

Large deviations and stochastic stability in the small noise double limit

WILLIAM H. SANDHOLM

Department of Economics, University of Wisconsin

MATHIAS STAUDIGL

Department of Quantitative Economics, Maastricht University

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. We use these results to perform a complete analysis of evolution in three-strategy coordination games that satisfy the marginal bandwagon property and that have an interior equilibrium, with agents following the logit choice rule.

KEYWORDS. Evolutionary game theory, equilibrium breakdown, stochastic stability, large deviations.

JEL CLASSIFICATION. C72, C73.

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 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

William H. Sandholm: whs@ssc.wisc.edu

Mathias Staudigl: mathias.staudigl@gmail.com

This article merges two working papers: “Large deviations and stochastic stability in the small noise double limit, I: Theory” and “Large deviations and stochastic stability in the small noise double limit, II: The logit model.” We thank a co-editor and a number of anonymous referees for valuable comments, and Daniel Liberzon for helpful discussions. Financial support from NSF Grants SES-0851580, SES-1155135, and SES-1458992, U.S. Air Force OSR Grant FA9550-09-0538, and the Vienna Science and Technology Fund (WWTF) under project fund MA 09-017 is gratefully acknowledged.

Copyright © 2016 William H. Sandholm and Mathias Staudigl. Licensed under the [Creative Commons Attribution-NonCommercial License 3.0](https://creativecommons.org/licenses/by-nc/3.0/). Available at <http://econtheory.org>.

DOI: [10.3982/TE1905](https://doi.org/10.3982/TE1905)

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 second step assesses the likelihoods of escapes from and transitions among equilibria and other recurrent classes. Finally, the third step uses graph-theoretic methods to distill the analysis of transitions between equilibria into a characterization of the limiting stationary distribution of the process.³

The second step in this analysis, which describes how an established equilibrium is upset and which (if any) new equilibrium is likely to arise, seems itself to be of inherent interest. But to date, this question of equilibrium breakdown has not attracted much attention in the game theory literature.

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 unlikelihood of each such choice. As a consequence, general results on transitions between equilibria and stochastic stability are only available for two-strategy games.⁵

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

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 \(2010c, Chapters 11 and 12\)](#).

²See, for instance, [Young \(2004\)](#) and [Sandholm \(2010c\)](#).

³See the previous references, [Freidlin and Wentzell \(1998\)](#), or [Catoni \(1999\)](#).

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

⁵[Blume \(2003\)](#) and [Sandholm \(2007, 2010b\)](#) study stochastic stability in two-strategy games using birth-death chain methods. [Staudigl \(2012\)](#) studies the case of two-population random matching in 2×2 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 \(2010c, Section 11.5\)](#).

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. Finally, to illustrate the applicability of these characterizations, we use control-theoretic methods to provide a complete analysis of long-run behavior in a class of three-strategy coordination games. To our knowledge, this work is 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.

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

⁶See Section 2.2. As we discuss below, the continuity assumption rules out the BRM model, in which unlikelihood functions are indicator functions.

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. In so doing, we study the behavior of the stochastic evolutionary process in the small noise limit when the population size is large.

In Sections 4 and 5, we develop our central result, which shows that as N grows large, the discrete path cost-minimization problems described above converge to continuous optimal control problems on the simplex. In Section 6, we combine this convergence result with graph-theoretic techniques to characterize various aspects of long-run behavior in the small noise double limit—expected exit times, stationary distribution asymptotics, and stochastic stability—in terms of solutions to these continuous control problems.

The control problems appearing in these characterizations are multidimensional and nonsmooth. Thus to demonstrate the utility of our results, we must show that these problems are nevertheless tractable in interesting cases.

We do so in Section 7. Our analysis there focuses on evolution under the logit choice rule, and on three-strategy coordination games that satisfy the marginal bandwagon property (Kandori and Rob 1998) and that admit an interior equilibrium. This class of games, which we call *simple three-strategy coordination games*, is large enough to allow some variety in its analysis, but small enough that the analysis remains manageable.

We analyze the control problems associated with two distinct kinds of large deviations properties. We first consider the *exit problem*, which is used to assess the expected time until the evolutionary process leaves the basin of attraction of a stable equilibrium and to determine the likely exit path. Solving this problem for the class of games under consideration, we show that the likely exit path proceeds along the boundary of the simplex, escaping the basin of attraction through a boundary mixed equilibrium.

To evaluate stationary distribution asymptotics and stochastic stability, one must instead consider the *transition problem*, which is used to assess the probable time until a transition between a given pair of stable equilibria and to determine the most likely path that this transition will follow. We solve the transition problem explicitly for simple three-strategy coordination games. We find that the nature of the problem's solution depends in a basic way on whether the game in question is also a potential game. When this is so, the optimal control problem is degenerate, in that there are open sets of states from which there are a continuum of minimal cost paths. Still, the optimal paths between equilibria always proceed directly along the relevant edge of the simplex. The control problem is not degenerate for games without a potential function, which we call *skewed games*. But unlike in the case of potential games, optimal paths between equilibria of skewed games need not be direct; instead, they may proceed along an alternate edge of the simplex, turn into the interior, and pass through the interior equilibrium.

By combining solutions to the control problems with our earlier results, we are able to characterize the long-run behavior of the evolutionary process in the small noise double limit. We use a parameterized class of examples to illustrate the effects of payoff-dependent mistake probabilities on equilibrium selection, and to contrast long-run behavior in the logit and BRM models. In addition, in the class of potential games we consider, we fully describe the asymptotic behavior of the stationary distribution in the small noise double limit, showing that the rate of decay of the stationary distribution mass at each state equals the difference between the value of the potential function at that state and the maximum value of potential. In contrast to those in previous work on logit choice in potential games,⁷ the assumptions we impose on the transition law of the evolutionary process are asymptotic in nature, and so do not allow us to express the stationary distribution in closed form. We instead build our analysis on large deviations estimates, and thereby obtain a clearer intuition about the form that the stationary distribution asymptotics take.

While the optimal control problems we solve have nonsmooth running costs, they are simple in other respects. If $L(x, u)$ represents the cost of choosing direction of motion u at state x , then L is piecewise linear in u regardless of the agents' choice rule. When agents employ the logit choice rule, L is also piecewise linear in x . Taking advantage of these properties, we use sufficient conditions for optimality due to [Boltysanski \(1966\)](#) and [Piccoli and Sussmann \(2000\)](#) to construct candidate value functions, and to verify that they are indeed the value functions for our problems. These sufficient conditions require the value function to be continuous, to be continuously differentiable except on a finite union of manifolds of positive codimension, and to satisfy the Hamilton–Jacobi–Bellman equation wherever the value function is smooth. In our case, for each fixed state x , the piecewise linearity of $L(x, u)$ in u means that only a small number of controls need to be considered, while the piecewise linearity of $L(x, u)$ in x makes it enough to check the Hamilton–Jacobi–Bellman equation at a small number of well chosen states.

These properties of the optimal control problem are not dependent on the class of games we consider, but only on the linearity of payoffs in the population state. Moreover, much of the structure of the problem is retained under alternatives to the logit choice rule. Thus as we explain in the final section of the paper, it should be possible to use the approach developed here to study long-run behavior in broader classes of games and under other choice rules.

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 \(2014\)](#), and the references therein. The present paper addresses

⁷See [Blume \(1993, 1997\)](#) and [Sandholm \(2010c, Sections 11.5 and 12.2\)](#), as well as [Section 7.6](#) below.

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 8.1 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 contrast, in the noisy best response models studied here, the least costly transition between a pair of pure equilibria may pass through the interior of the simplex.

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. Blume (2003) and Sandholm (2010b) 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 2×2 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. 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 describes the consequences for escape from equilibrium, limiting stationary distributions, and stochastic stability. Section 7 combines the foregoing analysis with optimal control techniques to study long-run behavior in a class of coordination games under the logit choice rule. Section 8 offers concluding discussions. Many proofs and auxiliary results are presented in the Appendix. See Table 1 for a listing of all notation used in the paper.

2. THE MODEL

2.1 Finite-population games

We consider games in which agents from a population of size N choose strategies from a 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 $\mathcal{X}^N = X \cap (1/N)\mathbb{Z}^n = \{x \in X : Nx \in \mathbb{Z}^n\}$. 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

Notation	Meaning	Section
A	Symmetric normal form game	2.1, 7.1
b	Pure best-response correspondence	4.1
b_i^N	Pure best-response correspondence	2.1
B	Mixed best-response correspondence	4.1
\mathcal{B}^i	Best-response region	4.2, 7.1
\mathcal{B}^{ij}	Boundary between best-response regions	7.1
$c_{x,y}^N$	Cost of a step	3.2
$c(\phi)$	Cost of path ϕ	4.3
$c^N(\phi^N)$	Cost of path ϕ^N	3.2
$C(K, \Xi)$	Minimum cost of a path from K to Ξ	5
$C^N(K^N, \Xi^N)$	Minimum cost of a path from K^N to Ξ^N	3.3
$C(\tau_K)$	Cost of tree τ_K	6.2
$C^N(\tau_{K^N})$	Cost of tree τ_{K^N}	6.2
e_i	Standard basis vector	2.1
f	Potential function	7.5
F	Continuous population game	4.1
F^N	Payoff function	2.1
$F_{i \rightarrow}^N$	Clever payoff function	2.1
ℓ^N	Length of a discrete path	2.3
K	Limiting recurrent class	4.2
\tilde{K}	Union of limiting recurrent classes other than K	6.1
K^N	Recurrent class	3.1
\tilde{K}^N	Union of recurrent classes other than K^N	3.1
\mathcal{K}	Set of limiting recurrent classes	4.2
\mathcal{K}^N	Set of recurrent classes	3.1
$L(x, u)$	Running cost function	7.2
N	Population size	2.1
$P_{x,y}^{N,\eta}$	Transition probability	2.3
Q^{ijk}, Q	Skew	7.1, 7.3
$R(K)$	Minimal cost of a K -tree	6.2
$R^N(K^N)$	Minimal cost of a K^N -tree	6.2
$s(x)$	Support of state x	3.1
S	Set of strategies	2.1
$S(x)$	$\{y \in X : s(y) \subseteq s(x)\}$	4.2
$\mathcal{S}^N(K^N)$	Strong basin of attraction	3.3
T^N	Duration of a discrete path	2.3
\mathcal{T}_K	Set of K -trees	6.2
TX	Tangent space of X	4.3
v^{ij}	$A' \zeta^{ij}$	7.3
$V(x)$	Value function	7.2
$V^N(x)$	$\{e_j - e_i : i \in s(x) \text{ and } j \in b_i^N(x)\}$	3.1
$\mathcal{W}^N(K^N)$	Weak basin of attraction	3.3
x	Population state	2.1
x^*	Completely mixed equilibrium	7.3
x^{ij}	Partially mixed equilibrium	7.3
X	Simplex	2.1
\mathcal{X}^N	Discrete state space	2.1

TABLE 1. Table of notation.

Notation	Meaning	Section
\mathcal{X}_i^N	Set of states at which strategy i is in use	2.1
$\mathbf{X}^{N,\eta}, X_k^{N,\eta}$	Stochastic evolutionary process	2.3
Z	$\text{conv}(\{e_j - e_i : i, j \in S\})$	4.3
$\gamma(x, y)$	Cost of the direct path from x to y	7.2
ζ^{ij}	$(x_k^*)^{-1}(x^{ij} - x^*)$	7.3
η	Noise level	2.2
$\mu^{N,\eta}$	Stationary distribution	2.3
$\nu(x)$	Feedback control	7.2
σ^η	Noisy best-response function	2.2
τ_K	K -tree	6.2
$\tau^{N,\eta}(\Xi^N)$	Hitting time	3.3
Y	Unlikelihood function	2.2
ϕ	Continuous path	4.3
ϕ^N	Discrete path	2.3
$\dot{\phi}^N$	Discrete right derivative of path ϕ^N	4.3
$\phi^{(N)}$	Interpolated path	4.3
$\Phi(K, \Xi)$	Set of paths from K to Ξ	5
$\Phi^N(K^N, \Xi^N)$	Set of paths from K^N to Ξ^N	3.3
$\mathbf{1}$	Vector of 1s	7.1
$ \cdot $	ℓ^1 norm	4.3
$[\cdot]_+$	Positive part vector	4.3
$[\cdot]_-$	Negative part vector	4.3
$\langle \cdot, \cdot \rangle$	Standard inner product	4.3

TABLE 1. *Continued.*

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 1. Suppose that N 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

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

◇

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 $y = x + (1/N)(e_j - e_i)$. 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

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

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

$$b_i^N(x) = \arg \max_{j \in S} F_{i \rightarrow j}^N(x). \tag{3}$$

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. 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

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

which with (1) is equivalent to

$$N e'_i A x - A_{ii} \geq N e'_j A x - A_{ji} \quad \text{and} \quad N e'_j A x - A_{jj} \geq N e'_i A x - A_{ij}. \tag{4}$$

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 $F_i^N(x) \geq F_j^N(x + (1/N)(e_j - e_i))$ for $j \neq i$, which from (4) is true if and only if $A_{ii} > A_{ji}$, as assumed. \diamond

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—in particular, by ensuring that in coordination games, only pure states are rest points—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:

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

Condition (P1) implies that if there is a unique optimal strategy, then this strategy is assigned a probability that approaches 1 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:

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

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 0. To do so, we introduce the *unlikelihood function* $Y: \mathbb{R}^n \rightarrow \mathbb{R}_+^n$, defined by

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

This definition can be expressed equivalently as

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

Either way, the unlikelihood $Y_j(\pi)$ represents the rate of decay of the probability that strategy j is chosen as η approaches 0. Because they are defined using logarithms of choice probabilities, the unlikelihoods of (conditionally) independent choices combine additively. This fact plays a basic role in the analysis; see Section 3.2.⁸

We maintain the following assumptions throughout the paper:

$$\text{The limit in (5) exists for all } \pi \in \mathbb{R}^n. \tag{U1}$$

$$\text{The unlikelihood function } Y \text{ is continuous.} \tag{U2}$$

$$\text{We have } Y_j(\pi) = 0 \text{ if and only if } j \in \operatorname{argmax}_{k \in S} \pi_k. \tag{U3}$$

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

$$\sigma_j^\eta(\pi) = \frac{\exp(\eta^{-1} \pi_j)}{\sum_{k \in S} \exp(\eta^{-1} \pi_k)}. \tag{6}$$

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

⁸Blume (2003) and Sandholm (2010b) 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 4 below.

strategies with control costs given by an entropy function.⁹ It is easy to verify that this protocol satisfies conditions (U1)–(U3) with piecewise linear unlikelihood function

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

EXAMPLE 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 function exists. Writing η for $1/m$, we obtain the protocol

$$\sigma_j^\eta(\pi) = \mathbb{P}\left(j \in \arg \max_{k \in S} (\pi_k + \bar{\varepsilon}_k^m)\right).$$

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

EXAMPLE 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 4**. **Dokumacı and Sandholm (2011)** provide an explicit, piecewise quadratic expression for the unlikelihood function Y . \diamond

The only noisy best-response 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. Under this protocol, any suboptimal strategy has unlikelihood 1 and a unique optimal strategy has unlikelihood 0, so condition (U2) must fail. For further discussion of the BRM protocol, see **Remark 13** and **Example 14**.

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 $1/N$ 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 $F_{i \rightarrow \cdot}^N(x)$ 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 k/N units of clock time. The transition probabilities

⁹See **Anderson et al. (1992)** or **Hofbauer and Sandholm (2002)**.

$P_{x,y}^{N,\eta}$ for the process $\mathbf{X}^{N,\eta}$ are given by

$$\begin{aligned}
 P_{x,y}^{N,\eta} &\equiv \mathbb{P}(X_{k+1}^{N,\eta} = y | X_k^{N,\eta} = x) \\
 &= \begin{cases} x_i \sigma_j^\eta(F_{i \rightarrow \cdot}^N(x)) & \text{if } y = x + \frac{1}{N}(e_j - e_i), j \neq i \\ \sum_{i=1}^n x_i \sigma_i^\eta(F_{i \rightarrow \cdot}^N(x)) & \text{if } y = x \\ 0 & \text{otherwise.} \end{cases}
 \end{aligned} \tag{7}$$

It is easy to verify that $\sum_{y \in \mathcal{X}^N} P_{x,y}^{N,\eta} = 1$ for all $x \in \mathcal{X}^N$.

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 $1/N$ 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 0, 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 with 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 $(1/N)V^N(x)$, where

$$V^N(x) = \{e_j - e_i : i \in s(x) \text{ and } j \in b_i^N(x)\} \tag{8}$$

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

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

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 .¹⁰

EXAMPLE 6. Let F^N be defined by random matching in the normal form coordination game A as in **Example 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 $y = x + (1/N)(e_i - e_j)$ is allowed under (DBR). Since $i \in b_j^N(x)$ is equivalent to $i \in b_i^N(x + (1/N)(e_i - e_j))$, 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 classes of (DBR) correspond to the pure states, $\mathcal{K}^N = \{e_1, \dots, e_n\}$, as shown by **Kandori and Rob (1998)**.¹¹ ◇

EXAMPLE 7. Again let F^N be defined by random matching in the normal form coordination game A . If $x \in \mathcal{X}^N$ is not Nash, there is a strategy j in the support of x satisfying $j \notin b_j^N(x)$. **Lemma 5** 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 5** inductively shows that from every state $x \in \mathcal{X}^N$, there is a solution to (DBR) that terminates at a pure state, implying that $\mathcal{K}^N = \{e_1, \dots, e_n\}$. ◇

¹⁰One 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 $\mathbf{X}^{N,*}$.

¹¹Unlike our model, the model of **Kandori and Rob (1995, 1998)** allows multiple revisions during each period; see also **footnote 18** below.

We conjecture that the set of recurrent classes of (DBR) is $\mathcal{K}^N = \{\{e_1\}, \dots, \{e_n\}\}$ for any coordination game as defined in Example 2. Example 8 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 $\mathbf{X}^{N,\eta}$ 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 $x \in \mathcal{X}^N$ to $y \in \mathcal{X}^N$ by

$$c_{x,y}^N = - \lim_{\eta \rightarrow 0} \eta \log P_{x,y}^{N,\eta},$$

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 (7), we can represent step costs in terms of the game's payoff function and the protocol's unlikelihood function:

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

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 opportunities.¹² 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:

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

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:

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

Definitions (7) 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

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

where \approx refers to the order of magnitude in η as η approaches 0. By statement (10), path ϕ^N has cost zero if and only if it is a solution to (DBR).

¹²The 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.

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

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

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 minimizing 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

$$\tilde{K}^N = \bigcup_{L^N \in \mathcal{K}^N \setminus \{K^N\}} L^N \tag{13}$$

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 .¹³ Proposition 1 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 1. *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 1 concerns behavior within the *strong basin of attraction* of K^N , the set of states $\mathcal{J}^N(K^N) = \mathcal{X}^N \setminus \mathcal{W}^N(\tilde{K}^N) \subseteq \mathcal{W}^N(K^N)$ 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 L^N is reached. But while the analogue of Proposition 1(ii) on the likely points of entry into

¹³Thus the cost of exit from K^N corresponds to the *radius* of K^N as defined by Ellison (2000).

$\mathcal{W}^N(L^N)$ is true, the analogue of Proposition 1(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.¹⁴

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 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

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

Thus $\Delta r^N(x)$ describes the exponential rate of decay of the stationary distribution weight on x as η approaches 0.

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.¹⁵ 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,¹⁶ 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

¹⁴This is the reason for the correction term appearing in Proposition 4.2 of Catoni (1999). See Freidlin and Wentzell (1998, pp. 197–198) for a clear discussion of this point.

¹⁵Explicitly, 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 (2010c, Section 12.A.5) explains this distinction in detail.

¹⁶See Young (1993, 1998), Kandori and Rob (1995), and Sandholm (2010c, Section 12.A).

a (continuous) population game, $F : X \rightarrow \mathbb{R}^n$, which we take to be a continuous function 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

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

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

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

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 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

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

Comparing the final expression above to definition (8), 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.

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):

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




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

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).¹⁷ 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 $\mathbf{0}$.¹⁸

In the classes of coordination games considered in Examples 6 and 7, the set \mathcal{K}^N of recurrent classes of the discrete best-response dynamic (DBR) is equal to $\mathcal{K} = \{\{e_1\}, \dots, \{e_n\}\}$ 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 8. Consider the continuous-population game $F(x) = Ax$ generated by a coordination game A (Example 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

¹⁷See Roth and Sandholm (2013) for a formal limit result.

¹⁸Kandori 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.

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. \diamond

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

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

where $\text{dist}(K_i^N, K_i)$ denotes the ℓ^1 -Hausdorff distance between K_i^N and K_i ¹⁹

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

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

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 $\phi_{k+1}^N = \phi_k^N + (1/N)(e_{j^N(k)} - e_{i^N(k)})$ and, hence,

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

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

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

¹⁹That is, $\text{dist}(K_i^N, K_i) = \max\{\max_{x \in K_i^N} \min_{y \in K_i} |x - y|, \max_{y \in K_i} \min_{x \in K_i^N} |x - y|\}$, where $|\cdot|$ is the ℓ^1 norm on \mathbb{R}^n .

This definition also accounts for each period in the N -agent process lasting $1/N$ 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

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

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 (18) and (20), any interpolated path $\phi^{(N)}$ satisfies the following bound on its speed:

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

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

$$c_{\phi_k^N, \phi_{k+1}^N}^N = Y_{j^N(k)}(F_{i^N(k) \rightarrow \cdot}^N(\phi_k^N)) = \langle Y(F_{i^N(k) \rightarrow \cdot}^N(\phi_k^N)), [\dot{\phi}_k^N]_+ \rangle. \tag{22}$$

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

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

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

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

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 nonpausing 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:

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

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

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

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:

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

This conclusion will justify definition (24) of the cost of a nonpausing 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, which is needed to obtain the lower bound in (27), 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 second assumption, which is needed to obtain the upper bound in (27), 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.

Motivated by bound (21), we say that piecewise monotone path ϕ *moves at full speed* if

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

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 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); in so doing, one verifies **Assumption 2** directly.

So as to appeal to **Assumption 2**, we assume in what follows that the target set Ξ is either a limiting recurrent class or a singleton.²⁰

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 $(1/N)C^N(K^N, \Xi^N)$.

The key to obtaining the lower bound is to show that if the normalized costs $(1/N)c^N(\phi^N)$ of a sequence of discrete paths of bounded durations converge, then the

²⁰For 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 10** to derive the asymptotics of the stationary distribution on the entire state space, rather than just its asymptotics on the recurrent classes.

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](#).

PROPOSITION 2. *Let $\{\phi^N\}_{N=N_0}^\infty$ be a sequence of paths with durations at most \bar{T} and whose costs satisfy $\lim_{N \rightarrow \infty} (1/N)c^N(\phi^N) = C^*$. 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 2](#), the costs of their linear interpolations $\phi^{(N)} \in \Phi(K^N, \Xi^N)$ 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 3. *We have $\liminf_{N \rightarrow \infty} (1/N)C^N(K^N, \Xi^N) \geq C(K, \Xi)$.*

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 2](#), 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 on 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 X^N are unable to move in this direction. At best, they can alternate between increments $(1/N)(e_2 - e_1)$ (i.e., switches by a single agent from 1 to 2) and $(1/N)(e_3 - e_1)$. 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.

Despite this difficulty, it is possible to construct discrete approximations ϕ^N whose costs approach those of the continuous path ϕ , provided that ϕ is piecewise monotone.²¹ To begin, Proposition 4 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 $2n/N$ of ϕ in the uniform norm.

PROPOSITION 4. *Suppose $\phi = \{\phi_t\}_{t \in [0, T]}$ is monotone and moves at full speed. If $N \geq 1/T$, there is an $s^N \in [0, 1/N)$ and a feasible monotone path $\phi^N = \{\phi_k^N\}_{k=0}^{\ell^N}$, $\ell^N = \lfloor N(T - s^N) \rfloor$ satisfying*

$$\max_{0 \leq k \leq \ell^N} |\phi_k^N - \phi_{s^N + k/N}| \leq \frac{2n}{N}. \tag{29}$$

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

Next, Proposition 5 shows that the normalized costs of the discrete paths so constructed converge to the cost of the original path ϕ .

PROPOSITION 5. *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 $\lim_{N \rightarrow \infty} (1/N)c^N(\phi^N) = c(\phi)$.*

The proof of Proposition 5 is presented in Appendix A.5, but we explain the logic of the proof here. By (23) and definition (17) of ϕ^N , we can express the normalized cost of path ϕ^N as

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

²¹As an aside, we note that the continuous and discrete paths in Figure 2 are both monotone with $S_+ = \{2, 3\}$ and $S_- = \{1\}$.

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

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

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. (86)) whose integrator ϕ is monotone. Since there are $O(\sqrt{N})$ rather than $O(N)$ summands, the $O(1/N)$ bound in (29) ensures that replacing ϕ^N with ϕ leads to an approximation of $(1/N)c^N(\phi^N)$ that is asymptotically correct.²² 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 4 and 5, 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 6. *We have $\limsup_{N \rightarrow \infty} (1/N)C^N(K^N, \Xi^N) \leq C(K, \Xi)$.*

Together, Propositions 3 and 6 establish the convergence of minimal path costs.

THEOREM 9 (Convergence theorem). *We have $\lim_{N \rightarrow \infty} (1/N)C^N(K^N, \Xi^N) = C(K, \Xi)$.*

6. CONSEQUENCES FOR LONG-RUN BEHAVIOR

We now use the convergence theorem to characterize exit times, stationary distribution asymptotics, and stochastic stability in the small noise double limit. These characterizations 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 $K^N \in \mathcal{K}^N$, (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.

²²One 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.

To characterize this expected waiting time, we let K be the limiting recurrent class corresponding to K^N (cf. (16)), and we 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 1(i) and Theorem 9 immediately yields the following result.

COROLLARY 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 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 vanishes is approximately $NC(K, \tilde{K})$. This quantity can be evaluated explicitly by solving control problem (26).

Turning to exit locations, Proposition 1(ii) showed that in the small noise limit, the exit point of $\mathbf{X}^{N,\eta}$ from the strong basin of attraction $\mathcal{B}^N(K^N) = \mathcal{X}^N \setminus \mathcal{W}^N(\tilde{K}^N)$ 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 $\mathcal{B}^N(K^N)$ 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 though \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:

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

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

$$r^N(x) = \min_{K^N \in \mathcal{K}^N} (R^N(K^N) + C^N(K^N, \{x\})). \tag{31}$$

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 .²³ 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 7 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 7. *The stationary distributions $\mu^{N,\eta}$ satisfy*

$$-\lim_{\eta \rightarrow 0} \eta \log \mu^{N,\eta}(x) = \Delta r^N(x) \quad \text{for all } x \in \mathcal{X}^N. \tag{15}$$

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

THEOREM 10. *The stationary distributions $\mu^{N,\eta}$ satisfy*

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

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 10**, one that does not require uniformity of the large N limit in x , would follow directly from **Theorem 9** and **Proposition 7**.²⁴ To prove **Theorem 10** as stated, we need to show that the limit in **Theorem 9** is uniform over all choices

²³State x need not be in the weak basin of the recurrent class that yields the minimum in (31).

²⁴In 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 9**.

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 10, 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.²⁵ Theorem 10 implies that state x is stochastically stable in the small noise double limit if and only if $\Delta r(x) = 0$.²⁶

7. AN ANALYSIS OF THE LOGIT MODEL

To move from the results in the previous section to analyses of specific examples, one needs to solve instances of the path cost-minimization problem (26). In this section, we show how to solve such problems using optimal control techniques, and we combine these solutions with the results below to describe long-run play in particular examples. Our focus is on evolution under the logit choice rule (Example 3), in three-strategy coordination games (Example 2) that satisfy the marginal bandwagon property (Example 6) and that have an interior equilibrium. In Section 8.2, we explain why it should be possible to carry out similar analyses in other settings.

7.1 Definitions

7.1.1 Notation and definitions for symmetric normal form games We begin by introducing a convenient new notation for working with symmetric normal form games $A \in \mathbb{R}^{n \times n}$. We use superscripts to refer to rows of A and subscripts to refer to columns. Thus A^i is the i th row of A , A_j is the j th column of A , and A^i_j is the (i, j) th entry. These objects can be obtained by pre- and postmultiplying A by standard basis vectors:

$$A^i = e'_i A, \quad A_j = A e_j, \quad A^i_j = e'_i A e_j.$$

In a similar fashion, we use super- and subscripts of the form $i - j$ to denote certain differences obtained from A :

$$A^{i-j} = A^i - A^j = (e_i - e_j)' A, \quad A^{i-j}_{k-\ell} = A^i_k - A^i_\ell - A^j_k + A^j_\ell = (e_i - e_j)' A (e_k - e_\ell).$$

In this notation, the best-response region for strategy i is described by

$$\mathcal{B}^i = \{x \in X : A^{i-j} x \geq 0 \text{ for all } j \in S\}.$$

The set $\mathcal{B}^{ij} = \mathcal{B}^i \cap \mathcal{B}^j$ is the boundary between the best-response regions for strategies i and j .

²⁵Logically, $\forall \delta > 0 \forall O \in \mathcal{O}(X, x) \exists N_0 \in \mathbb{N} \forall N > N_0 \exists \eta_0 > 0 \forall \eta < \eta_0 \mu^{N,\eta}(O) > \exp(-\eta\delta)$, where $\mathcal{O}(X, x)$ denotes the set of open subsets of X containing x .

²⁶This 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 $y \in O \cap \mathcal{X}^N$.

In the present notation, A is a coordination game (Example 2) if

$$A_i^i > A_i^j \quad \text{for all } i, j \in S \text{ with } j \neq i, \tag{32}$$

so that each pure state is a Nash equilibrium of F . This implies that

$$A_{i-j}^{i-j} > 0 \quad \text{for all } i, j \in S. \tag{33}$$

We call $A_{i-j}^{i-j} = A_{j-i}^{j-i}$ the (i, j) th *alignment* of A . This quantity, which corresponds to the denominator of the mixed equilibrium weights in the binary-choice game with strategies i and j , represents the strength of incentives to coordinate (or, if negative, to miscoordinate) in the restricted game with strategy set $\{i, j\}$.

Likewise, game A has the *marginal bandwagon property* (Example 6) if

$$A_{i-k}^{i-j} > 0 \quad \text{for all } i, j, k \in S \text{ with } i \notin \{j, k\}. \tag{34}$$

As noted earlier, this property requires that when opponents switch to strategy i from some other strategy, the payoffs to playing strategy i improve relative to those of all other strategies. In three-strategy coordination games with an interior equilibrium, property (34) has a simple geometric interpretation: it requires that the boundaries between best-response regions do not hit the boundary of the simplex at sharp angles; see Section 7.3.1, especially Figure 4.

The next definition for games with three or more strategies plays a basic role in our analysis. For an ordered triple of distinct strategies (i, j, k) , we define the (i, j, k) th *skew* of A by

$$\begin{aligned} Q^{ijk} &= A_{j-k}^i + A_{k-i}^j + A_{i-j}^k \\ &= A_{i-k}^{i-j} - A_{i-j}^{i-k} = A_{j-i}^{j-k} - A_{j-k}^{j-i} = A_{k-j}^{k-i} - A_{k-i}^{k-j}. \end{aligned} \tag{35}$$

Evidently skew is alternating, in the sense that it is preserved by even permutations of the index list and negated by odd ones:

$$Q^{ijk} = Q^{jki} = Q^{kij} = -Q^{kji} = -Q^{jik} = -Q^{ikj}. \tag{36}$$

We call A a *potential game* if $A = C + \mathbf{1}r'$ for some symmetric matrix $C \in \mathbb{R}^{n \times n}$ and some vector $r \in \mathbb{R}^n$, where $\mathbf{1} \in \mathbb{R}^n$ denotes the vector of 1s. Thus A is the sum of a *common interest game* C and a *passive game* $\mathbf{1}r'$ in which a player's payoff depends only on his opponent's strategy. Clearly, games A and C induce the same best-response correspondence and the same set of Nash equilibria.

Potential games admit a variety of characterizations. For instance, A is a potential game if and only if $\Phi A \Phi$ is a symmetric matrix, where $\Phi = I - (1/n)\mathbf{1}\mathbf{1}' \in \mathbb{R}^{n \times n}$ is the orthogonal projection onto the tangent space $TX = \{z \in \mathbb{R}^n : \sum_i z_i = 0\}$.²⁷ The latter condition says that A is a symmetric bilinear form on $TX \times TX$, meaning that $z'Az = \hat{z}'A\hat{z}$

²⁷The “only if” direction of this claim is obvious. Letting $\Xi = (1/n)\mathbf{1}\mathbf{1}' = I - \Phi$, the “if” direction follows from the decomposition $A = (\Phi A \Phi + (\Phi A \Xi + \Xi A' \Phi)) + \Xi(A - A' \Phi)$. Compare Sandholm (2010a).

for all $z, \hat{z} \in TX$. Alternatively, A is a potential game if and only if it satisfies the triangular integrability condition of Hofbauer (1985), which can be stated in terms of skews: $Q^{ijk} = 0$ for all distinct $i, j, k \in S$.²⁸

7.1.2 Path costs and the minimum cost path problems To determine the path cost function for the present context, recall that the unlikelihood function for the logit choice protocol (6) is

$$Y_i(\pi) = \max_{j \in S} \pi_j - \pi_i. \tag{37}$$

Plugging this expression into (24), we find that the cost of continuous path ϕ under the logit protocol in the linear population game $F(x) = Ax$ is

$$c(\phi) = \int_0^T [\dot{\phi}_t]'_+ (\mathbf{1}A^{\check{b}(\phi_t)} - A)\phi_t dt, \tag{38}$$

where $\check{b}(\cdot)$ is any selection from the game's pure best-response correspondence $b(\cdot)$.

The results in Section 6 described the long-run behavior of the process $\mathbf{X}^{N,\eta}$ in terms of the minimal costs of continuous paths from limiting recurrent classes to unions of these classes. In the case of coordination games with the marginal bandwagon property, Example 6 shows that the set of limiting recurrent classes is the set of pure equilibria: $\mathcal{X} = \{e_1, \dots, e_n\}$. Corollary 1 thus implies that in the small noise double limit, the expected time until the process $\mathbf{X}^{N,\eta}$ exits from equilibrium e_i to another equilibrium is captured by the *cost of exit*

$$C\left(\{e_i\}, \bigcup_{j \neq i} \{e_j\}\right) = \min \left\{ c(\phi) : \phi \in \Phi\left(\{e_i\}, \bigcup_{j \neq i} \{e_j\}\right) \right\}. \tag{39}$$

Theorem 10 shows that to evaluate limiting stationary distributions and stochastic stability in the small noise double limit, we must assess the *costs of transitions* between strict equilibria:

$$C(\{e_i\}, \{e_j\}) = \inf \{ c(\phi) : \phi \in \Phi(\{e_i\}, \{e_j\}) \}. \tag{40}$$

Example 8 shows that in coordination games, a straight-line path from any state in best-response region \mathcal{B}^j to equilibrium e_j has zero cost. Thus replacing $\{e_j\}$ with \mathcal{B}^j in (39) and (40) does not change the minimal cost in either case, and we will write the minimal cost path problems this way in what follows.

7.2 Preliminary analysis

7.2.1 A verification theorem To understand the long-term behavior of the processes $\mathbf{X}^{N,\eta}$ in the small noise double limit, we must solve the exit cost problems (39) and the transition cost problems (40). These problems have nonsmooth running costs and are multidimensional in games with more than two strategies. Nevertheless, these problems

²⁸See Sandholm (2009, Proposition 4.5).

can be solved explicitly. We now introduce the result from optimal control theory that we use to do so.

Let \mathcal{A} be an m -dimensional affine subspace of \mathbb{R}^n with tangent space $T\mathcal{A}$, and let the set $\Omega \subset \mathcal{A}$ be closed relative to \mathcal{A} and have piecewise smooth boundary. Let the function $L : \mathcal{A} \times T\mathcal{A} \rightarrow \mathbb{R}_+$ be Lipschitz continuous, and let $\mathcal{Z} \subset T\mathcal{A}$ be compact and convex. The control problem and its value function $V^* : \mathcal{A} \rightarrow \mathbb{R}_+$ are defined as

$$\begin{aligned}
 V^*(x) &= \min \int_0^T L(\phi_t, \nu_t) dt \\
 &\text{over } T \in [0, \infty), \quad \nu : [0, T] \rightarrow \mathcal{Z} \text{ measurable} \\
 &\text{subject to } \phi : [0, T] \rightarrow \mathcal{A} \text{ absolutely continuous} \tag{41} \\
 &\quad \phi_0 = x, \quad \phi_T \in \Omega \\
 &\quad \dot{\phi}_t = \nu_t \text{ for almost every } t \in [0, T].
 \end{aligned}$$

Theorem 11 provides sufficient conditions for a function $V : \mathcal{A} \rightarrow \mathbb{R}_+$ to be the value function of (41). The key requirement is that the *Hamilton–Jacobi–Bellman (HJB) equation*

$$\min_{u \in \mathcal{Z}} (L(x, u) + DV(x)u) = 0 \tag{42}$$

hold at almost every $x \in \mathcal{A}$.

THEOREM 11 (Verification theorem (Boltyanskii 1966, Piccoli and Sussmann 2000)). *Let $V : \mathcal{A} \rightarrow \mathbb{R}_+$ be a continuous function that is continuously differentiable except on the union $\mathcal{U} \subset \mathcal{A}$ of a finite number of manifolds, each of dimension less than m . Assume the following:*

- (i) *For every $x \in \mathcal{A}$, there is a time $T \in [0, \infty)$ and a measurable function $\nu : [0, T] \rightarrow \mathcal{Z}$ such that the corresponding controlled trajectory $\phi : [0, T] \rightarrow \mathcal{A}$ with $\phi_0 = x$ satisfies $\phi_T \in \Omega$ and $\int_0^T L(\phi_t, \nu_t) dt = V(x)$.*
- (ii) *The HJB equation (42) holds at all $x \in \mathcal{A} \setminus \mathcal{U}$.*
- (iii) *The boundary condition $V(x) = 0$ holds at all $x \in \Omega$.*

Then $V = V^*$.

Condition (i) of the theorem says that the values specified by the function V can all be achieved, and so implies that $V^* \leq V$. Establishing the opposite inequality is straightforward if V is C^1 . Suppose that this is the case, and that $\hat{T} \in [0, \infty)$ and $\hat{\nu} : [0, \hat{T}] \rightarrow \mathcal{Z}$ are feasible in problem (41), so that the controlled trajectory $\hat{\phi} : [0, \hat{T}] \rightarrow \mathcal{A}$ with $\hat{\phi}_0 = x$ satisfies $\hat{\phi}_{\hat{T}} \in \Omega$. Then the HJB equation (42) implies that

$$L(\hat{\phi}_t, \hat{\nu}_t) \geq -DV(\hat{\phi}_t)\hat{\nu}_t = -\frac{d}{dt}V(\hat{\phi}_t) \text{ for almost all } t \in (0, \hat{T}).$$




FIGURE 3. The original and extended versions of the transition problem (40).

Integrating and applying the boundary condition (iii) yields

$$\int_0^{\hat{T}} L(\phi_t, \nu_t) \geq -(V(\hat{\phi}_{\hat{T}}) - V(\hat{\phi}_0)) = V(x),$$

and so $V^* \geq V$.

To prove [Theorem 11](#) as stated, one establishes that the cost of any feasible controlled trajectory can be approximated arbitrarily well by the cost of a feasible controlled trajectory that only intersects the manifolds in \mathcal{U} at a finite set of times. The first result along these lines is due to [Boltyanskii \(1966\)](#), with various improvements culminating in the work of [Piccoli and Sussmann \(2000\)](#). [Theorem 11](#) above follows from the statement and proof of [Theorem 6.3.1](#) in the textbook treatment of [Schättler and Ledzewicz \(2012\)](#).²⁹

While our control problems are set in the simplex X , [Theorem 11](#) addresses problems whose state space is an affine subset of \mathbb{R}^n . To use the theorem, we redefine our problems by extending their state space to the affine hull $\text{aff}(X) = \{x \in \mathbb{R}^n : \sum_i x_i = 1\}$ of X . Since our target sets are defined by linear inequalities, we can define the target sets of our extended problem by imposing the same linear inequalities in $\text{aff}(X)$ rather than in X ([Figure 3](#)). If in this extended problem, the optimal paths from initial conditions in X to the extended target set are themselves contained in X , then these paths are optimal in the original problem; consequently, the restriction of the resulting value function to X is the value function of the original problem. This is precisely what happens in the minimum cost path problems for the games we focus on here. We discuss the general case in [Section 8.2](#).

²⁹In the statement of [Theorem 6.3.1](#) of [Schättler and Ledzewicz \(2012\)](#), \mathcal{A} is all of \mathbb{R}^n , the function L is C^1 , and the target set Ω is required to have a smooth boundary. However, inspection of their proof reveals that it goes through unchanged under the weaker requirements imposed in [Theorem 11](#) above.

Also, recall that under path cost function (24), reparameterizing a path—changing the speed at which the states in the path are traversed—does not affect its cost. Thus in looking for minimum cost paths between sets in $\text{aff}(X)$, it is without loss of generality to consider paths satisfying $\dot{\phi}_t \in Z$, where Z is the compact set $\text{conv}(\{e_i - e_j : i, j \in S\})$.

The control problem (41) is stated in terms of control trajectories $\nu : [0, T] \rightarrow Z$, which specify the control vectors as a function of time. It is convenient here to work with *feedback controls* $\hat{\nu} : \mathcal{A} \rightarrow Z$, which specify the control vectors as a function of the current state. The corresponding controlled trajectories are the Carathéodory solutions to the differential equation $\dot{\phi}_t = \hat{\nu}(\phi_t)$.³⁰ If ϕ is such a trajectory and if we define the control trajectory ν by $\nu_t = \hat{\nu}(\phi_t)$, then the pair (ϕ, ν) satisfies $\dot{\phi}_t = \nu_t$ for almost all times t , as required in problem (41).

7.2.2 A lemma for checking the HJB equation We now introduce a basic tool for verifying the HJB equation in our setting. When x is in $\mathcal{B}^i \subset \text{aff}(X)$, the HJB equation (42) becomes

$$\min_{u \in Z} ([u]'_+ (\mathbf{1}A^i - A)x + DV(x)u) = 0. \tag{43}$$

Since the function being minimized in (43) is linear in u on each orthant of \mathbb{R}^n , there must be a minimizer either at an extreme point of Z or at the origin, where the function evaluates to 0. Therefore, substituting $e_a - e_b$ for u , we see that (43) is equivalent to

$$\min_{e_a, e_b \neq e_a} ((e_i - e_a)'Ax + DV(x)(e_a - e_b)) \geq 0. \tag{44}$$

Lemma 1 provides a sufficient condition for the HJB equation (44) to be satisfied at a state in the (relative) interior of \mathcal{B}^i when A is a three-strategy game.

LEMMA 1. *Let A be a three-strategy game with $S = \{i, j, k\}$. Suppose that the candidate value function V is constructed from a feedback control that takes value $e_k - e_i$ at all states in a neighborhood of $x \in \text{int}(\mathcal{B}^i)$. If*

$$DV(x)(e_i - e_h) \geq 0 \quad \text{for } h \in \{j, k\} \tag{45}$$

$$(DV(x) - (Ax)')(e_j - e_k) \geq 0, \tag{46}$$

then V satisfies the HJB equation (44) at x .

The proof of **Lemma 1** is presented in **Appendix A.8**. We argue there that the assumption that the control is $e_k - e_i$ in a neighborhood of x implies that the function to be minimized in the HJB equation (44) equals 0 when $e_a = e_k$ and $e_b = e_i$. This equality can be restated as

$$(DV(x) - (Ax)')(e_k - e_i) = 0. \tag{47}$$

The proof then uses conditions (45)–(47) and the fact that $x \in \mathcal{B}^i$ to show that the function to be minimized in (44) is nonnegative for the remaining five choices of $e_a - e_b$.

³⁰A *Carathéodory solution* to an ordinary differential equation is an absolutely continuous trajectory that satisfies the equation at almost all times.

7.2.3 *Costs of direct paths* As a final preliminary, we present two simple formulas for path costs (38) in linear games under the logit rule. For $x, y \in \text{aff}(X)$, we let $\gamma(x, y)$ denote the cost of the direct (straight-line) path from x to y :

$$\gamma(x, y) = c(\phi), \quad \text{where } \phi: [0, 1] \rightarrow X \text{ is defined by } \phi_t = (1 - t)x + ty.$$

The first formula concerns a class of direct paths whose costs are easily expressed in terms of the paths' endpoints: those in which the motion of the state involves agents switching away from the current best response.

LEMMA 2. *Suppose that $x, y \in \mathcal{B}^i$, and that $y - x = d(\alpha - e_i)$ for some $\alpha \in X$ with $\alpha_i = 0$ and some $d \geq 0$. Then*

$$\gamma(x, y) = (x - y)' A \left(\frac{x + y}{2} \right).$$

PROOF. Since $\dot{\phi}_t = y - x = d(\alpha - e_i) \in TX$ and $[\dot{\phi}_t]'_-(\mathbf{1}A^i - A) = de'_i(\mathbf{1}A^i - A) = d(A^i - A^i) = \mathbf{0}'$,

$$\begin{aligned} \gamma(x, y) &= \int_0^1 [\dot{\phi}_t]'_+(\mathbf{1}A^i - A)\phi_t dt = \int_0^1 \dot{\phi}_t'(\mathbf{1}A^i - A)\phi_t dt = - \int_0^1 \dot{\phi}_t' A \phi_t dt \\ &= -(y - x)' A \int_0^1 \phi_t dt = (x - y)' A \left(\frac{x + y}{2} \right). \quad \square \end{aligned}$$

In some important cases this formula can be simplified further. The second formula describes the costs that are realized when the state moves from $x \in \mathcal{B}^i$ in direction $e_k - e_i$ until reaching a state y in the set $\mathcal{B}^{ij} = \mathcal{B}^i \cap \mathcal{B}^j$, where strategies i and j are both optimal.

LEMMA 3. *Let $x \in \mathcal{B}^i$, and suppose that*

$$y = x + d(e_k - e_i) \in \mathcal{B}^{ij} \quad \text{for some } d > 0 \tag{48}$$

and that $A_{i-k}^{i-j} \neq 0$. Then

$$d = \frac{A^{i-j}x}{A_{i-k}^{i-j}} \tag{49}$$

$$\gamma(x, y) = dA^{i-k}y + \frac{1}{2}d^2A_{i-k}^{i-k}. \tag{50}$$

In particular, if $j = k$, then $A^{i-k}y = 0$, so (50) becomes

$$\gamma(x, y) = \frac{1}{2}d^2A_{i-k}^{i-k} = \frac{1}{2} \frac{(A^{i-k}x)^2}{A_{i-k}^{i-k}}.$$

PROOF. Since $y \in \mathcal{B}^{ij}$, we have $A^{i-j}y = 0$, which with (48) implies (49). Also, combining (48) and Lemma 2 yields (50), since

$$\begin{aligned} \gamma(x, y) &= (x - y)' A \left(\frac{x + y}{2} \right) = d(e_i - e_k)' A \left(y + \frac{1}{2}d(e_i - e_k) \right) \\ &= dA^{i-k}y + \frac{1}{2}d^2 A_{i-k}^{i-k}. \end{aligned} \quad \square$$

7.3 Construction of value functions: The initial step

7.3.1 *Simple three-strategy coordination games* We now focus on three-strategy coordination games (32) that satisfy the marginal bandwagon property (34) and that admit a completely mixed equilibrium, a class of games we henceforth call *simple three-strategy coordination games*. The completely mixed equilibrium $x^* \in \text{int}(X)$ is the unique state in $\text{aff}(X)$ at which the payoffs to all strategies are equal: $Ax^* = c\mathbf{1}$ for some $c \in \mathbb{R}$. For distinct strategies $i, j \in S$, such games admit a unique mixed equilibrium x^{ij} with support $\{i, j\}$. This x^{ij} is the unique state in \mathcal{B}^{ij} with $x_k = 0$.

We now define two vectors that play basic roles in the analysis to come. For distinct strategies $i, j \in S$, we define the vector $\zeta^{ij} \in TX$ by

$$\zeta^{ij} = \frac{1}{x_k^*} (x^{ij} - x^*). \tag{51}$$

When drawn with its tail at x^* , ζ^{ij} points outward along the boundary \mathcal{B}^{ij} between best-response regions \mathcal{B}^i and \mathcal{B}^j (Figure 4). Since the vector $(A^{j-i})'$ is normal to \mathcal{B}^{ij} , ζ^{ij} is a multiple of the cross-product

$$(A^{j-i})' \times \mathbf{1} = A_{j-k}^{j-i} e_i + A_{i-k}^{i-j} e_j - A_{j-i}^{j-i} e_k.$$

Since (51) implies that $\zeta_k^{ij} = -1$, it follows that

$$\zeta^{ij} = \frac{A_{j-k}^{j-i}}{A_{j-i}^{j-i}} e_i + \frac{A_{i-k}^{i-j}}{A_{j-i}^{j-i}} e_j - e_k \equiv \beta^{ij} - e_k. \tag{52}$$

The equivalence in (52) defines the vector β^{ij} . Since A is a coordination game with the marginal bandwagon property, β^{ij} is an element of $\text{conv}(\{e_i, e_j\})$, and so ζ^{ij} is a convex combination of $e_i - e_k$ and $e_j - e_k$. Thus boundary \mathcal{B}^{ij} does not hit the boundary of the simplex at an angle of less than 60° , implying that mixed equilibrium x^{ij} is in the sextant northwest of mixed equilibrium x^* ; see Figure 4.³¹

Next, we define the vector $v^{ij} = v^{ji} \in \mathbb{R}^3$ by

$$(v^{ij})' = (\zeta^{ij})' A = \frac{1}{A_{j-i}^{j-i}} (A_{j-k}^{j-i} A^i + A_{i-k}^{i-j} A^j - A_{j-i}^{j-i} A^k).$$

³¹The sextants are six closed convex cones in $\text{aff}(X)$ with common origin x^* and 60° angles between their boundaries. In Figures 4 and 5, the portions of the sextants' boundaries lying in X are represented by dotted lines.




FIGURE 4. Multiples of the vectors ζ^{ij} , ζ^{jk} , and ζ^{ki} .

By definition (51) of ζ^{ij} , $(v^{ij})'x$ is positive if and only if mixed strategy x^{ij} earns a higher payoff than mixed strategy x^* at state x . Both the geometry and the importance of the vector v^{ij} will be made clear below.

7.3.2 *Construction of the value function near the target set* To solve the exit cost problem (39) and the transition cost problem (40) via dynamic programming, we first determine the form of the value function at states near the target set. We therefore consider the cost of reaching the set \mathcal{B}^k from nearby states in \mathcal{B}^i . It is natural to guess that there is a region $R^{ik} \subseteq \mathcal{B}^i$ whose boundary contains \mathcal{B}^{ik} , in which motion in direction $e_k - e_i$ leads to \mathcal{B}^{ik} and defines the optimal feedback control. By Lemma 3, this choice of control generates the candidate value function

$$V(x) = \frac{1}{2} \frac{(A^{i-k}x)^2}{A_{i-k}^{i-k}} \tag{53}$$

in region R^{ik} .

We use Lemma 1 to determine when this function satisfies the HJB equation (44) in R^{ik} . To start, we compute the derivative³² $DV : \text{aff}(X) \rightarrow L(TX, \mathbb{R})$ of V at points in the interior of R^{ik} :

$$DV(x)z = \frac{A^{i-k}x}{A_{i-k}^{i-k}} A^{i-k}z \quad \text{for } x \in \text{int}(R^{ik}). \tag{54}$$

³²Since V is defined on $\text{aff}(X)$, its derivative at x , $DV(x)$, is a linear map from TX to \mathbb{R} . There are many vectors $v \in \mathbb{R}^n$ that represent this map, in the sense that $DV(x)z = v'z$ for all $z \in TX$. The gradient of V at x , $\nabla V(x)$, is defined to be the unique representative of $DV(x)$ in TX ; it can be obtained by applying the orthogonal projection matrix $\Phi = I - \frac{1}{n}\mathbf{1}\mathbf{1}'$ to an arbitrary representative of $DV(x)$ in \mathbb{R}^n . See Sandholm (2010c, Section 3.C) for further discussion.

Since strategies i and k are both best responses at states in \mathcal{B}^{ik} , vectors tangent to \mathcal{B}^{ik} are orthogonal to $(A^{i-k})'$. Equation (54) implies that such vectors \hat{z} satisfy $DV(x)\hat{z} = 0$ and so are tangent to the level sets of the value function. Intuitively, moving the state in a direction tangent to \mathcal{B}^{ik} changes neither the distance needed to travel to \mathcal{B}^{ik} nor the payoff differences that must be overcome en route.

We now apply Lemma 1. To check condition (45), we first observe that

$$DV(x)(e_i - e_h) = \frac{A^{i-k}x}{A_{i-k}^{i-k}} A_{i-h}^{i-k}.$$

Now $A^{i-k}x \geq 0$ (since $x \in \mathcal{B}^i$), $A_{i-k}^{i-k} > 0$ (since A is a coordination game; see (33)), and $A_{i-j}^{i-k} \geq 0$ (by the marginal bandwagon property (34)). Thus $DV(x)(e_i - e_h) \geq 0$ for $h \in \{j, k\}$, establishing condition (45). To check condition (46), we compute

$$\begin{aligned} (DV(x) - (Ax)')(e_j - e_k) &= \frac{A^{i-k}x}{A_{i-k}^{i-k}} A_{j-k}^{i-k} - A^{j-k}x \\ &= \frac{1}{A_{i-k}^{i-k}} (A_{k-j}^{k-i} A^{i-k} - A_{k-i}^{k-i} A^{j-k})x \end{aligned} \tag{55}$$

$$\begin{aligned} &= \frac{1}{A_{i-k}^{i-k}} (A_{k-j}^{k-i} A^i - A_{k-i}^{k-i} A^j + A_{i-j}^{i-k} A^k)x \\ &= (v^{ki})'x. \end{aligned} \tag{56}$$

Lemma 1 thus yields the following result.

LEMMA 4. *Suppose that the function V is defined by (53) on a region $R^{ik} \subseteq \mathcal{B}^i$ as specified above. Then the HJB equation (44) for V is satisfied at $x \in \text{int}(R^{ik})$ if*

$$(v^{ki})'x \geq 0. \tag{57}$$

By our earlier interpretation of v^{ki} , inequality (57) requires that at state x , mixed strategy x^{ki} is a weakly better response than mixed strategy x^* .

7.3.3 *The geometry of the initial sufficient condition* We now describe the necessary condition (57) from Lemma 4 in geometric terms. In what follows, it is convenient to endow the strategy set $S = \{1, 2, 3\}$ with the cyclic order $1 < 2 < 3 < 1$. When we refer to the strategies generically, as i, j , and k , we require that this labeling satisfy $i < j < k < i$. We give the order a geometric meaning by labeling the vertices of the simplex X counterclockwise, as in Figure 4. If \mathbb{R}^3 is presented in right-handed coordinates, so that the cross-product obeys the right-hand rule, then our labeling of X corresponds to the view from the “outside,” with the origin lying behind the figure and the vector $\mathbf{1}$ pointing toward us.

Also, recalling the definition (35) and alternating property (36) of the skew, we abuse notation by writing $Q = Q^{ijk} = -Q^{kji}$. It follows from the discussion in Section 7.1.1 that the three-strategy games with zero skew are the potential games. Games with $Q > 0$ and

$Q < 0$ are said to have *clockwise skew* and *counterclockwise skew*.³³ Since the sign of the skew can be reversed by renaming the strategies, there is no loss of generality in focusing on games with zero or clockwise skew.

The following properties of the normal vector v^{ki} allow us to locate the states satisfying inequality (57), and hint at the effects of skew on solutions of our optimal control problems. It follows from expression (55) for v^{ki} , or from our interpretation of v^{ki} , that

$$(v^{ki})'x^* = 0, \tag{58}$$

implying that inequality (57) binds at the mixed equilibrium x^* . Moreover, expressions (55) and (56) for v^{ki} , and the fact that $A_{k-i}^{k-i} = A_{k-j}^{k-i} + A_{j-i}^{k-i}$ imply that

$$(v^{ki})'(e_i - e_k) = \frac{1}{A_{k-i}^{k-i}}(A_{k-j}^{k-i}A_{i-k}^{i-k} - A_{k-i}^{k-i}A_{i-k}^{j-k}) = Q \tag{59}$$

$$\begin{aligned} (v^{ki})'(e_k - e_j) &= \frac{1}{A_{k-i}^{k-i}}(A_{k-j}^{k-i}A_{k-j}^i - A_{k-i}^{k-i}A_{k-j}^j + A_{i-j}^{i-k}A_{k-j}^k) \\ &= \frac{1}{A_{k-i}^{k-i}}(A_{k-j}^{k-i}A_{j-k}^{j-i} + A_{j-i}^{k-i}A_{j-k}^{j-k}) > 0 \end{aligned} \tag{60}$$

$$(v^{ki})'(e_i - e_j) = \frac{1}{A_{k-i}^{k-i}}(A_{k-j}^{k-i}A_{i-j}^{i-k} + A_{k-i}^{k-i}A_{j-i}^{j-k}) > 0. \tag{61}$$

We illustrate the consequences of these relations in Figure 5. Figure 5(i) illustrates inequality (57) when $Q = 0$, so that A is a potential game. In this case, (59) says that the line on which (57) binds is parallel to $e_i - e_k$. Thus, by our interpretation of v^{ki} , whether mixed strategy x^{ki} or mixed strategy x^* is a better response to state x depends only on the value of x_j . Inequalities (60) and (61) imply that (57) is satisfied at states e_i and e_k , but not at state e_j , which also agrees with our interpretation of v^{ki} .

Figure 5(ii) illustrates inequality (57) when $Q > 0$, so that A has clockwise skew. In this case, (59) says that the line on which (57) binds is rotated counterclockwise through x^* relative to the unskewed case. Inequality (60) implies that this rotation is less than 60° , so that the line where (57) binds passes through the same sextant as mixed equilibrium x^{ij} . Finally, inequality (61) implies that (57) is satisfied at state e_i , but not at state e_j .³⁴

To complete the initial step of the analysis, let us consider states that are in the region $R^{ik} \subseteq \mathcal{B}^i$ introduced above and that are close to \mathcal{B}^{ik} . States in the latter set can be expressed as $x^* + d\zeta^{ki}$ with $d \geq 0$. Since (52) says that the vector ζ^{ki} is a convex combination of $e_k - e_j$ and $e_i - e_j$, equation (58) and inequalities (60) and (61) imply that

$$(v^{ki})'(x^* + d\zeta^{ki}) = d(v^{ki})'\zeta^{ki} \geq 0,$$

³³For motivation, note that $Q = A^i(e_j - e_k) + A^j(e_k - e_i) + A^k(e_i - e_j)$ represents a composite effect on payoffs of a clockwise circuit of the vertices of X .

³⁴If $Q < 0$, similar logic shows that the rotation of the line where (57) binds is clockwise relative to the unskewed case, and again less than 60° .




FIGURE 5. Skew and inequality (57) in coordination games with the marginal bandwagon property. The vector $\bar{v}^{ki} = \Phi v^{ki}$ is the orthogonal projection of the normal vector v^{ki} onto the tangent space TX .

with a strict inequality when $d > 0$. Thus Lemma 4 implies that at states in R^{ik} close to \mathcal{B}^{ik} , the value function (53) generated by control $e_k - e_i$ satisfies the HJB equation (44).

7.4 Characterization of exit costs

We now turn to the exit cost problem (39), whose solutions describe the expected time until the stochastic evolutionary process $\mathbf{X}^{N,\eta}$ escapes an equilibrium’s strong basin of attraction, as well as the likely point of exit.

To begin, we hypothesize that the optimal feedback control takes the form shown in Figure 6. There the best-response region \mathcal{B}^i is split into two regions; in one, the optimal control is $e_j - e_i$ and exit paths lead to \mathcal{B}^{ij} ; in the other, the optimal control is $e_k - e_i$ and exit paths lead to \mathcal{B}^{ik} . The boundary between the regions is a ray whose endpoint is the mixed equilibrium x^* , and that passes through a state \hat{x}^i determined below. From points on this ray, motion in either basic direction is optimal.

In Appendix A.9, we verify that the optimal feedback control takes this form and we present the corresponding value function. Lemma 7 determines the state \hat{x}^i . This state is uniquely defined by four properties: it lies on the boundary of X ; it places less weight on strategies j and k than does the mixed equilibrium x^* ; it equates the costs of moving in direction $e_j - e_i$ to \mathcal{B}^j and of moving in direction $e_k - e_i$ to \mathcal{B}^k ; and it ensures that the value function derived from the feedback control in Figure 6 satisfies the HJB equation (44). Proposition 12 provides explicit expressions for this feedback control and value function, and states that the latter is indeed the value function for the exit cost problem (39). The proposition is a direct consequence of Lemma 7, Lemma 4, and the verification theorem.




FIGURE 6. Optimal exit paths from \mathcal{B}^i when \hat{x}^i is on face $e_i e_j$.

This analysis yields the solution to exit problem (39). The optimal exit path from state e_i out of basin \mathcal{B}^i proceeds along a face of the simplex through either mixed equilibrium x^{ij} or mixed equilibrium x^{ki} , according to whether state \hat{x}^i lies on face $e_i e_k$ or on face $e_i e_j$; if $\hat{x}^i = e_i$, then both paths are optimal. We therefore have the following proposition.

PROPOSITION 8. *In a simple three-strategy coordination game,*

$$C(\{e_i\}, \mathcal{B}^j \cup \mathcal{B}^k) = \min\{\gamma(e_i, x^{ij}), \gamma(e_i, x^{ki})\} = \min\left(\frac{1}{2} \frac{(A^{i-j} e_i)^2}{A_{i-j}^{i-j}}, \frac{1}{2} \frac{(A^{i-k} e_i)^2}{A_{i-k}^{i-k}}\right).$$

7.5 Characterization of transition costs

In this section, we consider the transition cost problem (40), whose solutions are used to describe the global long-run behavior of the process $\mathbf{X}^{N, \eta}$.

Unlike that of exit costs, the analysis of transition costs depends in a basic way on whether the game at hand is a potential game. To see why, we recall the reasoning from Section 7.3.2, where we sought to define a region in \mathcal{B}^i from which optimal paths to \mathcal{B}^k proceed in direction $e_k - e_i$ to \mathcal{B}^{ik} , generating value function (53) in that region. By Lemma 4, this value function is consistent with the HJB equation (44) whenever $(v^{ki})'x \geq 0$.

Suppose first that A is a potential game, so that the skew Q equals 0. In this case, Figure 5(i) shows that states in \mathcal{B}^i satisfying $x_j \leq x_j^*$, from which motion in direction $e_k - e_i$ leads to \mathcal{B}^{ki} , also satisfy inequality (57). It is therefore consistent with the analysis so far for optimal paths to proceed in direction $e_k - e_i$ to \mathcal{B}^k whenever feasible. We analyze this case in Section 7.5.1.




FIGURE 7. Optimal transition paths to B^k in a potential game ($Q = 0$). In the cross-hatched regions, continuous sets of control directions are optimal.

If instead A has clockwise skew, so that $Q > 0$, Figure 5(ii) shows that the same conclusion about motion from B^i to B^k obtains. However, we cannot reach the analogous conclusion about motion from B^j to B^k . In the thin triangle to the left of x^* , motion in direction $e_k - e_j$ leads to B^{jk} . But since $(v^{jk})'x < 0$ here, this motion is not consistent with the HJB equation (44). Thus the optimal paths to B^k must take a different form, a form we determine in Section 7.5.2.

7.5.1 *Transition costs in potential games* Recall from Section 7.1.1 that the symmetric normal form game A is a potential game if $A = C + \mathbf{1}r'$ for some symmetric matrix $C \in \mathbb{R}^{n \times n}$ and some vector $r \in \mathbb{R}^n$. In this case, the function $f(x) = \frac{1}{2}x'Cx$ is a *potential function* for the population game $F(x) = Ax$ in the sense that $Df(x)z = F(x)z = z'Ax$ for all $z \in TX$ and $x \in X$.³⁵

In potential games, the value function for the transition cost problem (40) is easy to describe and is even smooth, but the optimal feedback controls are of a degenerate form. These controls are illustrated in Figure 7. At states in the sextant northwest of x^* other than those on the ray through x^{ij} , continuous ranges of control vectors are optimal. This degeneracy is particular to potential games, as we explain below.

Proposition 13 in Appendix A.10 provides explicit formulas for the optimal feedback controls and the corresponding value function, and states that this function is indeed the value function for the transition cost problem. The proof is a direct application of the verification theorem. A key step in the argument, Lemma 11, shows that the cost of any path is bounded below by the difference in potential at its initial and terminal points, and that the cost is equal to this difference if only optimal strategies lose mass along the path. Since this is true of all of the controlled trajectories pictured in Figure 7, the value function is entirely determined by such differences in potential.

³⁵See Hofbauer and Sigmund (1988) and Sandholm (2001, 2009).

This analysis provides the solution to the transition cost problem (40) in potential games. As shown in Figure 7, the optimal transition path from e_i to \mathcal{B}^k proceeds directly along the boundary through mixed equilibrium x^{ki} . As noted above, the cost of the path is given by the change in potential.

PROPOSITION 9. *If the simple three-strategy coordination game $A = C + \mathbf{1}r'$ is a potential game, so that $f(x) = \frac{1}{2}x'Cx$ is a potential function for $F(x) = Ax$, then*

$$C(\{e_i\}, \mathcal{B}^k) = \gamma(e_i, x^{ki}) = \frac{1}{2} \frac{(A^{i-k}e_i)^2}{A_{i-k}^{i-k}} = f(e_i) - f(x^{ki}).$$

REMARK 12. Because the integrand of the cost function (38) is piecewise linear in the control $u = \dot{\phi}_t$, it is natural to expect the optimal control vector in $\text{bd}(Z)$ to be unique at almost all states. That this is not true here is a consequence of the integrability properties that define potential games, a point we now consider from two points of view.

First, we noted above that along any controlled trajectory pictured in Figure 7, agents only switch from optimal strategies to suboptimal strategies, so that by Lemma 11, the minimal cost of reaching \mathcal{B}^k from each state x is the change in potential between state x and the terminal state of the controlled trajectory. When there are multiple controlled trajectories between the initial and terminal states, as in the cross-hatched region of Figure 7, each achieves this same minimal cost.³⁶

Second, we argue that A being a potential game is a necessary condition for having a region in \mathcal{B}^i where both $e_j - e_i$ and $e_k - e_i$ are optimal controls. Equation (47) implies that in the interior of such a region, the value function must satisfy

$$DV(x)(e_j - e_i) = A^{j-i}x \quad \text{and} \quad DV(x)(e_k - e_i) = A^{k-i}x.$$

Since $e_j - e_i$ and $e_k - e_i$ span TX , these equalities imply that

$$DV(x)z = z'Ax \quad \text{for all } z \in TX.$$

Thus the second derivative $D^2V(x)$ is given by

$$D^2V(x)(z, \hat{z}) = z' A \hat{z} \quad \text{for all } z, \hat{z} \in TX.$$

The first expression is symmetric in z and \hat{z} , by virtue of being a second derivative. Thus A is symmetric with respect to $TX \times TX$, and so is a potential game.

³⁶To address a possible misconception, let us consider an initial state $x = (1 - c)e_i + ce_j$ with $c \in (0, x_j^*)$. Figure 7 indicates that the optimal path from x to \mathcal{B}^k proceeds in direction $e_k - e_i$ until reaching the state $y \in \mathcal{B}^{ik}$ with $y_j = c$. The argument above shows that this path's cost is $f(x) - f(y)$. One might wonder why there is not a lower cost path that terminates at the mixed equilibrium x^{ki} : since $f(x^{ki})$ is greater than $f(y)$, $f(x) - f(x^{ki})$ is less than $f(x) - f(y)$. But along any path from x that first hits \mathcal{B}^{ik} at x^{ki} , some agents must abandon the suboptimal strategy j . Thus Lemma 11 does not apply and, indeed, the cost of such a path exceeds $f(x) - f(x^{ki})$. The cheapest path from x to x^{ki} goes first from x to y at cost $f(x) - f(y)$, and then from y to x^{ij} at zero cost. Proposition 13 implies that no path can reach x^{ij} more cheaply.




FIGURE 8. Optimal transition paths to \mathcal{B}^k in a clockwise skewed game when \hat{x}^{jk} is on face $e_i e_j$.

7.5.2 *Transition costs in skewed games* We now consider the transition cost problem (40) in games with clockwise skew: $Q > 0$.

It is natural to expect that if the skew Q is small, then the optimal control should resemble the one from the $Q = 0$ case from Figure 7. The previous discussion shows that once Q is positive, no region will have multiple optimal controls. Thus the form of the control in the sextant northwest of x^* must change.

At the start of this section, we argued that in clockwise-skewed games, motion from \mathcal{B}^i to \mathcal{B}^{ki} in direction $e_k - e_i$ is consistent with the HJB equation whenever such a path exists. We therefore hypothesize that motion is in direction $e_k - e_i$ throughout the interior of \mathcal{B}^i , even when such motion leads to boundary \mathcal{B}^{ij} . We also saw that motion from \mathcal{B}^j to \mathcal{B}^{jk} in direction $e_k - e_j$ is not always consistent with the HJB equation. This leads us to hypothesize that in a portion of \mathcal{B}^j close to \mathcal{B}^{ij} , motion will instead be in direction $e_i - e_j$.

The conjectured form of the optimal control is pictured in Figures 8 and 9. In the sextant northwest of x^* , the multiple optimal controls from Figure 7 have been replaced with selections from these controls. The boundary \mathcal{B}^{ij} is approached from states on both sides, but it is approached obliquely from the \mathcal{B}^i side and nearly squarely from the \mathcal{B}^j side. The figures differ in the position of state x^{jk} , determined below, which defines the boundary between the set of states where the feedback control is $e_k - e_j$ and the set where it is $e_i - e_j$.

In Appendix A.11, we verify that the optimal feedback control takes the form shown in Figures 8 and 9, and we describe the corresponding value functions. Lemma 13 determines the state x^{jk} , which is uniquely defined by four properties: it lies on the boundary of the simplex; it places less weight on strategies i and k than does the mixed equilibrium x^* ; it equates the costs of moving in direction $e_k - e_j$ to \mathcal{B}^k and of following the piecewise linear path through x^{ij} to x^* ; and it ensures that in the region below




FIGURE 9. Optimal transition paths to \mathcal{B}^k in a clockwise skewed game when \hat{x}^{jk} is on face $e_j e_k$.

segment $\hat{x}^{jk} x^*$, the value function derived from the proposed feedback control satisfies the HJB equation (44). Proposition 14 provides explicit expressions for the feedback control and value function, and states that the latter is indeed the value function for the transition cost problem. Its proof consists of a lengthy verification of the conditions of Theorem 11.

This analysis implies that in clockwise-skewed games, the possible optimal transition paths depend on the ordering of the strategy pair in question. For a clockwise transition, from e_i to \mathcal{B}^k , the optimal path is always the direct boundary path to x^{ki} . For a counterclockwise transition, from e_j to \mathcal{B}^k , the optimal path is either the direct boundary path to x^{jk} (Figure 8) or a two-segment path that proceeds first to mixed equilibrium x^{ij} and from there to interior equilibrium x^* (Figure 9). Proposition 10 provides a summary.

PROPOSITION 10. *In a simple three-strategy coordination game with clockwise skew,*

$$\begin{aligned}
 C(\{e_i\}, \mathcal{B}^k) &= \gamma(e_i, x^{ki}) = \frac{1}{2} \frac{(A^{i-k} e_i)^2}{A_{i-k}^{i-k}} \\
 C(\{e_j\}, \mathcal{B}^k) &= \min\{\gamma(e_j, x^{jk}), \gamma(e_j, x^{ij}) + \gamma(x^{ij}, x^*)\} \\
 &= \min\left\{ \frac{1}{2} \frac{(A^{j-k} e_j)^2}{A_{j-k}^{j-k}}, \frac{1}{2} \frac{(A^{j-i} e_i)^2}{A_{j-i}^{j-i}} + \frac{1}{2} (x_k^*)^2 (\zeta^{ij})' A \zeta^{ij} \right\}.
 \end{aligned}$$

REMARK 13. It is worth comparing the exit and transition costs for simple three-strategy coordination games under the logit protocol to those under the BRM protocol of Kandori et al. (1993), in which any switch to a suboptimal strategy has unlikelihood 1. Under the latter, the least cost exit path from e_i to $\mathcal{B}^j \cup \mathcal{B}^k$ follows a boundary to either mixed

equilibrium x^{ij} or mixed equilibrium x^{ki} , since these are the states in \mathcal{B}^{ij} and \mathcal{B}^{ki} at which x_i is largest. Thus exit costs under the BRM protocol are

$$C^{\text{BRM}}(\{e_i\}, \mathcal{B}^j \cup \mathcal{B}^k) = \min\{x_j^{ij}, x_k^{ki}\} = \min\left\{\frac{A^{i-j}e_i}{A_{i-j}^{i-j}}, \frac{A^{i-k}e_i}{A_{i-k}^{i-k}}\right\},$$

where the last expressions follow from Lemma 3. The candidate paths are the same as in the logit model, but since the cost of a given path differs in the two models, the identity of the optimal exit path may differ as well.

Turning to the transition problem, recall that since the unlikelihood function of the BRM protocol is discontinuous when multiple strategies are optimal, in violation of assumption (U2), our convergence theorem from Section 5 cannot be applied.³⁷ Nevertheless, results of Kandori and Rob (1998) imply that under the BRM protocol, the optimal path from e_i to \mathcal{B}^ℓ , $\ell \in \{j, k\}$ is the direct boundary path to mixed equilibrium $x^{i\ell}$.³⁸ Thus transition costs are given by

$$C^{\text{BRM}}(\{e_i\}, \mathcal{B}^\ell) = x_\ell^{i\ell} = \frac{A^{i-\ell}e_i}{A_{i-\ell}^{i-\ell}}.$$

In particular, in the games considered here, optimal BRM transition paths never pass through the interior of the simplex, as they may in the logit model.

7.6 Stationary distribution asymptotics and stochastic stability

We now combine results from Section 7.5 with Theorem 10 to draw conclusions about the global behavior of the stochastic evolutionary process $\mathbf{X}^{N,\eta}$ in the small noise double limit. As a first application, we characterize the asymptotic behavior of the stationary distributions $\mu^{N,\eta}$ when A is both a simple three-strategy coordination game and a potential game.

PROPOSITION 11. *Let $A = C + \mathbf{1}r'$ be a simple three-strategy coordination game and a potential game. Let $f(x) = \frac{1}{2}x'Cx$ be a potential function for $F(x) = Ax$, and let $\Delta^+f(x) = \max_{i \in S} f(e_i) - f(x)$. Then*

$$\lim_{N \rightarrow \infty} \lim_{\eta \rightarrow 0} \max_{x \in \mathcal{X}^N} \left| -\frac{\eta}{N} \log \mu^{N,\eta}(x) - \Delta^+f(x) \right| = 0. \tag{62}$$

In words, the proposition says that when N is large, the exponential rate of decay of $\mu^{N,\eta}(x)$ as η approaches 0 is approximately $N\Delta^+f(x)$, where $\Delta^+f(x) \geq 0$ is the deficit in

³⁷This was not an issue for the exit problem, since within a single best-response region the BRM protocol's unlikelihood function is constant.

³⁸For a proof, observe first that in simple three-strategy coordination games, $x_i^{i\ell} \geq x_i^*$ (see Figure 4). The previous paragraph showed that the direct boundary path from e_i to \mathcal{B}^ℓ is optimal among those that do not enter \mathcal{B}^h , $h \notin \{i, \ell\}$. This path's cost is $x_\ell^{i\ell} = 1 - x_i^{i\ell} \leq 1 - x_i^* = x_\ell^* + x_h^*$. But any transition path that enters \mathcal{B}^h must have at least this cost, since reaching \mathcal{B}^h entails a cost of at least $x_h^{ih} \geq x_h^*$, due to switches from i to h , plus a cost of at least x_ℓ^* , due to switches from either i or h to ℓ .

potential of state x relative to the maximizers of potential. Thus the latter states are the stochastically stable states in the small noise double limit.

PROOF OF PROPOSITION 11. We abuse notation in what follows by identifying singleton sets with their lone elements (e.g., by writing $C(e_j, e_i)$ in place of $C(\{e_j\}, \{e_i\})$).

We start by finding the minimal cost $R(e_i)$ of an e_i -tree. Since the three e_i -trees are $\{(e_j, e_i), (e_k, e_i)\}$, $\{(e_k, e_j), (e_j, e_i)\}$, and $\{(e_j, e_k), (e_k, e_i)\}$, **Proposition 9** implies that

$$\begin{aligned} R(e_i) &= \min\{C(e_j, e_i) + C(e_k, e_i), C(e_k, e_j) + C(e_j, e_i), C(e_j, e_k) + C(e_k, e_i)\} \\ &= \min\{(f(e_j) - f(x^{ij})) + (f(e_k) - f(x^{ki})), \\ &\quad (f(e_k) - f(x^{jk})) + (f(e_j) - f(x^{ij})), \\ &\quad (f(e_j) - f(x^{jk})) + (f(e_k) - f(x^{ki}))\} \\ &= f(e_j) + f(e_k) - \max\{f(x^{ij}) + f(x^{ki}), f(x^{jk}) + f(x^{ij}), f(x^{jk}) + f(x^{ki})\} \\ &= -f(e_i) + (f(e_i) + f(e_j) + f(e_k) \\ &\quad - \max\{f(x^{ij}) + f(x^{ki}), f(x^{jk}) + f(x^{ij}), f(x^{jk}) + f(x^{ki})\}). \end{aligned}$$

In the final expression, the term in parentheses, henceforth denoted K , does not depend on the choice of e_i .

Next, it follows from **Lemma 11** in **Appendix A.10** that for any $x \in X$,

$$-f(e_i) + C(e_i, x) \geq -f(e_i) + (f(e_i) - f(x)) = -f(x). \tag{63}$$

If $x \in \mathcal{B}^i$, then along the straight-line path from e_i to x , only the optimal strategy i loses mass, so **Lemma 11** implies that the inequality in (63) binds.

Combining these facts yields

$$r(x) \equiv \min_{i \in S} (R(e_i) + C(e_i, x)) = -f(x) + K.$$

Since A is a coordination game, the potential function f is maximized at a pure state, so

$$\Delta r(x) \equiv r(x) - \min_{y \in X} r(y) = -f(x) - \min_{y \in X} (-f(y)) = -f(x) + \max_{i \in S} f(e_i) = \Delta^+ f(x).$$

The proposition thus follows from **Theorem 10**. □

The close connection between stationary distributions and potential functions in potential games has been understood since the work of **Blume (1993, 1997)**. Building on Blume’s work, **Sandholm (2010c, Corollary 12.2.5)** derives statement (62) for a particular specification of the process $\mathbf{X}^{N, \eta}$. In this specification, not only the limit game F , but also all of the finite-population games F^N are assumed to be potential games. This definition ensures that $\mathbf{X}^{N, \eta}$ is reversible for each (N, η) pair, and so that each stationary

distribution $\mu^{N,\eta}$ admits a simple closed form.³⁹ Equation (62) is obtained by taking the limit of these explicit formulas.

In the present analysis, we only assume that the finite-population games F^N converge to a limiting potential game F .⁴⁰ This assumption does not require $\mathbf{X}^{N,\eta}$ to be reversible, and so explicit expressions for $\mu^{N,\eta}$ are generally unavailable. We describe the asymptotics of the stationary distribution under this weaker assumption by way of the large deviations properties of the stochastic processes. Doing so provides intuition about the forces behind the selection of the potential maximizer. Since transition costs are determined by differences in potential, the transitions used in every minimum cost tree pass through the same mixed equilibria, so that differences in the trees' costs are due to differences in potential at the trees' roots.

The next example provides explicit computations of stochastically stable states under the logit protocol, and compares these predictions with those under the BRM protocol.

EXAMPLE 14. Consider the game $F(x) = Ax$ with

$$A = \begin{pmatrix} 7 & 0 & 0 \\ 2 - q & 6 & 0 \\ 2 & 0 & 5 \end{pmatrix},$$

where $q \in [0, 5)$. For each such q , A is a simple coordination game⁴¹ with interior equilibrium $x^* = (6/(17+q), (5+q)/(17+q), 6/(17+q))$. The mixed equilibria on the boundary of X are $x^{12} = (6/(11+q), (5+q)/(11+q), 0)$, $x^{23} = (0, \frac{5}{11}, \frac{6}{11})$, and $x^{31} = (\frac{1}{2}, 0, \frac{1}{2})$. The parameter q is the skew of A . Thus when $q = 0$, A is a potential game.⁴²

To evaluate stochastic stability, we compute the costs of the direct paths from each pure state to the two adjacent mixed equilibria on the boundary of X , as well as the costs of the direct paths from the boundary mixed equilibria to the interior equilibrium x^* . We present these path costs in Figure 10(i).

Next, when q is positive, we determine whether the optimal path for each counter-clockwise transition from e_j to \mathcal{B}^k is the direct path to x^{jk} or the two-segment path via x^{ij} to x^* (see Proposition 10 and Figures 8 and 9). In the present example, the boundary paths are optimal for every $q \in (0, 5)$. Proposition 9 implies that they are also optimal when $q = 0$.

³⁹See Sandholm (2010c, Theorem 11.5.12).

⁴⁰Finite-population potential games are defined by equalities relating benefits from unilateral deviations to changes in potential, and so are nongeneric. Thus, a typical sequence of games F^N that converges to a limiting potential game F will not itself consist of potential games.

⁴¹For the marginal bandwagon property (34), note that $A_{3-1}^{3-2} = 5 - q$.

⁴²In this case, A admits the decomposition $A = C + \mathbf{1}r'$ with

$$C = \begin{pmatrix} 5 & 0 & 0 \\ 0 & 6 & 0 \\ 0 & 0 & 5 \end{pmatrix} \quad \text{and} \quad r' = (2 \ 0 \ 0).$$




FIGURE 10. The path costs needed to determine transition costs in Example 14.

We then determine the minimum cost $R(e_i)$ of an e_i tree for $i \in \{1, 2, 3\}$. Simple calculations show that

$$R(e_1) = C(e_3, e_2) + C(e_2, e_1) = \frac{25}{22} + \frac{18}{11+q}$$

$$R(e_2) = \begin{cases} C(e_3, e_2) + C(e_1, e_2) = \frac{25}{22} + \frac{(5+q)^2}{2(11+q)} & \text{if } q \leq \frac{1}{4}(-15 + \sqrt{265}) \approx 0.3197 \\ C(e_3, e_2) + C(e_1, e_3) = \frac{25}{22} + \frac{25}{20} & \text{otherwise} \end{cases}$$

$$R(e_3) = \begin{cases} C(e_1, e_2) + C(e_2, e_3) = \frac{(5+q)^2}{2(11+q)} + \frac{18}{11} & \text{if } q \leq \frac{5}{22} \\ C(e_2, e_1) + C(e_1, e_3) = \frac{18}{11+q} + \frac{25}{20} & \text{otherwise.} \end{cases}$$

Further calculations show that $R(e_2)$ is smallest when $q \in [0, \frac{17}{5}]$ and that $R(e_1)$ is smallest when $q \in [\frac{17}{5}, 5)$. Therefore, Theorem 10 implies that under the logit protocol, state e_2 is stochastically stable in the small noise double limit in the former case and state e_1 is in the latter; both are stochastically stable when $q = \frac{17}{5}$.

We now compare these selection results to those obtained under the BRM protocol.⁴³ Remark 13 states that under this protocol, optimal transition paths in simple coordination games are direct. The BRM costs of the six relevant paths can be read directly from the coordinates of the boundary equilibria; they are shown in Figure 10(ii).

⁴³We consider a version of the BRM protocol under which all optimal strategies are chosen with nonnegligible probability. Since the convergence results in Section 5 do not apply to the BRM model, we cannot appeal to them here. But in the present example, the intermediate results needed to establish stochastic stability follow from elementary considerations, provided that the minimal cost tree is unique. Compare Kandori and Rob (1995, 1998).

Calculations show that the minimal tree costs are

$$\begin{aligned}
 R^{\text{BRM}}(e_1) &= C^{\text{BRM}}(e_3, e_2) + C^{\text{BRM}}(e_2, e_1) = \frac{5}{11} + \frac{6}{11+q} \\
 R^{\text{BRM}}(e_2) &= \begin{cases} C^{\text{BRM}}(e_3, e_2) + C^{\text{BRM}}(e_1, e_2) = \frac{5}{11} + \frac{5+q}{11+q} & \text{if } q \leq 1 \\ C^{\text{BRM}}(e_3, e_2) + C^{\text{BRM}}(e_1, e_3) = \frac{5}{11} + \frac{5}{10} & \text{otherwise} \end{cases} \\
 R^{\text{BRM}}(e_3) &= \begin{cases} C^{\text{BRM}}(e_1, e_2) + C^{\text{BRM}}(e_2, e_3) = \frac{5+q}{11+q} + \frac{6}{11} & \text{if } q \leq \frac{11}{23} \\ C^{\text{BRM}}(e_2, e_1) + C^{\text{BRM}}(e_1, e_3) = \frac{6}{11+q} + \frac{5}{10} & \text{otherwise.} \end{cases}
 \end{aligned}$$

Finding the smallest of these costs, we conclude that under the BRM protocol, state e_2 is stochastically stable when $q \in [0, 1)$ and state e_1 is stochastically stable when $q \in (1, 5)$.

To compare predictions under the two protocols, it is useful to focus on the minimal cost trees themselves. Under the logit protocol, three trees have minimal cost for some $q \in [0, 5)$: the e_2 -tree $\{(e_3, e_2), (e_1, e_2)\}$ for $q \in [0, \hat{q})$, $\hat{q} \approx 0.3197$; the e_2 -tree $\{(e_3, e_2), (e_1, e_3)\}$ for $q \in [\hat{q}, \frac{17}{5})$; and the e_1 -tree $\{(e_3, e_2), (e_2, e_1)\}$ for $q \in [\frac{17}{5}, 5)$. Under the BRM protocol, only the first and last of these have minimal costs, according to whether $q \in [0, 1)$ or $q \in [1, 5)$. By way of explanation, notice that as q increases, so does the payoff disadvantage $7 - (2 - q) = 5 + q$ of strategy 2 at state e_1 . This causes the cost of the (e_1, e_2) transition to grow more rapidly under logit than under BRM, so that the optimal logit e_2 -tree abandons this transition earlier than the optimal BRM e_2 -tree.

Under both protocols, the stochastically stable state switches from equilibrium e_2 to efficient equilibrium e_1 as q increases. But the switch occurs sooner for BRM: for $q \in (1, \frac{17}{5})$, BRM selects e_1 , while logit selects e_2 . Under BRM, the selection switches once strategy 1 begins to pairwise risk dominate strategy 2. This would follow from classic results in the absence of strategy 3, and the fact that transition (e_3, e_2) , which heads away from e_3 , appears in all BRM minimal cost trees ensures that strategy 3's presence does not affect the selection. In contrast, as noted above, transition (e_1, e_3) , which heads into e_3 , is in the logit minimal cost tree for intermediate values of q . Its appearance there reflects the advantage of the indirect route from e_1 to e_2 via e_3 over the direct route, and explains why strategy 2 persists in being stochastically stable despite being pairwise risk dominated by strategy 1.

As q increases through $\frac{17}{5}$, the logit minimal cost tree replaces transition (e_1, e_3) with transition (e_2, e_1) , changing the stochastically stable state from e_2 to e_1 . The former transition must overcome an initial payoff disadvantage of $7 - 2 = 5$, compared to $6 - 0 = 6$ for the latter, leading the former to be less costly at low values of q . As q increases, mixed equilibrium x^{12} moves closer to state e_2 , causing the payoff advantage of strategy 2 over strategy 1 to dissipate more quickly as the state moves from e_2 toward e_1 . This reduces the cost of the (e_2, e_1) transition under the logit protocol, and leads to the replacement of e_2 by e_1 as the stochastically stable state. ◇

8. DISCUSSION

8.1 *Orders of limits and waiting times*

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 \(2010b\)](#) 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, ones that consider sequences of Markov processes that run on increasingly fine state spaces ([Dupuis 1988](#), [Dupuis and Ellis 1997](#)). For recent progress in this direction, see [Sandholm and Staudigl \(2015\)](#).

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 \(2010b\)](#) 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.

Stochastic stability models have been subject to the criticism that the amount of time required for their predictions to become relevant is too long for most economic applications. Since here we are taking multiple limits, this criticism holds additional force. To better understand the relevance of our analysis to applications, one could assess the extent to which versions of the model's predictions are correct when the noise level is not too small and the population size not too large. This could certainly be done numerically; whether analytical results along these lines can be established is a challenging open question.

8.2 *Analyzing other protocols and classes of games*

This paper characterized the long-run behavior of a class of stochastic evolutionary processes in the small noise double limit. Our explicit calculations in [Section 7](#) focused on evolution in simple three-strategy coordination games under the logit protocol. We conclude by discussing the prospects for extending our analysis to other games and choice rules.

To evaluate these prospects, recall that the running cost appearing in the path cost integral (24) is $L(x, u) = [u]_+^{\prime} Y(F(x))$, where Y is the unlikelihood function (5) of the revision protocol. The piecewise linearity of L in the control $u = \dot{\phi}_t$ ensures that at each state x , the optimal choices of u in the HJB equation (42) include extreme points of the control set $Z = \text{conv}(\{e_i - e_j : i, j \in S\})$. Thus for any game and revision protocol, we expect optimal feedback controls for the exit and transition problems (39) and (40) to partition the state space into regions in which the various basic directions $e_i - e_j$ are followed.

The logit protocol (6) is particularly convenient because its unlikelihood function (37) is piecewise linear in the payoff vector and, thus, is piecewise linear in the state when the limit payoff function $F(x) = Ax$ is linear. This leads the value functions for problems (39) and (40) to be piecewise quadratic; in particular, they are homogeneous of degree 2 in the displacement of the state from an interior equilibrium x^* . This ensures that the optimal feedback controls partition the state space into convex sets with common extreme point x^* , as shown in Figures 6–9. This structure should be preserved by certain other revision protocols. Under the probit protocol (Example 5), the unlikelihood function is piecewise quadratic. This should lead to value functions that are piecewise cubic—specifically, homogeneous of degree 3 in the displacement of the state from x^* —so that in the class of games studied here, the boundaries between control regions are again rays emanating from x^* .

Returning to the logit protocol, the piecewise linearity of running costs $L(x, u)$ in both the control u and the state x suggests that the exit and transition problems can be solved beyond the class of simple three-strategy coordination games studied here. The main new consideration in solving control problems (39) and (40) for general linear games is that the state constraints, which require controlled trajectories to stay in the state space X , may bind. The fact that these constraints are slack in the games studied here allowed us to appeal to a verification theorem, Theorem 11, that does not include such constraints. To handle more general cases, one would need to extend the verification theorem to allow for linear state constraints. For the class of problems generated by the logit protocol, we see no conceptual difficulty in obtaining this extension. Still, the proof of Theorem 11 is not simple, and extending it to incorporate state constraints is a challenge we leave for future research.

APPENDIX

A.1 Statement and proof of Lemma 5

The analysis of Example 7 requires the following lemma.

LEMMA 5. *Let F^N be a finite-population game defined by random matching in normal form coordination game A . Let $x \in \mathcal{X}_j^N$ 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 $(1/N)(e_{i_1} - e_j)$ 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),

$$\frac{N}{N-1}(e_{i^c} - e_k)'Ax^{c-1} - \frac{1}{N-1}(e_{i^c} - e_k)'Ae_j \geq 0 \quad \text{for all } k \in S. \tag{64}$$

By construction,

$$x^c = x^{c-1} + d^c(e_{i^c} - e_j) \quad \text{for some } d^c > 0. \tag{65}$$

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

$$\frac{N}{N-1}(e_{i^{c+1}} - e_k)'Ax^c - \frac{1}{N-1}(e_{i^{c+1}} - e_k)'Ae_j \geq 0 \quad \text{for all } k \in S, \tag{66}$$

and since $i^c \notin b_j^N(x^c)$, the inequality in (66) is strict when $k = i^c$. Combining (64) (with $k = i^{c+1}$) and the strict version of (66) (with $k = i^c$) with (65) yields

$$(e_{i^c} - e_{i^{c+1}})'A(e_{i^c} - e_j) < 0. \tag{67}$$

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 non-repeating 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

$$\sum_{c=1}^{C-1} (e_{i^c} - e_{i^{c+1}})'Ae_j = 0. \tag{68}$$

Summing (67) over $c \in \{1, \dots, C-1\}$ and substituting (68) yields

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

again contradicting that A is a coordination game.

A.2 Proof of Proposition 2

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

$$|\pi - \hat{\pi}| < \delta \text{ implies that } |Y_j(\pi) - Y_j(\hat{\pi})| < \varepsilon \quad \text{for all } \pi, \hat{\pi} \in F(X) \text{ and } j \in S. \tag{69}$$

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

$$N \geq N_0 \text{ implies that } |F_{i \rightarrow \cdot}^N(x) - F(x)| < \delta \text{ for all } x \in X_i^N \text{ and } i \in S \tag{70}$$

$$N \geq N_0 \text{ implies that } |F(\phi_t^{(N)}) - F(\phi_s^{(N)})| < \delta \text{ whenever } |t - s| \leq \frac{1}{N}. \tag{71}$$

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

$$\begin{aligned} \frac{1}{N}c^N(\phi^N) &= \frac{1}{N} \sum_{k=0}^{\ell^N-1} \langle Y(F_{i^N(k) \rightarrow \cdot}^N(\phi_k^N)), [\dot{\phi}_k^N]_+ \rangle \\ &= \frac{1}{N} \sum_{k=0}^{\ell^N-1} \langle Y(F(\phi_{k/N}^{(N)})), [\dot{\phi}_k^{(N)}]_+ \rangle + \alpha^N \frac{\ell^N}{N} \\ &= \int_0^{T^N} \langle Y(F(\phi_{\lfloor Nt \rfloor / N}^{(N)})), [\dot{\phi}_t^{(N)}]_+ \rangle dt + \alpha^N T^N \\ &= \int_0^{T^N} \langle Y(F(\phi_t^{(N)})), [\dot{\phi}_t^{(N)}]_+ \rangle dt + (\alpha^N + \beta^N) T^N \\ &= c(\phi^{(N)}) + (\alpha^N + \beta^N) T^N. \end{aligned} \tag{72}$$

The first equality is (23), the second follows from (69), (70), and (19), the third follows from (19) and (20), and the fourth follows from (69), (71), and (21). Since $\varepsilon > 0$ was chosen arbitrarily and the T^N are bounded, the proposition follows.

A.3 Proof of Proposition 3

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 $(1/N)c^N(\phi^N)$. There is a subsequence along which $(1/N)c^N(\phi^N)$ 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 $Y \circ F$ on the compact set X , we have that

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

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

Now fix $\varepsilon > 0$. The previous argument and (72) imply that for all N large enough, we have

$$c(\phi^{[N]}) \leq c(\phi^{(N)}) + \frac{2bd}{N} \leq \frac{1}{N}c^N(\phi^N) + 2\varepsilon\bar{T} + \frac{2bd}{N}. \tag{73}$$

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

A.4 Proof of Proposition 4

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

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

Summing these inequalities yields inequality (29).

Because $\phi = \{\phi_t\}_{t \in [0, T]}$ is monotone and moves at full speed, and since $1/N \leq T$, there is a time $s^N \in [0, 1/N)$ at which

$$\sum_{j \in S_+} \phi_{s^N,j} \in \frac{1}{N}\mathbb{Z} \quad \text{and, hence,} \quad \sum_{i \in S_-} \phi_{s^N,i} \in \frac{1}{N}\mathbb{Z}. \tag{74}$$

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. This assumption and (74) imply that there is a $\phi_0^N \in \mathcal{X}^N$ such that

$$\sum_{j \in S_+} \phi_{0,j} = \sum_{j \in S_+} \phi_{0,j}^N, \quad \sum_{i \in S_-} \phi_{0,i} = \sum_{i \in S_-} \phi_{0,i}^N \tag{75}$$

$$\sum_{i \in S} |\phi_{0,i} - \phi_{0,i}^N| < \frac{2}{N}. \tag{76}$$

Inequality (76) follows from the fact that every point in the simplex in \mathbb{R}^n is within ℓ^1 distance $2(n-1)/n$ of some vertex.

This inequality is the base of our inductive argument. To write the inductive step, let $x = \phi_{k/N}$, $y = \phi_{(k+1)/N}$, and $x = \phi_k^N$ be given, with $y = \phi_{k+1}^N$ to be determined. The inductive step says that if

$$\sum_{j \in S_+} |x_j - x_j| \leq \frac{2n_+}{N} \quad \text{and} \quad \sum_{i \in S_-} |x_i - x_i| \leq \frac{2n_-}{N},$$

then we can choose $y = x + (1/N)(e_{j^*} - e_{i^*})$ with $j^* \in S_+$ and $i^* \in S_-$ so that

$$\sum_{j \in S_+} |y_j - y_j| \leq \frac{2n_+}{N} \quad \text{and} \quad \sum_{i \in S_-} |y_i - y_i| \leq \frac{2n_-}{N}.$$

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

Our proof of the inductive step focuses on the claim for strategies in S_+ ; the proof of the claim for strategies in S_- is nearly identical. Monotonicity and the fact that ϕ moves at full speed imply that

$$y - x = \frac{1}{N}z \quad \text{for some } z \in Z \text{ with } z_j \geq 0 \text{ for } j \in S_+ \text{ and } \sum_{j \in S_+} z_j = 1. \tag{77}$$

Since $y = x + (1/N)(e_{j^*} - e_{i^*})$ for some $j^* \in S_+$ and $i^* \in S_-$, it follows that

$$\sum_{j \in S_+} |y_j - y_j| = \sum_{j \in S_+} \left| x_j - x_j + \frac{1}{N}(z_j - 1_{j=j^*}) \right| \leq \sum_{j \in S_+} |x_j - x_j| + \frac{2}{N}.$$

This establishes the claim for cases where $\sum_{j \in S_+} |x_j - x_j| \leq 2(n_+ - 1)/N$. The claim for the complementary case is a consequence of the following lemma.

LEMMA 6. *If*

$$\sum_{j \in S_+} |x_j - x_j| \geq \frac{2(n_+ - 1)}{N}, \tag{78}$$

then $y = x + (1/N)(e_{j^} - e_{i^*})$ can be chosen so that $\sum_{j \in S_+} |y_j - y_j| \leq \sum_{j \in S_+} |x_j - x_j|$.*

PROOF. Recall from (75) 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

$$0 = \sum_{j \in S_+} (x_j - x_j) = \sum_{j \in S_+} [x_j - x_j]_+ - \sum_{j \in S_+} [x_j - x_j]_-. \tag{79}$$

It follows that there are at most $n_+ - 1$ strategies $j \in S_+$ for which $x_j - x_j > 0$. Therefore, since (78) and (79) imply that

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

there is a strategy $j^* \in S_+$ with

$$x_{j^*} - x_{j^*} \geq \frac{1}{N}. \tag{80}$$

Since $y = x + (1/N)(e_{j^*} - e_{i^*})$ by definition, it follows from (80) and (77) that

$$y_{j^*} - y_{j^*} = x_{j^*} - x_{j^*} + \frac{1}{N}(z_{j^*} - 1) \geq 0.$$

This inequality and (77) yield

$$\begin{aligned}
 \sum_{j \in S_+} |y_j - y_j| &= \sum_{j \in S_+} [y_j - y_j]_+ + \sum_{j \in S_+} [y_j - y_j]_- \\
 &= \sum_{j \in S_+} \left[x_j - x_j + \frac{1}{N}(z_j - 1_{j=j^*}) \right]_+ + \sum_{j \in S_+} \left[x_j - x_j + \frac{1}{N}z_j \right]_- \\
 &\leq \sum_{j \in S_+} \left([x_j - x_j]_+ + \frac{1}{N}(z_j - 1_{j=j^*}) \right) + \sum_{j \in S_+} [x_j - x_j]_- \\
 &= \sum_{j \in S_+} [x_j - x_j]_+ + \sum_{j \in S_+} [x_j - x_j]_- \\
 &= \sum_{j \in S_+} |x_j - x_j|.
 \end{aligned}$$

□

This completes the proof of Proposition 4.

A.5 Proof of Proposition 5

Fix $\varepsilon > 0$. Choose $\delta > 0$ so that (69) holds, and then choose N_0 so that (70) and

$$N \geq N_0 \text{ implies that } |F(\phi_{k/N}) - F(\phi_k^N)| < \delta \text{ for all } k \in \{0, \dots, \lfloor NT \rfloor\} \tag{81}$$

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 4, we minimize notation by taking s^N to equal 0.

By the triangle inequality,

$$\begin{aligned}
 |Y_j(F_{i \rightarrow \cdot}^N(\phi_k^N)) - Y_j(F(\phi_{k/N}))| \\
 \leq |Y_j(F_{i \rightarrow \cdot}^N(\phi_k^N)) - Y_j(F(\phi_k^N))| + |Y_j(F(\phi_k^N)) - Y_j(F(\phi_{k/N}))|.
 \end{aligned}$$

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

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

The first equality here is (30), and the second follows from (69), (70), (81), and (82).

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

$$\lim_{N \rightarrow \infty} M^N = \infty \quad \text{and} \quad \lim_{N \rightarrow \infty} \frac{M^N}{N} = 0. \tag{84}$$

Also, choose $\tau > 0$ so that

$$|t - s| \leq 2\tau \text{ implies that } |F(\phi_t) - F(\phi_s)| < \delta. \tag{85}$$

Then continuing from (83), 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 $Y \circ F$ on X such that

$$\begin{aligned} \frac{1}{N}c^N(\phi^N) &= \sum_{m=0}^{M^N-1} \sum_{k=0}^{L^N-1} \langle Y(F(\phi_{mL^N})), [\phi_{mL^N+k+1}^N - \phi_{mL^N+k}^N]_+ \rangle + (2\alpha^N + \beta^N)T^N \\ &= \sum_{m=0}^{M^N-1} \langle Y(F(\phi_{mL^N})), [\phi_{(m+1)L^N}^N - \phi_{mL^N}^N]_+ \rangle + (2\alpha^N + \beta^N)T^N \\ &= \sum_{m=0}^{M^N-1} \langle Y(F(\phi_{mL^N})), [\phi_{(m+1)L^N/N} - \phi_{mL^N/N}]_+ \rangle + (2\alpha^N + \beta^N)T^N + \frac{bT^N}{\sqrt{N}}. \end{aligned} \tag{86}$$

The first equality uses (85) and (69), the second uses the monotonicity of ϕ^N , and the third uses the boundedness of $Y \circ F$ on X and the $O(1/N)$ convergence of ϕ^N to ϕ specified in (29).

The limits in (84) and the monotonicity of ϕ imply that as N approaches infinity, and the Riemann–Stieltjes sum in (86) converges to a Riemann integral. (To be more precise, writing the inner product in the initial term of (86) 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 ,

$$\begin{aligned} \frac{1}{N}c^N(\phi^N) &= \int_0^T \langle Y(F(\phi_t)), [\dot{\phi}_t]_+ \rangle dt + (2\alpha^N + \beta^N + \gamma^N)T^N + \frac{bT^N}{\sqrt{N}} \\ &= c(\phi) + (2\alpha^N + \beta^N + \gamma^N)T^N + \frac{bT^N}{\sqrt{N}}. \end{aligned} \tag{87}$$

Since $T^N \leq T$ (see the statement of Proposition 4), the last summand vanishes as N grows large. Thus since ε was arbitrary, we conclude that $\lim_{N \rightarrow \infty} (1/N)c^N(\phi^N) = c(\phi)$.

A.6 Proof of Proposition 6

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). 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 4; the length of this subpath is $\ell^{N,m} = \lfloor N(T_m - T_{m-1} - s^{N,m}) \rfloor$, where $s^{N,m} \in [0, 1/N)$ too is from Proposition 4.

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 $T_1 - \hat{s}^{N,1} = s^{N,1} + \ell^{N,1}/N$. Then $\hat{s}^{N,1} \in [0, 1/N)$, and we can bound the distance between $\phi_{\ell^{N,1}}^{N,1}$ and end at node $\phi_0^{N,2}$ as

$$\begin{aligned} |\phi_{\ell^{N,1}}^{N,1} - \phi_0^{N,2}| &\leq |\phi_{\ell^{N,1}}^{N,1} - \phi_{T_1 - \hat{s}^{N,1}}| + |\phi_{T_1 - \hat{s}^{N,1}} - \phi_{T_1 + s^{N,2}}| + |\phi_{T_1 + s^{N,2}} - \phi_0^{N,2}| \\ &\leq \frac{2n}{N} + \frac{4}{N} + \frac{2n}{N}. \end{aligned} \tag{88}$$

The bounds on the first and third terms are from Proposition 4, and the bound on the second term follow from the fact that $\hat{s}^{N,1}$ and $s^{N,2}$ are less than $1/N$ and the full speed requirement (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 $|\phi_0^{N,0} - \phi_0| \leq d/N$. We therefore have

$$|x^N - \phi_0^{N,1}| \leq |x_0^N - \phi_0| + |\phi_0 - \phi_{s^{N,1}}| + |\phi_{s^{N,1}} - \phi_0^{N,1}| \leq \frac{d}{N} + \frac{2}{N} + \frac{2n}{N}. \tag{89}$$

The bound on the second term follows from the fact that $s^{N,1} \leq 1/N$ and from the full speed requirement (28), and the bound on the third term follows from Proposition 4.

Observe that given any distinct $x, y \in \mathcal{X}^N$, there is a state \hat{x} adjacent to x such that $|\hat{x} - y| = |x - y| - 2/N$. These observations and inequalities (88) and (89) 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 $Y \circ F$ 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

$$\sum_{m=0}^M c^N(\psi^{N,m}) \leq 2b((M - 1)(2n + 2) + (d + 2 + 2n)) \leq 2b(3nM + d). \tag{90}$$

Since for each N , the total duration of subpaths $\phi^{N,1}, \dots, \phi^{N,M}$ is less than T , inequalities (87) and (90) imply that for all N large enough,

$$\frac{1}{N} c^N(\phi^N) \leq c(\phi) + 4\varepsilon T + \frac{bT}{\sqrt{N}} + \frac{2b(3nM + d)}{N}. \tag{91}$$

Since ε was arbitrary, it follows that

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

and, thus, that

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

A.7 Proof of Theorem 10

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 7, it is enough to show that for all large enough N ,

$$\max_{x \in \mathcal{X}^N} \left| \frac{1}{N} \Delta r^N(x) - \Delta r(x) \right| < \varepsilon. \tag{92}$$

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

$$\max_{x \in \mathcal{X}^N} \left| \frac{1}{N} r^N(x) - r(x) \right| < \varepsilon, \tag{93}$$

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 (92).

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

$$r^N(x) = \min_{K^N \in \mathcal{K}^N} \left(\min_{\tau_{K^N} \in \mathcal{T}_{K^N}} \sum_{(L^N, \hat{L}^N) \in \tau_{K^N}} C^N(L^N, \hat{L}^N) + C^N(K^N, \{x\}) \right), \tag{94}$$

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 (94): 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 9 is uniform: there is an N_0 such that for all $N \geq N_0$ and all choices of recurrent classes,

$$\begin{aligned} \left| \frac{1}{N} C^N(L^N, \hat{L}^N) - C(L, \hat{L}) \right| &< \frac{\varepsilon}{\kappa} \\ \left| \frac{1}{N} C^N(K^N, \{x\}) - C(K, \{x\}) \right| &< \frac{\varepsilon}{\kappa} \quad \text{if } x \in \mathcal{X}^N. \end{aligned} \tag{95}$$

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

To establish (93), 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 $(1/N)C^N(K^N, \{x\})$ to $C(K, \{x\})$ were uniform in x . To see that this is so, note that by inequalities (73) and (91), the choice of N_0 needed to ensure inequality (95) 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 $Y \circ F$ on X), \tilde{M} and \tilde{T} (since Assumption 2 requires that $M = M(x)$ and $T = T(x)$ from inequality (91) 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 depends on x , we can indeed choose N_0 so that (95) holds for all $N \geq N_0$ and for all $x \in \mathcal{X}^N$ simultaneously. This establishes (93) and so completes the proof of the theorem.

A.8 Proof of Lemma 1

We start by deriving (47). Since V is constructed from a feedback control that equals $e_k - e_i$ in a neighborhood of $x \in \text{int}(\mathcal{B}^i)$, Lemma 2 implies that

$$V(x + t(e_i - e_k)) - V(x) = \gamma(x + t(e_i - e_k), x) = tA^{i-k}x + \frac{1}{2}t^2 A_{i-k}^{i-k}$$

for t close to 0. Thus

$$DV(x)(e_k - e_i) = -\left(\frac{d}{dt}\gamma(x + t(e_i - e_k), x)\right)\Big|_{t=0} = A^{k-i}x,$$

which is equivalent to (47).

To verify the HJB equation (44), we must show that the function to be minimized,

$$H(e_a, e_b) = (e_i - e_a)'Ax + DV(x)'(e_a - e_b),$$

is nonnegative at each of the five choices of (e_a, e_b) other than (e_k, e_i) , and, indeed,

$$H(e_i, e_h) = DV(x)(e_i - e_h) \geq 0 \quad \text{for } h \in \{j, k\} \text{ by (45)}$$

$$H(e_j, e_i) = (DV(x) - (Ax)')(e_j - e_i) = (DV(x) - (Ax)')(e_j - e_k) \geq 0 \quad \text{by (47) and (46)}$$

$$H(e_k, e_j) = (e_i - e_k)'Ax + DV(x)(e_k - e_j) = DV(x)(e_i - e_j) \geq 0 \quad \text{by (47) and (45)}$$

$$H(e_j, e_k) = (e_i - e_j)'Ax + DV(x)(e_j - e_k) \geq (e_i - e_k)'Ax \geq 0 \quad \text{by (46) and the fact that } x \in \mathcal{B}^i.$$

A.9 The value function for the exit cost problem

We begin by determining the state \hat{x}^i that defines the boundary between the two control regions pictured in Figure 6. To start, we define

$$V^j(x) = \frac{1}{2} \frac{(A^{i-j}x)^2}{A_{i-j}^{i-j}} \quad \text{and} \quad V^k(x) = \frac{1}{2} \frac{(A^{i-k}x)^2}{A_{i-k}^{i-k}}.$$

By Lemma 3, $V^j(x)$ is the cost of a path from state x that moves through \mathcal{B}^i in direction $e_j - e_i$ until reaching boundary \mathcal{B}^{ij} ; $V^k(x)$ is interpreted analogously.




FIGURE 11. The construction of \hat{x}^i when it is on face $e_k e_i$ and $Q > 0$.

LEMMA 7. *There is a unique state $\hat{x}^i \in \mathcal{B}^i \cap \text{bd}(X)$ such that*

$$\hat{x}_j^i < x_j^* \quad \text{and} \quad \hat{x}_k^i < x_k^* \tag{96}$$

$$(v^{ij})' \hat{x}^i > 0 \quad \text{and} \quad (v^{ik})' \hat{x}^i > 0 \tag{97}$$

$$V^j(\hat{x}^i) = V^k(\hat{x}^i). \tag{98}$$

The interpretation of Lemma 7 is provided in Section 7.4, and its proof is presented at the end of this section. In brief, the proof considers the behavior of the difference $V^k - V^j$ on the lines $\ell^{ij} = \{se_i + (1 - s)e_j : s \in \mathbb{R}\}$ and $\ell^{ik} = \{se_i + (1 - s)e_k : s \in \mathbb{R}\}$ through $\text{aff}(X)$ (see Figure 11). It is easy to check that $V^k - V^j$ is quadratic on each of these lines, and that it is concave on ℓ^{ij} and convex on ℓ^{ik} . Computations show that $V^k - V^j$ admits two 0s on each line; the 0s of interest, denoted y^{ij} and y^{ik} , are those with the larger i components. By definition, these points satisfy condition (98), and further computations confirm that they satisfy conditions (96) and (97), and that y^{ij} , y^{ik} , and x^* are collinear. If y^{ij} and y^{ik} are both e_i , we set $\hat{x}^i = e_i$. If not, exactly one of y^{ij} and y^{ik} is in X , and that one is our \hat{x}^i .

With Lemma 7 in hand, we can describe the value function and the optimal feedback control for the exit problem. To do so, we define the cross-product

$$w^i = x^* \times (\hat{x}^i - x^*)$$

to be a vector normal to segment $\hat{x}^i x^*$. By the right-hand rule (see Section 7.3.3), states satisfying $(w^i)'x > 0$ appear to the left of segment $\hat{x}^i x^*$ in Figure 6. It is convenient to focus on controls from the boundary $\text{bd}(Z)$ of Z , since every nonzero element of Z is proportional to a point in $\text{bd}(Z)$. Note also that

$$\text{bd}(Z) = \{\alpha - \beta : \alpha, \beta \in X, \text{supp}(\alpha) \cap \text{supp}(\beta) = \emptyset\}.$$

For concision, the results to come do not say explicitly that the value function equals 0 on the target sets; neither do they specify that the optimal control on those sets is the null control.

PROPOSITION 12. *If A is a simple three-strategy coordination game, the value function $V^* : \mathcal{B}^i \rightarrow \mathbb{R}_+$ for the exit cost problem (39) with target set $\mathcal{B}^j \cup \mathcal{B}^k$ is the continuous function*

$$V^*(x) = \begin{cases} \frac{1}{2} \frac{(A^{i-k}x)^2}{A_{i-k}^{i-k}} & \text{if } (w^i)'x \leq 0 \\ \frac{1}{2} \frac{(A^{i-j}x)^2}{A_{i-j}^{i-j}} & \text{if } (w^i)'x > 0. \end{cases} \tag{99}$$

The optimal feedback controls with range $\text{bd}(Z)$ are

$$v^*(x) \begin{cases} = e_k - e_i & \text{if } (w^i)'x < 0 \\ \in \{e_k - e_i, e_j - e_i\} & \text{if } (w^i)'x = 0 \\ = e_j - e_i & \text{if } (w^i)'x > 0. \end{cases} \tag{100}$$

PROOF. We apply the verification theorem. The value function V^* in (99) is constructed from feedback controls (100) that generate feasible solutions to the exit problem, as required by condition (i) of Theorem 11. The continuity of V follows from Lemma 7 and the argument in the subsequent paragraph. The function V^* is clearly C^1 off the set $\{x \in \text{aff}(X) : (w^i)'x = 0\}$, and Lemmas 4 and 7 imply that the HJB equation holds away from this set. Thus condition (ii) of Theorem 11 is satisfied, and the proof is complete. \square

PROOF OF LEMMA 7. For concreteness, we assume that $Q \geq 0$. The proof when $Q < 0$ is essentially the same, but with strategies j and k interchanged.

Let $\tilde{x}^{ik} = x^* + x_k^*(e_i - e_k) = (x_i^* + x_k^*)e_i + x_j^*e_j$ and $\tilde{x}^{ij} = x^* + x_j^*(e_i - e_j) = (x_i^* + x_j^*)e_i + x_k^*e_k$ (see Figure 11). We start by establishing the following lemma.

LEMMA 8. *We have $V^k(\tilde{x}^{ik}) > V^j(\tilde{x}^{ik})$ and $V^k(\tilde{x}^{ij}) < V^j(\tilde{x}^{ij})$.*

PROOF. Observe that

$$\begin{aligned} V^k(\tilde{x}^{ik}) - V^j(\tilde{x}^{ik}) &= \frac{1}{2} \left(\frac{(A^{i-k}(x^* + x_k^*(e_i - e_k)))^2}{A_{i-k}^{i-k}} - \frac{(A^{i-j}(x^* + x_k^*(e_i - e_k)))^2}{A_{i-j}^{i-j}} \right) \\ &= \frac{1}{2} (x_k^*)^2 \left(A_{i-k}^{i-k} - \frac{(A_{i-k}^{i-j})^2}{A_{i-j}^{i-j}} \right). \end{aligned}$$

Thus to prove the first inequality, it suffices to show that $A_{i-j}^{i-j}A_{i-k}^{i-k} - (A_{i-k}^{i-j})^2 > 0$, and, indeed,

$$A_{i-j}^{i-j}A_{i-k}^{i-k} - A_{i-k}^{i-j}A_{i-k}^{i-j} = A_{i-j}^{i-j}A_{i-k}^{i-k} - A_{i-j}^{i-j}A_{i-k}^{i-j} - A_{j-k}^{i-j}A_{i-k}^{i-j} = A_{j-k}^{j-i}A_{i-k}^{i-j} + A_{i-j}^{i-j}A_{k-i}^{k-j} > 0.$$

Interchanging j and k in these calculations proves the second inequality. \triangleleft

Next, let $\ell^{ij} = \{se_i + (1 - s)e_j : s \in \mathbb{R}\}$. The directional derivative of the quadratic function $V^k - V^j$ along this line is evaluated as

$$\begin{aligned}
 (DV^k(x) - DV^j(x))(e_i - e_j) &= \left(\frac{A^{i-k}x}{A_{i-k}^{i-k}} A^{i-k} - \frac{A^{i-j}x}{A_{i-j}^{i-j}} A^{i-j} \right) (e_i - e_j) \\
 &= \frac{A^{i-k}x}{A_{i-k}^{i-k}} A^{i-k} - A^{i-j}x \\
 &= \frac{1}{A_{i-k}^{i-k}} (A_{i-j}^{i-k} A^{i-k} - A_{i-k}^{i-k} A^{i-j})x \tag{101} \\
 &= \frac{1}{A_{i-k}^{i-k}} (-A_{k-j}^{k-i} A^i + A_{i-k}^{i-k} A^j - A_{i-j}^{i-k} A^k)x \\
 &= -(v^{ki})'x.
 \end{aligned}$$

Thus on ℓ^{ij} , $V^k - V^j$ is concave and is maximized at the unique state \tilde{x}^{ik} satisfying $(v^{ki})'x = 0$ (see Figure 11).

Recall that $\tilde{x}^{ik} = x^* + x_k^*(e_i - e_k) = x_j^*e_j + (x_i^* + x_k^*)e_i \in \ell^{ij}$ and let $\tilde{x}^{jk} = x^* + x_k^*(e_j - e_k) = x_i^*e_i + (x_j^* + x_k^*)e_j \in \ell^{ij}$. Since $Q \geq 0$, (58) and (59) and inequality (60) imply that $(v^{ki})'\tilde{x}^{ik} = x_k^*Q \geq 0$ and that $(v^{ki})'\tilde{x}^{jk} < 0$. Thus \tilde{x}^{ik} lies between \tilde{x}^{ik} and \tilde{x}^{jk} , and is equal to the former if and only if $Q = 0$ (again, see Figure 11). Since $V^k(\tilde{x}^{ik}) > V^j(\tilde{x}^{ik})$ by Lemma 8 and since $V^k - V^j$ is concave quadratic on ℓ^{ij} , we have the following lemma.

LEMMA 9. *There is a unique state $y^{ij} \in \ell^{ij}$ with $y_i^{ij} > \tilde{x}_i^{ik} \geq x_i^*$ such that $V^k(y^{ij}) = V^j(y^{ij})$.*

Next, we consider the directional derivative of the quadratic function $V^k - V^j$ along line $\ell^{ik} = \{se_i + (1 - s)e_k : s \in \mathbb{R}\}$. A calculation similar to (101) shows that

$$(DV^k(x) - DV^j(x))(e_i - e_k) = \frac{1}{A_{i-j}^{i-j}} (A_{i-j}^{i-j} A^{i-k} - A_{i-k}^{i-j} A^{i-j})x = (v^{ij})'x.$$

Thus on ℓ^{ik} , $V^k - V^j$ is convex and is minimized at the unique state \tilde{x}^{ij} on ℓ^{ik} satisfying $(v^{ij})'x = 0$ (once again, see Figure 11). Since $Q \geq 0$, (58) and (59) and inequality (60) imply that $\tilde{x}^{ij} = \tilde{x}^{ij} + c(e_i - e_k)$ for some $c \geq 0$, with equality only if and only if $Q = 0$. Since $V^k(\tilde{x}^{ij}) < V^j(\tilde{x}^{ij})$ by Lemma 8 and since $V^k - V^j$ is convex quadratic on ℓ^{ik} , we have the following lemma.

LEMMA 10. *There is a unique state $y^{ik} \in \ell^{ik}$ with $y_i^{ik} > \tilde{x}_i^{ij} \geq \tilde{x}_i^{ij}$ such that $V^k(y^{ik}) = V^j(y^{ik})$.*

To complete the proof, we use the homogeneity of degree 2 of $V^j(x)$ and $V^k(x)$ in the displacement $z = x - x^*$ of x from x^* . Specifically, for $z \in TX$ and $s \in \mathbb{R}$, we have

$$\begin{aligned}
 V^k(x^* + sz) - V^j(x^* + sz) &= V^k(sz) - V^j(sz) = s^2(V^k(z) - V^j(z)) \\
 &= s^2(V^k(x^* + z) - V^j(x^* + z)).
 \end{aligned}$$

Thus if $V^k(x^* + z) = V^j(x^* + z)$, then $V^k(x^* + sz) = V^j(x^* + sz)$ for all $s \in \mathbb{R}$. It therefore follows from Lemmas 9 and 10 that y^{ij} and y^{ik} are collinear with x^* (see Figure 11), and so that both of these points satisfy (96), (97), and (98). It could be that $y^i = y^{ik} = e_i$, in which case we choose $\hat{x}^i = e_i$. Otherwise, exactly one of y^{ij} and y^{ik} is in X , in which case we choose \hat{x}^i to be this state. This completes the proof of Lemma 7. \square

A.10 *The value function for the transition cost problem in potential games*

The following proposition describes the optimal feedback controls (Figure 7) and value function for the transition cost problem (40) in potential games.

PROPOSITION 13. *Let A be a simple three-strategy coordination game and suppose that A is a potential game. Then the value function $V^* : \mathcal{B}^i \cup \mathcal{B}^j \rightarrow \mathbb{R}_+$ for the transition cost problem (40) with target set \mathcal{B}^k is the C^1 function*

$$V^*(x) = \begin{cases} \frac{1}{2} \frac{(A^{i-k}x)^2}{A^{i-k}} & \text{if } x_j < x_j^* \\ \frac{1}{2} (x - x^*)' A (x - x^*) & \text{if } x_j \geq x_j^* \text{ and } x_i \geq x_i^* \\ \frac{1}{2} \frac{(A^{j-k}x)^2}{A^{j-k}} & \text{if } x_i < x_i^*. \end{cases} \tag{102}$$

The optimal feedback controls with range $\text{bd}(Z)$ are

$$v^*(x) \begin{cases} = e_k - e_i & \text{if } x_j < x_j^* \\ \in \text{conv}(\{e_j - e_i, e_k - e_i\}) & \text{if } x_j \geq x_j^* \text{ and } A^{i-j}x > 0 \\ = -\zeta^{ij} & \text{if } A^{i-j}x = 0 \\ \in \text{conv}(\{e_i - e_j, e_k - e_j\}) & \text{if } x_i \geq x_i^* \text{ and } A^{i-j}x < 0 \\ = e_k - e_j & \text{if } x_i < x_i^*. \end{cases} \tag{103}$$

PROOF. To apply the verification theorem, we first show that V^* is C^1 . This is clearly true inside each of the three regions in the piecewise definition (102). It remains to consider the behavior of V^* at states satisfying $x_j = x_j^*$ or $x_i = x_i^*$. We focus on the former states. Such states satisfy $x = x^* + d(e_k - e_i)$ for some $d \geq 0$. It follows that V^* is continuous at such states, since

$$V^1(x) \equiv \frac{1}{2} \frac{(A^{i-k}x)^2}{A^{i-k}} = \frac{1}{2} d^2 A^{i-k} = \frac{1}{2} (x - x^*)' A (x - x^*) \equiv V^2(x).$$

To check that V^* is C^1 , recall from Section 7.1.1 that since A is a potential game, we can write $A = C + \mathbf{1}r'$ for some symmetric matrix $C \in \mathbb{R}^{n \times n}$ and some vector $r \in \mathbb{R}^n$. Using these facts and the fact that x^* is an interior Nash equilibrium of both A and C , we have

$$\begin{aligned} V^2(x) &= \frac{1}{2} (x - x^*)' A (x - x^*) = \frac{1}{2} (x - x^*)' C (x - x^*) = \frac{1}{2} (x - x^*)' C x \\ &= \frac{1}{2} (x' C x - x' C x^*) = \frac{1}{2} (x' C x - (x^*)' C x^*). \end{aligned} \tag{104}$$

Thus for $z \in TX$, these facts and the symmetry of A with respect to $TX \times TX$ yield

$$DV^2(x)z = x' C z = z' C x = z' A x = z' A (x - x^*) = (x - x^*)' A z. \tag{105}$$

But at states x with $x_j = x_j^*$, the fact that $A^{i-k} x^* = 0$ and the definition of d imply that

$$DV^1(x)z = \frac{A^{i-k} x}{A_{i-k}^{i-k}} A^{i-k} z = \frac{A^{i-k} (x - x^*)}{A_{i-k}^{i-k}} A^{i-k} z = d A^{k-i} z = (x - x^*)' A z,$$

so V^* is C^1 at these states.

Next, we show that the value function V^* is generated by the controls (103). For the first and third cases of definition (102), this follows from Lemma 3. To address the second case, we require the following lemma, which applies equally well to the other cases (see the discussion following this proof). The lemma uses the fact that $f(x) = \frac{1}{2} x' C x$ is a potential function for $F(x) = Ax$ on $\text{aff}(X)$, in the sense that $Df(x)z = F(x)'z = z' Ax$ for all $z \in TX$ and $x \in \text{aff}(X)$.

LEMMA 11. *The cost $c(\phi)$ of trajectory $\phi : [0, T] \rightarrow \text{aff}(X)$ satisfies $c(\phi) \geq f(\phi_0) - f(\phi_T)$. If for each $t \in (0, T)$, every strategy h with $(\dot{\phi}_t)_h < 0$ is optimal at ϕ_t , then $c(\phi) = f(\phi_0) - f(\phi_T)$.*

PROOF. By definition (38) of path costs and since $[\dot{\phi}_t]_+ \mathbf{1} = [\dot{\phi}_t]' \mathbf{1}$ and $\mathbf{1} A^{\check{b}(\phi_t)} \phi_t \geq A \phi_t$,

$$\begin{aligned} c(\phi) &= \int_0^T [\dot{\phi}_t]_+ (\mathbf{1} A^{\check{b}(\phi_t)} - A) \phi_t dt \geq \int_0^T ([\dot{\phi}_t]' A - [\dot{\phi}_t]_+ A) \phi_t dt \\ &= - \int_0^T \dot{\phi}_t' A \phi_t dt = - \int_0^T Df(\phi_t) \dot{\phi}_t dt = f(\phi_0) - f(\phi_T). \end{aligned}$$

If the assumption on the support of $[\dot{\phi}_t]_-$ holds, then $[\dot{\phi}_t]' \mathbf{1} A^{\check{b}(\phi_t)} \phi_t = [\dot{\phi}_t]'_- A \phi_t$, so the inequality in the display binds. □

Proceeding with our earlier argument, we note that any controlled path $\phi : [0, T] \rightarrow \text{aff}(X)$ starting from a state x with $x_j \geq x_j^*$ and $x_i \geq x_i^*$ and generated by controls satisfying (103) both satisfies the assumption of Lemma 11 and terminates at $\phi_T = x^*$ (see Figure 7). Thus Lemma 11, the definition of f , and (104) yield

$$c(\phi) = f(x) - f(x^*) = \frac{1}{2} x' C x - \frac{1}{2} (x^*)' C x^* = \frac{1}{2} (x - x^*)' A (x - x^*),$$

as specified in the second case of (102).

The proposition will follow from Theorem 11 if we can show that V^* satisfies the HJB equation (44) at all states. Since A is a potential game, the states with $(v^{ki})' x > 0$ are those satisfying $x_j < x_j^*$ (see Figure 5(i)), so Lemma 4 implies that V^* satisfies (44) at these states. Analogous reasoning shows that V^* satisfies (44) when $x_i < x_i^*$. It thus remains to check (44) at states satisfying $x_j \geq x_j^*$ and $x_i \geq x_i^*$. To do so, we show that (44) holds when $x_j \geq x_j^*$ and $A^{i-j} x > 0$; the argument when $x_i \geq x_i^*$ and $A^{i-j} x < 0$ is similar;

and then (44) must hold when $x_j = x_j^*$, $x_i = x_i^*$, or $A^{i-j}x = 0$ by virtue of the fact that V^* is C^1 .

So suppose that x satisfies $x_j \geq x_j^*$ and $A^{i-j}x > 0$. Since $DV^*(x)z = z'Ax$ at such states by (105), substitution into the HJB equation (44) yields

$$\min_{e_a, e_b \neq e_a} (e_i - e_b)'Ax = 0. \tag{106}$$

Since x is in \mathcal{B}^i but not in \mathcal{B}^j or \mathcal{B}^k (see Figure 7), minimization in (106) requires setting $e_b = e_i$. Then choosing e_a to be either e_j or e_k attains the minimum of 0.⁴⁴ This completes the proof of the proposition. \square

A.11 The value function for the transition cost problem in skewed games

A.11.1 Preliminary calculations Lemma 13 and Proposition 14 require a number of preliminary definitions and calculations. To begin, we introduce notation for the endpoints of paths that proceed in a basic direction until reaching a boundary between best-response regions. Using Lemma 3 and proceeding from top to bottom in Figure 8 or Figure 9, we have

for $x \in \mathcal{B}^i$ with $x_j \leq x_j^*$,

$$\text{let } \omega^{ik}(x) = x + (e_k - e_i)d^{ik}(x) \in \mathcal{B}^{ki}, \quad \text{where } d^{ik}(x) = \frac{A^{i-k}x}{A_{i-k}^{i-k}}$$

for $x \in \mathcal{B}^i$ with $x_j \geq x_j^*$,

$$\text{let } \chi^{ik}(x) = x + (e_k - e_i)\delta^{ik}(x) \in \mathcal{B}^{ij}, \quad \text{where } \delta^{ik}(x) = \frac{A^{i-j}x}{A_{i-k}^{i-j}}$$

for $x \in \mathcal{B}^j$ with $x_k \leq x_k^*$,

$$\text{let } \chi^{ji}(x) = x + (e_i - e_j)d^{ji}(x) \in \mathcal{B}^{ij}, \quad \text{where } d^{ji}(x) = \frac{A^{j-i}x}{A_{j-i}^{j-i}}$$

for $x \in \mathcal{B}^j$ with $x_i \leq x_i^*$,

$$\text{let } \omega^{jk}(x) = x + (e_k - e_j)d^{jk}(x) \in \mathcal{B}^{jk}, \quad \text{where } d^{jk}(x) = \frac{A^{j-k}x}{A_{j-k}^{j-k}}.$$

Using these definitions, we can define the pieces of the value function. Again proceeding from top to bottom in Figure 8 or Figure 9, we have

$$\begin{aligned} V^1(x) &= \gamma(x, \omega^{ik}(x)) \\ V^2(x) &= \gamma(x, \chi^{ik}(x)) + \gamma(\chi^{ik}(x), x^*) \end{aligned}$$

⁴⁴To consider all controls in $\text{bd}(Z)$, we must write the HJB equation in form (43); then the previous argument and the piecewise linearity of (43) imply that the set of optimal controls is $\text{conv}(\{e_j - e_i, e_k - e_i\})$, as described in the second case of (103).

$$V^3(x) = \gamma(x, \chi^{ji}(x)) + \gamma(\chi^{ji}(x), x^*)$$

$$V^4(x) = \gamma(x, \omega^{jk}(x)).$$

We next state a counterpart of Lemma 3 for paths along \mathcal{B}^{ij} to x^* .

LEMMA 12. *If $y \in \mathcal{B}^{ij}$, then*

$$y = x^* + d^\zeta(y)\zeta^{ij}, \quad \text{where } d^\zeta(y) = x_k^* - y_k$$

$$\gamma(y, x^*) = \frac{1}{2}d^\zeta(y)^2 A_\zeta^\zeta, \quad \text{where } A_\zeta^\zeta = (\zeta^{ij})' A \zeta^{ij} > 0. \tag{107}$$

PROOF. Since $y \in \text{conv}(\{x^{ij}, x^*\})$, (51) implies that we can write $y = x^* + d\zeta^{ij}$ for some $d > 0$. Then (52) implies that $y_k = x_k^* - d$, which yields (107), and Lemma 2 implies that

$$\begin{aligned} \gamma(y, x^*) &= (y - x^*)' A \left(\frac{1}{2}(y + x^*) \right) \\ &= d^\zeta(y)(\zeta^{ij})' A \left(x^* + \frac{1}{2}d^\zeta(y)\zeta^{ij} \right) = \frac{1}{2}d^\zeta(y)^2 A_\zeta^\zeta. \quad \square \end{aligned}$$

Lemma 12 gives an expression for $d^\zeta(y)$ that is affine in y . To match Lemma 3, one can instead write $d^\zeta(y) = (\zeta^{ij})' Ay / A_\zeta^\zeta$. The key point is that either way, $d^\zeta(x^* + z)$ is linear in the displacement z .

Next we give explicit expressions for each piece of the value function and their derivatives:

$$V^1(x) = \gamma(x, \omega^{ik}(x)) = \frac{1}{2}d^{ik}(x)^2 A_{i-k}^{i-k} = \frac{1}{2} \frac{(A^{i-k}x)^2}{A_{i-k}^{i-k}}, \quad \text{so}$$

$$DV^1(x) = \frac{A^{i-k}x}{A_{i-k}^{i-k}} A^{i-k}$$

$$V^2(x) = \gamma(x, \chi^{ik}(x)) + \gamma(\chi^{ik}(x), x^*)$$

$$= \delta^{ik}(x)A^{i-k} \chi^{ik}(x) + \frac{1}{2}\delta^{ik}(x)^2 A_{i-k}^{i-k} + \frac{1}{2}d^\zeta(\chi^{ik}(x))^2 A_\zeta^\zeta, \quad \text{so}$$

$$DV^2(x) = \delta^{ik}(x)A^{i-k} D\chi^{ik}(x) + A^{i-k} \chi^{ik}(x) D\delta^{ik}(x) + A_{i-k}^{i-k} \delta^{ik}(x) D\delta^{ik}(x)$$

$$+ A_\zeta^\zeta d^\zeta(\chi^{ik}(x)) Dd^\zeta(\chi^{ik}(x)) D\chi^{ik}(x)$$

$$V^3(x) = \gamma(x, \chi^{ji}(x)) + \gamma(\chi^{ji}(x), x^*) = \frac{1}{2}d^{ji}(x)^2 A_{j-i}^{j-i} + \frac{1}{2}d^\zeta(\chi^{ji}(x))^2 A_\zeta^\zeta, \quad \text{so}$$

$$DV^3(x) = A_{j-i}^{j-i} d^{ji}(x) Dd^{ji}(x) + A_\zeta^\zeta d^\zeta(\chi^{ji}(x)) Dd^\zeta(\chi^{ji}(x)) D\chi^{ji}(x)$$

$$V^4(x) = \gamma(x, \omega^{jk}(x)) = \frac{1}{2}d^{jk}(x)^2 A_{j-k}^{j-k} = \frac{1}{2} \frac{(A^{j-k}x)^2}{A_{j-k}^{j-k}}, \quad \text{so}$$

$$DV^4(x) = \frac{A^{j-k}x}{A_{j-k}^{j-k}} A^{j-k}.$$

The functions above are expressed in terms of derivatives of linear functions from Section 7.5.2 and Lemma 12. These derivatives are written explicitly as

$$\begin{aligned}
 D\delta^{ik}(x) &= \frac{A^{i-j}}{A_{i-k}^{i-j}} \\
 D\chi^{ik}(x) &= I + (e_k - e_i)D\delta^{ik}(x) = I + (e_k - e_i)\frac{A^{i-j}}{A_{i-k}^{i-j}} \\
 Dd^{ji}(x) &= \frac{A^{j-i}}{A_{j-i}^{j-i}} \\
 D\chi^{ji}(x) &= I + (e_i - e_j)Dd^{ji}(x) = I + (e_i - e_j)\frac{A^{j-i}}{A_{j-i}^{j-i}} \\
 Dd^{\zeta}(y) &= -e'_k \quad \text{for } y \in \mathcal{B}^{ij}.
 \end{aligned}$$

A.11.2 *Statement and proof of Lemma 13* Lemma 13 identifies state \hat{x}^{jk} from Figures 8 and 9. The requirements of the lemma were interpreted in Section 7.5.2, and its proof proceeds in similar fashion to that of Lemma 7.

LEMMA 13. *If A has clockwise skew, then there is a unique state $\hat{x}^{jk} \in \mathcal{B}^j \cap \text{bd}(X)$ such that*

$$\hat{x}_k^{jk} < x_k^* \tag{108}$$

$$(v^{jk})' \hat{x}^{jk} > 0 \tag{109}$$

$$V^3(\hat{x}^{jk}) = V^4(\hat{x}^{jk}). \tag{110}$$

PROOF. We consider the behavior of the quadratic function $V^4 - V^3$ on the line $\ell^{ij} = \{se_i + (1-s)e_j : s \in \mathbb{R}\}$. To begin, note that since $D\chi^{ji}(x)(e_j - e_i) = 0$, a calculation similar to (101) shows that

$$(DV^4(x) - DV^3(x))(e_j - e_i) = \frac{1}{A_{j-k}^{j-k}}(A_{j-i}^{j-k} A^{j-k} - A_{j-k}^{j-k} A^{j-i})x = -(v^{jk})'x.$$

Since $e_j - e_i$ is tangent to ℓ^{ij} , it follows that $V^4 - V^3$ is concave on ℓ^{ij} and reaches its maximum on this line at the unique state on ℓ^{ij} satisfying $(v^{jk})'x = 0$. We denote this state by \tilde{x}^{jk} (see Figure 12).

Let $\tilde{x}^{jk} = x^* + x_k^*(e_j - e_k) \in \ell^{ij}$. We show that $V^4(\tilde{x}^{jk}) - V^3(\tilde{x}^{jk}) > 0$. Observe that $d^{jk}(\tilde{x}^{jk}) = x_k^*$, $d^{\zeta}(\tilde{x}^{jk}) = x_k^*$, and

$$d^{ji}(\tilde{x}^{jk}) = \frac{A^{j-i}(x^* + x_k^*(e_j - e_k))}{A_{j-i}^{j-i}} = x_k^* \frac{A_{j-k}^{j-i}}{A_{j-i}^{j-i}},$$

$$= A_{j-k}^{i-i} Q > 0,$$

as claimed.

Since $(v^{jk})' \check{x}^{jk} = 0$ and $(v^{jk})'(e_i - e_j) > 0$ (see (61)), it follows that $\check{x}^{jk} = \tilde{x}^{jk} + c(e_j - e_i)$ for some $c > 0$. Thus as one proceeds along ℓ^{ij} in direction $e_j - e_i$ starting from \tilde{x}^{jk} , the function $V^4 - V^3$ starts at a positive value, increases until reaching its maximum at \check{x}^{jk} , and then decreases, ultimately approaching $-\infty$. Thus there is a unique point $y^{jk} = \check{x}^{jk} + b(e_j - e_i) \in \ell^{ij}$ with $b > 0$ at which $V^4(x) - V^3(x) = 0$ (see Figure 12).

If $y_i^{jk} \geq 0$, so that y^{jk} is in X , then we let $\hat{x}^{jk} = y^{jk}$, and this point clearly satisfies (108), (109), and (110). If instead $y_i^{jk} < 0$, we let

$$\hat{x}^{jk} = \frac{x_i^*}{x_i^* - y_i^{jk}} y_i^{jk} + \frac{-y_i^{jk}}{x_i^* - y_i^{jk}} x^*,$$

which is the point on the segment between y^{jk} and x^* whose i th component is 0 (see Figure 12). Since equality (110) and inequality (109) hold at y^{jk} and are preserved along rays from x^* , they continue to hold at \hat{x}^{jk} , with a strict inequality in the case of (109). And since \hat{x}_k^{jk} is a strictly convex combination of x_k^* and $-y_k^{jk} > 0$, we have $\hat{x}_k^{jk} < x_k^*$, which is inequality (108). This completes the proof of the lemma. \square

A.11.3 Statement of Proposition 14 Proposition 14 describes the optimal feedback controls (Figures 8 and 9) and value function for the transition cost problem (40) in skewed games. To state it, we define the vector w^{jk} to be the cross-product

$$w^{jk} = x^* \times (\hat{x}^{jk} - x^*).$$

In Figures 8 and 9, the states satisfying $(w^{jk})'x > 0$ are those below the ray from x^* through \hat{x}^{jk} .

PROPOSITION 14. *Let A be a simple three-strategy coordination game with clockwise skew. Then the value function $V^* : \mathcal{B}^i \cup \mathcal{B}^j \rightarrow \mathbb{R}_+$ for the transition cost problem (40) with target set \mathcal{B}^k is*

$$V^*(x) = \begin{cases} \gamma(x, \omega^{ik}(x)) & \text{if } x_j \leq x_j^* \\ \gamma(x, \chi^{ik}(x)) + \gamma(\chi^{ik}(x), x^*) & \text{if } x_j > x_j^* \text{ and } A^{i-j}x \geq 0 \\ \gamma(x, \chi^{ji}(x)) + \gamma(\chi^{ji}(x), x^*) & \text{if } A^{i-j}x < 0 \text{ and } (w^{jk})'x < 0 \\ \gamma(x, \omega^{jk}(x)) & \text{if } (w^{jk})'x \geq 0. \end{cases} \tag{113}$$

The optimal feedback controls with range $\text{bd}(Z)$ are

$$v^*(x) \begin{cases} = e_k - e_i & \text{if } A^{i-j}x > 0 \\ = -\zeta^{ij} & \text{if } A^{i-j}x = 0 \\ = e_i - e_j & \text{if } A^{i-j}x < 0 \text{ and } (w^{jk})'x < 0 \\ \in \{e_i - e_j, e_k - e_j\} & \text{if } (w^{jk})'x = 0 \\ = e_k - e_j & \text{if } (w^{jk})'x > 0. \end{cases}$$

To prove Proposition 14, we establish that the value function defined in (113) satisfies the conditions of the verification theorem. In Appendix A.11.4, we show that V is continuous and that it is differentiable except at states x at which $(w^{jk})'x = 0$. In Appendix A.11.5, we use Lemmas 1, 4, and 13 to show that the HJB equation holds at all other states. The proposition then follows from Theorem 11.

While the algebraic presentation below may look complicated, many of the arguments are quite simple when interpreted geometrically.

A.11.4 Continuity and piecewise smoothness of V^* Lemma 14 shows that the value function V^* is continuous on the boundary between the third and fourth cases of definition (113).

LEMMA 14. *If $x \in \mathcal{B}^j$ and $(w^{jk})'x = 0$, then $V^3(x) = V^4(x)$.*

PROOF. Since $(w^{jk})'x = 0$ and $w^{jk} = x^* \times (\hat{x}^{jk} - x^*)$, we can write $x = x^* + r(\hat{x}^{jk} - x^*)$ for some $r \in [0, 1]$. By condition (110) and the expressions for V^3 and V^4 above, it is enough to show that $d^{jk}(x) = rd^{jk}(\hat{x}^{jk})$, $d^{ji}(x) = rd^{ji}(\hat{x}^{jk})$, and $d^\zeta(\chi^{ji}(x)) = rd^\zeta(\chi^{ji}(x))$. And indeed, the fact that Ax^* is a multiple of $\mathbf{1}$ implies that

$$d^{jk}(x) = \frac{A^{j-k}x}{A_{j-k}^{j-k}} = r \frac{A^{j-k}\hat{x}^{jk}}{A_{j-k}^{j-k}} = rd^{jk}(\hat{x}^{jk})$$

$$d^{ji}(x) = \frac{A^{j-i}x}{A_{j-i}^{j-i}} = r \frac{A^{j-i}\hat{x}^{jk}}{A_{j-i}^{j-i}} = rd^{ji}(\hat{x}^{jk}),$$

while the third equality follows from the fact that

$$d^\zeta(\chi^{ji}(x)) = x_k^* - e'_k(x + (e_i - e_j)d^{ji}(x)) = x_k^* - x_k. \quad \square$$

Lemmas 15 and 16 establish differentiability of V^* on the boundaries between the first and second and the second and third cases of definition (113).

LEMMA 15. *If $\tilde{x} \in \mathcal{B}^i$ satisfies $\tilde{x} = x^* + d(e_i - e_k)$ for some $d \geq 0$, then $DV^1(\tilde{x}) = DV^2(\tilde{x}) = dA^{i-k}$.*

PROOF. Note first that

$$DV^1(\tilde{x}) = \frac{A^{i-k}\tilde{x}}{A_{i-k}^{i-k}} A^{i-k} = \frac{A^{i-k}(x^* + d(e_i - e_k))}{A_{i-k}^{i-k}} A^{i-k} = dA^{i-k}.$$

To compute $DV^2(\tilde{x})$, use the definition of χ^{ik} and Lemma 12 (or draw a picture) to show that $\delta^{ik}(\tilde{x}) = d$, $\chi^{ik}(\tilde{x}) = x^*$, and $d^\zeta(x^*) = 0$. Then since $A^{i-k}x^* = 0$, we have that

$$DV^2(\tilde{x}) = \delta^{ik}(\tilde{x})A^{i-k}D\chi^{ik}(\tilde{x}) + A_{i-k}^{i-k}\delta^{ik}(\tilde{x})D\delta^{ik}(\tilde{x})$$

$$= dA^{i-k} \left(I + (e_k - e_i) \frac{A^{i-j}}{A_{i-k}^{i-j}} \right) + A_{i-k}^{i-k}d \frac{A^{i-j}}{A_{i-k}^{i-j}}$$

$$\begin{aligned}
 &= dA^{i-k} + dA_{k-i}^{i-k} \frac{A^{i-j}}{A_{i-k}^{i-j}} + A_{i-k}^{i-k} d \frac{A^{i-j}}{A_{i-k}^{i-j}} \\
 &= dA^{i-k}.
 \end{aligned}$$

□

LEMMA 16. If $y \in \mathcal{B}^{ij}$, then $DV^2(y) = DV^3(y) = (Ay)'$ (as linear forms on TX).

PROOF. Since $y \in \mathcal{B}^{ij}$, we have $\delta^{ik}(y) = 0$, $\chi^{ik}(y) = y$, and $d^\zeta(y) = d$. Thus

$$\begin{aligned}
 DV^2(y) &= A^{i-k} \chi^{ik}(y) D\delta^{ik}(y) + A_\zeta^\zeta d^\zeta(\chi^{ik}(y)) Dd^\zeta(\chi^{ik}(y)) D\chi^{ik}(y) \\
 &= A^{i-k} y \frac{A^{i-j}}{A_{i-k}^{i-j}} + A_\zeta^\zeta(x_k^* - y_k) \left(-e'_k \left(I + (e_k - e_i) \frac{A^{i-j}}{A_{i-k}^{i-j}} \right) \right) \\
 &= A^{i-k} y \frac{A^{i-j}}{A_{i-k}^{i-j}} + A_\zeta^\zeta(x_k^* - y_k) \left(-e'_k - \frac{A^{i-j}}{A_{i-k}^{i-j}} \right) \\
 &= \frac{A^{i-k} y - A_\zeta^\zeta(x_k^* - y_k)}{A_{i-k}^{i-j}} A^{i-j} - A_\zeta^\zeta(x_k^* - y_k) e'_k.
 \end{aligned}$$

Thus

$$DV^2(y)(e_i - e_k) = A^{i-k} y - A_\zeta^\zeta(x_k^* - y_k) + A_\zeta^\zeta(x_k^* - y_k) = A^{i-k} y = (Ay)'(e_i - e_k),$$

and since $A^{i-j} \zeta^{ij} = 0$ (see (112)) and

$$(\zeta^{ij})' Ay = (\zeta^{ij})' A(x^* + d^\zeta(y) \zeta^{ij}) = A_\zeta^\zeta(x_k^* - y_k),$$

we have

$$DV^2(y) \zeta^{ij} = A_\zeta^\zeta(x_k^* - y_k) = (Ay)' \zeta^{ij}.$$

Since $e_i - e_k$ and ζ^{ij} span TX , we conclude that $DV^2(y) = (Ay)'$.

Again using $\delta^{ji}(y) = 0$ and $\chi^{ji}(y) = y$, we have

$$\begin{aligned}
 DV^3(y) &= A_\zeta^\zeta d^\zeta(\chi^{ji}(y)) Dd^\zeta(\chi^{ji}(y)) D\chi^{ji}(y) \\
 &= A_\zeta^\zeta(x_k^* - y_k) \left(-e'_i \left(I + (e_i - e_j) \frac{A^{j-i}}{A_{j-i}^{j-i}} \right) \right) \\
 &= -A_\zeta^\zeta(x_k^* - y_k) e'_k.
 \end{aligned}$$

Thus

$$\begin{aligned}
 DV^3(y)(e_i - e_j) &= 0 = (Ay)'(e_i - e_j) \\
 DV^3(y) \zeta^{ij} &= A_\zeta^\zeta(x_k^* - y_k) = (Ay)' \zeta^{ij}.
 \end{aligned}$$

Thus since $e_i - e_j$ and ζ^{ij} span TX , we conclude that $DV^3(y) = (Ay)'$.

□

A.11.5 *Checking the HJB equation* To complete the proof of Proposition 14, we need to show that the HJB equation (44) is satisfied at all states at which V^* is C^1 . In the first case of the definition (113) of V^* , this follows from Lemma 4 and the fact that $Q > 0$, since (59) implies that $(v^{ki})'x \geq 0$ when $x \in \mathcal{B}^i$ and $x_j \leq x_j^*$ (see Figure 5(ii)).

Similarly, in the fourth case of definition (113), the HJB equation follows from Lemma 4 (with the roles of i and j reversed) and Lemma 13, which ensures that $(w^{jk})'x \geq 0$ when $x \in \mathcal{B}^j$ and $(w^{jk})'x \geq 0$ (see Figure 12).

To handle the two remaining cases of definition (113), we apply Lemma 1. Observe that the regions defined by these cases are convex cones in $\text{aff}(X)$ emanating from x^* . Also, the expressions for $DV^2(x)$ and $DV^3(x)$ in Appendix A.11.1 imply that within each of these regions, the function to be minimized in the HJB equation (44) is linear in the displacement $z = x - x^*$ of x from x^* . Therefore, to establish inequalities (45) and (46) from Lemma 1 for all the states in one of these cones, it is enough to do so at three states: x^* and one state from each edge of the cone. Since we have shown that V^* is C^1 on the boundaries between the first and second and the second and third cases of (113), this analysis also establishes that the HJB equation (44) holds on these boundaries.

For the second case of definition (113), we show that inequalities (45) and (46) hold at states x^* , x^{ij} , and $\tilde{x}^{ik} = x^* + x_k^*(e_i - e_k) = (x_i^* + x_k^*)e_i + x_k^*e_j$:

$$\begin{aligned}
 DV^2(x^*) &= DV^2(x^*) - (Ax^*)' = \mathbf{0}' \quad (\text{as a linear form on } TX) \\
 DV^2(x^{ij})(e_i - e_k) &= A^{i-k}x^{ij} > 0 \\
 DV^2(x^{ij})(e_i - e_j) &= A^{i-j}x^{ij} = 0 \\
 DV^2(x^{ij}) - (Ax^{ij})' &= \mathbf{0}' \quad (\text{as a linear form on } TX) \\
 DV^2(\tilde{x}^{ik})(e_i - e_k) &= x_k^*A_{i-k}^{i-k} > 0 \\
 DV^2(\tilde{x}^{ik})(e_i - e_j) &= x_k^*A_{i-j}^{i-k} > 0 \\
 (DV^2(\tilde{x}^{ik}) - (A\tilde{x}^{ik})')(e_k - e_j) &= (DV^1(\tilde{x}^{ik}) - (A\tilde{x}^{ik})')(e_k - e_j) \leq 0.
 \end{aligned}$$

The final statement uses the fact that V^* is C^1 on the boundary between the second and third cases of (113), the fact that $(v^{ki})'\tilde{x}^{ik} > 0$, and the display before Lemma 4.

For the third case of definition (113), we show that inequalities (45) and (46) hold at states x^* , x^{ij} , and $y^{jk} \in \ell^{ij}$, the last of which was introduced in the proof of Lemma 13. The inequalities for the first two states are straightforward to check:

$$\begin{aligned}
 DV^3(x^*) &= DV^3(x^*) - (Ax^*)' = \mathbf{0}' \quad (\text{as a linear form on } TX) \\
 DV^3(x^{ij})(e_j - e_k) &= A^{j-k}x^{ij} > 0 \\
 DV^3(x^{ij})(e_j - e_i) &= A^{j-i}x^{ij} = 0 \\
 DV^3(x^{ij}) - (Ax^{ij})' &= \mathbf{0}' \quad (\text{as a linear form on } TX).
 \end{aligned}$$

It remains to check inequalities (45) and (46) for state y^{jk} . Since $D\chi^{ij}(x)(e_j - e_i) = 0$,

$$DV^3(y^{jk})(e_j - e_i) = A_{j-i}^{j-i}d^{jj}(y^{jk}) > 0.$$

Next, since

$$d^\zeta(\chi^{ji}(x)) = d^\zeta(x + (e_i - e_j)d^{ji}(x)) = x_k^* - x_k$$

and since for $y \in \mathcal{B}^{ij}$,

$$\begin{aligned} D\chi^{ji}(y)(e_j - e_k) &= (e_j - e_k) + (e_i - e_j) \frac{A_{j-k}^{j-i}}{A_{j-i}^{j-i}} \\ &= \frac{1}{A_{j-i}^{j-i}} ((e_j - e_k)A_{j-i}^{j-i} + (e_i - e_j)A_{j-k}^{j-i}) \\ &= \frac{1}{A_{j-i}^{j-i}} (A_{j-k}^{j-i}e_i + A_{i-k}^{i-j}e_j - A_{j-i}^{j-i}e_k) \\ &= \zeta^{ij}, \end{aligned}$$

we have

$$\begin{aligned} DV^3(x)(e_j - e_k) &= A_{j-k}^{j-i}d^{ji}(x) + A_\zeta^\zeta d^\zeta(\chi^{ji}(x))Dd^\zeta(\chi^{ji}(x))D\chi^{ji}(x)(e_j - e_k) \\ &= A_{j-k}^{j-i}d^{ji}(x) + A_\zeta^\zeta(x_k^* - x_k)(-e_k'\zeta^{ij}) \\ &= A_{j-k}^{j-i}d^{ji}(x) + A_\zeta^\zeta(x_k^* - x_k). \end{aligned}$$

Thus the fact that $y_k^{jk} = 0$ implies that

$$DV^3(y^{jk})(e_j - e_k) = A_{j-k}^{j-i}d^{ji}(x) + A_\zeta^\zeta x_k^* > 0.$$

This establishes the two cases of inequality (45) at state y^{jk} .

It remains to establish inequality (46) at state y^{jk} . Computing as above shows that for $y \in \mathcal{B}^{ij}$,

$$\begin{aligned} D\chi^{ji}(y)(e_i - e_k) &= \frac{1}{A_{j-i}^{j-i}} ((e_i - e_k)A_{j-i}^{j-i} + (e_i - e_j)A_{i-k}^{j-i}) = \zeta^{ij} \\ DV^3(x)(e_i - e_k) &= A_{i-k}^{j-i}d^{ji}(x) + A_\zeta^\zeta(x_k^* - x_k). \end{aligned}$$

Hence

$$\begin{aligned} (DV^3(x) - (Ax)')(e_i - e_k) &= A_{i-k}^{j-i} \frac{A_{j-i}^{j-i}x}{A_{j-i}^{j-i}} + A_\zeta^\zeta(x_k^* - x_k) - A^{i-k}x \\ &= \frac{1}{A_{j-i}^{j-i}} (A_{i-k}^{j-i}A^{j-i} - A_{j-i}^{j-i}A^{i-k})x + (\zeta^{ij})'A\zeta^{ij}(x_k^* - x_k) \\ &= \frac{1}{A_{j-i}^{j-i}} (-A_{j-k}^{j-i}A_i - A_{i-k}^{i-j}A_j + A_{j-i}^{j-i}A_k)x + \frac{1}{A_{j-i}^{j-i}}(v^{ij})'\zeta^{ij}(x_k^* - x_k) \end{aligned}$$

$$\begin{aligned}
 &= (v^{ij})'((x_k^* - x_k)\zeta^{ij} - x) \\
 &= (v^{ij})'\left(\frac{x_k^* - x_k}{x_k^*}(x^{ij} - x^*) - x\right) \\
 &= (v^{ij})'\left(\frac{x_k^* - x_k}{x_k^*}x^{ij} - x\right).
 \end{aligned}$$

The proof of Lemma 13 shows that $y^{jk} = x^{ij} + a(e_j - e_i)$ for some $a > 0$. Thus since $y_k^{jk} = 0$, we find that

$$(DV^3(y^{jk}) - (Ay^{jk})')(e_k - e_i) = (v^{ij})'(y^{jk} - x^{ij}) = a(v^{ij})'(e_j - e_i) = aQ > 0,$$

where the final equality follows from (59). This concludes the verification of the HJB equation at states where V^* is smooth, and so completes the proof of Proposition 14.

REFERENCES

Alós-Ferrer, Carlos and Nick Netzer (2010), “The logit response dynamics.” *Games and Economic Behavior*, 68, 413–427. [280]

Anderson, Simon P, Andre dePalma, and Jacques-Francois Thisse (1992), *Discrete Choice Theory of Product Differentiation*. MIT Press, Cambridge, Massachusetts. [289]

Binmore, Kenneth and Larry Samuelson (1997), “Muddling through: Noisy equilibrium selection.” *Journal of Economic Theory*, 74, 235–265. [284, 328]

Binmore, Kenneth G., Larry Samuelson, and Richard Vaughan (1995), “Musical chairs: Modeling noisy evolution.” *Games and Economic Behavior*, 11, 1–35; erratum: 21 (1997), 325. [284]

Blume, Lawrence E. (1993), “The statistical mechanics of strategic interaction.” *Games and Economic Behavior*, 5, 387–424. [280, 283, 288, 324]

Blume, Lawrence E. (1997), “Population games.” In *The Economy as an Evolving Complex System II* (W. Brian Arthur, Steven N. Durlauf, and David Lane, eds.), 425–460, Addison-Wesley, Reading, Massachusetts. [280, 283, 324]

Blume, Lawrence E. (2003), “How noise matters.” *Games and Economic Behavior*, 44, 251–271. [280, 284, 288]

Boltysanskii, Vladimir G. (1966), “Sufficient conditions for optimality and the justification of the dynamic programming method.” *SIAM Journal on Control*, 4, 326–361. [283, 309, 310]

Catoni, Olivier (1999), “Simulated annealing algorithms and Markov chains with rare transitions.” In *Séminaire de Probabilités XXXIII* (J. Azéma, M. Émery, M. Ledoux, and M. Yor, eds.), 69–119, Springer, Berlin. [280, 293, 294, 305]

Cho, In-Koo, Noah Williams, and Thomas J. Sargent (2002), “Escaping Nash inflation.” *Review of Economic Studies*, 69, 1–40. [283]

Dokumacı, Emin and William H. Sandholm (2011), “Large deviations and multinomial probit choice.” *Journal of Economic Theory*, 146, 2151–2158. [288, 289]

Dupuis, Paul (1988), “Large deviations analysis of some recursive algorithms with state dependent noise.” *Annals of Probability*, 16, 1509–1536. [328]

Dupuis, Paul and Richard S. Ellis (1997), *A Weak Convergence Approach to the Theory of Large Deviations*. Wiley, New York. [328]

Ellison, Glenn (2000), “Basins of attraction, long run equilibria, and the speed of step-by-step evolution.” *Review of Economic Studies*, 67, 17–45. [280, 284, 293]

Foster, Dean P. and H. Peyton Young (1990), “Stochastic evolutionary game dynamics.” *Theoretical Population Biology*, 38, 219–232. [280]

Freidlin, Mark I. and Alexander D. Wentzell (1998), *Random Perturbations of Dynamical Systems*, second edition. Springer, New York. [280, 293, 294, 304]

Fudenberg, Drew and Lorenz A. Imhof (2006), “Imitation processes with small mutations.” *Journal of Economic Theory*, 131, 251–262. [284]

Fudenberg, Drew and Lorenz A. Imhof (2008), “Monotone imitation dynamics in large populations.” *Journal of Economic Theory*, 140, 229–245. [284]

Gilboa, Itzhak and Akihiko Matsui (1991), “Social stability and equilibrium.” *Econometrica*, 59, 859–867. [295]

Hofbauer, Josef (1985), “The selection mutation equation.” *Journal of Mathematical Biology*, 23, 41–53. [308]

Hofbauer, Josef and William H. Sandholm (2002), “On the global convergence of stochastic fictitious play.” *Econometrica*, 70, 2265–2294. [289]

Hofbauer, Josef and Karl Sigmund (1988), *The Theory of Evolution and Dynamical Systems: Mathematical Aspects of Selection*. Cambridge University Press, Cambridge. [319]

Kandori, Michihiro and Rafael Rob (1995), “Evolution of equilibria in the long run: A general theory and applications.” *Journal of Economic Theory*, 65, 383–414. [280, 284, 291, 294, 296, 326]

Kandori, Michihiro and Rafael Rob (1998), “Bandwagon effects and long run technology choice.” *Games and Economic Behavior*, 22, 30–60. [280, 282, 284, 291, 296, 323, 326]

Kandori, Michihiro, George J. Mailath, and Rafael Rob (1993), “Learning, mutation, and long run equilibria in games.” *Econometrica*, 61, 29–56. [280, 289, 294, 322]

Myatt, David P. and Chris C. Wallace (2003), “A multinomial probit model of stochastic evolution.” *Journal of Economic Theory*, 113, 286–301. [280, 289]

Piccoli, Benedetto and Héctor J. Sussmann (2000), “Regular synthesis and sufficiency conditions for optimality.” *SIAM Journal on Control and Optimization*, 39, 359–410. [283, 309, 310]

- Roth, Grégory and William H. Sandholm (2013), “Stochastic approximations with constant step size and differential inclusions.” *SIAM Journal on Control and Optimization*, 51, 525–555. [296]
- Sandholm, William H. (2001), “Potential games with continuous player sets.” *Journal of Economic Theory*, 97, 81–108. [319]
- Sandholm, William H. (2007), “Simple formulas for stationary distributions and stochastically stable states.” *Games and Economic Behavior*, 59, 154–162. [280]
- Sandholm, William H. (2009), “Large population potential games.” *Journal of Economic Theory*, 144, 1710–1725. [308, 319]
- Sandholm, William H. (2010a), “Decompositions and potentials for normal form games.” *Games and Economic Behavior*, 70, 446–456. [307]
- Sandholm, William H. (2010b), “Orders of limits for stationary distributions, stochastic dominance, and stochastic stability.” *Theoretical Economics*, 5, 1–26. [280, 284, 288, 328]
- Sandholm, William H. (2010c), *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge, Massachusetts. [280, 283, 294, 314, 324, 325]
- Sandholm, William H. (2012), “Stochastic imitative game dynamics with committed agents.” *Journal of Economic Theory*, 147, 2056–2071. [284, 328]
- Sandholm, William H. and Mathias Staudigl (2015), “A sample path large deviation principle for a class of population processes.” Unpublished manuscript, University of Wisconsin and Maastricht University. [328]
- Schättler, Heinz and Urszula Ledzewicz (2012), *Geometric Optimal Control: Theory, Methods, and Examples*. Springer, New York. [310]
- Staudigl, Mathias (2012), “Stochastic stability in asymmetric binary choice coordination games.” *Games and Economic Behavior*, 75, 372–401. [280, 284]
- Williams, Noah (2014), “Escape dynamics in learning models.” Unpublished manuscript, University of Wisconsin. [283]
- Young, H. Peyton (1993), “The evolution of conventions.” *Econometrica*, 61, 57–84. [280, 294]
- Young, H. Peyton (1998), *Individual Strategy and Social Structure*. Princeton University Press, Princeton, New Jersey. [280, 294]
- Young, H. Peyton (2004), *Strategic Learning and Its Limits*. Oxford University Press, Oxford. [280]