(1 - (1 - x)e^{-kv}) + x(xe^{-kv}) = 1 - x + (2x - 1)e^{-kv}$. Therefore, the (nonpositive) relative payoff from choosing R converges to

$$\int_{v=t}^{\infty} me^{-(k+r)(v-t)} \left((1 - x + (2x - 1)e^{-kv})\Delta(R, b) + (x - (2x - 1)e^{-kv})\Delta(L, b) \right) dv$$

$$= \frac{m}{k + r} \left[\frac{rx + k}{2k + r}\Delta(R, b) + \left(1 - \frac{rx + k}{2k + r} \right) \Delta(L, b) \right]$$

As $k \to \infty$ the expression in square brackets converges to $\frac{1}{2}\Delta(R, b) + \frac{1}{2}\Delta(L, b)$. But $\frac{1}{2}\Delta(R, B^{1/2}) + \frac{1}{2}\Delta(L, B^{1/2}) = 0$. So for any $\varepsilon > 0$, if k is sufficiently large and τ is small enough given k, then b must be less than $B^{1/2} + \varepsilon/3$. Since $dQ(x)/dx = -\varepsilon/3$, $|b - Q(x)| < \varepsilon/3$, and R must be played to the right of Q, it follows that R must be played whenever $B_t > B^{1/2} + \varepsilon$. Q.E.D.

Proof of Theorem 2

The proof follows the lines of the proof of Theorem 1, so we simply sketch the differences. Instead of doing the iterations with lines of a fixed slope, we use translations of $B^{h(x)}$, where $h(x) = \frac{rx+k}{2k+r}$. (It is straightforward to check that $B^{h(x)}$ is strictly decreasing and continuously differentiable in x.) Suppose player A is willing to choose L at a state (b, x) satisfying $|b - Q(x)| < \varepsilon/2$ on the belief that others will play according to Q. By Lemma 2, in the limit as $\tau \to 0$ and then $\sigma, \mu \to 0$, player A expects an instantaneous bifurcation, with the probability $1 - x$ that it will be upwards. As in the proof of Theorem 1, her (nonpositive) relative payoff from choosing R is proportional to

$$\frac{rx + k}{2k + r}\Delta(R, b) + \left(1 - \frac{rx + k}{2k + r} \right) \Delta(L, b)$$

in the limit. This expression is zero precisely when $b = B^{h(x)}$. Hence, we can choose σ and μ sufficiently small, and τ small enough given σ and μ, so that $b \leq B^{h(x)} + \varepsilon/2$.

Since $|b - Q(x)| < \varepsilon/2$, Q must be within ε of $B^{h()}$. But R must be played to the right of Q. The proof that L must be played whenever $B_t < B^{h(X_t)} - \varepsilon$ is analogous. Q.E.D.

Proof of Lemma 2

Time of Bifurcation

Let $D_t = B_t - Q(X_t)$. Claim 1 says that a bifurcation is virtually guaranteed if D_t reaches a distance of order σ^2 from 0.

Claim 1 *There are $c < \infty$, $\underline{\mu} > 0$, and $\underline{\tau}_0(\cdot) > 0$ such that for any σ, if $|\mu| < \underline{\mu}$ and $\tau < \underline{\tau}_0(\sigma)$ then once $|D|$ exceeds $c\sigma^2$, the probability exceeds $1 - \varepsilon/2$ that X will reach a δ-neighborhood of 0 or 1 without the state (B, X) ever again being 'close to Q'.[31]*

Proof Let $\overline{Q}'$ and $\underline{Q}'$ be (strictly positive and finite) upper and lower bounds on $|Q'(x)|$. These bounds exist because Q' is strictly negative and continuous over the compact interval $[0, 1]$. When on one side of Q, the mean of $|D_t|$ changes at a rate no less than $\lambda = k\underline{Q}'\delta - |\mu|$, since $\min(1 - X_t, X_t) > \delta$. Let $\underline{\mu} = k\underline{Q}'\delta/2$, so that λ is strictly positive if $|\mu| < \underline{\mu}$. Then if $|D_t| = c\sigma^2$, the chance of hitting Q before X reaches a δ-neighborhood of 0 or 1 is less than $e^{-2c\lambda}$ in the limit as $\tau \to 0$ (see Karlin and Taylor [16, p. 362]).[32] So letting $c = \ln(2/\varepsilon)/(k\underline{Q}'\delta)$, once $|D_t|$ equals $c\sigma^2$, the probability exceeds $1 - \varepsilon/2$ that the state will reach a δ-neighborhood of 0 or 1 without ever again being 'close to Q'. Q.E.D.$_{\text{(Claim 1)}}$

We now show that for small σ^2, $|D|$ reaches $c\sigma^2$ quickly with high probability. Claim 2 shows that independently of the state at time t, by time $t + \sigma^2$ the noise in B takes us on the order of σ^2 away from Q with positive probability:

Claim 2 *There are constants $p_1 > 0$ and $c_1 > 0$, and a function $\underline{\tau}_1(\cdot) > 0$, such that for any σ^2, μ, B_t and X_t, if $\tau < \underline{\tau}_1(\sigma)$ then the probability that $|B_{t+\sigma^2} - Q(X_{t+\sigma^2})| > c_1\sigma^2$ is greater than p_1.*

---

[31]For brevity, Claim 1 only addresses the case in which X_t is not 0 or 1. Otherwise, the same proof guarantees a bifurcation only if D reaches $+\underline{d}$ (in the case $X_t = 0$) or $-\underline{d}$ (in the case $X_t = 1$). This is not a problem since in Part 2 we will show that the probability of reaching the 'correct' $\underline{d}$ is nearly 1 in the limit.

[32]This uses the fact, from Billingsley [1], that as $\tau \to 0$, B converges to a Brownian motion with variance σ^2 and drift μ.

Proof

$$\Pr\left(\left|B_{t+\sigma^2} - Q(X_{t+\sigma^2})\right| > c_1\sigma^2\right) \geq \Pr\left(\left|B_{t+\sigma^2} - Q(X_t)\right| - \left|Q(X_{t+\sigma^2}) - Q(X_t)\right| > c_1\sigma^2\right)$$

$$\geq \Pr\left(\left|B_{t+\sigma^2} - Q(X_t)\right| > (c_1 + k\overline{Q}')\sigma^2\right)$$

since $|Q(X_{t+\sigma^2}) - Q(X_t)| \leq k\overline{Q}'\sigma^2$. But the probability that the difference between a random variable $B_{t+\sigma^2}$ and a given number $Q(X_t)$ exceeds some number $(c_1 + k\overline{Q}')\sigma^2$ is at least the probability that the difference between the random variable and its *mean* exceeds this number. Thus,

$$\Pr\left(\left|B_{t+\sigma^2} - Q(X_t)\right| > (c_1 + k\overline{Q}')\sigma^2\right) \geq \Pr\left(\left|B_{t+\sigma^2} - E(B_{t+\sigma^2}|B_t)\right| > (c_1 + k\overline{Q}')\sigma^2\right)$$

As τ shrinks the increment $B_{t+\sigma^2} - E(B_{t+\sigma^2}|B_t)$ becomes normal with standard deviation $\sqrt{\sigma^2 \cdot \sigma^2} = \sigma^2$ (by the central limit theorem). Hence, the last term exceeds some fixed probability p_1 if τ is sufficiently small. Q.E.D.$_{\text{(Claim 2)}}$

Once the state is a distance of order σ^2 from Q, there is some chance that it will reach $c\sigma^2$ in time of order σ^2, at which point a bifurcation is almost guaranteed by Claim 1. Claim 3 shows that this probability is independent of σ^2.

Claim 3 *For any $c > 0$, there are positive constants p_2 and c_2, independent of σ, and a function $\underline{\tau}_2(\cdot) > 0$, such that if $|D_v| = c_1\sigma^2$ and $\tau < \underline{\tau}_2(\sigma)$, then the probability that $|D_{v+c_2\sigma^2}| > c\sigma^2$ exceeds p_2.*

Proof The given probability exceeds the probability that $|D_{v+c_2\sigma^2}| > c\sigma^2$ *and* that D is never close to zero between the times v and $v + c_2\sigma^2$. While D remains on one side of zero, $|D|$ is a random walk with a strictly positive trend. Let $\widetilde{D}$ be a random walk with the same variance as D but with zero trend, and assume $\widetilde{D}_v = c_1\sigma^2$. The probability that $|D_{v+c_2\sigma^2}| > c\sigma^2$ exceeds the probability that $\widetilde{D}$ exceeds $c\sigma^2$ at time $v + c_2\sigma^2$ and is never close to zero between the two times. This probability is approximately $\Pr\left(\widetilde{D}_{v+c_2\sigma^4} > c\sigma^2\right) - \Pr\left(\widetilde{D}_{v+c_2\sigma^4} < -c\sigma^2\right)$, since for small τ the reflection principle can be used. But $\widetilde{D}_{v+c_2\sigma^4}$ is approximately normal for small τ, so this probability is close to $\Pr\left(N(c_1\sigma^2, c_2\sigma^4) > c\sigma^2\right) - \Pr\left(N(c_1\sigma^2, c_2\sigma^4) < -c\sigma^2\right)$, which equals $\Pr\left(\frac{c-c_1}{\sqrt{c_2}} < N(0,1) < \frac{c+c_1}{\sqrt{c_2}}\right)$, where $N(a, b)$ denotes a normal random variable with mean a and variance b. This is a positive constant that does not depend on σ. Q.E.D.$_{\text{(Claim 3)}}$

Let the times $s_0, s_1, \ldots$ be given by $s_i = t + i(\sigma^2 + c_2\sigma^2)$. The chance that the state reaches a horizontal distance $c\sigma^2$ from Q at or before time s_i is at least $1 - (1 - p_1p_2)^i$.

This holds because conditional on not reaching $c\sigma^2$ before time s_i, the probability of reaching $c\sigma^2$ by time s_{i+1} is at least $p_1 p_2$. Why? Regardless of the state's position at time s_i, at time $s_i + \sigma^2$ the probability that horizontal distance between the state and Q exceeds $c_1 \sigma^2$ is greater than p_1. Given that this happens, the probability that horizontal distance between the state and Q exceeds $c\sigma^2$ at time $s_i + \sigma^2 + c_2\sigma^2$ is greater than p_2.

As σ shrinks, we reach times s_i for arbitrarily large i before time $t + \varepsilon$. Hence, we can find $\underline{\sigma}$ small enough that the probability of reaching $c\sigma^2$ before time $t + \varepsilon$ is at least $1 - \varepsilon/2$ if $\sigma < \underline{\sigma}$ and $\tau < \min\{\underline{\tau}_1(\sigma), \underline{\tau}_2(\sigma)\}$. If, additionally, $|\mu| < \underline{\mu}$, and $\tau < \underline{\tau}_0(\sigma)$ then by Claim 1 the chance of a δ-bifurcation before time $t + \varepsilon$ must be at least $1 - \varepsilon$.

We have shown that for any $\varepsilon > 0$ and any large enough c,[33] there are $\underline{\sigma} > 0$, $\underline{\tau}(\cdot) > 0$, and $\underline{\mu} > 0$ such that if $\sigma < \underline{\sigma}$, $\tau < \underline{\tau}(\sigma)$, and $|\mu| < \underline{\mu}$, then $|D|$ reaches a distance $c\sigma^2$ with probability $1 - \varepsilon/2$ within time ε, and once it reaches this distance, a δ-bifurcation has occurred with probability at least $1 - \varepsilon/2$. This proves part 1.

Direction of Bifurcation

Since for c large enough, a bifurcation occurs with probability $1 - \varepsilon/2$ once a distance $c\sigma^2$ from Q is reached, it suffices to compute the ratio of probabilities of reaching $+c\sigma^2$ first vs. reaching $-c\sigma^2$ first. For any $\varepsilon' > 0$ we can choose σ, μ, and τ small enough that the distance $c\sigma^2$ is reached with probability $1 - \varepsilon'/2$ within time ε'. During this time, X can change by no more than $k\varepsilon'$ (taking τ to be small). Hence, the absolute value of the trend in D_t is bounded between two constants $\underline{\lambda}_R < \overline{\lambda}_R$ when D_t is positive and two other constants $\underline{\lambda}_L < \overline{\lambda}_L$ when it is negative, where these bounds converge to $\lambda_R = k(1 - X_t)|Q'(X_t)| + \mu$ and $\lambda_L = kX_t|Q'(X_t)| - \mu$, respectively, as ε' shrinks. The ratio of bifurcation probabilities lies between the ratio corresponding to constant trends $\underline{\lambda}_R$ and $\overline{\lambda}_L$ and the ratio that comes from $\overline{\lambda}_R$ and $\underline{\lambda}_L$. This is because more positive trends make D higher for any given realization of the random component of B (i.e., given any sequence of signs of B's jumps). Hence, ε' shrinks, the ratio of bifurcation probabilities converges to that computed using the constant absolute trends λ_R and

---

[33] By the proof of claim 1, c is 'large enough' if it exceeds $\ln(2/\varepsilon)/(k\underline{Q}'\delta)$.

λ_L. For the remainder of the proof we will assume constant trends and show that for large enough c, the ratio of the probabilities is approximately λ_R/λ_L, which converges to $(1 - X_t)/X_t$ as μ shrinks.

Claim 4 shows that for small τ, the relative probability of reaching $+c\sigma^2$ first vs. reaching $-c\sigma^2$ first depends little on the precise value of D_t as long as it is 'close to 0'. Let $P^\uparrow$ and $P^\downarrow$, respectively, be the probabilities of reaching $+c\sigma^2$ first and $-c\sigma^2$ first on the assumption that D_t equals 0.

Claim 5 *For any $\varepsilon' > 0$ there is a $\underline{\tau} > 0$ such that if $\tau < \underline{\tau}$ then whenever $D_t \in [-\sigma\sqrt{\tau} + \mu\tau, +\sigma\sqrt{\tau} + \mu\tau]$, the probability of reaching $+c\sigma^2$ ($-c\sigma^2$) first is within ε' of $P^\uparrow$ ($P^\downarrow$).*

Proof For any D_t in the given range and for any small $\nu > 0$, the c.d.f. Γ of $D_{t+\nu}$ is between two bounds. The lower bound is the time $t + \nu$ c.d.f. $\underline{\Gamma}$ of a random walk that starts at $-\sigma\sqrt{\tau} - |\mu|\tau$ and has the constant trend $-\lambda_L$. The upper bound is the time $t + \nu$ c.d.f. $\overline{\Gamma}$ of a random walk that starts at $\sigma\sqrt{\tau} + |\mu|\tau$ and has the trend λ_R. When τ is small, these two bounds are approximately normal with means $-\sigma\sqrt{\tau} - |\mu|\tau - \lambda_L\nu$ and $\sigma\sqrt{\tau} + |\mu|\tau + \lambda_R\nu$, respectively, and common standard deviation $\sigma\sqrt{\nu}$. For small ν, $\sqrt{\nu}/\nu$ is arbitrarily large, so the two bounds converge to each other (as ν and τ shrink, with τ sufficiently small given ν). Let $P^\uparrow(z)$ be the probability of an upwards bifurcation when $D_{t+\nu} = z$. Then the probability of an upward bifurcation is between $\int_{z=-\infty}^\infty P^\uparrow(z)d\underline{\Gamma}(z)$ and $\int_{z=-\infty}^\infty P^\uparrow(z)d\overline{\Gamma}(z)$. Since $P^\uparrow(z)$ is always between 0 and 1, the two integrals converge to each other (see Royden [20, p. 232]). Q.E.D.$_{\text{(Claim 4)}}$

The next claim shows that to compute the probability ratio $P^\uparrow/P^\downarrow$ we can focus on what happens in an arbitrarily short time interval after time t. Given any $\eta > 0$, let $P^\uparrow_\eta$ and $P^\downarrow_\eta$, respectively, be the probabilities of an upward and downward $c\sigma^2$-bifurcation occurring at some time $v \in [t, t + \eta)$, on the assumption that D_t equals 0. (A '$c\sigma^2$-bifurcation' occurs at the last time at which D is close to zero before reaching $\pm c\sigma^2$.)

Claim 5 *For any $\eta > 0$ and $\varepsilon' > 0$, there is a $\underline{\tau} > 0$ such that if $\tau < \underline{\tau}$, then $P^\uparrow/P^\downarrow$ is within ε' of $P^\uparrow_\eta/P^\downarrow_\eta$.*

Proof Let $P_\eta = P^\uparrow_\eta + P^\downarrow_\eta$. If there is no $c\sigma^2$-bifurcation in the interval $[t, t + \eta)$, then we must have D_s close to 0 for some $s \geq t + \eta$. By Claim 4, for any $\varepsilon > 0$, if τ is

31

small enough the probability of an upward $c\sigma^2$-bifurcation following time s is in the range $[P^\uparrow - \varepsilon, P^\uparrow + \varepsilon]$. Thus, $P^\uparrow$ is between $P^\uparrow_\eta + (1 - P_\eta)(P^\uparrow - \varepsilon)$ and $P^\uparrow_\eta + (1 - P_\eta)(P^\uparrow + \varepsilon)$. This shows that $P^\uparrow$ is between $\frac{P^\uparrow_\eta + \varepsilon(1 - P_\eta)}{P_\eta}$ and $\frac{P^\uparrow_\eta - \varepsilon(1 - P_\eta)}{P_\eta}$. Likewise, $P^\downarrow$ is between $\frac{P^\downarrow_\eta + \varepsilon(1 - P_\eta)}{P_\eta}$ and $\frac{P^\downarrow_\eta - \varepsilon(1 - P_\eta)}{P_\eta}$. Therefore, $P^\uparrow/P^\downarrow$ is between $\frac{P^\uparrow_\eta - \varepsilon(1 - P_\eta)}{P^\downarrow_\eta + \varepsilon(1 - P_\eta)}$ and $\frac{P^\uparrow_\eta + \varepsilon(1 - P_\eta)}{P^\downarrow_\eta - \varepsilon(1 - P_\eta)}$. But given η, we can take τ small enough that ε is negligible relative to the sum $P_\eta = P^\uparrow_\eta + P^\downarrow_\eta$. Q.E.D.(Claim 5)

By Claim 5, $\lim_{\tau \to 0} P^\uparrow/P^\downarrow = \lim_{\tau \to 0} P^\uparrow_\eta/P^\downarrow_\eta$ for any $\eta > 0$. Thus, $\lim_{\tau \to 0} P^\uparrow/P^\downarrow$ equals $\lim_{\eta \to 0} \lim_{\tau \to 0} P^\uparrow_\eta/P^\downarrow_\eta$, which we now compute. For a $c\sigma^2$-bifurcation to occur in $[t, t + \eta)$, two things must happen. First, D must not be 'close to zero' at time $t + \eta$. Second, D subsequently must reach a distance $c\sigma^2$ while remaining not 'close to zero.' Let $\Phi_{M,S}$ denote the c.d.f. of a normally distributed random variable with mean M and standard deviation S. By the law of large numbers,

$$\lim_{\tau \to 0} \Pr(D_{t+\eta} < z) \in [\Phi_{\lambda_R \eta, \sigma\sqrt{\eta}}(z), \Phi_{-\lambda_L \eta, \sigma\sqrt{\eta}}(z)] \tag{7}$$

since $D_t = 0$ by assumption. Claim 6 states that, for almost all values of $D_{t+\eta}$, the probability $P(D_{t+\eta})$ that D reaches a distance $c\sigma^2$ before once again being close to zero is approximately proportional to the initial level $D_{t+\eta}$ times the trend (which is λ_R if $D_{t+\eta}$ is positive and λ_L if negative). This approximation becomes exact as the constant c in $c\sigma^2$ grows.

Claim 6 *There is a $c_3 > 0$ such that for any $a > 0$, $\lim_{c \to \infty} \lim_{\eta \to 0} \lim_{\tau \to 0} \frac{P(a\sqrt{\eta})}{c_3 a\sqrt{\eta}\lambda_R} = 1$ and $\lim_{c \to \infty} \lim_{\eta \to 0} \lim_{\tau \to 0} \frac{P(-a\sqrt{\eta})}{c_3 a\sqrt{\eta}\lambda_L} = 1$.*

Proof We will show the first equality; the second follows analogously. In the given limit, $a\sqrt{\eta}$ becomes arbitrarily large relative to the jumps in D and arbitrarily small relative to $c\sigma^2$. Suppose $D_{t+\eta} = a\sqrt{\eta}$ and let $a' \in (0, a)$. To reach zero, D must first reach $a'\sqrt{\eta}$. Since D has stationary increments as long as it remains positive, $\lim_{\eta \to 0} \lim_{\tau \to 0} \frac{P(a\sqrt{\eta})}{P([a-a']\sqrt{\eta}) + (1 - P([a-a']\sqrt{\eta}))P(a'\sqrt{\eta})} = 1$. But $\lim_{\eta \to 0} \lim_{\tau \to 0} P(a'\sqrt{\eta}) = \lim_{\eta \to 0} \lim_{\tau \to 0} P([a - a']\sqrt{\eta}) = 0$, so $\lim_{\eta \to 0} \lim_{\tau \to 0} \frac{P(a\sqrt{\eta})}{P([a-a']\sqrt{\eta}) + P(a'\sqrt{\eta})} = 1$. Since this holds for any $a' \in (0, a)$, $P(a\sqrt{\eta})$ must become proportional to $a\sqrt{\eta}$. As for λ_R, as long as D remains positive, it is simply a random walk with constant trend λ_R. For any

32

$\ell > 1$, consider the process $\hat{D}_v = \frac{1}{\ell}D_{\ell^2 v}$. In the limit as $\tau \to 0$,

$$\text{Var}(\hat{D}_w - \hat{D}_v) = \text{Var}(\frac{1}{\ell}D_{\ell^2 w} - \frac{1}{\ell}D_{\ell^2 v}) = \frac{1}{\ell^2}\sigma^2\left(\ell^2 w - \ell^2 v\right) = \sigma^2(w - v) = \text{Var}(D_w - D_v)$$

By the change of variables, $P(a\sqrt{\eta})$ equals the probability that, given that $\hat{D}_{\ell^2(t+\eta)} = a\sqrt{\eta}/\ell$, $\hat{D}$ subsequently reaches $c\sigma^2/\ell$ before zero. The trend of $\hat{D}$ is $\ell\lambda$. But

$$\lim_{c\to\infty}\lim_{\eta\to 0}\lim_{\tau\to 0}\text{Pr}(\hat{D} \text{ reaches } c\sigma^2 \text{ before } 0 \mid \hat{D} \text{ starts at } c\sigma^2/\ell) =$$
$$\lim_{c\to\infty}\lim_{\eta\to 0}\lim_{\tau\to 0}\text{Pr}(D \text{ reaches } \ell c\sigma^2 \text{ before } 0 \mid D \text{ starts at } c\sigma^2) = 1$$

since the probability in the second line is greater than the probability that D will *never* return to zero if it starts at $c\sigma^2$, which converges to $1 - e^{-2c\lambda_R}$ as $\tau \to 0$ (see Karlin and Taylor [16]). So

$$\lim_{c\to\infty}\lim_{\eta\to 0}\lim_{\tau\to 0}\frac{\text{Pr}(\widehat{D} \text{ reaches } c\sigma^2 \text{ before } 0 \mid \hat{D}_{\ell^2(t+\eta)} = a\sqrt{\eta}/\ell)}{P(a\sqrt{\eta})} = 1$$

But by prior arguments,

$$\lim_{\eta\to 0}\lim_{\tau\to 0}\frac{\text{Pr}(\widehat{D} \text{ reaches } c\sigma^2 \text{ before } 0 \mid \hat{D}_{\ell^2(t+\eta)} = a\sqrt{\eta}/\ell)}{\frac{1}{\ell}\cdot\text{Pr}(\widehat{D} \text{ reaches } c\sigma^2 \text{ before } 0 \mid \hat{D}_{\ell^2(t+\eta)} = a\sqrt{\eta})} = 1$$

so that

$$\lim_{c\to\infty}\lim_{\eta\to 0}\lim_{\tau\to 0}\frac{\ell\cdot P(a\sqrt{\eta})}{\text{Pr}(\widehat{D} \text{ reaches } c\sigma^2 \text{ before } 0 \mid \hat{D}_{\ell^2(t+\eta)} = a\sqrt{\eta})} = 1$$

Since $\widehat{D}$ has the variance of D but trend $\ell\lambda_R$, $P(a\sqrt{\eta})$ must become proportional to the trend as well. Q.E.D.$_{\text{(Claim 6)}}$

By Claim 6,

$$\lim_{c\to\infty}\lim_{\eta\to 0}\lim_{\tau\to 0}\frac{P_\eta^\uparrow}{P_\eta^\downarrow} = \lim_{c\to\infty}\lim_{\eta\to 0}\lim_{\tau\to 0}\frac{\int_{a=0}^\infty P(a\sqrt{\eta})d\Psi_\eta(a\sqrt{\eta})}{\int_{a=0}^\infty P(-a\sqrt{\eta})d\Psi_\eta(-a\sqrt{\eta})}$$

where Ψ_η is the c.d.f. of $D_{t+\eta}$. Since $\lim_{\eta\to 0}\sqrt{\eta}/\eta = \infty$ and by (7), Ψ_η converges to a c.d.f. of a normally distributed variable with mean 0 and standard deviation $\sigma\sqrt{\eta}$. Hence, by Claim 6, $\lim_{c\to\infty}\lim_{\eta\to 0}\lim_{\tau\to 0}P_\eta^\uparrow/P_\eta^\downarrow = \lambda_R/\lambda_L$, so $\lim_{c\to\infty}\lim_{\tau\to 0}P^\uparrow/P^\downarrow = \lambda_R/\lambda_L$. This shows that for c large enough, the relative probability of reaching $+c\sigma^2$ vs. $-c\sigma^2$ first with constant trends λ_R and λ_L is approximately λ_R/λ_L for small τ. Q.E.D.$_{\text{(Lemma 2)}}$

Proof of Lemma 1

We replace B_t with $\hat{B}_t = B_{t/k}$ and X_t with $\hat{X}_t = X_{t/k}$. The variance of $\hat{B}$ is σ^2/k and the trend is μ/k. Since these go to zero as k grows, the odds of an upward vs. a downward bifurcation are as in Lemma 2. A bifurcation happens with probability $1 - \varepsilon$ before time ε in the new time units; in the original time units, it happens with probability $1 - \varepsilon$ before time ε/k. The time between periods is τk in the new time units. To apply the prior argument we require $\tau k < \underline{\tau}^{\text{Lem 2}}(\sigma)$ (where $\underline{\tau}^{\text{Lem 2}}$ denotes the $\underline{\tau}$ of Lemma 2), so we can let $\underline{\tau}(k) = \underline{\tau}^{\text{Lem 2}}(\sigma)/k$. Since σ is fixed in Lemma 2, $\underline{\tau}(k)$ is a function of k alone.

References

[1] Billingsley, P. *Convergence of Probability Measures*. New York: Wiley, 1968.

[2] Blume, L. 1993. "The Statistical Mechanics of Strategic Interaction." *Games and Economic Behavior* 5:387-424.

[3] Boylan, R.T. 1992. "Laws of Large Numbers for Dynamical Systems with Randomly Matched Individuals." *Journal of Economics Theory* 57:473-504.

[4] Burdzy, K., D.M. Frankel, and A. Pauzner. 1997. "Fast Equilibrium Selection by Rational Players Living in a Changing World." Foerder Institute w.p. 7-97.

[5] Burdzy, Krzysztof, David M. Frankel, and Ady Pauzner. 1998. "On the Time and Direction of Stochastic Bifurcation." In B. Szyszkowicz, ed. *Asymptotic Methods in Probability and Statistics: A Volume in Honour of Miklos Csorgo*. Holland: Elsevier.

[6] Carlsson, H., and E. van Damme. 1993. "Global Games and Equilibrium Selection." *Econometrica* 61:989-1018.

[7] Ellison, G. 1993. "Learning, Local Interaction, and Coordination." *Econometrica* 61:1047-1071.

[8] Foster, D., and P.H. Young. 1990. "Stochastic Evolutionary Game Dynamics." *Theoretical Population Biology* 38:219-232.

[9] Frankel, David M., and Ady Pauzner. 1999. "Resolving Indeterminacy in Dynamic Settings: the Role of Shocks." Forthcoming, *Quarterly Journal of Economics*. Available at http://econ.tau.ac.il/~ady/papers.html

[10] Fudenberg, D., and J. Tirole. 1991. *Game Theory*. Cambridge: MIT Press.

[11] Fudenberg, D., and C. Harris. 1992. "Evolutionary Dynamics with Aggregate Shocks." *Journal of Economic Theory* 57:420-441.

[12] Gilboa, I., and A. Matsui. 1992. "A Model of Random Matching." *Journal of Mathematical Economics* 21:185-197.

[13] Harsanyi, J., and R. Selten. 1988. *A General Theory of Equilibrium Selection in Games*. Cambridge: MIT Press.

[14] Judd, K.L. 1985. "The Law of Large Numbers with a Continuum of IID Random Variables." *Journal of Economic Theory* 35:19-25.

[15] Kandori, M., G.J. Mailath, and R. Rob. 1993. "Learning, Mutation, and Long Run Equilibria in Games." *Econometrica* 61:29-56.

[16] Karlin, S., and H.M. Taylor. *A First Course in Stochastic Processes, 2nd ed.* New York: Academic Press, 1975.

[17] Matsui, A.. 1998. "Multiple Lenders and Credit Squeeze." Mimeo, U. Tokyo.

[18] Matsui, A., and K. Matsuyama. 1995. "An Approach to Equilibrium Selection." *Journal of Economic Theory* 65:415-434.

[19] Morris, S., R. Rob, and H.S. Shin. 1995. "p-Dominance and Belief Potential." *Econometrica* 63:145-157.

[20] Royden, H. L. 1968. *Real Analysis (2nd ed.)* Toronto: MacMillan.

[21] Young, P. 1993. "The Evolution of Conventions." *Econometrica* 61:57-84.