

Choice Paralysis in Evolutionary Games

Brendon G. Anderson*

Abstract—In this paper, we consider finite-strategy approximations of infinite-strategy evolutionary games. We prove that such approximations converge to the true dynamics over finite-time intervals, under mild regularity conditions which are satisfied by classical examples, e.g., the replicator dynamics. We identify and formalize novel characteristics in evolutionary games: *choice mobility*, and its complement *choice paralysis*. Choice mobility is shown to be a key sufficient condition for the long-time limiting behavior of finite-strategy approximations to coincide with that of the true infinite-strategy game. An illustrative example is constructed to showcase how choice paralysis may lead to the infinite-strategy game getting “stuck,” even though every finite approximation converges to equilibrium.

I. INTRODUCTION

Evolutionary games are models in which a (typically large) number of players strategically interact and revise their strategies, in order to maximize their own expected payoffs. Such models arise in a variety of applications ranging from traffic congestion to economic markets [27]. Extensive rigorous analyses of these models have been developed in the case that the players have access to finitely many strategies, with a particular focus on (often control-theoretic) stability guarantees ensuring that the population of players converges to a (Nash) equilibrium state [3, 12, 16, 20]. However, many realistic games have a continuum of choices available to players, e.g., in reinforcement learning [21], optimization [1], power system pricing and generation [24], and games of timing [4]. The study of such infinite-strategy evolutionary games becomes significantly more challenging, as the dynamics amount to an infinite-dimensional differential equation in the space of probability measures [2].

A handful of works have analyzed infinite-strategy models with specific dynamics—most commonly the replicator dynamics [6, 7, 9, 10, 14, 15, 22, 23], but also the Brown-von Neumann-Nash, pairwise comparison, logit, and perturbed best response dynamics [8, 17–19], among others. However, even such special cases of infinite-strategy evolutionary games cannot be explicitly solved or numerically simulated, requiring the use of finite-strategy approximations as a tractable surrogate. Unfortunately, such approximations are in general unreliable, as the finite-dimensional dynamics may exhibit asymptotic stability, even though the true infinite-strategy dynamics are unstable [2, 23].

In this paper, we study when finite-strategy approximations of infinite-strategy evolutionary dynamics can and cannot be trusted, both over finite-time intervals and in the long-time limit (e.g., convergence to equilibria). A few works have considered finite-strategy approximations for

computing Nash equilibria of *static* infinite-strategy games [13, 26]. Though, to the best of our knowledge, the only work to consider the *dynamics* of finite approximations of evolutionary games is Oechssler and Riedel [23], who show that approximations of the replicator dynamics with compact real number interval strategy sets weakly converge to the true dynamics over finite-time intervals. Our work studies the problem for more general dynamics and strategy sets—recovering Oechssler and Riedel [23, Theorem 4] as a special case of our Theorem 1—and we provide novel characterizations of approximations in the long-time limit. Specifically, our primary contributions are as follows.

A. Contributions

- 1) We prove that, under mild regularity conditions, the trajectories of finite-dimensional approximations weakly converge to the true infinite-dimensional evolutionary dynamics on compact time intervals.
- 2) We introduce the key notion of *choice mobility*, and its complement *choice paralysis*.
- 3) We prove that choice mobility is sufficient to ensure coincidence between the long-time limiting behavior of finite approximations and the infinite-dimensional dynamics (again under mild regularity assumptions).
- 4) We show that choice paralysis arises if strategy switching rates decrease as more strategies are introduced.
- 5) We construct an example demonstrating that choice paralysis may manifest in convergence of every finite approximation to equilibrium, despite the infinite-strategy game being stuck at a different state.

Proofs are deferred to the appendix.

II. INFINITE-STRATEGY EVOLUTIONARY GAMES

We consider a game in which a large population of players choose strategies from an infinite compact metric space S with metric d_S . The sets of Borel probability measures and continuous real-valued functions on S are denoted by $\mathcal{P}(S)$ and $C(S)$, respectively. The Borel σ -algebra on S is $\mathcal{B}(S)$. The *state* of the population is encoded by a probability measure, $\mu \in \mathcal{P}(S)$, over the strategy set S , representing the frequency with which each strategy is played. The *game*, $F: \mathcal{P}(S) \rightarrow C(S)$, is defined by *average payoffs* $(F(\mu))(s)$ to strategies $s \in S$ when the population is at state μ .

The population state evolves according to an evolutionary dynamics model (EDM) [2, 25] taking the form

$$\dot{\mu}(t) = v(\mu(t), \rho(t)), \quad \rho(t) = F(\mu(t)), \quad (1)$$

where $v: \mathcal{P}(S) \times C(S) \rightarrow T\mathcal{P}(S)$, with $T\mathcal{P}(S)$ being the tangent space of $\mathcal{P}(S)$, defines the dynamics, and ρ encodes

*Mechanical Engineering Department, California Polytechnic State University. Email: bga@calpoly.edu.

the game's payoffs that steer the dynamics. The EDM (1) can be viewed as a state feedback control system with v defining the open-loop dynamics, and F being the controller. The time-derivative $\dot{\mu}(t)$ is defined in the Fréchet sense. Throughout, we assume that (1) admits a unique and weakly continuous solution $\mu: [0, \infty) \rightarrow \mathcal{P}(S)$.

We consider a general class of dynamics, termed *mean dynamics* [8, 27], where the dynamics map v takes the form

$$v(\mu, \rho)(B) = \int_S \int_B \varphi(s, s', \mu, \rho) d\lambda(s') d\mu(s) - \int_S \int_B \varphi(s', s, \mu, \rho) d\mu(s') d\lambda(s) \quad (2)$$

for some reference probability measure $\lambda \in \mathcal{P}(S)$ and some adequately integrable map $\varphi: S \times S \times \mathcal{P}(S) \times C(S) \rightarrow \mathbb{R}_+$, with \mathbb{R}_+ denoting the set of nonnegative real numbers. The map φ is known as the *revision protocol*, as it characterizes the rate at which players switch from one strategy to another, given a particular state and payoff profile of the game. The first term in $v(\mu, \rho)(B)$ represents the “inflow” of mass into the strategies within $B \in \mathcal{B}(S)$, and the second term represents the corresponding “outflow.”

Example 1. Popular examples of mean dynamics include the replicator dynamics [6, 7, 22, 23], the Brown-von Neumann-Nash (BNN) dynamics [17], and the pairwise comparison dynamics [8]. Specifically, the replicator dynamics correspond to setting $\lambda = \mu$ and $\varphi(s, s', \mu, \rho) = \max\{0, \rho(s') - \rho(s)\}$. The BNN and pairwise comparison dynamics correspond to a fixed-in-time λ (typically uniform), and are respectively defined by $\varphi(s, s', \mu, \rho) = \max\{0, \rho(s') - \int_S \rho d\mu\}$ and $\varphi(s, s', \mu, \rho) = \gamma(s, s', \rho)$ for some function γ independent of μ , assumed to satisfy *sign-preservation*: the switch rate is positive if and only if the new strategy has higher payoff than the old strategy, i.e., $\text{sign}(\gamma(s, s', \rho)) = \text{sign}(\max\{0, \rho(s') - \rho(s)\})$ for all $s, s' \in S$ and all $\rho \in C(S)$. The pairwise comparison dynamics are called *impartial* if, for all $s' \in S$, there exists a continuous function $\phi_{s'}: \mathbb{R} \rightarrow \mathbb{R}_+$ satisfying $\gamma(s, s', \rho) = \phi_{s'}(\rho(s') - \rho(s))$ for all $s \in S$ and all $\rho \in C(S)$.

The EDM (1) is an infinite-dimensional differential equation evolving in the space of probability measures, and hence cannot be tractably solved in general. This motivates the problem studied in this paper: introduce and analyze finite approximations, and characterize when such approximations accurately reflect the true behavior of the evolutionary game.

III. FINITE-STRATEGY APPROXIMATIONS

A. Formulation

We now explicitly formulate finite-dimensional approximations of the EDM (1) with the mean dynamics map (2). Consider $n \in \mathbb{N}$, and let $\lambda_n, \mu_n(0) \in \mathcal{P}(S)$ be discrete measures with respect to some reference measure ν on S (e.g., Lebesgue measure when S is a subset of Euclidean space). Suppose that $\text{supp}(\lambda_n) = \text{supp}(\mu_n(0)) = S_n := \{s_1, \dots, s_n\} \subseteq S$, so that $\lambda_n = \sum_{i=1}^n \lambda_n(s_i) \delta_{s_i}$ and $\mu_n(0) = \sum_{i=1}^n (\mu_n(0))(s_i) \delta_{s_i}$. Here, we write $\lambda_n(s_i)$ to mean $\lambda_n(\{s_i\})$ and similarly for other measures evaluated on

singleton sets. Let $\mu_n: [0, \infty) \rightarrow \mathcal{P}(S)$ denote the solution to the EDM with reference measure λ_n and initial state $\mu_n(0)$. The measures λ_n and $\mu_n(t)$ respectively serve as approximations to λ and $\mu(t)$, which may have full support. Writing the mean dynamics for the approximation yields

$$\begin{aligned} (\dot{\mu}_n(t))(B) &= v_n(\mu_n(t), F(\mu_n(t)))(B) \\ &= \int_S \int_B \varphi(s, s', \mu_n(t), F(\mu_n(t))) d\lambda_n(s') d(\mu_n(t))(s) \\ &\quad - \int_S \int_B \varphi(s', s, \mu_n(t), F(\mu_n(t))) d(\mu_n(t))(s') d\lambda_n(s) \end{aligned}$$

for all $B \in \mathcal{B}(S)$. Since $\lambda_n(B) = (\mu_n(0))(B) = 0$ for $B = S \setminus S_n$, we immediately see that $(\dot{\mu}_n(0))(B) = 0$ and hence $(\mu_n(t))(B) = 0$ for all $t \in [0, \infty)$ for this set B . Thus, $\mu_n(t)$ is discrete with $\text{supp}(\mu_n(t)) \subseteq S_n$ for all $t \in [0, \infty)$, implying that it takes the form

$$\mu_n(t) = \sum_{i=1}^n (\mu_n(t))(s_i) \delta_{s_i}. \quad (3)$$

Thus, we may uniquely identify μ_n with the curve $x_n: [0, \infty) \rightarrow \Delta^{n-1}$ defined by $(x_n(t))_i = (\mu_n(t))(s_i)$, where Δ^{n-1} denotes the probability simplex in \mathbb{R}^n .

We define a finite-dimensional game $\hat{F}_n: \Delta^{n-1} \rightarrow \mathbb{R}^n$ by restricting F to discrete measures on S_n :

$$(\hat{F}_n(\theta))_i = F\left(\sum_{j=1}^n \theta_j \delta_{s_j}\right)(s_i), \quad i \in \{1, \dots, n\}.$$

Similarly, we define a finite-dimensional revision protocol $\hat{\varphi}_n: \{1, \dots, n\} \times \{1, \dots, n\} \times \Delta^{n-1} \rightarrow \mathbb{R}_+$ by

$$\hat{\varphi}_n(i, j, \theta) = \varphi\left(s_i, s_j, \sum_{k=1}^n \theta_k \delta_{s_k}, F\left(\sum_{k=1}^n \theta_k \delta_{s_k}\right)\right).$$

Under these constructions, the mean dynamics of the finite approximation are equivalently rewritten as

$$\begin{aligned} (\dot{x}_n(t))_i &= \lambda_n(s_i) \sum_{j=1}^n \hat{\varphi}_n(j, i, x_n(t)) (x_n(t))_j \\ &\quad - (x_n(t))_i \sum_{j=1}^n \hat{\varphi}_n(i, j, x_n(t)) \lambda_n(s_j), \end{aligned} \quad (4)$$

for all $i \in \{1, \dots, n\}$. Note that the approximation (4) is a system of n ordinary differential equations, and hence can be solved efficiently using off-the-shelf numerical solvers.

Example 2. For the replicator dynamics of Example 1, the above constructions give $\lambda_n(s_i) = (x_n(t))_i$ and $\hat{\varphi}_n(i, j, \theta) = \max\{0, (\hat{F}_n(\theta))_j - (\hat{F}_n(\theta))_i\}$, and therefore (4) reduces to

$$(\dot{x}_n(t))_i = (x_n(t))_i \left((\hat{F}_n(x_n(t)))_i - x_n(t)^\top \hat{F}_n(x_n(t)) \right),$$

for all $i \in \{1, \dots, n\}$, which coincides with the classical replicator dynamics on the finite strategy set $\{1, \dots, n\}$ (see Sandholm [27, Example 4.3.1]). Similar recovery of the classical finite-dimensional BNN dynamics [27, Example 4.3.4] and pairwise comparison dynamics [27, Example 4.3.5] are also achieved (specifically, under the fixed uniform reference measure given by $\lambda_n(s_i) = \frac{1}{n}$ for all i).

B. Finite-Time Analysis

In this section, we show that, under some mild regularity conditions on φ , approximations to the EDM (1) weakly converge to the true infinite-dimensional dynamics on compact time intervals. Throughout our analyses, we denote the bounded-Lipschitz metric on $\mathcal{P}(S)$ by $d_{\text{BL}}: (\nu, \eta) \mapsto \sup \{ |\int_S g d\nu - \int_S g d\eta| : g \in \text{Lip}_1(S), \|g\|_\infty \leq 1 \}$, where $\text{Lip}_L(S)$ denotes the set of all L -Lipschitz real-valued functions on S and $\|\cdot\|_\infty$ denotes the supremum norm. Recall that d_{BL} metrizes the weak topology on $\mathcal{P}(S)$ [5, Theorem 8.3.2]. We start by formally stating the regularity conditions on φ :

Assumption 1. The following all hold:

- 1) **Bounded.** There exists $M \in \mathbb{R}_+$ such that $|\varphi(s, s', \mu, F(\mu))| \leq M$ for all $s, s' \in S, \mu \in \mathcal{P}(S)$.
- 2) **Lipschitz on S .** There exist $L_1, L_2 \in \mathbb{R}_+$ such that $\varphi(\cdot, s', \mu, F(\mu)) \in \text{Lip}_{L_1}(S)$ for all $s' \in S, \mu \in \mathcal{P}(S)$, $\varphi(s, \cdot, \mu, F(\mu)) \in \text{Lip}_{L_2}(S)$ for all $s \in S, \mu \in \mathcal{P}(S)$.
- 3) **Lipschitz on $\mathcal{P}(S)$.** There exists $L_3 \in \mathbb{R}_+$ such that, for all $s, s' \in S, \mu, \mu' \in \mathcal{P}(S)$, it holds that $|\varphi(s, s', \mu, F(\mu)) - \varphi(s, s', \mu', F(\mu'))| \leq L_3 d_{\text{BL}}(\mu, \mu')$.

Notice that all conditions in Assumption 1 are with respect to the closed-loop dynamics, i.e., they are required to hold for state-payoff pairs (μ, ρ) satisfying $\rho = F(\mu)$. In general, this is much less stringent than requiring them to hold for all states $\mu \in \mathcal{P}(S)$ and all payoff functions $\rho \in C(S)$.

The replicator, BNN, and impartial pairwise comparison dynamics (recall Example 1) all satisfy Assumption 1 for typical games, and thus the conditions are considered mild:

Proposition 1. Consider either the replicator dynamics, the BNN dynamics, or the impartial pairwise comparison dynamics with $(s', p) \mapsto \phi_{s'}(p)$ Lipschitz with respect to the metric $d_{S \times \mathbb{R}}: ((s', p), (\tilde{s}', \tilde{p})) \mapsto \|(d_S(s', \tilde{s}'), |p - \tilde{p}|\)| where $\|\cdot\|$ denotes an arbitrary norm on \mathbb{R}^2 (e.g., the Smith dynamics [8]). If F is Lipschitz continuous between $(\mathcal{P}(S), d_{\text{BL}})$ and $(C(S), \|\cdot\|_\infty)$ and there exists $L \in \mathbb{R}_+$ such that $F(\mu) \in \text{Lip}_L(S)$ for all $\mu \in \mathcal{P}(S)$, then Assumption 1 holds. In particular, this is the case for the commonly considered linear games F taking the form$

$$F(\mu)(s) = \int_S f(s, s') d\mu(s')$$

with $f: S \times S \rightarrow \mathbb{R}$ bounded and Lipschitz with respect to the metric $d_{S \times S}: ((s, s'), (\tilde{s}, \tilde{s}')) \mapsto \|(d_S(s, \tilde{s}), d_S(s', \tilde{s}'))\|$.

The proof of Proposition 1 involves routine Lipschitz analysis and hence is omitted due to space constraints. We now state our finite-time approximation result.

Theorem 1. Assume Assumption 1. For all $n \in \mathbb{N}$, consider approximations $\lambda_n, \mu_n(0) \in \mathcal{P}(S)$ to λ and $\mu(0)$, with associated solutions μ_n to the mean dynamics EDM (1)–(2). Assume that λ and all λ_n are fixed in time. If $\lambda_n \xrightarrow{n \rightarrow \infty} \lambda$ weakly and $\mu_n(0) \xrightarrow{n \rightarrow \infty} \mu(0)$ weakly, then

$$\lim_{n \rightarrow \infty} \sup_{t \in [0, T]} d_{\text{BL}}(\mu_n(t), \mu(t)) = 0 \text{ for all } T \in (0, \infty). \quad (5)$$

Remark 1. Theorem 1 gives a strong type of convergence: the approximation converges to the true dynamics uniformly on compact time intervals. This implies the pointwise convergence that $\mu_n(t) \xrightarrow{n \rightarrow \infty} \mu(t)$ weakly for all $t \in [0, \infty)$.

Remark 2. Although Theorem 1 holds for general approximating measures $\lambda_n, \mu_n(0) \in \mathcal{P}(S)$ —not only finite ones—it is most useful to consider the finite case, which guarantees that the computationally tractable system of ODEs (4) accurately reflects the true infinite-dimensional dynamics in finite time. We will see in Section III-C that this does *not* necessarily imply accuracy in the long-time limit.

Remark 3. Theorem 1 still holds for time-varying reference measures λ and λ_n , under the additional assumptions that 1) $t \mapsto \lambda(t)$ and all $t \mapsto \lambda_n(t)$ are weakly continuous, and 2) $\lambda_n(t) \xrightarrow{n \rightarrow \infty} \lambda(t)$ weakly for all $t \in [0, \infty)$. The proof becomes more complicated, relying on a generalized version of Grönwall's inequality [11, Theorem 9] and some additional continuity arguments, which we omit for conciseness.

Remark 4. Given a measure $\lambda \in \mathcal{P}(S)$, we may always construct an approximating measure $\lambda_n := \frac{1}{n} \sum_{i=1}^n \delta_{s_i} \in \mathcal{P}(S)$ with $s_1, \dots, s_n \in S$ sampled independently and identically from λ , and in this case, the law of large numbers ensures that the weak convergence assumption in Theorem 1 holds λ -almost surely. The same goes for approximating $\mu(0)$.

C. Long-Time Analysis and Convergence to Equilibria

Although Theorem 1 shows that finite-dimensional approximations accurately reflect the evolutionary game's behavior over any finite-time interval under rather weak regularity conditions, it does *not* automatically imply accurate approximations in the long-time limit. Specifically, Theorem 1 cannot be applied to assess the game's convergence to Nash equilibria using the finite approximations. To rigorously analyze such asymptotic behavior, we identify and introduce the key notion of “choice paralysis.”

Definition 1. Consider a collection $\mathcal{X} = \{x_n : n \in \mathbb{N}\}$ of trajectories $x_n: [0, \infty) \rightarrow \Delta^{n-1}$ such that, for all $n \in \mathbb{N}$, there exists $\bar{x}_n \in \Delta^{n-1}$ such that $\lim_{t \rightarrow \infty} x_n(t) = \bar{x}_n$. The collection \mathcal{X} is said to be *choice-mobile* if

$$\lim_{t \rightarrow \infty} \sup_{n \in \mathbb{N}} \|x_n(t) - \bar{x}_n\|_1 = 0. \quad (6)$$

Otherwise, \mathcal{X} is said to *suffer from choice paralysis*.

Remark 5. The definition of choice-mobility can be extended to collections of finite-dimensional EDMs, such as (4), by requiring all associated trajectories to be choice-mobile.

Remark 6. The ℓ_1 -norm in (6) cannot hastily be replaced with another norm, even if they induce the same topology on \mathbb{R}^n for a fixed value of n , since the tightness of the bounds between equivalent norms on \mathbb{R}^n in general depends on n . Using $\|\cdot\|_1$ induces a natural comparison between approximations of different dimensionality, as $\|x\|_1 = 1$ for all $x \in \Delta^{n-1}$, irrespective of the number of strategies $n \in \mathbb{N}$.

We now consider what choice-mobility represents in a game theoretic context. Mathematically, (6) corresponds to

uniform convergence $x_n(t) \xrightarrow{t \rightarrow \infty} \bar{x}_n$. The uniformity in n requires every finite-dimensional approximation to converge to its associated equilibrium at some speed independent of the number of strategies available. This corresponds to finite-dimensional games that remain sufficiently “mobile” even as more and more strategies become available, i.e., giving the players more strategies to choose from does not significantly slow down the overall dynamics of the game. Contrarily, if a population suffers from choice paralysis, progression towards equilibrium may slow down upon increasing the (finite) number of strategies; the more choices there are, the longer it takes players to figure out what they want to do. If the evolution rate is slowed down enough relative to the convergence rates toward equilibria at a fixed number of strategies, then convergence can no longer be guaranteed as the number of strategies increases towards infinity. From this perspective, it should be expected that the choice mobility condition (6) is key in ensuring that the infinite-dimensional state $\mu(t)$ actually “makes progress” towards an equilibrium. We show that this is indeed the case in Theorem 2.

Theorem 2. Assume Assumption 1. Let $\theta_n, x_n(0) \in \Delta^{n-1}$ for all $n \in \mathbb{N}$, and consider the approximation (4) with $\lambda_n := \sum_{i=1}^n (\theta_n)_i \delta_{s_i}$ and initial condition $x_n(0)$. Assume that λ and all λ_n are fixed in time. If $\lambda_n \xrightarrow{n \rightarrow \infty} \lambda$ weakly and $\sum_{i=1}^n (x_n(0))_i \delta_{s_i} \xrightarrow{n \rightarrow \infty} \mu(0)$ weakly, then

$$\lim_{t \rightarrow \infty} \lim_{n \rightarrow \infty} d_{\text{BL}} \left(\sum_{i=1}^n (\bar{x}_n)_i \delta_{s_i}, \mu(t) \right) = 0 \quad (7)$$

whenever $\mathcal{X} = \{x_n : n \in \mathbb{N}\}$ is choice-mobile with every $\lim_{t \rightarrow \infty} x_n(t) = \bar{x}_n$ for some collection $\{\bar{x}_n \in \Delta^{n-1} : n \in \mathbb{N}\}$. In this case,

- 1) $\mu(t) \xrightarrow{t \rightarrow \infty} \bar{\mu}$ weakly whenever $\sum_{i=1}^n (\bar{x}_n)_i \delta_{s_i} \xrightarrow{n \rightarrow \infty} \bar{\mu}$ weakly for some $\bar{\mu} \in \mathcal{P}(S)$, and
- 2) $\sum_{i=1}^n (\bar{x}_n)_i \delta_{s_i} \xrightarrow{n \rightarrow \infty} \bar{\mu}$ weakly whenever $\mu(t) \xrightarrow{t \rightarrow \infty} \bar{\mu}$ weakly for some $\bar{\mu} \in \mathcal{P}(S)$.

Remark 7. As was the case for Theorem 1, Theorem 2 also still holds for time-varying reference measures λ and λ_n , under the same additional assumptions outlined in Remark 3. The proof of Theorem 2 is essentially unchanged in this case.

At this point, let us dissect Theorem 2. The result (7) ensures that, as the resolution of the approximation (4) becomes finer, the approximation accurately reflects the long-time behavior of the true infinite-strategy game. The two final enumerated results ensure coincidence between an asymptotically stable equilibrium $\bar{\mu}$ of the infinite-dimensional game and the limit of the asymptotically stable equilibria \bar{x}_n of the approximations. In particular, they assert that, when the discrete distribution generated by \bar{x}_n converges towards some limiting distribution $\bar{\mu}$, then $\mu(t)$ must approach this same distribution, and, on the other hand, when $\mu(t)$ approaches some limiting distribution $\bar{\mu}$, the discrete distribution generated by \bar{x}_n must converge towards the same distribution. These consequences show that we may “trust” finite-dimensional approximations of infinite strategy games when the approximations satisfy choice mobility.

D. Emergence of Choice Paralysis

In the following result, we show that one way that choice paralysis may occur is through decaying revision protocol values; paralysis occurs if the rate of strategy switching decreases towards zero as the number of strategies increases.

Proposition 2. Let $\theta_n, x_n(0) \in \Delta^{n-1}$ for all $n \in \mathbb{N}$, and consider the approximation (4) with $\lambda_n := \sum_{i=1}^n (\theta_n)_i \delta_{s_i}$ and initial condition $x_n(0)$. Suppose for all $n \in \mathbb{N}$ that $\lim_{t \rightarrow \infty} x_n(t) = \bar{x}_n$ for some $\bar{x}_n \in \Delta^{n-1}$. If there exist $\epsilon > 0$ and a sequence $\{M_n \in \mathbb{R}_+ : n \in \mathbb{N}\}$ such that $M_n \xrightarrow{n \rightarrow \infty} 0$ and, for all $n \in \mathbb{N}$, it holds that $\|x_n(0) - \bar{x}_n\|_1 > \epsilon$ and

$$\|\dot{x}_n(t)\|_1 \leq M_n \text{ for all } t \in [0, \infty), \quad (8)$$

then $\mathcal{X} = \{x_n : n \in \mathbb{N}\}$ suffers from choice paralysis. In particular, (8) holds when the revision protocol $\hat{\varphi}_n$ satisfies $|\hat{\varphi}_n(i, j, \theta)| \leq \frac{1}{2} M_n$ for all $i, j \in \{1, \dots, n\}$, $\theta \in \Delta^{n-1}$.

Proposition 2 is not restricted to fixed reference measures; it holds for time-varying λ_n , e.g., in the replicator dynamics.

IV. EXAMPLE

We now construct an explicit example illustrating choice paralysis and inconsistency between the infinite-strategy dynamics and the finite approximations. Consider the game on $S = [0, 1]$ defined by $F(\mu)(s) = \int_S f(s, s') d\mu(s')$ with

$$f(s, s') = -1 \text{ if } s = s', \quad f(s, s') = 0 \text{ otherwise,}$$

evolving according to the replicator dynamics (recall Example 1). This game penalizes players that choose the same strategy as another player, and hence is often referred to as the “anticoordination game” [27]. It is straightforward to see that $\dot{\mu}(t)(B) = \int_B (F(\mu(t))(s) - \int_S F(\mu(t))(s') d\mu(t)(s')) d\mu(t)(s) = 0$ for every set $B \in \mathcal{B}(S)$, whenever $\mu(t)$ is absolutely continuous with respect to Lebesgue measure. In particular, this shows that the infinite-strategy evolutionary game gets completely stuck at such population states; $\mu(t) = \mu(0)$ for all $t \in [0, \infty)$, for every $\mu(0)$ with (Lebesgue) density.

On the other hand, the approximation (4) reduces to

$$(\dot{x}_n(t))_i = -(x_n(t))_i^2 + (x_n(t))_i \sum_{j=1}^n (x_n(t))_j^2$$

for all $i \in \{1, \dots, n\}$ for every $n \in \mathbb{N}$, which corresponds to the finite-dimensional game defined by $\hat{F}_n(\theta) = -\theta$. It is straightforward to show that $\bar{x}_n = \frac{1}{n} \mathbf{1}_n$ is the only stationary point of these dynamics in the relative interior of Δ^{n-1} , where $\mathbf{1}_n$ denotes the n -vector of all ones, and that \bar{x}_n is also this (strictly concave) game’s unique Nash equilibrium [27, Corollary 3.1.4]. As such, the stationary point $\bar{x}_n = \frac{1}{n} \mathbf{1}_n$ is globally asymptotically stable over the relative interior of Δ^{n-1} [27, Theorem 7.2.4] for every $n \in \mathbb{N}$, suggesting that the game converges towards the uniform distribution on S when initialized with full support. Therefore, the approximations fail to accurately model the game’s behavior in the long-time limit. Indeed, the dynamics suffer from choice paralysis. Intuitively, this is because, in the vicinity

of the stationary point, $\|\dot{x}_n(t)\|_1 \leq 2\|x_n(t)\|_2^2 \approx 2\|\bar{x}_n\|_2^2 = \frac{2}{n} \xrightarrow{n \rightarrow \infty} 0$, so paralysis emerges (recall Proposition 2). This can occur even at states that are uniformly bounded away from \bar{x}_n in ℓ_1 . For instance, let $\epsilon > 0$, and for all $n \in \mathbb{N}$ define $m_n \in \mathbb{N}$ to be the ceiling of $\frac{\epsilon n}{2}$ and consider $x_n(0) = ((\frac{1}{n} + \frac{\epsilon}{2m_n})\mathbf{1}_{m_n}, (\frac{1}{n} - \frac{\epsilon}{2m_n})\mathbf{1}_{m_n}, \frac{1}{n}\mathbf{1}_{n-2m_n})$. It holds that $x_n(0) \in \Delta^{n-1}$, that $\|x_n(0) - \bar{x}_n\|_1 = \epsilon$, and that $\|\dot{x}_n(0)\|_1 \leq 2\|x_n(0)\|_2^2 = \frac{2}{n} + \frac{\epsilon^2}{m_n} \xrightarrow{n \rightarrow \infty} 0$, so the collection of approximations exhibits paralysis immediately at $t = 0$ despite the fact that every individual approximation converges to its associated equilibrium.

While the above example illustrates the correlation between choice paralysis and decaying convergence rates of the approximations, it utilizes a discontinuous payoff function f that violates Assumption 1. An interesting open problem is to characterize the relationship between choice paralysis and the continuity properties of f .

APPENDIX

Proof of Theorem 1. To simplify the exposition, we denote $\varphi(s, s', \mu(\tau), F(\mu(\tau)))$ by $\varphi(s, s', \tau)$ and $\varphi(s, s', \mu_n(\tau), F(\mu_n(\tau)))$ by $\varphi_n(s, s', \tau)$ for all $\tau \in [0, \infty)$.

Let $T \in (0, \infty)$. Furthermore, let $n \in \mathbb{N}$, let $t \in [0, T]$, and let $g \in \text{Lip}_1(S)$ be such that $\|g\|_\infty \leq 1$. It holds that

$$\begin{aligned} \int_S g d\mu(t) &= \int_{[0,t]} \int_S g(s) d(\dot{\mu}(\tau))(s) d\tau + \int_S g(s) d(\mu(0))(s) \\ &= \int_{[0,t]} \int_S g(s) \int_S \varphi(s', s, \tau) d(\mu(\tau))(s') d\lambda(s) d\tau \\ &\quad - \int_{[0,t]} \int_S g(s) \int_S \varphi(s, s', \tau) d\lambda(s') d(\mu(\tau))(s) d\tau \\ &\quad + \int_S g(s) d(\mu(0))(s), \\ \int_S g d\mu_n(t) &= \int_{[0,t]} \int_S g(s) d(\dot{\mu}_n(\tau))(s) d\tau + \int_S g(s) d(\mu_n(0))(s) \\ &= \int_{[0,t]} \int_S g(s) \int_S \varphi_n(s', s, \tau) d(\mu_n(\tau))(s') d\lambda_n(s) d\tau \\ &\quad - \int_{[0,t]} \int_S g(s) \int_S \varphi_n(s, s', \tau) d\lambda_n(s') d(\mu_n(\tau))(s) d\tau \\ &\quad + \int_S g(s) d(\mu_n(0))(s). \end{aligned}$$

Our goal is to estimate $\int_S g d\mu(t) - \int_S g d\mu_n(t)$ by comparing corresponding “inflows,” “outflows,” and initial states.

a) Inflows: We have that

$$\begin{aligned} \Delta_1(\tau) &:= \int_S g(s) \int_S \varphi(s', s, \tau) d(\mu(\tau))(s') d\lambda(s) \\ &\quad - \int_S g(s) \int_S \varphi_n(s', s, \tau) d(\mu_n(\tau))(s') d\lambda_n(s) \\ &= \int_S g(s) \int_S \varphi(s', s, \tau) d(\mu(\tau))(s') d\lambda(s) \\ &\quad - \int_S g(s) \int_S \varphi_n(s', s, \tau) d(\mu(\tau))(s') d\lambda(s) \\ &\quad + \int_S g(s) \int_S \varphi_n(s', s, \tau) d(\mu(\tau))(s') d\lambda(s) \\ &\quad - \int_S g(s) \int_S \varphi_n(s', s, \tau) d(\mu_n(\tau))(s') d\lambda(s) \\ &\quad + \int_S g(s) \int_S \varphi_n(s', s, \tau) d(\mu_n(\tau))(s') d\lambda(s) \\ &\quad - \int_S g(s) \int_S \varphi_n(s', s, \tau) d(\mu_n(\tau))(s') d\lambda_n(s). \end{aligned}$$

Comparing neighboring terms, we have that

$$\begin{aligned} &|\int_S g(s) \int_S (\varphi(s', s, \tau) - \varphi_n(s', s, \tau)) d(\mu(\tau))(s') d\lambda(s)| \\ &= |\int_S g(s) \int_S (\varphi(s', s, \tau) - \varphi_n(s', s, \tau)) d(\mu(\tau))(s') d\lambda(s)| \\ &\leq \int_S |g(s)| \int_S L_3 d_{\text{BL}}(\mu(\tau), \mu_n(\tau)) d(\mu(\tau))(s') d\lambda(s) \\ &\leq L_3 d_{\text{BL}}(\mu(\tau), \mu_n(\tau)), \end{aligned}$$

which follows from the bound $\|g\|_\infty \leq 1$ and the Lipschitzness of φ on $\mathcal{P}(S)$. Next,

$$\begin{aligned} &|\int_S g(s) \int_S \varphi_n(s', s, \tau) d(\mu(\tau) - \mu_n(\tau))(s') d\lambda(s)| \\ &= |\int_S g(s) \int_S \varphi_n(s', s, \tau) d(\mu(\tau) - \mu_n(\tau))(s') d\lambda(s)| \\ &\leq \max\{L_1, M\} d_{\text{BL}}(\mu(\tau), \mu_n(\tau)) \end{aligned}$$

follows from boundedness and Lipschitzness of φ in its first argument together with the bound $\|g\|_\infty \leq 1$. Furthermore,

$$\begin{aligned} &|\int_S g(s) \int_S \varphi_n(s', s, \tau) d(\mu_n(\tau))(s') d(\lambda - \lambda_n)(s)| \\ &= |\int_S \int_S g(s) \varphi_n(s', s, \tau) d(\lambda - \lambda_n)(s) d(\mu_n(\tau))(s')| \\ &\leq (L_2 + M) d_{\text{BL}}(\lambda, \lambda_n), \end{aligned}$$

since $s \mapsto g(s) \varphi_n(s', s, \tau)$ is bounded with constant M and is Lipschitz with constant $L_2 + M$ (which intimately relies on boundedness of both g and φ ; in general, the product of Lipschitz functions is not Lipschitz). Thus, we conclude that

$$\begin{aligned} |\Delta_1(\tau)| &\leq (L_3 + \max\{L_1, M\}) d_{\text{BL}}(\mu(\tau), \mu_n(\tau)) \\ &\quad + (L_2 + M) d_{\text{BL}}(\lambda, \lambda_n). \end{aligned}$$

b) Outflows: From a nearly identical analysis as above:

$$\begin{aligned} \Delta_2(\tau) &:= \int_S g(s) \int_S \varphi(s, s', \tau) d\lambda(s') d(\mu(\tau))(s) \\ &\quad - \int_S g(s) \int_S \varphi_n(s, s', \tau) d\lambda_n(s') d(\mu_n(\tau))(s), \\ |\Delta_2(\tau)| &\leq (L_1 + L_3 + M) d_{\text{BL}}(\mu(\tau), \mu_n(\tau)) \\ &\quad + \max\{L_2, M\} d_{\text{BL}}(\lambda, \lambda_n). \end{aligned}$$

c) Initial states: As $g \in \text{Lip}_1(S)$, $\|g\|_\infty \leq 1$, we have

$$\begin{aligned} &|\int_S g(s) d(\mu(0))(s) - \int_S g(s) d(\mu_n(0))(s)| \\ &\leq d_{\text{BL}}(\mu(0), \mu_n(0)). \end{aligned}$$

d) Completing the proof: Given our above estimates, we complete the proof as follows. First, we see that

$$\begin{aligned} &|\int_S g d\mu(t) - \int_S g d\mu_n(t)| \\ &\leq \int_{[0,t]} (|\Delta_1(\tau)| + |\Delta_2(\tau)|) d\tau + d_{\text{BL}}(\mu(0), \mu_n(0)) \\ &\leq K \int_{[0,t]} (d_{\text{BL}}(\mu(\tau), \mu_n(\tau)) + d_{\text{BL}}(\lambda, \lambda_n)) d\tau \\ &\quad + d_{\text{BL}}(\mu(0), \mu_n(0)), \end{aligned}$$

where $K := \max\{2L_3 + 3 \max\{L_1, M\}, 3 \max\{L_2, M\}\}$. Taking the supremum on the left-hand side over $g \in \text{Lip}_1(S)$, $\|g\|_\infty \leq 1$, we find that

$$\begin{aligned} &d_{\text{BL}}(\mu(t), \mu_n(t)) \\ &\leq K \int_{[0,t]} (d_{\text{BL}}(\mu(\tau), \mu_n(\tau)) + d_{\text{BL}}(\lambda, \lambda_n)) d\tau \\ &\quad + d_{\text{BL}}(\mu(0), \mu_n(0)). \end{aligned}$$

Applying Grönwall's inequality gives that

$$\begin{aligned} &d_{\text{BL}}(\mu(t), \mu_n(t)) \\ &\leq -d_{\text{BL}}(\lambda, \lambda_n) + (d_{\text{BL}}(\lambda, \lambda_n) + d_{\text{BL}}(\mu(0), \mu_n(0))) e^{Kt} \\ &\leq -d_{\text{BL}}(\lambda, \lambda_n) + (d_{\text{BL}}(\lambda, \lambda_n) + d_{\text{BL}}(\mu(0), \mu_n(0))) e^{KT}. \end{aligned}$$

Thus,

$$\begin{aligned} &\sup_{t \in [0,T]} d_{\text{BL}}(\mu(t), \mu_n(t)) \\ &\leq -d_{\text{BL}}(\lambda, \lambda_n) + (d_{\text{BL}}(\lambda, \lambda_n) + d_{\text{BL}}(\mu(0), \mu_n(0))) e^{KT}. \end{aligned}$$

Since $n \in \mathbb{N}$ is arbitrary and $\lim_{n \rightarrow \infty} d_{\text{BL}}(\lambda, \lambda_n) = \lim_{n \rightarrow \infty} d_{\text{BL}}(\mu(0), \mu_n(0)) = 0$ as both $\lambda_n \xrightarrow{n \rightarrow \infty} \lambda$ weakly and $\mu_n(0) \xrightarrow{n \rightarrow \infty} \mu(0)$ weakly, the conclusion (5) follows. \square

Proof of Theorem 2. Suppose that $\lim_{t \rightarrow \infty} \sup_{n \in \mathbb{N}} \|x_n(t) - \bar{x}_n\|_1 = 0$ for some collection $\{\bar{x}_n \in \Delta^{n-1} : n \in \mathbb{N}\}$. Let $T \in (0, \infty)$. By the triangle inequality and Theorem 1,

$$\begin{aligned} & \limsup_{n \rightarrow \infty} d_{\text{BL}}(\sum_{i=1}^n (\bar{x}_n)_i \delta_{s_i}, \mu(T)) \\ & \leq \limsup_{n \rightarrow \infty} d_{\text{BL}}(\sum_{i=1}^n (\bar{x}_n)_i \delta_{s_i}, \sum_{i=1}^n (x_n(T))_i \delta_{s_i}) \\ & \quad + \limsup_{n \rightarrow \infty} \sup_{\tau \in [0, T]} d_{\text{BL}}(\sum_{i=1}^n (x_n(\tau))_i \delta_{s_i}, \mu(\tau)) \\ & = \limsup_{n \rightarrow \infty} d_{\text{BL}}(\sum_{i=1}^n (\bar{x}_n)_i \delta_{s_i}, \sum_{i=1}^n (x_n(T))_i \delta_{s_i}) \\ & \leq \sup_{n \in \mathbb{N}} d_{\text{BL}}(\sum_{i=1}^n (\bar{x}_n)_i \delta_{s_i}, \sum_{i=1}^n (x_n(T))_i \delta_{s_i}). \end{aligned}$$

Since $T \in (0, \infty)$ is arbitrary, we conclude that

$$\begin{aligned} & \limsup_{T \rightarrow \infty} \limsup_{n \rightarrow \infty} d_{\text{BL}}(\sum_{i=1}^n (\bar{x}_n)_i \delta_{s_i}, \mu(T)) \\ & \leq \limsup_{T \rightarrow \infty} \sup_{n \in \mathbb{N}} d_{\text{BL}}(\sum_{i=1}^n (\bar{x}_n)_i \delta_{s_i}, \sum_{i=1}^n (x_n(T))_i \delta_{s_i}) \\ & = \limsup_{T \rightarrow \infty} \sup_{n \in \mathbb{N}} \sup_{\substack{g \in \text{Lip}_1(S) \\ \|g\|_\infty \leq 1}} |\sum_{i=1}^n ((\bar{x}_n)_i - (x_n(T))_i) g(s_i)| \\ & \leq \limsup_{T \rightarrow \infty} \sup_{n \in \mathbb{N}} \sum_{i=1}^n |(\bar{x}_n)_i - (x_n(T))_i| \\ & = \limsup_{T \rightarrow \infty} \sup_{n \in \mathbb{N}} \|\bar{x}_n - x_n(T)\|_1 = 0. \end{aligned}$$

Hence, $\lim_{t \rightarrow \infty} \lim_{n \rightarrow \infty} d_{\text{BL}}(\sum_{i=1}^n (\bar{x}_n)_i \delta_{s_i}, \mu(t)) = 0$, which proves (7). The two final enumerated results follow immediately from the triangle inequality. \square

Proof of Proposition 2. Let $t \in [0, \infty)$. For all $n \in \mathbb{N}$,

$$\begin{aligned} \|x_n(t) - \bar{x}_n\|_1 &= \left\| \int_{[0, t]} \dot{x}_n(\tau) d\tau + x_n(0) - \bar{x}_n \right\|_1 \\ &\geq \|x_n(0) - \bar{x}_n\|_1 - \int_{[0, t]} \|\dot{x}_n(\tau)\|_1 d\tau. \end{aligned}$$

Therefore,

$$\begin{aligned} \limsup_{n \rightarrow \infty} \|x_n(0) - \bar{x}_n\|_1 &\leq \limsup_{n \rightarrow \infty} \|x_n(t) - \bar{x}_n\|_1 \\ &\quad + \limsup_{n \rightarrow \infty} \int_{[0, t]} \|\dot{x}_n(\tau)\|_1 d\tau, \end{aligned}$$

implying that

$$\begin{aligned} \sup_{n \in \mathbb{N}} \|x_n(t) - \bar{x}_n\|_1 &\geq \limsup_{n \rightarrow \infty} \|x_n(0) - \bar{x}_n\|_1 \\ &\quad - \limsup_{n \rightarrow \infty} \int_{[0, t]} \|\dot{x}_n(\tau)\|_1 d\tau \\ &\geq \epsilon - \limsup_{n \rightarrow \infty} M_n t = \epsilon, \end{aligned}$$

so $\mathcal{X} = \{x_n : n \in \mathbb{N}\}$ suffers from choice paralysis.

If $|\hat{\varphi}_n(i, j, \theta)| \leq \frac{1}{2} M_n$ for all $i, j \in \{1, \dots, n\}$, $\theta \in \Delta^{n-1}$, then (4) gives that

$$\begin{aligned} |(\dot{x}_n(t))_i| &\leq \frac{1}{2} M_n \sum_{j=1}^n (\lambda_n(s_i)(x_n(t))_j + (x_n(t))_i \lambda_n(s_j)) \\ &= \frac{1}{2} M_n (\lambda_n(s_i) + (x_n(t))_i) \end{aligned}$$

for all $n \in \mathbb{N}$ and all $i \in \{1, \dots, n\}$, and therefore

$$\|\dot{x}_n(t)\|_1 \leq \frac{1}{2} M_n \sum_{i=1}^n (\lambda_n(s_i) + (x_n(t))_i) = M_n,$$

so (8) holds since $t \in [0, \infty)$ is arbitrary. \square

ACKNOWLEDGMENTS

The author thanks Murat Arcak for their constructive feedback and insightful discussions.

REFERENCES

- [1] B. G. Anderson, S. Pfrommer, and S. Sojoudi. Approximately Gaussian replicator flows: Nonconvex optimization as a Nash-convergent evolutionary game. In *IEEE Conference on Decision and Control*, 2024.
- [2] B. G. Anderson, J. Li, S. Sojoudi, and M. Arcak. Evolutionary games on infinite strategy sets: Convergence to Nash equilibria via dissipativity. *IEEE Transactions on Automatic Control*, 2025.
- [3] M. Arcak and N. C. Martins. Dissipativity tools for convergence to Nash equilibria in population games. *IEEE Transactions on Control of Network Systems*, 2021.
- [4] D. T. Bishop and C. Cannings. A generalized war of attrition. *Journal of Theoretical Biology*, 1978.
- [5] V. I. Bogachev. *Measure Theory*, volume 2. Springer, 2007.
- [6] I. M. Bomze. Dynamical aspects of evolutionary stability. *Monatshefte für Mathematik*, 1990.
- [7] I. M. Bomze. Cross entropy minimization in uninhabitable states of complex populations. *Journal of Mathematical Biology*, 1991.
- [8] M.-W. Cheung. Pairwise comparison dynamics for games with continuous strategy space. *Journal of Economic Theory*, 2014.
- [9] R. Cressman. Stability of the replicator equation with continuous strategy space. *Mathematical Social Sciences*, 2005.
- [10] R. Cressman and J. Hofbauer. Measure dynamics on a one-dimensional continuous trait space: Theoretical foundations for adaptive dynamics. *Theoretical Population Biology*, 2005.
- [11] S. S. Dragomir. *Some Gronwall Type Inequalities and Applications*. Nova Science Pub Inc, 2003.
- [12] M. J. Fox and J. S. Shamma. Population games, stable games, and passivity. *Games*, 2013.
- [13] S. Ganzfried. Algorithm for computing approximate Nash equilibrium in continuous games with application to continuous Blotto. *Games*, 2021.
- [14] D. Hingu, K. S. Mallikarjuna Rao, and A. J. Shaiju. Evolutionary stability of polymorphic population states in continuous games. *Dynamic Games and Applications*, 2018.
- [15] D. Hingu, K. S. Mallikarjuna Rao, and A. J. Shaiju. On superiority and weak stability of population states in evolutionary games. *Annals of Operations Research*, 2020.
- [16] J. Hofbauer and W. H. Sandholm. Stable games and their dynamics. *Journal of Economic Theory*, 2009.
- [17] J. Hofbauer, J. Oechssler, and F. Riedel. Brown-von Neumann-Nash dynamics: The continuous strategy case. *Games and Economic Behavior*, 2009.
- [18] R. Lahkar and F. Riedel. The logit dynamic for games with continuous strategy sets. *Games and Economic Behavior*, 2015.
- [19] R. Lahkar, S. Mukherjee, and S. Roy. Generalized perturbed best response dynamics with a continuum of strategies. *Journal of Economic Theory*, 2022.
- [20] M. A. Mabrok. Passivity analysis of replicator dynamics and its variations. *IEEE Transactions on Automatic Control*, 2021.
- [21] E. Mazumdar, L. J. Ratliff, M. I. Jordan, and S. S. Sastry. Policy-gradient algorithms have no guarantees of convergence in linear quadratic games. In *International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 2020.
- [22] J. Oechssler and F. Riedel. Evolutionary dynamics on infinite strategy spaces. *Journal of Economic Theory*, 2001.
- [23] J. Oechssler and F. Riedel. On the dynamic foundation of evolutionary stability in continuous models. *Journal of Economic Theory*, 2002.
- [24] J.-B. Park, B. H. Kim, J.-H. Kim, M.-H. Jung, and J.-K. Park. A continuous strategy game for power transactions analysis in competitive electricity markets. *IEEE Transactions on Power Systems*, 2001.
- [25] S. Park, N. C. Martins, and J. S. Shamma. From population games to payoff dynamics models: A passivity-based approach. In *IEEE Conference on Decision and Control*, 2019.
- [26] P. J. Reny. Strategic approximations of discontinuous games. *Economic Theory*, 2011.
- [27] W. H. Sandholm. *Population Games and Evolutionary Dynamics*. MIT Press, 2010.