

Sample Path Large Deviations for Stochastic Evolutionary Game Dynamics

Citation for published version (APA):

Sandholm, W. H., & Staudigl, M. (2018). Sample Path Large Deviations for Stochastic Evolutionary Game Dynamics. *Mathematics of Operations Research*, 43(4), 1348-1377. <https://doi.org/10.1287/moor.2017.0908>

Document status and date:

Published: 01/11/2018

DOI:

[10.1287/moor.2017.0908](https://doi.org/10.1287/moor.2017.0908)

Document Version:

Publisher's PDF, also known as Version of record

Document license:

Taverne

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Sample Path Large Deviations for Stochastic Evolutionary Game Dynamics

William H. Sandholm,^a Mathias Staudigl^b

^aDepartment of Economics, University of Wisconsin–Madison, Madison, Wisconsin 53706; ^bDepartment of Quantitative Economics, Maastricht University, NL-6200 MD Maastricht, Netherlands

Contact: whs@ssc.wisc.edu,  <http://orcid.org/0000-0003-4731-419X> (WHS); mathias.staudigl@gmail.com,

 <http://orcid.org/0000-0003-2481-0019> (MS)

Received: May 16, 2016

Revised: February 8, 2017; August 9, 2017

Accepted: September 19, 2017

Published Online in Articles in Advance:
July 27, 2018

MSC2010 Subject Classification: 60F10,
60J10, 60J20, 90B20, 91A22, 92D25

OR/MS Subject Classification: Games,
noncooperative; probability; Markov processes;
probability, stochastic model applications

<https://doi.org/10.1287/moor.2017.0908>

Copyright: © 2018 INFORMS

Abstract. We study a model of stochastic evolutionary game dynamics in which the probabilities that agents choose suboptimal actions are dependent on payoff consequences. We prove a sample path large deviation principle, characterizing the rate of decay of the probability that the sample path of the evolutionary process lies in a prespecified set as the population size approaches infinity. We use these results to describe excursion rates and stationary distribution asymptotics in settings where the mean dynamic admits a globally attracting state, and we compute these rates explicitly for the case of logit choice in potential games.

Funding: Financial support from the National Science Foundation [Grants SES-1155135, SES-1458992, and SES-1728853], U.S. Army Research Office [Grant MSN201957], and U.S. Air Force Office of Scientific Research [Grant FA9550-09-0538] is gratefully acknowledged.

Keywords: sample path large deviations • evolutionary game theory • stochastic stability • Markov chains • potential games • congestion games

1. Introduction

Evolutionary game theory concerns the dynamics of aggregate behavior of populations of strategically interacting agents. To define these dynamics, one specifies the population size N , the game being played, and the revision protocol agents follow when choosing new actions. Together, these objects define a Markov chain X^N over the set of population states—that is, of distributions of agents across actions.

From this common origin, analyses of evolutionary game dynamics generally proceed in one of two directions. One possibility is to consider deterministic dynamics, which describe the evolution of aggregate behavior using an ordinary differential equation. More precisely, a *deterministic evolutionary dynamic* is a map that assigns population games to dynamical systems on the set of population states. The replicator dynamic (Taylor and Jonker [43]), the best-response dynamic (Gilboa and Matsui [21]), and the logit dynamic (Fudenberg and Levine [20]) are prominent examples.

To derive deterministic dynamics from the Markovian model of individual choice posited above, one can consider the limiting behavior of the Markov chains as the population size N approaches infinity. Kurtz [28], Benaïm [4], Benaïm and Weibull [5, 6], and Roth and Sandholm [33] show that if this limit is taken, then over any finite time span, it becomes arbitrarily likely that the Markov chain is very closely approximated by solutions to a differential equation—the *mean dynamic*—defined by the Markov chain's expected increments. Different revision protocols generate different deterministic dynamics: for instance, the replicator dynamic can be obtained from a variety of protocols based on imitation, whereas the best response and logit dynamics are obtained from protocols based on exact and perturbed optimization, respectively.¹ These deterministic dynamics describe typical behavior in a large population, specifying how the population settles upon a stable equilibrium, a stable cycle, or a more complex stable set. They thus provide theories of how equilibrium behavior is attained or of how it may fail to be attained.

At the same time, if the process X^N is ergodic—for instance, if there is always a small chance of a revising player choosing any available action—then any stable state or other stable set of the mean dynamic is only temporarily so: the equilibrium must break down, and new equilibria must emerge. Behavior over very long time spans is summarized by the stationary distribution of the process. This distribution is typically concentrated near a single stable set, the identity of which is determined by the relative probabilities of transitions between stable sets.

This last question is the subject of the large literature on stochastic stability under evolutionary game dynamics.² The most commonly employed framework in this literature is that of Kandori et al. [27] and Kandori and Rob [25, 26]. These authors consider a population of fixed size, and they suppose that agents employ the *best response with mutations* rule: with high probability, a revising agent plays an optimal action, and with the complementary probability, the agent chooses an action uniformly at random. They then study the long-run behavior of the stochastic game dynamic as the probability of mutation approaches zero. The assumption that all mistakes are equally likely makes the question of equilibrium breakdown simple to answer, as the unlikelihood of a given sample path depends only on the number of suboptimal choices it entails. This eases the determination of the stationary distribution, which is accomplished by means of the well-known Markov chain tree theorem.³

To connect these two branches of the literature, one can consider the questions of equilibrium breakdown and stochastic stability in the large population limit, describing the behavior of the processes X^N when this behavior differs substantially from that of the mean dynamic. Taking a key first step in this direction, this paper establishes a *sample path large deviation principle*: for any prespecified set of sample paths Φ of a fixed duration, we characterize the rate of decay of the probability that the sample path of X^N lies in Φ as N grows large. This large deviation principle is the basic preliminary to obtaining characterizations of the expected waiting times before transitions between equilibria and of stationary distribution asymptotics.

As we noted earlier, most work in stochastic evolutionary game theory has focused on evolution under the best response with mutations rule and on properties of the small noise limit. In some contexts, it seems more realistic to follow the approach taken here, in which the probabilities of mistakes depend on their costs.⁴ Concerning the choice of limits, Binmore and Samuelson [8] argue that the large population limit is more appropriate than the small noise limit for most economic modeling. However, the technical demands of this approach have restricted previous analyses to the two-action case.⁵ This paper provides a necessary first step toward obtaining tractable analyses of large population limits in many-action environments.

To move from the large deviation principle to statements about the long-run behavior of the stochastic process, one can adapt the analyses of Freidlin and Wentzell [19] of diffusions with vanishing noise parameters to our setting of sequences of Markov chains running on increasingly fine grids in the simplex. In Section 4, we explain how the large deviation principle can be used to estimate the waiting times to reach sets of states away from an attractor and to describe the asymptotic behavior of the stationary distribution in cases where the mean dynamic admits a globally attracting state. We prove that when agents playing a potential game make decisions using the logit choice rule, the control problems in the statement of the large deviation principle can be solved explicitly. We illustrate the implications of these results by using them to characterize long-run behavior in a model of traffic congestion.

Our work here is closely connected to developments in two branches of the stochastic processes literature. Large deviation principles for environments quite close to those considered here have been established by Azencott and Ruget [1], Dupuis [16], and Dupuis and Ellis [17]. In these works, the sequences of processes under consideration are defined on open sets in \mathbb{R}^n , and have transition laws that allow for motion in all directions from every state. These results do not apply directly to the evolutionary processes considered here, which necessarily run on a compact set. Thus relative to these works, our contribution lies in addressing behavior at and near boundary states.

There are also close links to work on interacting particle systems with long-range interactions. In game-theoretic terms, the processes studied in this literature describe the *individual choices* of each of N agents as they evolve in continuous time, with the stochastic changes in each agent's action being influenced by the aggregate behavior of all agents. Two large deviation principles for such systems are proved by Léonard [29]. The first describes large deviations of the sequence of probability distributions on the set of empirical distributions, where the latter distributions anonymously describe the N agents' sample paths through the finite set of actions $\mathcal{A} = \{1, \dots, n\}$.⁶ The second describes large deviations of the sequence of probability distributions over paths on discrete grids X^N in the n -simplex, paths that represent the evolution of *aggregate* behavior in the N agent particle system.⁷ The Freidlin–Wentzell theory for particle systems with long-range interactions has been developed by Borkar and Sundaresan [13], who provide many further references to this literature.

The large deviation principle we prove here is a discrete-time analogue of the second result of Léonard [29] noted above. Unlike Léonard [29], we allow individuals' transition probabilities to depend in a vanishing way on the population size, as is natural in our game-theoretic context (see Examples 1–4). Also, our discrete-time framework obviates the need to address large deviations in the arrival of revision opportunities.⁸ But the central advantage of our approach is its simplicity. Describing the evolution of the choices of each of N individual agents requires a complicated stochastic processes. Understanding the proofs (and even the statements) of large

deviation principles for these processes requires substantial background knowledge. Here, our interest is in aggregate behavior. By making the aggregate behavior process our primitive, we are able to state our large deviation principle with a minimum of preliminaries. Likewise, our proof of this principle, which follows the weak convergence approach of Dupuis and Ellis [17], is relatively direct, and in Section 5 we explain its main ideas in a straightforward manner. These factors may make the work to follow accessible to researchers in economics, biology, engineering, and other fields.

This paper is part of a larger project on large deviations and stochastic stability under evolutionary game dynamics with payoff-dependent mistake probabilities and arbitrary numbers of actions. In Sandholm and Staudigl [39], we considered the case of the small noise double limit, in which the noise level in agents' decisions is first taken to zero and then the population size to infinity. The initial analysis of the small noise limit concerns a sequence of Markov chains on a fixed finite state space; the relevant characterizations of large deviations properties in terms of discrete minimization problems are simple and well known. Taking the second limit as the population size grows large turns these discrete minimization problems into continuous optimal control problems. We show that the latter problems possess a linear structure that allows them to be solved analytically.

The present paper begins the analysis of large deviations and stochastic stability when the population size is taken to infinity for a fixed noise level. This analysis concerns a sequence of Markov chains on ever finer grids in the simplex, making the basic large deviations result—our main result here—considerably more difficult than its small-noise counterpart. Future work will provide a full development of the Freidlin–Wentzell theory for the large population limit, allowing for mean dynamics with multiple stable states. It will then introduce the second limit as the noise level vanishes and determine the extent to which the agreement of the two double limits agree. Further discussion of this research agenda is offered in Section 4.3.

2. The Model

We consider a model in which all agents are members of a single population. The extension to multipopulation settings requires only more elaborate notation.

2.1. Finite-Population Games

We consider games in which the members of a population of N agents choose actions from the common finite action set $\mathcal{A} = \{1, \dots, n\}$. We describe the population's aggregate behavior 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 action i .

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 action i when the population state is $x \in \mathcal{X}^N$.

Example 1 (Matching in Normal Form Games). Assume that agents are matched in pairs to play a symmetric two-player normal-form game $A \in \mathbb{R}^{n \times n}$, where A_{ij} is the payoff obtained by an i player who is matched with a j player. If each agent is matched with all other agents (but not with himself), then average payoffs in the resulting population game are given by $F_i^N(x) = (1/(N-1))(A(Nx - e_i))_i = (Ax)_i + (1/(N-1))((Ax)_i - A_{ii})$. \square

Example 2 (Congestion Games). To define a *congestion game* (Beckmann et al. [2], Rosenthal [32]), one specifies a collection of facilities Λ (e.g., links in a highway network) and associates with each facility $\lambda \in \Lambda$ a function $l_\lambda^N: \{0, 1/N, \dots, 1\} \rightarrow \mathbb{R}$ describing the cost (or benefit, if $l_\lambda^N < 0$) of using the facility as a function of the fraction of the population that uses it. Each action $i \in \mathcal{A}$ (a path through the network) requires the facilities in a given set $\Lambda_i \subseteq \Lambda$ (the links on the path), and the payoff to action i is the negative of the sum of the costs accruing from these facilities. Payoffs in the resulting population game are given by $F_i^N(x) = -\sum_{\lambda \in \Lambda_i} l_\lambda^N(u_\lambda(x))$, where $u_\lambda(x) = \sum_{i: \lambda \in \Lambda_i} x_i$ denotes the total utilization of facility λ at state x . \square

Because the population size is finite, the payoff vector an agent considers when revising may depend on his current action. To allow for this possibility, we let $F_{i \rightarrow j}^N: \mathcal{X}^N \rightarrow \mathbb{R}^n$ denote the payoff vector considered at state x by an action i player.

Example 3 (Simple Payoff Evaluation). Under *simple payoff evaluation*, all agents' decisions are based on the current vector of payoffs: $F_{i \rightarrow j}^N(x) = F_j^N(x)$ for all $i, j \in \mathcal{A}$. \square

Example 4 (Clever Payoff Evaluation). Under *clever payoff evaluation*, an action i player accounts for the fact that by switching to action j at state x , he changes the state to the adjacent state $y = x + (1/N)(e_j - e_i)$. To do so, he evaluates payoffs according to the *clever payoff vector* $F_{i \rightarrow j}^N(x) = F_j^N(x + (1/N)(e_j - e_i))$.⁹ \square

As the assumptions in Section 2.4 will make clear, our results are the same whether simple or clever payoff evaluation is assumed.

2.2. Revision Protocols

In our model of evolution, each agent occasionally receives opportunities to switch actions. At such moments, an agent decides which action to play next by employing a *protocol* $\rho^N: \mathbb{R}^n \times \mathcal{X}^N \rightarrow X^n$, with the choice probabilities of a current action i player being described by $\rho_{ij}^N: \mathbb{R}^n \times \mathcal{X}^N \rightarrow X$. Specifically, if a revising action i player faces payoff vector $\pi \in \mathbb{R}^n$ at population state $x \in \mathcal{X}^N$, then the probability that he proceeds by playing action j is $\rho_{ij}^N(\pi, x)$. We will assume shortly that this probability is bounded away from zero, so that there is always a nonnegligible probability of the revising agent playing each of the actions in \mathcal{A} ; see condition (6).

Example 5 (The Logit Protocol). A fundamental example of a revision protocol with positive choice probabilities is the *logit protocol*, defined by

$$\rho_{ij}^N(\pi, x) = \frac{\exp(\eta^{-1}\pi_j)}{\sum_{k \in \mathcal{A}} \exp(\eta^{-1}\pi_k)} \quad (1)$$

for some *noise level* $\eta > 0$. When η is small, an agent using this protocol is very likely to choose an optimal action but places positive probability on every action, with lower probabilities being placed on worse-performing actions. \square

Example 6 (Perturbed Best-Response Protocols). One can generalize (1) by assuming that agents choice probabilities maximize the difference between their expected base payoff and a convex penalty:

$$\rho_{ij}^N(\pi, x) = \arg \max_{x \in \text{int}(X)} \left(\sum_{k \in \mathcal{A}} \pi_k x_k - h(x) \right),$$

where $h: \text{int}(X) \rightarrow \mathbb{R}$ is strictly convex and steep, in the sense that $|\nabla h(x)|$ approaches infinity whenever x approaches the boundary of X . The logit protocol (1) is recovered when h is the negated entropy function $\eta^{-1} \sum_{k \in \mathcal{A}} x_k \log x_k$. \square

Example 7 (The Pairwise Logit Protocol). Under the *pairwise logit protocol*, a revising agent chooses a candidate action at random and then applies the logit rule (1) only to his current action and the candidate action:

$$\rho_{ij}^N(\pi, x) = \frac{\exp(\eta^{-1}\pi_j)}{\exp(\eta^{-1}\pi_i) + \exp(\eta^{-1}\pi_j)}. \quad \square$$

Example 8 (Imitation with “Mutations”). Suppose that with probability $1 - \varepsilon$, a revising agent picks an opponent at random and switches to her action with probability proportional to the opponent’s payoff, and that with probability $\varepsilon > 0$ the agent chooses an action at random. If payoffs are normalized to take values between 0 and 1, the resulting protocol takes the form

$$\rho_{ij}^N(\pi, x) = \begin{cases} (1 - \varepsilon) \frac{N}{N-1} x_j \pi_j + \frac{\varepsilon}{n}, & \text{if } j \neq i, \\ (1 - \varepsilon) \left(\frac{N x_i - 1}{N} + \sum_{k \neq i} \frac{N}{N-1} x_k (1 - \pi_k) \right) + \frac{\varepsilon}{n}, & \text{if } j = i. \end{cases}$$

The positive mutation rate ensures that all actions are chosen with positive probability. \square

For many further examples of revision protocols, see Sandholm [37].

2.3. The Stochastic Evolutionary Process

Together, a population game F^N and a revision protocol ρ^N define a discrete-time stochastic process $\mathbf{X}^N = \{X_k^N\}_{k=0}^\infty$, which is defined informally as follows: During each period, a single agent is selected at random and given a revision opportunity. The probabilities with which he chooses each action are obtained by evaluating the protocol ρ^N at the relevant payoff vector and population state. Each period of the process \mathbf{X}^N takes $1/N$ units of clock time, as this fixes at one the expected number of revision opportunities that each agent receives during one unit of clock time.

More precisely, the process \mathbf{X}^N is a Markov chain with initial condition $X_0^N \in \mathcal{X}^N$ and transition law

$$\mathbb{P}(X_{k+1}^N = y \mid X_k^N = x) = \begin{cases} x_i \rho_{ij}^N(F_{i \rightarrow \cdot}^N(x), x), & \text{if } y = x + \frac{1}{N}(e_j - e_i) \text{ and } j \neq i, \\ \sum_{i=1}^n x_i \rho_{ii}^N(F_{i \rightarrow \cdot}^N(x), x), & \text{if } y = x, \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

When a single agent switches from action i to action j , the population state changes from x to $y = x + (1/N) \cdot (e_j - e_i)$. This requires that the revision opportunity be assigned to a current action i player, which occurs with probability x_i , and that this player choose to play action j , which occurs with probability $\rho_{ij}^N(F_{i \rightarrow \cdot}^N(x), x)$. Together, these yield the law (2).

Example 9. Suppose that N agents are matched to play the normal-form game $A \in \mathbb{R}^{n \times n}$, using clever payoff evaluation and the logit choice rule with noise level $\eta > 0$. Then if the state in period k is $x \in \mathcal{X}^N$, then by Examples 1, 4, and 5 and Equation (2), the probability that the state in period $k + 1$ is $x + (1/N)(e_j - e_i) \neq x$ is equal to¹⁰

$$\mathbb{P}\left(X_{k+1}^N = x + \frac{1}{N}(e_j - e_i) \mid X_k^N = x\right) = x_i \cdot \frac{\exp(\eta^{-1} \cdot (1/(N-1))(A(Nx - e_i))_j)}{\sum_{l \in \mathcal{A}} \exp(\eta^{-1} \cdot (1/(N-1))(A(Nx - e_i))_l)}. \quad \square$$

2.4. A Class of Population Processes

It will be convenient to consider an equivalent class of Markov chains defined using a more parsimonious notation. All Markov chains to come are defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and we sometimes use the notation \mathbb{P}_x to indicate that the Markov chain \mathbf{X}^N under consideration is run from initial condition $x \in \mathcal{X}^N$.

The Markov chain $\mathbf{X}^N = \{X_k^N\}_{k=0}^\infty$ runs on the discrete grid $\mathcal{X}^N = X \cap (1/N)\mathbb{Z}^n$, with each period taking $1/N$ units of clock time, so that each agent expects to receive one revision opportunity per unit of clock time (see Section 3.1). We define the law of \mathbf{X}^N by setting an initial condition $X_0^N \in \mathcal{X}^N$ and specifying subsequent states via the recursion

$$X_{k+1}^N = X_k^N + \frac{1}{N} \zeta_{k+1}^N. \tag{3}$$

The normalized increment ζ_{k+1}^N follows the conditional law $\nu^N(\cdot \mid X_k^N)$, defined by

$$\nu^N(z \mid x) = \begin{cases} x_i \sigma_{ij}^N(x), & \text{if } z = e_j - e_i \text{ and } j \neq i, \\ \sum_{i=1}^n x_i \sigma_{ii}^N(x), & \text{if } z = \mathbf{0}, \\ 0, & \text{otherwise,} \end{cases} \tag{4}$$

where the function $\sigma^N: \mathcal{X}^N \rightarrow \mathbb{R}_+^{n \times n}$ satisfies $\sum_{j \in \mathcal{A}} \sigma_{ij}^N(x) = 1$ for all $i \in \mathcal{A}$ and $x \in \mathcal{X}^N$. The *switch probability* $\sigma_{ij}^N(x)$ is the probability that an agent playing action i who receives a revision opportunity proceeds by playing action j . The model described in the previous sections can be expressed in the present notation as $\sigma_{ij}^N(x) = \rho_{ij}^N(F_{i \rightarrow \cdot}^N(x), x)$.

We observe that the support of the transition measure $\nu^N(\cdot \mid x)$ is contained in the set of *raw increments* $\mathcal{Z} = \{e_j - e_i: i, j \in \mathcal{A}\}$. Since an unused action cannot become less common, the support of $\nu^N(\cdot \mid x)$ is contained in $\mathcal{Z}(x) = \{e_j - e_i: i, j \in \mathcal{A}, x_i > 0\}$.

Our large deviations results concern the behavior of sequences $\{\mathbf{X}_N^N\}_{N=N_0}^\infty$ of Markov chains defined by (3) and (4). To allow for finite population effects, we permit the switch probabilities $\sigma_{ij}^N(x)$ to depend on N in a manner that becomes negligible as N grows large. Specifically, we assume that there is a Lipschitz continuous function $\sigma: \mathcal{X}^N \rightarrow \mathbb{R}_+^{n \times n}$ that describes the limiting switch probabilities, in the sense that

$$\lim_{N \rightarrow \infty} \max_{x \in \mathcal{X}^N} \max_{i, j \in \mathcal{A}} |\sigma_{ij}^N(x) - \sigma_{ij}(x)| = 0. \tag{5}$$

In the game model, this assumption holds when the sequences of population games F^N and revision protocols ρ^N have at most vanishing finite population effects, in that they converge to a limiting population game $F: X \rightarrow \mathbb{R}^n$ and a limiting revision protocol $\rho: \mathbb{R}^n \times X \rightarrow \mathbb{R}$, both of which are Lipschitz continuous.

In addition, we assume that limiting switch probabilities are bounded away from zero: there is a $\zeta > 0$ such that

$$\min_{x \in X} \min_{i, j \in \mathcal{A}} \sigma_{ij}(x) \geq \zeta. \tag{6}$$

This assumption is satisfied in the game model when the choice probabilities $\rho_{ij}^N(\pi, x)$ are bounded away from zero.¹¹ This is so under all of the revision protocols from Section 2.2. Assumption (6) and the transition law (4) imply that the Markov chain \mathbf{X}^N is aperiodic and irreducible for N large enough.

Assumptions (5) and (6) imply that the transition kernels (4) of the Markov chains \mathbf{X}^N approach a limiting kernel $\nu: X \rightarrow \Delta(\mathcal{Z})$, defined by

$$\nu(z | x) = \begin{cases} x_i \sigma_{ij}(x), & \text{if } z = e_j - e_i \text{ and } j \neq i, \\ \sum_{i \in \mathcal{A}} x_i \sigma_{ii}(x), & \text{if } z = \mathbf{0}, \\ 0, & \text{otherwise.} \end{cases} \quad (7)$$

Condition (5) implies that the convergence of ν^N to ν is uniform:

$$\lim_{N \rightarrow \infty} \max_{x \in \mathcal{X}^N} \max_{z \in \mathcal{Z}} |\nu^N(z | x) - \nu(z | x)| = 0. \quad (8)$$

The probability measures $\nu(\cdot | x)$ depend Lipschitz continuously on x , and by virtue of condition (6), each measure $\nu(\cdot | x)$ has support $\mathcal{Z}(x)$.

3. Sample Path Large Deviations

3.1. Deterministic Approximation

Before considering the large deviations properties of the processes \mathbf{X}^N , we describe their typical behavior. By definition, each period of the process $\mathbf{X}^N = \{X_k^N\}_{k=0}^\infty$ takes $1/N$ units of clock time and leads to a random increment of size $1/N$. Thus when N is large, each brief interval of clock time contains a large number of periods during which the transition measures $\nu^N(\cdot | X_k^N)$ vary little. Intuition from the law of large numbers then suggests that over this interval, and hence over any finite concatenation of such intervals, the Markov chain \mathbf{X}^N should follow an almost deterministic trajectory—namely, the path determined by the process’s expected motion.

To make this statement precise, note that the expected increment of the process \mathbf{X}^N from state x during a single period is

$$\mathbb{E}(X_{k+1}^N - X_k^N | X_k^N = x) = \frac{1}{N} \mathbb{E}(\zeta_k^N | X_k^N = x) = \frac{1}{N} \sum_{z \in \mathcal{Z}} z \nu^N(z | x). \quad (9)$$

Since there are N periods per time unit, the expected increment per time unit is obtained by multiplying (9) by N . Doing so and taking N to its limit defines the *mean dynamic*,

$$\dot{x} = \sum_{z \in \mathcal{Z}} z \nu(z | x) = \mathbb{E} \zeta_x, \quad (10a)$$

where ζ_x is a random variable with law $\nu(\cdot | x)$. By construction, $\mathbb{E} \zeta_x$ lies in the set $Z = \text{conv}(\mathcal{Z})$. Substituting definition (7) into (10a) and simplifying yields the coordinate formula:

$$\dot{x}_i = \sum_{j \in \mathcal{A}} x_j \sigma_{ji}(x) - x_i. \quad (10b)$$

Assumption (6) implies that the boundary of the simplex is repelling under (10). Since the right-hand side of (10) is Lipschitz continuous, it admits a unique forward solution $\{x_t\}_{t \geq 0}$ from every initial condition $x_0 = x$ in X , and this solution does not leave X .

A version of the deterministic approximation result to follow was first proved by Kurtz [28], with the exponential rate of convergence established by Benaïm [4]; see also Benaïm and Weibull [5, 6]. To state the result, we let $|\cdot|$ denote the l^1 norm on \mathbb{R}^n , and we define $\hat{\mathbf{X}}^N = \{\hat{X}_t^N\}_{t \geq 0}$ to be the piecewise affine interpolation of the process \mathbf{X}^N :

$$\hat{X}_t^N = X_{\lfloor Nt \rfloor}^N + (Nt - \lfloor Nt \rfloor)(X_{\lfloor Nt \rfloor + 1}^N - X_{\lfloor Nt \rfloor}^N).$$

Theorem 1. *Suppose that $\{X_k^N\}$ has initial condition $x^N \in \mathcal{X}^N$, and let $\lim_{N \rightarrow \infty} x^N = x \in X$. Let $\{x_t\}_{t \geq 0}$ be the solution to (10) with $x_0 = x$. For any $T < \infty$ there exists a constant $c > 0$ independent of x such that for all $\varepsilon > 0$ and N large enough,*

$$\mathbb{P}_{x^N} \left(\sup_{t \in [0, T]} |\hat{X}_t^N - x_t| \geq \varepsilon \right) \leq 2n \exp(-c\varepsilon^2 N).$$

3.2. The Cramér Transform and Relative Entropy

Stating our large deviations results requires some additional machinery.¹² Let $\mathbb{R}_0^n = \{z \in \mathbb{R}^n : \sum_{i \in \mathcal{A}} z_i = 0\}$ denote the set of vectors tangent to the simplex. The *Cramér transform* $L(x, \cdot) : \mathbb{R}_0^n \rightarrow [0, \infty]$ of probability distribution $\nu(\cdot | x) \in \Delta(\mathcal{Z})$ is defined by

$$L(x, z) = \sup_{u \in \mathbb{R}_0^n} (\langle u, z \rangle - H(x, u)), \quad \text{where } H(x, u) = \log \left(\sum_{z \in \mathcal{Z}} e^{\langle u, z \rangle} \nu(z | x) \right). \tag{11}$$

In other words, $L(x, \cdot)$ is the convex conjugate of the log moment generating function of $\nu(\cdot | x)$. It is well known that $L(x, \cdot)$ is convex, lower semicontinuous, and nonnegative and that $L(x, z) = 0$ if and only if $z = \mathbb{E}\zeta_x$; moreover, $L(x, z) < \infty$ if and only if $z \in Z(x)$, where $Z(x) = \text{conv}(\mathcal{Z}(x))$ is the convex hull of the support of $\nu(\cdot | x)$.¹³

To help interpret what is to come, we recall *Cramér’s theorem*: Let $\{\zeta_x^k\}_{k=1}^\infty$ be a sequence of independent and identically distributed (i.i.d.) random variables with law $\nu(\cdot | x)$, and let $\bar{\zeta}_x^N$ be the sequence’s N th sample mean. Then for any set $V \subseteq Z(x)$,

$$\limsup_{N \rightarrow \infty} \frac{1}{N} \log \mathbb{P}(\bar{\zeta}_x^N \in V) \leq - \inf_{z \in \text{cl}(V)} L(x, z), \quad \text{and} \tag{12a}$$

$$\liminf_{N \rightarrow \infty} \frac{1}{N} \log \mathbb{P}(\bar{\zeta}_x^N \in V) \geq - \inf_{z \in \text{int}(V)} L(x, z), \tag{12b}$$

where $\text{cl}(V)$ and $\text{int}(V)$ denote the relative closure of V and the relative interior of V , respectively. Thus for “nice” sets V , those for which the right-hand sides of the upper and lower bounds (12a) and (12b) are equal, this common value is the exponential rate of decay of the probability that $\bar{\zeta}_x^N$ lies in V .

Our analysis relies heavily on a well-known characterization of the Cramér transform as a constrained minimum of relative entropy, a characterization that also provides a clear intuition for Cramér’s theorem. Recall that the *relative entropy* of probability measure $\lambda \in \Delta(\mathcal{Z})$ given probability measure $\pi \in \Delta(\mathcal{Z})$ is the extended real number

$$R(\lambda || \pi) = \sum_{z \in \mathcal{Z}} \lambda(z) \log \frac{\lambda(z)}{\pi(z)},$$

with the conventions that $0 \log 0 = 0 \log \frac{0}{0} = 0$. It is well known that $R(\cdot || \cdot)$ is convex, lower semicontinuous, and nonnegative, that $R(\lambda || \pi) = 0$ if and only if $\lambda = \pi$, and that $R(\lambda || \pi) < \infty$ if and only if $\text{supp}(\lambda) \subseteq \text{supp}(\pi)$.

A basic interpretation of relative entropy is provided by *Sanov’s theorem*,¹⁴ which concerns the asymptotics of the empirical distributions \mathcal{E}_x^N of the sequence $\{\zeta_x^k\}_{k=1}^\infty$, defined by $\mathcal{E}_x^N(z) = (1/N) \sum_{k=1}^N 1(\zeta_x^k = z)$. This theorem says that for every set of distributions $\Lambda \subseteq \Delta(\mathcal{Z})$,

$$\limsup_{N \rightarrow \infty} \frac{1}{N} \log \mathbb{P}(\mathcal{E}_x^N \in \Lambda) \leq - \inf_{\lambda \in \text{cl}(\Lambda)} R(\lambda || \nu(\cdot | x)), \quad \text{and} \tag{13a}$$

$$\liminf_{N \rightarrow \infty} \frac{1}{N} \log \mathbb{P}(\mathcal{E}_x^N \in \Lambda) \geq - \inf_{\lambda \in \text{int}(\Lambda)} R(\lambda || \nu(\cdot | x)). \tag{13b}$$

Thus for “nice” sets Λ , the probability that the empirical distribution lies in Λ decays at an exponential rate given by the minimal value of relative entropy on Λ . Displays (13a) and (13b) enable us to interpret $R(\lambda || \nu(\cdot | x))$ as an asymmetric measure of the “distance” from probability measure $\nu(\cdot | x)$ to probability measure λ .

The intuition behind Sanov’s theorem and relative entropy is straightforward. We can express the probability that the N th empirical distribution is the feasible distribution $\lambda \in \Lambda$ as the product of the probability of obtaining a particular realization of $\{\zeta_x^k\}_{k=1}^\infty$ with empirical distribution λ and the number of such realizations:

$$\mathbb{P}(\mathcal{E}_x^N = \lambda) = \prod_{z \in \mathcal{Z}} \nu(z | x)^{N\lambda(z)} \times \frac{N!}{\prod_{z \in \mathcal{Z}} (N\lambda(z))!}.$$

Then applying Stirling’s approximation $n! \approx n^n e^{-n}$ yields

$$\frac{1}{N} \log \mathbb{P}(\mathcal{E}_x^N = \lambda) \approx \sum_{z \in \mathcal{Z}} \lambda(z) \log \nu(z | x) - \sum_{z \in \mathcal{Z}} \lambda(z) \log \lambda(z) = -R(\lambda || \nu(\cdot | x)).$$

The rate of decay of $\mathbb{P}(\mathcal{E}_x^N \in \Lambda)$ is then determined by the “most likely” empirical distribution in Λ : that is, by the one whose relative entropy is smallest.¹⁵

The representation of the Cramér transform in terms of relative entropy is obtained by a variation on the final step above: given Sanov’s theorem, the rate of decay of obtaining a sample mean \bar{c}_x^N in $V \subset \mathbb{R}^n$ should be determined by the smallest relative entropy associated with a probability distribution whose mean lies in V .¹⁶ Combining this idea with (12a) and (12b) suggests the representation¹⁷

$$L(x, z) = \min_{\lambda \in \Delta(z)} \left\{ R(\lambda \| \nu(\cdot | x)) : \sum_{z \in Z} z \lambda(z) = z \right\}. \quad (14)$$

If $z \in Z(x)$, so that $L(x, z) < \infty$, then the minimum in (14) is attained uniquely.

3.3. Path Costs

To state the large deviation principle for the sequence of interpolated processes $\{\hat{X}^N\}_{N=N_0}^\infty$, we must introduce a function that characterizes the rates of decay of the probabilities of sets of sample paths through the simplex X . Doing so requires some preliminary definitions. For $T \in (0, \infty)$, let $\mathcal{C}[0, T]$ denote the set of continuous paths $\phi: [0, T] \rightarrow X$ through X over time interval $[0, T]$, endowed with the supremum norm. Let $\mathcal{C}_x[0, T]$ denote the set of such paths with initial condition $\phi_0 = x$, and let $\mathcal{AC}_x[0, T]$ be the set of absolutely continuous paths in $\mathcal{C}_x[0, T]$.

We define the *path cost function* (or *rate function*) $c_{x,T}: \mathcal{C}[0, T] \rightarrow [0, \infty]$ by

$$c_{x,T}(\phi) = \begin{cases} \int_0^T L(\phi_t, \dot{\phi}_t) dt, & \text{if } \phi \in \mathcal{AC}_x[0, T], \\ \infty, & \text{otherwise.} \end{cases} \quad (15)$$

By Cramér’s theorem, $L(\phi_t, \dot{\phi}_t)$ describes the “difficulty” of proceeding from state ϕ_t in direction $\dot{\phi}_t$ under the transition laws of our Markov chains. Thus the path cost $c_{x,T}(\phi)$ represents the *difficulty* of following the entire path ϕ . Since $L(x, z) = 0$ if and only if $z = \mathbb{E}\zeta_x$, path $\phi \in \mathcal{C}_x[0, T]$ satisfies $c_{x,T}(\phi) = 0$ if and only if it is the solution to the mean dynamic (10) from state x . In light of definition (15), we sometimes refer to the function $L: X \times \mathbb{R}_0^n \rightarrow [0, \infty]$ as the *running cost function*.

As illustrated by Cramér’s theorem, the rates of decay described by large deviation principles are defined in terms of the smallest value of a function over the set of outcomes in question. This makes it important for such functions to satisfy lower semicontinuity properties. The following result, which follows Dupuis and Ellis [17, proposition 6.2.4], provides such a property.

Proposition 1. *The function $c_{x,T}$ is a (good) rate function: its lower-level sets $\{\phi \in \mathcal{C}: c_{x,T}(\phi) \leq M\}$ are compact.*

3.4. A Sample Path Large Deviation Principle

Our main result, Theorem 2, shows that the sample paths of the interpolated processes \hat{X}^N satisfy a large deviation principle with rate function (15). To state this result, we use the notation $\hat{X}_{[0,T]}^N$ as shorthand for $\{\hat{X}_t^N\}_{t \in [0,T]}$.

Theorem 2. *Suppose that the processes $\{\hat{X}^N\}_{N=N_0}^\infty$ have initial conditions $x^N \in X^N$ satisfying $\lim_{N \rightarrow \infty} x^N = x \in X$. Let $\Phi \subseteq \mathcal{C}[0, T]$ be a Borel set. Then*

$$\limsup_{N \rightarrow \infty} \frac{1}{N} \log \mathbb{P}_{x^N}(\hat{X}_{[0,T]}^N \in \Phi) \leq - \inf_{\phi \in \text{cl}(\Phi)} c_{x,T}(\phi), \quad \text{and} \quad (16a)$$

$$\liminf_{N \rightarrow \infty} \frac{1}{N} \log \mathbb{P}_{x^N}(\hat{X}_{[0,T]}^N \in \Phi) \geq - \inf_{\phi \in \text{int}(\Phi)} c_{x,T}(\phi). \quad (16b)$$

We refer to inequality (16a) as the *large deviation principle upper bound* and to (16b) as the *large deviation principle lower bound*.

While Cramér’s theorem concerns the probability that the sample mean of N i.i.d. random variables lies in a given subset of \mathbb{R}^n as N grows large, Theorem 2 concerns the probability that the sample path of the process $\hat{X}_{[0,T]}^N$ lies in a given subset of $\mathcal{C}[0, T]$ as N grows large. If $\Phi \subseteq \mathcal{C}[0, T]$ is a set of paths for which the infima in (16a) and (16b) are equal and attained at some path ϕ^* , then Theorem 2 shows that the probability that the sample path of $\hat{X}_{[0,T]}^N$ lies in Φ is of order $\exp(-Nc_{x,T}(\phi^*))$.

3.5. Uniform Results

Applications to the Freidlin–Wentzell [19] theory (see Section 4.3) require uniform versions of the previous two results, allowing for initial conditions $x^N \in \mathcal{X}^N$ that take values in compact subsets of X . We therefore note the following extensions of Proposition 1 and Theorem 2.

Proposition 2. *For any compact set $K \subseteq X$ and any $M < \infty$, the sets $\bigcup_{x \in K} \{\phi \in \mathcal{C}[0, T]: c_{x, T}(\phi) \leq M\}$ are compact.*

Theorem 3. *Let $\Phi \subseteq \mathcal{C}[0, T]$ be a Borel set. For every compact set $K \subseteq X$,*

$$\limsup_{N \rightarrow \infty} \frac{1}{N} \log \left(\sup_{x^N \in K \cap \mathcal{X}^N} \mathbb{P}_{x^N}(\hat{\mathbf{X}}_{[0, T]}^N \in \Phi) \right) \leq - \inf_{x \in K} \inf_{\phi \in \text{cl}(\Phi)} c_{x, T}(\phi), \quad \text{and} \tag{17a}$$

$$\liminf_{N \rightarrow \infty} \frac{1}{N} \log \left(\inf_{x^N \in K \cap \mathcal{X}^N} \mathbb{P}_{x^N}(\hat{\mathbf{X}}_{[0, T]}^N \in \Phi) \right) \geq - \sup_{x \in K} \inf_{\phi \in \text{int}(\Phi)} c_{x, T}(\phi). \tag{17b}$$

The proof of Proposition 2 is an easy extension of that of Dupuis and Ellis [17, proposition 6.2.4]; compare Dupuis and Ellis [17, p. 165], where the property being established is called *compact level sets uniformly on compacts*. Theorem 3, the *uniform large deviation principle*, follows from Theorem 2 and an elementary compactness argument; compare Dembo and Zeitouni [15, proof of corollary 5.6.15].

4. Applications

To be used in most applications, the results above must be combined with ideas from Freidlin–Wentzell theory. In this section, we use the large deviation principle to study the frequency of excursions from a globally attracting rest point and the asymptotics of the stationary distribution, with a full analysis of the case of logit choice in potential games. We then remark on future applications to games with multiple stable equilibria and the wider scope for analytical solutions that may arise by introducing a second limit.

4.1. Excursions from a Globally Attracting State and Stationary Distributions

In this section, we describe results on excursions from globally attracting rest points that can be obtained by combining the results above with the work of Freidlin and Wentzell [19] and refinements attributable to Dembo and Zeitouni [15], who consider this question in the context of diffusion processes with a vanishingly small noise parameter. To prove the results in our setting, one must adapt arguments for diffusions to sequences of processes running on increasingly fine finite state spaces. As an illustration of the difficulties involved, observe that although a diffusion process is a Markov process with continuous sample paths, our original process \mathbf{X}^N is Markov but with discrete sample paths, whereas our interpolated process $\hat{\mathbf{X}}^N$ has continuous sample paths but is not Markov.¹⁸ Since neither process possesses both desirable properties, a complete analysis of the problem is quite laborious. We therefore only sketch the main arguments here and will present a detailed analysis in future work.

Consider a sequence of processes \mathbf{X}^N satisfying the assumptions above and whose mean dynamic (10) has a globally attracting rest point x^* , meaning that solutions to (10) from every initial condition in X converge to x^* . We would like to estimate the time until the process exits a given open set $O \subset X$ containing x^* . By the large deviations logic described in Section 3.2, we expect that this time should be determined by the cost of the least cost path that starts at x^* and leaves O . With this in mind, let ∂O denote the boundary of O relative to X , and define

$$C_y = \inf_{T > 0} \inf_{\phi \in \mathcal{C}_{x^*, [0, T]}: \phi(T) = y} c_{x^*, T}(\phi), \quad \text{for } y \in X, \tag{18}$$

$$C_{\partial O} = \inf_{y \in \partial O} C_y. \tag{19}$$

Thus C_y is the lowest cost of a path from x^* to y , and the *exit cost* $C_{\partial O}$ is the lowest cost of a path that leaves O .

Now define $\hat{\tau}_{\partial O}^N = \inf\{t \geq 0: \hat{X}_t^N \in \partial O\}$ to be the random time at which the interpolated process $\hat{\mathbf{X}}^N$ hits ∂O of O . If this boundary satisfies a mild regularity condition,¹⁹ we can show that for all $\varepsilon > 0$ and all sequences of $x^N \in \mathcal{X}^N$ converging to some $x \in O$, we have

$$\lim_{N \rightarrow \infty} \mathbb{P}_{x^N} \left(C_{\partial O} - \varepsilon < \frac{1}{N} \log \hat{\tau}_{\partial O}^N < C_{\partial O} + \varepsilon \right) = 1 \quad \text{and} \tag{20}$$

$$\lim_{N \rightarrow \infty} \frac{1}{N} \log \mathbb{E}_{x^N} \hat{\tau}_{\partial O}^N = C_{\partial O}. \tag{21}$$

That is, the time until exit from O is of approximate order $\exp(NC_{\partial O})$ with probability near 1, and the expected time until exit from O is of this order as well. Since stationary distribution weights are inversely proportional to expected return times, Equation (21) can be used to show that the rates of decay of stationary distribution weights are also determined by minimal costs of paths. If we let $B_\delta(y) = \{x \in X: |y - x| < \delta\}$, then for all $y \in X$ and $\varepsilon > 0$, there is a δ sufficiently small that

$$-C_y - \varepsilon \leq \frac{1}{N} \log \mu^N(B_\delta(y)) \leq -C_y + \varepsilon \quad (22)$$

for all large enough N .

The main ideas of the proofs of (20) and (21) are as follows. To prove the upper bounds, we use the large deviation principle (LDP) lower bound to show that there is a finite duration T such that the probability of reaching ∂O in T time units starting from any state in O is at least $q_T^N = \exp(-N(C_{\partial O} + \varepsilon))$. It then follows from the strong Markov property that the probability of failing to reach ∂O within kT time units is at most $(1 - q_T^N)^k$. Put differently, if we define the random variable R_T^N to equal k if ∂O is reached between times $(k - 1)T$ and kT , then the distribution of R_T^N is stochastically dominated by the geometric distribution with parameter q_T^N . It follows that the expected time until ∂O is reached is at most $T \cdot \mathbb{E}R_T^N \leq T/q_T^N = T \exp(N(C_{\partial O} + \varepsilon))$, yielding the upper bound in (21). The upper bound in (20) then follows from Chebyshev's inequality.

To prove the lower bounds in (20) and (21), we again view the process \hat{X}^N as making a series of attempts to reach ∂O . Each attempt requires at least $\delta > 0$ units of clock time, and the LDP upper bound implies that for N large enough, an attempt succeeds with probability less than $\exp(-N(C_{\partial O} - \varepsilon/2))$. Thus to reach ∂O within $k\delta$ time units, one of the first k attempts must succeed, and this has probability less than $k \exp(-N(C_{\partial O} - \varepsilon/2))$. Choosing $k \approx \delta^{-1} \exp(N(C_{\partial O} - \varepsilon))$, we conclude that the probability of exiting O in $\exp(N(C_{\partial O} - \varepsilon))$ time units is less than $k\delta \approx \delta^{-1} \exp(-N\varepsilon/2)$. This quantity vanishes as N grows large, yielding the lower bound in (20); then Chebyshev's inequality gives the lower bound in (21).

4.2. Logit Evolution in Potential Games

We now apply the results above in a context for which the exit costs C_O can be computed explicitly: that of evolution in potential games under the logit choice rule. Consider a sequence of stochastic evolutionary processes $\{\mathbf{X}^N\}_{N=N_0}^\infty$ derived from population games F^N and revision protocols ρ^N that converge uniformly to Lipschitz continuous limits F and ρ (see Section 2.4), where ρ is the logit protocol with noise level $\eta > 0$ (Example 5). Theorem 1 implies that when N is large, the process \hat{X}^N is well approximated over fixed time spans by solutions to the mean dynamic (10), which in the present case is the *logit dynamic*

$$\dot{x} = M^\eta(F(x)) - x, \quad \text{where } M_j^\eta(\pi) = \frac{\exp(\eta^{-1}\pi_j)}{\sum_{k \in \mathcal{A}} \exp(\eta^{-1}\pi_k)}. \quad (23)$$

Now suppose in addition that the limiting population game F is a *potential game* (Sandholm [34]), meaning that there is a function $f: \mathbb{R}_+^n \rightarrow \mathbb{R}$ such that $\nabla f(x) = F(x)$ for all $x \in X$.²⁰ In this case, Hofbauer and Sandholm [24] establish the following global convergence result.

Proposition 3. *If F is a potential game, then the logit potential function*

$$f^\eta(x) = \eta^{-1} f(x) - h(x), \quad h(x) = \sum_{i \in \mathcal{A}} x_i \log x_i \quad (24)$$

is a strict global Lyapunov function for the logit dynamic (23). Thus solutions of (23) from every initial condition converge to connected sets of rest points of (23).

We provide a concise proof of this result in Section 8. Together, Theorem 1 and Proposition 3 imply that for large N , the typical behavior of the process \hat{X}^N is to follow a solution of the logit dynamic (23), ascending the function f^η and approaching a rest point of (23).

We now use the large deviation principle to describe the excursions of the process \hat{X}^N from stable rest points, focusing as before on cases where the mean dynamic (23) has a globally attracting rest point x^* , which by Proposition 3 is the unique local maximizer of f^η on X .

According to the results from the previous section, the time required for the process to exit an open set O containing x^* is characterized by the exit cost (19), which is the infimum over paths from x^* to some $y \notin O$ of the path cost

$$c_{x^*, T}(\phi) = \int_0^T L(\phi_t, \dot{\phi}_t) dt = \int_0^T \sup_{u_t \in \mathbb{R}_+^n} (u_t \dot{\phi}_t - H(\phi_t, u_t)) dt, \quad (25)$$

where the expression of the path cost in terms of the log moment generating functions $H(x, \cdot)$ follows from definition (11) of the Cramér transform. In the logit/potential model, we are able to use properties H to compute the minimal costs (18) and (19) exactly.

Proposition 4. *In the logit/potential model, when (23) has a globally attracting rest point x^* , we have $c_y^* = f^\eta(x^*) - f^\eta(y)$, and so $C_{\partial O}^* = \min_{y \in \partial O} (f^\eta(x^*) - f^\eta(y))$.*

In other words, the minimal cost of a path from x^* to state $y \neq x^*$ is equal to the decrease in the logit potential. Combined with Equations (20)–(22), Proposition 4 implies that the waiting times $\tau_{\partial O}^N$ to escape the set O are described by the smallest decrease in f^η required to reach the boundary of O , and that f^η also governs the rates of decay in the stationary distributions weights $\mu^N(B_\delta(y))$.

Proof. We prove the proposition using tools from the calculus of variations (cf. Freidlin and Wentzell [19, section 5.4]) and these two basic facts about the function H in the logit/potential model, which we prove in Section 8.

Lemma 1. *Suppose that x is in the (relative) interior of X . Then in the logit/potential model,*

$$H(x, -\nabla f^\eta(x)) = 0 \quad \text{and} \tag{26}$$

$$\nabla_u H(x, -\nabla f^\eta(x)) = -(M^\eta(F(x)) - x). \tag{27}$$

Equation (26) is the Hamilton–Jacobi equation associated with path cost minimization problem (18), and it shows that changes in potential provide a lower bound on the cost of reaching any state y from state x^* along an interior path $\phi \in \mathcal{C}_{x^*}[0, T]$ with $\phi_T = y$:

$$c_{x^*, T}(\phi) = \int_0^T \sup_{u_t \in \mathbb{R}_0^n} (u_t' \dot{\phi}_t - H(\phi_t, u_t)) dt \geq \int_0^T -\nabla f^\eta(\phi_t)' \dot{\phi}_t dt = f^\eta(x^*) - f^\eta(y). \tag{28}$$

In Section 8, we prove a generalization of (26) to boundary states, which lets us extend inequality (28) to paths with boundary segments—see Equation (90). These inequalities give us the lower bound

$$c_y^* \geq f^\eta(x^*) - f^\eta(y). \tag{29}$$

Equation (27) is the first-order condition for the first integrand in (28) for paths that are reverse-time solutions to the logit dynamic (23). Thus if $\psi: (-\infty, 0] \rightarrow X$ satisfies $\psi_0 = y$ and $\dot{\psi}_t = -(M^\eta(F(\psi_t)) - \psi_t)$ for all $t \leq 0$, then Proposition 3 and the assumption that x^* is globally attracting imply that $\lim_{t \rightarrow -\infty} \psi_t = x^*$, which with (26) and (27) yields

$$\int_{-\infty}^0 \sup_{u_t \in \mathbb{R}_0^n} (u_t' \dot{\psi}_t - H(\psi_t, u_t)) dt = \int_{-\infty}^0 -\nabla f^\eta(\psi_t)' \dot{\psi}_t dt = f^\eta(x^*) - f^\eta(y). \tag{30}$$

This equation and a continuity argument imply that lower bound (29) is tight. \square


Congestion games are the most prominent example of potential games appearing in applications, and the logit protocol is a standard model of decision making in this context (Ben-Akiva and Lerman [3]). We now illustrate how the results above can be used to describe excursions of the process \mathbf{X}^N from the stationary state of the logit dynamic and the stationary distribution μ^N of the process. We consider a network with three parallel links to simplify the exposition, as our analysis can be conducted just as readily in any congestion game with increasing delay functions.

Example 10. Consider a network consisting of three parallel links with delay functions $l_1(u) = 1 + 8u$, $l_2(u) = 2 + 4u$, and $l_3(u) = 4$. The links are numbered in increasing order of congestion-free travel times (lower-numbered links are shorter in distance) but in decreasing order of congestion-based delays (higher-numbered links have greater capacity). The corresponding continuous-population game has payoff functions $F_i(x) = -l_i(x_i)$ and concave potential function

$$f(x) = - \sum_{i \in \mathcal{A}} \int_0^{x_i} l_i(u) du = -(x_1 + 4(x_1)^2 + 2x_2 + 2(x_2)^2 + 4x_3).$$

The unique Nash equilibrium of this game, $x^* = (\frac{3}{8}, \frac{1}{2}, \frac{1}{8})$, is the state at which travel times on each link are equal ($l_1(x_1^*) = l_2(x_2^*) = l_3(x_3^*) = 4$), and it is the maximizer of f on X .

Figure 1. Solution trajectories of the logit dynamics and level sets of f^η in a congestion game. In both panels, lighter shades represent higher values of f^η , and increments between level sets are 0.5 units. For any point y on a solution trajectory, the most likely excursion path from the rest point to a neighborhood of y follows the trajectory backward from the rest point. The values of f^η also describe the rates of decay of mass in the stationary distribution.



Suppose that a large but finite population of agents repeatedly plays this game, with agents occasionally revising their strategies by applying the logit rule M^η with noise level η . Then in the short term, aggregate behavior evolves according to the logit dynamic (23), ascending the logit potential function $f^\eta = \eta^{-1}f + h$ until closely approaching its global maximizer x^η . Thereafter, (20) and (21) imply that excursions from f^η to other states y occur at rate $\exp(N(f^\eta(x^\eta) - f^\eta(y)))$. The values of N and $f^\eta(y)$ also describe the proportions of time spent at each state: by virtue of (22), $\mu^N(B_\delta(y)) \approx \exp(-N(f^\eta(x^\eta) - f^\eta(y)))$.

Figure 1 presents solutions of the logit dynamic (23) and level sets of the logit potential function f^η in the congestion game above for noise levels $\eta = 0.25$ (panel (i)) and $\eta = 0.1$ (panel (ii)). In both cases, all solutions of (23) ascend the logit potential function and converge to its unique maximizer, $x^{(0.25)} \approx (0.3563, 0.4482, 0.1956)$ in (i) and $x^{(0.1)} = (0.3648, 0.4732, 0.1620)$ in (ii). The latter rest point is closer to the Nash equilibrium on account of the smaller amount of noise in agents’ decisions.

In each panel, the “major axes” of the level sets of f^η correspond to exchanges of agents playing strategy 3 for agents playing strategies 2 and 1 in fixed shares, with a slightly larger share for strategy 2. That deviations of this sort are the most likely is explained by the lower sensitivity of delays on higher numbered links to fluctuations in usage. In both panels, the increments between the displayed level sets of f^η are 0.5 units. Many more level sets are drawn in panel (ii) than in panel (i):²¹ when there is less noise in agents’ decisions, excursions from equilibrium of a given unlikelihood are generally smaller, and excursions of a given size and direction are less common. □

4.3. Discussion

The analyses above rely on the assumption that the mean dynamic (10) admits a globally attracting state. If, instead, this dynamic has multiple attractors, then the time $\hat{\tau}_{\partial O}^N$ to exit O starting from a stable rest point $x^* \in O$ need only satisfy properties (20) and (21) when the set O is contained in the basin of attraction of x^* . Beyond this case, the most likely amount of time required to escape O may disagree with the expected amount of time to do so, since the latter may be driven by a small probability of becoming stuck near another attractor in O . Likewise, when the global structure of (10) is nontrivial, the asymptotics of the stationary distribution are more complicated, being driven by the relative likelihoods of transitions between the different attractors. To study these questions in our context, one must not only address the complications noted in Section 4.1 but also employ the graph-theoretic arguments developed by Freidlin and Wentzell [19, chap. 6] to capture the structure of transitions among the attractors. Because the limiting stationary distribution is the basis for the approach to equilibrium selection discussed in the introduction, carrying out this analysis is an important task for future work.

We have shown that the control problems appearing in the statement of the large deviation principle can be solved explicitly in the case of logit choice in potential games. They can also be solved in the context of

two-action games, in which the state space X is one dimensional. Beyond these two cases, the control problems do not appear to admit analytical solutions.

To contend with this, and to facilitate comparisons with other analyses in the literature, one can consider the *large population double limit*, studying the behavior of the large population limit as the noise level in agents' decisions is taken to zero. There are strong reasons to expect this double limit to be analytically tractable. In Sandholm and Staudigl [39], we study the reverse order of limits, under which the noise level η is first taken to zero and then the population size N to infinity. For this order of limits, we show that large deviations properties are determined by the solutions to piecewise linear control problems and that these problems can be solved analytically. Moreover, Sandholm [36] uses birth-death chain methods to show that in the two-action case, large deviations properties under the two orders of limits are identical. These results and our preliminary analyses suggest that the large population double limit is tractable and that in typical cases, conclusions for the two orders of limits will agree. Although we are a number of steps away from reaching these ends, the analysis here provides the tools required for work on this program to proceed.

5. Analysis

The proof of Theorem 2 follows the weak convergence approach of Dupuis and Ellis [17] (henceforth referred to as DE). As noted in the introduction, the main novelty we must contend with is the fact that our processes run on a compact set X . This necessitates a delicate analysis of the behavior of the process on and near the boundary of X . At the same time, the fact that the conditional laws (4) have finite support considerably simplifies a number of the steps from DE's approach. Proofs of auxiliary results that would otherwise interrupt the flow of the argument are relegated to Sections 6 and 7. At some points where the arguments mirror those from DE, we present them somewhat informally; the omitted details are provided in the working version of this paper, Sandholm and Staudigl [40] (henceforth referred to as SS).

Before entering into the details of our analysis, we provide an overview. In Section 5.1, we use the representation (14) of the Cramér transform to establish joint continuity properties of the running cost function $L(\cdot, \cdot)$. To start, we provide examples of discontinuities that this function exhibits at the boundary of X . We then show that excepting these discontinuities, the running cost function is "as continuous as possible" (Proposition 5).

The remaining sections follow the line of argument in DE, with modifications that use Proposition 5 to contend with boundary issues. Section 5.2 describes how the large deviation principle upper and lower bounds can be deduced from corresponding Laplace principle upper and lower bounds. The latter bounds concern the limits of expectations of continuous functions, making them amenable to analysis using weak convergence arguments. Section 5.3 explains how the expectations appearing in the Laplace principle can be expressed as solutions to stochastic optimal control problems (47), the running costs of which are relative entropies defined with respect to the transition laws $\nu^N(\cdot | x)$ of \mathbf{X}^N . Section 5.4 describes the limit properties of the controlled processes as N grows large. Finally, Sections 5.5 and 5.6 use the foregoing results to prove the Laplace principle upper and lower bounds; here, the main novelty is in Section 5.6, where we show that the control problem appearing on the right-hand side of the Laplace principle admits ε -optimal solutions that initially obey the mean dynamic and remain in the interior of the simplex thereafter (Proposition 9).

5.1. Joint Continuity of Running Costs

Representation (14) implies that for each $x \in X$, the Cramér transform $L(x, \cdot)$ is continuous on its domain $Z(x)$ (see the beginning of the proof of Proposition 5). The remainder of this section uses this representation to establish joint continuity properties of the running cost function $L(\cdot, \cdot)$.

The difficulty lies in establishing these properties at states on the boundary of X . Fix $x \in X$, and let $i \in \text{supp}(x)$ and $j \neq i$. Since $e_j - e_i$ is an extreme point of $Z(x)$, the point mass $\delta_{e_j - e_i}$ is the only distribution in $\Delta(\mathcal{Z})$ with mean $e_j - e_i$. Thus representation (14) implies that

$$L(x, e_j - e_i) = R(\delta_{e_j - e_i} \| \nu(\cdot | x)) = -\log x_i \sigma_{ij}(x) \geq -\log x_i. \quad (31)$$

Thus $L(x, e_j - e_i)$ grows without bound as x approaches the face of X on which $x_i = 0$, and $L(x, e_j - e_i) = \infty$ when $x_i = 0$. Intuitively, reducing the number of action i players reduces the probability that such a player is selected to revise; when there are no such players, selecting one becomes impossible.

A more serious difficulty is that running costs are not continuous at the boundary of X even when they are finite. For example, suppose that $n \geq 3$, let x be in the interior of X , and let $z_\alpha = e_3 - (\alpha e_1 + (1 - \alpha)e_2)$. Since the unique λ with mean z_α has $\lambda(e_3 - e_1) = \alpha$ and $\lambda(e_3 - e_2) = 1 - \alpha$, Equation (14) implies that

$$L(x, z_\alpha) = \alpha \log \frac{\alpha}{x_1 \sigma_{13}(x)} + (1 - \alpha) \log \frac{1 - \alpha}{x_2 \sigma_{23}(x)}.$$

If we set $\alpha(x) = -(\log(x_1\sigma_{13}(x)))^{-1}$ and let x approach some x^* with $x_1^* = 0$, then $L(x, z_{\alpha(x)})$ approaches $1 - \log(x_2\sigma_{23}(x^*))$; however, $z_{\alpha(x)}$ approaches $e_3 - e_2$, and $L(x^*, e_3 - e_2) = -\log(x_2\sigma_{23}(x^*))$.

These observations leave open the possibility that the running cost function is continuous as a face of X is approached, provided that one restricts attention to displacement directions $z \in Z = \text{conv}(\mathcal{Z}) = \text{conv}(\{e_j - e_i; i, j \in \mathcal{A}\})$ that remain feasible on that face. Proposition 5 shows that this is indeed the case.

For any nonempty $I \subseteq \mathcal{A}$, define $X(I) = \{x \in X: I \subseteq \text{supp}(x)\}$, $\mathcal{Z}(I) = \{e_j - e_i; j \in \mathcal{A}, i \in I\} = \{z \in \mathcal{Z}: z_j \geq 0 \text{ for all } j \notin I\}$, and $Z(I) = \text{conv}(\mathcal{Z}(I)) = \{z \in Z: z_j \geq 0 \text{ for all } j \notin I\}$ (cf. Endnote 13).

Proposition 5. (i) *The function $L(\cdot, \cdot)$ is continuous on $\text{int}(X) \times Z$.*

(ii) *For any nonempty $I \subseteq \mathcal{A}$, $L(\cdot, \cdot)$ is continuous on $X(I) \times Z(I)$.*

Proof. For any $\lambda \in \Delta(\mathcal{Z}(I))$ and $x \in X(I)$, we have $\text{supp}(\lambda) \subseteq \mathcal{Z}(I) \subseteq \text{supp}(\nu(\cdot | x))$. Thus by the definition of relative entropy, the function $\mathcal{L}: X(I) \times \Delta(\mathcal{Z}(I)) \rightarrow [0, \infty]$ defined by $\mathcal{L}(x, \lambda) = R(\lambda \| \nu(\cdot | x))$ is real valued and continuous.

Let

$$\Lambda_{Z(I)}(z) = \left\{ \lambda \in \Delta(\mathcal{Z}): \text{supp}(\lambda) \subseteq \mathcal{Z}(I), \sum_{z \in \mathcal{Z}} z \lambda(z) = z \right\} \quad (32)$$

be the set of distributions on \mathcal{Z} with support contained in $\mathcal{Z}(I)$ and with mean z . Then the correspondence $\Lambda_{Z(I)}: Z(I) \Rightarrow \Delta(\mathcal{Z}(I))$ defined by (32) is clearly continuous and compact valued. Thus if we define $L_I: X(I) \times Z(I) \rightarrow [0, \infty)$ by

$$L_I(x, z) = \min\{R(\lambda \| \nu(\cdot | x)): \lambda \in \Lambda_{Z(I)}(z)\}, \quad (33)$$

then the theorem of the maximum (Berge [7]) implies that L_I is continuous.

By representation (14),

$$L(x, z) = \min\{R(\lambda \| \nu(\cdot | x)): \lambda \in \Lambda_{\mathcal{Z}}(z)\}. \quad (34)$$

Since $\mathcal{Z}(\mathcal{A}) = \mathcal{Z}$, (33) and (34) imply that $L_S(x, z) = L(x, z)$, establishing part (i).

To begin the proof of part (ii), we eliminate redundant cases using an inductive argument on the cardinality of I . Part (i) establishes the base case in which $\#I = n$. Suppose that the claim in part (ii) is true when $\#I > k \in \{1, \dots, n-1\}$; we must show that this claim is true when $\#I = k$.

Suppose that $\text{supp}(x) = J \supset I$, so that $\#J > k$. Then the inductive hypothesis implies that the restriction of L to $X(J) \times Z(J)$ is continuous at (x, z) . Since $X(J) \subset X(I)$ is open relative to X and since $Z(J) \supset Z(I)$, the restriction of L to $X(I) \times Z(I)$ is also continuous at (x, z) .

It remains to show that the restriction of L to $X(I) \times Z(I)$ is continuous at all $(x, z) \in X(I) \times Z(I)$ with $\text{supp}(x) = I$. Since $\mathcal{Z}(I) \subset \mathcal{Z}$, (33) and (34) imply that for all $(x, z) \in X(I) \times Z(I)$,

$$L(x, z) \leq L_I(x, z). \quad (35)$$

If, in addition, $\text{supp}(x) = I$, then $\mathcal{L}(x, \lambda) = \infty$ whenever $\text{supp}(\lambda) \not\subseteq \mathcal{Z}(I)$, so (33) and (34) imply that inequality (35) binds. Since L_I is continuous, our remaining claim follows directly from this uniform approximation.

Lemma 2. *For any $\varepsilon > 0$, there exists a $\delta > 0$ such that for any $x \in X$ with $\max_{k \in \mathcal{A} \setminus I} x_k < \delta$ and any $z \in Z(I)$, we have*

$$L(x, z) \geq L_I(x, z) - \varepsilon. \quad (36)$$

A constructive proof of Lemma 2 is provided in Section 6.

5.2. The Large Deviation Principle and the Laplace Principle

While Theorem 2 is stated for the finite time interval $[0, T]$, we assume without loss of generality that $T = 1$. In what follows, \mathcal{C} denotes the set of continuous functions from $[0, 1]$ to X endowed with the supremum norm, $\mathcal{C}_x \subset \mathcal{C}$ denotes the set of paths in \mathcal{C} starting at x , and $\mathcal{AC} \subset \mathcal{C}$ and $\mathcal{AC}_x \subset \mathcal{C}_x$ are the subsets consisting of absolutely continuous paths.

Following DE, we deduce Theorem 2 from the Laplace principle. The processes $\{\hat{\mathbf{X}}^N\}_{N=N_0}^\infty$ are said to satisfy the *Laplace principle* if for any bounded continuous function $h: \mathcal{C} \rightarrow \mathbb{R}$ we have

$$\lim_{N \rightarrow \infty} \frac{1}{N} \log \mathbb{E}_{x^N} [\exp(-Nh(\hat{\mathbf{X}}^N))] = - \inf_{\phi \in \mathcal{C}} (c_x(\phi) + h(\phi)). \quad (37)$$

Equation (37) can be decomposed into two inequalities:

$$\limsup_{N \rightarrow \infty} \frac{1}{N} \log \mathbb{E}_{x^N} [\exp(-Nh(\hat{\mathbf{X}}^N))] \leq - \inf_{\phi \in \mathcal{C}} (c_x(\phi) + h(\phi)), \quad \text{and} \quad (38a)$$

$$\liminf_{N \rightarrow \infty} \frac{1}{N} \log \mathbb{E}_{x^N} [\exp(-Nh(\hat{\mathbf{X}}^N))] \geq - \inf_{\phi \in \mathcal{C}} (c_x(\phi) + h(\phi)). \quad (38b)$$

Inequality (38a) is called the *Laplace principle upper bound*, and inequality (38b) is called the *Laplace principle lower bound*.

Because c_x is a rate function (Proposition 1), the large deviation principle (Theorem 2) and the Laplace principle (Theorem 4) each imply the other.

Theorem 4. *Suppose that the processes $\{\hat{\mathbf{X}}^N\}_{N=N_0}^\infty$ have initial conditions $x^N \in \mathcal{X}^N$ satisfying $\lim_{N \rightarrow \infty} x^N = x \in X$. Then the large deviation principle (16) holds if and only if the Laplace principle (37) holds.*

The forward implication of Theorem 4 is known as *Varadhan’s integral lemma* (DE, theorem 1.2.1). For intuition, express the large deviation principle loosely as $\mathbb{P}_x(\mathbf{X}^N \approx \phi) \approx \exp(-Nc_x(\phi))$ for $\phi \in \mathcal{C}$ and argue heuristically that

$$\mathbb{E}_x[\exp(-Nh(\hat{\mathbf{X}}^N))] \approx \int_{\Phi} \exp(-Nh(\phi)) \mathbb{P}_x(\mathbf{X}^N \approx \phi) d\phi \approx \int_{\Phi} \exp(-N(h(\phi) + c_x(\phi))) d\phi \approx \exp\left(-N \inf_{\phi \in \mathcal{C}} (c_x(\phi) + h(\phi))\right),$$

where the final approximation uses the Laplace method for integrals (de Bruijn [14]).

Our analysis requires the reverse implication of Theorem 4 (DE, theorem 1.2.3). For a heuristic derivation, let $\Phi \subset \mathcal{C}$, and let the (extended real-valued, discontinuous) function h_Φ be the indicator function of $\Phi \subset \mathcal{C}$ in the sense of convex analysis:

$$h_\Phi(\phi) = \begin{cases} 0 & \text{if } \phi \in \Phi, \\ +\infty & \text{otherwise.} \end{cases}$$

Then plugging h_Φ into (37) yields

$$\lim_{N \rightarrow \infty} \frac{1}{N} \log \mathbb{P}_{x^N}(\hat{\mathbf{X}}^N \in \Phi) = - \inf_{\phi \in \Phi} c_x(\phi),$$

which is the large deviation principle. The proof of DE, theorem 1.2.3, proceeds by considering well-chosen approximations of h_Φ by bounded continuous functions.

The left-hand sides of the inequalities (38a) and (38b) that together form the Laplace principle concern limits of expectations of continuous functions, raising the possibility that the inequalities can be evaluated by means of weak convergence arguments. We return to this point at the end of the next section.

5.3. The Stochastic Optimal Control Problem

For a given function $h: \mathcal{C} \rightarrow \mathbb{R}$, we define

$$V^N(x^N) = -\frac{1}{N} \log \mathbb{E}_{x^N}[\exp(-Nh(\hat{\mathbf{X}}^N))] \tag{39}$$

to be the negation of the expression from the left-hand side of the Laplace principle. This section, which follows DE, sections 3.2 and 4.3, shows how V^N can be expressed as the solution of a stochastic optimal control problem. The running costs of this problem are relative entropies, and its terminal costs are determined by the function h .

For each $k \in \{0, 1, 2, \dots, N\}$ and sequence $(x_0, \dots, x_k) \in (\mathcal{X}^N)^{k+1}$, we define the period k value function

$$V_k^N(x_0, \dots, x_k) = -\frac{1}{N} \log \mathbb{E}[\exp(-Nh(\hat{\mathbf{X}}^N)) \mid X_0^N = x_0, \dots, X_k^N = x_k]. \tag{40}$$

Note that $V_0^N \equiv V^N$. If we define the map $\hat{\phi} (= \hat{\phi}^N)$ from sequences x_0, \dots, x_N to paths in \mathcal{C} by

$$\hat{\phi}_t(x_0, \dots, x_N) = x_k + (Nt - k)(x_{k+1} - x_k) \quad \text{for all } t \in [k/N, (k+1)/N], \tag{41}$$

then (40) implies that

$$V_N^N(x_0, \dots, x_N) = h(\hat{\phi}(x_0, \dots, x_N)). \tag{42}$$

Note also that $\hat{X}_t^N(\omega) = \hat{\phi}_t(X_0^N(\omega), \dots, X_N^N(\omega))$; this can be expressed concisely as $\hat{\mathbf{X}}^N = \hat{\phi}(\mathbf{X}^N)$.

Proposition 6 shows that the value functions V_k^N satisfy a dynamic programming functional equation, with running costs given by relative entropy functions and with terminal costs given by $h(\hat{\phi}(\cdot))$. To read Equation (43), recall that ν^N is the transition kernel for the Markov chain $\{X_k^N\}$.

Proposition 6. *For $k \in \{0, 1, \dots, N - 1\}$ and $(x_0, \dots, x_k) \in (\mathcal{X}^N)^{k+1}$, we have*

$$V_k^N(x_0, \dots, x_k) = \inf_{\lambda \in \Delta(\mathcal{Z})} \left(\frac{1}{N} R(\lambda \parallel \nu^N(\cdot \mid x_k)) + \sum_{z \in \mathcal{Z}} V_{k+1}^N\left(x_0, \dots, x_k + \frac{1}{N}z\right) \lambda(z) \right). \tag{43}$$

For $k = N$, V_N^N is given by the terminal condition (42).

The key idea behind Proposition 6 is the following observation (DE, proposition 1.4.2), which provides a variational formula for expressions such as (39) and (40).

Observation 1. For any probability measure $\pi \in \Delta(\mathcal{Z})$ and function $\gamma: \mathcal{Z} \rightarrow \mathbb{R}$, we have

$$-\log \sum_{z \in \mathcal{Z}} e^{-\gamma(z)} \pi(z) = \min_{\lambda \in \Delta(\mathcal{Z})} \left(R(\lambda \| \pi) + \sum_{z \in \mathcal{Z}} \gamma(z) \lambda(z) \right). \quad (44)$$

The minimum is attained at $\lambda^*(z) = \pi(z) e^{-\gamma(z)} / \sum_{y \in \mathcal{Z}} \pi(y) e^{-\gamma(y)}$. In particular, $\lambda^* \ll \pi$.

Above, the notation $\lambda^* \ll \pi$ means that the measure λ is absolutely continuous with respect to the measure π .

Equation (44) expresses the log expectation on its left-hand side as the minimized sum of two terms: a relative entropy term that only depends on the probability measure π and an expectation that only depends on the function γ . This additive separability and the Markov property lead to Equation (43). Specifically, observe that

$$\begin{aligned} \exp(-NV_k^N(x_0, \dots, x_k)) &= \mathbb{E}[\exp(-Nh(\hat{\phi}(\mathbf{X}^N))) \mid X_0^N = x_0, \dots, X_k^N = x_k] \\ &= \mathbb{E}[\mathbb{E}[\exp(-Nh(\hat{\phi}(\mathbf{X}^N))) \mid X_0^N, \dots, X_{k+1}^N] \mid X_0^N = x_0, \dots, X_k^N = x_k] \\ &= \mathbb{E}[-NV_{k+1}^N(X_0^N, \dots, X_{k+1}^N) \mid X_0^N = x_0, \dots, X_k^N = x_k] \\ &= \sum_{z \in \mathcal{Z}} \exp\left(-NV_{k+1}^N\left(x_0, \dots, x_k, x_k + \frac{1}{N}z\right)\right) v^N(z \mid x_k), \end{aligned}$$

where the last line uses the Markov property. This equality and Observation 1 yield

$$\begin{aligned} V_k^N(x_0, \dots, x_k) &= -\frac{1}{N} \log \sum_{z \in \mathcal{Z}} \exp\left(-NV_{k+1}^N\left(x_0, \dots, x_k, x_k + \frac{1}{N}z\right)\right) v^N(z \mid x_k) \\ &= \frac{1}{N} \inf_{\lambda \in \Delta(\mathcal{Z})} \left(R(\lambda \| v^N(\cdot \mid x_k)) + \sum_{z \in \mathcal{Z}} NV_{k+1}^N\left(x_0, \dots, x_k + \frac{1}{N}z\right) \lambda(z) \right), \end{aligned}$$

which is Equation (43).

Since the value functions V_k^N satisfy the dynamic programming functional Equation (43), they also can be represented by describing the same dynamic program in sequence form. To do so, we define for $k \in \{0, \dots, N-1\}$ a *period k control* $\lambda_k^N: (\mathcal{X}^N)^k \rightarrow \Delta(\mathcal{Z})$, which for each sequence of states (x_0, \dots, x_k) specifies a probability distribution $\lambda_k^N(\cdot \mid x_0, \dots, x_k)$ —namely, the minimizer of problem (43). Given the sequence of controls $\{\lambda_k^N\}_{k=0}^{N-1}$ and an initial condition $\xi_0^N = x^N \in \mathcal{X}^N$, we define the *controlled process* $\xi^N = \{\xi_k^N\}_{k=0}^N$ by $\xi_0^N = x^N \in \mathcal{X}^N$ and the recursive formula

$$\xi_{k+1}^N = \xi_k^N + \frac{1}{N} \zeta_k^N, \quad (45)$$

where ζ_k^N is a \mathcal{Z} -valued random variable with law $\lambda_k^N(\cdot \mid \xi_0^N, \dots, \xi_k^N)$. We also define the piecewise affine interpolation $\tilde{\xi}^N = \{\tilde{\xi}_t^N\}_{t \in [0,1]}$ by $\tilde{\xi}_t^N = \hat{\phi}_t(\xi^N)$, where $\hat{\phi}$ is the interpolation function (41). We then have the following.

Proposition 7. For $k \in \{0, 1, \dots, N-1\}$ and $(x_0, \dots, x_k) \in (\mathcal{X}^N)^{k+1}$, $V_k^N(x_0, \dots, x_k)$ equals

$$\inf_{\lambda_k^N, \dots, \lambda_{N-1}^N} \mathbb{E} \left[\frac{1}{N} \sum_{j=k}^{N-1} R(\lambda_j^N(\cdot \mid \xi_0^N, \dots, \xi_j^N) \| v^N(\cdot \mid \xi_j^N)) + h(\tilde{\xi}^N) \mid \xi_0^N = x_0, \dots, \xi_k^N = x_k \right]. \quad (46)$$

Since Observation 1 implies that the infimum in the functional Equation (43) is always attained, Proposition 7 follows from standard results (cf. DE, theorem 1.5.2), and moreover, the infimum in (46) is always attained.

Since $V_0^N \equiv V^N$ by construction, Proposition 7 yields the representation

$$V^N(x^N) = \inf_{\lambda_0^N, \dots, \lambda_{N-1}^N} \mathbb{E}_{x^N} \left[\frac{1}{N} \sum_{j=0}^{N-1} R(\lambda_j^N(\cdot \mid \xi_0^N, \dots, \xi_j^N) \| v^N(\cdot \mid \xi_j^N)) + h(\tilde{\xi}^N) \right]. \quad (47)$$

The running costs in (47) are relative entropies of control distributions with respect to transition distributions of the Markov chain \mathbf{X}^N and so reflect how different the control distribution is from the law of the Markov

chain at the relevant state. Note as well that the terminal payoff $h(\hat{\xi}^N)$ may depend on the entire path of the controlled process ξ^N .

With this groundwork in place, we can describe our application of Dupuis and Ellis’s weak convergence approach to large deviations. Equation (47) represents expression $V^N(x^N)$ from the Laplace principle as the optimal solution to a stochastic control problem. Section 5.4 uses compactness arguments to show that as N grows large, interpolated versions of the stochastic controls $\{\lambda_k^N\}$ and controlled processes ξ^N converge in distribution, where with probability 1, the limiting controlled process depends only on the limiting stochastic control by way of its means (see Proposition 8, especially (49)). By applying this result to the optimal controls and using Fatou’s lemma, Section 5.5 establishes the Laplace principle upper bound (50). In retrospect, it is this calculation that determines the form of the rate function $c_x(\cdot)$. Section 5.6 derives the Laplace principle lower bound (52) by first constructing a path ψ that approximately minimizes $c_x(\cdot) + h(\cdot)$, then introducing delicately constructed controlled processes that mirror ψ , and at last applying Proposition 8 and the dominated convergence theorem to establish the bound. A key argument in this proof, Proposition 9, shows that the path ψ can be chosen to have properties that justify our appeal to dominated convergence despite the discontinuity of the running cost function $L(\cdot, \cdot)$ at the boundary of the state space.

5.4. Convergence of the Controlled Processes

The increments of the controlled process ξ^N are determined in two steps: first, the history of the process determines the measure $\lambda_k^N(\cdot | \xi_0^N, \dots, \xi_k^N) \in \Delta(\mathcal{Z})$, and then the increment itself is determined by a draw from this measure. With some abuse of notation, one can write $\lambda_k^N(\cdot) = \lambda_k^N(\cdot | \xi_0^N, \dots, \xi_k^N)$ and thus view λ_k^N as a random measure. Then, using compactness arguments, one can show that as N grows large, certain subsequences of the random measures λ_k^N on $\Delta(\mathcal{Z})$ converge in a suitable sense to limiting random measures. Because the increments of ξ^N become small as N grows large (cf. (45)), intuition from the law of large numbers—specifically, Theorem 1—suggests that the idiosyncratic part of the randomness in these increments should be averaged away. Thus in the limit, the evolution of the controlled process should still depend on the realizations of the random measures, but it should only do so by way of their means.

The increments of the controlled process ξ^N are determined in two steps: first, the history of the process determines the measure $\lambda_k^N(\cdot | \xi_0^N, \dots, \xi_k^N) \in \Delta(\mathcal{Z})$, and then the increment itself is determined by a draw from this measure. With some abuse of notation, one can write $\lambda_k^N(\cdot) = \lambda_k^N(\cdot | \xi_0^N, \dots, \xi_k^N)$ and thus view λ_k^N as a random measure. Then, using compactness arguments, one can show that as N grows large, certain subsequences of the random measures λ_k^N on $\Delta(\mathcal{Z})$ converge in a suitable sense to limiting random measures. Because the increments of ξ^N become small as N grows large (cf. (45)), intuition from the law of large numbers—specifically Theorem 1—suggests that the idiosyncratic part of the randomness in these increments should be averaged away. Thus in the limit, the evolution of the controlled process should still depend on the realizations of the random measures, but it should only do so by way of their means.

To make this argument precise, we introduce continuous-time interpolations of the controlled processes $\xi^N = \{\xi_k^N\}_{k=0}^N$ and the sequence of controls $\{\lambda_k^N\}_{k=0}^{N-1}$. The piecewise affine interpolation $\bar{\xi}^N = \{\bar{\xi}_t^N\}_{t \in [0,1]}$ was introduced above; it takes values in the space $\mathcal{C} = \mathcal{C}([0,1]: X)$. The piecewise constant interpolation $\bar{\xi}^N = \{\bar{\xi}_t^N\}_{t \in [0,1]}$ is defined by

$$\bar{\xi}_t^N = \begin{cases} \xi_k^N, & \text{if } t \in [k/N, (k+1)/N) \text{ and } k = 0, 1, \dots, N-2, \\ \xi_{N-1}^N, & \text{if } t \in [(N-1)/N, 1]. \end{cases}$$

The process $\bar{\xi}^N$ takes values in the space $\mathcal{D} = \mathcal{D}([0,1]: X)$ of left-continuous functions with right limits, which we endow with the Skorokhod topology. Finally, the piecewise constant control process $\{\bar{\lambda}_t^N\}_{t \in [0,1]}$ is defined by

$$\bar{\lambda}_t^N(\cdot) = \begin{cases} \lambda_k^N(\cdot | \xi_0^N, \dots, \xi_k^N), & \text{if } t \in [k/N, (k+1)/N) \text{ and } k = 0, 1, \dots, N-2, \\ \lambda_{N-1}^N(\cdot | \xi_0^N, \dots, \xi_{N-1}^N), & \text{if } t \in [(N-1)/N, 1]. \end{cases}$$

Using these definitions, we can rewrite formulation (47) of $V^N(x^N)$ as

$$V^N(x^N) = \inf_{\lambda_0^N, \dots, \lambda_{N-1}^N} \mathbb{E}_{x^N} \left(\int_0^1 R(\bar{\lambda}_t^N \| \nu^N(\cdot | \bar{\xi}_t^N)) dt + h(\hat{\xi}^N) \right). \tag{48}$$

As noted after Proposition 7, the infimum in (48) is always attained by some choice of the control sequence $\{\lambda_k^N\}_{k=0}^{N-1}$.

Let $\Delta(\mathcal{Z} \times [0, 1])$ denote the space of probability measures on $\mathcal{Z} \times [0, 1]$. For a collection $\{\pi_t\}_{t \in [0, 1]}$ of measures $\pi_t \in \Delta(\mathcal{Z})$ that is Lebesgue measurable in t , we define the measure $\pi_t \otimes dt \in \Delta(\mathcal{Z} \times [0, 1])$ by

$$(\pi_t \otimes dt)(\{z\} \times B) = \int_B \pi_t(z) dt$$

for all $z \in \mathcal{Z}$ and all Borel sets B of $[0, 1]$. Using this definition, we can represent the piecewise constant control process $\{\bar{\lambda}_t^N\}_{t \in [0, 1]}$ as the *control measure* $\Lambda^N = \bar{\lambda}_t^N \otimes dt$. Evidently, Λ^N is a random measure taking values in $\Delta(\mathcal{Z} \times [0, 1])$, a space we endow with the topology of weak convergence.

Proposition 8, a direct consequence of DE, theorem 5.3.5 and p. 165, formalizes the intuition expressed in the first paragraph of this section. It shows that along certain subsequences, the control measures Λ^N and the interpolated controlled processes $\bar{\xi}^N$ and $\bar{\xi}^N$ converge in distribution to a random measure $\Lambda = \lambda_t \otimes dt$ and a random process ξ , and moreover, that the evolution of ξ is almost surely determined by the means of the measures λ_t .

Proposition 8. *Suppose that the initial conditions $x^N \in \mathcal{X}^N$ converge to $x \in X$ and that the control sequence $\{\lambda_k^N\}_{k=0}^{N-1}$ is such that $\sup_{N \geq N_0} V^N(x^N) < \infty$.*

(i) *Given any subsequence of $\{(\Lambda^N, \bar{\xi}^N, \bar{\xi}^N)\}_{N=N_0}^\infty$, there exists a $\Delta(\mathcal{Z} \times [0, 1])$ -valued random measure Λ and \mathcal{C} -valued random process ξ (both possibly defined on a new probability space) such that some subsubsequence converges in distribution to (Λ, ξ, ξ) in the topologies specified above.*

(ii) *There is a collection of $\Delta(\mathcal{Z})$ -valued random measures $\{\lambda_t\}_{t \in [0, 1]}$, measurable with respect to t , such that with probability one, the random measure Λ can be decomposed as $\Lambda = \lambda_t \otimes dt$.*

(iii) *With probability 1, the process ξ satisfies $\xi_t = x + \int_0^t (\sum_{z \in \mathcal{Z}} z \lambda_s(z)) ds$ for all $t \in [0, 1]$, and is absolutely continuous in t . Thus with probability 1,*

$$\dot{\xi}_t = \sum_{z \in \mathcal{Z}} z \lambda_t(z) \tag{49}$$

almost surely with respect to Lebesgue measure.

5.5. Proof of the Laplace Principle Upper Bound

In this section, we consider the Laplace principle upper bound (38a), which definition (39) allows us to express as

$$\liminf_{N \rightarrow \infty} V^N(x^N) \geq \inf_{\phi \in \mathcal{C}} (c_x(\phi) + h(\phi)). \tag{50}$$

The argument here follows DE, section 6.2. Let $\{\lambda_k^N\}_{k=0}^{N-1}$ be the optimal control sequence in representation (47), and let $\bar{\xi}^N$ be the corresponding controlled process. Define the triples $\{(\Lambda^N, \bar{\xi}^N, \bar{\xi}^N)\}_{N=N_0}^\infty$ of interpolated processes as in Section 5.4. Proposition 8 shows that for any subsequence of these triples, there is a subsubsequence that converges in distribution to some triple $(\lambda_t \otimes dt, \xi, \xi)$ satisfying (49). Then one argues that along this subsubsequence,

$$\begin{aligned} \liminf_{N \rightarrow \infty} V^N(x^N) &\geq \mathbb{E}_x \left(\int_0^1 R(\lambda_t \| \nu(\cdot | \xi_t)) dt + h(\xi) \right) \geq \mathbb{E}_x \left(\int_0^1 L \left(\xi_t, \sum_{z \in \mathcal{Z}} z \lambda_t(z) \right) dt + h(\xi) \right) \\ &= \mathbb{E}_x \left(\int_0^1 L(\xi_t, \dot{\xi}_t) dt + h(\xi) \right) \geq \inf_{\phi \in \mathcal{C}} (c_x(\phi) + h(\phi)). \end{aligned}$$

The key ingredients needed to establish the initial inequality are Equation (47), Skorokhod's theorem, Equation (51), the lower semicontinuity of relative entropy, and Fatou's lemma. Then the second inequality follows from representation (14) of the Cramér transform, the equality from Equation (49), and the final inequality from the definition (15) of the cost function c_x . Since the subsequence chosen initially was arbitrary, inequality (50) is proved.

The details of this argument can be found in section 9.1 of SS, which largely follows DE, section 6.2. But while in DE the transition kernels $\nu^N(\cdot | x)$ of the Markov chains are assumed to be independent of N , here we allow for a vanishing dependence on N (cf. Equation (8)). Thus we require a simple additional argument, Lemma 9.1, that uses lower bound (6) to establish the uniform convergence of relative entropies: namely, that if $\lambda^N: \mathcal{X}^N \rightarrow \Delta(\mathcal{Z})$ are transition kernels satisfying $\lambda^N(\cdot | x) \ll \nu^N(\cdot | x)$ for all $x \in \mathcal{X}^N$, then

$$\lim_{N \rightarrow \infty} \max_{x \in \mathcal{X}^N} |R(\lambda^N(\cdot | x) \| \nu^N(\cdot | x)) - R(\lambda^N(\cdot | x) \| \nu(\cdot | x))| = 0. \tag{51}$$

5.6. Proof of the Laplace Principle Lower Bound

Finally, we consider the Laplace principle lower bound (38b), which definition (39) lets us express as

$$\limsup_{N \rightarrow \infty} V^N(x^N) \leq \inf_{\phi \in \mathcal{C}} (c_x(\phi) + h(\phi)). \tag{52}$$

The argument here largely follows DE, sections 6.2 and 6.4. Their argument begins by choosing a path that is ε -optimal in the minimization problem from the right-hand side of (52). To account for our processes running on a set with a boundary, we show that this path can be chosen to start with a brief segment that follows the mean dynamic and then stays in the interior of X thereafter (Proposition 9). With this choice of path, the joint continuity properties of the running costs $L(\cdot, \cdot)$ established in Proposition 5 are sufficient to complete the dominated convergence argument in display (61), which establishes that inequality (52) is violated by no more than ε . Since ε was arbitrary, (52) follows.

For a path $\phi \in \mathcal{C} = \mathcal{C}([0, 1]; X)$ and an interval $I \subseteq [0, 1]$, write ϕ_I for $\{\phi_t; t \in I\}$. Define the set of paths

$$\tilde{\mathcal{C}} = \{\phi \in \mathcal{C}: \text{for some } \alpha \in (0, 1], \phi_{[0, \alpha]} \text{ solves (10) and } \phi_{[\alpha, 1]} \subset \text{int}(X)\}.$$

Let $\tilde{\mathcal{C}}_x$ denote the set of such paths that start at x .

Proposition 9. For all $x \in X$, $\inf_{\phi \in \mathcal{C}} (c_x(\phi) + h(\phi)) = \inf_{\phi \in \tilde{\mathcal{C}}_x} (c_x(\phi) + h(\phi))$.

The proof of Proposition 9 is deferred until Section 7, but it can be summarized as follows. Given a path $\phi \in \mathcal{C}_x$ and a small enough $\alpha > 0$, we construct a new path $\tilde{\phi}^\alpha \in \mathcal{C}_x$ that follows the solution to the mean dynamic (10) from initial condition x until time α and whose increments starting at time α mimic those of ϕ starting at time 0 but are scaled down by a factor that approaches 1 as α approaches 0. By construction, $\tilde{\phi}^\alpha$ converges uniformly to ϕ as α approaches 0, implying that $h(\tilde{\phi}^\alpha)$ converges to $h(\phi)$. Since $\tilde{\phi}^\alpha_{[0, \alpha]}$ is a solution to (10), the cost of this portion of path $\tilde{\phi}^\alpha$ is 0. Along the remainder of path $\tilde{\phi}^\alpha$, the components $(\tilde{\phi}^\alpha_t)_i$ are bounded below by a multiple of α and by a multiple of $(\phi_{t-\alpha})_i$, and the derivatives $\dot{\tilde{\phi}}^\alpha_t$ are approximately equal to the derivatives $\dot{\phi}_{t-\alpha}$ (Lemma 5). Combining these facts with the definition (15) of path costs in terms of the Cramér transform L , the boundedness of $L(\cdot, z)$ on compact subsets of $\text{int}(X)$, the convexity of $L(x, \cdot)$, and the variational representation (14) of $L(x, z)$ in terms of relative entropies $R(\cdot \| \nu(\cdot | x))$, we argue that the cost of path segment $\tilde{\phi}^\alpha_{[\alpha, 1]}$ converges to the cost of path ϕ as α approaches 0.

The next proposition, a version of DE, lemma 6.5.5, allows us to further restrict our attention to paths having convenient regularity properties. We let $\mathcal{C}^* \subset \tilde{\mathcal{C}}$ denote the set of paths $\phi \in \tilde{\mathcal{C}}$ such that after the time $\alpha > 0$ such that $\phi^\alpha_{[0, \alpha]}$ solves (10), the derivative $\dot{\phi}$ is piecewise constant and takes values in Z .

Proposition 10. $\inf_{\phi \in \tilde{\mathcal{C}}} (c_x(\phi) + h(\phi)) = \inf_{\phi \in \mathcal{C}^*} (c_x(\phi) + h(\phi))$.

The proof of Proposition 10 mimics that of DE, lemma 6.5.5; see SS, section 9.2 for details.

Now fix $\varepsilon > 0$. By the previous two propositions, we can choose an $\alpha > 0$ and a path $\psi \in \mathcal{C}^*$ such that $\psi_{[0, \alpha]}$ solves (10) and

$$c_x(\psi) + h(\psi) \leq \inf_{\phi \in \mathcal{C}} (c_x(\phi) + h(\phi)) + \varepsilon. \tag{53}$$

We now introduce a controlled process χ^N that follows ψ in expectation as long as it remains in a neighborhood of ψ . Representation (14) implies that for each $k \in \{0, \dots, N - 1\}$ and $x \in \mathcal{X}^N$, there is a transition kernel $\pi_k^N(\cdot | x)$ that minimizes relative entropy with respect to $\nu(\cdot | x)$ subject to the aforementioned constraint on its expectation:

$$R(\pi_k^N(\cdot | x) \| \nu(\cdot | x)) = L(x, \dot{\psi}_{k/N}) \quad \text{and} \quad \sum_{z \in Z} z \pi_k^N(z | x) = \dot{\psi}_{k/N}. \tag{54}$$

To ensure that this definition makes sense for all k , we replace the piecewise constant function $\dot{\psi}$ with its right continuous version.

Let $Z(x) = \text{conv}(Z(x)) = \{z \in Z: z_j \geq 0 \text{ for all } j \notin \text{supp}(x)\}$ be the set of feasible expected increments of the original process \mathbf{X}^N from state x (cf. Endnote 13). Since $\psi_{[0, \alpha]}$ solves (10), inequality (6) implies that

$$\varsigma - (\psi_t)_i \leq (\dot{\psi}_t)_i \leq 1 \tag{55}$$

for all $t \in [0, \alpha]$ and $i \in \mathcal{A}$. It follows that there is an $\hat{\alpha} \in (0, \alpha]$ such that

$$\dot{\psi}_t \in Z(x) \quad \text{whenever } t \in [0, \hat{\alpha}]. \tag{56}$$

In addition, property (55) implies that for all $t \in [0, \alpha]$ and $i \in \text{supp}(x)$, we have

$$(\psi_t)_i = (\psi_0)_i + \int_0^t (\dot{\psi}_s)_i ds \geq x_i + \int_0^t (\zeta - (\psi_s)_i) ds = e^{-t} x_i + (1 - e^{-t}) \zeta \geq x_i \wedge \zeta,$$

where the second equality follows from the fact that $v_t = e^{-t} x_i + (1 - e^{-t}) \zeta$ is the solution to $\dot{v} = \zeta - v$ from initial condition $v_0 = x_i$.

Now choose a $\delta > 0$ satisfying

$$\delta < \min(\{\zeta\} \cup \{x_i : i \in \text{supp}(x)\} \cup \{\frac{1}{2} \text{dist}(\psi_i, \partial X) : t \in [\hat{\alpha}, 1]\}). \quad (57)$$

For future reference, note that if $y \in X$ satisfies $|y - \psi_t| < \delta$, then $|y_i - (\psi_t)_i| < \delta/2$ for all $i \in \mathcal{A}$ (by the definition of the l^1 norm and the fact that $y - \psi_t \in \mathbb{R}_0^n$), and so if $t \in [0, \alpha]$, we also have

$$y \in \bar{X}_x \equiv \{\hat{x} \in X : \hat{x}_i \geq \frac{1}{2}(x_i \wedge \zeta) \text{ for all } i \in \text{supp}(x)\}. \quad (58)$$

For each $(x_0, \dots, x_k) \in (\mathcal{X}^N)^{k+1}$, define the sequence of controls $\{\lambda_k^N\}_{k=0}^{N-1}$ by

$$\lambda_k^N(x | x_0, \dots, x_k) = \begin{cases} \pi_k^N(x | x_k), & \text{if } \max_{0 \leq j \leq k} |x_j - \psi_{j/N}| \leq \delta, \\ \nu^N(x | x_k), & \text{if } \max_{0 \leq j \leq k} |x_j - \psi_{j/N}| > \delta. \end{cases} \quad (59)$$

Finally, define the controlled process $\chi^N = \{\chi_k^N\}_{k=0}^N$ by setting $\chi_0^N = x^N$ and using the recursive formula $\chi_{k+1}^N = \chi_k^N + (1/N)\zeta_k^N$, where ζ_k^N has law $\lambda_k^N(\cdot | \chi_0^N, \dots, \chi_k^N)$. Under this construction, the process χ^N evolves according to the transition kernels π_k^N and so follows ψ in expectation, so long as it stays δ -synchronized with ψ . If this ever fails to be true, the evolution of χ^N proceeds according to the kernel ν^N of the original process \mathbf{X}^N . This implies that until the stopping time

$$\tau^N = \frac{1}{N} \min\{k \in \{0, 1, \dots, N\} : |\chi_k^N - \psi_{k/N}| > \delta\} \wedge 1, \quad (60)$$

the relative entropies of transitions are given by (54) and that after τ^N these relative entropies are zero.

Define the pair $\{(\Lambda^N, \hat{\xi}^N)\}_{N=N_0}^\infty$ of interpolated processes as in Section 5.4. Proposition 8 shows that for any subsequence of these pairs, there is a subsubsequence that converges in distribution to some pair $(\lambda_t \otimes dt, \xi)$ satisfying (49). By Prokhorov's theorem, τ^N can be assumed to converge in distribution on this subsubsequence to some $[0, 1]$ -valued random variable τ . Indeed, DE, lemma 6.4.2 and proposition 5.3.8 imply that $\tau = 1$ and $\chi = \psi$ with probability 1.

For the subsubsequence specified above, we argue as follows:

$$\begin{aligned} \limsup_{N \rightarrow \infty} V^N(x^N) &\leq \lim_{N \rightarrow \infty} \mathbb{E}_{x^N} \left[\frac{1}{N} \sum_{j=0}^{N-1} R(\lambda_j^N(\cdot | \chi_0^N, \dots, \chi_j^N) \| \nu^N(\cdot | \chi_j^N)) + h(\hat{\chi}^N) \right] \\ &= \lim_{N \rightarrow \infty} \mathbb{E}_{x^N} \left[\frac{1}{N} \sum_{j=0}^{N\tau^N-1} L(\chi_j^N, \dot{\psi}_{j/N}) + h(\hat{\chi}^N) \right] \\ &= \lim_{N \rightarrow \infty} \mathbb{E}_{x^N} \left[\frac{1}{N} \sum_{j=0}^{(N\tau^N \wedge [N\hat{\alpha}]) - 1} L(\hat{\chi}_{j/N}^N, \dot{\psi}_{j/N}) + \frac{1}{N} \sum_{j=N\tau^N \wedge [N\hat{\alpha}]}^{N\tau^N-1} L(\hat{\chi}_{j/N}^N, \dot{\psi}_{j/N}) + h(\hat{\chi}^N) \right] \\ &= \int_0^{\hat{\alpha}} L(\psi_t, \dot{\psi}_t) dt + \int_{\hat{\alpha}}^1 L(\psi_t, \dot{\psi}_t) dt + h(\psi) = c_x(\psi) + h(\psi) \leq \inf_{\phi \in \mathcal{E}} (c_x(\phi) + h(\phi)) + \varepsilon. \end{aligned} \quad (61)$$

The initial inequality follows from representation (47); the second line from the uniform convergence in (51), along with Equations (54), (59), and (60); the fifth line from the definition of c_x ; and the sixth line from inequality (53). The fourth line is a consequence of the continuity of h , the convergence of τ^N to τ , the uniform convergence of $\hat{\chi}^N$ to ψ , the piecewise continuity and right continuity of $\dot{\psi}$, Skorokhod's theorem, and the dominated convergence theorem. For the application of dominated convergence to the first sum, we use the fact that when $j/N < \tau^N \wedge \hat{\alpha}$, we have $\hat{\chi}_{j/N}^N \in \bar{X}_x$ (see (58)) and $\dot{\psi}_{j/N} \in Z(x)$ (see (56)), along with the fact that $L(\cdot, \cdot)$ is continuous, and hence uniformly continuous and bounded, on $\bar{X}_x \times Z(x)$, as follows from Proposition 5(ii). For the application of dominated convergence to the second sum, we use the fact when $\hat{\alpha} \leq j/N < \tau^N$, we have $\text{dist}(\hat{\chi}_{j/N}^N, \partial X) \geq \delta/2$ (see (57)), and the fact that $L(\cdot, \cdot)$ is continuous and bounded on $Y \times Z$ when Y is a closed subset of $\text{int}(X)$, as follows from Proposition 5(i).

Since every subsequence has a subsubsequence that satisfies (61), the sequence as a whole must satisfy (61). Therefore, since $\varepsilon > 0$ was arbitrary, (61) and (39) establish inequality (52), and hence the lower bound (38b).

6. Proof of Lemma 2

Lemma 2 follows from Equation (34) and Lemma 4, which in turn requires Lemma 3. Lemma 3 shows that for any distribution λ on $\mathcal{Z} = \{e_j - e_i : i, j \in \mathcal{A}\}$ with mean $z \in Z(I) = \text{conv}(Z(I))$, there is a distribution $\bar{\lambda}$ on $Z(I) = \{e_j - e_i : j \in \mathcal{A}, i \in I\}$ whose mean is also z , with the variational distance between $\bar{\lambda}$ and λ bounded by a fixed multiple of the mass that λ places on components outside of $Z(I)$. The lemma also specifies some equalities that λ and $\bar{\lambda}$ jointly satisfy. Lemma 4 shows that if x puts little mass on actions outside I , then the reduction in relative entropy obtained by switching from $\bar{\lambda}$ to λ is small at best, uniformly over the choice of displacement vector $z \in Z(I)$.

Both lemmas require additional notation. Throughout what follows, we write K for $\mathcal{A} \setminus I$. For $\lambda \in \Delta(\mathcal{Z})$, we write λ_{ij} for $\lambda(e_j - e_i)$ when $j \neq i$. We write $\lambda_{[i]} = \sum_{j \neq i} \lambda_{ij}$ for the i th “row sum” of λ and $\lambda^{[j]} = \sum_{i \neq j} \lambda_{ij}$ for the j th “column sum.” (Remember that λ has no “diagonal components” but instead has a single null component $\lambda_0 = \lambda(\mathbf{0})$.) For $I \subseteq \mathcal{A}$, we write $\lambda_I = \sum_{i \in I} \lambda_{[i]}$ for the sum over all elements of λ from rows in I . If $\lambda, \bar{\lambda} \in \Delta(\mathcal{Z})$, we apply the same notational devices to $\Delta\lambda = \bar{\lambda} - \lambda$ and to $|\Delta\lambda|$, the latter of which is defined by $|\Delta\lambda|_{ij} = |(\Delta\lambda)_{ij}|$. Finally, if $\chi \in \mathbb{R}_+^I$, we write $\chi_{[i]}$ for $\sum_{i \in I} \chi_i$.

Lemma 3. Fix $z \in Z(I)$ and $\lambda \in \Lambda_z(z)$. Then there exist a distribution $\bar{\lambda} \in \Lambda_{Z(I)}(z)$ and a vector $\chi \in \mathbb{R}_+^I$ satisfying

- (i) $\Delta\lambda_{[i]} = \Delta\lambda^{[i]} = -\chi_i$ for all $i \in I$,
- (ii) $\Delta\lambda_0 = \lambda_{[K]} + \chi_{[I]}$,
- (iii) $\chi_{[I]} \leq \lambda_{[K]}$, and
- (iv) $|\Delta\lambda|_{[S]} \leq 3\lambda_{[K]}$,

where $\Delta\lambda = \bar{\lambda} - \lambda$.

Lemma 4. Fix $\varepsilon > 0$. There exists a $\delta > 0$ such that for any $x \in X$ with $\bar{x}_K \equiv \max_{k \in K} x_k < \delta$, any $z \in Z(I)$, and any $\lambda \in \Lambda_{Z(\text{supp}(x))}(z)$, we have

$$R(\lambda \| \nu(\cdot | x)) \geq R(\bar{\lambda} \| \nu(\cdot | x)) - \varepsilon,$$

where $\bar{\lambda} \in \Lambda_{Z(I)}(z)$ is the distribution determined for λ in Lemma 3.

To see that Lemma 4 implies Lemma 2, fix $x \in X$ with $\bar{x}_K < \delta$ and $z \in Z(I)$, and let $\lambda \in \Lambda_z(z)$ and $\lambda^I \in \Lambda_{Z(I)}(z)$ be the minimizers in (34) and (33), respectively; then since $\bar{\lambda} \in \Lambda_{Z(I)}(z)$,

$$L(x, z) = R(\lambda \| \nu(\cdot | x)) \geq R(\bar{\lambda} \| \nu(\cdot | x)) - \varepsilon \geq R(\lambda^I \| \nu(\cdot | x)) - \varepsilon = L_I(x, z) - \varepsilon.$$

6.1. Proof of Lemma 3

To prove Lemma 3, we introduce an algorithm that incrementally constructs the pair $(\bar{\lambda}, \chi) \in \Lambda_{Z(I)}(z) \times \mathbb{R}_+^I$ from the pair $(\lambda, \mathbf{0}) \in \Lambda_z(z) \times \mathbb{R}_+^I$. All intermediate states (ν, π) of the algorithm are in $\Lambda_z(z) \times \mathbb{R}_+^I$.

The intuition behind the algorithm is as follows. Suppose that distribution $\lambda \in \Lambda_z(z)$ has mean $z \in Z(I)$, so that $z_k \geq 0$ for all $k \notin I$. Suppose too that $\lambda(e_i - e_k) > 0$ for some $e_i - e_k \notin Z(I)$ (i.e., with $k \notin I$ and $i \neq k$), so that $\lambda \notin \Lambda_{Z(I)}(z)$. We would like to reduce the mass that distribution λ places on outcome $e_i - e_k$ while keeping the mean of the distribution fixed. To do so, we observe that there must be some $j \in \mathcal{A}$ such that $\lambda(e_k - e_j) > 0$, as otherwise, z_k would be negative. Now let $c = \min\{\lambda(e_i - e_k), \lambda(e_k - e_j)\}$. If $i \neq j$, we can remove mass c from $\lambda(e_i - e_k)$ and from $\lambda(e_k - e_j)$ and add mass c to $\lambda(e_i - e_j)$ and to $\lambda(\mathbf{0})$ without affecting the distribution’s mean (see Case 1). If, instead, $i = j$, then we can remove mass c from $\lambda(e_i - e_k)$ and from $\lambda(e_k - e_j)$ and add mass $2c$ to $\lambda(\mathbf{0})$ (see Cases 2 and 3). Either way, this adjustment reduces the mass placed on some outcome involving action k to zero, and it does not increase the mass on any outcome involving action k . Thus adjustments of this sort will eliminate all mass on increments outside of $Z(I)$ in a finite number of steps and so will result in a new distribution $\bar{\lambda}$ in $\Lambda_{Z(I)}(z)$.

We proceed with the formal definition and analysis of the algorithm. Suppose, without loss of generality, that $K = \{1, \dots, k\}$. The algorithm first reduces all elements of the first row of λ to 0, then all elements of the second row, and eventually all elements of the k th row.

Suppose that at some stage of the algorithm, the state is $(\nu, \pi) \in \Lambda_z(z) \times \mathbb{R}_+^I$, rows 1 through $k - 1$ have been zeroed, and row k has not been zeroed:

$$\nu_{[h]} = 0, \quad \text{for all } h < k, \quad \text{and} \tag{62}$$

$$\nu_{[k]} > 0. \tag{63}$$

Since $v \in \Lambda_z(z)$ and $z \in Z(I)$,

$$v^{[i]} - v_{[i]} = \sum_{z \in Z} x_i v(z) = z_i, \quad \text{for all } i \in I, \quad \text{and} \quad (64)$$

$$v^{[l]} - v_{[l]} = \sum_{z \in Z} x_l v(z) = z_l \geq 0, \quad \text{for all } l \in K. \quad (65)$$

Equations (63) and (65) together imply that $v^{[k]} > 0$. Thus there exist $j \neq k$ and $i \neq k$ such that

$$v_{kj} \wedge v_{ik} \equiv c > 0. \quad (66)$$

Using (66), we now construct the algorithm's next state $(\hat{v}, \hat{\pi})$ from the current state (v, π) using one of three mutually exclusive and exhaustive cases, as described next; only the components of v and π whose values change are noted explicitly.

<i>Case 1</i> ($i \neq j$):	<i>Case 2</i> ($i = j \in K$):	<i>Case 3</i> ($i = j \in I$):
$\hat{v}_{kj} = v_{kj} - c,$	$\hat{v}_{kj} = v_{kj} - c,$	$\hat{v}_{kj} = v_{kj} - c,$
$\hat{v}_{ik} = v_{ik} - c,$	$\hat{v}_{jk} = v_{jk} - c,$	$\hat{v}_{jk} = v_{jk} - c,$
$\hat{v}_{ij} = v_{ij} + c,$		
$\hat{v}_0 = v_0 + c.$	$\hat{v}_0 = v_0 + 2c.$	$\hat{v}_0 = v_0 + 2c,$
		$\hat{\pi}_i = \pi_i + c.$

In what follows, we confirm that following the algorithm to completion leads to a final state $(\bar{\lambda}, \chi)$ with the desired properties.

Write $\Delta v = \hat{v} - v$ and $\Delta \pi = \hat{\pi} - \pi$, and define $|\Delta v|$ componentwise as described above. The following statements are immediate:

<i>Case 1</i> ($i \neq j$):	<i>Case 2</i> ($i = j \in K$):	<i>Case 3</i> ($i = j \in I$):	
$\Delta v_{[k]} = \Delta v^{[k]} = -c,$	$\Delta v_{[k]} = \Delta v^{[k]} = -c,$	$\Delta v_{[k]} = \Delta v^{[k]} = -c,$	(67a)

$\Delta v_{[j]} = \Delta v^{[j]} = -c,$	(67b)
---	-------

$\Delta v_{[l]} = \Delta v^{[l]} = 0, l \neq k,$	$\Delta v_{[l]} = \Delta v^{[l]} = 0, l \notin \{k, j\},$	$\Delta v_{[l]} = \Delta v^{[l]} = 0, l \notin \{k, j\},$	(67c)
--	---	---	-------

$\Delta v_0 = c,$	$\Delta v_0 = 2c,$	$\Delta v_0 = 2c,$	(67d)
-------------------	--------------------	--------------------	-------

$\Delta \pi_{[i]} = 0,$	$\Delta \pi_{[i]} = 0,$	$\Delta \pi_{[i]} = c,$	(67e)
-------------------------	-------------------------	-------------------------	-------

$\Delta v_{[K]} = -c,$	$\Delta v_{[K]} = -2c,$	$\Delta v_{[K]} = -c,$	(67f)
------------------------	-------------------------	------------------------	-------

$ \Delta v _{[A]} = 3c.$	$ \Delta v _{[A]} = 2c.$	$ \Delta v _{[S]} = 2c.$	(67g)
--------------------------	--------------------------	--------------------------	-------

The initial equalities in (67a)–(67c) imply that if v is in $\Lambda_z(z)$, then so is \hat{v} . Equations (67a)–(67c) also imply that no step of the algorithm increases the mass in any row of v . Moreover, (66) and the definition of the algorithm imply that during each step, no elements of the k th row or the k th column of v are increased and that at least one such element is zeroed. It follows that the first row is zeroed in at most $2n - 3$ steps, followed by the second row, and ultimately followed by row \bar{k} . Thus a terminal state with $\bar{\lambda}_{[K]} = 0$, and hence with $\bar{\lambda} \in \Lambda_{z(I)}(z)$, is reached in a finite number of steps. For future reference, we note that

$$\Delta \lambda_{[K]} = \bar{\lambda}_{[K]} - \lambda_{[K]} = -\lambda_{[K]}. \quad (68)$$

We now verify conditions (i)–(iv) from the statement of the lemma. First, for any $i \in I$, (67b), (67c), and (67e) imply that in all three cases of the algorithm, $\Delta v_{[i]} = \Delta v^{[i]} = -\Delta \pi_{[i]}$; the common value is 0 in Cases 1 and 2 and $-c$ in Case 3. Thus aggregating over all steps of the algorithm yields $\Delta \lambda_{[i]} = \Delta \lambda^{[i]} = -\chi_i$, which is condition (i).

Second, (67d)–(67f) imply that in all three cases, $\Delta v_0 = -\Delta v_{[K]} + \Delta \pi_{[I]}$. Aggregating over all steps of the algorithm yields $\Delta \lambda_0 = -\Delta \lambda_{[K]} + \chi_{[I]}$. Then substituting (68) yields $\Delta \lambda_0 = \lambda_{[K]} + \chi_{[I]}$, which is condition (ii).

Third, (67e) and (67f) imply that in all three cases, $\Delta \pi_{[i]} \leq -\Delta v_{[K]}$. Aggregating and using (68) yields $\chi_{[I]} \leq -\Delta \lambda_{[K]} = \lambda_{[K]}$, establishing (iii).

Fourth, (67f) and (67g) imply that in all three cases, $|\Delta v|_{[A]} \leq -3\Delta v_{[K]}$. Aggregating and using (68) yields $|\Delta \lambda|_{[A]} \leq -3\Delta \lambda_{[K]} = 3\lambda_{[K]}$, establishing (iv).

This completes the proof of Lemma 3.

6.2. Proof of Lemma 4

To prove Lemma 4, it is natural to introduce the notation $d = \lambda - \bar{\lambda} = -\Delta\lambda$ to represent the increment generated by a move from distribution $\bar{\lambda} \in \Lambda_{z(I)}(z)$ to distribution $\lambda \in \Lambda_z(z)$. We will show that when $\bar{x}_K = \max_{k \in I} x_i$ is small, such a move can only result in a slight reduction in relative entropy.

To start, observe that

$$d_{[i]} = d^{[i]} = \chi_i, \quad \text{for all } i \in I, \tag{69}$$

$$d_{[k]} = d^{[k]} = \lambda_{[k]}, \quad \text{for all } k \in K, \tag{70}$$

$$d_0 = -\lambda_{[K]} - \chi_{[I]} \geq -2\lambda_{[K]}, \quad \text{and} \tag{71}$$

$$|d|_{[A]} \leq 3\lambda_{[K]}. \tag{72}$$

Display (69) follows from part (i) of Lemma 3, display (71) from parts (ii) and (iii), and display (72) from part (iv). For display (70), note first that since λ and $\bar{\lambda}$ are both in $\Lambda_z(z)$, we have

$$\lambda^{[k]} - \lambda_{[k]} = \sum_{z \in Z} z_k \lambda(z) - \sum_{z \in Z} z_k \bar{\lambda}(z) = \bar{\lambda}^{[k]} - \bar{\lambda}_{[k]},$$

and hence

$$d^{[k]} = \lambda^{[k]} - \bar{\lambda}^{[k]} = \lambda_{[k]} - \bar{\lambda}_{[k]} = d_{[k]},$$

then (70) follows from the fact that $\bar{\lambda}_{[k]} = 0$, which is true since $\bar{\lambda} \in \Lambda_{z(I)}(z)$.

By definition,

$$R(\lambda \| \nu(\cdot | x)) = \sum_{i \in A} \sum_{j \neq i} (\lambda_{ij} \log \lambda_{ij} - \lambda_{ij} \log x_i \sigma_{ij}) + \lambda_0 \log \lambda_0 - \lambda_0 \log \Sigma,$$

where $\sigma_{ij} = \sigma_{ij}(x)$ and $\Sigma = \sum_{j \in A} x_j \sigma_{jj}$. Thus, writing $l(p) = p \log p$ for $p \in (0, 1]$ and $l(0) = 0$, we have

$$R(\lambda \| \nu(\cdot | x)) - R(\bar{\lambda} \| \nu(\cdot | x)) = \sum_{i \in A} \sum_{j \neq i} (l(\lambda_{ij}) - l(\bar{\lambda}_{ij})) + (l(\lambda_0) - l(\bar{\lambda}_0)) \tag{73}$$

$$- \left(\sum_{i \in A} \sum_{j \neq i} d_{ij} \log x_i \sigma_{ij} + d_0 \log \Sigma \right).$$

We first use (69)–(71), Lemma 3(iii), and the facts that $\chi \geq 0$ and $\Sigma \geq \varsigma$ to obtain a lower bound on the final term of (73):

$$- \left(\sum_{i \in A} \sum_{j \neq i} d_{ij} \log x_i \sigma_{ij} + d_0 \log \Sigma \right)$$

$$= - \sum_{i \in I} \chi_i \log x_i - \sum_{k \in K} \lambda_{[k]} \log x_k - \sum_{i \in A} \sum_{j \neq i} d_{ij} \log \sigma_{ij} + \left(\sum_{i \in I} \chi_i + \sum_{k \in K} \lambda_{[k]} \right) \log \Sigma$$

$$\geq \sum_{i \in I} \chi_i \log \Sigma + \sum_{k \in K} \lambda_{[k]} \log \frac{\Sigma}{x_k} + \sum_{i \in A} \sum_{j \neq i} |d_{ij}| \log \sigma_{ij}$$

$$\geq \chi_{[I]} \log \varsigma + \lambda_{[K]} \log \frac{\varsigma}{\bar{x}_K} + |d|_{[A]} \log \varsigma$$

$$\geq \lambda_{[K]} \log \varsigma + \lambda_{[K]} \log \frac{\varsigma}{\bar{x}_K} + 3\lambda_{[K]} \log \varsigma$$

$$\geq \lambda_{[K]} \log \frac{\varsigma^5}{\bar{x}_K}.$$

To bound the initial terms on the right-hand side of (73), observe that the function $l: [0, 1] \rightarrow \mathbb{R}$ is convex, nonpositive, and minimized at e^{-1} , where $l(e^{-1}) = -e^{-1}$. Define the convex function $\hat{l}: [0, 1] \rightarrow \mathbb{R}$ by $\hat{l}(p) = l(p)$ if $p \leq e^{-1}$ and $\hat{l}(p) = -e^{-1}$ otherwise. We now argue that for any $p, q \in [0, 1]$, we have

$$-|l(p) - l(q)| \geq \hat{l}(|p - q|). \tag{74}$$

Since l is nonpositive with minimum $-e^{-1}$, $-|l(p) - l(q)| \geq -e^{-1}$ always holds. If $|p - q| \leq e^{-1}$, then $-|l(p) - l(q)| \geq \min\{l(|p - q|), l(1 - |p - q|)\} = l(|p - q|)$; the inequality follows from the convexity of l , and the equality obtains

because $l(r) - l(1-r) \leq 0$ for $r \in [0, \frac{1}{2}]$, which is true because $r \mapsto l(r) - l(1-r)$ is convex on $[0, \frac{1}{2}]$ and equals 0 at the end points. Together, these claims yield (74).

Together, inequality (74), Jensen's inequality, and inequality (72) imply that

$$\begin{aligned} \sum_{i \in \mathcal{A}} \sum_{j \neq i} (l(\lambda_{ij}) - l(\bar{\lambda}_{ij})) &\geq - \sum_{i \in \mathcal{A}} \sum_{j \neq i} |l(\lambda_{ij}) - l(\bar{\lambda}_{ij})| \geq \sum_{i \in \mathcal{A}} \sum_{j \neq i} \hat{l}(|\lambda_{ij} - \bar{\lambda}_{ij}|) \\ &\geq (n^2 - n) \hat{l} \left(\frac{\sum_{i \in \mathcal{A}} \sum_{j \neq i} |\lambda_{ij} - \bar{\lambda}_{ij}|}{n^2 - n} \right) \geq (n^2 - n) \hat{l} \left(\frac{3}{n^2 - n} \lambda_{[K]} \right). \end{aligned}$$

Finally, display (71) implies that $d_0 = \lambda_0 - \bar{\lambda}_0 \in [-2\lambda_{[K]}, 0]$. Since $l: [0, 1] \rightarrow \mathbb{R}$ is convex with $l(1) = 0$ and $l'(1) = 1$, it follows that

$$l(\lambda_0) - l(\bar{\lambda}_0) \geq l(1 + d_0) - l(1) \geq d_0 \geq -2\lambda_{[K]}.$$

Substituting three of the last four displays into (73), we obtain

$$R(\lambda \| \nu(\cdot | x)) - R(\bar{\lambda} \| \nu(\cdot | x)) \geq (n^2 - n) \hat{l} \left(\frac{3}{n^2 - n} \lambda_{[K]} \right) + \lambda_{[K]} \left(\log \frac{\zeta^5}{\bar{x}_K} - 2 \right).$$

To complete the proof of the lemma, it is enough to show that if $\bar{x}_K \leq \zeta^5 e^{-2}$, then

$$(n^2 - n) \hat{l} \left(\frac{3}{n^2 - n} \lambda_{[K]} \right) + \lambda_{[K]} \left(\log \frac{\zeta^5}{\bar{x}_K} - 2 \right) \geq -(n^2 - n) \left(\frac{\bar{x}_K}{e \zeta^5} \right)^{1/3}. \quad (75)$$

We do so by computing the minimum value of the left-hand side of (75) over $\lambda_{[K]} \geq 0$. Let $a = n^2 - n$ and $c = \log \zeta^5 / \bar{x}_K - 2$. The assumption that $\bar{x}_K \leq \zeta^5 e^{-2}$ then becomes $c \geq 0$. Thus if $s > a/(3e)$, then

$$a \hat{l} \left(\frac{3}{a} s \right) + cs = -ae^{-1} + cs \geq -ae^{-1} + c \cdot \frac{a}{3e} = a \hat{l} \left(\frac{3}{a} \cdot \frac{a}{3e} \right) + c \cdot \frac{a}{3e}.$$

Thus if $s \leq a/(3e)$ minimizes $a \hat{l}((3/a)s) + cs$ over $s \geq 0$, it also minimizes $a \hat{l}((3/a)s) + cs$, and the minimized values are the same. To minimize $a \hat{l}((3/a)s) + cs$, note that it is concave in s ; solving the first-order condition yields the minimizer, $s^* = (a/3) \exp(-c/3 - 1)$. This is less than or equal to $a/(3e)$ when $c \geq 0$. Plugging s^* into the objective function yields $-a \exp(-c/3 - 1)$, and substituting in the values of a and c and simplifying yields the right-hand side of (75).

This completes the proof of Lemma 4.

7. Proof of Proposition 9

It remains to prove Proposition 9. Inequality (6) implies that solutions $\bar{\phi}$ to the mean dynamic (10) satisfy

$$\zeta - \bar{\phi}_i \leq \dot{\bar{\phi}}_i \leq 1 \quad (76)$$

for every action $i \in \mathcal{A}$. Thus if $(\bar{\phi}_0)_i \leq \zeta/2$, then for all $t \in [0, \zeta/4]$, the upper bound in (76) yields $(\bar{\phi}_t)_i \leq 3\zeta/4$. Then the lower bound yields $(\bar{\phi}_t)_i \geq \zeta - 3\zeta/4 = \zeta/4$, and thus

$$(\bar{\phi}_0)_i \leq \frac{1}{2}\zeta \text{ implies that } (\bar{\phi}_t)_i - (\bar{\phi}_0)_i \geq \frac{1}{4}\zeta t, \text{ for all } t \in [0, \frac{1}{4}\zeta]. \quad (77)$$

In addition, it follows easily from (76) and (77) that solutions $\bar{\phi}$ of (10) from every initial condition in X satisfy

$$(\bar{\phi}_t)_i \geq \frac{1}{4}\zeta t \text{ for all } t \in [0, \frac{1}{4}\zeta]. \quad (78)$$

Fix $\phi \in \mathcal{C}_x$ with $c_x(\phi) < \infty$, so that ϕ is absolutely continuous, with $\dot{\phi}_t \in Z$ at all $t \in [0, 1]$ where ϕ is differentiable. Let $\bar{\phi} \in \mathcal{C}_x$ be the solution to (10) starting from x . Then for $\alpha \in (0, \zeta/4]$, define trajectory $\phi^\alpha \in \tilde{\mathcal{C}}_x$ as follows:

$$\phi_t^\alpha = \begin{cases} \bar{\phi}_t, & \text{if } t \leq \alpha, \\ \bar{\phi}_\alpha + (1 - (2/\zeta)\alpha)(\phi_{t-\alpha} - x), & \text{if } t > \alpha. \end{cases} \quad (79)$$

Thus ϕ^α follows the solution to the mean dynamic from x until time α ; then, the increments of ϕ^α starting at time α mimic those of ϕ starting at time 0 but are slightly scaled down.

The next lemma describes the key properties of ϕ^α . In part (ii) and hereafter, $|\cdot|$ denotes the l^1 norm on \mathbb{R}^n .

Lemma 5. *If $\alpha \in (0, \varsigma/4]$ and $t \in [\alpha, 1]$, then*

- (i) $\dot{\phi}_t^\alpha = (1 - (2/\varsigma)\alpha)\dot{\phi}_{t-\alpha}$;
- (ii) $|\dot{\phi}_t^\alpha - \dot{\phi}_{t-\alpha}| \leq (2 + 4/\varsigma)\alpha$;
- (iii) for all $i \in \mathcal{A}$, $(\phi_t^\alpha)_i \geq (\varsigma/4)\alpha$; and
- (iv) for all $i \in \mathcal{A}$, $(\phi_t^\alpha)_i \geq (\varsigma/12)(\phi_{t-\alpha})_i$.

Proof. Part (i) is immediate. For part (ii), combine the fact that $\bar{\phi}_t$ and ϕ_t both move at an l^1 speed of at most 2 (since both have derivatives in Z) and the identity $\phi_{t-\alpha} = \phi_0 + (\phi_{t-\alpha} - \phi_0)$ with the definition of ϕ_t^α to obtain

$$|\phi_t^\alpha - \phi_{t-\alpha}| \leq |\bar{\phi}_\alpha - x| + \frac{2}{\varsigma}\alpha(\phi_{t-\alpha} - x) \leq 2\alpha + \frac{2}{\varsigma}\alpha \cdot 2(t - \alpha) \leq \left(2 + \frac{4}{\varsigma}\right)\alpha.$$

We turn to part (iii). If $(\phi_{t-\alpha} - x)_i \geq 0$, then it is immediate from definition (79) and inequality (78) that $(\phi_t^\alpha)_i \geq (\varsigma/4)\alpha$. So suppose instead that $(\phi_{t-\alpha} - x)_i < 0$. Then (79) and the fact that $(\phi_{t-\alpha})_i \geq 0$ imply that

$$(\phi_t^\alpha)_i \geq (\bar{\phi}_\alpha)_i - \left(1 - \frac{2}{\varsigma}\alpha\right)x_i = ((\bar{\phi}_\alpha)_i - x_i) + \frac{2}{\varsigma}\alpha x_i. \tag{80}$$

If $x_i \leq \varsigma/2$, then (80) and (77) yield

$$(\phi_t^\alpha)_i \geq ((\bar{\phi}_\alpha)_i - x_i) + \frac{2}{\varsigma}\alpha x_i \geq \frac{\varsigma}{4}\alpha + 0 = \frac{\varsigma}{4}\alpha.$$

If $x_i \in [\varsigma/2, \varsigma]$, then (80) and (76) yield

$$(\phi_t^\alpha)_i \geq ((\bar{\phi}_\alpha)_i - x_i) + \frac{2}{\varsigma}\alpha x_i \geq 0 + \frac{2}{\varsigma}\alpha \cdot \frac{\varsigma}{2} \geq \alpha.$$

And if $x_i \geq \varsigma$, then (80), (76), and the fact that $\dot{\phi}_i \geq -1$ yield

$$(\phi_t^\alpha)_i \geq ((\bar{\phi}_\alpha)_i - x_i) + \frac{2}{\varsigma}\alpha x_i \geq -\alpha + \frac{2}{\varsigma}\alpha \cdot \varsigma \geq \alpha.$$

It remains to establish part (iv). If $(\phi_{t-\alpha})_i = 0$, then there is nothing to prove. If $(\phi_{t-\alpha})_i \in (0, 3\alpha]$, then part (iii) implies that

$$\frac{(\phi_t^\alpha)_i}{(\phi_{t-\alpha})_i} \geq \frac{(\varsigma/4)\alpha}{3\alpha} = \frac{\varsigma}{12}.$$

And if $(\phi_{t-\alpha})_i \geq 3\alpha$, then definition (79) and the fact that $\dot{\phi}_i \geq -1$ and $\alpha \leq \varsigma/4$ imply that

$$\frac{(\phi_t^\alpha)_i}{(\phi_{t-\alpha})_i} \geq \frac{(\bar{\phi}_\alpha - x)_i + (1 - (2/\varsigma)\alpha)(\phi_{t-\alpha})_i + (2/\varsigma)\alpha x_i}{(\phi_{t-\alpha})_i} \geq -\frac{\alpha}{3\alpha} + 1 - \frac{2}{\varsigma}\alpha \geq \frac{2}{3} - \frac{1}{2} = \frac{1}{6}. \quad \square$$

Each trajectory ϕ^α is absolutely continuous, and Lemma 5(ii) and the fact that (10) is bounded imply that ϕ^α converges uniformly to ϕ as α approaches 0. This uniform convergence implies that

$$\lim_{\alpha \rightarrow 0} h(\phi^\alpha) = h(\phi). \tag{81}$$

Since $\phi_{[0,\alpha]}^\alpha$ is a solution to (10), and thus has cost zero, it follows from Lemma 5(i) and the convexity of $L(x, \cdot)$ that

$$c_x(\phi^\alpha) = \int_\alpha^1 L(\phi_t^\alpha, \dot{\phi}_t^\alpha) dt \leq \int_\alpha^1 L(\phi_t^\alpha, \dot{\phi}_{t-\alpha}) dt + \int_\alpha^1 \frac{2}{\varsigma}\alpha(L(\phi_t^\alpha, \mathbf{0}) - L(\phi_t^\alpha, \dot{\phi}_{t-\alpha})) dt. \tag{82}$$

To handle the second integral in (82), fix $t \in [\alpha, 1]$. Since $\phi_t^\alpha \in \text{int}(X)$, $\nu(\cdot | \phi_t^\alpha)$ has support Z , a set with extreme points $\text{ext}(Z) = \{e_j - e_i : j \neq i\} = Z \setminus \{0\}$. Therefore, the convexity of $L(x, \cdot)$, the final equality in (31), the lower bound (6), and Lemma 5(iii) imply that for all $z \in Z$,

$$L(\phi_t^\alpha, z) \leq \max_{i \in \mathcal{A}} \max_{j \neq i} L(\phi_t^\alpha, e_j - e_i) \leq -\log \left(\varsigma \min_{i \in \mathcal{A}} (\phi_t^\alpha)_i \right) \leq -\log \frac{\varsigma^2}{4} \alpha.$$

Thus since L is nonnegative and since $\lim_{\alpha \rightarrow 0} \alpha \log \alpha = 0$, the second integrand in (82) converges uniformly to zero, and so

$$\lim_{\alpha \rightarrow 0} \int_{\alpha}^1 \frac{2}{\zeta} \alpha \left(L(\phi_t^\alpha, \mathbf{0}) - L(\phi_t^\alpha, \dot{\phi}_{t-\alpha}) \right) dt = 0. \quad (83)$$

To bound the first integral in (82), note first that by representation (14), for each $t \in [0, 1]$ there is a probability measure $\lambda_t \in \Delta(\mathcal{Z})$ with $\lambda_t \ll \nu(\cdot | \phi_t)$ such that

$$L(\phi_t, \dot{\phi}_t) = R(\lambda_t \| \nu(\cdot | \phi_t)) \quad \text{and} \quad (84)$$

$$L(\phi_{t+\alpha}^\alpha, \dot{\phi}_t) \leq R(\lambda_t \| \nu(\cdot | \phi_{t+\alpha}^\alpha)), \quad \text{for all } \alpha \in (0, 1]. \quad (85)$$

DE, lemma 6.2.3 ensures that $\{\lambda_t\}_{t \in [0, 1]}$ can be chosen to be a measurable function of t . (Here and below, we ignore the measure zero set on which either $\dot{\phi}_t$ is undefined or $L(\phi_t, \dot{\phi}_t) = \infty$.)

Lemma 6 and expressions (84) and (85) imply that

$$\begin{aligned} \limsup_{\alpha \rightarrow 0} \int_{\alpha}^1 L(\phi_t^\alpha, \dot{\phi}_{t-\alpha}) dt &= \limsup_{\alpha \rightarrow 0} \int_0^{1-\alpha} L(\phi_{t+\alpha}^\alpha, \dot{\phi}_t) dt \leq \limsup_{\alpha \rightarrow 0} \left(\int_0^{1-\alpha} R(\lambda_t \| \nu(\cdot | \phi_{t+\alpha}^\alpha)) dt + \int_{1-\alpha}^1 0 dt \right) \\ &= \int_0^1 R(\lambda_t \| \nu(\cdot | \phi_t)) dt = \int_0^1 L(\phi_t, \dot{\phi}_t) dt = c_x(\phi), \end{aligned}$$

where the third line follows from the dominated convergence theorem and Lemma 6. Combining this inequality with (81), (82), and (83), we see that

$$\inf_{\phi \in \mathcal{E}^o} (c_x(\phi) - h(\phi)) \leq \limsup_{\alpha \rightarrow 0} (c_x(\phi^\alpha) - h(\phi^\alpha)) \leq c_x(\phi) - h(\phi).$$

Since $\phi \in \mathcal{E}$ was arbitrary, the result follows.

It remains to prove the following lemma.

Lemma 6. Write $R_{t+\alpha}^\alpha = R(\lambda_t \| \nu(\cdot | \phi_{t+\alpha}^\alpha))$ and $R_t = R(\lambda_t \| \nu(\cdot | \phi_t))$. Then

- (i) for all $t \in [0, 1)$, $\lim_{\alpha \rightarrow 0} R_{t+\alpha}^\alpha = R_t$; and
- (ii) for all $\alpha > 0$ small enough and $t \in [0, 1 - \alpha]$, $R_{t+\alpha}^\alpha \leq R_t + \log(12/\zeta) + 1$.

Proof. Definition (7) implies that

$$\begin{aligned} R_{t+\alpha}^\alpha - R_t &= \sum_{z \in \mathcal{Z}(\phi_t)} \lambda_t(z) \log \frac{\nu(z | \phi_t)}{\nu(z | \phi_{t+\alpha}^\alpha)} \\ &= \sum_{i \in \text{supp}(\phi_t)} \sum_{j \in \mathcal{A} \setminus \{i\}} \lambda_t(e_j - e_i) \left(\log \frac{(\phi_t)_i}{(\phi_{t+\alpha}^\alpha)_i} + \log \frac{\sigma_{ij}(\phi_t)}{\sigma_{ij}(\phi_{t+\alpha}^\alpha)} \right) + \lambda_t(\mathbf{0}) \log \frac{\sum_{i \in \mathcal{A}} (\phi_t)_i \sigma_{ii}(\phi_t)}{\sum_{i \in \mathcal{A}} (\phi_{t+\alpha}^\alpha)_i \sigma_{ii}(\phi_{t+\alpha}^\alpha)}. \quad (86) \end{aligned}$$

The uniform convergence from Lemma 5(ii) and the continuity of σ imply that for each $t \in [0, 1)$, the denominators of the fractions in (86) converge to their numerators as α vanishes, implying statement (i). The lower bound (6) then implies that the second and third logarithms in (86) themselves converge uniformly to zero as α vanishes; in particular, for α small enough and all $t \in [0, 1 - \alpha]$, these logarithms are bounded above by 1. Moreover, Lemma 5(iv) implies that when α is small enough and $t \in [0, 1 - \alpha]$ is such that $i \in \text{supp}(\phi_t)$, the first logarithm is bounded above by $\log(12/\zeta)$. Together, these claims imply statement (ii). This completes the proof of the lemma and hence the proof of Proposition 9.

8. Proofs and Auxiliary Results for Section 4

In the analyses in this section, we are often interested in the action of a function's derivative in directions $z \in \mathbb{R}_0^n$ that are tangent to the simplex. With this in mind, we let $\mathbf{1} \in \mathbb{R}^n$ be the vector of ones, and let $P = I - (1/n)\mathbf{1}\mathbf{1}'$ be the matrix that orthogonally projects \mathbb{R}^n onto \mathbb{R}_0^n . Given a function $g: \mathbb{R}^n \rightarrow \mathbb{R}$, we define the gradient of g with respect to \mathbb{R}_0^n by $\nabla_0 g(x) = P \nabla g(x)$, so that for $z \in \mathbb{R}_0^n$, we have $\nabla g(x)'z = \nabla g(x)'Pz = (P \nabla g(x))'z = \nabla_0 g(x)'z$.

Proof of Proposition 3. It is immediate from the definition of M^η that $M^\eta(\pi) = M^\eta(P\pi)$ for all $\pi \in \mathbb{R}^n$, leading us to introduce the notation $\bar{M}^\eta \equiv M^\eta|_{\mathbb{R}_0^n}$. Recalling that $h(x) = \sum_{k \in S} x_k \log x_k$ denotes the negated entropy function, one can verify by direct substitution that

$$\bar{M}^\eta: \mathbb{R}_0^n \rightarrow \text{int}(X) \quad \text{and} \quad \eta \nabla_0 h: \text{int}(X) \rightarrow \mathbb{R}_0^n \text{ are inverse functions.} \tag{87}$$

Now let $x_t \in \text{int}(X)$ and $y_t = M^\eta(F(x_t)) = M^\eta(PF(x_t))$. Then display (87) implies that $\eta \nabla_0 h(y_t) = PF(x_t)$. Since $f^\eta(x) = \eta^{-1}f(x) - h(x)$, $\nabla f(x) = F(x)$, and $\dot{x}_t = M^\eta(F(x_t)) - x_t \in \mathbb{R}_0^n$, we can compute as follows:

$$\frac{d}{dt} f^\eta(x_t) = \nabla_0 f^\eta(x)' \dot{x}_t = (\eta^{-1}PF(x_t) - \nabla_0 h(x_t))'(M^\eta(F(x_t)) - x_t) = (\nabla_0 h(y_t) - \nabla_0 h(x_t))'(y_t - x_t) \leq 0$$

strictly so whenever $M^\eta(F(x_t)) \neq x_t$, by the strict convexity of h . Since the boundary of X is repelling under (23), the proof is complete. \square

We now turn to Lemma 1 and subsequent claims. We start by stating the generalization of the Hamilton–Jacobi Equation (26). For each $R \subseteq S$ with $\#R \geq 2$, let $X_R = \{x \in X: \text{supp}(x) \subseteq R\}$, and define $f_R^\eta: X \rightarrow \mathbb{R}$ by

$$f_R^\eta(x) = \eta^{-1}f(x) - \left(\sum_{i \in R} x_i \log x_i + \sum_{j \in S \setminus R} x_j \right).$$

Evidently,

$$f_R^\eta(x) = f^\eta(x) \text{ when } \text{supp}(x) = R. \tag{88}$$

Our generalization of Equation (26) is

$$H(x, -\nabla f_R^\eta(x)) \leq 0 \text{ when } \text{supp}(x) = R, \text{ with equality if and only if } R = S. \tag{89}$$

To use (88) and (89) to establish the upper bound $c_y^* \leq -f^\eta(y)$ for paths $\phi \in \mathcal{C}_{x^*}[0, T]$, $\phi(T) = y$ that include boundary segments, define $S_t = \text{supp}(\phi_t)$. At any time t at which ϕ is differentiable, $\dot{\phi}_t$ is tangent to the face of X corresponding to S_t , and so (88) implies that $(d/dt)f^\eta(\phi_t) = \nabla f_{S_t}^\eta(\phi_t)' \dot{\phi}_t$. We therefore have

$$\begin{aligned} c_{x^*, T}(\phi) &= \int_0^T \sup_{u_t \in \mathbb{R}_0^n} (u_t' \dot{\phi}_t - H(\phi_t, u_t)) dt \geq \int_0^T (-\nabla f_{S_t}^\eta(\phi_t)' \dot{\phi}_t - H(\phi_t, -\nabla f_{S_t}^\eta(\phi_t))) dt \\ &\geq \int_0^T -\nabla f_{S_t}^\eta(\phi_t)' \dot{\phi}_t dt = f^\eta(x^*) - f^\eta(y) = -f^\eta(y), \end{aligned} \tag{90}$$

establishing the lower bound.

Derivation of Property (89). Let $x \in X$ have support $R \subseteq S$, $\#R \geq 2$. Then since $P\mathbf{1} = \mathbf{0}$,

$$\nabla_0 f_R^\eta(x) = P \left(\eta^{-1}F(x) - \sum_{i \in R} e_i(1 + \log x_i) - \sum_{j \in S \setminus R} e_j \right) = P \left(\eta^{-1}F(x) - \sum_{i \in R} e_i \log x_i \right). \tag{91}$$

Recalling definition (11) of H , letting ζ_x be a random variable with distribution $\nu(\cdot | x)$, and using the fact that $P(e_j - e_i) = e_j - e_i$, we compute as follows:

$$\begin{aligned} &\exp(H(x, -\nabla f_R^\eta(x))) \\ &= \mathbb{E} \exp(-\nabla f_R^\eta(x)' \zeta_x) \\ &= \sum_{i \in S} \sum_{j \neq i} \exp(-\nabla f_R^\eta(x)'(e_j - e_i)) \mathbb{P}(\zeta_x = e_j - e_i) + \mathbb{P}(\zeta_x = 0) \\ &= \sum_{i \in R} \sum_{j \in S \setminus \{i\}} \exp(-\eta^{-1}F_j(x) + \eta^{-1}F_i(x) + \log x_j - \log x_i) x_i \frac{\exp(\eta^{-1}F_j(x))}{\sum_{k \in S} \exp(\eta^{-1}F_k(x))} + \sum_{i \in S} x_i \frac{\exp(\eta^{-1}F_i(x))}{\sum_{k \in S} \exp(\eta^{-1}F_k(x))} \\ &= \sum_{i \in R} \frac{\exp(\eta^{-1}F_i(x))}{\sum_{k \in S} \exp(\eta^{-1}F_k(x))} (1 - x_i) + \sum_{i \in R} \frac{\exp(\eta^{-1}F_i(x))}{\sum_{k \in S} \exp(\eta^{-1}F_k(x))} x_i = \frac{\sum_{i \in R} \exp(\eta^{-1}F_i(x))}{\sum_{k \in S} \exp(\eta^{-1}F_k(x))}. \end{aligned}$$

Since the final expression equals 1 when $R = S$ and is less than 1 when $R \subset S$, property (89) follows. \square

Derivation of Equation (27). Let $x \in \text{int}(X)$, and observe that

$$\frac{\partial H}{\partial u_i}(x, u) = \frac{\sum_{j \neq i} (\exp(u_i - u_j) x_j \exp(\eta^{-1} F_i(x)) - \exp(u_j - u_i) x_i \exp(\eta^{-1} F_j(x)))}{\mathbb{E} \exp(u' \zeta_x) \sum_{k \in S} \exp(\eta^{-1} F_k(x))}. \quad (92)$$

Recall from the previous derivation that $\mathbb{E} \exp(-\nabla f^\eta(x)' \zeta_x) = 1$. Thus since $u_i - u_j = (e_i - e_j)' u = (P(e_i - e_j))' u$, it follows from (92) that $(\partial H / \partial u_i)(x, u) = (\partial H / \partial u_i)(x, Pu)$, so we can use Equation (91) with $R = S$ to compute as follows:

$$\begin{aligned} \frac{\partial H}{\partial u_i}(x, -\nabla f^\eta(x)) &= \frac{\partial H}{\partial u_i}(x, -\nabla_0 f^\eta(x)) \\ &= \frac{1}{\sum_{k \in \mathcal{A}} \exp(\eta^{-1} F_k(x))} \sum_{j \neq i} (\exp(-\eta^{-1} F_i(x) + \eta^{-1} F_j(x) + \log x_i - \log x_j) x_j \exp(\eta^{-1} F_i(x)) \\ &\quad - \exp(-\eta^{-1} F_j(x) + \eta^{-1} F_i(x) + \log x_j - \log x_i) x_i \exp(\eta^{-1} F_j(x))) \\ &= \frac{1}{\sum_{k \in \mathcal{A}} \exp(\eta^{-1} F_k(x))} \sum_{j \neq i} (x_i \exp(\eta^{-1} F_j(x)) - x_j \exp(\eta^{-1} F_i(x))) \\ &= x_i \frac{\sum_{j \neq i} \exp(\eta^{-1} F_j(x))}{\sum_{k \in \mathcal{A}} \exp(\eta^{-1} F_k(x))} - (1 - x_i) \frac{\exp(\eta^{-1} F_i(x))}{\sum_{k \in \mathcal{A}} \exp(\eta^{-1} F_k(x))} \\ &= x_i (1 - M_i^\eta(F(x))) - (1 - x_i) M_i^\eta(F(x)) = x_i - M_i^\eta(F(x)). \quad \square \end{aligned}$$

Acknowledgments

The authors thank Michel Benaïm for extensive discussions about this paper and related topics, and they thank two anonymous referees and an associate editor for helpful comments.

Endnotes

¹For imitative dynamics, see Helbing [22], Weibull [44], Björnerstedt and Weibull [10], Hofbauer [23], and Schlag [41]; for exact and perturbed best-response dynamics, see Roth and Sandholm [33] and Hofbauer and Sandholm [24], respectively. For surveys, see Sandholm [37, 38].

²Key early contributions include Foster and Young [18], Kandori et al. [27], and Young [45]; for surveys, see Young [46] and Sandholm [37].

³See Freidlin and Wentzell [19, lemma 6.3.1] or Young [45, theorem 4].

⁴Important early work featuring this assumption includes the logit model of Blume [11, 12] and the probit model of Myatt and Wallace [31]. For more recent work and references, see Sandholm [35, 36], Staudigl [42], and Sandholm and Staudigl [39].

⁵See Binmore et al. [9], Binmore and Samuelson [8], Blume [11], and Sandholm [35, 36].

⁶In more detail, each sample path of the N agent particle system specifies the action $i \in \mathcal{A}$ played by each agent as a function of time $t \in [0, T]$. Each sample path generates an empirical distribution D^N over the set of paths $\mathcal{J} = \{u: [0, T] \rightarrow \mathcal{A}\}$, where with probability 1, D^N places mass $1/N$ on N distinct paths in \mathcal{J} . The random draw of a sample path of the particle system then induces a probability distribution \mathcal{P}^N over empirical distributions D^N on the set of paths \mathcal{J} . The large deviation principle noted above concerns the behavior of the probability distributions \mathcal{P}^N as N grows large.

⁷In parlance of the particle systems literature, the first result concerns the “empirical distributions” (or “empirical measures”) of the system, whereas the latter concerns the “empirical process.”

⁸Under a continuous-time process, the number of revision opportunities received by N agents over a short but fixed time interval $[t, t + dt]$ follows a Poisson distribution with mean $N dt$. As the population size grows large, the number of arrivals per agent over this interval becomes almost deterministic. However, a large deviation principle for the evolutionary process must account for exceptional realizations of the number of arrivals. For a simple example illustrating how random arrivals influence large deviations results, see Dembo and Zeitouni [15, exercise 2.3.18].

⁹This adjustment is important in finite population models: see Sandholm [37, section 11.4.2] or Sandholm and Staudigl [39].

¹⁰Since the logit protocol is parameterized by a noise level and since clever payoff evaluation is used, this example satisfies the assumptions of our analysis of the small noise double limit in Sandholm and Staudigl [39].

¹¹More specifically, the bound on choice probabilities must hold uniformly over the payoff vectors π that may arise in the population games F^N .

¹²For further background on this material, see Dembo and Zeitouni [15] or Dupuis and Ellis [17].

¹³Here and throughout, $\mathcal{Z} = \{e_j - e_i; i, j \in \mathcal{A}\}$ and $\mathcal{Z}(x) = \{e_j - e_i; i, j \in \mathcal{A}, x_i > 0\}$ are *discrete sets of normalized increments* of the process \mathbf{X}^N , and $Z = \text{conv}(\mathcal{Z})$ and $Z(x) = \text{conv}(\mathcal{Z}(x))$ are *convex sets of feasible expected increments* of \mathbf{X}^N .

¹⁴See Dembo and Zeitouni [15], theorem 2.1.10.

¹⁵Since number of empirical distributions for sample size N only grows polynomially (it is less than $(N+1)^{|\mathcal{Z}|}$), the rate of decay of $\mathbb{P}(\mathcal{E}_x^N \in \Lambda)$ cannot be driven by a large set of distributions in Λ with higher relative entropies.

¹⁶This general idea—the preservation of large deviation principles under continuous functions—is known as the *contraction principle*. See Dembo and Zeitouni [15].

- ¹⁷See Dembo and Zeitouni [15], section 2.1.2, or Dupuis and Ellis [17], lemma 6.2.3(f).
- ¹⁸When the interpolated process \tilde{X}^N is halfway between adjacent states x and y at time $(k + 1/2)/N$, its position at time k/N determines its position at time $(k + 1)/N$.
- ¹⁹The condition requires that for all $\delta > 0$ small enough, there is a nonempty closed set $K_\delta \subset X$ disjoint from $\text{cl}(O)$ such that for all $x \in \partial O$, there exists a $y \in K_\delta$ satisfying $|x - y| = \delta$.
- ²⁰The analysis to follow only requires the limiting game F to be a potential game. In particular, there is no need for the convergent sequence $\{F^N\}$ of finite-population games to consist of potential games (as defined in Monderer and Shapley [30]) or to assume that any of the processes \tilde{X}^N are reversible (cf. Blume [12]).
- ²¹In panel (i) the size of the range of f^η is $f^{(0.25)}(x^{(0.25)}) - f^{(0.25)}(e_1) \approx -10.73 - (-20) = 9.27$, whereas in panel (ii) it is $f^{(0.1)}(x^{(0.1)}) - f^{(0.1)}(e_1) \approx -28.38 - (-50) = 21.62$.

References

- [1] Azencott R, Ruget G (1977) Mélanges d'équations différentielles et grands écarts à la loi des grands nombres. *Zeitschrift Wahrscheinlichkeitstheorie verwandte Gebiete* 38(1):1–54.
- [2] Beckmann M, McGuire CB, Winsten CB (1956) *Studies in the Economics of Transportation* (Yale University Press, New Haven, CT).
- [3] Ben-Akiva M, Lerman SR (1985) *Discrete Choice Analysis: Theory and Application to Travel Demand* (MIT Press, Cambridge, MA).
- [4] Benaïm M (1998) Recursive algorithms, urn processes, and chaining number of chain recurrent sets. *Ergodic Theory Dynam. Systems* 18(1):53–87.
- [5] Benaïm M, Weibull JW (2003) Deterministic approximation of stochastic evolution in games. *Econometrica* 71(3):873–903.
- [6] Benaïm M, Weibull JW (2009) Mean-field approximation of stochastic population processes in games. unpublished manuscript, Université de Neuchâtel, Neuchâtel, Switzerland.
- [7] Berge C (1963) *Topological Spaces* (Oliver & Boyd, Edinburgh, Scotland).
- [8] Binmore K, Samuelson L (1997) Muddling through: Noisy equilibrium selection. *J. Econom. Theory* 74(2):235–265.
- [9] Binmore K, Samuelson L, Vaughan R (1995) Musical chairs: Modeling noisy evolution. *Games Econom. Behav.* 11(1):1–35. [Erratum (1997) 21:325.]
- [10] Björnerstedt J, Weibull JW (1996) Nash equilibrium and evolution by imitation. Arrow KJ, Colombatto E, Perlman M, Schmidt C, eds. *The Rational Foundations of Economic Behavior* (St. Martin's Press, New York), 155–181.
- [11] Blume LE (1993) The statistical mechanics of strategic interaction. *Games Econom. Behav.* 5(3):387–424.
- [12] Blume LE (1997) Population games. Arthur WB, Durlauf SN, Lane DA, eds. *The Economy as an Evolving Complex System II* (Addison-Wesley, Reading, MA), 425–460.
- [13] Borkar VS, Sundaresan R (2012) Asymptotics of the invariant measure in mean field models with jumps. *Stochastic Systems* 2(2): 322–380.
- [14] de Bruijn NG (1970) *Asymptotic Methods in Analysis*, 3rd ed. (North Holland, Amsterdam).
- [15] Dembo A, Zeitouni O (1998) *Large Deviations Techniques and Applications*, 2nd ed. (Springer, New York).
- [16] Dupuis P (1988) Large deviations analysis of some recursive algorithms with state dependent noise. *Ann. Probab.* 16(4):1509–1536.
- [17] Dupuis P, Ellis RS (1997) *A Weak Convergence Approach to the Theory of Large Deviations* (John Wiley & Sons, New York).
- [18] Foster DP, Young HP (1990) Stochastic evolutionary game dynamics. *Theoret. Population Biol.* 38(2):219–232.
- [19] Freidlin MI, Wentzell AD (1998) *Random Perturbations of Dynamical Systems*, 2nd ed. (Springer, New York).
- [20] Fudenberg D, Levine DK (1998) *The Theory of Learning in Games* (MIT Press, Cambridge, MA).
- [21] Gilboa I, Matsui A (1991) Social stability and equilibrium. *Econometrica* 59(3):859–867.
- [22] Helbing D (1992) A mathematical model for behavioral changes by pair interactions. Haag G, Mueller U, Troitzsch KG, eds. *Economic Evolution and Demographic Change: Formal Models in Social Sciences* (Springer, Berlin), 330–348.
- [23] Hofbauer J (1995) Imitation dynamics for games, unpublished manuscript, University of Vienna, Vienna.
- [24] Hofbauer J, Sandholm WH (2007) Evolution in games with randomly disturbed payoffs. *J. Econom. Theory* 132(1):47–69.
- [25] Kandori M, Rob R (1995) Evolution of equilibria in the long run: A general theory and applications. *J. Econom. Theory* 65(2):383–414.
- [26] Kandori M, Rob R (1998) Bandwagon effects and long run technology choice. *Games Econom. Behav.* 22(1):84–120.
- [27] Kandori M, Mailath GJ, Rob R (1993) Learning, mutation, and long run equilibria in games. *Econometrica* 61(1):29–56.
- [28] Kurtz TG (1970) Solutions of ordinary differential equations as limits of pure jump Markov processes. *J. Appl. Probab.* 7(1):49–58.
- [29] Léonard C (1995) Large deviations for long range interacting particle systems with jumps. *Annales de l'Institut Henri Poincaré B* 31(2): 289–323.
- [30] Monderer D, Shapley LS (1996) Potential games. *Games Econom. Behav.* 14(1):124–143.
- [31] Myatt DP, Wallace CC (2003) A multinomial probit model of stochastic evolution. *J. Econom. Theory* 113(2):286–301.
- [32] Rosenthal RW (1973) A class of games possessing pure strategy Nash equilibria. *Internat. J. Game Theory* 2(1):65–67.
- [33] Roth G, Sandholm WH (2013) Stochastic approximations with constant step size and differential inclusions. *SIAM J. Control Optim.* 51(1):525–555.
- [34] Sandholm WH (2001) Potential games with continuous player sets. *J. Econom. Theory* 97(1):81–108.
- [35] Sandholm WH (2007) Simple formulas for stationary distributions and stochastically stable states. *Games Econom. Behav.* 59(1):154–162.
- [36] Sandholm WH (2010) Orders of limits for stationary distributions, stochastic dominance, and stochastic stability. *Theoret. Econom.* 5(1):1–26.
- [37] Sandholm WH (2010) *Population Games and Evolutionary Dynamics* (MIT Press, Cambridge, MA).
- [38] Sandholm WH (2015) Population games and deterministic evolutionary dynamics. Young HP, Zamir S, eds. *Handbook of Game Theory*, Vol. 4 (Elsevier, Amsterdam), 703–778.
- [39] Sandholm WH, Staudigl M (2016) Large deviations and stochastic stability in the small noise double limit. *Theoret. Econom.* 11(1): 279–355.
- [40] Sandholm WH, Staudigl M (2017) Sample path large deviations for stochastic evolutionary game dynamics. Working paper, University of Wisconsin–Madison, Madison. arXiv:1511.07897v4.

- [41] Schlag KH (1998) Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits. *J. Econom. Theory* 78(1):130–156.
- [42] Staudigl M (2012) Stochastic stability in asymmetric binary choice coordination games. *Games Econom. Behav.* 75(1):372–401.
- [43] Taylor PD, Jonker L (1978) Evolutionarily stable strategies and game dynamics. *Math. Biosci.* 40(1–2):145–156.
- [44] Weibull JW (1995) *Evolutionary Game Theory* (MIT Press, Cambridge, MA).
- [45] Young HP (1993) The evolution of conventions. *Econometrica* 61(1):57–84.
- [46] Young HP (1998) *Individual Strategy and Social Structure* (Princeton University Press, Princeton, NJ).