

Fast Equilibrium Selection by Rational Players Living in a Changing World

Krzysztof Burdzy,
David M. Frankel,
and Ady Pauzner¹

November 18, 1998

¹Burdzy: Department of Mathematics, University of Washington, Seattle, WA 98195, USA (e-mail: burdzy@math.washington.edu). Frankel and Pauzner: Eitan Berglas School of Economics, Tel Aviv University, Ramat Aviv, Tel Aviv 69978, Israel (e-mail: dfrankel@econ.tau.ac.il, ady@econ.tau.ac.il.). We are grateful for the helpful comments of Elchanan Ben-Porath, Eddie Dekel, Drew Fudenberg, Philip Reny, Ariel Rubenstein, 3 anonymous referees, and seminar participants at Berkeley, Brown, Cornell, Harvard, Hebrew U., LSE, Northwestern, Penn, Pittsburgh, Princeton, Stanford, Tel Aviv, the Technion, and U. Washington. Burdzy was supported in part by NSF grant 9322689.

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 over time according to some Brownian motion. 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.

1 Introduction

Games with multiple strict Nash equilibria present a major difficulty 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. Nonuniqueness not only limits the predictive power of the theory; it casts doubt on its entire validity. Without an explanation of how players reach coordination, it is hard to justify why equilibrium play is, at all, a reasonable prediction.

The difficulty of generating predictions from the normal form game need not be considered a failure of game theory. It may simply indicate that the game sometimes needs to be specified in greater detail. One important goal of game theory is to find out which details are most relevant in determining whether equilibrium behavior is reasonable and, if so, which equilibrium is the better prediction. A way to approach this goal is to study the models that result from filling in the missing details of the normal form game in various ways.

In *population models*, a large population is randomly matched, from time to time, to play a given normal-form game. Rather than simply postulating that players guess correctly what others will do, these models provide a dynamic framework in which players can observe the actions of others in the process of picking their own. This gives rise to an interactive process that can potentially lead to equilibrium play. The study of population models also enlarges the predictive scope of game theory to the process of convergence to equilibrium play itself, rather than limiting it to the ultimate result. A number of important issues can be studied, including rates of convergence and the role of initial conditions.

Most of the population literature has focused on the evolution of play in a fixed world. In this paper we study the evolution of play when the world changes over time. In many of the natural applications of population models, 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.

We model the changing world using an exogenous stochastic process that affects the payoff matrix of the static game. We assume that, in the (perhaps very distant) future, these changes in payoffs have the potential to make any action strictly dominant. This leads to an unraveling effect that yields sharp predictions about what players will do in the present. With small frictions in changing actions, *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 most population models, in which the selection either depends on initial conditions or is only a long run prediction, taking the form of an ergodic distribution with most of its weight put on one equilibrium.

Players in our model are fully rational. Models with rational players have typically been solved using the assumption of *rational expectations equilibrium*². In contrast, our result is established using only iterated conditional dominance, an extension of backwards induction to infinite horizon games.

One should not consider our use of iterated conditional dominance as a refinement of Nash equilibrium; in our model, every Nash equilibrium outcome survives the iterative procedure. Rather, we use iterated dominance so as to avoid the assumption of equilibrium play: i.e., that players' strategies happen to be best responses to each other. A primary motivation for looking at population models is to explain why players playing a normal form game coordinate on an equilibrium. Assuming equilibrium play in the dynamic game would only shift the problem of justifying coordination from one model to another. In contrast, iterated dominance traces a plausible process by which players reason about what others will do. In our model, this process alone leads players to coordinate their beliefs. Accordingly, our use of weaker assumptions in the dynamic game

²This is an assignment of strategies and beliefs to each player, such that each player's belief over the expected evolution of the environment is correct given everybody's strategies, and such that each player's strategy is optimal given that belief.

helps justify the prediction of equilibrium play in the normal form game.

This paper is closely related to papers by Carlsson and van Damme [10] and Matsui and Matsuyama [24]. In the former, two players play a one shot game of incomplete information. Each receives a slightly noisy signal of the game's payoffs. Iterated strict dominance leads to the selection of the risk-dominant equilibrium through a contagion argument. In Matsui and Matsuyama, a large population of rational players is randomly matched to play a fixed 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 leads to equilibrium selection through a contagion argument that is akin to that of Carlsson and van Damme. The relations among these three papers are discussed in detail in section 5.

The rest of this paper is organized as follows. The model and results are presented in sections 2 and 3. We explain the intuition for the results in section 4. Section 5 is a literature review. Section 6 concludes with a discussion of how the results depend on the various features of the model. The formal proofs are collected in an appendix, with references to mathematical results that appear in Burdzy, Frankel, and Pauzner [9].

2 The Model

The static game we study is a 2×2 , symmetric coordination game, with two strict Nash equilibria. The population literature has typically focussed on these games since, while being extremely simple, they exhibit the problems associated with multiple equilibria to their full extent. In addition, many economic interactions can be analyzed using models of large populations that are randomly matched to play such games.

To understand our model, it may be helpful to keep in mind the following story. The professors in a certain university work on two types of computers: IBM and Macintosh. From time to time, two professors need to share files. Their payoffs from this interaction

depend on the type of computer that each has; in particular, imperfect compatibility creates a benefit to coordination. The following table gives an example.

	IBM	Macintosh
IBM	3, 3	2, 0
Macintosh	0, 2	4, 4

Occasionally, a professor's computer breaks down and she must choose a new one. To decide which to buy, she needs to consider several issues. At the current state of technology, how do the two computers compare? How many other professors currently use each standard? Which are they likely to choose when their own computers break down? How will their choices be affected by possible technological developments and by their own predictions of how people will choose after them?

One may wonder whether the professors eventually coordinate on a given computer standard and, if so, on which and how soon. These are the types of questions we seek to answer in our more general model. We first describe the static game, which specifies the players' payoffs from a single match. We then present the dynamic context in which players are matched and choose their actions.

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, t : 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 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 continuous and strictly increasing in B_t at a bounded rate: there exists a positive constant \bar{w} such that for all $b > \hat{b}$,

$$0 < \Delta(a, b) - \Delta(a, \hat{b}) < \bar{w}[b - \hat{b}] \quad (1)$$

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	$0, 2 + B_t$	$4, 4$

We assume that B_t follows a Brownian motion. This is essentially the continuous time version of a random walk and may also have a deterministic trend. The Brownian motion has two parameters. The variance σ^2 tells us how fast the Brownian motion spreads out. The trend μ gives the rate at which its mean changes over time. More precisely, a Brownian motion has the following properties (Billingsley [3, p. 522]):

1. It is continuous with probability one.
2. For any $t > \hat{t} > 0$, the random variable $B_t - B_{\hat{t}}$ (which takes values in \Re) is normally distributed with mean $\mu(t - \hat{t})$ and variance $\sigma^2(t - \hat{t})$.
3. Its increments are independent. For any $t > \hat{t} \geq v > \hat{v}$, the random variable $B_t - B_{\hat{t}}$ is independent of $B_v - B_{\hat{v}}$.

In our game, an action is *p-dominant* (Morris, Rob and Shin [25]) 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*.) For example, in the above game, when $B_t = 0$, R is exactly 0.4-dominant. Clearly, R is exactly p -dominant if and only if L is exactly $(1 - p)$ -dominant.

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

⁴For 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.

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:

$$pu(R, R, B^p) + (1 - p)u(R, L, B^p) = pu(L, R, B^p) + (1 - p)u(L, L, B^p) \quad (2)$$

In the above game, for instance, $0 = B^{0.4}$. 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 [16]) 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 $B_t \leq B^{1/2}$.

2.2 The Dynamic Context

A continuum of players are randomly matched from time to time to play the static game. A player's matches arrive according to a Poisson process with common arrival rate $m > 0$. When a player is matched, she cannot instantaneously change actions. Rather, she is locked into an action she chose before. Her opportunities to revise actions arrive according to a Poisson process with common arrival rate k . When k is high we say that frictions are small, since players are locked into their actions for less time. We assume that all of the Poisson processes are independent and that there is no aggregate uncertainty.⁵

The players observe the evolution of both B_t , which we call the “state of the world”, and X_t , the proportion of players currently committed to playing R (rather than L), which we call the “state of play”. We refer to 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 that time, $(B_v, X_v)_{v \in [0, t]}$. A player's *private history* at time t consists of her actions and the details of her matches up to time t . A player's information set at time t is given by the public history, together with her private history. *Strategies* are functions from the

⁵Judd [17] discusses some technical problems that arise with a continuum of i.i.d. variables. Boylan [7], and Gilboa and Matsui [15] offer possible solutions in the context of random matching.

set of all information sets to the action set $\{R, L\}$ that indicate the action a player will choose at any information set, should she have an action revision opportunity.⁶

Note that the actions of any player will be observed by only a countable number of other players. Since there is a continuum of players, the probability is zero that a player's past actions or those of her opponents will have any effect on the actions taken by any future opponent. This means that a player's payoffs from different actions can in no way depend on her private history. For this reason we may assume, without loss of generality, that a strategy is simply a map from the set H of all possible public histories⁷ to the action set $\{R, L\}$.

Suppose an agent receives an opportunity to revise her action at time t , after the history h_t . Denote the (realized) times of her subsequent matches by t_1, t_2, \dots . Suppose that at the time of the agent's n th subsequent match the state of the world is B_{t_n} , the agent is playing the action $a_n \in \{R, L\}$, and the agent's partner in that match is playing $b_n \in \{R, L\}$. Then the agent's time t continuation payoff is

$$\sum_{n=0}^{\infty} e^{-r(t_n-t)} u(a_n, b_n, B_{t_n}) \quad (3)$$

where $r \geq 0$ is the constant discount rate.⁸ When a player has a revision opportunity, she maximizes the expectation of (3) with respect to the probability distribution over paths $(B_v)_{v>t}$ and her beliefs about the path of play $(X_v)_{v>t}$ that will result from any given realization of $(B_v)_{v>t}$.

Because a single player has no influence over which path $(X_v)_{v>t}$ will occur, the best she can do is to pick the action that maximizes her discounted payoff for the (random)

⁶The restriction to pure strategies is without loss of generality. Our iterated dominance argument applies almost unchanged if a player can choose a mixed strategy.

⁷This is the set of all paths $(B_v, X_v)_{v \in [0, t]}$, for any t and (B_0, X_0) .

⁸Alternatively, the player's payoff may depend on the state of the world at the time that she picked her action, B_t , rather than on the states at the times of the matches, B_{t_n} . Such a specification may be more appropriate in our motivating example, since the properties of a computer depend largely on the state of technology at the time of purchase. All of our results hold for this specification as well.

period in which she will be committed to that action. In any period $[v, v + dv]$, a player is matched with probability mdv to an opponent who plays R with probability X_v and L with probability $1 - X_v$. With probability $e^{-k(v-t)}$, the player is still locked into the action she chose at time t . Since her pure discount factor is $e^{-r(v-t)}$, her effective discount factor is $e^{-(k+r)(v-t)}$. Therefore, the relative payoff to playing R is:

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

A player chooses R if this is positive and L if it is negative.

Finally, to give our iterated dominance argument a place to start, we assume that for B_t large enough, R is strictly dominant: that $E \left[\int_{v=t}^{\infty} e^{-(r+k)(v-t)} \Delta(L, B_v) dv \mid B_t \right] > 0$. For sufficiently low values of B_t , an analogous condition makes L 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 [13, pp. 128 ff.]). This is essentially the extension of backwards induction to infinite horizon games.

It is important to note that, in our model, iterated conditional dominance is not a refinement of Nash equilibrium.⁹ This is because players are small, 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. But every subgame perfect equilibrium survives iterated conditional dominance.¹⁰

⁹We thank Philip Reny for suggesting the argument that follows.

¹⁰To be more precise, let s be a Nash equilibrium strategy profile, and let $s(b, x)$ denote the play pre-

In our model, the iterative procedure works as follows (ignoring some technical points that are addressed in the proofs). 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. This is depicted in Figure 1. To the right of f^0 , we know that the player must play R; to the left of f^0 we cannot yet say what the player does. In the first step we eliminate all the strategies in which a player plays 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 play R when they are to the right of f^0 . With this belief, there is a new boundary, f^1 , such that a player must play R when she is to the right of f^1 . f^1 must lie to the left of f^0 , since knowing that other players will sometimes play R makes 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, \dots . Whenever (B_t, X_t) is on the right side of F , any player who is called to act must play 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 bound G , such that

scribed by s for histories that begin with $(B_0, X_0) = (b, x)$. We construct a subgame perfect equilibrium \hat{s} that has the same distribution of equilibrium paths as s as follows. If h_t is consistent with s , players continue to play according to s . Otherwise, let $v \leq t$ be the earliest time such that h_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 \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.

any strategy that survives the iterative elimination prescribes playing L to the left of G .

The iterative elimination procedure divides the environment space into three regions. In all the surviving strategies, R is played to the right of F and L is played to the left of G . We do not know what happens in the “?” region between the two lines; 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

Our main result states that as frictions disappear, the “?” region shrinks to 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 ($k \rightarrow \infty$).

Theorem 1

Fix σ , μ , and r . For any $\epsilon > 0$, there is a \underline{k} such that if $k > \underline{k}$,

R must be played whenever $B_t > B^{1/2} + \epsilon$ and L whenever $B_t < B^{1/2} - \epsilon$.

Proof See appendix.

This result has strong implications for the evolution of aggregate play in the population. Consider the initial normal-form game, and assume that the actions L and R are not exactly 1/2 dominant. If frictions are small enough, the population will immediately converge to the action that is risk-dominant, regardless of the initial state of play

X_0 . 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.

Theorem 2 considers the case in which lock-in remains positive and instead the world changes more and more slowly. Once again, the (B_t, X_t) 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 strictly downward sloping. 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 initial state of play determines the equilibrium on which players eventually coordinate. If enough players are initially 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

Fix k and r , and let $h(x) = \frac{rx+k}{r+2k}$.

For any $\epsilon > 0$, there is a $\bar{\sigma} > 0$ and a $\bar{\mu} > 0$ such that if $\sigma < \bar{\sigma}$ and $\mu < \bar{\mu}$,

R must be played whenever $B_t > B^{h(X_t)} + \epsilon$ and L whenever $B_t < B^{h(X_t)} - \epsilon$.

As players become more and more patient ($r \rightarrow 0$), $h(x) \rightarrow 1/2$, as in Theorem 1.¹¹

Proof See appendix.

¹¹Theorem 2 holds k fixed and takes σ and μ to zero. This is actually not necessary; the result holds whenever k/σ and k/μ both go to infinity. This generalization of Theorem 2 implies Theorem 1 since if k goes to ∞ , $h(x)$ goes to $1/2$.

4 Intuition

Outline

We first explain the intuition for the case of vanishing frictions, and then discuss how things change with positive frictions and a slowly changing world. Why must F (and, analogously, G) converge to a vertical line at $B^{1/2}$? Consider a player who chooses an action at time t at some point on F and who believes that all other players will choose R to the right of F and L to the left. We show that such a player must be indifferent between R and L. We can then infer B_t by counting the proportion p of R opponents she expects to meet: since she is indifferent, B_t must equal B^p .

We then show that, if $X_t = 1$, the player expects to meet at least one half R players. Since she is indifferent, the upper endpoint of F must be to the left of $B^{1/2}$. Similarly, the lower endpoint must lie to the right. Hence, if F were not identically $B^{1/2}$, there would be some point where F had a finite, negative slope and was not equal to $B^{1/2}$.¹² The heart of the proof shows that whenever F has such a slope, the player expects to meet exactly one half R players. Since she is indifferent, F must in fact equal $B^{1/2}$.

The Argument in Greater Detail

Let us say that a player plays “according to F ” if she always chooses L to the left of F and R to the right. If a player believes that all other players play according to F , then she also wants to. Why? To compute f^n from f^{n-1} we use the belief that makes R the least desirable: that all players play L to the left of f^{n-1} (and R to the right). This is just the belief that other players play according to f^{n-1} . With this belief, a player wants to play according to f^n . Since one more iteration from F gives F , a player who believes that all other players play according to F will also want to. A continuity argument implies that, if the player is on F , she is indifferent between R and L.

Suppose our player revises her action at time t . While she is locked into this action

¹²In the intuition we assume that F is continuous. In fact, we don't know this about F , and the formal proof does not assume it.

she will be matched to a sequence of players, whom we divide into two groups: “new” players and “old” players. A player is “old” if, at the time of his match with her, he is locked into an action that he chose before time t . A “new” player is one who chose his action after time t . The expected number of “old” players equals the expected number of “new” players. Why? Over a player’s lifetime, half of the players she is matched with will have chosen their actions before she chose hers, because all players have the same rate of revision opportunities. But a lifetime is just a sequence of commitments to actions.¹³ Since the period of commitment that begins at time t is, *ex ante*, the same as any other, the expected numbers of “new” and “old” players in the period must also be equal.¹⁴

The probability that an “old” player plays R is given by X_t . This immediately gives bounds on $F(0)$ and $F(1)$. (We now treat F as a function from X_t to B_t .) Suppose that a player chooses an action on F when $X_t = 1$. Since all old players are playing R, she cannot expect to meet more than one half L players. Since she is indifferent on F , she must be to the left of $B^{1/2}$. This shows that $F(1) \leq B^{1/2}$. Analogously, $F(0) \geq B^{1/2}$.

Now suppose that, in the limit as frictions vanish, F is not a vertical line at $B^{1/2}$. Then there must be a point away from $B^{1/2}$ where F has a finite, negative slope. Suppose a player receives an action revision opportunity on F at this point at time t . Assume she believes that all other players play according to F . As argued above, she is indifferent between R and L. Let’s count the proportion of R players she expects to meet while locked into her action. Half of them will be old players; of these, X_t play R. Half of them will be new players; we will show that the probability that a new player plays

¹³For a discussion, see Hare Krishna (1742).

¹⁴More formally, let O_i and N_i be the (random) numbers of old and new players that a player will meet while locked into her i th action. Over a player’s lifetime she meets equal proportions of new and old players: $\text{plim}_{I \rightarrow \infty} \left[\sum_{i=1}^I N_i / \sum_{i=1}^I O_i \right] = 1$. Divide both the numerator and the denominator by I . Since the expectation of N_i is independent of i , the law of large numbers implies that $\text{plim}_{I \rightarrow \infty} \left[\sum_{i=1}^I N_i \right] / I = E(N_i)$, and likewise for O_i . Thus, $E(N_i) = E(O_i)$.

R is $1 - X_t$ in the limit. Therefore, the total proportion of R players converges to $X_t/2 + (1 - X_t)/2 = 1/2$. Since the player is indifferent, $F(X_t)$ must be close to $B^{1/2}$ when k is large.

How does a player forecast how the new players will play? To answer this, we need to know how her belief over other players strategies combines with the exogenous stochastic process B to give a prediction of how the state of play X will evolve. Suppose all players do play according to F . Then X satisfies the following differential equation:

$$\dot{X}_t = \begin{cases} k(1 - X_t) & \text{if } B_t > F(X_t) \\ -kX_t & \text{if } B_t < F(X_t) \end{cases}$$

Why? When $B_t > F(X_t)$, all players currently playing L switch to R when they have the chance. The proportion of L players is $1 - X_t$, and they get chances to change actions at the rate k , so $\dot{X}_t = k(1 - X_t)$. Similarly, when $B_t < F(X_t)$, players switch from R to L, and the proportion of R players is X_t , so $\dot{X}_t = -kX_t$.




Figure 4: Local dynamics around F .

Figure 4 illustrates the dynamics of the system. Suppose we are initially on F . If B increases slightly, players start switching to R, so X rises. This causes us to move away from F . X will continue to rise unless a reverse movement of B brings us to the other side of the indifference line. If this happens, X will begin to fall—again, until a sufficiently large reverse movement of B takes us back to F . The longer we stay on one side of F , the less chance that we will return, since the changes in X always take us further away from F . Thus, sooner or later there will be a *bifurcation*: a time beyond

which we stay on one side of the line until X (almost) reaches either 0 or 1.¹⁵

As frictions shrink to zero, two things happen. First, the bifurcation happens almost instantaneously, so that almost all new players choose their actions *after* the bifurcation. This means that either nearly all new players choose R (if there is an upwards bifurcation) or nearly all choose L (if the bifurcation is downwards). Second, the ratio of the probabilities of bifurcating up vs. down converges to the ratio of the speeds at which the population moves to R versus L on the two sides of F . This ratio is just $(1 - X_t)/X_t$. The intuition for these two properties is explained below.

The probability that an “old” player plays R is given by X_t . With small frictions, the probability that a “new” player will play R is just the probability that X bifurcates upwards, which is $1 - X_t$. In total, the player must believe that her probability of being matched to an R player is

$$\frac{1}{2} \cdot X_t + \frac{1}{2} \cdot (1 - X_t) = \frac{1}{2} \quad (4)$$

Since the player is indifferent between R and L given her belief, $F(X_t)$ cannot be different from $B^{1/2}$ in the limit as k goes to infinity. This completes the intuition for Theorem 1.

The Case of a Slowly Changing World

The intuition for Theorem 2 is exactly the same, with one exception. With non-vanishing frictions, a player is locked into her action for a positive amount of time. 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 (4), she puts a weight of more than one half on the X_t and less than one half on the $1 - X_t$. Thus, a higher X_t makes R a more appealing choice. As a result, F is not vertical but rather slopes downward. 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 .

¹⁵Note that X is never exactly equal to 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.

Intuition for Bifurcation Properties

We first need to show that *the bifurcation occurs almost instantly*. This is easiest to see in the second case, of fixed frictions and a slowly changing world. After a short time on either side of F , X takes us away from the line at a fixed rate; the slower B is, the smaller is the chance that a reverse movement in B will take us back, and thus the earlier is the bifurcation.

In the other case, of shrinking frictions, we need to show something stronger. In order for the player to ignore how new players play before the bifurcation, the bifurcation must happen very early *relative* to the time $1/k$ that she expects to be locked into her action, which itself shrinks to zero as frictions disappear. To show this, we stretch the time scale so as to keep the expected lock-in time fixed. (I.e., we replace B_t with $\hat{B}_t = B_{t/k}$ and X_t with $\hat{X}_t = X_{t/k}$.) The Brownian motion in the new scaling becomes slower and slower, whereupon the prior argument can be applied.

The second thing we need to show is that *the relative chance of bifurcating up vs. down equals $(1 - X_t)/X_t$* . To see this, it helps to transform the problem to one dimension: the horizontal distance between B_t and $F(X_t)$. Locally, F is approximately a straight line. Thus, the distance $D_t = B_t - F(X_t)$ looks locally like a Brownian motion with two different constant trends, each pulling D away from zero (see Figure 5). Denote the trends pulling D to the right and left by $\lambda_R = k(1 - X_t)|F'(X_t)| + \mu$ and $\lambda_L = kX_t|F'(X_t)| - \mu$.




Figure 5: Linearization of $D_v = B_v - F(X_v)$ around $v = t$.

In this approximation, a bifurcation occurs at time t if D_t equals zero but D_v is nonzero for all $v > t$. An upwards bifurcation in (B_t, X_t) space corresponds to a positive bifurcation of D . We will show that the ratio of probabilities of positive and negative

bifurcations of D is approximately the ratio of the respective trends, λ_R/λ_L , which converges to $(1 - X_t)/X_t$ as k grows.

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. Let P_ϵ^\uparrow and P_ϵ^\downarrow , respectively, be the probabilities of a positive and negative bifurcation occurring at some time $v \in [t, t + \epsilon]$. We claim that the ratio P^\uparrow/P^\downarrow of bifurcation probabilities equals the ratio $P_\epsilon^\uparrow/P_\epsilon^\downarrow$. Why? Let $P_\epsilon = P_\epsilon^\uparrow + P_\epsilon^\downarrow$. If there is no bifurcation in the interval $[t, t + \epsilon]$, then we must have $D_v = 0$ for some $v > t + \epsilon$. As of time v , the probability of an upwards bifurcation is once again P^\uparrow . Thus, $P^\uparrow = P_\epsilon^\uparrow + (1 - P_\epsilon)P^\uparrow$. This shows that $P^\uparrow = P_\epsilon^\uparrow/P_\epsilon$. Likewise, $P^\downarrow = P_\epsilon^\downarrow/P_\epsilon$. Therefore, $P^\uparrow/P^\downarrow = P_\epsilon^\uparrow/P_\epsilon^\downarrow$.

Since this equality holds for all ϵ , it also holds as ϵ goes to zero. This limit turns out to be easy to compute. For an upwards bifurcation to occur in $[t, t + \epsilon]$, two things must happen. First, D must be positive at time $t + \epsilon$. Second, D must remain positive forever after. Where is D at $t + \epsilon$? Since $D_t = 0$, the value of D at time $t + \epsilon$ is dominated by the noise in B . This is because the standard deviation of $B_{t+\epsilon} - B_t$ is proportional to $\sqrt{\epsilon}$, while the linear trends λ_R and λ_L produce a change of order only ϵ . For small ϵ , $\sqrt{\epsilon}/\epsilon$ is arbitrarily large. So for small ϵ we can treat the distribution Ψ_ϵ of $D_{t+\epsilon}$ as approximately symmetric around zero. (In particular, it is approximately normal with mean 0 and variance $\sigma^2\epsilon$).¹⁶

Given that D is positive at $t + \epsilon$, what is the probability that it remains so forever after? If $D_{t+\epsilon} = z$ and z is small, this probability is proportional to the distance z times the trend λ_R . To see this, let p be this probability, and consider what happens if $D_{t+\epsilon}$ starts twice as far away, at $2z$. What is the probability that D never hits zero?

¹⁶This property of Brownian motions, that the noise swamps the trend over short intervals, is an implication of independent increments. The change in the Brownian motion over a given interval of length, say, 1, is the sum of N i.i.d. changes over intervals of length $1/N$. The only way this sum can retain a nontrivial variance as N grows is for the variance over each subinterval to remain relatively large; i.e., at least proportional to $1/N$.

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 proportional to z for small z .

Why is this probability is also linear in λ_R ? Note that as long as $D > 0$, D is simply a Brownian motion with trend λ_R . Let us multiply the time scale by 4 and the space scale (the horizontal axis) by 2. This gives a new Brownian motion $D_v = 2D_{v/4}$, which has the same variance as the old one:

$$\text{Var}(\hat{D}_w - \hat{D}_v) = \text{Var}(2D_{w/4} - 2D_{v/4}) = 4\sigma^2(w/4 - v/4) = \sigma^2(w - v) = \text{Var}(D_w - D_v)$$

The new process \hat{D} begins at $2z$. Since the time scale is stretched by twice the space scale, the trend of \hat{D} is $\lambda_R/2$, half the trend of D . But a change in the scaling cannot affect the probability that \hat{D} never hits 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 $D_{t+\epsilon} = z$ and z is small, the probability that D_v remains positive for all $v > t + \epsilon$ is proportional to $z\lambda_R$. Since most of the weight of Ψ_ϵ is on small z 's, the probability P_ϵ^\uparrow is approximately proportional to

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

The probability P_ϵ^\downarrow that D bifurcates to the left in $[t, t + \epsilon]$ is approximately proportional to

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

because Ψ_ϵ is approximately symmetric. Therefore, the ratio $P_\epsilon^\uparrow/P_\epsilon^\downarrow$ equals λ_R/λ_L as ϵ goes to zero. Since $\lambda_R = k(1 - X_t)|F'(X_t)| + \mu$ and $\lambda_L = kX_t|F'(X_t)| - \mu$, this ratio converges to $(1 - X_t)/X_t$ as k goes to infinity. This shows that the relative probability of bifurcating up vs. down equals $(1 - X_t)/X_t$ as frictions vanish.

5 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 population models.

The first program began with Nash [26, 27, 28]. He showed that his bargaining game has a unique solution if one imposes conditions on how the solution can vary when the game is changed.¹⁷ Later contributions used a strategic rather than axiomatic approach; important examples include Carlsson and van Damme [10] and Morris, Rob and Shin [25].

Our framework is closely related to Carlsson and van Damme’s. They study a one-shot 2×2 game whose payoffs are not common knowledge. Rather, each player receives a noisy signal of the true payoffs. The space of possible payoffs includes regions where each action is strictly dominant. Because of the noise, there is no common knowledge among the players that the true game is *not* in one of these regions. Iterated 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.

In both their paper and ours, the connection between different decision problems gives rise to a contagion effect. A player’s optimal action depends on what she thinks her opponent will do. In Carlsson and van Damme, a player does not know exactly which signal her opponent received, so she must take into account a distribution of possible types of opponents. The action of each of these opponents is the solution to a decision problem in a slightly different game. In our game, a player does not know what the state of the environment will be when her future opponent chooses his own action. Again, there is a distribution of possible “types” of opponents, each seeing a different state.

However, it is not possible to apply the technique of Carlsson and van Damme directly

¹⁷See Carlsson and van Damme [10, pp. 1007-1008] for a discussion.

to our framework. A naive attempt to do iterated dominance on the payoff variable B_t alone leads to a much 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 $B^{1/4}$. This is because the iterative procedure tells us only how new players will play, and these constitute just half of a player’s potential opponents.¹⁸

The second research program, on population models, divides into three branches: “mechanical”, “boundedly rational” and “rational”. The “mechanical” branch (e.g., Foster and Young [12], Fudenberg and Harris [14]) postulates some law of motion for the whole population, such as the replicator dynamics. While well suited to biological evolution, this paradigm has been criticized as an economic model of human behavior because it ignores the individual’s decision problem. In contrast, the “boundedly rational” branch derives the population’s law of motion from a rule of thumb used by individual players. The rule of thumb need not be optimal but has to be “reasonable”. For example, players might play a best response to the actions that others have used in the past. The classic papers in this branch are Kandori, Mailath and Rob [18] and Young [30]. Other contributions include Bergin and Lipman [2], Binmore and Samuelson [4], and Binmore, Samuelson, and Vaughan [5].

One limitation of both mechanical and boundedly rational models is that equilibrium selection occurs only in the long run and takes the form of a ergodic distribution with most of its weight on the selected equilibrium.¹⁹ Some have also criticized models with

¹⁸To see this, suppose that iterated dominance from the region where R is strictly dominant “goes as far as b ”. That is, if a player believes that others choose R when $B_t > b$ and L when $B_t < b$, then she is always willing to play R at b and is sometimes indifferent. When is she indifferent? In the worst case, when all old players are playing L. With small frictions, the player cares only about the very near future, when the Brownian motion is equally likely to be above or below b . Therefore, given her belief, she expects that half of the new players will play R and half L. Since she is indifferent and one quarter of her opponents play R, b must be $B^{1/4}$.

¹⁹Models of bounded rationality with local interactions in place of random matching yield faster convergence; see, for example, Ellison [11] and Blume [6].

boundedly rational players on the grounds that it is hard to justify the choice of a particular rule of thumb. This is especially true when a player could do better by recognizing that others are using the given rule.

In models with rational players, equilibrium selection is typically fast and irreversible. Moreover, a clever player cannot “outsmart” her opponents if they are fully rational as well. Most importantly, the “rational” approach makes it possible to analyze the individual player’s decision problem and the process of forming beliefs about what others will do.

Matsui and Matsuyama [24] study a “rational” model like ours but with a fixed world. Both steady states in which all players play the same action are rational expectations equilibria of the dynamic game. However, Matsui and Matsuyama find support for the prediction that the risk-dominant equilibrium of the static game is more likely to be played: from any initial state of play there is a rational expectations equilibrium leading to it; and from initial states close enough to the risk-dominant equilibrium, play must converge to it.

Lagunoff and Matsui [21] consider models with large (rational) players and lock-in. In pure coordination games with patient players, every subgame perfect equilibrium must lead to the *Pareto dominant* equilibrium of the static game. This happens because players are large enough to influence aggregate play, so they can “steer” others into the Pareto dominant equilibrium by playing it themselves. Although they may lose in the short run, they do not care if they are sufficiently patient. Matsui and Rob [23] and Lagunoff and Matsui [22] also assume large, rational players but do not use equilibrium reasoning. Instead, they identify conditions (including restrictions on beliefs) under which play converges to the Pareto dominant equilibrium.

Our paper also belongs to the “rational” branch, but differs in two ways from the above models. First, we use backwards induction rather than equilibrium notions or *ex ante* restrictions on beliefs. Second, we assume that payoffs change stochastically over time. Fudenberg and Harris [14] also assume that payoffs follow a Brownian motion,

but in a mechanical model. They show that this perturbation alone does not give rise to equilibrium selection under the replicator dynamics. Changes in payoffs appear also in the model of Ben-Porath, Dekel and Rustichini [1], who study how different mutation rates perform in a changing world.

6 Concluding Remarks

We view a normal form game as an abstraction that captures aspects of strategic interaction that are common to a variety of contexts. Since the normal form game is often not enough to make a prediction, we would like to learn which omitted features affect the outcome. An analysis of which assumptions are critical for our results can shed some light on conditions that make coordination on the risk-dominant equilibrium more likely.

Small Players

We assumed that players are small enough that they do not think that their actions will be observed and affect the evolution of play. If this assumption is relaxed, punishments can be devised that can sustain other types of behavior or, as in Lagunoff and Matsui [21], a player may incur short term losses in order to steer play towards the Pareto dominant equilibrium.

The Contagion Argument

A contagion argument works by connecting nearby decision problems, which enables players to predict what others will do. A few assumptions seem to be critical. For the argument to have a place to start, there must be regions of the parameter space where each action is strictly dominant. In addition, for the argument to work throughout the parameter space, all of a player's opponents must face decision problems that are *different* from (but close to) her own. Otherwise there may be a region of the parameter space with multiple equilibria, since players can have different self-sustaining beliefs about the behavior of those who face the same decision problem. In our model, we ensure that decision problems differ by breaking the simultaneity of moves: the measure of players

who change their actions simultaneously (and thus who see the same the same value of B_t) is always zero. Carlsson and van Damme [10] achieve a similar effect by assuming a continuous distribution of payoff signals, so that the two players never observe exactly the same signal.

Our contagion argument also uses the assumption that the relative payoff to R against either action is strictly increasing in B_t . One might be interested in cases where this does not hold. For example, in some range, a higher B_t may raise the payoff to *coordinating* on either action; i.e., it may raise the relative payoff to R if others choose R, but lower it if others choose L. More generally, the relative payoff to R may depend in an arbitrary way on B_t . This may prevent the contagion from spreading throughout the parameter space. Nevertheless, the following weaker version of Theorem 1 can still 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 *Lipschitz* in B_t . (A function is Lipschitz if its rate of change is bounded.²⁰) Let b be any value of B_t at which R is strictly dominant. Let (\underline{b}, \bar{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 frictions vanish, R must be played at all points in this interval. An analogous result can be proved for L.

The Stochastic Process

Brownian motion is a natural process to study because it is the only continuous process that has independent, stationary increments [29, pp. 146, 157].²¹ However, one may still wonder how our assumption of Brownian motion limits the generality of our results. Here we consider two ways in which this assumption can be weakened.

The first is continuity. There are many discontinuous processes with independent, stationary increments. A discrete process with this property is a random walk; it makes a sequence of i.i.d. jumps at fixed intervals. Alternatively, the time between jumps may

²⁰Formally, g is Lipschitz if there exists a $c > 0$ such that $|g(t) - g(s)| \leq c|t - s|$ for all t and s .

²¹Stationarity means that the distribution of increments over a given time interval can depend only on its length.

itself be a random variable; e.g., jumps may occur at Poisson intervals. Brownian motion is the limit of all such processes as the (expected) time between jumps goes to zero.²² Proposition 1 shows that our results hold for all such processes in a neighborhood of Brownian motion in the case of vanishing frictions (Theorem 1). The time between jumps must be small relative to the frictions so that only a small fraction of players face the same decision problem. An analogous result can be proved for the case of a slowly changing world (Theorem 2).

Consider a right-continuous process $(A_t)_{t \geq 0}$. We will say that $(A_t)_{t \geq 0}$ has *i.i.d. jumps* if there is a (random) sequence of times $\{t_0 = 0, t_1, t_2, \dots\}$ such that: (1) the process A_t is constant on every interval $[t_{i-1}, t_i)$; and (2) the random vectors $(t_i - t_{i-1}, A_{t_i} - A_{t_{i-1}})$ are independent and identically distributed. (In particular, each jump has a distribution that can depend only on the time since the prior jump.)

Proposition 1

For each i , let $(A_t^i)_{t \geq 0}$ be a process with i.i.d. jumps. Suppose that, as i goes to infinity, (A_t^i) converges in distribution to a Brownian motion (B_t) with trend μ and variance σ^2 . Let Γ_i be the game described in section 2, except that the payoff parameter changes according to (A_t^i) instead of (B_t) .

For any $\epsilon > 0$, there is a \underline{k} and a function $\psi(\cdot)$ such that if $k > \underline{k}$ and $i > \psi(k)$, R must be played in Γ_i whenever $A_t^i > B^{1/2} + \epsilon$ and L must be played whenever $A_t^i < B^{1/2} - \epsilon$.

Proof *See appendix.*

The assumption of independent, stationary increments can also be weakened. First, our results hold if B is any strictly increasing Lipschitz function of a Brownian motion. This is because any such transformation of B is equivalent to a change in the

²²More precisely, both the time between jumps and the jumps themselves must shrink to zero in an appropriate way. If the jumps shrink too slowly relative to the time between them, the increments of the resulting process over a fixed time period will not have a finite mean and variance. If they shrink too quickly, the process will become a straight line. Otherwise, the end result must be a Brownian motion.

utility function: instead of $u(a, a', \tilde{B}_t)$ where $\tilde{B}_t = g(B_t)$, we use the utility function $\tilde{u}(a, a', B_t) = u(a, a', g(B_t))$. With \tilde{u} , the relative payoff to playing R is still a Lipschitz function of B_t . This extension is important because the primitives that affect payoffs in the real world (such as prices or temperatures) may be bounded.

The results also hold if the trend parameter, rather than being the constant μ , is a bounded Lipschitz function of t . For instance, B_t may be the price of oil, which has a seasonal component. The trend can also be a Lipschitz function of B_t itself. One important example is the mean-reverting process $d\tilde{B}_t = dB_t - \mu(B_t - b)dt$, where B_t is a Brownian motion with no trend. Our results may be of greater interest with such processes. This is because Brownian motions tend to wander away from $B^{1/2}$ to regions where one action is strictly dominant; a mean reverting process with b close to $B^{1/2}$ spends a positive fraction of its time in the area where the static game has multiple equilibria.²³

Homogeneity of Players

How robust are our results to the assumption that all players have the same payoff function? Suppose that the payoff of player $i \in [0, 1]$ is $u(a, a', B_t + \theta_i)$, where θ_i is a personal taste parameter. Assume that most of the players have taste parameters 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 $1 - \epsilon$ of the players play R when $B_t > B^{1/2-\epsilon/2} - \theta$ and at least $1 - \epsilon'$ play L when $B_t < B^{1/2+\epsilon'/2} - \theta'$. This shows that there is continuity: when most of the players have very similar tastes, our results hold approximately. It remains an open question whether one can determine what happens for intermediate values of B_t . This is especially interesting when there is a large degree of heterogeneity.

²³In the long run \tilde{B}_t has a stationary distribution that is normal with mean b and variance $\sigma^2/2\mu$ (see Proposition 5.1 in Karlin and Taylor [20, p. 219]). A Brownian motion does not have a stationary distribution, as its variance goes to infinity.

Discounting

We have assumed that the discount rate r is positive. However, there are relevant situations in which a player has some reason to put more weight on the future. For example, a technology (such as the Internet) may be used by a growing number of people. A user thus might expect to interact with other users more often as time passes.

With a negative discount rate, a player's payoff while committed to a given action may still be well defined. This is because, when the player takes into account her chance of obtaining a new revision opportunity, her effective discount rate becomes $k + r$. Therefore, a player's payoff *during her commitment period* is finite as long as $r > -k$. In the case of vanishing frictions ($k \rightarrow \infty$), this is true for any r , and the result of Theorem 1 can be proved in this case. With a slowly changing world, payoffs are well defined only if $r > -k$. Even in this range, iterated dominance implies a weaker result than that of Theorem 2. If $r \in (-k, 0)$, players must play R when $B_t > B^{1/2-\lambda}$ and L when $B_t < B^{1/2+\lambda}$, where $\lambda = |r/(4k + 2r)|$. What happens for $B_t \in (B^{1/2+\lambda}, B^{1/2-\lambda})$ remains an open question.

A Proofs

Mathematical Preliminaries

We first collect several mathematical results that are proved in Burdzy, Frankel, and Pauzner [9] (henceforth, BFP). We begin with two lemmas. The first states that the differential equation that governs the evolution of the state of play has a unique solution that depends on certain parameters of the model in a continuous way.

Lemma 1 (BFP) *Let B_t be a Brownian motion with drift μ and variance σ^2 . Let $k > 0$. Fix $x_0 \in (0, 1)$ and $b_0 \in \mathfrak{R}$ and assume that $(B_0, X_0) = (b_0, x_0)$. Assume that f is a decreasing Lipschitz function. Consider the following differential equation:*

$$dX_t/dt = \begin{cases} k(1 - X_t) & \text{if } B_t > f(X_t), \\ -kX_t & \text{if } B_t < f(X_t). \end{cases} \quad (5)$$

1. *This equation has a unique Lipschitz solution $(X_t)_{t \geq 0}$ for almost every path $(B_t)_{t \geq 0}$.*
2. *Over any closed time interval $[0, T]$, the solution X_t is a uniformly continuous function of b_0 and x_0 . If $f(X_t)$ is replaced by $f(X_t) + \lambda$, the solution X_t is also uniformly continuous in λ .*

Since equation (5) does not specify what happens when $B_t = f(X_t)$, uniqueness applies only to Lipschitz solutions. For example, X_t may be identically equal to $f^{-1}(B_t)$ in some time interval; this gives a solution to (5) that is not Lipschitz. However, only Lipschitz solutions are consistent with the model, since no more than kdt of the agents can change strategies in a period of length dt .

A trivial corollary to Lemma 1 shows that the relative payoff to playing R in the dynamic game is also a continuous function of initial conditions and of f . Let $\phi(b, x; f)$ be the relative payoff to playing R that a player expects if she believes that (5) will hold and chooses her action at $(B_0, X_0) = (b, x)$, which is

$$E \left[\int_{t=0}^{\infty} e^{-(r+k)t} (X_t \Delta(R, B_t) + (1 - X_t) \Delta(L, B_t)) dt \mid (B_0, X_0) = (b, x) \right]$$

Corollary 1 *Assume that f is Lipschitz. Then the function $\phi(b, x; f)$ is continuous in b and x . Moreover, $\phi(b, x; f + \gamma)$ is continuous in $\gamma \in \mathbb{R}$.*

Proof

Without loss of generality, we assume that Brownian motions B_t and \tilde{B}_t starting from any two different points $B_0 = b$ and $\tilde{B}_0 = \tilde{b}$ are related by $B_t - b = \tilde{B}_t - \tilde{b}$ and suppose that (\tilde{b}, \tilde{x}) converges to (b, x) . Note that

$$\begin{aligned} \phi(b, x; f) - \phi(\tilde{b}, \tilde{x}; f) &= E \left[\int_{t=0}^{\infty} e^{-(r+k)t} \left(X_t \Delta(R, B_t) - \tilde{X}_t \Delta(R, \tilde{B}_t) \right) dt \right. \\ &\quad \left. + \int_{t=0}^{\infty} e^{-(r+k)t} \left((1 - X_t) \Delta(L, B_t) - (1 - \tilde{X}_t) \Delta(L, \tilde{B}_t) \right) dt \right. \\ &\quad \left. \mid (B_0, X_0) = (b, x); (\tilde{B}_0, \tilde{X}_0) = (\tilde{b}, \tilde{x}) \right]. \end{aligned}$$

By Lemma 1, \tilde{X}_t converges uniformly to X_t on any closed time interval $[0, T]$ as $(\tilde{b}, \tilde{x}) \rightarrow (b, x)$. This implies that $\phi(b, x; f) - \phi(\tilde{b}, \tilde{x}; f)$ goes to zero. Thus, $\phi(b, x; f)$ is continuous in its first two variables. Continuity in f follows in a similar manner. Q.E.D.

Lemma 2 shows that bifurcation times go to zero and that the relative chance of bifurcating up vs. down equals the relative speed of X on the two sides of F .

Lemma 2 (BFP) *For each $n > 0$, let B_t^n be a Brownian motion with drift μ_n and variance σ_n^2 , where $\lim_{n \rightarrow \infty} \mu_n = \lim_{n \rightarrow \infty} \sigma_n^2 = 0$. Let f_n be a continuously differentiable decreasing function. Suppose that $x_0^n \in [0, 1]$ converge to some fixed $x_0 \in (0, 1)$ as $n \rightarrow \infty$. Assume that $\lim_{n \rightarrow \infty} f'_n(x_0)$ exists and is nonzero, and the derivatives are asymptotically uniformly continuous at x_0 , i.e., for every $\epsilon > 0$ there exists an $n_0 < \infty$ and a $\delta > 0$ such that $|f'_n(x) - f'_n(x_0)| < \epsilon$ for all $x \in [x_0 - \delta, x_0 + \delta]$ and all $n > n_0$. Let X_t^n be the solution to the following differential equation, with $X_0^n = x_0^n$ and $B_0^n = f_n(x_0^n)$,*

$$dX_t^n/dt = \begin{cases} \lambda_n(1 - X_t^n) & \text{if } B_t^n > f_n(X_t^n), \\ -\hat{\lambda}_n X_t^n & \text{if } B_t^n < f_n(X_t^n), \end{cases}$$

Assume that $\lim_{n \rightarrow \infty} \lambda_n = \lambda \in (0, \infty)$ and $\lim_{n \rightarrow \infty} \hat{\lambda}_n = \hat{\lambda} \in (0, \infty)$. Let $c_0 \in (0, \min(x_0, 1 - x_0))$, and let $T_1^n = \inf\{t > 0 : X_t^n \notin (c_0, 1 - c_0)\}$ and $T_0^n = \sup\{t < T_1^n : B_t^n = f_n(X_t^n)\}$.

1. The random bifurcation times T_0^n converge to 0 in distribution as $n \rightarrow \infty$.
2. The probability that the bifurcation is positive (i.e., that $dX^n/dt > 0$ for all $t \in (T_0^n, T_1^n)$) converges to $\frac{\lambda(1-x_0)}{\lambda(1-x_0)+\lambda x_0}$.

For brevity, Lemma 2 is stated only for the case $x_0 \in (0, 1)$. It also holds when $x_0 \in \{0, 1\}$, with a slight change in the definition of T_1^k . Let $c_0 \in (0, 1)$. If $x_0 = 0$, let $T_1^k = \inf\{t > 0 : X_t^k \geq 1 - c_0\}$. If $x_0 = 1$, let $T_1^k = \inf\{t > 0 : X_t^k \leq c_0\}$.

Lemma 3 concerns the case of processes with i.i.d. jumps (see definition on p. 24). It is used to show that, as the processes converge to a Brownian motion, so does the relative payoff to playing R. Fix a discount rate $c > 0$ and let $\nu(b, x)$ be any Lipschitz function. Let A be either a Brownian motion or a discrete process with independent, stationary increments. Suppose that $A_0 = f(X_0)$ and that X satisfies

$$\dot{X}_t = \begin{cases} k(1 - X_t) & \text{if } A_t > f(X_t) \\ -kX_t & \text{if } A_t < f(X_t) \end{cases} \quad (6)$$

Because equation (6) does not specify what happens if $A_t = f(X_t)$, it does not pin down what happens to X in the period after 0 before A jumps: X can either remain unchanged, move up, or move down. Let \overline{X}_t and \underline{X}_t be the maximal and the minimal Lipschitz solutions to (6). Let

$$\overline{\Phi}(x, b; A, f) = E \left[\int_{t=0}^{\infty} e^{-ct} \nu(A_t, \overline{X}_t) dt \mid (A_0, \overline{X}_0) = (b, x) \right]$$

and

$$\underline{\Phi}(x, b; A, f) = E \left[\int_{t=0}^{\infty} e^{-ct} \nu(A_t, \underline{X}_t) dt \mid (A_0, \underline{X}_0) = (b, x) \right]$$

Note that if $c = r + k$ and $\nu(b, x) = x\Delta(R, b) + (1 - x)\Delta(L, b)$, $\overline{\Phi}$ and $\underline{\Phi}$ are the highest and lowest relative payoffs to playing R that a player can expect.

Lemma 3 (BFP)

Let A^i be a sequence of processes with i.i.d. jumps that converges in distribution to the Brownian motion B as $i \rightarrow \infty$. Let f^i be a sequence of strictly decreasing Lipschitz functions that converges to f as $i \rightarrow \infty$. Suppose that (x^i, b^i) converges to (x, b) . Then $\overline{\Phi}(x^i, b^i; A^i, f^i)$ and $\underline{\Phi}(x^i, b^i; A^i, f^i)$ both converge to $\overline{\Phi}(x, b; B, f) = \underline{\Phi}(x, b; B, f)$.

Proof of Theorem 1

Since we know little about the shapes of the f^n 's, it is technically difficult to work with them directly. Instead, we work with sequences of functions whose form we do know. Fix some $\rho > 0$ and let $q_0(x) = -\rho x + \lambda_0$, where λ_0 is the smallest constant such that R is strictly dominant at every (b, x) for which $b > q_0(x)$. Then inductively define $q_n(x)$ for $n \geq 1$ by letting $q_n(x) = -\rho x + \lambda_n$, where λ_n is the smallest number such that $\phi(b, x; q_{n-1}) \geq 0$ for all (b, x) with $b = q_n(x)$.

Let $Q_k(x)$ be the infimum of $q_n(x)$ as n goes to infinity. (We write Q_k to denote the dependence on k .) By iterated elimination, R must be played to the right of Q_k . We first show that Q_k must have the *indifference property*: that there is an $x \in [0, 1]$ such that $\phi(Q_k(x), x; Q_k) = 0$. Suppose otherwise. Then we have $\phi(Q_k(x), x; Q_k) > 0$ for all x . By the continuity of ϕ and compactness of $[0, 1]$, there is a band of strictly positive width to the left of Q_k such that $\phi(\tilde{b}, \tilde{x}; Q_k) > 0$ for all (\tilde{b}, \tilde{x}) within this band. But $\phi(b, x; Q_k + \gamma)$ is continuous in $\gamma \in \Re$ by Corollary 1. Thus, for q_n sufficiently close to Q_k , $\phi(\tilde{b}, \tilde{x}; q_n) > 0$ for all (\tilde{b}, \tilde{x}) in some band to the left of q_n that includes Q_k . This implies that $q_{n+1} < Q_k$, a contradiction. So Q_k must have the indifference property. For any k , let x_k be one such point of indifference. That is,

$$(r+k)E \left[\int_{t=0}^{\infty} e^{-(r+k)t} (X_t \Delta(R, B_t) + (1 - X_t) \Delta(L, B_t)) dt \mid (B_0, X_0) = (Q_k(x_k), x_k) \right] = 0 \quad (7)$$

where $(X_t)_{t \geq 0}$ satisfies:

$$\dot{X}_t = \begin{cases} k(1 - X_t) & \text{if } B_t > Q_k(X_t) \\ -kX_t & \text{if } B_t < Q_k(X_t) \end{cases} \quad (8)$$

By passing to a subsequence if necessary, we may assume that the indifference points x_k converge to a point $x_\infty \in [0, 1]$ as $k \rightarrow \infty$.

We now make a simplification: in taking the limit of (7) as $k \rightarrow \infty$, B_t can be replaced by B_0 . The intuition is that as $k \rightarrow \infty$, players care increasingly about the payoffs they get in smaller and smaller neighborhoods of $t = 0$, where B_t is closer and closer to B_0 (since B_t is continuous with probability one). By (1), for any $a \in \{R, L\}$,

$$(r+k)E \left[\int_{t=0}^{\infty} e^{-(r+k)t} |\Delta(a, B_t) - \Delta(a, B_0)| dt \right]$$

is no greater than

$$(r+k) \int_{t=0}^{\infty} e^{-(r+k)t} \overline{w} E[|B_t - B_0|] dt \leq \overline{w}(r+k) \int_{t=0}^{\infty} e^{-(r+k)t} \sqrt{E[|B_t - B_0|^2]} dt$$

Since $B_t - B_0$ is normal with mean μt and variance $\sigma^2 t$, this is no greater than

$$\begin{aligned} \overline{w}(r+k) \int_{t=0}^{\infty} e^{-(r+k)t} \sqrt{\sigma^2 t + (\mu t)^2} dt &\leq \overline{w}(r+k) \int_{t=0}^{\infty} e^{-(r+k)t} [\sigma \sqrt{t} + |\mu| t] dt \\ &= \overline{w}(r+k) \left[\sigma \frac{\sqrt{\pi}}{2(r+k)^{3/2}} + |\mu| \frac{1}{(r+k)^2} \right] \end{aligned}$$

which goes to zero as $k \rightarrow \infty$. Hence, we can substitute $B_0 = Q_k(x_k)$ for B_t when we take the limit of (7) as $k \rightarrow \infty$.

We now want to show that the time of bifurcation goes to zero as k goes to ∞ , even relative to the agent's time horizon (which also goes to zero). We rescale time so that the agent's horizon is independent of k . We then need to show that the time of bifurcation in the new units goes to zero. Let $v = (r+k)t$, $\tilde{X}_v = X_{v/(r+k)}$, and $\tilde{B}_v = B_{v/(r+k)}$. Then (7), with $B_0 = Q_k(x_k)$ substituted for B_v , becomes

$$\lim_{k \rightarrow \infty} E \left[\int_{v=0}^{\infty} e^{-v} \left(\tilde{X}_v \Delta(R, Q_k(x_k)) + (1 - \tilde{X}_v) \Delta(L, Q_k(x_k)) \right) dv \mid \tilde{X}_0 = x_k \right] = 0$$

where \tilde{X}_v satisfies

$$\dot{\tilde{X}}_v = \begin{cases} \frac{k}{r+k}(1 - \tilde{X}_v) & \text{if } \tilde{B}_v > Q_k(\tilde{X}_v) \\ -\frac{k}{r+k}\tilde{X}_v & \text{if } \tilde{B}_v < Q_k(\tilde{X}_v) \end{cases}$$

and $(\tilde{B}_v)_{v \geq 0}$ is a Brownian motion with drift parameter $\mu/(r+k)$ and variance parameter $\sigma^2/(r+k)$ that begins at $\tilde{B}_0 = Q_k(x_k)$.

We now define the time of bifurcation. Fix an arbitrarily small $c_0 > 0$. If $x_\infty \in (0, 1)$, we will assume moreover that $c_0 < \min(x_\infty, 1 - x_\infty)$. Let T_1^k be the first time v at which $\tilde{B}_v \notin (Q_k(1 - c_0), Q_k(c_0))$. Let T_0^k be the largest $v \leq T_1^k$ at which $\tilde{B}_v = Q_k(\tilde{X}_v)$. The time of bifurcation is defined as T_0^k ; it is the last time at which $\tilde{B}_v = Q_k(\tilde{X}_v)$ before \tilde{X}_v reaches an c_0 -neighborhood of either 0 or 1. If $x_\infty = 0$, we let T_1^k be the first time v at which $\tilde{B}_v \leq Q_k(1 - c_0)$. If $x_\infty = 1$, we let T_1^k be the first time v at which $\tilde{B}_v \geq Q_k(c_0)$. The definition of T_0^k is unchanged in these cases.

By Lemma 2, as $k \rightarrow \infty$,

- (i) The bifurcation times T_0^k converge to zero in distribution; and
- (ii) The probability that $\dot{\tilde{X}}_v = \frac{k}{r+k}(1 - \tilde{X}_v)$ for all $v \in (T_0^k, T_1^k)$, converges to $1 - x_\infty$.

The probability that $\dot{\tilde{X}}_v = -\frac{k}{r+k}\tilde{X}_v$ for all $v \in (T_0^k, T_1^k)$, converges to x_∞ .

By taking c_0 arbitrarily close to zero, we obtain

$$\begin{aligned} \lim_{k \rightarrow \infty} E(\tilde{X}_v \mid \hat{X}_0 = x_k) &= (1 - x_\infty)(1 - (1 - x_\infty)e^{-v}) + x_\infty(x_\infty e^{-v}) \\ &= 1 - x_\infty + (2x_\infty - 1)e^{-v} \end{aligned}$$

so that

$$\begin{aligned} \lim_{k \rightarrow \infty} E \left[\int_{v=0}^{\infty} e^{-v} \left(\tilde{X}_v \Delta(R, Q_k(x_k)) + (1 - \tilde{X}_v) \Delta(L, Q_k(x_k)) \right) dv \mid \tilde{X}_0 = x_k \right] \\ &= \lim_{k \rightarrow \infty} \int_{v=0}^{\infty} e^{-v} \left(E(\tilde{X}_v \mid \tilde{X}_0 = x_k) [\Delta(R, Q_k(x_k)) - \Delta(L, Q_k(x_k))] + \Delta(L, Q_k(x_k)) \right) dv \\ &= \lim_{k \rightarrow \infty} \int_{v=0}^{\infty} e^{-v} \left([1 - x_\infty + (2x_\infty - 1)e^{-v}] [\Delta(R, Q_k(x_k)) - \Delta(L, Q_k(x_k))] + \Delta(L, Q_k(x_k)) \right) dv \\ &= \lim_{k \rightarrow \infty} \frac{\Delta(R, Q_k(x_k)) + \Delta(L, Q_k(x_k))}{2} \end{aligned}$$

Since the first line is equal to zero by the indifference property, it must be that

$$\Delta(R, \lim_{k \rightarrow \infty} Q_k(x_k)) + \Delta(L, \lim_{k \rightarrow \infty} Q_k(x_k)) = 0$$

By (2), $\lim_{k \rightarrow \infty} Q_k(x_k) = B^{1/2}$. This concludes the proof for R, since the slope $-\rho$ of the function Q_k can be arbitrarily close to zero. The proof for L is analogous.

Proof of Theorem 2

The proof follows the lines of the proof of Theorem 1, so we simply sketch the main steps. We need to show that the indifference line is given by $B^{h(x)}$. We construct a sequence of functions q_n as follows. Let $q_0(x) = B^{h(x)} + \lambda_0$, where λ_0 is large enough that $\phi(b, x; q_0) > 0$ for all (b, x) with $b > q_0(x)$. Then inductively define $q_n(x) = B^{h(x)} + \lambda_n$, where λ_n is the smallest number such that $\phi(b, x; q_{n-1}) \geq 0$ for all (b, x) with $b = q_n(x)$.

Notice that each q_n is Lipschitz. This is because $B^{h(x)}$ is given implicitly by

$$x\Delta(R, B^{h(x)}) + (1-x)\Delta(L, B^{h(x)}) = 0$$

By differentiating this with respect to x , we obtain

$$\frac{dB^{h(x)}}{dx} = -\frac{x\Delta_2(R, B^{h(x)}) + (1-x)\Delta_2(L, B^{h(x)})}{\Delta(R, B^{h(x)}) - \Delta(L, B^{h(x)})}$$

where $\Delta_2(\cdot, \cdot)$ denotes the derivative of $\Delta(\cdot, \cdot)$ with respect to the second argument. Since $\Delta(R, B_t) > \Delta(L, B_t)$ for all B_t , the denominator is bounded away from zero over the compact interval $[B^{h(1)}, B^{h(0)}]$. By (1), the numerator is bounded above, so $B^{h(x)}$ is Lipschitz, as are its translations q_n .

Let $Q_{\sigma, \mu}$ be the infimum of the q_n 's. (The subscripts of $Q_{\sigma, \mu}$ indicates its dependence on the parameters of the Brownian motion.) By iterated dominance, R must be played to the right of $Q_{\sigma, \mu}$. As in Theorem 1, we can show the “indifference property” for $Q_{\sigma, \mu}$, i.e., that there exists an $x_{\sigma, \mu}$ with $\phi(Q_{\sigma, \mu}(x_{\sigma, \mu}), x_{\sigma, \mu}; Q_{\sigma, \mu}) = 0$. Again, we can assume that $x_{\sigma, \mu}$ converges to x_∞ as σ and μ go to zero.

As before, we can substitute $Q_{\sigma, \mu}(x_{\sigma, \mu})$ for B_t in computing the limit of $\phi(Q_{\sigma, \mu}(x_{\sigma, \mu}), x_{\sigma, \mu}; Q_{\sigma, \mu})$. The indifference property then implies

$$\begin{aligned} \lim_{\mu, \sigma \rightarrow 0} (r+k)E \left[\int_{t=0}^{\infty} e^{-(r+k)t} (X_t \Delta(R, Q_{\sigma, \mu}(x_{\sigma, \mu})) + (1-X_t) \Delta(L, Q_{\sigma, \mu}(x_{\sigma, \mu}))) dt \right. \\ \left. \mid (B_0, X_0) = (Q_{\sigma, \mu}(x_{\sigma, \mu}), x_{\sigma, \mu}) \right] = 0 \end{aligned}$$

where X_t satisfies an equation analogous to (8). Using the same change of variables $v = (r+k)t$, we obtain

$$\lim_{\mu, \sigma \rightarrow 0} E \left[\int_{v=0}^{\infty} e^{-v} (\tilde{X}_v \Delta(R, Q_{\sigma, \mu}(x_{\sigma, \mu})) + (1-\tilde{X}_v) \Delta(L, Q_{\sigma, \mu}(x_{\sigma, \mu}))) dt \mid \tilde{X}_0 = x_{\sigma, \mu} \right] = 0$$

where \tilde{X}_v satisfies

$$\dot{\tilde{X}}_v = \begin{cases} \frac{k}{r+k}(1 - \tilde{X}_v) & \text{if } \tilde{B}_v > Q_{\sigma,\mu}(\tilde{X}_v) \\ -\frac{k}{r+k}\tilde{X}_v & \text{if } \tilde{B}_v < Q_{\sigma,\mu}(\tilde{X}_v) \end{cases}$$

and $(\tilde{B}_v)_{v \geq 0}$ is a Brownian motion with drift parameter $\mu/(r+k)$ and variance parameter $\sigma^2/(r+k)$ that begins at $\tilde{B}_0 = Q_{\sigma,\mu}(x_{\sigma,\mu})$.

Fix an arbitrarily small $c_0 > 0$, and let $T_0^{\sigma,\mu}$ be the corresponding bifurcation time. By Lemma 2, as $\sigma, \mu \rightarrow 0$, the time of bifurcation converges to zero and the chances of bifurcating to 1 and 0 converge to $1 - x_\infty$ and x_∞ , respectively. Therefore,

$$\begin{aligned} \lim_{\mu, \sigma \rightarrow 0} E(\tilde{X}_v \mid \tilde{X}_0 = x_{\sigma,\mu}) &= (1 - x_\infty)(1 - (1 - x_\infty)e^{-\frac{k}{r+k}v}) + x_\infty(x_\infty e^{-\frac{k}{r+k}v}) \\ &= 1 - x_\infty + (2x_\infty - 1)e^{-\frac{k}{r+k}v} \end{aligned}$$

so that

$$\begin{aligned} \lim_{\mu, \sigma \rightarrow 0} E \left[\int_{v=0}^{\infty} e^{-v} \left(\tilde{X}_v \Delta(R, Q_{\sigma,\mu}(x_{\sigma,\mu})) + (1 - \tilde{X}_v) \Delta(L, Q_{\sigma,\mu}(x_{\sigma,\mu})) \right) dv \mid \tilde{X}_0 = x_{\sigma,\mu} \right] \\ = \lim_{\mu, \sigma \rightarrow 0} \left[\frac{rx_\infty + k}{r + 2k} \Delta(R, Q_{\sigma,\mu}(x_{\sigma,\mu})) + \frac{-rx_\infty + r + k}{r + 2k} \Delta(L, Q_{\sigma,\mu}(x_{\sigma,\mu})) \right] \end{aligned}$$

Since the first line is equal to zero by the indifference property, and $\Delta(\cdot, \cdot)$ is strictly increasing in its second argument by (1), it must be that

$$\frac{rx_\infty + k}{r + 2k} \Delta(R, \lim_{\mu, \sigma \rightarrow 0} Q_k(x_k)) + \frac{-rx_\infty + r + k}{r + 2k} \Delta(L, \lim_{\mu, \sigma \rightarrow 0} Q_k(x_k)) = 0$$

By (2), $\lim_{\mu, \sigma \rightarrow 0} Q_{\sigma,\mu}(x_{\sigma,\mu}) = B^{h(x_\infty)}$ where $h(x) = \frac{rx+k}{r+2k}$. The proof for L is analogous.

Proof of Proposition 1

To find out where R must be played, we do the iterative elimination using the lowest possible relative payoff to playing R. This equals the function $\underline{\Phi}$ (defined on p. 30), where we let $c = r + k$ and $\nu(b, x) = x\Delta(R, b) + (1 - x)\Delta(L, b)$. Let $q_0^i(x) = -\rho x + \lambda_0$, where R is strictly dominant at every (b, x) for which $b > q_0^i(x)$ and b is a possible value of A^i . Let $q_n^i(x) = -\rho x + \lambda_n$, where λ_n is the smallest number such that $\underline{\Phi}(b, x; A^i, q_{n-1}^i) \geq 0$ for all (b, x) such that $b = q_n^i(x)$ is a possible value of A^i . Let Q_k^i be the infimum of the

q_n^i 's over n . Q_k^i must have the *indifference property*: there is an $x_k^i \in [0, 1]$ such that $\underline{\Phi}(Q_k^i(x_k^i), x_k^i; A^i, Q_k^i) = 0$. Moreover, R must be played to the right of Q_k^i if the payoff perturbation follows A^i .

By passing to a subsequence if necessary, we may assume that the limits $x_k = \lim_{i \rightarrow \infty} x_k^i$ and $Q_k = \lim_{i \rightarrow \infty} Q_k^i$ exist. By Lemma 3, $\underline{\Phi}(Q_k^i(x_k^i), x_k^i; A^i, Q_k^i)$ converges to $\underline{\Phi}(Q_k(x_k), x_k; B, Q_k)$, which therefore must also be zero. This means (following the argument of Theorem 1) that $Q_k(x_k)$ converges to $B^{1/2}$ as $k \rightarrow \infty$. Hence there is a \underline{k} such that, if $k > \underline{k}$, $Q_k(x_k) \leq B^{1/2} + \epsilon/3$. Since ρ was chosen arbitrarily, we can assume that it is less than $\epsilon/3$; this guarantees that $Q_k(x) \leq B^{1/2} + 2\epsilon/3$ for all x . Finally, choose $\psi(k)$ such that if $i > \psi(k)$, Q_k^i is no further than $\epsilon/3$ to the right of Q_k (i.e., $Q_k^i \leq Q_k + \epsilon/3$). It follows that $Q_k^i(x) \leq B^{1/2} + \epsilon$ for all x . Accordingly, if $k > \underline{k}$ and $i > \psi(k)$, R must be played under the discrete process A^i whenever $A_t^i > B^{1/2} + \epsilon$. An analogous proof holds for L using $\overline{\Phi}$.

References

- [1] Ben-Porath, Elchanan, Eddie Dekel, and Aldo Rustichini. 1993. "On the Relationship Between Mutation Rates and Growth Rates in a Changing Environment." *Games and Economic Behavior* 5:576-603.
- [2] Bergin, J., and B. Lipman. 1996. "Evolution with State Dependent Mutations." *Econometrica* 64:943-956.
- [3] Billingsley, Patrick. 1986. *Probability and Measure*. New York: John Wiley & Sons.
- [4] Binmore, K.G., and L. Samuelson. 1994. "Muddling Through: Noisy Equilibrium Selection." Mimeo.
- [5] Binmore, K.G., L. Samuelson, and R. Vaughan. 1995. "Musical Chairs: Modeling Noisy Evolution." *Games and Economic Behavior* 11:1-35.
- [6] Blume, Lawrence. 1993. "The Statistical Mechanics of Strategic Interaction." *Games and Economic Behavior* 5:387-424.
- [7] Boylan, R.T. 1992. "Laws of Large Numbers for Dynamical Systems with Randomly Matched Individuals." *Journal of Economics Theory* 57:473-504.
- [8] Burdzy, Krzysztof. 1987. *Multidimensional Brownian Excursions and Potential Theory*. Essex, England: Longman.
- [9] Burdzy, Krzysztof, David M. Frankel, and Ady Pauzner. 1997. "On the Time and Direction of Stochastic Bifurcation." Tel-Aviv University: Foerder Institute working paper no. 8-97.
- [10] Carlsson, Hans, and Eric van Damme. 1993. "Global Games and Equilibrium Selection." *Econometrica* 61:989-1018.
- [11] Ellison, Glenn. 1993. "Learning, Local Interaction, and Coordination." *Econometrica* 61:1047-1071.

- [12] Foster, D., and P.H. Young. 1990. "Stochastic Evolutionary Game Dynamics." *Theoretical Population Biology* 38:219-232.
- [13] Fudenberg, Drew, and Jean Tirole. 1991. *Game Theory*. Cambridge: MIT Press.
- [14] Fudenberg, Drew, and Chris Harris. 1992. "Evolutionary Dynamics with Aggregate Shocks." *Journal of Economic Theory* 57:420-441.
- [15] Gilboa, I., and A. Matsui. 1992. "A Model of Random Matching." *Journal of Mathematical Economics* 21:185-197.
- [16] Harsanyi, J., and R. Selten. 1988. *A General Theory of Equilibrium Selection in Games*. Cambridge: MIT Press.
- [17] Judd, Kenneth L. 1985. "The Law of Large Numbers with a Continuum of IID Random Variables." *Journal of Economic Theory* 35:19-25.
- [18] Kandori, Michihiro, George J. Mailath, and Rafael Rob. 1993. "Learning, Mutation, and Long Run Equilibria in Games." *Econometrica* 61:29-56.
- [19] Karlin, S., and H.M. Taylor. *A First Course in Stochastic Processes*. New York: Academic Press, 1975.
- [20] Karlin, S., and H.M. Taylor. *A Second Course in Stochastic Processes*. New York: Academic Press, 1981.
- [21] Lagunoff, Roger, and Akihiko Matsui. 1995. "An 'Anti-Folk Theorem' for a Class of Asynchronously Repeated Games." Mimeo, University of Pennsylvania.
- [22] Lagunoff, Roger, and Akihiko Matsui. 1995. "Evolution in Mechanisms for Public Projects." *Economic Theory* 6:195-223.
- [23] Matsui, Akihiko, and Rafael Rob. 1993. "Evolution with Forward-Looking Players." Mimeo, University of Pennsylvania.

- [24] Matsui, Akihiko, and Kiminori Matsuyama. 1995. "An Approach to Equilibrium Selection." *Journal of Economic Theory* 65:415-434.
- [25] Morris, Stephen, Rob, Rafael, and Hyun Song Shin. 1995. "p-Dominance and Belief Potential." *Econometrica* 63:145-157.
- [26] Nash, John F. 1950. "The Bargaining Problem." *Econometrica* 18:155-162.
- [27] Nash, John F. 1951. "Non-Cooperative Games." *Annals of Mathematics* 54:286-295.
- [28] Nash, John F. 1953. "Two-Person Cooperative Games." *Econometrica* 21:128-140.
- [29] Stroock, Daniel. 1993. *Probability Theory: An Analytic View*. Cambridge University Press.
- [30] Young, Peyton. 1993. "The Evolution of Conventions." *Econometrica* 61:57-84.