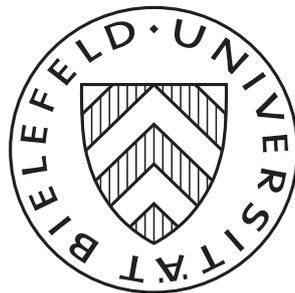


March 2014

Large Deviations and Stochastic Stability in the Small Noise Double Limit, I: Theory

Willian H. Sandholm and Mathias Staudigl



Large Deviations and Stochastic Stability in the Small Noise Double Limit, I: Theory*

William H. Sandholm[†] and Mathias Staudigl[‡]

March 21, 2014

Abstract

We consider a model of stochastic evolution under general noisy best response protocols, allowing the probabilities of suboptimal choices to depend on their payoff consequences. Our analysis focuses on behavior in the small noise double limit: we first take the noise level in agents' decisions to zero, and then take the population size to infinity. We show that in this double limit, escape from and transitions between equilibria can be described in terms of solutions to continuous optimal control problems. These are used in turn to characterize the asymptotics of the stationary distribution, and so to determine the stochastically stable states. The control problems are tractable in certain interesting cases, allowing analytical descriptions of the escape dynamics and long run behavior of the stochastic evolutionary process.

1. Introduction

Evolutionary game theory studies the behavior of strategically interacting agents whose decisions are based on simple myopic rules. Together, a game, a decision rule, and a population size define a stochastic aggregate behavior process on the set of population states. How one should analyze this process depends on the time span of interest. Over short to moderate time spans, the process typically settles on a small set of

*Financial support from NSF Grants SES-0851580 and SES-1155135, US Air Force OSR Grant FA9550-09-0538, and the Vienna Science and Technology Fund (WWTF) under project fund MA 09-017 is gratefully acknowledged.

[†]Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706, USA. e-mail: whs@ssc.wisc.edu; website: www.ssc.wisc.edu/~whs.

[‡]Center for Mathematical Economics, Bielefeld University, Germany. e-mail: mathias.staudigl@uni-bielefeld.de; website: mwpweb.eu/MathiasStaudigl.

population states, most often near a Nash equilibrium of the underlying game. If agents sometimes choose suboptimal strategies, then over longer time spans, transitions between equilibria are inevitable, with some occurring more readily than others. This variation in the difficulties of transitions ensures that a single equilibrium—the stochastically stable equilibrium—will be played in a large proportion of periods over long enough time spans. Thus noise in individuals' decisions can generate unique predictions of play for interactions of long duration.¹

While stochastic stability analysis is valued for its conclusions about equilibrium selection, the intermediate steps of this analysis are themselves of direct interest. The first step, which identifies equilibria and other recurrent classes of the aggregate behavior process, can be viewed as a part of a large literature on the convergence and nonconvergence of disequilibrium learning processes to Nash equilibrium.² The next step assesses the likelihoods of escapes from and transitions among equilibria and other recurrent classes. The questions of how an established equilibrium is upset, and which (if any) new equilibrium is likely to arise, seem themselves to be of general import. But to date, the question of equilibrium breakdown has not attracted much attention in the game theory literature.³

Most work on stochastic stability follows Kandori et al. (1993) by considering the best response with mutations (BRM) model, in which the probability of a suboptimal choice is independent of its payoff consequences.⁴ This model eases the determination of stochastically stable states, as the difficulty of transiting from one equilibrium to another can be determined by counting the number of mutations needed for the transition to occur.

Of course, this simplicity of analysis owes to a polar stance on the nature of suboptimal choices. In some applications, it may be more realistic to suppose that the probability of a suboptimal choice depends on its payoff consequences, as in the logit model of Blume (1993, 2003), and the probit model of Myatt and Wallace (2003). When mistake probabilities are payoff-dependent, the probability of a transition between equilibria becomes more difficult to assess, depending now not only on the number of suboptimal choices required, but also on the unlikeliness of each such choice. As a consequence, general results on transitions between equilibria and stochastic stability are only available for two-strategy games.⁵

¹Stochastic stability analysis was introduced to game theory by Foster and Young (1990), Kandori et al. (1993), and Young (1993), and since these early contributions has developed into a substantial literature. For surveys, see Young (1998) and Sandholm (2010b, Ch. 11–12).

²See, for instance, Young (2004) and Sandholm (2010b).

³The final step in stochastic stability analysis uses graph-theoretic methods to distill the analysis of transitions between equilibria into a characterization of the limiting stationary distribution of the process. See the references above, Freidlin and Wentzell (1998), Catoni (1999), or Section 6.2 below.

⁴Kandori and Rob (1995, 1998) and Ellison (2000) provide key contributions to this approach.

⁵Blume (2003) and Sandholm (2007, 2010a) study stochastic stability in two-strategy games using birth-

In this paper, we consider a model of stochastic evolution under general noisy best response protocols. To contend with the complications raised by the sensitivity of mistakes to payoffs, we study behavior in the small noise double limit, first taking the noise level in agents' decisions to zero, as in the works referenced above, and then taking the population size to infinity. We thereby evaluate the small noise limit when the population size is large.

We show that in this double limit, transitions between equilibria can be described in terms of solutions to continuous optimal control problems. By combining this analysis with standard graph-theoretic techniques, we characterize the asymptotics of the stationary distribution and the stochastically stable states. In a companion paper, Sandholm and Staudigl (2014), we show that the control problems derived here can be solved analytically in interesting classes of examples. To our knowledge, these papers are the first to provide tractable analyses of transition dynamics and stochastic stability when mistake probabilities depend on payoff consequences and agents choose among more than two strategies.

We consider stochastic evolution in a population of size N . The population recurrently plays an n -strategy population game F^N , which specifies the payoffs to each strategy as a function of the population state. In each period, a randomly chosen agent receives an opportunity to switch strategies. The agent's choice is governed by a *noisy best response protocol* σ^η with *noise level* η , which places most probability on strategies that are currently optimal, but places positive probability on every strategy.

We assume that for any given vector of payoffs, the probability with which a given strategy is chosen vanishes at a well-defined rate as the noise level approaches zero. This rate, called the *strategy's unlikelihood*, is positive if and only if the strategy is suboptimal, and is assumed to depend continuously on the vector of payoffs. For instance, under the logit choice model, a strategy's unlikelihood is the difference between its current payoff and the current optimal payoff.⁶

A population game F^N and a protocol σ^η generate a stochastic evolutionary process $\mathbf{X}^{N,\eta}$. In Section 3, we use standard techniques to evaluate the behavior of this process as the noise level η approaches zero. We start by introducing a *discrete best response dynamic*, which describes the possible paths of play when only optimal strategies are chosen. The *recurrent classes* of this dynamic are the minimal sets of states from which the dynamic cannot escape.

death chain methods. Staudigl (2012) studies the case of two-population random matching in 2x2 normal form games. Results are also available for certain specific combinations of games and choice protocols, most notably potential games under logit choice: see Blume (1993, 1997), Alós-Ferrer and Netzer (2010), and Sandholm (2010b, Sec. 11.5).

⁶See Section 2.2. The continuity assumption rules out the BRM model, in which unlikelihood functions are indicator functions. We discuss this model in Section 7.2.

To evaluate the probabilities of transitions between recurrent classes in the small noise limit, we define the *cost* of a path as the sum of the unlikelihoods associated with the changes in strategy along the path. Thus a path's cost is the exponential rate of decay of its probability as the noise level vanishes.

According to a well-known principle from the theory of large deviations, the probability of a transition between equilibria should be governed by the minimum cost path that effects the transition. These *transition costs*, if they can be determined, provide the inputs to a graph-theoretic analysis—the construction of certain trees on the set of recurrent classes—that characterizes the behavior of the stationary distribution in the small-noise limit, and so determines the stochastically stable states.

Solving these minimum cost path problems is computationally intensive if the number of agents is not small. In the case of the BRM model, this difficulty is mitigated by the fact that all mistakes are equally likely, so that the cost of a path is determined by its length. But when probabilities of mistakes depend on their consequences, this simplification is no longer available.

We overcome this difficulty by considering the small noise double limit: after taking the noise level η to zero, we take the population size N to infinity. We prove that as N grows large, the discrete path cost minimization problem converges to a continuous optimal control problem on the simplex. Although this problem is not smooth, it has a simple enough structure to be tractable even when the state space is multidimensional. To demonstrate this point, our companion paper provides a complete analysis of large deviations and stochastic stability for a range of three-strategy coordination games under the logit choice rule.

While work in stochastic evolutionary game theory typically focuses on stochastic stability and equilibrium selection, we feel that the dynamics of transitions between equilibria are themselves of inherent interest. Just as theories of disequilibrium learning offer explanations of how and when equilibrium play may arise, models of transition dynamics suggest how equilibrium is likely to break down. The importance of this question has been recognized in macroeconomics, where techniques from large deviations theory have been used to address this possibility in a variety of applications; see Cho et al. (2002), Williams (2009), and the references therein. The present paper addresses this question in an environment where the stochastic process arises endogenously as a description of the aggregate behavior of a population of strategically interacting agents.

A number of earlier papers on stochastic evolution have considered small noise double limits. Binmore et al. (1995) and Binmore and Samuelson (1997) (see also Sandholm (2012)) analyze models of imitation with mutations, focusing on two-strategy games; see Section

7.3 for a discussion. Fudenberg and Imhof (2006, 2008) extend these analyses to the many-strategy case. The key insight of the latter papers is that under imitation with mutations, the stochastic evolutionary process is nearly always at vertices or on edges of the simplex. Because of this, transitions between equilibria can be analyzed as one-dimensional problems using birth-death chain methods. In the noisy best response models studied here, the least costly transition between a pair of equilibria need not follow an edge of the simplex, a point illustrated explicitly in our companion paper.

Turning to noisy best response models, Kandori and Rob (1995, 1998) and Ellison (2000) analyze stochastic evolution under the BRM rule in the small noise double limit, as we discuss in Section 7.2. Blume (2003) and Sandholm (2010a) use birth-death chain techniques to study this limit in two-strategy games when mistake probabilities are payoff dependent. In the work closest to the present one, Staudigl (2012) studies the small noise double limit when two populations are matched to play 2x2 coordination games. The analysis uses optimal control methods to evaluate the probabilities of transitions between equilibria. It takes advantage of the fact that each population’s state variable is scalar, and only affects the payoffs of members of the opposing population; this causes the control problem to retain a one-dimensional flavor absent from the general case.

The paper proceeds as follows. Section 2 introduces our class of stochastic evolutionary processes. Section 3 reviews stochastic stability in the small noise limit for the present context. The following three sections study the small noise double limit. Section 4 provides definitions, Section 5 presents the main technical results on the convergence of exit and transition costs, and Section 6 gives the consequences for escape from equilibrium, limiting stationary distributions, and stochastic stability. Section 7 offers concluding discussion. Proofs are presented in the Appendix.

2. The Model

2.1 Finite-population games

We consider games in which agents from a population of size N choose strategies from the common finite strategy set S . The population’s aggregate behavior is described by a *population state* x , an element of the simplex $X = \{x \in \mathbb{R}_+^n : \sum_{i=1}^n x_i = 1\}$, or more specifically, the grid $\mathcal{X}^N = X \cap \frac{1}{N}\mathbb{Z}^n = \{x \in X : Nx \in \mathbb{Z}^n\}$. The standard basis vector $e_i \in X \subset \mathbb{R}^n$ represents the *pure population state* at which all agents play strategy i . States that are not pure are called *mixed population states*.

We identify a *finite-population game* with its payoff function $F^N: \mathcal{X}^N \rightarrow \mathbb{R}^n$, where

$F_i^N(x) \in \mathbb{R}$ is the payoff to strategy i when the population state is $x \in \mathcal{X}^N$. Only the values that the function F_i^N takes on the set $\mathcal{X}_i^N = \{x \in \mathcal{X}^N : x_i > 0\}$ are meaningful, since at the remaining states in \mathcal{X}^N strategy i is unplayed.

Example 2.1. Suppose that $N \geq 2$ agents are matched to play a symmetric two-player normal form game $A \in \mathbb{R}^{n \times n}$. If self-matching is not allowed, then payoffs take the form

$$(1) \quad F_i^N(x) = \frac{1}{N-1} e_i' A(Nx - e_i) = (Ax)_i + \frac{1}{N-1}((Ax)_i - A_{ii}). \quad \blacklozenge$$

In a finite-population game, an agent who switches from strategy i to strategy j when the state is x changes the state to the adjacent state $y = x + \frac{1}{N}(e_j - e_i)$. Thus at any given population state, players playing different strategies face slightly different incentives. To account for this, we use the *clever payoff function* $F_{i \rightarrow \cdot}^N : \mathcal{X}_i^N \rightarrow \mathbb{R}^n$ to denote the payoff opportunities faced by i players at each state $x \in \mathcal{X}_i^N$. The j th component of the vector $F_{i \rightarrow \cdot}^N(x)$ is thus

$$(2) \quad F_{i \rightarrow j}^N(x) = F_j^N(x + \frac{1}{N}(e_j - e_i)).$$

Clever payoffs allow one to describe Nash equilibria of finite-population games in a simple way. The *pure best response correspondence* for strategy $i \in S$ in finite-population game F^N is denoted by $b_i^N : \mathcal{X}_i^N \Rightarrow S$, and is defined by

$$(3) \quad b_i^N(x) = \operatorname{argmax}_{j \in S} F_{i \rightarrow j}^N(x).$$

State $x \in \mathcal{X}^N$ is a *Nash equilibrium* of F^N if no agent can obtain a higher payoff by switching strategies: that is, if $i \in b_i^N(x)$ whenever $x_i > 0$.

Example 2.2. The normal form game $A \in \mathbb{R}^{n \times n}$ is a *coordination game* if $A_{ii} > A_{ji}$ for all distinct $i, j \in S$, so that if one's match partner plays i , one is best off playing i oneself. If F^N is the population game obtained by matching in A without self-matching, then the Nash equilibria of F^N are precisely the pure population states. Thus finite-population matching differs from continuous-population matching, under which the Nash equilibria of the population game correspond to the pure and mixed symmetric Nash equilibria of A .

To see that no mixed population state of F^N can be Nash, suppose that $x \in \mathcal{X}_i^N \cap \mathcal{X}_j^N$ is a Nash equilibrium. Then

$$F_i^N(x) \geq F_j^N(x + \frac{1}{N}(e_j - e_i)) \quad \text{and} \quad F_j^N(x) \geq F_i^N(x + \frac{1}{N}(e_i - e_j)),$$

which with (1) is equivalent to

$$(4) \quad Ne'_i Ax - A_{ii} \geq Ne'_j Ax - A_{ji} \quad \text{and} \quad Ne'_j Ax - A_{jj} \geq Ne'_i Ax - A_{ij}.$$

Summing these inequalities and rearranging yields $(A_{ii} - A_{ji}) + (A_{jj} - A_{ij}) \leq 0$, contradicting that A is a coordination game. Furthermore, pure state e_i is a Nash equilibrium if $F_i^N(x) \geq F_j^N(x + \frac{1}{N}(e_j - e_i))$ for $j \neq i$, which from (4) is true if and only if $A_{ii} > A_{ji}$, as assumed. \blacklozenge

It is convenient to assume that revising agents make decisions by considering clever payoffs, as it ensures that all agents are content if and only if the current state is a Nash equilibrium. The previous example shows that in a coordination game, such a state must be pure. While the use of clever payoffs simplifies the finite population dynamics, it does not affect our results on large population limits in an essential way.

2.2 Noisy best response protocols and unlikelihood functions

In our model of stochastic evolution, agents occasionally receive opportunities to switch strategies. Upon receiving a revision opportunity, an agent selects a strategy by employing a *noisy best response protocol* $\sigma^\eta: \mathbb{R}^n \rightarrow \text{int}(X)$ with *noise level* $\eta > 0$, a function that maps vectors of payoffs to probabilities of choosing each strategy.

To justify its name, the protocol σ^η should recommend optimal strategies with high probability when the noise level is small:

$$(P1) \quad j \notin \operatorname{argmax}_{k \in S} \pi_k \Rightarrow \lim_{\eta \rightarrow 0} \sigma_j^\eta(\pi) = 0.$$

Condition (P1) implies that a unique optimal strategy is assigned a probability that approaches one as the noise level vanishes. For simplicity, we also require that when there are multiple optimal strategies, each retains positive probability in the small noise limit:

$$(P2) \quad j \in \operatorname{argmax}_{k \in S} \pi_k \Rightarrow \lim_{\eta \rightarrow 0} \sigma_j^\eta(\pi) > 0.$$

To analyze large deviations and stochastic stability, we must impose regularity conditions on the rates at which the probabilities of choosing suboptimal strategies vanish as the noise level η approaches zero. To do so, we introduce the *unlikelihood function* $\Upsilon: \mathbb{R}^n \rightarrow \mathbb{R}_+^n$, defined by

$$(5) \quad \Upsilon_j(\pi) = -\lim_{\eta \rightarrow 0} \eta \log \sigma_j^\eta(\pi).$$

This definition can be expressed equivalently as

$$\sigma_j^\eta(\pi) = \exp\left(-\eta^{-1}(\Upsilon_j(\pi) + o(1))\right).$$

Either way, the unlikelihood $\Upsilon_j(\pi)$ represents the rate of decay of the probability that strategy j is chosen as η approaches zero.⁷

We maintain the following assumptions throughout the paper:

- (U1) The limit in (5) exists for all $\pi \in \mathbb{R}^n$.
- (U2) Υ is continuous.
- (U3) $\Upsilon_j(\pi) = 0$ if and only if $j \in \operatorname{argmax}_{k \in S} \pi_k$.

Note that the “if” direction of condition (U3) is implied by condition (P2), and that condition (U1) and the “only if” direction of condition (U3) refine condition (P1).

We proceed with three examples that satisfy the conditions above.

Example 2.3. Logit choice. The logit choice protocol with noise level η , introduced to evolutionary game theory by Blume (1993), is defined by

$$\sigma_j^\eta(\pi) = \frac{\exp(\eta^{-1}\pi_j)}{\sum_{k \in S} \exp(\eta^{-1}\pi_k)}.$$

It is well known that this protocol can be derived from an additive random utility model with extreme-value distributed shocks, or from a model of choice among mixed strategies with control costs given by an entropy function.⁸ It is easy to verify that this protocol satisfies conditions (U1)–(U3) with piecewise linear unlikelihood function

$$\Upsilon_j(\pi) = \max_{k \in S} \pi_k - \pi_j. \quad \blacklozenge$$

Example 2.4. Random utility with averaged shocks. Consider an additive random utility model in which the payoff vector π is perturbed by adding the sample average $\bar{\varepsilon}^m$ of an i.i.d. sequence $\{\varepsilon^\ell\}_{\ell=1}^m$ of random vectors, where the n components of ε^ℓ are drawn from a continuous distribution with unbounded convex support and whose moment generating

⁷Blume (2003) and Sandholm (2010a) place assumptions on the rates of decay of choice probabilities in the context of two-strategy games. Unlikelihood functions for choice problems with many alternatives are introduced by Dokumacı and Sandholm (2011); see Example 2.4 below.

⁸See Anderson et al. (1992) or Hofbauer and Sandholm (2002).

function exists. Writing η for $\frac{1}{m}$, we obtain the protocol

$$\sigma_j^\eta(\pi) = \mathbb{P}\left(j \in \operatorname{argmax}_{k \in S} (\pi_k + \varepsilon_k^m)\right)$$

Dokumacı and Sandholm (2011) show that the limit (5) exists for each $\pi \in \mathbb{R}^n$, and characterize the function Υ in terms of the Cramér transform of ε^ℓ . They also show that Υ_j is nonincreasing in π_j , nondecreasing in π_k for $k \neq j$, and convex (and hence continuous) in π . ♦

Example 2.5. Probit choice. Following Myatt and Wallace (2003), consider an additive random utility model in which the payoff vector π is perturbed by a multivariate normal random vector whose components are independent with common variance η . Since the average of independent normal random variables is normal, the probit choice model is a special case of Example 2.4. Dokumacı and Sandholm (2011) provide an explicit, piecewise quadratic expression for the unlikelihood function Υ . ♦

The only parameterized protocol commonly considered in the literature that does not satisfy our assumptions is the best response with mutations (BRM) protocol of Kandori et al. (1993), the focus of much of the literature to date.⁹ Under this protocol, any suboptimal strategy has unlikelihood 1, and a unique optimal strategy has unlikelihood 0, so condition (U2) must fail.¹⁰ We discuss the BRM protocol in Section 7.2.

2.3 The stochastic evolutionary process

A population game F^N and a revision protocol σ^η define a stochastic evolutionary process. The process runs in discrete time, with each period taking $\frac{1}{N}$ units of clock time.

During each period, a single agent is chosen at random from the population. This agent updates his strategy by applying the noisy best response protocol σ^η . As discussed in Section 2.1, we assume that agents are clever, so that an i player evaluates payoffs using the clever payoff vector $F_{i \rightarrow}^N(x)$ defined in (2).

This procedure described above generates a Markov chain $\mathbf{X}^{N,\eta} = \{X_k^{N,\eta}\}_{k=0}^\infty$ on the state space \mathcal{X}^N . The index k denotes the number of revision opportunities that have occurred to date, and corresponds to $\frac{k}{N}$ units of clock time. The transition probabilities $P_{x,y}^{N,\eta}$ for the

⁹See especially Kandori and Rob (1995, 1998) and Ellison (2000).

¹⁰There is more than one way to specify choice probabilities under the BRM protocol at payoff vectors with multiple optimal strategies—see Sandholm (2010a).

process $\mathbf{X}^{N,\eta}$ are given by

$$(6) \quad P_{x,y}^{N,\eta} \equiv \mathbb{P}\left(X_{k+1}^{N,\eta} = y \mid X_k^{N,\eta} = x\right) = \begin{cases} x_i \sigma_j^\eta(F_{i \rightarrow \cdot}^N(x)) & \text{if } y = x + \frac{1}{N}(e_j - e_i), j \neq i \\ \sum_{i=1}^n x_i \sigma_i^\eta(F_{i \rightarrow \cdot}^N(x)) & \text{if } y = x, \\ 0 & \text{otherwise.} \end{cases}$$

A realization of the process $\mathbf{X}^{N,\eta}$ over its first $\ell^N < \infty$ periods is described by a *path* through \mathcal{X}^N of length ℓ^N , a sequence $\phi^N = \{\phi_k^N\}_{k=0}^{\ell^N}$ in which successive states either are identical or are adjacent in \mathcal{X}^N . Since each period lasts $\frac{1}{N}$ time units, the *duration* of this path in clock time is $T^N = \ell^N/N$.

Since revising agents are chosen at random and play each strategy in S with positive probability, the Markov chain $\mathbf{X}^{N,\eta}$ is irreducible and aperiodic, and so admits a unique stationary distribution, $\mu^{N,\eta}$. It is well known that the stationary distribution is the limiting distribution of the Markov chain, as well as its limiting empirical distribution along almost every sample path.

3. The Small Noise Limit

We now consider the behavior of the stochastic process $\mathbf{X}^{N,\eta}$ as the noise level η approaches zero, proceeding from short run through very long run behavior. Over short to medium time scales, $\mathbf{X}^{N,\eta}$ is nearly a discrete best response process. We introduce this best response process and its recurrent classes in Section 3.1. Over longer periods, runs of suboptimal choices occasionally occur, leading to transitions between the recurrent classes of the best response process. We consider these in Sections 3.2 and 3.3. Finally, over very long time spans, $\mathbf{X}^{N,\eta}$ spends the vast majority of periods at the stochastically stable states, which we define in Section 3.4. Most of the ideas presented in this section can be found in the evolutionary game literature, though not always in an explicit form.

3.1 The discrete best response dynamic and its recurrent classes

In the literature on stochastic evolution in games, the Markov chain $\mathbf{X}^{N,\eta}$ is typically viewed as a perturbed version of some “unperturbed” process $\mathbf{X}^{N,0}$ based on exact best responses. To define the latter process as a Markov chain, one must specify the probability which each best response is chosen when more than one exists. Here we take a more general approach, defining $\mathbf{X}^{N,0}$ not as a Markov chain, but by way of a difference inclusion—in other words, using set-valued deterministic dynamics.

Fix a population size N and a game F^N . Suppose that during each discrete time period, a single agent is chosen from the population, and that he selects a strategy that is optimal given the current population state and his current strategy. If the current state is $x \in \mathcal{X}^N$, then the set of increments in the state that are possible under this procedure is $\frac{1}{N}V^N(x)$, where

$$(7) \quad V^N(x) = \{e_j - e_i : i \in s(x) \text{ and } j \in b_i^N(x)\},$$

and where $s(x) = \{i \in S : x_i > 0\}$ denotes the support of state x . The paths through \mathcal{X}^N that can arise under this procedure are the solutions to the difference inclusion

$$(DBR) \quad x_{k+1}^N - x_k^N \in \frac{1}{N}V^N(x_k^N).$$

We call (DBR) the *discrete best response dynamic*.

We call the set $K^N \subseteq \mathcal{X}^N$ *strongly invariant* under (DBR) if no solution to (DBR) starting in K^N ever leaves K^N . A set that is minimal with respect to this property is called a *recurrent class* of (DBR). We denote the collection of such recurrent classes by \mathcal{K}^N .¹¹

Example 3.1. Let F^N be defined by random matching in the normal form coordination game A as in Example 2.2, so that the Nash equilibria of F^N are the pure states. Suppose in addition that A has the *marginal bandwagon property* of Kandori and Rob (1998): $A_{ii} - A_{ik} > A_{ji} - A_{jk}$ for all $i, j, k \in S$ with $i \notin \{j, k\}$. This property requires that when some agent switches to strategy i from any other strategy k , current strategy i players benefit most. An easy calculation shows that in games with this property, $i \in b_i^N(x)$ implies that $i \in b_k^N(x)$ for all $k \in s(x)$; this is a consequence of the fact that a strategy i player has one less opponent playing strategy i than a strategy $k \neq i$ player.

Now suppose that state $x \in \mathcal{X}^N$ is not a Nash equilibrium. Then there are distinct strategies i and j such that $j \in s(x)$ (j is in use) and $i \in b_j^N(x)$ (i is optimal for agents playing j), so that a step from x to $y = x + \frac{1}{N}(e_i - e_j)$ is allowed under (DBR). Since $i \in b_j^N(x)$ is equivalent to $i \in b_i^N(x + \frac{1}{N}(e_i - e_j))$, the marginal bandwagon property (specifically, the claim ending the previous paragraph) implies that $i \in b_k^N(y)$ for all $k \in s(y)$. Repeating this argument shows that any path from y along which the number of strategy i players increases until pure state e_i is reached is a solution to (DBR). We conclude that the recurrent

¹¹One can represent the solutions and the recurrent classes of (DBR) using a suitably chosen Markov chain $\mathbf{X}^{N,*}$. Define $\mathbf{X}^{N,*}$ by supposing that during each period, a randomly chosen agent receives a revision opportunity and switches to a best response, choosing each with equal probability (or, more generally, with any positive probability). Then a finite-length path is a solution to (DBR) if and only if it has positive probability under $\mathbf{X}^{N,*}$, and the recurrent classes of (DBR) as defined above are the recurrent classes of $\mathbf{X}^{N,*}$.

classes of (DBR) correspond to the pure states, $\mathcal{X}^N = \{\{e_1\}, \dots, \{e_n\}\}$, as shown by Kandori and Rob (1998).¹² ♦

Example 3.2. Again let F^N be defined by random matching in the normal form coordination game A . If $x \in X^N$ is not Nash, there is a strategy j in the support of x satisfying $j \notin b_j^N(x)$. Lemma A.1 in Appendix A.1 shows that in this case, there is a solution to (DBR) starting from x along which the number of j players decreases until j is unused.

Now suppose further that in game F^N , switching to an unused strategy is never optimal: $j \in b_j^N(x)$ implies that $x_j > 0$. In this case, applying Lemma A.1 inductively shows that from every state $x \in X^N$, there is a solution to (DBR) that terminates at a pure state, implying that $\mathcal{X}^N = \{\{e_1\}, \dots, \{e_n\}\}$. ♦

We conjecture that the set of recurrent classes of (DBR) is $\mathcal{X}^N = \{\{e_1\}, \dots, \{e_n\}\}$ for any coordination game as defined in Example 2.2. Example 4.1 establishes a version of this claim for the large population limit.

3.2 Step costs and path costs

When the noise level η is small, the process $X^{N,\eta}$ will linger in recurrent classes, but will occasionally transit between them. We now work toward describing the probabilities of these transitions in the small noise limit.

To begin, we define the *cost of a step* from $x \in X^N$ to $y \in X^N$ by

$$(8) \quad c_{x,y}^N = -\lim_{\eta \rightarrow 0} \eta \log P_{x,y}^{N,\eta}.$$

with the convention that $-\log 0 = +\infty$. Thus $c_{x,y}^N$ is the exponential rate of decay of the probability of a step from x to y as η approaches 0. Using definitions (5) and (6), we can represent step costs in terms of the game's payoff function and the protocol's unlikelihood function:

$$(9) \quad c_{x,y}^N = \begin{cases} \Upsilon_j(F_{i \rightarrow \cdot}^N(x)) & \text{if } y = x + \frac{1}{N}(e_j - e_i) \text{ and } j \neq i, \\ \min_{i \in S(x)} \Upsilon_i(F_{i \rightarrow \cdot}^N(x)) & \text{if } y = x, \text{ and} \\ +\infty & \text{otherwise.} \end{cases}$$

The important case in (9) is the first one, which says that the cost of a step in which an i player switches to strategy j is the unlikelihood of strategy j given i 's current payoff

¹²Unlike our model, the model of Kandori and Rob (1995, 1998) allows multiple revisions during each period; see also footnote 19 below.

opportunities.¹³ By virtue of (9) and condition (U3), a step has cost zero if and only if it is feasible under the discrete best response dynamic:

$$(10) \quad c_{x,y}^N = 0 \Leftrightarrow y - x \in V^N(x).$$

The cost of path $\phi^N = \{\phi_k^N\}_{k=0}^{\ell^N}$ of length $\ell^N < \infty$ is the sum of the costs of its steps:

$$(11) \quad c^N(\phi^N) = \sum_{k=0}^{\ell^N-1} c_{\phi_k^N, \phi_{k+1}^N}^N.$$

Definitions (6) and (9) imply that the cost of a path is the rate at which the probability of following this path decays as the noise level vanishes: for fixed N , we have

$$\mathbb{P}\left(X_k^{N,\eta} = \phi_k^N, k = 0, \dots, \ell^N \mid X_0^{N,\eta} = \phi_0^N\right) = \prod_{k=0}^{\ell^N-1} P_{\phi_k^N, \phi_{k+1}^N}^{N,\eta} \approx \exp(-\eta^{-1} c^N(\phi^N)).$$

where \approx refers to the order of magnitude in η as η approaches zero. By statement (10), path ϕ^N has cost zero if and only if it is a solution to (DBR).

3.3 Exit costs and transition costs

We now consider escape from and transitions between recurrent classes. Let $K^N \in \mathcal{K}^N$ be a recurrent class of (DBR), and let $\Xi^N \subset \mathcal{X}^N$ be a set of states. We define $\Phi^N(K^N, \Xi^N)$ to be the set of finite-length paths through \mathcal{X}^N with initial state in K^N and terminal state in Ξ^N , so that

$$(12) \quad C^N(K^N, \Xi^N) = \min\{c^N(\phi^N) : \phi^N \in \Phi^N(K^N, \Xi^N)\}$$

is the minimal cost of a path from K^N to Ξ^N .

If Ξ^N is a union of recurrent classes from \mathcal{K}^N , we define the *weak basin of attraction* of Ξ^N , denoted $\mathcal{W}^N(\Xi^N)$, to be the set of states in \mathcal{X}^N from which there is a zero-cost path that terminates at a state in Ξ^N . Notice that by definition,

$$C^N(K^N, \Xi^N) = C^N(K^N, \mathcal{W}^N(\Xi^N)).$$

We also define $\Omega^N(K^N, \mathcal{W}^N(\Xi^N)) \subseteq \mathcal{W}^N(\Xi^N)$ to be the set of terminal states of cost mini-

¹³The second case of (9) indicates that at a state where no agent is playing a best response, staying still is costly. Since staying still does not facilitate transitions between recurrent classes, this possibility is not realized on minimum cost paths, but we must account for it carefully in what follows—see Section 4.3.2.

mizing paths from K^N to $\mathcal{W}^N(\Xi^N)$ that do not hit $\mathcal{W}^N(\Xi^N)$ until their final step.

Two specifications of the target set Ξ^N are of particular interest. First, let

$$(13) \quad \tilde{K}^N = \bigcup_{L^N \in \mathcal{K}^N \setminus \{K^N\}} L^N$$

be the union of the recurrent classes other than K^N . We call $C^N(K^N, \tilde{K}^N)$ the *cost of exit* from K^N .¹⁴ Proposition 3.3 provides an interpretation of this quantity. Here $\tau^{N,\eta}(\Xi^N) = \min\{k: X_k^{N,\eta} \in \Xi^N\}$ denotes the time at which the process $\mathbf{X}^{N,\eta}$ first hits Ξ^N .

Proposition 3.3. *Suppose that $X_0^{N,\eta} = x^N \in K^N$ for all η . Then*

- (i) $\lim_{\eta \rightarrow 0} \eta \log \mathbb{E} \tau^{N,\eta}(\tilde{K}^N) = \lim_{\eta \rightarrow 0} \eta \log \mathbb{E} \tau^{N,\eta}(\mathcal{W}^N(\tilde{K}^N)) = C^N(K^N, \tilde{K}^N);$
- (ii) $\lim_{\eta \rightarrow 0} \eta \log \mathbb{P}(X_{\tau^{N,\eta}(\mathcal{W}^N(\tilde{K}^N))}^{N,\eta} = y) = 0$ if and only if $y \in \Omega^N(K^N, \mathcal{W}^N(\tilde{K}^N))$.

Part (i) of the proposition shows that when η is small, the expected time required for the process to escape from K^N to another recurrent class is of order $\exp(\eta^{-1} C^N(K^N, \tilde{K}^N))$. Part (ii) shows that the states in $\mathcal{W}^N(\tilde{K}^N)$ most likely to be reached first are the terminal states of cost minimizing paths from K^N to $\mathcal{W}^N(\tilde{K}^N)$. Both parts follow by standard arguments from Proposition 4.2 of Catoni (1999), which provides a discrete-state analogue of the Freidlin and Wentzell (1998) theory.

Proposition 3.3 concerns behavior within the *strong basin of attraction* of K^N , the set of states $\mathcal{S}^N(K^N) = \mathcal{X}^N \setminus \mathcal{W}^N(\tilde{K}^N) \subseteq \mathcal{W}^N(K^N)$ from which there is no zero-cost path to any other recurrent class. But to understand the global behavior of the process, we must also consider transitions from K^N to each other individual recurrent class in \mathcal{K}^N .

When $L^N \in \mathcal{K}^N$, we call $C^N(K^N, L^N)$ the *cost of a transition* from K^N to L^N . Intuitively, $C^N(K^N, L^N)$ describes the likely order of magnitude of the time until Ξ^N is reached. But while the analogue of Proposition 3.3(ii) on the likely points of entry into $\mathcal{W}(L^N)$ is true, the analogue of Proposition 3.3(i) on the expected hitting time of L^N is false in general, since this expectation may be driven by a low probability of becoming stuck in some third recurrent class.¹⁵

3.4 Stationary distribution asymptotics and stochastic stability

The transition costs $C^N(K^N, L^N)$ are the basic ingredient in Freidlin and Wentzell's (1998) graph-theoretic characterization of limiting stationary distributions and stochastic

¹⁴Thus the cost of exit from K^N corresponds to the *radius* of K^N as defined by Ellison (2000).

¹⁵This is the reason for the correction term appearing in Proposition 4.2 of Catoni (1999). See Freidlin and Wentzell (1998, p. 197–198) for a clear discussion of this point.

stability. According to this characterization, there is a function $\Delta r^N: \mathcal{X}^N \rightarrow \mathbb{R}_+$, defined in terms of the aggregate costs of certain graphs on \mathcal{X}^N , such that

$$(14) \quad -\lim_{\eta \rightarrow 0} \eta \log \mu^{N,\eta}(x) = \Delta r^N(x) \text{ for all } x \in \mathcal{X}^N.$$

Thus $\Delta r^N(x)$ describes the exponential rate of decay of the stationary distribution weight on x as η approaches zero.

We call state $x \in \mathcal{X}^N$ *stochastically stable in the small noise limit* if as η approaches 0, its stationary distribution weight $\mu^{N,\eta}(x)$ does not vanish at an exponential rate.¹⁶ By virtue of (14), state x is stochastically stable in this sense if and only if $\Delta r^N(x) = 0$. Since these ideas are well known in evolutionary game theory,¹⁷ we postpone the detailed presentation until Section 6.2.

4. The Small Noise Double Limit

The exit costs and transition costs introduced in Section 3.3, defined in terms of minimum cost paths between sets of states in \mathcal{X}^N , describe the transitions of the process $\mathbf{X}^{N,\eta}$ between recurrent classes when the noise level η is small. When step costs depend on payoffs, finding these minimum cost paths is a challenging computational problem.

We contend with this difficulty by taking a second limit: after taking the noise level η to 0, we take the population size N to infinity, thus evaluating behavior in the small noise limit when the population size is large. In the remainder of this paper, we show how one can evaluate this double limit by approximating the discrete constructions from the previous section by continuous ones. In particular, taking the second limit here turns the path cost minimization problem (12) into an optimal control problem. Although this problem is nonsmooth and multidimensional, it is nevertheless simple enough to admit analytical solutions in interesting cases.

4.1 Limits of finite-population games

To consider large population limits, we must specify a notion of convergence for sequences $\{F^N\}_{N=N_0}^\infty$ of finite-population games. If such a sequence converges, its limit is a (*continuous*) *population game*, $F: X \rightarrow \mathbb{R}^n$, which we take to be a continuous function

¹⁶Explicitly, this means that for all $\delta > 0$ there is an $\eta_0 > 0$ such that for all $\eta < \eta_0$, $\mu^{N,\eta}(x) > \exp(-\eta\delta)$. This definition of stochastic stability is slightly less demanding than the one appearing in Kandori et al. (1993) and Young (1993); Sandholm (2010b, Sec. 12.A.5) explains this distinction in detail.

¹⁷See Young (1993, 1998), Kandori and Rob (1995), and Sandholm (2010b, Section 12.A).

from the compact set X to \mathbb{R} . The pure and mixed best response correspondences for the population game F are denoted by $b: X \Rightarrow S$ and $B: X \Rightarrow X$, and are defined by

$$b(x) = \operatorname{argmax}_{i \in S} F_i(x) \quad \text{and} \quad B(x) = \{y \in X: \operatorname{supp}(y) \subseteq b(x)\} = \operatorname{argmax}_{y \in X} y'F(x).$$

State x is a *Nash equilibrium* of F if $i \in b(x)$ whenever $x_i > 0$, or, equivalently, if $x \in B(x)$.

The notion of convergence we employ for the sequence $\{F^N\}_{N=N_0}^\infty$ is uniform convergence, which asks that

$$(15) \quad \lim_{N \rightarrow \infty} \max_{x \in \mathcal{X}^N} |F^N(x) - F(x)| = 0,$$

where $|\cdot|$ denotes the ℓ^1 norm on \mathbb{R}^n . It is easy to verify that under this notion of convergence, the Nash equilibrium correspondences for finite-population games are “upper hemicontinuous at infinity”: if the sequence of games $\{F^N\}$ converges to F , the sequence of states $\{x^N\}$ converges to x , and if each x^N is a Nash equilibrium of the corresponding F^N , then x is a Nash equilibrium of F .

When agents are matched to play a symmetric two-player normal form game $A \in \mathbb{R}^{n \times n}$ (Example 2.1), it is easy to verify that uniform convergence obtains with the limit game given by $F(x) = Ax$. It is also easy to verify that if a sequence of population games converges uniformly, then the clever payoff functions associated with that game all converge uniformly to the same limit.

4.2 The complete best response dynamic and limiting recurrent classes

The solutions of the discrete best response dynamic (DBR) are the paths through \mathcal{X}^N that can be traversed at zero cost. To define the analogous dynamic for the large population limit, let $S(x) = \{y \in X: s(y) \subseteq s(x)\}$ be the set of states whose supports are contained in the support of x . Then the *complete best response dynamic* is the differential inclusion

$$\begin{aligned} \text{(CBR)} \quad \dot{x} &\in B(x) - S(x) \\ &= \{\beta - \alpha: \beta \in B(x), \alpha \in S(x)\} \\ &= \operatorname{conv}(\{e_j - e_i \mid i \in s(x), j \in b(x)\}). \end{aligned}$$

Comparing the final expression above to definition (7), we see that (CBR) is the continuous-time analogue of the discrete best response dynamic (DBR), obtained by taking the large N limit of (DBR) and convexifying the result. We will soon see that solutions to (CBR) correspond to zero-cost continuous paths under our limiting path cost function.

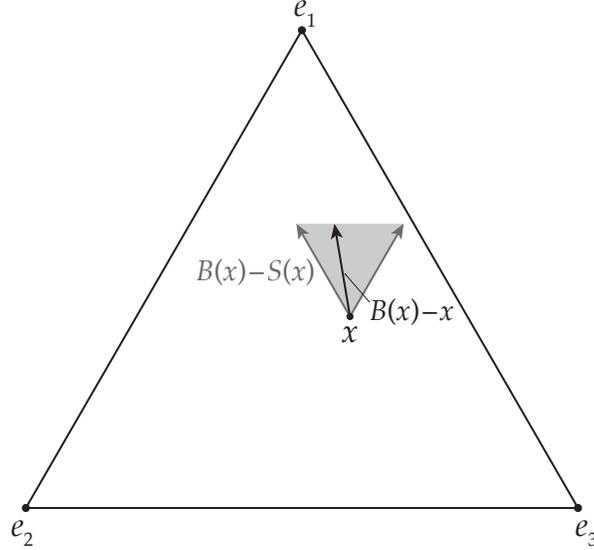


Figure 1: The dynamics (BR) and (CBR) and in a three-strategy game from a state x with $b(x) = \{1\}$.

For intuition, we contrast (CBR) with the standard model of best response strategy revision in a large population—the *best response dynamic* of Gilboa and Matsui (1991):

$$\text{(BR)} \quad \dot{x} \in B(x) - x.$$

To obtain (BR) as the limit of finite-population dynamics, one assumes that in each discrete time period, an agent is chosen at random from the population and then updates to a best response. As the population size grows large, the law of large numbers ensures that the rates at which the various strategies are abandoned are proportional to the prevalences of the strategies in the population, generating the $-x$ outflow term in (BR).¹⁸ Thus at states where the best response is unique, (BR) specifies a single vector of motion, as shown in Figure 1 at a state at which the unique best response is strategy 1. Under (DBR), there is no presumption that revision opportunities are assigned at random. Thus, in the large population limit (CBR), the strategies present in the population can be abandoned at any relative rates, leading to the $-S(x)$ outflow term in (CBR). In Figure 1, the set of vectors of motion under (CBR) is the convex hull of the vectors $e_1 - e_2$, $e_1 - e_3$, and $\mathbf{0}$.¹⁹

In the classes of coordination games considered in Examples 3.1 and 3.2, the set \mathcal{K}^N of recurrent classes of the discrete best response dynamic (DBR) is equal to $\mathcal{K} = \{\{e_1\}, \dots, \{e_n\}\}$

¹⁸See Roth and Sandholm (2013) for a formal limit result.

¹⁹Kandori and Rob (1995, 1998) consider a discrete-time best response dynamic in which any subset of the players may revise during any period; for instance, the entire population may switch to a current best response simultaneously. Figure 1 of Kandori and Rob (1995), used to illustrate this discrete-time dynamic, resembles Figure 1 above, but the processes these figures represent are different.

for every population size. We now show that in any coordination game, we can view this \mathcal{K} as the set of “recurrent classes” of the complete best response dynamic (CBR).

Example 4.1. Consider the continuous population game $F(x) = Ax$ generated by a coordination game A (Example 2.2). Since each pure state e_i is a strict equilibrium, the unique solution to (CBR) starting at e_i is the stationary trajectory. At any state ξ in the best response region $\mathcal{B}_i = \{x \in X: (Ax)_i \geq (Ax)_j \text{ for all } j \in S\}$, the vector $e_i - \xi$, which points directly from ξ to e_i , is a feasible direction of motion under the best response dynamic (BR). Since \mathcal{B}_i is convex and contains e_i , motion can continue toward e_i indefinitely: the trajectory $\{x_t\}_{t \geq 0}$, defined by $x_t = e^{-t}\xi + (1 - e^{-t})e_i$, is a solution to (BR), and hence a solution to (CBR). Thus for any coordination game $F(x) = Ax$, starting from any initial condition $\xi \in X$, there is solution to (CBR) that converges to a pure, and hence stationary, population state. ♦

More generally, the exact positions of the recurrent classes of the discrete best response dynamic (DBR) will vary with the population size. To allow for this, we assume that there is a set $\mathcal{K} = \{K_1, \dots, K_\kappa\}$ of disjoint closed subsets of X called *limiting recurrent classes*. To justify this name, we require that for some constant $d > 0$ and all large enough population sizes N , the dynamic (DBR) has κ recurrent classes, $\mathcal{K}^N = \{K_1^N, \dots, K_\kappa^N\}$, and that

$$(16) \quad \text{dist}(K_i^N, K_i) \leq \frac{d}{N} \text{ for all } i \in \{1, \dots, \kappa\},$$

where $\text{dist}(K_i^N, K_i)$ denotes the Hausdorff distance between K_i^N and K_i .²⁰

4.3 Costs of continuous paths

To evaluate stochastic stability in the small noise double limit, we need to determine the costs $C^N(K^N, \Xi^N)$, defined by the discrete cost minimization problems (12) on X^N , for large values of N . To prepare for our continuous approximation of these problems, we now introduce a definition of costs for continuous paths through the simplex X .

4.3.1 Discrete paths, derivatives, and interpolations

Let $\phi^N = \{\phi_k^N\}_{k=0}^{\ell^N}$ be a path for the N -agent process. Since each period of this process takes $\frac{1}{N}$ units of clock time, we define

$$(17) \quad \dot{\phi}_k^N = N(\phi_{k+1}^N - \phi_k^N)$$

²⁰That is, $\text{dist}(K_i^N, K_i) = \max \left\{ \max_{x \in K_i^N} \min_{y \in K_i} |x - y|, \max_{y \in K_i} \min_{x \in K_i^N} |x - y| \right\}$.

to be the discrete right derivative of path ϕ^N at time k . Let $i^N(k) \in S$ and $j^N(k) \in S$ denote the pre- and post-revision strategies of the agent who revises in period k . Then $\dot{\phi}_{k+1}^N = \dot{\phi}_k^N + \frac{1}{N}(e_{j^N(k)} - e_{i^N(k)})$, and hence

$$(18) \quad \dot{\phi}_k^N = e_{j^N(k)} - e_{i^N(k)}.$$

Note that if $i^N(k) = j^N(k)$, so that the revising agent does not switch strategies, then $\dot{\phi}_k^N = \mathbf{0}$.

Each discrete path $\{\phi_k^N\}_{k=0}^{\ell^N}$ induces a continuous path $\{\phi_t^{(N)}\}_{t \in [0, \ell^N/N]}$ via piecewise affine interpolation:

$$(19) \quad \phi_t^{(N)} = \phi_{\lfloor Nt \rfloor}^N + (Nt - \lfloor Nt \rfloor)(\phi_{\lfloor Nt \rfloor + 1}^N - \phi_{\lfloor Nt \rfloor}^N).$$

This definition too accounts for each period in the N -agent process lasting $\frac{1}{N}$ units of clock time. Evidently, the derivative $\dot{\phi}^{(N)}$ of this process agrees with the discrete derivative $\dot{\phi}^N$ defined in (17), in the sense that

$$(20) \quad \dot{\phi}_t^{(N)} = \dot{\phi}_{\lfloor Nt \rfloor}^N \text{ whenever } Nt \notin \mathbb{Z}.$$

Speed of motion along a continuous path $\{\phi_t\}_{t \in [0, T]}$ is measured most naturally by evaluating the ℓ^1 norm $|\dot{\phi}_t| = \sum_{i \in S} |(\dot{\phi}_t)_i|$ of $\dot{\phi}_t \in TX \equiv \{z \in \mathbb{R}^n : \sum_{i \in S} z_i = 0\}$, as this norm makes it easy to separate the contributions of strategies that are gaining players from those of strategies that are losing players. If for $z \in \mathbb{R}^n$ we define $[z]_+ \in \mathbb{R}_+^n$ and $[z]_- \in \mathbb{R}_+^n$ by $([z]_+)_i = [z]_i$ and $([z]_-)_i = [z]_i$, then by virtue of equations (18) and (20), any interpolated path $\phi^{(N)}$ satisfies the following bound on its speed:

$$(21) \quad |[\dot{\phi}_t^{(N)}]_+| \equiv |[\dot{\phi}_t^{(N)}]_-| \leq 1, \text{ and thus } |\dot{\phi}_t^{(N)}| \leq 2.$$

4.3.2 Costs of continuous paths

To work toward our definition of the cost of a continuous path, we now express the path cost function (11) in a more suggestive form. Let $\langle \cdot, \cdot \rangle$ denote the standard inner product on \mathbb{R}^n , and let $\phi^N = \{\phi_k^N\}_{k=0}^{\ell^N}$ be a discrete path. If $j^N(k) \neq i^N(k)$, then definitions (9) and (17) imply that the cost of step k is

$$(22) \quad c_{\phi_k^N, \phi_{k+1}^N}^N = \Upsilon_{j^N(k)}(F_{i^N(k)}^N(\phi_k^N)) = \langle \Upsilon(F_{i^N(k)}^N(\phi_k^N)), [\dot{\phi}_k^N]_+ \rangle.$$

If $j^N(k) = i^N(k)$, so that the revising agent does not switch strategies, then $\dot{\phi}_k^N$ equals $\mathbf{0}$; thus the rightmost expression of (22) evaluates to 0 for such null steps. This disagrees

with the second case of (9) when there is no best response to ϕ_k^N is in its support. Since this discrepancy only arises when a path lingers at some such state, it is inconsequential when determining the minimal cost of a path between subsets of \mathcal{X}^N , as there is always a least cost path that does not linger at all.

Summing up the step costs, the cost (11) of a discrete path ϕ^N without null steps can be expressed as

$$(23) \quad c^N(\phi^N) = \sum_{k=0}^{\ell^N-1} c_{\phi_k^N, \phi_{k+1}^N}^N = \sum_{k=0}^{\ell^N-1} \langle \Upsilon(F_{i^N(k) \rightarrow}^N(\phi_k^N)), [\dot{\phi}_k^N]_+ \rangle.$$

Now let $\phi: [0, T] \rightarrow X$ be absolutely continuous and *non-pausing*, meaning that $|\dot{\phi}_t| \neq 0$ for almost all $t \in [0, T]$. Since $F_{i \rightarrow}^N \approx F$ for large N , the form of the path costs in (23) suggests that the cost of ϕ should be defined as

$$(24) \quad c(\phi) = \int_0^T \langle \Upsilon(F(\phi_t)), [\dot{\phi}_t]_+ \rangle dt.$$

This derivation is informal; the formal justification of definition (24) is provided by the approximation results to follow.

While the discrete path cost function (23) only concerns paths with discrete derivatives of the basic form $\dot{\phi}_k^N = e_{j^N(k)} - e_{i^N(k)}$, definition (24) allows any absolutely continuous path with derivatives $\dot{\phi}_t$ in $Z = \text{conv}(\{e_j - e_i : i, j \in S\})$, or indeed in the tangent space TX . This extension combines two new ingredients. First, allowing $\dot{\phi}_t$ to be the weighted average of a number of vectors $e_j - e_i$ makes it possible to approximate the cost of a continuous path by the costs of rapidly oscillating discrete paths, a point we discuss further in Section 5.3. Second, by virtue of the linear homogeneity of the integrand of (24) in $\dot{\phi}_t$, the cost of a continuous path is independent of the speed at which it is traversed.

Finally, we observe that a non-pausing absolutely continuous path ϕ has zero cost under (24) if and only if it is a solution of the complete best response dynamic (CBR).

5. The Convergence Theorem

In Section 3.3, we defined the minimal cost $C^N(K^N, \Xi^N)$ of a discrete path from recurrent class $K^N \in \mathcal{K}^N$ to set $\Xi^N \subset \mathcal{X}^N$. We now consider a sequence of such problems, where the recurrent classes K^N converge to the limiting recurrent class $K \in \mathcal{K}$ as in condition (16),

and where the target sets $\Xi^N \subset \mathcal{X}^N$ converge to a closed set $\Xi \subset X$ in the same sense:

$$(25) \quad \text{dist}(\Xi^N, \Xi) \leq \frac{d}{N}$$

for some $d > 0$ and all N large enough.

Let $\Phi(K, \Xi)$ be the set of absolutely continuous paths of arbitrary duration through X from K to Ξ , and define

$$(26) \quad C(K, \Xi) = \min\{c(\phi) : \phi \in \Phi(K, \Xi)\}$$

to be the minimal cost of a continuous path from K to Ξ . Our aim in this section is to show that the normalized optimal values of the discrete problems converge to the optimal value of the continuous problem:

$$(27) \quad \lim_{N \rightarrow \infty} \frac{1}{N} C^N(K^N, \Xi^N) = C(K, \Xi).$$

This conclusion will justify the definition (24) of the cost of a non-pausing absolutely continuous path, and will provide the tool needed to evaluate exit times, stationary distribution asymptotics, and stochastic stability in the large population double limit.

5.1 Assumptions

We prove our results under two assumptions about the minimum cost path problems (12) and (26). To state the first assumption, we recall that the *duration* $T^N = \ell^N/N$ of the discrete path $\{\phi_k^N\}_{k=0}^{\ell^N}$ is the number of units of clock time it entails.

Assumption 1. *There exists a constant $\bar{T} < \infty$ such that for all $K^N \in \mathcal{K}^N$, $\Xi^N \subset \mathcal{X}^N$, and N , there is a path of duration at most \bar{T} that achieves the minimum in (12).*

Since state space \mathcal{X}^N is finite, cost-minimizing paths between subsets of Ξ^N can always be assumed to have finite length. Assumption 1 imposes a uniform bound on the amount of clock time that these optimal paths require. It thus requires that cost-minimizing paths not become extremely convoluted as the population size grows large, as might be possible if despite the uniform convergence of payoff functions in (15), the step costs $c_{x,y}^N$ defined in (9) became highly irregular functions of the current population state.

To introduce our assumption for the continuous problem (26), we need some additional definitions. Let $\phi : [0, T] \rightarrow X$ be a continuous path. We call ϕ *monotone* if we can express the strategy set S as the disjoint union $S_+ \cup S_-$, with ϕ_j nondecreasing for $j \in S_+$ and ϕ_i

nonincreasing for $i \in S_-$. If M is a positive integer, we say that ϕ is M -piecewise monotone if its domain $[0, T]$ can be partitioned into M subintervals such that ϕ is monotone on each; if this is true for some M , we say that ϕ is *piecewise monotone*. Monotonicity and piecewise monotonicity for discrete paths are defined analogously.

In addition, we say that piecewise monotone path ϕ *obeys the speed limit* if

$$(28) \quad |[\dot{\phi}_t]_+| \equiv |[\dot{\phi}_t]_-| \leq 1, \text{ and thus } |\dot{\phi}_t| \leq 2 \text{ for almost all } t \in [0, T].$$

Equivalently, ϕ obeys the speed limit if $\dot{\phi}_t \in Z = \text{conv}(\{e_j - e_i : i, j \in S\})$ for almost all t . If the inequalities in (28) bind, then ϕ is said to *move at full speed*. By the linear homogeneity of the integrand of cost function (24), there is no loss if the minimum in (26) is taken over paths in $\Phi(K, \Xi)$ that obey the speed limit, or over paths in $\Phi(K, \Xi)$ that move at full speed.

Assumption 2. *There exist constants $\tilde{T} < \infty$ and $\tilde{M} < \infty$ such that for all $K \in \mathcal{K}$ and $\Xi \in \mathcal{K} \cup \{\{x\}\}_{x \in X}$, there is an \tilde{M} -piecewise monotone, full speed path of duration at most \tilde{T} that achieves the minimum in (26).*

Since the state space X is compact and the integrand of the cost function is (24) continuous, and since we may work with the compact, convex set of controls Z , it is reasonable to expect the minimum in (26) to be achieved by some finite-duration path. Piecewise monotonicity is a mild regularity condition on the form of the minimizer. In practice, one applies the results developed below by explicitly solving control problem (26) (see Sandholm and Staudigl (2014)). In so doing, one verifies Assumption 2 directly.

In order to appeal to Assumption 2, we assume in what follows that the target set Ξ is either a limiting recurrent class or a singleton.²¹

5.2 The lower bound

To establish the convergence claim in (27), we must show that $C(K, \Xi)$ provides both a lower and an upper bound on the limiting behavior of $\frac{1}{N}C^N(K^N, \Xi^N)$.

The key to obtaining the lower bound is to show that if the normalized costs $\frac{1}{N}c^N(\phi^N)$ of a sequence of discrete paths of bounded durations converge, then the costs $c(\phi^{(N)})$ of the corresponding linear interpolations converge to the same limit. This is the content of the following proposition. Its proof, which is based on continuity arguments, is presented in Appendix A.2.

²¹For the results to follow that only concern recurrent classes, it is enough in Assumption 2 to consider target sets in \mathcal{K} . Singleton target sets are needed in Theorem 6.3 to derive the asymptotics of the stationary distribution on the entire state space, rather than just its asymptotics on the recurrent classes.

Proposition 5.1. *Let $\{\phi^N\}_{N=N_0}^\infty$ be a sequence of paths with durations at most \bar{T} and whose costs satisfy $\lim_{N \rightarrow \infty} \frac{1}{N} c^N(\phi^N) = C^*$. Then the corresponding sequence $\{\phi^{(N)}\}_{N=N_0}^\infty$ of linear interpolations satisfies $\lim_{N \rightarrow \infty} c(\phi^{(N)}) = C^*$.*

Now consider a sequence (or, if necessary, a subsequence) of optimal discrete paths $\phi^N \in \Phi^N(K^N, \Xi^N)$ for problem (12) with durations $T^N \leq \bar{T}$ (cf. Assumption 1) and whose costs converge to C^* . By Proposition 5.1, the costs of their linear interpolations $\phi^{(N)} \in \Phi(K, \Xi)$ also converge to C^* . We can extend these to paths in $\Phi(K, \Xi)$ by adding subpaths linking K to $\phi_0^{(N)} \in K^N$ and $\phi_{T^N}^{(N)} \in L^N$ to L . Conditions (16) and (25) imply that this can be done at negligible cost. This argument yields the following result, whose proof is presented in Appendix A.3.

Proposition 5.2. $\liminf_{N \rightarrow \infty} \frac{1}{N} C^N(K^N, \Xi^N) \geq C(K, \Xi)$.

5.3 The upper bound

The key to obtaining the upper bound is to show that given a continuous path ϕ with cost $c(\phi)$, we can find a sequence of discrete paths $\{\phi^N\}$ whose normalized costs approach $c(\phi)$. The natural approach to this problem is to define each ϕ^N as a suitable discrete approximation of ϕ , and then to use continuity arguments to establish the convergence of normalized costs. But unlike the argument behind Proposition 5.1, the cost convergence argument here is not straightforward. The earlier argument took advantage of the fact that every discrete path induces a continuous path via linear interpolation. Here, the discrete approximation of the continuous path must be constructed explicitly.

Moreover, there are limits to what a discrete approximation can achieve. As definitions (23) and (24) state, a path's cost depends its derivatives at each point of time; these derivatives specify the sequence of revisions that occur over the course of the path. However, one cannot always construct discrete approximations ϕ^N whose derivatives approximate those of the continuous path ϕ .

As an illustration, consider Figure 2(i), which presents a continuous path ϕ through X from vertex e_1 to the barycenter $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. As this path is followed, the state moves in direction $\frac{1}{2}(e_2 + e_3) - e_1$: the mass playing strategy 1 falls over time, while the masses playing strategies 2 and 3 rise at equal rates. But discrete paths through \mathcal{X}^N are unable to move in this direction. At best, they can alternate between increments $\frac{1}{N}(e_2 - e_1)$ (i.e., switches by a single agent from 1 to 2) and $\frac{1}{N}(e_3 - e_1)$. The states in the resulting discrete paths are all close to states in ϕ . But the alternation of increments needed to stay close to ϕ prevents the derivatives of the discrete paths from converging as N grows large.

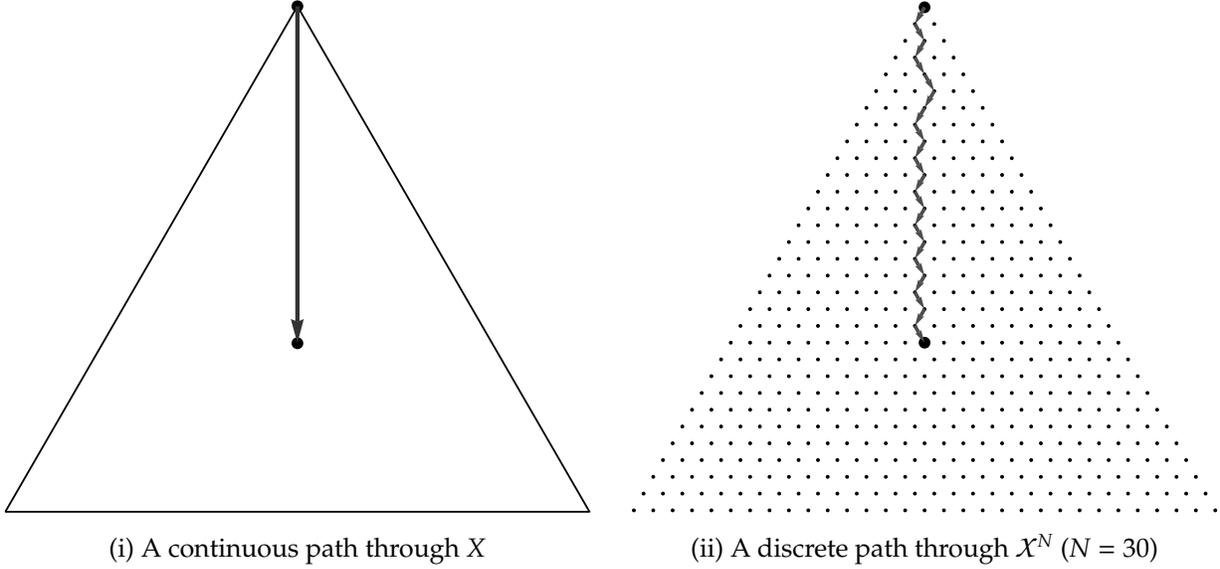


Figure 2: Discrete approximation of continuous paths.

As we now show, it is possible to overcome this difficulty so long as path ϕ is piecewise monotone.²² Proposition 5.3 shows that if ϕ is monotone and moves at full speed, then we can find discrete paths ϕ^N that are also monotone and that closely approximate ϕ , in that ϕ^N is within $\frac{2n}{N}$ of ϕ in the uniform norm.

Proposition 5.3. *Suppose $\phi = \{\phi_t\}_{t \in [0, T]}$ is monotone and moves at full speed. If $N \geq \frac{1}{T}$, there is an $s^N \in [0, \frac{1}{N})$ and a feasible monotone path $\phi^N = \{\phi_k^N\}_{k=0}^{\ell^N}$, $\ell^N = \lfloor N(T - s^N) \rfloor$, satisfying*

$$(29) \quad \max_{0 \leq k \leq \ell^N} \left| \phi_k^N - \phi_{s^N + \frac{k}{N}} \right| \leq \frac{2n}{N}.$$

A constructive proof of this proposition is presented in Appendix A.4.

Next, Proposition 5.4 shows that the normalized costs of the discrete paths so constructed converge to the cost of the original path ϕ .

Proposition 5.4. *Suppose that the path $\{\phi_t\}_{t \in [0, T]}$ is monotone and moves at full speed, and that the paths $\{\{\phi_k^N\}_{k=0}^{\ell^N}\}_{N=N_0}^{\infty}$ are monotone and approximate ϕ in the sense of (29). Then $\lim_{N \rightarrow \infty} \frac{1}{N} c^N(\phi^N) = c(\phi)$.*

The proof of Proposition 5.4 is presented in Appendix A.5, but we explain the logic of the proof here. By equation (23) and definition (17) of $\dot{\phi}^N$, we can express the normalized

²²As an aside, we note that the continuous and discrete paths in Figure 2 are both monotone with $S_+ = \{2, 3\}$ and $S_- = \{1\}$.

cost of path ϕ^N as

$$(30) \quad \frac{1}{N}c^N(\phi^N) = \sum_{k=0}^{\ell^N-1} \langle \Upsilon(F_{t^N(k)}^N, (\phi_k^N)), [\phi_{k+1}^N - \phi_k^N]_+ \rangle.$$

Because path ϕ^N is monotone, the second term in the inner product telescopes:

$$(31) \quad [\phi_b^N - \phi_a^N]_+ = \sum_{k=a}^{b-1} [\phi_{k+1}^N - \phi_k^N]_+$$

This property allows us to approximate (30) by a sum with only $O(\sqrt{N})$ summands, each of which corresponds to $O(\sqrt{N})$ terms in the original expression. This sum can be approximated in turn by replacing values of ϕ^N with values of ϕ . Doing so yields a Riemann-Stieltjes sum (cf. (56)) whose integrator ϕ is monotone. Since there are $O(\sqrt{N})$ rather than $O(N)$ summands, the $O(\frac{1}{N})$ bound in (29) ensures that replacing ϕ^N with ϕ leads to an approximation of $\frac{1}{N}c^N(\phi^N)$ that is asymptotically correct.²³ But since the number of summands still grows without bound in N , the Riemann-Stieltjes sums converge as N grows large; their limit is the integral that defines $c(\phi)$.

By Assumption 2, there is a full speed, piecewise monotone optimal path $\phi \in \Phi(K, \Xi)$ for problem (26). By Propositions 5.3 and 5.4, there are monotone discrete approximations of each monotone segment of ϕ with total cost close to $c(\phi)$. To construct a path $\phi^N \in \Phi(K^N, \Xi^N)$, we patch together these monotone discrete approximations, and also add segments from K^N to $\phi_0 \in K$ and from $\phi_T \in \Xi$ to Ξ^N . As before, conditions (16) and (25) ensure that this can be done at negligible cost. This argument yields the following upper bound, whose proof is presented in Appendix A.6.

Proposition 5.5. $\limsup_{N \rightarrow \infty} \frac{1}{N}C^N(K^N, \Xi^N) \leq C(K, \Xi)$.

Together, Propositions 5.2 and 5.5 establish the convergence of minimal path costs.

Theorem 5.6 (Convergence theorem). $\lim_{N \rightarrow \infty} \frac{1}{N}C^N(K^N, \Xi^N) = C(K, \Xi)$.

6. Consequences

We now use the convergence theorem to characterize exit times, stationary distribution asymptotics, and stochastic stability in the small noise double limit. These characteriza-

²³One can make an equivalent point in terms of derivatives: while $\dot{\phi}^N$ does not converge to $\dot{\phi}$, the local averages of $\dot{\phi}^N$ over time intervals of length $O(\sqrt{N})$ converge to the corresponding local averages of $\dot{\phi}$. Compare Figure 2.

tions are stated in terms of solutions to the continuous control problems (26). As these problems are tractable in certain interesting cases, the results here allow one to obtain explicit descriptions of the long run behavior of the stochastic evolutionary processes.

6.1 Expected exit times and exit locations

Given a recurrent class $K^N \in \mathcal{K}^N$, equation (13) defined \tilde{K}^N as the union of the recurrent classes in \mathcal{K}^N other than K^N . Thus if the process $\mathbf{X}^{N,\eta}$ starts in K^N , then $\mathbb{E}\tau^{N,\eta}(\tilde{K}^N)$ is the expected time until it reaches another recurrent class.

To characterize this expected waiting time, we let K be the limiting recurrent class corresponding to K^N (cf. equation (16)), and define

$$\tilde{K} = \bigcup_{L \in \mathcal{K} \setminus \{K\}} L$$

to be the union of the limiting recurrent classes other than K . Combining Proposition 3.3(i) and Theorem 5.6 immediately yields the following result.

Corollary 6.1. *Let $X_0^{N,\eta} = x^N \in K^N$ for all $\eta > 0$ and $N \geq N_0$. Then*

$$\lim_{N \rightarrow \infty} \lim_{\eta \rightarrow 0} \frac{\eta}{N} \log \mathbb{E}\tau^{N,\eta}(\tilde{K}^N) = C(K, \tilde{K}).$$

In words, Corollary 6.1 says that when N is sufficiently large, the exponential growth rate of the expected waiting time $\mathbb{E}\tau^{N,\eta}(\tilde{K}^N)$ as η^{-1} vanishes is approximately $N C(K, \tilde{K})$. This quantity can be evaluated explicitly by solving control problem (26).

Turning to exit locations, Proposition 3.3(ii) showed that in the small noise limit, the exit point of $\mathbf{X}^{N,\eta}$ from the strong basin of attraction $\mathcal{S}^N(K^N) = \mathcal{X}^N \setminus \mathcal{W}^N(\tilde{K}^N)$ is very likely to be the terminal state of a minimum cost path from K^N to $\mathcal{W}^N(\tilde{K}^N)$. Although the statements of the main results in Section 5 focus on costs, their proofs establish that optimal discrete paths can be approximated arbitrarily well by nearly optimal continuous paths, and vice versa. It follows that the likely exit points of $\mathbf{X}^{N,\eta}$ from $\mathcal{S}^N(K^N)$ can be approximated by the terminal points of the optimal solutions of the appropriate control problems (26).

6.2 Stationary distribution asymptotics and stochastic stability

6.2.1 The small noise limit

To state our results on stationary distribution asymptotics and stochastic stability in the small noise double limit, we first review the well-known results for the small noise

limit alluded to in Section 3.4. The analysis, which follows Freidlin and Wentzell (1998), is cast in terms of graphs on the set of recurrent classes \mathcal{K}^N .

A *tree* on \mathcal{K}^N with *root* K^N , sometimes called a K^N -*tree*, is a directed graph on \mathcal{K}^N with no outgoing edges from K^N , exactly one outgoing edge from each $L^N \neq K^N$, and a unique path through \mathcal{K}^N from each $L^N \neq K^N$ to K^N . Denote a typical K^N -tree by τ_{K^N} , and let \mathcal{T}_{K^N} denote the set of K^N -trees. The *cost of tree* τ_{K^N} on \mathcal{K}^N is the sum of the costs of the transitions over its edges:

$$(32) \quad C^N(\tau_{K^N}) = \sum_{(L^N, \hat{L}^N) \in \tau_{K^N}} C^N(L^N, \hat{L}^N).$$

Let $R^N: \mathcal{K}^N \rightarrow \mathbb{R}_+$ assign each recurrent class $K^N \in \mathcal{K}^N$ the minimal cost of a K^N -tree:

$$R^N(K^N) = \min_{\tau_{K^N} \in \mathcal{T}_{K^N}} C^N(\tau_{K^N}).$$

Then define the function $r^N: \mathcal{X}^N \rightarrow \mathbb{R}_+$ by

$$(33) \quad r^N(x) = \min_{K^N \in \mathcal{K}^N} (R^N(K^N) + C^N(K^N, \{x\})).$$

If x is in recurrent class K^N , then $r^N(x) = R^N(K^N)$. Otherwise, $r^N(x)$ is the sum of the cost of some K^N -tree and the cost of a path from K^N to x .²⁴ Finally, let $\Delta r^N: \mathcal{X}^N \rightarrow \mathbb{R}$ be a version of r^N whose values have been shifted to have minimum 0:

$$\Delta r^N(x) = r^N(x) - \min_{y \in \mathcal{X}^N} r^N(y).$$

Proposition 6.2 shows that the function Δr^N describes the exponential rates of decay of the stationary distribution weights $\mu^{N,\eta}(x)$ in the small noise limit. It is an easy consequence of Proposition 4.1 of Catoni (1999).

Proposition 6.2. *The stationary distributions $\mu^{N,\eta}$ satisfy*

$$(14) \quad -\lim_{\eta \rightarrow 0} \eta \log \mu^{N,\eta}(x) = \Delta r^N(x) \text{ for all } x \in \mathcal{X}^N.$$

6.2.2 The small noise double limit

To describe the asymptotics of the stationary distribution in the small noise double limit, we repeat the construction above using the set of limit recurrent classes \mathcal{K} and the

²⁴State x need not be in the weak basin of the recurrent class that yields the minimum in (33).

limit costs C . Denote a typical K -tree on the set of limiting recurrent classes \mathcal{K} by τ_K , and let \mathcal{T}_K denote the set of K -trees. Define the cost of tree τ_K by

$$C(\tau_K) = \sum_{(L, \hat{L}) \in \tau_K} C(L, \hat{L}).$$

Then define the functions $R: \mathcal{K} \rightarrow \mathbb{R}_+$, $r: X \rightarrow \mathbb{R}_+$, and $\Delta r: X \rightarrow \mathbb{R}$ by

$$R(K) = \min_{\tau_K \in \mathcal{T}_K} C(\tau_K), \quad r(x) = \min_{K \in \mathcal{K}} (R(K) + C(K, \{x\})), \quad \text{and} \quad \Delta r(x) = r(x) - \min_{y \in X} r(y).$$

Theorem 6.3 describes the asymptotics of the stationary distributions $\mu^{N,\eta}$ in the small noise double limit.

Theorem 6.3. *The stationary distributions $\mu^{N,\eta}$ satisfy*

$$\lim_{N \rightarrow \infty} \limmax_{\eta \rightarrow 0} \max_{x \in \mathcal{X}^N} \left| -\frac{\eta}{N} \log \mu^{N,\eta}(x) - \Delta r(x) \right| = 0.$$

In words, the theorem says that when N is sufficiently large, the exponential rate of decay of $\mu^{N,\eta}(x)$ as η^{-1} approaches infinity is approximately $N\Delta r(x)$.

A weaker version of Theorem 6.3, one that did not require uniformity of the large N limit in x , would follow directly from Theorem 5.6 and Proposition 6.2.²⁵ In order to prove Theorem 6.3 as stated, we need to show that the limit in Theorem 5.6 is uniform over all choices of the target set $\Xi \in \mathcal{K} \cup \{\{x\}\}_{x \in X}$ (cf. Assumption 2). We accomplish this in Appendix A.7 by bounding the rate of convergence in the results from Section 5 independently of the specific paths and target sets under consideration. This uniform convergence in these earlier results directly yields the uniform asymptotics for the stationary distributions.

In view of Theorem 6.3, we call state $x \in X$ *stochastically stable in the small noise double limit* if for any open set $O \subset X$ containing x , probability mass $\mu^{N,\eta}(O)$ does not vanish at an exponential rate in η once N is large enough.²⁶ Theorem 6.3 implies that state x is stochastically stable in the small noise double limit if and only if $\Delta r(x) = 0$.²⁷

²⁵In fact, since the number of recurrent classes is finite, a version of the theorem that focused only on these would also follow directly from Theorem 5.6.

²⁶Logically: $\forall \delta > 0 \forall O \in \mathcal{O}(X, x) \exists N_0 \in \mathbb{N} \forall N > N_0 \exists \eta_0 > 0 \forall \eta < \eta_0 \mu^{N,\eta}(O) > \exp(-\eta\delta)$, where $\mathcal{O}(X, x)$ denotes the set of open subsets of X containing x .

²⁷This characterization remains true under a more demanding definition of stochastic stability, requiring that for every $\delta > 0$, there exist an $O \in \mathcal{O}(X, x)$ such that (leaving the quantifiers on N and η in place) $\mu^{N,\eta}(y) > \exp(-\eta\delta)$ for every $y \in O \cap \mathcal{X}^N$.

7. Discussion

7.1 Solving the optimal control problems

It is well known that for a fixed population size N , as the noise level η approaches zero, the long run behavior of the process $\mathbf{X}^{N,\eta}$ can be described in terms of solutions to path cost minimization problems on the discrete set \mathcal{X}^N . The results in this paper show that after taking a second limit in the population size, the long run behavior of $\mathbf{X}^{N,\eta}$ can be described in terms of solutions of continuous optimal control problems on the simplex X . The value of these results lies in the fact that while the discrete optimization problems are computationally intensive once N is not small, the continuous problems can be approached analytically.

This is the topic of our companion paper, Sandholm and Staudigl (2014), which applies the results developed here to study the evolution of play under the logit choice rule in certain three-strategy coordination games. Using dynamic programming methods developed by Boltyanskii (1966) and Piccoli and Susmann (2000), we explicitly compute both the optimal exit paths from the basins of each pure Nash equilibrium and the optimal transition paths between these equilibria. Combined with the results here, these analyses of the optimal control problems provide us with full descriptions of the long run behavior of the stochastic evolutionary process. While the companion paper focuses on a particular class of games and choice rule, the properties of the optimal control problems that make them amenable to analysis hold more generally.

7.2 The best response with mutations model

Most work on long run behavior in stochastic evolutionary models has focused on the BRM protocol of Kandori et al. (1993), under which every suboptimal choice has unlikelihood 1. Because the unlikelihood function is an indicator function for suboptimal actions, the cost of a path in the small noise limit equals the number of suboptimal choices it entails. Taking the second limit in the population size, and considering any path along which all switches are to suboptimal strategies, this number of mutations is proportional to the path's ℓ^1 length. In many interesting examples, the resulting least-cost paths can be found by direct inspection: see especially Kandori and Rob (1995, 1998).

Since the BRM protocol has a discontinuous unlikelihood function, it falls outside the direct purview of our analysis.²⁸ The simplicity of the unlikelihood function suggests

²⁸Evaluating the cost of exiting the best response region of a strict equilibrium is an exception to this rule, since within a best response region the BRM protocol's unlikelihood function is constant.

that analogues of our results could be established for this case. Even so, the discontinuity of unlikelihood means that the relevant optimal control problems have discontinuous running costs, making them less amenable to analysis. Thus in the case of the BRM model, there is limited advantage in applying the analytical approach developed here.

7.3 Orders of limits

This paper investigates long run behavior in stochastic evolutionary models in the small noise double limit, taking η to zero before taking N to infinity. This order of limits emphasizes the consequences of the rareness of mistakes for long run play.

Following work by Binmore and Samuelson (1997) and Sandholm (2010a) on the two-strategy case, one can instead investigate the averaging effects of large population sizes on long run play by focusing on the large population limit, either by itself, or followed by the small noise limit. With just two strategies, birth-death chain methods can be used to carry this analysis to its completion. To obtain results in more general environments, one needs to use more sophisticated tools from the theory of sample path large deviations. For recent progress in this direction, see Benaïm et al. (2014).

It is natural to ask whether the conclusions about long run play are independent of the order in which the limits in η and N are taken, so that the force driving the large deviations analysis does not change the form our predictions takes. In the case of two-strategy games, for which birth-death chain methods are available, the effects of orders of limits on the limiting stationary distribution and stochastic stability are well understood. In the case of imitative dynamics with mutations, Binmore and Samuelson (1997) show that reversing the order of limits can alter the set of stochastically stable states in Hawk-Dove games, although Sandholm (2012) shows that this dependence can be eliminated by vanishingly small changes in the specification of the model. For noisy best response rules, Sandholm (2010a) shows that the asymptotic behavior of the stationary distributions, and hence the identity of the stochastically stable states, is the same for both orders of limits. Whether these conclusions extend to games with more than two strategies is an intriguing open question.

A. Appendix

A.1 Statement and proof of Lemma A.1

The analysis of Example 3.2 requires the following lemma.

Lemma A.1. Let F^N be a finite-population game defined by random matching in normal form coordination game A . Let $x \in X_j^N$ satisfy $x_j > 0$ and $j \notin b_j^N(x)$. Then there is a solution to (DBR) that begins at x and reaches a state at which j is unused in Nx_j steps.

Proof. We construct a solution to (DBR) as follows. The initial state is $x^0 = x$. We choose $i^1 \in b_j^N(x^0)$ to be a best response for a j player at this state, and then advance in increments $\frac{1}{N}(e_{i^1} - e_j)$ until reaching a state $x^1 = x^0 + d^1(e_{i^1} - e_j)$ where either j is unused or $i^1 \notin b_j^N(x^0)$. In the latter case, we choose $i^2 \in b_j^N(x^1)$ and continue the procedure until reaching a state x^c at which j is unused.

To prove the lemma, it is enough to show that upon reaching state x^c , $c < C$, the best response $i^{c+1} \in b_j^N(x^c)$ cannot be j itself. To do so, recall from definition (3) that $i^c \in b_j^N(x^{c-1})$ means that $F_{j \rightarrow i^c}^N(x^c) \geq F_{j \rightarrow k}^N(x^{c-1})$ for all $k \in S$, or equivalently, by (1) and (2),

$$(34) \quad \frac{N}{N-1}(e_{i^c} - e_k)'Ax^{c-1} - \frac{1}{N-1}(e_{i^c} - e_k)'Ae_j \geq 0 \text{ for all } k \in S.$$

By construction,

$$(35) \quad x^c = x^{c-1} + d^c(e_{i^c} - e_j) \text{ for some } d^c > 0,$$

Since $i^{c+1} \in b_j^N(x^c)$,

$$(36) \quad \frac{N}{N-1}(e_{i^{c+1}} - e_k)'Ax^c - \frac{1}{N-1}(e_{i^{c+1}} - e_k)'Ae_j \geq 0 \text{ for all } k \in S,$$

and since $i^c \notin b_j^N(x^c)$, the inequality in (36) is strict when $k = i^c$. Combining (34) (with $k = i^{c+1}$) and the strict version of (36) (with $k = i^c$) with (35) yields

$$(37) \quad (e_{i^c} - e_{i^{c+1}})'A(e_{i^c} - e_j) < 0.$$

Since A is a coordination game, we conclude that $i^{c+1} \neq j$, as we aimed to show. ■

A few additional steps show that the sequence of best responses $\{i^1, \dots, i^C\}$ is nonrepeating, and hence that $C < n$. Suppose to the contrary that two elements of the sequence are the same; for definiteness, let $i^1 = i^C$. Then

$$(38) \quad \sum_{c=1}^{C-1} (e_{i^c} - e_{i^{c+1}})'Ae_j = 0.$$

Summing (37) over $c \in \{1, \dots, C-1\}$ and substituting (38) yields

$$\sum_{c=1}^{C-1} (e_{i^c} - e_{i^{c+1}})' A e_{i^c} < 0,$$

again contradicting that A is a coordination game.

A.2 Proof of Proposition 5.1

Fix $\varepsilon > 0$. Since F and Υ are continuous, Υ is uniformly continuous on $F(X)$, so we can choose $\delta > 0$ so that

$$(39) \quad |\pi - \hat{\pi}| < \delta \text{ implies that } |\Upsilon_j(\pi) - \Upsilon_j(\hat{\pi})| < \varepsilon \text{ for all } \pi, \hat{\pi} \in F(X) \text{ and } j \in S.$$

Moreover, since $\{F^N\}$ is uniformly convergent, F uniformly continuous, and each $\phi^{(N)}$ is Lipschitz continuous with Lipschitz constant 2 (by (21)), we can choose N_0 so that

$$(40) \quad N \geq N_0 \text{ implies that } |F_{i \rightarrow \cdot}^N(x) - F(x)| < \delta \text{ for all } x \in X_i^N \text{ and } i \in S, \text{ and}$$

$$(41) \quad N \geq N_0 \text{ implies that } |F(\phi_t^{(N)}) - F(\phi_s^{(N)})| < \delta \text{ whenever } |t - s| \leq \frac{1}{N}.$$

It follows that for $N \geq N_0$, there exist α^N and β^N with $|\alpha^N| < \varepsilon$ and $|\beta^N| < \varepsilon$ such that

$$\begin{aligned} \frac{1}{N} c^N(\phi^{(N)}) &= \frac{1}{N} \sum_{k=0}^{\ell^N-1} \langle \Upsilon(F_{i^N(k) \rightarrow \cdot}^N(\phi_k^{(N)})), [\dot{\phi}_k^N]_+ \rangle \\ &= \frac{1}{N} \sum_{k=0}^{\ell^N-1} \langle \Upsilon(F(\phi_{\frac{k}{N}}^{(N)})), [\dot{\phi}_k^N]_+ \rangle + \alpha^N \frac{\ell^N}{N} \\ &= \int_0^{T^N} \langle \Upsilon(F(\phi_{\frac{[Nt]}{N}}^{(N)})), [\dot{\phi}_t^{(N)}]_+ \rangle dt + \alpha^N T^N \\ &= \int_0^{T^N} \langle \Upsilon(F(\phi_t^{(N)})), [\dot{\phi}_t^{(N)}]_+ \rangle dt + (\alpha^N + \beta^N) T^N \\ (42) \quad &= c(\phi^{(N)}) + (\alpha^N + \beta^N) T^N. \end{aligned}$$

The first equality is (23), the second follows from (39), (40), and (19), the third from (19) and (20), and the fourth from (39), (41), and (21). Since $\varepsilon > 0$ was chosen arbitrarily and the T^N are bounded, the proposition follows. ■

A.3 Proof of Proposition 5.2

By Assumption 1, there are paths $\phi^N = \{\phi_k^N\}_{k=0}^{\ell^N} \in \Phi^N(K^N, \Xi^N)$ of durations $T^N = \ell^N/N < \bar{T} < \infty$ that are optimal in problem (12), so that $C^N(K^N, \Xi^N) = c^N(\phi^N)$. Let C^* be the liminf of $\frac{1}{N}c^N(\phi^N)$. There is a subsequence along which $\frac{1}{N}c^N(\phi^N)$ converges to C^* , which we take without loss of generality to be the entire sequence.

For each ϕ^N , we construct a corresponding continuous path $\phi^{[N]} \in \Phi(K, \Xi)$ by concatenating three subpaths: a subpath $\phi^{[N],0}$ from a point in K_i to ϕ_0^N , the linear interpolation $\phi^{(N)}$ defined in (19), which leads from ϕ_0^N to $\phi_{\ell^N}^N$, and a subpath $\phi^{[N],1}$ from $\phi_{\ell^N}^N$ to a point in K_j .

To construct $\phi^{[N],0}$, recall from condition (16) that since $\phi_0^N \in K_i^N$, there is an $x_0^{[N]} \in K_i$ such that $|\phi_0^N - x_0^{[N]}| \leq \frac{d}{N}$. Define $\{\phi_t^{[N],0}\}_{t \in [0,1]}$ by $\phi_t^{[N],0} = (1-t)x_0^{[N]} + t\phi_0^N$. Then letting $b < \infty$ be the maximum of the continuous function $\Upsilon \circ F$ on the compact set X , we have that

$$c(\phi^{[N],0}) = \int_0^1 \langle \Upsilon(F(\phi_t^{[N],0})), [\dot{\phi}_t^{[N],0}]_+ \rangle dt \leq \frac{bd}{N}.$$

Subpath $\phi^{[N],1}$ is constructed analogously and satisfies the same bound.

Now fix $\varepsilon > 0$. The previous argument and equation (42) imply that for all N large enough, we have

$$(43) \quad c(\phi^{[N]}) \leq c(\phi^{(N)}) + \frac{2bd}{N} \leq \frac{1}{N}c^N(\phi^N) + 2\varepsilon\bar{T} + \frac{2bd}{N}.$$

Since ε was arbitrary, we conclude that $\lim_{N \rightarrow \infty} c(\phi^{[N]}) \leq \lim_{N \rightarrow \infty} \frac{1}{N}c^N(\phi^N) = C^*$, and hence that $C(K, \Xi) \leq C^*$. ■

A.4 Proof of Proposition 5.3

Fix N , and write $n_+ = \#S_+$ and $n_- = \#S_-$. To prove the proposition, we construct for all N large enough a monotone path ϕ^N that satisfies

$$\max_k \sum_{j \in S_+} \left| \phi_{k,j}^N - \phi_{s^N + \frac{k}{N}, j} \right| \leq \frac{2n_+}{N} \quad \text{and} \quad \max_k \sum_{i \in S_-} \left| \phi_{k,i}^N - \phi_{s^N + \frac{k}{N}, i} \right| \leq \frac{2n_-}{N}.$$

Summing these inequalities yields inequality (29).

Because ϕ is monotone and moves at full speed, there is a time $s^N \in [0, \frac{1}{N})$ at which

$$(44) \quad \sum_{j \in S_+} \phi_{s^N, j} \in \frac{1}{N} \mathbb{Z}, \text{ and hence } \sum_{i \in S_-} \phi_{s^N, i} \in \frac{1}{N} \mathbb{Z}.$$

This is the s^N introduced in the statement of the theorem. To minimize notation in what follows we will take s^N to be 0.

The point in the simplex X that maximizes the minimum ℓ^1 distance from all vertices is the barycenter, which is ℓ^1 distance $\frac{2(n-1)}{n}$ from every vertex. This and (44) imply that there is a $\phi_0^N \in \mathcal{X}^N$ such that

$$(45) \quad \sum_{j \in S_+} \phi_{0, j} = \sum_{j \in S_+} \phi_{0, j}^N, \quad \sum_{i \in S_-} \phi_{0, i} = \sum_{i \in S_-} \phi_{0, i}^N, \text{ and}$$

$$\sum_{i \in S} [\phi_{0, i} - \phi_{0, i}^N]_+ = \sum_{i \in S} [\phi_{0, i} - \phi_{0, i}^N]_- = \frac{n-1}{n} \cdot \frac{1}{N}.$$

This is the base of our inductive argument. The inductive step says that if

$$\sum_{j \in S_+} \left| \phi_{\frac{k}{N}, j} - \phi_{k, j}^N \right| \leq \frac{2n_+}{N} \quad \text{and} \quad \sum_{i \in S_-} \left| \phi_{\frac{k}{N}, i} - \phi_{k, i}^N \right| \leq \frac{2n_-}{N},$$

then we can choose $\phi_{\frac{k+1}{N}}^N = \phi_k^N + \frac{1}{N}(e_j - e_i)$ with $j \in S_+$ and $i \in S_-$ so that

$$\sum_{j \in S_+} \left| \phi_{\frac{k+1}{N}, j} - \phi_{k+1, j}^N \right| \leq \frac{2n_+}{N}, \quad \text{and} \quad \sum_{i \in S_-} \left| \phi_{\frac{k+1}{N}, i} - \phi_{k+1, i}^N \right| \leq \frac{2n_-}{N}.$$

Note that this procedure ensures that ϕ^N is also monotone, with the same partition $S = S_+ \cup S_-$ as ϕ .

Let $x = \phi_{\frac{k}{N}}$, $y = \phi_{\frac{k+1}{N}}$, and $\chi = \phi_k^N$, be given, with $y = \phi_{k+1}^N$ to be determined. Our analysis focuses on strategies in S_+ . Monotonicity and the fact that ϕ moves at full speed imply that

$$(46) \quad y - x = z \text{ for some } z \text{ with } z_j \geq 0 \text{ for } j \in S_+ \text{ and } \sum_{j \in S_+} z_j = \frac{1}{N}.$$

This and the fact that $y = \chi + \frac{1}{N}(e_j - e_i)$ for some $j \in S_+$ and $i \in S_-$ imply that

$$(47) \quad \sum_{j \in S_+} [y_j - \chi_j]_+ \leq \frac{1}{N} + \sum_{j \in S_+} [x_j - \chi_j]_+.$$

Furthermore,

Lemma A.2. *If*

$$(48) \quad \sum_{j \in S_+} [x_j - \chi_j]_+ \geq \frac{1}{N}(n_+ - 1),$$

then y can be chosen so that $\sum_{j \in S_+} [y_j - y_j]_+ \leq \sum_{j \in S_+} [x_j - \chi_j]_+$.

To prove the lemma, recall from (45) that ϕ_0 and ϕ_0^N place equal total mass on strategies in S_+ . Thus, since ϕ and ϕ^N move at full speed and are monotone with respect to the same partition $S = S_+ \cup S_-$, it follows that this equality is maintained at all corresponding points on paths ϕ and ϕ^N . In particular, we have

$$(49) \quad 0 = \sum_{j \in S_+} (x_j - \chi_j) = \sum_{j \in S_+} [x_j - \chi_j]_+ - \sum_{j \in S_+} [x_j - \chi_j]_-.$$

Thus, there are at most $n_+ - 1$ strategies $j \in S_+$ for which $x_j - \chi_j > 0$. Consequently, (48) implies that there is at least one strategy $j^* \in S_+$ with

$$(50) \quad x_{j^*} - \chi_{j^*} \geq \frac{1}{N}.$$

If we choose $y = \chi + \frac{1}{N}(e_{j^*} - e_i)$ (for any $i \in S_-$), then applying (50) and (46) shows that

$$\begin{aligned} \sum_{j \in S_+} [y_j - y_j]_+ &= [x_{j^*} + z_{j^*} - (\chi_{j^*} + \frac{1}{N})]_+ + \sum_{j \in S_+ \setminus \{j^*\}} [x_j + z_j - \chi_j]_+ \\ &\leq x_{j^*} + z_{j^*} - (\chi_{j^*} + \frac{1}{N}) + \sum_{j \in S_+ \setminus \{j^*\}} ([x_j - \chi_j]_+ + z_j) \\ &= \sum_{j \in S_+} ([x_j - \chi_j]_+ + z_j) - \frac{1}{N} \\ &= \sum_{j \in S_+} [x_j - \chi_j]_+, \end{aligned}$$

proving the lemma.

Together, equation (47) and Lemma A.2 imply that

$$\sum_{j \in S_+} [y_j - y_j]_+ \leq \max \left\{ \frac{n_+}{N}, \sum_{j \in S_+} [x_j - \chi_j]_+ \right\}.$$

Applying (49) to both sides of this inequality, and then summing the resulting inequality and the original one yields

$$\sum_{j \in S_+} |y_j - y_j| \leq \max \left\{ \frac{2n_+}{N}, \sum_{j \in S_+} |x_j - x_j| \right\}.$$

This establishes the part of the inductive step concerning S_+ ; a similar argument proves the part concerning S_- . This completes the proof of the proposition. ■

A.5 Proof of Proposition 5.4

Fix $\varepsilon > 0$. Choose $\delta > 0$ so that (39) holds, and then choose N_0 so that (40) and

$$(51) \quad N \geq N_0 \text{ implies that } |F(\phi_{\frac{k}{N}}) - F(\phi_k^N)| < \delta \text{ for all } k \in \{0, \dots, \lfloor NT \rfloor\}$$

hold; the latter is possible because F is uniformly continuous and because ϕ^N converges uniformly to ϕ , as described in (29); as in the proof of Proposition 5.3, we minimize notation by taking s^N to equal 0.

By the triangle inequality,

$$(52) \quad \left| \Upsilon_j(F_{i \rightarrow \cdot}^N(\phi_k^N)) - \Upsilon_j(F(\phi_{\frac{k}{N}})) \right| \leq \left| \Upsilon_j(F_{i \rightarrow \cdot}^N(\phi_k^N)) - \Upsilon_j(F(\phi_k^N)) \right| + \left| \Upsilon_j(F(\phi_k^N)) - \Upsilon_j(F(\phi_{\frac{k}{N}})) \right|.$$

Thus if $N \geq N_0$, there exists α^N with $|\alpha^N| < \varepsilon$ such that

$$(53) \quad \begin{aligned} \frac{1}{N} c^N(\phi^N) &= \sum_{k=0}^{\ell^N-1} \langle \Upsilon(F_{i \rightarrow \cdot}^N(\phi_k^N)), [\phi_{k+1}^N - \phi_k^N]_+ \rangle \\ &= \sum_{k=0}^{\ell^N-1} \langle \Upsilon(F(\phi_{\frac{k}{N}})), [\phi_{k+1}^N - \phi_k^N]_+ \rangle + 2\alpha^N T^N. \end{aligned}$$

The first equality here is (30), and the second follows from (39), (40), (51), and (52).

Now let $L^N = \lfloor \sqrt{N} \rfloor$ and let $M^N = \lfloor NT \rfloor / L^N$, so that

$$(54) \quad \lim_{N \rightarrow \infty} M^N = \infty \text{ and } \lim_{N \rightarrow \infty} \frac{M^N}{N} = 0.$$

Also, choose $\tau > 0$ so that

$$(55) \quad |t - s| \leq 2\tau \text{ implies that } |F(\phi_t) - F(\phi_s)| < \delta.$$

Then continuing from (53), considering $N \geq N_0$ large enough that $L^N/N \leq \tau$, and taking M^N to be an integer for notational convenience only, there exist β^N with $|\beta^N| < \varepsilon$ and a constant $b > 0$ whose value depends on the maximum of $\Upsilon \circ F$ on X such that

$$\begin{aligned}
\frac{1}{N}c^N(\phi^N) &= \sum_{m=0}^{M^N-1} \sum_{k=0}^{L^N-1} \langle \Upsilon(F(\phi_{mL^N})), [\phi_{mL^N+k+1}^N - \phi_{mL^N+k}^N]_+ \rangle + (2\alpha^N + \beta^N)T^N \\
&= \sum_{m=0}^{M^N-1} \langle \Upsilon(F(\phi_{mL^N})), [\phi_{(m+1)L^N}^N - \phi_{mL^N}^N]_+ \rangle + (2\alpha^N + \beta^N)T^N \\
(56) \quad &= \sum_{m=0}^{M^N-1} \langle \Upsilon(F(\phi_{mL^N})), [\phi_{\frac{(m+1)L^N}{N}} - \phi_{\frac{mL^N}{N}}]_+ \rangle + (2\alpha^N + \beta^N)T^N + \frac{bT^N}{\sqrt{N}}.
\end{aligned}$$

The first equality uses (55) and (39), the second uses the monotonicity of ϕ^N , and the third uses the boundedness of $\Upsilon \circ F$ on X and the $O(\frac{1}{N})$ convergence of ϕ^N to ϕ specified in (29).

The limits in (54) and the monotonicity of ϕ imply that as N approaches infinity, and the Riemann-Stieltjes sum in (56) converges to a Riemann integral. (To be more precise, writing the inner product in the initial term of (56) as a sum and then reversing the order of summation yields a sum of n Riemann-Stieltjes sums, which converges to a sum of n Riemann integrals.) Accounting explicitly for the approximation error, there exist γ^N with $|\gamma^N| < \varepsilon$ such that for large enough N ,

$$\begin{aligned}
\frac{1}{N}c^N(\phi^N) &= \int_0^T \langle \Upsilon(F(\phi_t)), [\dot{\phi}_t]_+ \rangle dt + (2\alpha^N + \beta^N + \gamma^N)T^N + \frac{bT^N}{\sqrt{N}} \\
(57) \quad &= c(\phi) + (2\alpha^N + \beta^N + \gamma^N)T^N + \frac{bT^N}{\sqrt{N}}.
\end{aligned}$$

Since $T^N \leq T$ (see the statement of Proposition 5.3), the last summand vanishes as N grows large. Thus since ε was arbitrary, we conclude that $\lim_{N \rightarrow \infty} \frac{1}{N}c^N(\phi^N) = c(\phi)$. ■

A.6 Proof of Proposition 5.5

By Assumption 2, there is a continuous, piecewise monotone path $\phi = \{\phi_t\}_{t \in [0, T]} \in \Phi(K, \Xi)$ with cost $c(\phi) = C(K, \Xi)$. As noted in Section 5.1, we can assume that path ϕ moves at full speed, as in (28) with the inequalities binding. Fix $\varepsilon > 0$. We will construct a sequence of discrete paths with $\phi^N \in \Phi^N(K^N, \Xi^N)$ whose normalized costs converge to the sum of $c(\phi)$ and terms that vanish with ε .

As ϕ is piecewise monotone, there is an $M < \infty$ and times $0 = T_0 < T_1 < \dots < T_M = T$

such that ϕ is monotone on each subinterval $[T_{m-1}, T_m]$. The discrete path ϕ^N is the concatenation of $2M + 1$ subpaths: $\psi^{N,0}, \phi^{N,1}, \psi^{N,1}, \phi^{N,2}, \dots, \phi^{N,M}, \psi^{N,M}$. For $m \in \{1, \dots, M\}$, subpath $\phi^{N,m}$ is the discrete approximation of $\phi|_{[T_{m-1}, T_m]}$ constructed in Proposition 5.3; the length of this subpath is $\ell^{N,m} = \lfloor N(T_m - T_{m-1} - s^{N,m}) \rfloor$, where $s^{N,m} \in [0, \frac{1}{N}]$ too is from Proposition 5.3.

For $m \in \{1, \dots, M-1\}$, subpath $\psi^{N,m}$ must begin at node $\phi_{\ell^{N,m}}^{N,m}$ and end at node $\phi_0^{N,m+1}$. We focus for notational convenience on $m = 1$, although the bound we establish next applies for all $m \in \{1, \dots, M-1\}$. Define $\hat{s}^{N,1}$ by $T_1 - \hat{s}^{N,1} = s^{N,1} + \frac{\ell^{N,1}}{N}$. Then $\hat{s}^{N,1} \in [0, \frac{1}{N})$, and we can bound the distance between $\phi_{\ell^{N,1}}^{N,1}$ and end at node $\phi_0^{N,2}$ as follows:

$$(58) \quad \left| \phi_{\ell^{N,1}}^{N,1} - \phi_0^{N,2} \right| \leq \left| \phi_{\ell^{N,1}}^{N,1} - \phi_{T_1 - \hat{s}^{N,1}} \right| + \left| \phi_{T_1 - \hat{s}^{N,1}} - \phi_{T_1 + s^{N,2}} \right| + \left| \phi_{T_1 + s^{N,2}} - \phi_0^{N,2} \right| \leq \frac{2n}{N} + \frac{4}{N} + \frac{2n}{N}.$$

The bounds on the first and third terms are from Proposition 5.3, and the bound on the second term follow from the fact that $\hat{s}^{N,1}$ and $s^{N,2}$ are less than $\frac{1}{N}$ and the speed bound (28) on ϕ .

The initial subpath $\psi^{N,0}$ begins at a state in K^N and ends at $\phi_0^{N,1}$, and the final subpath $\psi^{N,M}$ begins at $\phi_{\ell^{N,M}}^{N,M}$ and ends at a state in Ξ^N . Focusing for convenience on the former, note that since $\phi_0 \in K$, condition (16) ensures that we can choose $\phi_0^{N,0} = x^N \in K^N$ with $|\phi_0^{N,0} - \phi_0| \leq \frac{d}{N}$. We therefore have

$$(59) \quad \left| x^N - \phi_0^{N,1} \right| \leq \left| x^N - \phi_0 \right| + \left| \phi_0 - \phi_{s^{N,1}} \right| + \left| \phi_{s^{N,1}} - \phi_0^{N,1} \right| \leq \frac{d}{N} + \frac{2}{N} + \frac{2n}{N}.$$

The bound on the second term follows from the fact that $s^{N,1} \leq \frac{1}{N}$ and from the speed bound (28), and the bound on the third term follows from Proposition 5.3.

Observe that given any distinct $\chi, y \in \mathcal{X}^N$, there is a state $\hat{\chi}$ adjacent to χ such that $|\hat{\chi} - y| = |\chi - y| - \frac{2}{N}$. These observations and inequalities (58) and (59) imply that each subpath $\psi^{N,m}$, $m \in \{1, \dots, M-1\}$ can be constructed to have length no greater than $2n + 2$, and that subpaths $\psi^{N,0}$ and $\psi^{N,M}$ can each be constructed to have length no greater than $\frac{1}{2}(d + 2 + 2n)$.

As before, let $b < \infty$ be the maximum of the continuous function $\Upsilon \circ F$ on the compact set X . Since each $F_{i \rightarrow \cdot}^N$ converges uniformly to F , for all N large enough the maximum cost of a feasible step in the N th process is at most $2b$. This fact and the arguments from the previous paragraph show that for such N , the total cost of the subpaths $\psi^{N,m}$ satisfies

$$(60) \quad \sum_{m=0}^M c^N(\psi^{N,m}) \leq 2b \left((M-1)(2n+2) + (d+2+2n) \right) \leq 2b(3nM+d).$$

Since for each N the total duration of subpaths $\phi^{N,1}, \dots, \phi^{N,M}$ is less than T , inequalities (57) and (60) imply that for all N large enough,

$$(61) \quad \frac{1}{N}c^N(\phi^N) \leq c(\phi) + 4\varepsilon T + \frac{bT}{\sqrt{N}} + \frac{2b(3nM+d)}{N}.$$

Since ε was arbitrary, it follows that

$$\lim_{N \rightarrow \infty} \frac{1}{N}c^N(\phi^N) \leq c(\phi),$$

and thus that

$$\limsup_{N \rightarrow \infty} \frac{1}{N}C^N(K^N, \Xi^N) \leq C(K, \Xi). \quad \blacksquare$$

A.7 Proof of Theorem 6.3

Fix $\varepsilon > 0$. We need to show that for all large enough N ,

$$\lim_{\eta \rightarrow 0} \max_{x \in \mathcal{X}^N} \left| -\frac{\eta}{N} \log \mu^{N,\eta}(x) - \Delta r(x) \right| < \varepsilon.$$

By Proposition 6.2, it is enough to show that for all large enough N ,

$$(62) \quad \max_{x \in \mathcal{X}^N} \left| \frac{1}{N} \Delta r^N(x) - \Delta r(x) \right| < \varepsilon.$$

In fact, it is enough to show that for all large enough N ,

$$(63) \quad \max_{x \in \mathcal{X}^N} \left| \frac{1}{N} r^N(x) - r(x) \right| < \varepsilon,$$

since this uniform convergence of r^N to r implies that the minimum of r^N converges to the minimum of r , and together these imply (62).

Combining the definitions of r^N , R^N , and C^N yields

$$(64) \quad r^N(x) = \min_{K^N \in \mathcal{K}^N} \left(\min_{\tau_{K^N} \in \mathcal{T}_{K^N}} \sum_{(L^N, \hat{L}^N) \in \tau_{K^N}} C^N(L^N, \hat{L}^N) + C^N(K^N, \{x\}) \right),$$

and $r(x)$ can be expressed analogously. Now fix a population size N and a state $x \in \mathcal{X}^N$. For this fixed x , there are κ^2 transition costs that need to be found to evaluate (64): specifically, there are $\kappa^2 - \kappa$ terms of the form $C^N(L^N, \hat{L}^N)$, where (L^N, \hat{L}^N) is an ordered pair of distinct

recurrent classes, and there are κ terms of the form $C^N(K^N, \{x\})$. Since κ^2 is finite, the convergence of these costs guaranteed by Theorem 5.6 is uniform: there is an N_0 such that for all $N \geq N_0$ and all choices of recurrent classes,

$$(65) \quad \begin{aligned} & \left| \frac{1}{N} C^N(L^N, \hat{L}^N) - C(L, \hat{L}) \right| < \frac{\varepsilon}{\kappa} \quad \text{and} \\ & \left| \frac{1}{N} C^N(K^N, \{x\}) - C(K, \{x\}) \right| < \frac{\varepsilon}{\kappa} \quad \text{if } x \in \mathcal{X}^N. \end{aligned}$$

Thus $|\frac{1}{N} r^N(x) - r(x)| < \varepsilon$, and hence $\lim_{N \rightarrow \infty} \frac{1}{N} r^N(x) = r(x)$, where the limit is taken over N such that $x \in \mathcal{X}^N$.

In order to establish (63), we must show that the limit just obtained holds uniformly over x . By the previous logic, this would follow if we could show that convergence of $\frac{1}{N} C^N(K^N, \{x\})$ to $C(K, \{x\})$ were uniform in x . To see that this is so, note that by inequalities (43) and (61), the choice of N_0 needed to ensure inequality (65) for all $N \geq N_0$ can be determined as a function of the following constants: d (from condition (16)), \bar{T} (from Assumption 1), b (the maximum of $\Upsilon \circ F$ on X), \tilde{M} and \tilde{T} (since Assumption 2 requires that $M = M(x)$ and $T = T(x)$ from inequality (61) satisfy $M(x) \leq \tilde{M}$ and $T(x) \leq \tilde{T}$ for all x), and n (the number of strategies). Since none of these constants depend on x , we can indeed choose N_0 so that (65) holds for all $N \geq N_0$ and for all $x \in \mathcal{X}^N$ simultaneously. This establishes (63), and so completes the proof of the theorem. ■

References

- Alós-Ferrer, C. and Netzer, N. (2010). The logit response dynamics. *Games and Economic Behavior*, 68:413–427.
- Anderson, S. P., de Palma, A., and Thisse, J.-F. (1992). *Discrete Choice Theory of Product Differentiation*. MIT Press, Cambridge.
- Benaïm, M., Sandholm, W. H., and Staudigl, M. (2014). Large deviations and stochastic stability in the large population limit. Unpublished manuscript, Université de Neuchâtel, University of Wisconsin, and Bielefeld University.
- Binmore, K. and Samuelson, L. (1997). Muddling through: Noisy equilibrium selection. *Journal of Economic Theory*, 74:235–265.
- Binmore, K., Samuelson, L., and Vaughan, R. (1995). Musical chairs: Modeling noisy evolution. *Games and Economic Behavior*, 11:1–35. Erratum, 21 (1997), 325.

- Blume, L. E. (1993). The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5:387–424.
- Blume, L. E. (1997). Population games. In Arthur, W. B., Durlauf, S. N., and Lane, D. A., editors, *The Economy as an Evolving Complex System II*, pages 425–460. Addison-Wesley, Reading, MA.
- Blume, L. E. (2003). How noise matters. *Games and Economic Behavior*, 44:251–271.
- Boltyanskii, V. G. (1966). Sufficient conditions for optimality and the justification of the dynamic programming method. *SIAM Journal on Control*, 4:326–361.
- Catoni, O. (1999). Simulated annealing algorithms and Markov chains with rare transitions. In Azéma, J., Émery, M., Ledoux, M., and Yor, M., editors, *Séminaire de Probabilités XXXIII*, pages 69–119. Springer, Berlin.
- Cho, I.-K., Williams, N., and Sargent, T. J. (2002). Escaping Nash inflation. *Review of Economic Studies*, 69:1–40.
- Dokumacı E. and Sandholm, W. H. (2011). Large deviations and multinomial probit choice. *Journal of Economic Theory*, 146:2151–2158.
- Ellison, G. (2000). Basins of attraction, long run equilibria, and the speed of step-by-step evolution. *Review of Economic Studies*, 67:17–45.
- Foster, D. P. and Young, H. P. (1990). Stochastic evolutionary game dynamics. *Theoretical Population Biology*, 38:219–232. Corrigendum, 51 (1997), 77-78.
- Freidlin, M. I. and Wentzell, A. D. (1998). *Random Perturbations of Dynamical Systems*. Springer, New York, second edition.
- Fudenberg, D. and Imhof, L. A. (2006). Imitation processes with small mutations. *Journal of Economic Theory*, 131:251–262.
- Fudenberg, D. and Imhof, L. A. (2008). Monotone imitation dynamics in large populations. *Journal of Economic Theory*, 140:229–245.
- Gilboa, I. and Matsui, A. (1991). Social stability and equilibrium. *Econometrica*, 59:859–867.
- Hofbauer, J. and Sandholm, W. H. (2002). On the global convergence of stochastic fictitious play. *Econometrica*, 70:2265–2294.
- Kandori, M., Mailath, G. J., and Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica*, 61:29–56.
- Kandori, M. and Rob, R. (1995). Evolution of equilibria in the long run: A general theory and applications. *Journal of Economic Theory*, 65:383–414.

- Kandori, M. and Rob, R. (1998). Bandwagon effects and long run technology choice. *Games and Economic Behavior*, 22:84–120.
- Myatt, D. P. and Wallace, C. C. (2003). A multinomial probit model of stochastic evolution. *Journal of Economic Theory*, 113:286–301.
- Piccoli, B. and Sussmann, H. J. (2000). Regular synthesis and sufficiency conditions for optimality. *SIAM Journal on Control and Optimization*, 39:359–410.
- Roth, G. and Sandholm, W. H. (2013). Stochastic approximations with constant step size and differential inclusions. *SIAM Journal on Control and Optimization*, 51:525–555.
- Sandholm, W. H. (2007). Simple formulas for stationary distributions and stochastically stable states. *Games and Economic Behavior*, 59:154–162.
- Sandholm, W. H. (2010a). Orders of limits for stationary distributions, stochastic dominance, and stochastic stability. *Theoretical Economics*, 5:1–26.
- Sandholm, W. H. (2010b). *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge.
- Sandholm, W. H. (2012). Stochastic imitative game dynamics with committed agents. *Journal of Economic Theory*, 147:2056–2071.
- Sandholm, W. H. and Staudigl, M. (2014). Stochastic stability in the small noise double limit, II: The logit model. Unpublished manuscript, University of Wisconsin and Bielefeld University.
- Staudigl, M. (2012). Stochastic stability in asymmetric binary choice coordination games. *Games and Economic Behavior*, 75:372–401.
- Williams, N. (2009). Escape dynamics in learning models. Unpublished manuscript, University of Wisconsin.
- Young, H. P. (1993). The evolution of conventions. *Econometrica*, 61:57–84.
- Young, H. P. (1998). *Individual Strategy and Social Structure*. Princeton University Press, Princeton.
- Young, H. P. (2004). *Strategic Learning and Its Limits*. Oxford University Press, Oxford.