Hamilton-Jacobi Equations with Semilinear Costs and State Constraints, with Applications to Large Deviations in Games*

William H. Sandholm†, Hung V. Tran‡, and Srinivas Arigapudi§

July 4, 2018

Abstract

We characterize solutions of a class of time-homogeneous optimal control problems with semilinear running costs and state constraints as maximal viscosity subsolutions to Hamilton-Jacobi equations, and show that optimal solutions to these problems can be constructed explicitly. We present applications to large deviations problems arising in evolutionary game theory.

1. Introduction

This paper considers a class of time-homogeneous optimal control problems. The state variable in these problems is required to stay in a compact convex subset X of $\mathbb{R}^n$, and the cost of motion from state x in feasible direction v takes the semilinear form

$$L(x, v) = \sum_{i=1}^{n} \Psi_i(x)\,[v_i]_+$$

for some Lipschitz continuous function $\Psi\colon X \to \mathbb{R}^n$ with $\Psi = (\Psi_1, \ldots, \Psi_n)$. We consider both a source problem, which specifies a set of allowable initial states, as well as a target

*Financial support from Army Research Office Grant MSN201957 and NSF Grants DMS-1664424 and SES-1728853 is gratefully acknowledged.

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

‡Department of Mathematics, University of Wisconsin, 480 Lincoln Drive, Madison, WI 53706. e-mail: hung@math.wisc.edu; website: www.math.wisc.edu/~hung.

§Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706. e-mail: arigapudi@ssc.wisc.edu; website: sites.google.com/view/srinivasarigapudi/home.

problem, which specifies a set of allowable terminal states to be reached within an arbitrary finite time span. Problems of these sorts arise naturally in large deviations analyses of stochastic processes arising in evolutionary game theory.

Our main results characterize the solutions to these problems as maximal viscosity subsolutions of suitable Hamilton-Jacobi equations, and we provide a corresponding verification theorem. We also show that the subsolution condition needs only be checked at almost every state in X. Thus in applying our results, there is no need to work with superdifferentials or test functions, as it is sufficient to evaluate the Hamilton-Jacobi equation at states where the candidate value function is differentiable.

The class of problems we consider has a variety of interesting features. First, the semilinearity of running costs implies that paths that differ only by a reparameterization of time have the same cost. Thus in defining a feedback control, the length of the control vector plays no role; only its direction matters. Second, we are able to obtain explicit solutions to the problems we consider. As we argue below, semilinearity ensures that when maximizing the Hamiltonian at a given state, it is enough to consider controls from a suitable finite set, and that in typical problems, the state space will be divided into regions on which the optimal control is constant.

From the point of view of PDE theory, the fact that we are able to use Hamilton-Jacobi equations with state constraints to obtain explicit solutions to control problems—solutions of finite duration, along which the state constraints are active—is noteworthy. Work on Hamilton-Jacobi equations often focuses on questions related to well-posedness, and says little about the nature of optimal paths, and, in the case of state-constrained problems, about whether state constraints are active. We are unaware of previous examples of control problems with explicitly computed finite-time optimal paths along which state constraints bind.

In Sections 2 and 3, we state our main results and apply them in some simple examples. In these examples, intuitive reasoning suggests the form that the optimal solutions must take, which perhaps makes it surprising that existing results do not provide a method for verifying these solutions. Since our objective function does not include discounting, results of Soner (1986) on state-constrained control problems do not apply. Likewise, the semilinearity of the Lagrangians and linear growth of the Hamiltonians put the problems we study outside the class of state-constrained problems considered by Capuzzo-Dolcetta and Lions (1990). In our simple examples, our results allow us to verify easily that the intuitive solutions are indeed optimal. In particular, our verification theorem (Theorem 3.4) is very useful here, and is not covered in the two aforementioned papers. Section 4 presents the proofs of the main results.

We apply our results to large deviations problems arising in stochastic dynamic models from evolutionary game theory. These models concern behavior in populations of $N < \infty$ strategically interacting agents, where each agent's payoffs depend both on his own actions and on the distribution of choices of other agents. The behavior of the population is modeled as a Markov chain. Each agent occasionally receives opportunities to switch actions, and decides whether to switch actions and which action to choose by applying a simple myopic rule parameterized by a noise level $\eta > 0$. These noisy best response rules place most probability on the currently optimal action, but place positive probabilities on all actions. The stochastic processes described here are formally defined in Section 5 below.[1]

Following Sandholm and Staudigl (2016), we study the behavior of these Markov chains $\mathbf{X}^{N,\eta}$, which run on discrete grids $\mathcal{X}^N$ of mesh size $\frac{1}{N}$ in the probability simplex, in the *small noise double limit*. That is, we first characterize behavior as the noise level η in agents' decisions is taken to zero, and then describe the behavior of this limit as the population size N is taken to infinity.

When the noise level is small, so that agents nearly always choose actions that are currently optimal, the process $\mathbf{X}^{N,\eta}$ quickly moves toward a recurrent class of the zero-noise process, which often corresponds to a Nash equilibrium of the underlying game. However, the noise in agents' decisions ensures that the process will occasionally transit between Nash equilibria. Results of Freidlin and Wentzell (1984) and Catoni (1999) describe the growth rate in η of the waiting times until transitions between equilibria in terms of path-cost minimization problems on the discrete state space $\mathcal{X}^N$. These waiting times can in turn be used to determine the limiting behavior of the stationary distribution of $\mathbf{X}^{N,\eta}$, which describes the proportion of time the process spends in each state over very long time spans. Typically, differences in transition rates ensure that the process spends the majority of time at a single Nash equilibrium, the so-called *stochastically stable state*, which receives the preponderance of mass in the stationary distribution of $\mathbf{X}^{N,\eta}$ as η approaches zero.[2]

The discrete path-cost minimization problems mentioned above typically do not admit analytical solutions. To contend with this, Sandholm and Staudigl (2016) show that if one follows the limit as η approaches 0 with the limit as N approaches infinity, then the discrete path-cost minimization problems on the state spaces $\mathcal{X}^N$ converge to continuous

[1]The processes studied in evolutionary game theory are quite distinct from mean field games, in which a population of agents simultaneously solve dynamic optimization problems that each depend on the population's aggregate behavior. For background on stochastic dynamics in evolutionary game theory, see Young (1998) and Sandholm (2010b).

[2]See Foster and Young (1990), Kandori et al. (1993), and Young (1993, 1998).

path-cost minimization problems on the probability simplex X, which can in principle be solved using optimal control methods. However, existing results on Hamilton-Jacobi equations only permitted the computation of solutions to these problems in cases where the constraint that the state remains in X is not binding. The main results in this paper allow us to verify optimal value functions whether or not the state constraints bind, greatly expanding the class of environments in which large deviations and stochastic stability properties can be explicitly determined. Section 6 presents examples that illustrate the application of our main results.

Section 7 presents some concluding remarks, and the appendices contain certain proofs and computations omitted from the main text.

2. Optimal control problems with semilinear running costs

2.1 Definitions

Let $X \subset \mathbb{R}^n$ be an m-dimensional polytope. Let TX denote the set of tangent vectors from states in the relative interior X° of X. If X is n-dimensional, then $TX = \mathbb{R}^n$; more generally, if the affine hull of X is a translation of subspace $Y \subset \mathbb{R}^n$, then $TX = Y$. In what follows, topological statements are made with respect to the relative topology on X; for instance, we will refer to X° as the interior of X. Likewise, derivatives of functions $f\colon X \to \mathbb{R}^m$ will be understood as maps $Df\colon X \to \mathcal{L}(TX, \mathbb{R}^m)$ from X to linear functions from TX to $\mathbb{R}^m$.

Let $TX(x) = \{t(y - x)\colon y \in X, t \geq 0\} \subseteq TX$ be the tangent cone of X at x. $TX(x)$ is the set of feasible controls at state x, though as we note shortly, it will be enough to restrict attention to controls in a compact subset of $TX(x)$.

We assume that the running cost function $L\colon X \times TX \to [0, +\infty]$ takes a semilinear form. Specifically, we assume that for some Lipschitz continuous function $\Psi\colon X \to \mathbb{R}^n_+$, we have

$$(2.1) \qquad L(x, v) = \begin{cases} \displaystyle\sum_{i=1}^{n} \Psi_i(x)\,[v_i]_+ & \text{if } v \in TX(x), \\ +\infty & \text{otherwise.} \end{cases}$$

Thus the constraint that the state remains in X is built into the definition of running costs.

Let Φ_T be the set of Lipschitz continuous paths $\phi\colon [0, T] \to X$, and let $\Phi = \bigcup_{T\geq 0} \Phi_T$. Then the *path cost function* $c\colon \Phi \to \mathbb{R}_+$ is defined by

$$c(\phi) = \int_0^T L(\phi(t), \dot{\phi}(t))\,\mathrm{d}t \quad \text{when } \phi \in \Phi_T. \tag{2.2}$$

The *source problem* for the compact set $X_0 \subset X$ is that of finding the minimal cost of reaching each state in X from a free initial condition in X_0. The value function for the source problem for X_0 is

$$V(x) = \inf\{c(\phi)\colon \phi \in \Phi_T \text{ for some } T \geq 0, \phi(0) \in X_0, \phi(T) = x\}. \tag{SP}$$

Likewise, the *target problem* for the compact set $Y \subset X$ is that of finding the minimal cost of reaching a free state in Y from initial condition x. The value function for the target problem for Y is

$$W(x) = \inf\{c(\phi)\colon \phi \in \Phi_T \text{ for some } T \geq 0, \phi(0) = x, \phi(T) \in Y\}. \tag{TP}$$

Remark 2.1. It follows immediately from the semilinearity of the running cost function (2.1) that the cost of a path does not depend on the speed at which it is traversed: if $\phi \in \Phi_T$ and $\hat{\phi} \in \Phi_{\hat{T}}$ differ only by a reparameterization of time, then $c(\phi) = c(\hat{\phi})$. Because of this, the solutions to problems (SP) and (TP) do not change if we restrict the control variable u to a compact convex set whose conical hull is TX. Viewed through the prism of feedback controls, semilinearity implies that we need only determine the optimal *directions* of motion from each state; the speed of motion in an optimal direction is irrelevant.

We will take advantage of this property by introducing a convenient restriction on the control variable. Let $|\cdot|$ denote the ℓ_1 norm on $\mathbb{R}^n$, so that $|u| = \sum_{i=1}^n |u_i|$. For $r > 0$, let $B_r = \{u \in TX\colon |u| \leq r\}$ be the closed radius-r ball in $\mathbb{R}^n$. The foregoing discussion shows that in solving (SP) and (TP), there is no loss in restricting attention to paths $\phi \in \Phi$ with $\dot{\phi}(t) \in B_r$ for almost all t. This restriction is captured in the reformulation (3.1) of the running cost function in Section 3.

2.2 Motivating examples

As motivation, we introduce two simple examples of the target problem (TP).

Example 2.2. Consider the path-cost minimization problem in $\mathbb{R}^2$ defined by

$$X = \{x \in [0,1]^2\colon x_1 + x_2 \leq \tfrac{3}{2}\},$$
$$Y = \{x \in [0,1]^2\colon x_1 + x_2 = \tfrac{3}{2}\},$$

$$L(x,v) = \begin{cases} (\frac{3}{2} - x_1 - x_2)\,([v_1]_+ + 2[v_2]_+) & \text{if } v \in TX(x), \\ +\infty & \text{otherwise.} \end{cases}$$

Under running cost function L, costs of motion are proportional to $\frac{3}{2} - x_1 - x_2$, and so decrease linearly as the target set Y is approached orthogonally. In addition, the cost of increasing x_2 is always twice as large as the cost of increasing x_1.

The solution to this problem is intuitively clear. Since motion is most costly at states far from the target set, one should increase x_1 first with x_2 held fixed to take advantage of the low relative cost of doing so. If the initial condition satisfies $x_2 \geq \frac{1}{2}$, then x_1 can be increased until Y is reached. If not, then the state constraint will bind once $x_1 = 1$, so one should then increase x_2 until Y is reached.

The value function $W\colon X \to \mathbb{R}_+$ corresponding to this feedback control is easily computed as

$$\text{(2.3)} \qquad W(x) = \begin{cases} \int_{x_1}^{\frac{3}{2}-x_2} (\frac{3}{2} - t - x_2)\,\mathrm{d}t & \text{if } x_2 \geq \frac{1}{2}, \\ \int_{x_1}^{1} (\frac{3}{2} - t - x_2)\,\mathrm{d}t + \int_{x_2}^{\frac{1}{2}} (1 - 2u)\,\mathrm{d}u & \text{if } x_2 < \frac{1}{2}, \end{cases}$$
$$= \begin{cases} \frac{1}{8}(2x_1 + 2x_2 - 3)^2 & \text{if } x_2 \geq \frac{1}{2}, \\ \frac{1}{2}x_1^2 + x_2^2 + x_1x_2 - \frac{3}{2}x_1 - 2x_2 + \frac{5}{4} & \text{if } x_2 < \frac{1}{2}. \end{cases}$$

Thus the cost of reaching Y from the origin is $\frac{5}{4}$. The derivative of W is

$$\text{(2.4)} \qquad DW(x) = \begin{cases} (x_1 + x_2 - \frac{3}{2}, x_1 + x_2 - \frac{3}{2}) & \text{if } x_2 \geq \frac{1}{2}, \\ (x_1 + x_2 - \frac{3}{2}, x_1 + 2x_2 - 2) & \text{if } x_2 < \frac{1}{2}. \end{cases}$$

Inspection of (2.4) reveals that DW is continuous along the line segment $\{x \in X\colon x_2 = \frac{1}{2}\}$, and hence W is continuously differentiable.

While the nature of the solution is transparent, and the value function is pleasingly smooth, the lack of smoothness of the running costs and the presence of the state constraint prevent optimality from being verified by textbook methods. ◆

Example 2.3. Consider the following control problem in $\mathbb{R}^3$:

$$X = \{x \in [0,1]^3\colon x_1 + x_2 + x_3 \leq \tfrac{5}{2}\},$$
$$Y = \{x \in [0,1]^3\colon x_1 + x_2 + x_3 = \tfrac{5}{2}\},$$

$$L(x,u) = \begin{cases} (\frac{5}{2} - x_1 - x_2 - x_3)\,([v_1]_+ + 2[v_2]_+ + 3[v_3]_+) & \text{if } v \in TX(x), \\ +\infty & \text{otherwise.} \end{cases}$$

The optimal feedback control is similar to that in Example 2.2: First increase x_1 until Y is reached or $x_1 = 1$; in the latter event, increase x_2 until Y is reached or $x_2 = 1$; and in this latter event, increase x_3 until Y is reached. It is straightforward to compute the corresponding value function:

$$W(x) = \begin{cases} \int_{x_1}^{\frac{5}{2}-x_2-x_3} (\frac{5}{2} - t - x_2 - x_3)\,\mathrm{d}t & \text{if } x_2 + x_3 \geq \frac{3}{2}, \\ \int_{x_1}^{1} (\frac{5}{2} - t - x_2 - x_3)\,\mathrm{d}t + \int_{x_2}^{\frac{3}{2}-x_3} (3 - 2u - 2x_3)\,\mathrm{d}u & \text{if } x_2 + x_3 < \frac{3}{2},\ x_3 \geq \frac{1}{2}, \\ \int_{x_1}^{1} (\frac{5}{2} - t - x_2 - x_3)\,\mathrm{d}t + \int_{x_2}^{1} (3 - 2u - 2x_3)\,\mathrm{d}u + \int_{x_3}^{\frac{1}{2}} (\frac{3}{2} - 3v)\,\mathrm{d}v & \text{if } x_2 + x_3 < \frac{3}{2},\ x_3 < \frac{1}{2}. \end{cases}$$

One can verify that this function too is continuously differentiable. But as above, optimality cannot be verified by textbook methods. The nature of the optimal solution is somewhat more complicated than in Example 2.2, since choice of the optimal control on the $x_1 = 1$ boundary of X is nontrivial. ◆

3. Main results

To take advantage of Remark 2.1, we replace the running cost function (2.1) with one in which controls outside of B_r are infeasible:

$$(3.1) \qquad L(x,v) = \begin{cases} \sum_{i=1}^{n} \Psi_i(x)\,[v_i]_+ & \text{if } v \in TX(x) \cap B_r, \\ +\infty & \text{otherwise.} \end{cases}$$

For $v \in \mathbb{R}^n$, we define the componentwise positive part function $[v]_+$ by $([v_+])_i = [v_i]_+$, and we define $[v]_-$ analogously. Using this notation, we can write the first case of (3.1) concisely as

$$(3.2) \qquad L(x,v) = \langle \Psi(x), [v]_+ \rangle \quad \text{if } v \in TX(x) \cap B_r.$$

Let $\mathcal{H}\colon X \times \mathbb{R}^m \times TX \to \mathbb{R}$ denote the Hamiltonian corresponding to running cost function (3.1):

$$\mathcal{H}(x,u,v) = \langle u, v \rangle - L(x,v)$$

$$= \langle u, v\rangle - \langle \Psi(x), [v]_+\rangle$$

$$= \langle u - \Psi(x), [v]_+\rangle - \langle u, [v]_-\rangle\,. \tag{3.3}$$

Let $H\colon X \times \mathbb{R}^m \to \mathbb{R}$ denote the maximized Hamiltonian:

$$H(x, u) = \max_{v \in TX \cap B_r} \mathcal{H}(x, u, v). \tag{3.4}$$

As usual, $H(x, \cdot)$ is the convex conjugate of $L(x, \cdot)$.

Let $\|\Psi\|_\infty = \max_{x\in X} |\Psi(x)|$ denote the L^∞ norm of Ψ. Then it is easy to verify that $H(x, \cdot)$ satisfies the following linear lower bound:

$$H(x, u) \geq \frac{r}{n}|u| - r\|\Psi\|_\infty. \tag{3.5}$$

By contrast, Capuzzo-Dolcetta and Lions (1990, p. 678) consider state-constrained Hamilton-Jacobi equations in which $H(x, \cdot)$ satisfies a superlinear lower bound.

Remark 3.1. Because $L(x, \cdot)$ is linear on each orthant of $\mathbb{R}^n$, the maximization problem in (3.4) only requires us to consider controls from a well-chosen finite set. Let $\mathcal{K}^n$ be the collection of closed orthants of $\mathbb{R}^n$, and define

$$\mathcal{E} = \bigcup_{K \in \mathcal{K}^n} \operatorname{ext}(K \cap TX \cap B_r). \tag{3.6}$$

The linearity of $L(x, \cdot)$ on each orthant implies that for $x \in X^\circ$, it is enough to perform the maximization in (3.4) over the set of extreme points $\mathcal{E}$.

For instance, if X has full dimension, so that $TX = \mathbb{R}^n$, we choose $r = 1$ and obtain

$$\mathcal{E} = \{e_1, -e_1, \ldots, e_n, -e_n, \mathbf{0}\} \tag{3.7}$$

where e_i is the ith standard basis vector. In our applications to evolutionary game theory, X will be the simplex, so that $TX = \mathbb{R}^n_0 = \{v \in \mathbb{R}^n\colon \sum_i z_i = 0\}$. In this case we choose $r = 2$ and obtain

$$\mathcal{E} = \{e_j - e_i\colon i, j \in \{1, \ldots, n\}\} = \{e_j - e_i\colon i \neq j\} \cup \{\mathbf{0}\}. \tag{3.8}$$

To state the main results, we say that a function $V\colon X \to \mathbb{R}$ is *maximal* with respect to given properties if for any other function $V_0\colon X \to \mathbb{R}$ satisfying the properties, we have $V_0 \leq V$ on X.

Theorems 3.2 and 3.3 characterize the solutions to the source and target problems in terms of subsolutions in the almost everywhere sense. Thus the subsolution inequalities need not be checked at boundary states, or at states where the candidate function is not differentiable.

Theorem 3.2. *The solution to the source problem* (SP) *is the maximal Lipschitz continuous function $V\colon X \to \mathbb{R}$ satisfying*

$$\begin{cases} H(x, DV(x)) \le 0 & \text{for almost all } x \in X^\circ; \\ V(x^*) = 0 & \text{for all } x^* \in X_0. \end{cases} \tag{3.9}$$

Theorem 3.3. *The solution to the target problem* (TP) *is the maximal Lipschitz continuous function $W\colon X \to \mathbb{R}$ satisfying*

$$\begin{cases} H(x, -DW(x)) \le 0 & \text{for almost all } x \in X^\circ; \\ W(y^*) = 0 & \text{for all } y^* \in Y. \end{cases} \tag{3.10}$$

Theorems 3.2 and 3.3 lead directly to verification theorems for the source and target problems, which replace the requirement of maximality with the requirement of attainability of the values of V and W. We only state the theorem for the target problem; the statement for the source problem is similar.

Theorem 3.4. *Suppose that $W\colon X \to \mathbb{R}$ is Lipschitz continuous with $W(y^*) = 0$ for all $y^* \in Y$. In addition, suppose that for each $x \in X \setminus Y$, there is a time $T > 0$ and a path $\phi \in \Phi_T$ with $\phi(0) = x$ and $\phi(T) \in Y$ such that $W(x) = c(\phi)$. If W satisfies* (3.10), *then W is the solution to* (TP).

In general environments with Hamilton-Jacobi PDE, it is not known whether for two given points x and y, there exists an optimal finite-time path $\phi \in \Phi_T$ connecting them, even in cases without state constraints. It is thus appealing that Theorem 3.4 can be used to constructively establish the existence of optimal finite-time paths, as the examples and applications illustrate.

Before presenting the proofs of Theorems 3.2–3.4, we return to the minimization problems from Section 2.2.

Example 3.5. Example 2.2 introduced the problem of reaching target set $Y = \{x \in [0,1]^2 \colon x_1 + x_2 = \frac{3}{2}\}$ from states in $X = \{x \in [0,1]^2 \colon x_1 + x_2 \le \frac{3}{2}\}$ when running costs in feasible directions are given by

$$\begin{aligned} L(x, v) &= \Psi_1(x)\,[v_1]_+ + \Psi_2(x)\,[v_2]_+ \\ &\equiv (\tfrac{3}{2} - x_1 - x_2)\,[v_1]_+ + (3 - 2x_1 - 2x_2)\,[v_2]_+. \end{aligned} \tag{3.11}$$

The value function W proposed in (2.3) was computed by evaluating the costs of the proposed optimal paths to the target set—first increase x_1, then increase x_2 if necessary. By Theorem 3.4, verifying that this function is the optimal value function only requires us to check that W satisfies the subsolution condition in (3.10).

To do this, we use (3.3) to write the Hamiltonian at $u = -DW(x)$ as

$$\mathcal{H}(x, -DW(x), v) = \langle -DW(x) - \Psi(x), [v]_+ \rangle - \langle -DW(x), [v]_- \rangle . \tag{3.12}$$

By Remark 3.1, it is enough to show that for almost all $x \in X^\circ$,

$$\max_{v \in \mathcal{E}} \mathcal{H}(x, -DW(x), v) \leq 0, \quad \text{where } \mathcal{E} = \{e_1, -e_1, e_2, -e_2, \mathbf{0}\}. \tag{3.13}$$

So fix $x \in X^\circ$. Clearly $\mathcal{H}(x, -DW(x), \mathbf{0}) = 0$. Expression (2.4) shows that $DW(\hat{x}) < 0$ for $\hat{x} \in X \setminus Y^*$ (in terms of the partial order on $\mathbb{R}^2$), which with (3.12) implies that $\mathcal{H}(x, -DW(x), -e_i) < 0$ for $i \in \{1, 2\}$. To check the remaining values of v in $\mathcal{E}$, we use (2.4), (3.11), and (3.12) to write

$$\mathcal{H}(x, -DW(x), e_1) = -\frac{\partial W}{\partial x_1}(x) - \Psi_1(x) = 0,$$
$$\mathcal{H}(x, -DW(x), e_2) = -\frac{\partial W}{\partial x_2}(x) - \Psi_2(x) = \left\{ \begin{matrix} x_1 + x_2 - \frac{3}{2} & \text{if } x_2 \geq \frac{1}{2} \\ x_1 - 1 & \text{if } x_2 < \frac{1}{2} \end{matrix} \right\} \leq 0.$$

Thus (3.13) holds at all $x \in X^\circ$, implying that the function W defined in (2.3) is indeed the optimal value function.

The maximizers in (3.13) indicate the form of the optimal feedback control. Only controls v in $\{e_1, \mathbf{0}\}$ achieve the maximum in (3.13) on X°, implying that the optimal control at interior states requires motion in direction e_1 (at an arbitrary speed). When $x_1 = 1$, controls in $\{e_1, e_2, \mathbf{0}\}$ achieve the maximum in (3.13), but choosing $v = e_1$ would violate the state constraint, suggesting that motion in direction e_2 is optimal on this boundary of X. This is the optimal feedback control we originally proposed. ◆

Example 3.6. Similar calculations show that the value function proposed in Example 2.3 is the optimal value function. Likewise, determining the controls in $\mathcal{E}$ that maximize the Hamiltonian and accounting for state constraints identifies the optimal feedback control.

It is noteworthy that condition (3.10) only requires the subsolution condition to be checked in the interior of X, even though the form of the optimal control on the boundary of X is nontrivial. Of course, the definition of the value function in X° incorporates the definition of the control on $\text{bd}(X)$, so that suboptimal choices of the control vector on the boundary would lead to failures of the subsolution condition near the boundary. ◆

4. Analysis

To prove Theorem 3.2, we first show in Section 4.1 that the solution V to the state-constrained source problem (SP) is the maximal viscosity solution to the associated Hamilton-Jacobi equation (4.1)—that is, V is both a viscosity subsolution on X° and a viscosity supersolution on X—and that it is maximal with respect to this pair of properties. The fact that H is nonnegative implies that the supersolution condition is vacuous, and we show that V is a subsolution using a direct argument based on the dynamic programming principle. The coercivity of $H(x, \cdot)$ implies that subsolutions are Lipschitz continuous (see equation (4.7)). Taking advantage of this fact and the convexity of $H(x, \cdot)$, we use a convolution argument, followed by continuity arguments that link the costs of interior and general paths, to prove that V is the maximal viscosity solution.

To complete the proof of Theorem 3.2, we use another convolution argument and the stability of viscosity subsolutions to conclude that the subsolution condition need only be checked on a full measure subset of X°, and in particular only at states where the candidate value function is differentiable (Section 4.2). Theorem 3.3 follows from Theorem 3.2 and a simple manipulation that relates the source and target problems. Then Theorem 3.4 follows directly from Theorem 3.3 and the definition (TP) of the target problem (Section 4.3).

The state-constrained system of interest to us here is

$$\begin{cases} H(x, DV(x)) \leq 0 & \text{in } X^\circ; \\ H(x, DV(x)) \geq 0 & \text{on } X; \\ V(x^*) = 0 & \text{for all } x^* \in X_0. \end{cases} \tag{4.1}$$

Our analysis first considers the viscosity solutions of this system.

Definition 1. *Assume that $V \in C(X)$ and $V(x^*) = 0$ for all $x^* \in X_0$.*

Firstly, we say that V is a viscosity subsolution of (4.1) *if for any test function $\varphi \in C^\infty(X)$ such that $V - \varphi$ has a maximum at $x_0 \in X^\circ \setminus X_0$, then $H(x_0, D\varphi(x_0)) \leq 0$.*

Secondly, we say that V is a viscosity supersolution of (4.1) *if for any test function $\varphi \in C^\infty(X)$ such that $V - \varphi$ has a minimum at $x_0 \in X \setminus X_0$, then $H(x_0, D\varphi(x_0)) \geq 0$.*

Finally, $V \in C(X)$ is a viscosity solution of (4.1) *if it is both a viscosity subsolution and supersolution of* (4.1).

When considering optimal control problems with discounting, one obtains *monotone* Hamilton-Jacobi equations, in which the left-hand sides of the inequalities in (4.1) are replaced by $H(x, DV(x)) + \lambda V(x)$ for some $\lambda > 0$. Monotone Hamilton-Jacobi equations with state constraints, first studied by Soner (1986), commonly possess unique viscosity solutions. In contrast, the system (4.1) admits multiple viscosity solutions; for example, $V \equiv 0$ is always a viscosity solution of (4.1). This naturally leads us to consider the notion of a maximal solution:

Definition 2. *Let $V \in C(X)$ be a viscosity solution of* (4.1)*. We say that V is a maximal viscosity solution if for any viscosity solution $\hat{V}$ of* (4.1)*, $V \geq \hat{V}$ on X.*

See Capuzzo-Dolcetta and Lions (1990), Mitake (2008), Armstrong and Tran (2015) and the references therein for the use of maximal solutions in other PDE contexts.

We define the minimum cost of traveling from x_0 to x by

$$C(x_0, x) = \inf\left\{c(\phi) \,:\, \phi \in \Phi_T \text{ for some } T \geq 0, \phi(0) = x_0, \phi(T) = x\right\}.$$

Then V in (SP) can be written as

$$V(x) = \inf\{C(x_0, x) \colon\ x_0 \in X_0\}.$$

Let $B_r(x)$ denote the closed unit ball in X with center x. Then the previous definitions immediately yield:

Observation 4.1 (The dynamic programming principle)**.**
Let V be defined as in (SP)*. Fix $x \in X$. For $r > 0$, the following holds*

$$V(x) = \min_{y \in B_r(x) \cap X} \left(V(y) + C(y, x)\right). \tag{4.2}$$

4.1 Proof that optimal solutions are maximal viscosity solutions of (4.1)

Our first preliminary result characterizes solutions to the source problem (SP) in terms of system (4.1).

Theorem 4.2. *The solution to the source problem* (SP) *is the unique maximal solution of* (4.1)*.*

The proof of Theorem 4.2 is divided into two steps. We first show that V is a viscosity solution to (4.1). To do this, we use the dynamic programming principle to verify that V is a viscosity subsolution of (4.1). Since H is always nonnegative, we do not need to check the viscosity supersolution property.

Proposition 4.3. *The function V defined in* (SP) *is a viscosity subsolution of* (4.1)*.*

Proof. Take $\varphi \in C^\infty(X)$ and assume that $V - \varphi$ has a strict maximum at $x \in X^\circ \setminus X_0$. By adding a constant to φ, we can assume further that $V(x) = \varphi(x)$ and that $V(y) < \varphi(y)$ for $y \neq x$. As $x \in X^\circ$, for $r > 0$ sufficiently small, $B_r(x) \subset X$. By (4.2), for $y \in B_r(x)$,

$$\varphi(x) = V(x) \leq V(y) + C(y, x) \leq \varphi(y) + C(y, x). \tag{4.3}$$

Now pick an arbitrary displacement vector $v \in TX(x) \cap B_r$. Without loss of generality, we can assume that $r < 1$. Set $y = x - rv$, and define the path $\phi : [0, r] \to X$ by $\phi(t) = y + tv$. Then $\phi \in \Phi$ and

$$C(y, x) \leq \int_0^r \left\langle \Psi(\phi(t)), [\dot{\phi}(t)]_+ \right\rangle dt. \tag{4.4}$$

Combining (4.3) and (4.4) yields

$$\varphi(x) - \varphi(y) = \varphi(\phi(r)) - \varphi(\phi(0)) \leq \int_0^r \left\langle \Psi(\phi(t)), [\dot{\phi}(t)]_+ \right\rangle dt. \tag{4.5}$$

Divide both sides by r and let $r \to 0$ to deduce further that

$$\langle D\varphi(x), v\rangle - \langle \Psi(x), [v]_+\rangle \leq 0.$$

(In defining $\|D\varphi\|_\infty$, we identify $D\varphi(x) \in \mathcal{L}(TX, \mathbb{R}^m)$ with the unique $w \in TX$ satisfying $DV(x)z = \langle w, z\rangle$ for all $z \in TX$.) Since the direction v can be chosen freely in $TX(x) \cap B_r$, and since $L(x, v) = +\infty$ for other choices of v, we conclude that

$$\max_{v \in TX \cap B_r} (\langle D\varphi(x), v\rangle - L(x, v)) \leq 0. \tag{4.6}$$

Thus $H(x, D\varphi(x)) \leq 0$. □

We next check the maximal property.

Proposition 4.4. *The function V defined in* (SP) *is the maximal solution of* (4.1)*.*

Proof. Take an arbitrary viscosity solution $\hat{V}$ of (4.1). In light of the linear lower bound (3.5) on $H(x, \cdot)$, we claim that $\hat{V}$ is Lipschitz continuous with

$$\|D\hat{V}\|_\infty \leq n\|\Psi\|_\infty. \tag{4.7}$$

See Appendix A for a proof of this fact. By Rademacher's theorem, $\hat{V}$ is differentiable almost everywhere in X°. At each point $x \in X^\circ$ where $\hat{V}$ is differentiable, we have (see Crandall et al. (1984); Le et al. (2017))

$$H(x, D\hat{V}(x)) \leq 0. \tag{4.8}$$

Fix $x, y \in X^\circ$. Let $\phi : [0, T] \to X^\circ$ be a Lipschitz continuous path satisfying $\phi \in \Phi$, $\phi(0) = y$, and $\phi(T) = x$. We claim that

$$\hat{V}(x) \leq \hat{V}(y) + \int_0^T L(\phi(t), \dot{\phi}(t))\, dt. \tag{4.9}$$

To prove (4.9), we construct smoothed versions of $\hat{V}$ by convolution with standard mollifiers (see Le et al. (2017, Proposition 4.10)). Pick $r > 0$ sufficiently small such that $B_r(\phi(t)) \subset X^\circ$ for all $t \in [0, T]$, and define

$$X^r = \{x \in X \,:\, B_r(x) \subset X^\circ\}.$$

Then $\phi([0, T]) \subset X^r \subset X^\circ$. Let $\rho \in C_c^\infty(\mathbb{R}^m, [0, \infty))$ be a standard convolution kernel, that is, ρ is smooth and nonnegative, (compactly) supported in B_1, and satisfies $\int_{\mathbb{R}^m} \rho(x)\, dx = 1$. For $\varepsilon \in (0, r)$, define $\rho^\varepsilon(x) = \varepsilon^{-m}\rho(\varepsilon^{-1}x)$, so that

$$\int_{\mathbb{R}^m} \rho^\varepsilon(x)\, dx = 1 \quad \text{and} \quad \operatorname{supp}(\rho^\varepsilon) \subset B_\varepsilon.$$

Now define the smooth function $V^\varepsilon \colon X^r \to \mathbb{R}_+$ by

$$V^\varepsilon(x) = \rho^\varepsilon * \hat{V}(x) = \int_{B_\varepsilon(x)} \rho^\varepsilon(x - y)\hat{V}(y)\, dy = \int_{B_\varepsilon} \rho^\varepsilon(y)\hat{V}(x - y)\, dy.$$

For $x \in X^r$, use (4.8), Jensen's inequality, and the Lipschitz continuity of Ψ to argue as follows:

$$\begin{aligned} H(x, DV^\varepsilon(x)) &= H\left(x, \int_{B_\varepsilon} \rho^\varepsilon(y) D\hat{V}(x - y)\, dy\right) \\ &\leq \int_{B_\varepsilon} \rho^\varepsilon(y) H(x, D\hat{V}(x - y))\, dy \\ &\leq \int_{B_\varepsilon} \rho^\varepsilon(y) \Big(H(x - y, D\hat{V}(x - y)) + K|y|\Big)\, dy \end{aligned}$$

$$
\begin{aligned}
&\leq \int_{B_\varepsilon} \rho^\varepsilon(y) H(x-y, D\hat{V}(x-y))\, dy + K\varepsilon \\
&\leq K\varepsilon.
\end{aligned}
$$

Thus V^ε is smooth on X^r and satisfies the following inequality in the classical sense:

$$
H(x, DV^\varepsilon(x)) \leq K\varepsilon \quad \text{for } x \in X^r. \tag{4.10}
$$

In particular, for almost all $t \in [0, T]$,

$$
\left\langle DV^\varepsilon(\phi(t)), \dot{\phi}(t) \right\rangle - L(\phi(t), \dot{\phi}(t)) \leq H(x, DV^\varepsilon(x)) \leq K\varepsilon.
$$

Integrating over $t \in [0, T]$ yields

$$
V^\varepsilon(x) = V^\varepsilon(\phi(T)) \leq V^\varepsilon(\phi(0)) + \int_0^T L(\phi(t), \dot{\phi}(t))\, dt + KT\varepsilon.
$$

Then taking $\varepsilon \to 0$ gives us (4.9).

To proceed, take the infimum of (4.9) over all paths $\phi \in \Phi$ to obtain

$$
\hat{V}(x) \leq \hat{V}(y) + \inf\left\{c(\phi) \colon \phi \in \Phi_T \text{ for some } T \geq 0, \phi(0) = y, \phi(T) = x, \phi([0, T]) \subset X^\circ\right\}.
$$

Although the infimum above is only taken over all paths ϕ staying in X°, it is still equal to the infimum taken over all paths ϕ staying on X because of the continuity of the path cost function (2.2). Specifically, for each $\phi \in \Phi_T$ with $\phi(0) = y, \phi(T) = x$, one can choose a sequence of paths $\{\phi_k\} \subset \Phi_T$ such that $\phi_k([0, T]) \subset X^\circ$ for all $k \in \mathbb{N}$ and

$$
\lim_{k \to \infty} \left(\|\phi_k - \phi\|_\infty + \|\dot{\phi}_k - \dot{\phi}\|_1 \right) = 0,
$$

where $\|\cdot\|_1$ denotes the L^1 norm; then

$$
\begin{aligned}
|c(\phi_k) - c(\phi)| &= \left| \int_0^T \left(\left\langle \Psi(\phi_k(t)), [\dot{\phi}_k(t)]_+ \right\rangle - \left\langle \Psi(\phi(t), [\dot{\phi}(t)]_+ \right\rangle \right) dt \right| \\
&\leq \int_0^T \left(|\Psi(\phi_k(t)) - \Psi(\phi(t))| \cdot |[\dot{\phi}(t)]_+| + |\Psi(\phi_k(t))| \cdot |[\dot{\phi}_k(t)]_+ - [\dot{\phi}(t)]_+| \right) dt \to 0
\end{aligned}
$$

as $k \to \infty$. This approximation implies that

$$
\hat{V}(x) \leq \hat{V}(y) + C(y, x). \tag{4.11}
$$

Finally, let y tend to X_0 and use the continuity of the path cost function again to conclude that

$$\hat{V}(x) \leq \inf\{C(y,x)\colon\ y \in X_0\} = V(x). \qquad \square$$

Remark 4.5. We note that the Lipschitz continuity assumption of Ψ on X can be relaxed to a continuity assumption of Ψ on X. If Ψ is continuous on X, then H is continuous on $X \times \mathbb{R}^m$. For $\varepsilon > 0$, let ω denote the modulus of continuity of H on $X \times B_{n\|\Psi\|_\infty}$: that is, for $\varepsilon > 0$, define

$$\omega(\varepsilon) = \max\left\{|H(x,u) - H(y,u)|\colon x, y \in X, |u| \leq n\|\Psi\|_\infty, |x-y| \leq \varepsilon\right\}.$$

It is clear that $\lim_{\varepsilon\to 0+}\omega(\varepsilon) = 0$. Then, in the above proof, we replace the step using Lipschitz continuity by

$$\begin{aligned}\int_{B_\varepsilon} \rho^\varepsilon(y)H(x, DV'(x-y))\,dy &\leq \int_{B_\varepsilon} \rho^\varepsilon(y)\left(H(x-y, DV'(x-y)) + \omega(\varepsilon)\right)\,dy \\ &\leq \int_{B_\varepsilon} \rho^\varepsilon(y)H(x-y, DV'(x-y))\,dy + \omega(\varepsilon) \\ &\leq \omega(\varepsilon).\end{aligned}$$

4.2 Proof of Theorem 3.2

We are now ready to prove Theorem 3.2.

Proof of Theorem 3.2. Let V be the solution to the source problem (SP). By Theorem 4.2, V is the maximal solution to (4.1). By (3.5), V is Lipschitz continuous with $\|DV\|_\infty \leq n\|\Psi\|_\infty$. In particular, V is differentiable almost everywhere in X° and satisfies (3.9).

We therefore just need to show that V is the maximal solution to (3.9). Take an arbitrary Lipschitz solution $\hat{V}$ of (3.9). Repeating the proof of Proposition 4.4, for $r > 0$, $\varepsilon \in (0, r)$ and $V^\varepsilon = \rho^\varepsilon * \hat{V}$, we have that V^ε is smooth on X^r and V^ε satisfies in the classical sense that

$$H(x, DV^\varepsilon(x)) \leq K\varepsilon \quad \text{for } x \in X^r.$$

As V^ε converges uniformly to $\hat{V}$ in X^r, it follows from the stability of viscosity subsolutions (Crandall et al. (1984)), that $\hat{V}$ is a viscosity subsolution to

$$H(x, D\hat{V}(x)) \leq 0 \quad \text{on } X^r.$$

Since $X^\circ = \bigcup_{r>0} X^r$, $\hat{V}$ is a viscosity subsolution to (4.1). Thus since V is a maximal solution to (4.1), we conclude that $V \geq \hat{V}$. □

4.3 Proofs of Theorems 3.3 and 3.4

Next, we deduce Theorem 3.3 from Theorem 3.2.

Proof of Theorem 3.3. For any $\phi \in \Phi_T$ with $T \geq 0$, define $\tilde{\phi} \in \Phi_T$ by $\tilde{\phi}(t) = \phi(T-t)$ for all $t \in [0,T]$. Define

$$\tilde{L}(x,v) = \begin{cases} \displaystyle\sum_{i=1}^{n} \Psi_i(x)\,[v_i]_- & \text{if } v \in TX \cap B_r, \\ +\infty & \text{otherwise.} \end{cases}$$

Then $\tilde{L}(x,v) = L(x,-v)$ if $x \in X^\circ$, or if $x \in \mathrm{bd}(X)$ and $v \in TX(x) \cap (-TX(x))$. Since ϕ is Lipschitz continuous it is differentiable almost everywhere, and if $t \in [0,T]$ is such that $\phi(t) \in \mathrm{bd}(X)$ and $\dot{\phi}(t)$ exists, then $\dot{\phi}(t) \in TX(x) \cap (-TX(x))$. Since this last set is a subspace, $\tilde{\phi}$ also satisfies these properties. Thus if we define

$$\tilde{c}(\tilde{\phi}) = \int_0^T \tilde{L}(\tilde{\phi}(t), \dot{\tilde{\phi}}(t))\,dt,$$

then $\tilde{c}(\tilde{\phi}) = c(\phi)$. This fact and the definition (TP) of $W(x)$ imply that

$$\begin{aligned} W(x) &= \inf\{c(\phi)\colon \phi \in \Phi_T \text{ for some } T \geq 0, \phi(0) = x, \phi(T) \in Y\} \\ &= \inf\{\tilde{c}(\tilde{\phi})\colon \tilde{\phi} \in \Phi_T \text{ for some } T \geq 0, \tilde{\phi}(0) \in Y, \tilde{\phi}(T) = x\} = \tilde{V}(x), \end{aligned} \tag{4.12}$$

where $\tilde{V}$ denotes the solution to the source problem (SP) with running cost function $\tilde{L}$ and source set Y.

To complete the proof, let $\tilde{H}$ be the maximized Hamiltonian corresponding to $\tilde{L}$. By equation (4.12) and Theorem 3.2, $\tilde{V}$ is the maximal solution to (3.9) with $\tilde{H}$ and Y replacing H and X_0. But for $x \in X$,

$$\begin{aligned} \tilde{H}(x,u) &= \max_{v \in TX \cap B_r} \Big(\langle u, v\rangle - \tilde{L}(x,v)\Big) \\ &= \max_{v \in TX \cap B_r} (\langle -u, -v\rangle - \langle \Psi(x), [-v]_+\rangle) = H(x,-u). \end{aligned} \tag{4.13}$$

Comparing (3.9) and (3.10) and using this fact, we conclude that $W = \tilde{V}$ is the maximal solution to (3.10). This proves Theorem 3.3. □

We now use Theorem 3.3 to prove Theorem 3.4.

Proof of Theorem 3.4. Let W satisfy the conditions of the theorem, and let $\overline{W}$ be the maximal Lipschitz continuous solution to (3.10), or equivalently (by Theorem 3.3), the solution to (TP). We show that $W = \overline{W}$. Clearly $W \le \overline{W}$, since W and $\overline{W}$ solve (3.10) and $\overline{W}$ is maximal. For the reverse inequality, the conditions of the theorem require that for each $x \in X \smallsetminus Y$, there is a $T > 0$ and path $\hat{\phi} \in \Phi_T$ such that $\hat{\phi}(0) = x, \hat{\phi}(T) \in Y$ and $W(x) = c(\hat{\phi})$, implying that

$$W(x) = c(\hat{\phi}) \ge \inf \{c(\phi) \colon \phi \in \Phi_T \text{ for some } T \ge 0, \phi(0) = x, \phi(T) \in Y\} = \overline{W}(x). \qquad \square$$

Remark 4.6. As noted above, the viscosity supersolution test for (4.1) is vacuous. By the proof of Theorem 3.2, being a Lipschitz viscosity subsolution of (4.1) is equivalent to being a Lipschitz almost everywhere subsolution.

These two points allow us to deduce that any convex combination of viscosity solutions of (4.1) is itself a viscosity solution of (4.1). In particular, let V be the maximal solution of (4.1). Since 0 is a trivial solution, λV is also a solution for any $\lambda \in [0, 1]$. Hence, (4.1) has infinitely many solutions, which is an interesting phenomenon mathematically, and the question on finding the unique maximal solution is clearly not easy in general. The verification result (Theorem 3.4) allows us to handle this difficulty efficiently in our settings. In other PDE settings, in particular those in which both the Lagrangian and Hamiltonian are superlinear, direct analogues of Theorem 3.4 may not be useful, as the existence of optimal finite-time paths may not be guaranteed.

5. Large deviations in evolutionary game theory

We now apply the results developed above to solve control problems arising in large deviations analyses of dynamic models from evolutionary game theory. We define a class of Markov chains $\mathbf{X}^{N,\eta} = \{X_k^{N,\eta}\}_{k=0}^{\infty}$ parameterized by a population size N and a noise level $\eta > 0$, and which run on discrete grids $\mathcal{X}^N$ of mesh size $\frac{1}{N}$ in the simplex X. These Markov chains describe the evolution of aggregate behavior in a population of N strategically interacting agents. Each agent adjusts his actions over time by following a noisy best response rule, under which the probabilities of choosing suboptimal actions vanish at exponential rates in η^{-1}.

In the classes of games we focus on here, the Markov chains typically approach and then remain near "pure" states corresponding to strict Nash equilibria of the underlying game, but the ergodicity of these processes ensure that transitions between such states must occur. Sandholm and Staudigl (2016) (henceforth SS16) develops large deviations results that describe the waiting times until and likely paths of transitions between strict Nash equilibria. The analysis concerns the *small noise double limit*, meaning that the noise level η is first taken to zero, and then the population size N to infinity.[3] Large deviation properties of $\mathbf{X}^{N,\eta}$ are described in terms of solutions to optimal control problems with semilinear running costs and state constraints. Because existing verification theorems do not allow for state constraints, analyses of examples in this paper were limited to cases in which the state constraints did not bind.[4] The results developed above allow us to complete the large deviations analysis in considerably more general environments.

5.1 The model

5.1.1 Finite-population games

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

We identify a *finite-population game* with its payoff function $F^N \colon \mathcal{X}^N \to \mathbb{R}^n$, where $F^N_i(x) \in \mathbb{R}$ is the payoff to action i when the population state is $x \in \mathcal{X}^N$. Only the values that the function F^N_i takes on the set $\mathcal{X}^N_i = \{x \in \mathcal{X}^N \colon x_i > 0\}$ are meaningful, since at the remaining states in $\mathcal{X}^N$ action i is unplayed.

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

$$(5.1) \qquad F^N_{i\to j}(x) = F^N_j(x + \tfrac{1}{N}(e_j - e_i)).$$

Using this notation, we define a state $x \in \mathcal{X}^N$ to be a *Nash equilibrium* of F^N if no agent can

[3] We discuss the reverse order of limits in Section 7.

[4] More precisely, the relevant control problems were extended to be defined on the affine hull of the simplex. Solutions to these problems were obtained using verification theorems due to Boltyanskii (1966) and Piccoli and Sussmann (2000) (see also Schättler and Ledzewicz (2012)), and the optimal controlled trajectories for source states in the simplex were shown to reach the target set without leaving the simplex.

obtain a higher payoff by switching actions:

$$(5.2)\qquad i \in \underset{j\in S}{\operatorname{argmax}}\, F^N_{i\to j}(x) \;\text{ whenever } x_i > 0.$$

In the examples we focus on here, agents are matched against all opponents to play a symmetric two-player normal form game $A \in \mathbb{R}^{n\times n}$, where A_{ij} is the payoff that an agent playing i obtains when matched against an agent playing j. As self-matching is not allowed, then the average payoffs obtained by agents playing action i is

$$(5.3)\qquad F^N_i(x) = \tfrac{1}{N-1}\, e_i' A(Nx - e_i) = (Ax)_i + \tfrac{1}{N-1}((Ax)_i - A_{ii}).$$

In our examples, the normal form game A is a *coordination game*, meaning that

$$(5.4)\qquad A_{ii} > A_{ji} \text{ for all distinct } i, j \in \mathscr{A},$$

so that if one's match partner plays i, one is best off playing i oneself. One can verify that if F^N is the population game obtained by matching in A without self-matching, then the Nash equilibria of F^N are precisely the pure population states.

5.1.2 Noisy best response protocols and unlikelihood functions

In our model of stochastic evolution, agents occasionally receive opportunities to switch actions. Upon receiving a revision opportunity, an agent selects a action by employing a *noisy best response protocol* $\sigma^\eta \colon \mathbb{R}^n \to \operatorname{int}(X)$ with *noise level* $\eta > 0$, a function that maps vectors of payoffs to probabilities of choosing each action. When a revising agent's evaluation of the actions' payoffs is described by the vector π, then the probability that this agent proceeds by playing action j is $\sigma_j(\pi)$.

Under a noisy best response protocol, the probability of playing an action that is suboptimal at π vanishes at a well-defined rate as η approaches zero. These rates are captured by the *unlikelihood function* $\Upsilon \colon \mathbb{R}^n \to \mathbb{R}^n_+$ associated with the protocols $\{\sigma^\eta\}_{\eta\in(0,\bar\eta)}$, defined by

$$(5.5)\qquad \Upsilon_j(\pi) = -\lim_{\eta\to 0} \eta \log \sigma^\eta_j(\pi),$$

Since (5.5) is equivalent to

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

$\Upsilon_j(\pi)$ is the rate of decay of the probability that action j is chosen as η approaches zero. We assume that the unlikelihood function Υ is Lipschitz continuous, and that the actions with unlikelihood zero at π are the optimal actions at π.

$$\Upsilon_j(\pi) = 0 \text{ if and only if } j \in \operatorname*{argmax}_{k \in S} \pi_k. \tag{5.6}$$

Condition (5.6) implies that the probability with which optimal actions are chosen under $\sigma^\eta(\pi)$ approaches one as η approaches zero.

We focus on protocols derived from the *additive random utility model*, in which choices are made after payoffs are perturbed by i.i.d. random shocks ε_k^η:

$$\sigma_j^\eta(\pi) = \mathbb{P}\left(j \in \operatorname*{argmax}_{k \in S}\, (\pi_k + \varepsilon_k^\eta)\right) \tag{5.7}$$

Example 5.1. Logit choice. If each ε_k^η has a Gumbel distribution with mean zero and variance $\frac{\eta^2\pi^2}{6}$, then (5.7) generates the *logit choice protocol* with noise level η (Blume (1997)). Choice probabilities under this protocol can be stated in closed form as

$$\sigma_j^\eta(\pi) = \frac{\exp(\eta^{-1}\pi_j)}{\sum_{k \in S} \exp(\eta^{-1}\pi_k)}. \tag{5.8}$$

(see Anderson et al. (1992)). It is easy to verify that logit choice generates the piecewise linear unlikelihood function

$$\Upsilon_j(\pi) = \max_{k \in S} \pi_k - \pi_j. \tag{5.9}$$

For later convenience, we express (5.9) in vector form as

$$\Upsilon(\pi) = \left(\max_{k \in S} \pi_k\right)\mathbf{1} - \pi, \tag{5.10}$$

where $\mathbf{1}$ is the column vector of 1s. ◆

Example 5.2. Probit choice. If each ε_k^η has a normal distribution with mean 0 and variance η, then (5.7) becomes the *probit choice protocol* (Myatt and Wallace (2003)). There is no elementary closed-form expression for this protocol. But by applying Cramér's theorem on large deviations of sums of i.i.d. random variables (Dembo and Zeitouni (1998)), Dokumacı and Sandholm (2011) derive a piecewise quadratic expression for the protocol's unlikelihood function. When there are $n = 3$ actions, this unlikelihood function is

$$\Upsilon_1(\pi) = \begin{cases} 0 & \text{if } \pi_1 \geq \pi_2 \vee \pi_3, \\ \frac{(\pi_1-\pi_2)^2}{2} & \text{if } \pi_2 \geq \pi_1 \geq \pi_3, \text{ or if } \pi_2 \geq \pi_3 \geq \pi_1 \text{ and } \frac{\pi_1+\pi_2}{2} \geq \pi_3, \\ \frac{(\pi_1-\pi_3)^2}{2} & \text{if } \pi_3 \geq \pi_1 \geq \pi_2, \text{ or if } \pi_3 \geq \pi_2 \geq \pi_1 \text{ and } \frac{\pi_1+\pi_3}{2} \geq \pi_2, \\ \frac{(\pi_1-\pi_2)^2+(\pi_1-\pi_3)^2+(\pi_2-\pi_3)^2}{6} & \text{otherwise.} \end{cases} \tag{5.11}$$ ◆

The differences between the logit and probit unlikelihood functions (5.9) and (5.11) can be traced to the tail properties of the Gumbel and normal distributions. The thick right tail of the Gumbel distribution leads logit unlikelihoods to grow linearly, and to depend only on underperformance relative to the optimal payoff. In contrast, normal noise leads probit unlikelihoods to grow quadratically, and allows them to depend on the payoffs of three or more actions—for instance, when the action at issue generates a markedly lower payoff than at least two other actions. For further discussion, see Dokumacı and Sandholm (2011) and Arigapudi (2018).

5.1.3 The stochastic evolutionary process

A population game F^N and a revision protocol σ^η define a stochastic evolutionary process. The process runs in discrete time, with each period taking $\frac{1}{N}$ units of clock time. During each period, a single agent is chosen at random from the population to update his action using the noisy best response protocol σ^η.

This procedure described above generates a Markov chain $\mathbf{X}^{N,\eta} = \{X_k^{N,\eta}\}_{k=0}^\infty$ on the state space $\mathcal{X}^N$. The transition probabilities $P_{x,y}^{N,\eta}$ for the process $\mathbf{X}^{N,\eta}$ are given by

$$P_{x,y}^{N,\eta} \equiv \mathbb{P}\left(X_{k+1}^{N,\eta} = y \,\middle|\, X_k^{N,\eta} = x\right) = \begin{cases} x_i\, \sigma_j^\eta(F_{i\to\cdot}^N(x)) & \text{if } y = x + \frac{1}{N}(e_j - e_i),\, j \neq i \\ \sum_{i=1}^n x_i\, \sigma_i^\eta(F_{i\to\cdot}^N(x)) & \text{if } y = x. \end{cases} \tag{5.12}$$

If the revision opportunity is assigned to an agent playing action i (probability x_i), and this agent chooses to switch to action $j \neq i$ (probability $\sigma_j^\eta(F_{i\to\cdot}^N(x))$), then the increment in the state is $\frac{1}{N}(e_j - e_i)$: there is one fewer i player and one more j player.

5.2 Large deviations analysis

When the noise level η is small, typical sample paths of the process $\mathbf{X}^{N,\eta}$ are sample paths of an exact-best-response process $\mathbf{X}^{N,0}$, under which revising agents either switch to or continue playing optimal actions. Since A is a coordination game, we expect the recurrent classes of the exact-best-response process to be the pure population states e_i, corresponding to the pure Nash equilibria of A.[5]

To evaluate transitions between Nash equilibria, we study the behavior of the process $\mathbf{X}^{N,\eta}$ in the *small noise double limit*. First, we take the noise level η to zero, leaving the population size N and hence the state space $\mathcal{X}^N$ fixed. Following Freidlin and Wentzell (1984) and Catoni (1999), we describe the large deviation properties in this small noise limit as solutions to discrete optimal control problems on $\mathcal{X}^N$. As the population size is next made large, the state spaces $\mathcal{X}^N$ become ever-finer grids in the simplex X, and the discrete optimal control problems are approximated arbitrarily well by continuous optimal control problems. It is to these latter problems that we apply the results on Hamilton-Jacobi equations developed above.

5.2.1 The small noise limit

Large deviation properties in the small noise limit are described in terms of solutions to path cost minimization problems on $\mathcal{X}^N$, where the source is a pure population state. To start the large deviations analysis, we define the *cost of a step* from state $x \in \mathcal{X}^N$ to state $y \in \mathcal{X}^N$ to be the exponential rate of decay of the conditional probability observing this increment as η approaches 0:

$$c^N_{x,y} = -\lim_{\eta \to 0} \eta \log P^{N,\eta}_{x,y}. \tag{5.13}$$

Using definitions (5.5) and (5.12), we can represent step costs in terms of the game's payoff function and the protocol's unlikelihood function:

$$c^N_{x,y} = \begin{cases} \Upsilon_j(F^N_{i\to\cdot}(x)) & \text{if } y = x + \frac{1}{N}(e_j - e_i) \text{ and } j \neq i, \\ \min\limits_{i:\, x_i > 0} \Upsilon_i(F^N_{i\to\cdot}(x)) & \text{if } y = x. \end{cases} \tag{5.14}$$

The key case in (5.14) is the first one, which says that the cost of a step in which an i player switches to action j is the unlikelihood of action j given i's current payoff opportunities.

Next we consider finite-length paths through the state space $\mathcal{X}^N$, which are sequences $\phi^N = \{\phi^N_k\}^{\ell}_{k=0}$ of *length* $\ell < \infty$ in which successive states are either adjacent in $\mathcal{X}^N$ or identical. We let Φ^N_ℓ denote the set of length-ℓ paths. Since each period lasts $\frac{1}{N}$ time units, the *duration* of a length-ℓ path in clock time is $T = \ell/N$.

The *cost of a path* $\phi^N = \{\phi^N_k\}^{\ell}_{k=0}$ of length ℓ is the sum of the costs of its steps:

[5] While we conjecture that these pure states are the only recurrent classes of this process, we do not have a proof that this is the case. The difficulty lies in the discreteness of the state space $\mathcal{X}^N$, as the expected result holds in the large population limit. See SS16 Sections 3.1 and 4.2 for discussions. In the examples in Section 6, one can verify that the pure states are the recurrent states of the exact-best-response process.

(5.15) $$c^N(\phi^N) = \sum_{k=0}^{\ell-1} c^N_{\phi^N_k, \phi^N_{k+1}}.$$

Definitions (5.12) and (5.14) imply that the cost of a path is the rate at which the probability of following this path decays as the noise level vanishes: for fixed N, we have

$$\mathbb{P}\left(X^{N,\eta}_k = \phi^N_k, k = 0, \ldots, \ell \,\middle|\, X^{N,\eta}_0 = \phi^N_0\right) = \prod_{k=0}^{\ell-1} P^{N,\eta}_{\phi^N_k, \phi^N_{k+1}} \approx \exp(-\eta^{-1} c^N(\phi^N)).$$

where $\approx$ refers to the order of magnitude in η as η approaches zero.

Finally, we define the *transition cost* from state e_i to state e_j and the *exit cost* from state e_i for population size N:

(5.16) $$C^N(e_i, e_j) = \min\{c^N(\phi^N) \colon \phi^N \in \Phi^N_\ell \text{ for some } \ell \geq 0, \phi^N_0 = e_i, \phi^N_\ell = e_j\},$$

(5.17) $$C^N(e_i, e_{-i}) = \min_{j \neq i} C^N(e_i, e_j).$$

Freidlin and Wentzell (1984) and Catoni (1999) show that these quantities describe the large deviations properties of $\mathbf{X}^{N,\eta}$ in the small noise limit (cf. SS16 Sections 3.4 and 6.2). For instance, suppose the process $\mathbf{X}^{N,\eta}$ is initialized at state e_i, and let $\tau^{N,\eta}$ denote the time at which the process first reaches the set $\{e_j \colon j \neq i\}$. Then Proposition 4.2 of Catoni (1999) implies that

(5.18) $$\lim_{\eta \to 0} \eta \log \mathbb{E}\tau^{N,\eta} = C^N(e_i, e_{-i}).$$

Thus the exit cost $C^N(e_i, e_{-i})$ is the exponential rate of increase of the expected time for the process travel from e_i to another Nash equilibrium as η approaches zero. For its part, the transition cost $C^N(e_i, e_j)$ is useful for determining the limiting behavior of the stationary distributions of $\mathbf{X}^{N,\eta}$, as we explain in Example 6.2 below.

5.2.2 The small noise double limit

In general, the discrete transition cost problem (5.16) and the discrete exit cost problem (5.17) are computationally demanding. SS16 show that taking the second limit in the population size allows one to approximate the solutions to these discrete problems with solutions to continuous, state-constrained optimal control problems.

The first step in doing so is to propose an analogue of the discrete path cost function (5.15) for continuous paths. Equations (5.14) and (5.15) show that the cost generated by a switch to action j from any other action at state x is $\Upsilon_j(F^N_{i\to\cdot}(x))$, and that the cost of a path is the sum of the costs of its steps. This suggests the following definition of the cost of a continuous path:[6]

$$c(\phi) = \int_0^T \langle \Upsilon(F(\phi(t))), [\dot{\phi}(t)]_+\rangle \, \mathrm{d}t, \quad \text{where } F(x) = Ax. \tag{5.19}$$

We turn now to constraints. To arise as the limit of discrete paths ϕ^N whose transitions are between adjacent states in $\mathcal{X}^N$, a continuous path ϕ should satisfy the state space restriction $\phi(t) \in X$ for all t, and the speed restriction $|\dot{\phi}(t)| \leq 2$ for almost all t. We incorporate these restrictions explicitly into the following expression for the running costs that define the path cost function (5.19):

$$L(x, v) = \begin{cases} \langle \Upsilon(F(x)), [v]_+\rangle & \text{if } v \in TX(x) \cap B_2, \\ +\infty & \text{otherwise.} \end{cases} \tag{5.20}$$

Of course, Remark 2.1 tells us that the speed restriction in (5.20) is inconsequential.

To justify focusing on continuous control problems defined by (5.20), we require an approximation that relates the minimal costs in the discrete minimization problems (5.16) and (5.17) with discrete path cost function (5.15) to solutions of continuous minimization problems with cost function (5.19). With this aim in mind, we define

$$C(e_i, e_j) = \min\{c(\phi) \colon \phi \in \Phi_T \text{ for some } T \geq 0, \phi(0) = e_i, \phi(T) = e_j\}, \tag{5.21}$$

$$C(e_i, e_{-i}) = \min_{j \neq i} C(e_i, e_j). \tag{5.22}$$

Theorem 5.6 of SS16 shows that under mild regularity conditions, the following approximations hold, where (5.24) is an immediate consequence of (5.23):

$$\lim_{N\to\infty} \frac{1}{N} C^N(e_i, e_j) = C(e_i, e_j), \tag{5.23}$$

$$\lim_{N\to\infty} \frac{1}{N} C^N(e_i, e_{-i}) = C(e_i, e_{-i}). \tag{5.24}$$

[6]Definition (5.19) ignores the the second case of the step cost (5.14), which indicates that staying still is costly at states where no agent is playing a best response. Because staying still does not help advance the state toward the target set, ignoring this possibility is without loss. For more on this, and for a heuristic derivation of (5.19), see SS16 Section 4.3.

5.2.3 Optimal paths when state constraints do not bind

SS16 compute exit costs (5.22) and transition costs (5.21) for three-action coordination games (5.4) under the logit choice protocol (5.8). To ensure that the state constraints in problems (5.22) and (5.21) do not bind, they focus on coordination games satisfying two conditions. To state these conditions, we let A^i denote the ith row of A, so that the payoff to action i at state x is $F_i(x) = A^i x$.

(5.25) there is a unique $x^* \in X^\circ$ such that $A^i x^* = A^j x^*$ for all $i, j \in \mathscr{A}$,

(5.26) $A^i(e_i - e_k) > A^j(e_i - e_k)$ for all $i, j, k \in \mathscr{A}$ with $i \notin \{j, k\}$.

Condition (5.25) says that normal form game A admits an interior Nash equilibrium x^*, at which all three actions earn equal payoffs. Condition (5.26), the *marginal bandwagon property* of Kandori and Rob (1998), says that switches to action i from other actions benefit action i more than other actions.

Figure 1 illustrates both the assumptions above and the resulting form of the optimal feedback controls for problems (5.22) and (5.21). The diagrams divide the simplex X into *best response regions*,

$$\mathscr{B}^i = \{x \in X \colon A^i x \geq A^j x \text{ for all } j \in \mathscr{A}\},$$

in which each action i is optimal. The interior Nash equilibrium x^* posited in condition (5.25) is the point where the three regions intersect. Let $\mathscr{B}^{ij} = \mathscr{B}^i \cap \mathscr{B}^j$ denote the boundary between the best response regions for actions i and j. Condition (5.26) amounts to the assumption that these boundaries do not emanate from x^* at "sharp angles".[7]

In coordination games, the direct path from any state in best response region $\mathscr{B}^i$ to equilibrium e_i has zero cost under (5.19). Thus to solve the exit problem (5.22), it is enough to determine the least cost path from e_i to the complement of $\mathscr{B}^i$. Under the assumptions from SS16, the optimal feedback control for the exit problem takes the form shown in Figure 1(i): $\mathscr{B}^1$ is divided into subregions in which motion is in direction $e_2 - e_1$ (i.e., agents switch from action 1 to action 2) or in direction $e_3 - e_1$, and the optimal exit path from e_1 depends on whether the boundary between the subregions is to the left or right of e_1. Likewise, to solve transition problems with target state e_k, it is enough to determine the least cost path from each state outside of $\mathscr{B}^i$ to a state in $\mathscr{B}^i$. Figure 1(ii) illustrates a typical optimal feedback control. Notice that in both pictures, the optimal feedback control at

[7] In Figure 1(ii), $\mathscr{B}^{12}$ is the segment from x^* to x^{ij}. Condition (5.26) requires that the angle at which $\mathscr{B}^{12}$ proceeds from x^* be in a 60° range, that which restricts x^{ij} to lie between points $\tilde{x}^{ik}$ and $\tilde{x}^{jk}$.

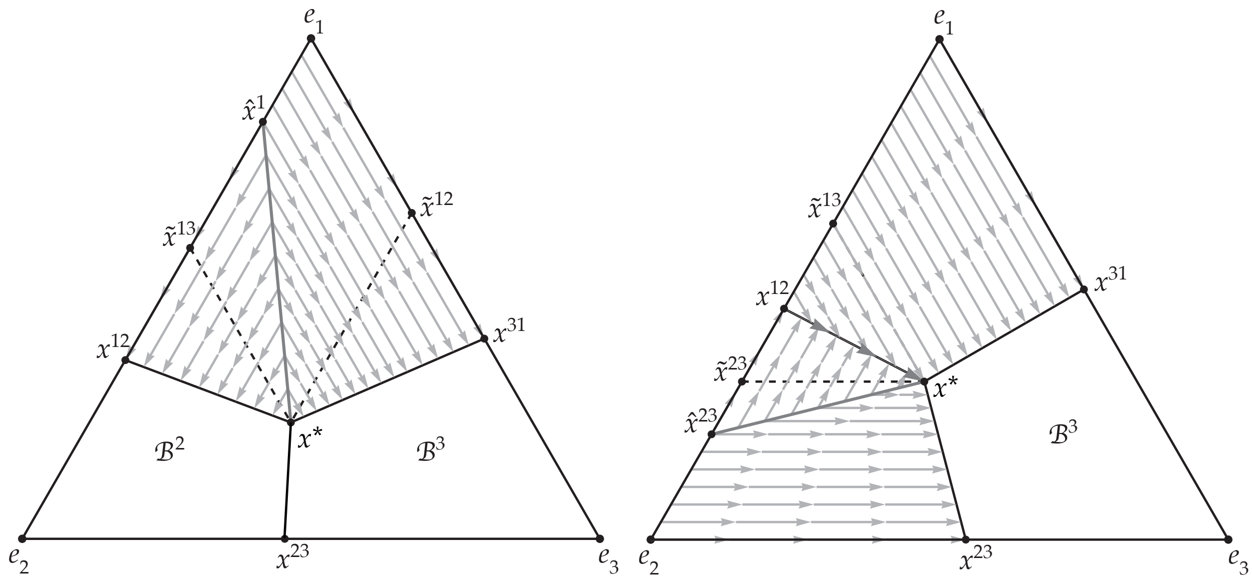

(i) exit from e_1 (ii) transitions to e_3

Figure 1: Optimal feedback controls under conditions (5.25) and (5.26).

almost all states is in one of the basic directions $e_j - e_i$, as suggested in Remark 3.1; an exception is the forced motion on segment from state x^{12} to state x^* in Figure 1(ii).

For present purposes, what is most important about conditions (5.25) and (5.26) is that they ensure that the state constraints in problems (5.22) and (5.21) do not bind. SS16 establish this by extending these problems to be defined on aff(X) in a suitable way, solving the extended problems using results of Boltyanskii (1966) and Piccoli and Sussmann (2000), and showing that the restrictions of these solutions to X are forward invariant on X. But in games in which conditions (5.25) and (5.26) fail, or if choice protocols other than logit (5.8) are used, then state constraints may bind, and this approach cannot succeed.

6. Optimal exit and transition paths with binding state constraints

In this section, we use Theorem 3.4 to solve exit and transition problems without placing restrictions on the positions of the best response regions. Indeed, along many of the optimal paths we compute, the constraint to remain in the simplex binds. Here we confine ourselves to discussing key features of the feedback controls and the solutions to problems (5.21) and (5.22). The formulas for the optimal value functions and all details about their verification are presented in Appendices B and C.

Example 6.1. An exit problem. Consider evolution under the logit rule (5.8) in the following

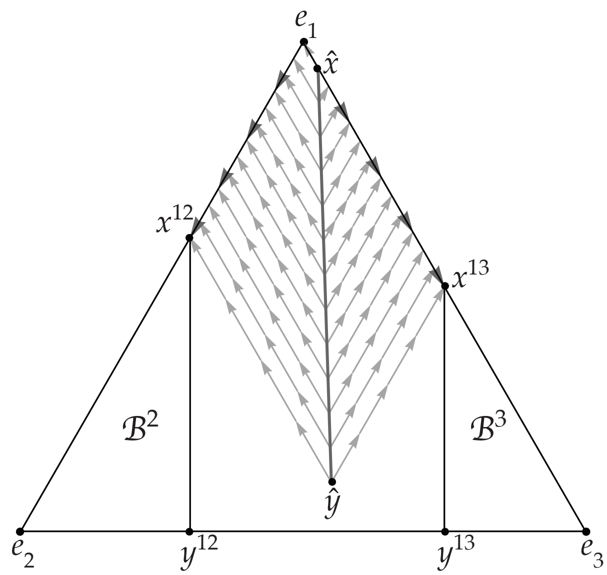

Figure 2: The optimal feedback control for the problem of exit from e_1 in games of form (6.1).

class of coordination games:

$$A = \begin{pmatrix} a & 0 & 0 \\ 0 & b & -(2a+b) \\ 0 & -(2a+c) & c \end{pmatrix}, \quad \text{with } a + c \geq b > c > 0. \tag{6.1}$$

Games of form (6.1) have five Nash equilibria, three at the vertices, and two at boundary states x^{12} and x^{13} (see Figure 2 and equation (B.3)); in particular, they do not admit interior Nash equilibria. Best response regions $\mathcal{B}^2$ and $\mathcal{B}^3$ do not intersect, and boundaries $\mathcal{B}^{12}$ and $\mathcal{B}^{13}$ are vertical lines. It is easy to verify that region $\mathcal{B}^2$ is larger than region $\mathcal{B}^3$. Evidently, games in this class satisfy neither condition (5.25) nor condition (5.26) from Section 5.2.3.

We first consider the problem of exit from equilibrium e_1. As discussed above, it is enough to determine the least cost paths from all states in best response region $\mathcal{B}^1$ to states outside this region. Since within this region motion in basic directions that increases the use of action 1 has zero cost, there are zero-cost paths out of $\mathcal{B}^1$ from states below segment $x^{12}\hat{y}$, from which $\mathcal{B}^2$ can be reached at zero cost by proceeding northwest in direction $e_1 - e_3$, as well as from states below segment $\hat{y}x^{13}$, from which $\mathcal{B}^3$ can be reached at zero cost by proceeding northeast in direction $e_1 - e_2$.

The optimal feedback control from the subset of $\mathcal{B}^1$ that remains is illustrated in Figure 2. From the interior states in this subset motion proceeds in one of two basic directions. But unlike in Figure 1(i), the directions followed here *increase* the use of optimal action 1 rather than decreasing it: basic direction $e_1 - e_3$ or $e_1 - e_2$ is followed at zero cost until a state constraint binds. From that moment forward, motion proceeds along a boundary of X, in direction $e_2 - e_1$ or direction $e_3 - e_1$, until $\mathcal{B}^2$ or $\mathcal{B}^3$ is reached. Thus as in Figure 1(i), the optimal exit path from e_1 proceeds along the boundary of X, but here the boundary path arises because of a binding state constraint.

In the figure, the optimal exit path proceeds along the boundary from e_1 to e_2. Let $\phi^{12}\colon [0,1] \to X$ be the parameterization of this path with constant speed $|\dot{\phi}^{12}(t)| = 2$; that is, $\phi^{12}(t) = (1-t, t, 0)'$. This path stays in best response region $\mathcal{B}^1$ until time $t = x_2^{12} = \frac{a}{a+b}$, when it reaches state $x^{12} \in \mathcal{B}^{12}$; after this, the remaining motion to state e_2 has zero cost. Thus using (5.19) and (5.10), we can compute the exit cost from equilibrium e_1:

$$\begin{aligned} C(e_1, e_{-1}) = c(\phi^{12}) &= \int_0^1 \langle \Upsilon(A\phi^{12}(t)), [\dot{\phi}^{12}(t)]_+\rangle \,\mathrm{d}t = \int_0^{x_2^{12}} \langle (\mathbf{1}A^1 - A)\phi^{12}(t)), e_2\rangle \,\mathrm{d}t \\ &= \int_0^{x_2^{12}} (A^1 - A^2)\phi^{12}(t)\,\mathrm{d}t = \int_0^{a/(a+b)} (a(1-t) - bt)\,\mathrm{d}t = \frac{1}{2}\cdot\frac{a^2}{a+b}. \end{aligned} \tag{6.2}$$

The reason that the optimal control requires "retreating" followed by motion along a boundary can be explained as follows. In an extended version of this problem without state constraints, with the boundaries $\mathcal{B}^{12}$ and $\mathcal{B}^{13}$ of the best response regions extending vertically beyond the triangle, motion in direction $e_1 - e_3$ or $e_1 - e_2$ would ultimately reach a boundary, allowing escape from $\mathcal{B}^1$ at zero cost. The constraint that the state remains in X precludes following such a path. However, by proceeding, say, northwest, in direction $e_1 - e_3$, one reduces the payoff difference between actions 2 and 1, and pays no cost in order to do so. Thus the optimal paths reduce this payoff difference as much as possible before at last having agents switch to suboptimal action 2. ◆

Example 6.2. Transition problems and stochastic stability. Following Foster and Young (1990), Kandori et al. (1993), and Young (1993), much work on stochastic evolutionary game dynamics concerns the identification of *stochastically stable states*. These are the points in the simplex where the mass in the stationary distribution of the process $\mathbf{X}^{N,\eta}$ becomes concentrated as the parameters of the process approach their limiting values. Most work on stochastic stability in games focuses on processes in which the probabilities of suboptimal choices are state independent, so that the improbability of a transition between equilibria is determined by the number of suboptimal choices that the transition requires. The characterization of costs of transitions in terms of control problems in SS16 allows one to evaluate stochastic stability in the small noise double limit when mistake probabilities are payoff dependent, provided that the control problems can be solved. The results in the present paper allow us to solve these control problems even in environments in which state constraints bind.

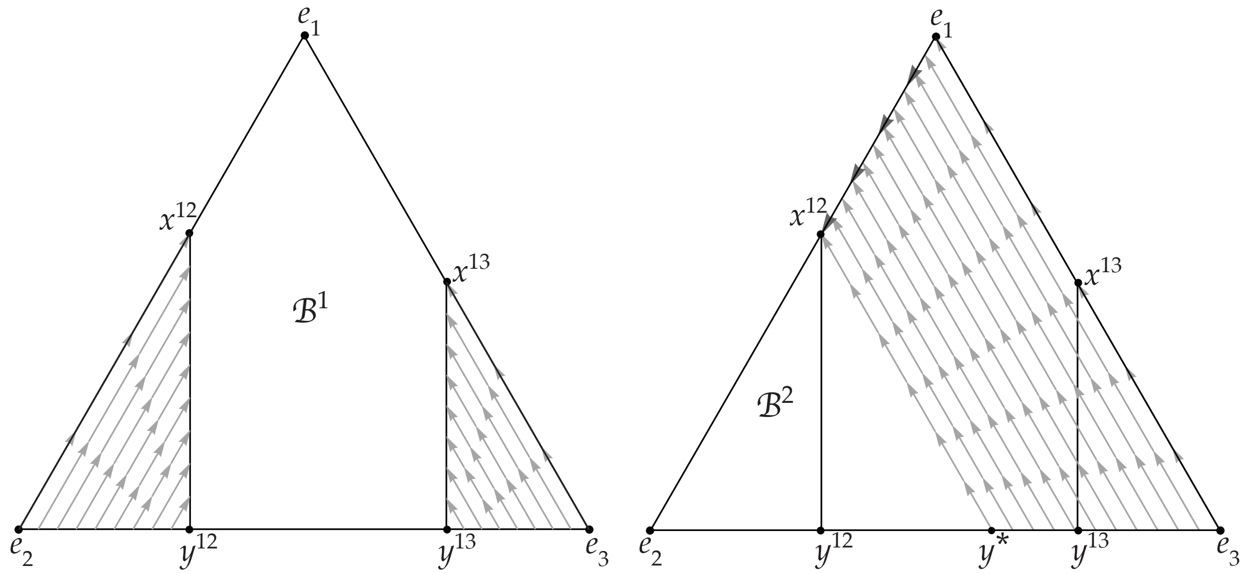

(i) transitions to e_1 (ii) transitions to e_2

Figure 3: Optimal feedback controls for transition problems in games of form (6.1). The diagram of optimal transition paths to e_3 resembles the mirror image of panel (ii).

We illustrate these ideas using the class of games (6.1) from the previous example. To begin the stochastic stability analysis, we determine the costs of the least-cost transition paths between each ordered pair of distinct equilibria. As noted in Section 5.2.3, it is enough to solve the target problem for each equilibrium e_i, or, equivalently, for each best response region $\mathcal{B}^i$.

The optimal feedback controls for the transition problems are presented in Figure 3. Panel (i) shows that the optimal paths to region $\mathcal{B}^1$ from region $\mathcal{B}^2$ are in direction $e_1 - e_2$, and that the optimal paths to $\mathcal{B}^1$ from $\mathcal{B}^3$ are in direction $e_1 - e_3$. The state constraints do not appear to bind in this problem, and indeed the problem could be solved using the approach from SS16.

Panel (ii) of the figure shows the optimal paths to $\mathcal{B}^2$. Here motion is in direction $e_1 - e_3$ until either $\mathcal{B}^2$ or the $x_3 = 0$ boundary is reached; in the latter case, the trajectory proceeds along the boundary in direction $e_2 - e_1$ until $\mathcal{B}^2$ is reached. The optimal paths to $\mathcal{B}^3$ take the same form, mutatis mutandis. Evidently, in these problems the state constraints bind.

To determine the stochastically stable state, we name and indicate the costs of the optimal transition paths. For each ordered pair (i, j) of distinct actions, let ϕ^{ij} denote the direct boundary path from e_i to e_j.[8] Then the foregoing analysis shows that the transition costs are as follows:

[8] Of course, the portion of this path in $\mathcal{B}^j$ generates zero cost.

$$
\begin{aligned}
&C(e_2,e_1) = c(\phi^{21}), \qquad && C(e_1,e_2) = c(\phi^{12}), \qquad && C(e_1,e_3) = c(\phi^{13}),\\
&C(e_3,e_1) = c(\phi^{31}), \qquad && C(e_3,e_2) = c(\phi^{31}) + c(\phi^{12}), \qquad && C(e_2,e_3) = c(\phi^{21}) + c(\phi^{13}).
\end{aligned}
$$

Next, following Freidlin and Wentzell (1984) and Young (1998), for each equilibrium e_i we determine the minimal cost of a tree with root e_i on nodes $\{e_1, e_2, e_3\}$, where the cost of a tree is the sum of the transition costs of its directed edges. Since the optimal transitions between e_2 and e_3 follow indirect routes through e_1, it is easy to verify that the minimum cost trees are the unique ones that do not use these transitions. Letting $R(e_i)$ denote the minimum cost of an e_i tree, we have

$$
\begin{aligned}
R(e_1) &= C(e_2,e_1) + C(e_3,e_1) = \frac{1}{2}\left(\frac{b^2}{a+b} + \frac{c^2}{a+c}\right),\\
R(e_2) &= C(e_1,e_2) + C(e_3,e_1) = \frac{1}{2}\left(\frac{a^2}{a+b} + \frac{c^2}{a+c}\right),\\
R(e_3) &= C(e_2,e_1) + C(e_1,e_3) = \frac{1}{2}\left(\frac{b^2}{a+b} + \frac{a^2}{a+c}\right),
\end{aligned}
\tag{6.3}
$$

(cf. Lemma B.1 and equation (B.3)). By condition (6.1), $R(e_3) > R(e_2)$, and $R(e_1) < R(e_2)$ if and only if $a > b$. Thus in games of form (6.1) under the logit rule, display (6.3) and SS16 Theorem 6.3 imply that the equilibrium with the highest payoff is stochastically stable.

An intuition for this result is as follows. Since all transition costs are determined by paths through e_1, the minimum tree costs appearing in display (6.3) are determined by the two-action pure coordination games with actions 1 and 2 and with actions 1 and 3. Since $A_{33} = c < b = A_{22}$, there is more selection pressure against action 3 in the latter game than against action 2 in the former, so action 3 cannot be stochastically stable.[9] Since action 3 affects both $R(e_1)$ and $R(e_2)$ through the addition of the edge from e_3 to e_1, the selection between e_1 and e_2 comes down to the selection between them in the pure coordination game with actions 1 and 2 only, where the equilibrium with the higher payoff is selected. ◆

Example 6.3. A transition problem in a potential game. We next consider the problem of transitions to equilibrium e_3 in the following class of games:

$$
A = \begin{pmatrix} a & 0 & 0 \\ 0 & b & -d \\ 0 & -d & c \end{pmatrix} \text{ with } a,b,c,d > 0 \text{ and } a > d > \sqrt{bc}. \tag{6.4}
$$

[9] To be more precise, note that the inequality $c < b$ is enough to conclude that $R(e_2) < R(e_3)$.

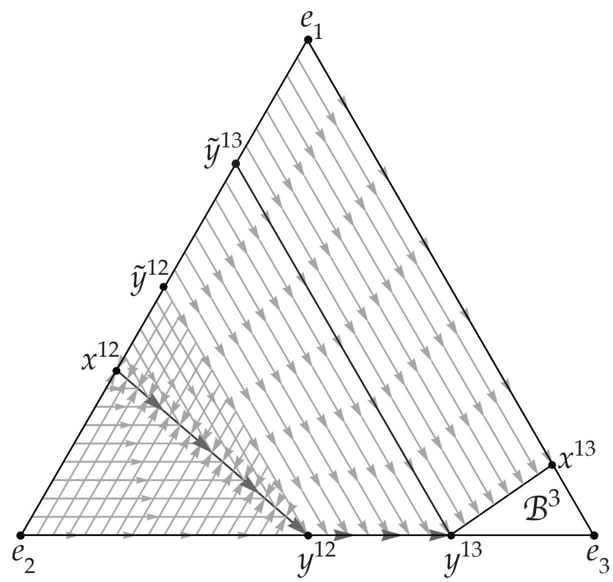

Figure 4: Optimal feedback controls for the problem of transitions to state e_3 in potential games of form (6.4).

Like those in class (6.1), games in class (6.4) admit five Nash equilibria, three at the vertices, and two at boundary states x^{12} and x^{13} (see Figure 4 and equation (C.1)). The condition $d^2 > bc$ implies that the point $x^* \in \text{aff}(X)$ at which all actions earn the same payoff has $x_1^* < 0$ (equation (C.2)). But the condition $a > d$ ensures that games of form (6.4) do satisfy the marginal bandwagon property (5.26). In addition, the symmetry of matrix A—or, more precisely, the symmetry of A as a bilinear form on TX—defines A as a (*symmetric two-player*) *potential game*.[10]

Figure 4 presents the optimal feedback control for the transition problem to equilibrium e_3. We observe that the optimal paths from states in triangle $e_2 y^{13} \tilde{y}^{13}$ each have a portion on segment $y^{12} y^{13}$, where the nonnegativity constraint for action 1 binds.

The combination of the potential game (6.4) with the logit protocol (5.8) causes the Markov chain $\mathbf{X}^{N,\eta}$ to have a variety of special properties.[11] One such property is illustrated in the southwest portion of Figure 4. From almost all states in the triangle $e_2 y^{12} \tilde{y}^{12}$ there are multiple optimal directions of motion: above segment $x^{12} y^{12}$, all convex combinations of $e_2 - e_1$ and $e_3 - e_1$ are optimal directions; below this segment, all convex combinations of $e_1 - e_2$ and $e_3 - e_2$ are optimal. If a path ever hits segment $x^{12} y^{12}$, there is forced motion along this segment until state y^{12} is reached.

The equality of cost integrals that is implied by the existence of multiple optimal controls, along with the symmetry property that defines potential games, suggests that

[10] See Monderer and Shapley (1996) and Sandholm (2010a, Example 3.4).

[11] For instance, $\mathbf{X}^{N,\eta}$ is reversible: see Blume (1997) and Sandholm (2010b, Section 11.5).

the combination of potential games and logit choice generates a running cost function that satisfies an integrability property. This is indeed the case. At states x in best response region $\mathcal{B}^i$, we can use the logit unlikelihood function (5.10) to express the cost of motion (5.20) in direction $e_j - e_i$, $j \neq i$, as

$$L(x, e_j - e_i) = \langle \Upsilon(Ax), [e_j - e_i]_+ \rangle = \langle (\mathbf{1}A^i - A)x, e_j \rangle = \langle (\mathbf{1}A^i - A)x, e_j - e_i \rangle.$$

Thus the costs of paths ϕ moving from states $\phi(t) = x \in \mathcal{B}^i$ in directions $\dot{\phi}(t) \in \mathrm{conv}(\{e_j - e_i \colon j \neq i\})$ behave like line integrals with respect to the vector field

$$\hat{\Psi}(x) = (\mathbf{1}A^i - A)x.$$

Since $\mathrm{conv}(\{e_j - e_i \colon j \neq i\}) \subset TX$, the relevant integrability condition for vector field $\hat{\Psi}$ only concerns displacement directions in the tangent space TX. And because the columns of $\mathbf{1}A^i$ are constant vectors, the restriction of $D\hat{\Psi}(x)$ to a bilinear form on TX is

$$D\hat{\Psi}(x) = -A \quad \text{as an element of } \mathcal{L}^2(TX, \mathbb{R}). \tag{6.5}$$

Potential games are defined by the property that the bilinear form (6.5) is symmetric, giving us the anticipated integrability condition on the vector field $\hat{\Psi}$.[12] ◆

The previous examples show that retaining the logit choice rule (5.8) but dropping the payoff assumptions (5.25) and (5.26) from SS16 can cause the state constraints in problems (5.22) and (5.21) to bind. Applying the results in the present paper, Arigapudi (2018) studies the exit problem in coordination games satisfying (5.25) and (5.26), but under the assumption that agents follow the probit choice rule (Example 5.2). He finds that the piecewise quadratic specification of the probit unlikelihood function (5.11) can lead state constraints to bind in the exit problem even when restrictions (5.25) and (5.26) on the game are maintained. Thus changes in any of the assumptions used in SS16 can lead to control problems in which explicit accounting for state constraints is necessary.

7. Concluding remarks

This paper analyzes optimal control problems with semilinear cost functions and state constraints, characterizing solutions to source and target problems as maximal viscosity

[12] See SS16 Section 7.5.1 for a related discussion, and see Appendix C—specifically, equation (C.5) and the surrounding arguments—to see how integrability is used to verify the optimality of the value function in triangle $e_2 y^{12} \tilde{y}^{12}$.

subsolutions to Hamilton-Jacobi equations, and providing a corresponding verification theorem. These results are used to analyze large deviations problems arising in stochastic dynamic models of play in large population games, providing explicit descriptions of the waiting times until exit from and transitions between stable equilibrium states.

Following SS16, the large deviations analysis in this paper takes place in the small noise double limit, in which the noise level η is first taken to zero, and then the population size N to infinity. It is at least as interesting to consider the large population double limit, which by taking the limit in N first emphasizes the role of averaging effects in equilibrium breakdown and transition. We conjecture that in games with simple payoff structures, predictions based on the two distinct orders of limits agree.

Sandholm and Staudigl (2018, 2019) study large deviations properties of the process $\mathbf{X}^{N,\eta}$ in the large N limit for fixed η, characterizing these properties in terms of solutions to certain convex optimal control problems. In light of these results and those of the present paper, the key argument needed to prove agreement of the two double limits is a selection theorem, one showing that the solutions to the state-constrained Hamilton-Jacobi equations arising for fixed positive η converge to the maximal solutions to the Hamilton-Jacobi equations arising in the limit. The latter are precisely the equations studied in the present paper. Such selection results and their consequences will be developed in future research.

Appendix

A. Derivation of equation (4.7)

We present here the derivation of equation (4.7). Let $|\cdot|_2$ be the ℓ^2 norm on $\mathbb{R}^n$. Fix $y \in X^\circ$. There exists $r > 0$ such that $U_{3r}(y) \subset X^\circ$, where $U_{3r}(y)$ is the open ℓ^2 ball with center y and radius $3r$. To derive (4.7), we use the fact that $U_{2r}(y) \subset X^\circ$ to show that

$$\hat{V}(x) - \hat{V}(y) \leq K|x-y|_2 \quad \text{for all } x \in U_r(y), \tag{A.1}$$

for any fixed $K > n\|\Psi\|_\infty$. The corresponding inequality for $\hat{V}(y) - \hat{V}(x)$ then follows by applying (A.1) and noting that $U_{2r}(x) \subset X^\circ$ as $x \in U_r(y)$.

To establish (A.1), construct a smooth function $f : [0, 2r) \to \mathbb{R}$ such that

$$\begin{cases} f(s) = Ks & \text{for } s \in [0, r], \\ f'(s) \geq K & \text{for } s \in (0, 2r), \\ \lim_{s \to 2r} f(s) = +\infty. \end{cases}$$

Let $\varphi(x) = \hat{V}(y) + f(|x-y|_2)$ for all $x \in U_{2r}(y)$ be our test function. Geometrically, φ has a cone shape, blows up on the boundary of $U_{2r}(y)$, and is smooth in $U_{2r}(y) \setminus \{y\}$. It is clear that $\hat{V} - \varphi$ has a maximizer z in $U_{2r}(y)$. Suppose that $z \neq y$. Then by the viscosity subsolution test, we have $H(z, D\varphi(z)) \leq 0$ with

$$D\varphi(z) = f'(|z-y|_2)\frac{z-y}{|z-y|_2} \quad \text{and} \quad |D\varphi(z)| \geq |D\varphi(z)|_2 = f'(|z-y|_2) \geq K,$$

which contradicts (3.5). Thus y is the unique maximizer of $\hat{V} - \varphi$ in $U_{2r}(y)$, which implies (A.1).

B. Analyses of Examples 6.1 and 6.2

B.1 Generalities

This appendix and the next use the verification theorem, Theorem 3.4, to solve the exit and transition problems for the examples from Section 6. We start by introducing a convenient notation for working with symmetric normal form games $A \in \mathbb{R}^{n\times n}$. We use superscripts to refer to rows of A and subscripts to refer to columns. Thus A^i is the ith row of A, A_j is the jth column of A, and A^i_j is the (i,j)th entry. These objects can be obtained by pre- and post-multiplying A by standard basis vectors:

$$A^i = e_i'A, \quad A_j = Ae_j, \quad A^i_j = e_i'Ae_j.$$

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

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

If $x \in \mathcal{B}^i$ and $x + \varepsilon(e_k - e_j) \in \mathcal{B}^i$ for some $\varepsilon > 0$, then (5.20) and (5.10) imply that the cost of motion from x in direction $e_k - e_j$ with $k \neq j$ under logit choice is

(B.1) $$L(x, e_k - e_j) = \langle \Upsilon(Ax), [e_k - e_j]_+ \rangle = \langle A^i x\,\mathbf{1} - Ax, e_k \rangle = A^{i-k}x.$$

Likewise, Lemma B.1 presents two formulas for the costs of paths that move in a basic direction within a single best response region under logit choice.

Lemma B.1. *Let $x \in \mathcal{B}^i$, and let $y = x + d\,(e_k - e_j) \in \mathcal{B}^i$, where $d \geq 0$. Define $\phi : [0,1] \to X$ by $\phi(t) = (1-t)x + ty$.*

(*i*) *The cost of path ϕ is $c(\phi) = \frac{1}{2}dA^{i-k}(y+x)$.*

(*ii*) *If in addition $y \in \mathcal{B}^{ik}$ and $A^{i-k}_{j-k} \neq 0$, then $c(\phi) = \dfrac{1}{2}\dfrac{(A^{i-k}x)^2}{A^{i-k}_{j-k}}$.*

Proof. For part (i), use formulas (5.19) and (5.10) and the fact that segment xy lies in $\mathcal{B}^i$ to compute as follows:

$$c(\phi) = \int_0^1 [\dot{\phi}_t]'_+(\mathbf{1}A^i - A)\phi_t \,\mathrm{d}t = \int_0^1 de'_k(\mathbf{1}A^i - A)\phi_t \,\mathrm{d}t = dA^{i-k}\int_0^1 \phi_t \,\mathrm{d}t = \frac{1}{2}dA^{i-k}(y + x).$$

For part (ii), note that if $y \in \mathcal{B}^{ik} = \mathcal{B}^i \cap \mathcal{B}^k$, then $A^{i-k}y = 0$, and since $y = x + d(e_k - e_j)$ we have $d = A^{i-k}x/A^{i-k}_{j-k}$; thus substitution into (i) yields the result. □

Next, we introduce a concise notation for Hamiltonians arising in our analysis. By Theorem 3.4, verifying that a function W is a solution to a target problem (TP) requires showing that W satisfies the subsolution condition $H(x, -DW(x)) \leq 0$ at almost all $x \in X^\circ$. By definitions (3.3) and (3.4) and Remark 3.1, this condition reduces to

$$\max_{v \in \mathcal{E}} \mathcal{H}(x, -DW(x), v) \leq 0, \text{ where } \mathcal{E} = \{e_j - e_i : i \neq j\} \cup \{\mathbf{0}\}.$$

We therefore introduce the notation

$$H^W(x, v) = \mathcal{H}(x, -DW(x), v) = -DW(x)v - L(x, v).$$

In this notation, checking the subsolution condition at state $x \in X^\circ$ requires showing that

$$\text{(B.2)} \qquad H^W(x, v) \leq 0 \text{ for } v \in \mathcal{E}.$$

This inequality will bind when v is an optimal control at x, and of course when $v = \mathbf{0}$.

B.2 Facts about games of form (6.1)

We proceed with the analysis of exit and transition costs for coordination of games of the form (6.1) under the logit choice rule (5.8). Figures 2 and 3 show the best response regions $\mathcal{B}^1$, $\mathcal{B}^2$, and $\mathcal{B}^3$ for this game. The vertices of these regions besides e_1, e_2, and e_3 are

$$\text{(B.3)} \qquad \begin{aligned} x^{12} &= (\tfrac{b}{a+b}, \tfrac{a}{a+b}, 0), & x^{13} &= (\tfrac{c}{a+c}, 0, \tfrac{a}{a+c}), \\ y^{12} &= (0, \tfrac{2a+b}{2(a+b)}, \tfrac{b}{2(a+b)}), & y^{13} &= (0, \tfrac{c}{2(a+c)}, \tfrac{2a+c}{2(a+c)}). \end{aligned}$$

We also note these facts computed from (6.1) for later use:

$$\text{(B.4)} \qquad \begin{aligned} &A^{1-2} = (a \;\; -b \;\; 2a+b), \;\; A^{1-3} = (a \;\; 2a+c \;\; -c), \;\; A^{2-3} = (0 \;\; 2a+b+c \;\; -(2a+b+c)), \\ &A^{1-2}_{1-2} = a+b, \quad A^{1-2}_{3-1} = a+b, \quad A^{1-2}_{3-2} = 2(a+b), \\ &A^{1-3}_{1-3} = a+c, \quad A^{1-3}_{2-1} = a+c, \quad A^{1-3}_{2-3} = 2(a+c), \\ &A^{2-3}_{1-3} = 2a+b+c, \quad A^{2-3}_{2-1} = 2a+b+c, \quad A^{2-3}_{2-3} = 4a+2b+2c. \end{aligned}$$

B.3 Exit from state e_1

Here we verify that the optimal value function W to the target problem associated with the exit problem for state e_1 is that induced by the feedback control in Figure 2. At states in $\mathcal{B}^2 \cup \mathcal{B}^3$, $W = 0$. To define W at states in $\mathcal{B}^1$, let

$$\text{(B.5)} \qquad W^2(x) = \frac{1}{2}\frac{(A^{1-2}z(x))^2}{A^{1-2}_{1-2}}, \text{ where } z(x) = x + x_3(e_1 - e_3), \text{ and}$$

$$W^3(x) = \frac{1}{2}\frac{(A^{1-3}y(x))^2}{A^{1-3}_{1-3}}, \text{ where } y(x) = x + x_2(e_1 - e_2).$$

For $x \in \mathcal{B}^1$ with $x_2 \le x_2^{12}$, consider a path that first proceeds from x in direction $e_1 - e_3$ until reaching boundary state $z(x)$, and then proceeds southwest along the face of X to state x^{12}. The first segment of this path has cost 0, and Lemma B.1(ii) implies that second segment has cost $W^2(x)$. Likewise, for $x \in \mathcal{B}^1$ with $x_3 \le x_3^{13}$, consider a path that first proceeds from x in direction $e_1 - e_2$ until reaching boundary state $y(x)$, and then proceeds southeast along the face of X to state x^{13}. This path has cost $W^3(x)$.

The line in $\text{aff}(X)$ on which $W^2(x) = W^3(x)$ is defined by the linear equation

$$\text{(B.6)} \qquad \frac{(x_1 + x_3)a - bx_2}{\sqrt{a+b}} = \frac{(x_1 + x_2)a - cx_3}{\sqrt{a+c}}.$$

To show that this line crosses the $x_2 = 0$ face of X, set $x_2 = 0$ and $x_3 = 1 - x_1$ and solve (B.6) for x_1 to obtain

$$\hat{x}_1 = \frac{a}{\sqrt{(a+b)(a+c)}} + \frac{c}{a+c} \in (0, 1),$$

where $\hat{x}_1 < 1$ because $b > c$ (by assumption (6.1)). This calculation defines state $\hat{x}$, which appears on the e_1e_3 boundary of X as shown in Figure 2. Equation (B.6) is also satisfied at

$$\hat{y} = (1 - x_2^{12} - x_3^{13}, x_2^{12}, x_3^{13}).$$

If $\hat{y}_1 \ge 0$, then $\hat{y} \in X$ is as shown in Figure 2; otherwise, $\hat{y}$ is below the e_2e_3 boundary. Our analysis covers both cases.

Now define $R^2 = \text{conv}(\{e_1, x^{12}, \hat{y}, \hat{x}\}) \cap X$ and $R^3 = \text{conv}(\{x^{13}, \hat{y}, \hat{x}\}) \cap X$. Then the candidate value function $W : X \to \mathbb{R}$ is defined by

$$\text{(B.7)} \qquad W(x) = \begin{cases} W^2(x) & \text{if } x \in R^2, \\ W^3(x) & \text{if } x \in R^3, \\ 0 & \text{otherwise.} \end{cases}$$

It is clear from the foregoing that W is Lipschitz continuous, that $W(e_2) = W(e_3) = 0$, and that all values $W(x)$ can be achieved using the paths shown in Figure 2 along with zero-cost paths. It remains to check the subsolution condition (B.2). To verify this condition in int(R^2), observe from (B.7) and (B.5) that

$$W(x) = \frac{1}{2}\frac{(A^{1-2}z(x))^2}{A^{1-2}_{1-2}} \qquad \text{(where } z(x) = x + x_3(e_1 - e_3)\text{)} \tag{B.8}$$

$$= \frac{1}{2}\frac{(A^{1-2}x + A^{1-2}_{1-3}e_3'x)^2}{A^{1-2}_{1-2}}.$$

It follows that

$$DW(x) = \frac{1}{A^{1-2}_{1-2}}A^{1-2}z(x)(A^{1-2} + A^{1-2}_{1-3}e_3'), \tag{B.9}$$

and hence, by (B.4), that

$$\begin{aligned} DW(x)(e_1 - e_3) &= \frac{1}{A^{1-2}_{1-2}}A^{1-2}z(x)(A^{1-2} + A^{1-2}_{1-3}e_3')(e_1 - e_3) = 0, \\ DW(x)(e_1 - e_2) &= \frac{1}{A^{1-2}_{1-2}}A^{1-2}z(x)(A^{1-2} + A^{1-2}_{1-3}e_3')(e_1 - e_2) = A^{1-2}z(x), \\ DW(x)(e_3 - e_2) &= \frac{1}{A^{1-2}_{1-2}}A^{1-2}z(x)(A^{1-2} + A^{1-2}_{1-3}e_3')(e_3 - e_2) = A^{1-2}z(x). \end{aligned} \tag{B.10}$$

Using these facts and those from (B.1), (B.3), and (B.4), we verify (B.2) in int(R^2) as follows:

$$\begin{aligned} H^W(x, e_1 - e_3) &= -L(x, e_1 - e_3) - DW(x)(e_1 - e_3) \\ &= 0, \\ H^W(x, e_3 - e_1) &= -L(x, e_3 - e_1) - DW(x)(e_3 - e_1) \\ &= -A^{1-3}x + 0 \\ &\leq 0 \qquad \text{(since } x \in \mathcal{B}^1\text{);} \\ H^W(x, e_1 - e_2) &= -L(x, e_1 - e_2) - DW(x)(e_1 - e_2) \\ &= 0 - A^{1-2}z(x) \\ &\leq 0 \qquad \text{(since } z(x) \in \mathcal{B}^1\text{),} \\ H^W(x, e_2 - e_1) &= -L(x, e_2 - e_1) - DW(x)(e_2 - e_1) \\ &= -A^{1-2}x + A^{1-2}z(x) \\ &= -A^{1-2}x + A^{1-2}x + x_3A^{1-2}_{1-3} \\ &= -x_3(a + b) \\ &\leq 0; \\ H^W(x, e_3 - e_2) &= -L(x, e_3 - e_2) - DW(x)(e_3 - e_2) \\ &= -A^{1-3}x - A^{1-2}z(x) \\ &\leq 0 \qquad \text{(since } x \in \mathcal{B}^1 \text{ and } z(x) \in \mathcal{B}^1\text{),} \end{aligned}$$

$$
\begin{aligned}
H^W(x, e_2 - e_3) &= -L(x, e_2 - e_3) - DW(x)(e_2 - e_3)\\
&= -A^{1-2}x + A^{1-2}z(x)\\
&= -x_3(a+b)\\
&\leq 0.
\end{aligned}
$$

This completes the verification of the subsolution condition (B.2) in $\mathrm{int}(R^2)$. Very similar calculations establish the subsolution condition in $\mathrm{int}(R^3)$. We thus conclude from Theorem 3.4 that (B.7) is the optimal value function for the exit problem from state e_1.

B.4 Transitions to state e_1

We next verify that the optimal value function for the transition problem to state e_1 corresponds to the feedback control in Figure 3(i). By Lemma B.1(ii), the candidate value function corresponding to this figure is

$$
\text{(B.11)} \qquad W(x) = \begin{cases} W^2(x) \equiv \dfrac{1}{2}\dfrac{(A^{2-1}x)^2}{A^{2-1}_{2-1}} & \text{if } x \in \mathcal{B}^2,\\ W^3(x) \equiv \dfrac{1}{2}\dfrac{(A^{3-1}x)^2}{A^{3-1}_{3-1}} & \text{if } x \in \mathcal{B}^3,\\ 0 & \text{if } x \in \mathcal{B}^1. \end{cases}
$$

It is clear that W is Lipschitz continuous, that $W(e_1) = 0$, and that each value $W(x)$ is achieved by a suitable choice of path.

It remains to check the subsolution condition (B.2). Here we establish (B.2) in $\mathrm{int}(\mathcal{B}^2)$. For x in this set we have

$$
DW(x) = \frac{A^{2-1}x}{A^{2-1}_{2-1}} A^{2-1},
$$

and hence, by (B.4),

$$
\begin{aligned}
DW(x)(e_1 - e_2) &= \frac{A^{2-1}x}{A^{2-1}_{2-1}} A^{2-1}_{1-2} = -A^{2-1}x,\\
DW(x)(e_1 - e_3) &= \frac{A^{2-1}x}{A^{2-1}_{2-1}} A^{2-1}_{1-3} = A^{2-1}x,\\
DW(x)(e_3 - e_2) &= \frac{A^{2-1}x}{A^{2-1}_{2-1}} A^{2-1}_{3-2} = -2A^{2-1}x.
\end{aligned}
$$

Combining these with (B.1), (B.3), and (B.4), along with the standing assumptions $a + c \geq b > c > 0$ from (6.1), we verify (B.2) as follows:

$$
H^W(x, e_1 - e_2) = -L(x, e_1 - e_2) - DW(x)(e_1 - e_2)
$$

$$
\begin{aligned}
&= -A^{2-1}x + A^{2-1}x \\
&= 0, \\
H^W(x, e_2 - e_1) &= -L(x, e_2 - e_1) - DW(x)(e_2 - e_1) \\
&= 0 - A^{2-1}x \\
&\leq 0 && \text{(since } x \in \mathcal{B}^2\text{);} \\
H^W(x, e_1 - e_3) &= -L(x, e_1 - e_3) - DW(x)(e_1 - e_3) \\
&= -A^{2-1}x - A^{2-1}x \\
&\leq 0 && \text{(since } x \in \mathcal{B}^2\text{);} \\
H^W(x, e_3 - e_1) &= -L(x, e_3 - e_1) - DW(x)(e_3 - e_1) \\
&= -A^{2-3}x + A^{2-1}x \\
&= -(ax_1 + (2a + c)x_2 - cx_3) \\
&= -a - (a + c)(x_2 - x_3) && \text{(since } x \in X\text{)} \\
&\leq 0 && \text{(since } x_2 \geq x_3 \text{ for } x \in \mathcal{B}^2\text{);} \\
H^W(x, e_3 - e_2) &= -L(x, e_3 - e_2) - DW(x)(e_3 - e_2) \\
&= -A^{2-3}x + 2A^{2-1}x \\
&= -2ax_1 - (2a - b + c)x_2 - (2a + b - c)x_3 \\
&\leq -2ax_1 - ax_2 - 2ax_3 && \text{(since } a + c \geq b \text{ and } b > c\text{),} \\
&\leq 0 \\
H^W(x, e_2 - e_3) &= -L(x, e_2 - e_3) - DW(x)(e_2 - e_3) \\
&= 0 - 2A^{2-1}x \\
&\leq 0 && \text{(since } x \in \mathcal{B}^2\text{).}
\end{aligned}
$$

Similar calculations show that the subsolution condition (B.2) holds in int($\mathcal{B}^3$), with the main novelty being that the assumption that $a + c \geq b$ is not needed. We thus conclude from Theorem 3.4 that (B.11) is the optimal value function for the transition problem to state e_1.

B.5 Transitions to state e_2

We now the optimal value function for the transition problem to state e_2. Figure 3(ii) illustrates the optimal feedback control for cases in which $x_2^{12} \geq y_2^{13}$, which we analyze explicitly here.[13]

In region $R = \text{conv}(\{x^{12}, e_1, x^{13}, y^{13}, y^*\})$, the costs of the paths illustrated in Figure 3(ii) are given by equation (B.8) in Section B.3. In region $\mathcal{B}^3$, the path cost is the sum of the cost of reaching $\mathcal{B}^{23}$, which we compute using Lemma B.1(ii), and the cost of proceeding from

[13] If $x_2^{12} < y_2^{13}$, then there is a region northeast of y^{13} from which paths moving in direction $e_1 - e_3$ hit $\mathcal{B}^2$ rather than segment $x^{12}e_1$. In this region, the subsolution condition (B.2) follows from the analysis of the transition problem to state e_1.

there to x^{12}, given by (B.8). Combining these cases yields the candidate value function

$$W(x) = \begin{cases} \dfrac{1}{2}\dfrac{(A^{1-2}z(x))^2}{A^{1-2}_{1-2}} & \text{if } x \in R, \\ \dfrac{1}{2}\dfrac{(A^{3-1}x)^2}{A^{3-1}_{3-1}} + \dfrac{1}{2}\dfrac{((A^{1-2}z(x))^2}{A^{1-2}_{1-2}} & \text{if } x \in \mathcal{B}^3, \\ 0 & \text{otherwise,} \end{cases} \tag{B.12}$$

where $z(x) = x + x_3(e_1 - e_3)$.

As in the previous analyses, we need only check the subsolution condition (B.2). In fact, the analysis in Section B.3 does so for states in region R, so we need only verify the condition for states in $\text{int}(\mathcal{B}^3)$.

To do so, observe that for x in this region, $DW(x)$ is the sum of expression (B.9) and

$$D\left(\frac{1}{2}\frac{(A^{3-1}x)^2}{A^{3-1}_{3-1}}\right) = \frac{A^{3-1}x}{A^{3-1}_{3-1}}A^{3-1}.$$

We can thus take advantage of displays (B.10) and (B.4) to compute the directional derivatives of DW at x:

$$\begin{aligned} DW(x)(e_1 - e_3) &= \frac{A^{3-1}x}{A^{3-1}_{3-1}}A^{3-1}_{1-3} + 0 = -A^{3-1}x, \\ DW(x)(e_1 - e_2) &= \frac{A^{3-1}x}{A^{3-1}_{3-1}}A^{3-1}_{1-2} + A^{1-2}z(x) = A^{3-1}x + A^{1-2}z(x), \\ DW(x)(e_2 - e_3) &= \frac{A^{3-1}x}{A^{3-1}_{3-1}}A^{3-1}_{2-3} - A^{1-2}z(x) = -2A^{3-1}x - A^{1-2}z(x). \end{aligned}$$

With these expressions in hand, we can evaluate the conditions on the Hamiltonian:

$$\begin{aligned} H^W(x, e_1 - e_3) &= -L(x, e_1 - e_3) - DW(x)(e_1 - e_3) \\ &= -A^{3-1}x + A^{3-1}x \\ &= 0, \\ H^W(x, e_3 - e_1) &= -L(x, e_3 - e_1) - DW(x)(e_3 - e_1) \\ &= 0 - A^{3-1}x \\ &\leq 0 \qquad \text{(since } x \in \mathcal{B}^3\text{)}, \\ H^W(x, e_1 - e_2) &= -L(x, e_1 - e_2) - DW(x)(e_1 - e_2) \\ &= -A^{3-1}x - A^{3-1}x - A^{1-2}z(x) \\ &\leq 0 \qquad \text{(since } x \in \mathcal{B}^3 \text{ and } z(x) \in \mathcal{B}^1\text{);} \\ H^W(x, e_2 - e_1) &= -L(x, e_2 - e_1) - DW(x)(e_2 - e_1) \\ &= -A^{3-2}x + A^{3-1}x + A^{1-2}z(x) \end{aligned}$$

$$
\begin{aligned}
&= -A^{1-2}(x - z(x)) \\
&= -A^{1-2}_{3-1}x_3 \\
&= -(a+b)x_3 \\
&\leq 0; \\
H^W(x, e_2 - e_3) &= -L(x, e_2 - e_3) - DW(x)(e_2 - e_3) \\
&= -A^{3-2}x + 2A^{3-1}x + A^{1-2}z(x) \\
&= -A^{1-2}(x - z(x)) + A^{3-1}x \\
&= -A^{1-2}_{3-1}x_3 + A^{3-1}x \\
&= -(a+b)x_3 - ax_1 - (2a+c)x_2 + cx_3 \\
&= -ax_1 - (2a+c)x_2 - (a+b-c)x_3 \\
&\leq 0 \qquad\qquad \text{(since } b > c\text{)}, \\
H^W(x, e_3 - e_2) &= -L(x, e_3 - e_2) - DW(x)(e_3 - e_2) \\
&= 0 - 2A^{3-1}x - A^{1-2}z(x) \\
&\leq 0 \qquad\qquad \text{(since } x \in \mathcal{B}^3 \text{ and } z(x) \in \mathcal{B}^1\text{)}.
\end{aligned}
$$

We conclude from Theorem 3.4 that (B.12) is the optimal value function for the transition problem to state e_2.

C. Analysis of Example 6.3

In this section we verify that the optimal value function for the transition problem to e_3 in potential games of the form (6.4) under the logit rule (5.8) corresponds to the feedback controls illustrated in Figure 4.

We start with some facts about potential games of the form (6.4). Clearly such games are coordination games (5.4), and it is easy to verify that they satisfy the marginal bandwagon property (5.26). The labeled points in Figure 4 are

$$
\text{(C.1)} \qquad
\begin{aligned}
x^{12} &= (\tfrac{b}{a+b}, \tfrac{a}{a+b}, 0), & x^{13} &= (\tfrac{c}{a+c}, 0, \tfrac{a}{a+c}), \\
y^{12} &= (0, \tfrac{d}{b+d}, \tfrac{b}{b+d}), & y^{13} &= (0, \tfrac{c}{c+d}, \tfrac{d}{c+d}), \\
\tilde{y}^{12} &= (\tfrac{b}{b+d}, \tfrac{d}{b+d}, 0), & \tilde{y}^{13} &= (\tfrac{d}{c+d}, \tfrac{c}{c+d}, 0).
\end{aligned}
$$

One can verify that $\mathcal{B}^2 = \text{conv}(\{e_2, x^{12}, y^{12}\})$ and that $\mathcal{B}^3 = \text{conv}(\{e_3, x^{13}, y^{13}\})$. Finally, the unique point in $\text{aff}(X)$ at which all three actions receive equal payoffs is

$$
\text{(C.2)} \qquad x^* = \frac{1}{bc - d^2 + ab + ac + 2ad}(bc - d^2, a(c+d), a(b+d)).
$$

If drawn in Figure 4, x^* is the point below boundary e_2e_3 where rays $x^{12}y^{12}$ and $x^{13}y^{13}$ meet; in particular, (C.2) and assumption (6.4) imply that $x^*_1 < 0$.

Define $\mathcal{B}^1_1 = \text{conv}(\{e_1, x^{13}, y^{13}, \tilde{y}^{13}\})$, $\mathcal{B}^1_2 = \text{conv}(\{\tilde{y}^{13}, \tilde{y}^{13}, y^{12}, \tilde{y}^{12}\})$, and $\mathcal{B}^1_3 = \text{conv}(\{\tilde{y}^{12}, y^{12}, x^{12}\})$.

The value function corresponding to the controls in Figure 4 is

$$W(x) = \begin{cases} \dfrac{1}{2}\dfrac{(A^{1-3}x)^2}{A^{1-3}_{1-3}} & \text{if } x \in \mathcal{B}^1_1, \\ \left(A^{1-3}x x_1 - \dfrac{1}{2}A^{1-3}_{1-3}(x_1)^2\right) + \dfrac{1}{2}\dfrac{(A^{1-3}y(x))^2}{A^{1-3}_{2-3}} & \text{if } x \in \mathcal{B}^1_2, \\ \left(\dfrac{1}{2}x'Ax - \dfrac{1}{2}(y^{12})'Ay^{12}\right) + \dfrac{1}{2}\dfrac{(A^{1-3}y^{12})^2}{A^{1-3}_{2-3}} & \text{if } x \in \mathcal{B}^1_3 \cup \mathcal{B}^2, \\ 0 & \text{if } x \in \mathcal{B}^3. \end{cases} \tag{C.3}$$

$$\text{where } y(x) = x + x_1(e_3 - e_1). \tag{C.4}$$

For $x \in \mathcal{B}^1_1$, the formula follows from Lemma B.1(ii). For $x \in \mathcal{B}^1_2$, the path proceeds from x in direction $e_3 - e_1$ to the boundary segment e_2e_3, and then along this segment to y^{13}; the cost of the first piece is obtained from Lemma B.1(i), and the cost of the second piece from Lemma B.1(ii). The formula when $x \in \mathcal{B}^1_3 \cup \mathcal{B}^2$ will be derived below.

We first verify the subsolution condition (B.2) in $\text{int}(\mathcal{B}^1_1)$. The first four inequalities use the facts that $x \in \mathcal{B}^1$, that A is a coordination game (5.4), and that A satisfies the marginal bandwagon property (5.26).

$$H^W(x, e_3 - e_1) = -L(x, e_3 - e_1) - DW(x)(e_3 - e_1) = -A^{1-3}x - \frac{A^{1-3}x}{A^{1-3}_{1-3}}A^{1-3}_{3-1} = 0,$$

$$H^W(x, e_1 - e_3) = -L(x, e_1 - e_3) - DW(x)(e_1 - e_3) = 0 - \frac{A^{1-3}x}{A^{1-3}_{1-3}}A^{1-3}_{1-3} = -A^{1-3}x \le 0,$$

$$\begin{aligned} H^W(x, e_3 - e_2) &= -L(x, e_3 - e_2) - DW(x)(e_3 - e_2) \\ &= -A^{1-3}x - \frac{A^{1-3}x}{A^{1-3}_{1-3}}A^{1-3}_{3-2} \\ &= -A^{1-3}x\frac{A^{1-3}_{1-2}}{A^{1-3}_{1-3}} \le 0, \end{aligned}$$

$$H^W(x, e_2 - e_3) = -L(x, e_2 - e_3) - DW(x)(e_2 - e_3) = -A^{1-2}x - \frac{A^{1-3}x}{A^{1-3}_{1-3}}A^{1-3}_{2-3} \le 0,$$

$$H^W(x, e_1 - e_2) = -L(x, e_1 - e_2) - DW(x)(e_1 - e_2) = 0 - \frac{A^{1-3}x}{A^{1-3}_{1-3}}A^{1-3}_{1-2} \le 0,$$

$$\begin{aligned} H^W(x, e_2 - e_1) &= -L(x, e_2 - e_1) - DW(x)(e_2 - e_1) \\ &= -A^{1-2}x - \frac{A^{1-3}x}{A^{1-3}_{1-3}}A^{1-3}_{2-1} \\ &= -(ax_1 - bx_2 + dx_3) + \left(\frac{ax_1 + dx_2 - cx_3}{a + c}\right)(a - d) \end{aligned}$$

$$
\begin{aligned}
&= \frac{-a(c+d)(x_1+x_3)-(d^2-ab-bc-ad)x_2}{a+c} \\
&= \frac{-a(c+d)+(ab+bc+ca+2ad-d^2)x_2}{a+c} \\
&\leq 0 \qquad \text{(since } x_2 < x_2^* \text{ for } x \in \mathcal{B}_1^1\text{)}.
\end{aligned}
$$

Verifying the subsolution condition (B.2) in $\text{int}(\mathcal{B}_2^1)$ requires some preliminary calculations. Use (C.4) to rewrite the second case of (C.3) as

$$
W(x) = A^{1-3}x(e_1'x) - \frac{1}{2}A_{1-3}^{1-3}(e_1'x)^2 + \frac{1}{2}\frac{(A^{1-3}x - A_{1-3}^{1-3}e_1'x)^2}{A_{2-3}^{1-3}}.
$$

From this we compute the derivative of W as

$$
DW(x) = A^{1-3}xe_1' + x_1A^{1-3} - A_{1-3}^{1-3}x_1e_1' + \frac{1}{A_{2-3}^{1-3}}A^{1-3}y(x)(A^{1-3} - A_{1-3}^{1-3}e_1').
$$

Thus

$$
\begin{aligned}
DW(x)(e_1-e_3) &= A^{1-3}x, \\
DW(x)(e_1-e_2) &= A^{1-3}x + A_{1-2}^{1-3}x_1 - A_{1-3}^{1-3}x_1 + \frac{1}{A_{2-3}^{1-3}}A^{1-3}y(x)(A_{1-2}^{1-3} - A_{1-3}^{1-3}) \\
&= A^{1-3}x + A_{1-2}^{1-3}x_1 - A_{1-3}^{1-3}x_1 - A^{1-3}(x - x_1(e_1-e_3)) \\
&= A_{1-2}^{1-3}x_1, \\
DW(x)(e_3-e_2) &= -A_{2-3}^{1-3}x_1 - A^{1-3}y(x) \\
&= -A_{2-3}^{1-3}x_1 - (A^{1-3}x - A_{1-3}^{1-3}x_1) \\
&= A_{1-2}^{1-3}x_1 - A^{1-3}x.
\end{aligned}
$$

Using these formulas, we check the subsolution condition (B.2), repeatedly using the facts that $x \in \mathcal{B}^1$ and that A satisfies the marginal bandwagon property (5.26).

$$
\begin{aligned}
H^W(x, e_3-e_1) &= -L(x, e_3-e_1) - DW(x)(e_3-e_1) = -A^{1-3}x + A^{1-3}x = 0, \\
H^W(x, e_1-e_3) &= -L(x, e_1-e_3) - DW(x)(e_1-e_3) = 0 - A^{1-3}x \leq 0, \\
H^W(x, e_1-e_2) &= -L(x, e_1-e_2) - DW(x)(e_1-e_2) = 0 - x_1A_{1-2}^{1-3} \leq 0, \\
H^W(x, e_2-e_1) &= -L(x, e_2-e_1) - DW(x)(e_2-e_1) \\
&= -A^{1-2}x + x_1A_{1-2}^{1-3} \\
&= -(ax_1 - bx_2 + dx_3) + x_1(a-d) \\
&= (b+d)x_2 - d \qquad \text{(since } x \in X\text{)} \\
&\leq 0 \qquad \text{(since } x_2 \leq y_2^{12} = d/(b+d) \text{ when } x \in \mathcal{B}_2^1\text{)}, \\
H^W(x, e_3-e_2) &= -L(x, e_3-e_2) - DW(x)(e_3-e_2)
\end{aligned}
$$

$$
\begin{aligned}
&= -A^{1-3}x - x_1 A^{1-3}_{1-2} + A^{1-3}x \\
&= -x_1 A^{1-3}_{1-2} \\
&\leq 0, \\
H^W(x, e_2 - e_3) &= -L(x, e_2 - e_3) - DW(x)(e_2 - e_3) \\
&= -A^{1-2}x + x_1 A^{1-3}_{1-2} + A^{1-3}x \\
&= -A^{1-3}x + H^W(x, e_2 - e_1) \\
&\leq 0.
\end{aligned}
$$

Lastly, we consider region $\mathcal{B}^1_3$; the analysis of $\mathcal{B}^2$ is nearly identical. To derive the formula for $W(x)$ from the third case of (C.3), note that second term is just the cost of the boundary path from y^{12} to y^{13} (cf. the second case of (C.3)). For the first term, let $x \in \mathcal{B}^1_3$, and let $\phi : [0,1] \to X$ be any path from x to y^{12} that is always in $\mathcal{B}^1_3$ and that satisfies $[\dot{\phi}(t)]_- = x_1 e_1$ for all $t \in [0,1]$. These are precisely the paths through $\mathcal{B}^1_3$ shown in Figure 4. Since $[\dot{\phi}(t)]'_-(\mathbf{1}A^1 - A) = x_1 e'_1(\mathbf{1}A^1 - A) = \mathbf{0}'$, the cost of any path ϕ of the form above is

$$
\begin{aligned}
c(\phi) &= \int_0^1 [\dot{\phi}_t]'_+(\mathbf{1}A^i - A)\phi_t \,\mathrm{d}t = \int_0^1 (\dot{\phi}_t)'(\mathbf{1}A^i - A)\phi_t \,\mathrm{d}t = -\int_0^1 (\dot{\phi}_t)' A\phi_t \,\mathrm{d}t \\
&= \frac{1}{2}x'Ax - \frac{1}{2}(y^{12})'Ay^{12}. \qquad \text{(C.5)}
\end{aligned}
$$

This agrees with the third case of (C.3). Thus the derivative of W in $\mathrm{int}(\mathcal{B}^1_3)$ is

$$
DW(x) = Ax.
$$

Using this fact, we can evaluate the subsolution condition (B.2) in $\mathrm{int}(\mathcal{B}^1_3)$:

$$
\begin{aligned}
H^W(x, e_3 - e_1) &= -L(x, e_3 - e_1) - DW(x)(e_3 - e_1) = -A^{1-3}x - A^{3-1}x = 0, \\
H^W(x, e_2 - e_1) &= -L(x, e_2 - e_1) - DW(x)(e_2 - e_1) = -A^{1-2}x - A^{2-1}x = 0, \\
H^W(x, e_1 - e_3) &= -L(x, e_1 - e_3) - DW(x)(e_1 - e_3) = 0 - A^{1-3}x \leq 0, \\
H^W(x, e_1 - e_2) &= -L(x, e_1 - e_2) - DW(x)(e_1 - e_2) = 0 - A^{1-2}x \leq 0, \\
H^W(x, e_3 - e_2) &= -L(x, e_3 - e_2) - DW(x)(e_3 - e_2) = -A^{1-3}x - A^{3-2}x = -A^{1-2}x \leq 0, \\
H^W(x, e_2 - e_3) &= -L(x, e_2 - e_3) - DW(x)(e_2 - e_3) = -A^{1-2}x - A^{2-3}x = -A^{1-3}x \leq 0.
\end{aligned}
$$

We thus conclude from Theorem 3.4 that (C.3) is the optimal value function for the transition problem to state e_3.

References

Anderson, S. P., de Palma, A., and Thisse, J.-F. (1992). *Discrete Choice Theory of Product Differentiation*. MIT Press, Cambridge.

Arigapudi, S. (2018). Exit from equilibrium in coordination games under probit choice. Unpublished manuscript, University of Wisconsin.

Armstrong, S. N. and Tran, H. V. (2015). Viscosity solutions of general viscous Hamilton-Jacobi equations. *Math. Ann.*, 361(3-4):647–687.

Blume, L. E. (1997). Population games. In Arthur, W. B., Durlauf, S. N., and Lane, D. A., editors, *The Economy as an Evolving Complex System II*, pages 425–460. Addison-Wesley, Reading, MA.

Boltyanskii, V. G. (1966). Sufficient conditions for optimality and the justification of the dynamic programming method. *SIAM Journal on Control*, 4:326–361.

Capuzzo-Dolcetta, I. and Lions, P.-L. (1990). Hamilton-Jacobi equations with state constraints. *Trans. Amer. Math. Soc.*, 318(2):643–683.

Catoni, O. (1999). Simulated annealing algorithms and Markov chains with rare transitions. In Azéma, J., Émery, M., Ledoux, M., and Yor, M., editors, *Séminaire de Probabilités XXXIII*, pages 69–119. Springer, Berlin.

Crandall, M. G., Evans, L. C., and Lions, P.-L. (1984). Some properties of viscosity solutions of Hamilton-Jacobi equations. *Trans. Amer. Math. Soc.*, 282(2):487–502.

Dembo, A. and Zeitouni, O. (1998). *Large Deviations Techniques and Applications*. Springer, New York, second edition.

Dokumacı E. and Sandholm, W. H. (2011). Large deviations and multinomial probit choice. *Journal of Economic Theory*, 146:2151–2158.

Foster, D. P. and Young, H. P. (1990). Stochastic evolutionary game dynamics. *Theoretical Population Biology*, 38:219–232. Corrigendum, 51 (1997), 77-78.

Freidlin, M. I. and Wentzell, A. D. (1984). *Random Perturbations of Dynamical Systems*. Springer, New York.

Kandori, M., Mailath, G. J., and Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica*, 61:29–56.

Kandori, M. and Rob, R. (1998). Bandwagon effects and long run technology choice. *Games and Economic Behavior*, 22:84–120.

Le, N. Q., Mitake, H., and Tran, H. V. (2017). *Dynamical and geometric aspects of Hamilton-Jacobi and linearized Monge-Ampère equations—VIASM 2016*, volume 2183 of *Lecture Notes in Mathematics*. Springer, Cham. Edited by Mitake and Tran.

Mitake, H. (2008). Asymptotic solutions of Hamilton-Jacobi equations with state constraints. *Appl. Math. Optim.*, 58(3):393–410.

Monderer, D. and Shapley, L. S. (1996). Potential games. *Games and Economic Behavior*, 14:124–143.

Myatt, D. P. and Wallace, C. C. (2003). A multinomial probit model of stochastic evolution. *Journal of Economic Theory*, 113:286–301.

Piccoli, B. and Sussmann, H. J. (2000). Regular synthesis and sufficiency conditions for optimality. *SIAM Journal on Control and Optimization*, 39:359–410.

Sandholm, W. H. (2010a). Decompositions and potentials for normal form games. *Games and Economic Behavior*, 70:446–456.

Sandholm, W. H. (2010b). *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge.

Sandholm, W. H. and Staudigl, M. (2016). Large deviations and stochastic stability in the small noise double limit. *Theoretical Economics*, 11:279–355.

Sandholm, W. H. and Staudigl, M. (2018). Sample path large deviations for stochastic evolutionary game dynamics. *Mathematics of Operations Research*, forthcoming.

Sandholm, W. H. and Staudigl, M. (2019). Large deviations and stochastic stability in the large population limit. In preparation.

Schättler, H. and Ledzewicz, U. (2012). *Geometric Optimal Control: Theory, Methods, and Examples*. Springer, New York.

Soner, H. M. (1986). Optimal control with state-space constraint. I. *SIAM J. Control Optim.*, 24(3):552–561.

Young, H. P. (1993). The evolution of conventions. *Econometrica*, 61:57–84.

Young, H. P. (1998). *Individual Strategy and Social Structure*. Princeton University Press, Princeton.