

Fast Convergence in Semi-Anonymous Potential Games

Holly Borowski and Jason R. Marden

Abstract—Log-linear learning has been extensively studied in both the game theoretic and distributed control literature. It is appealing for many applications because it often guarantees that the agents’ collective behavior will converge in probability to the optimal system configuration. However, the worst case convergence time can be prohibitively long, i.e., exponential in the number of players. Building upon the work in [22], we formalize a modified log-linear learning algorithm whose worst case convergence time is roughly linear in the number of players. We prove this characterization for a class of potential games where agents’ utility functions can be expressed as a function of aggregate behavior within a finite collection of populations. Finally, we show that the convergence time remains roughly linear in the number of players even when the players are permitted to enter and exit the game over time.

I. INTRODUCTION

Game theoretic learning algorithms have gained traction as a design tool for distributed control systems [9], [10], [17], [23], [25]. Here, a static game is repeated over time, and agents revise their strategies based on their objective functions and on observations of other agents’ behavior. Emergent collective behavior for such revision strategies has been studied extensively in the literature, e.g., fictitious play [8], [15], [18], regret matching [12], and log-linear learning [1], [5], [22]. Although many of these learning rules have desirable asymptotic guarantees, their convergence times either remain uncharacterized or are prohibitively long [7], [11], [13], [22]. Characterizing convergence rates is key to determining whether a distributed algorithm is desirable for system control.

In many multi-agent systems, the agent objective functions can be designed to align with the system-level objective function, yielding a *potential game* [19] whose potential function is precisely the system objective function. Here, the optimal collective behavior of a multi-agent system corresponds to the Nash equilibrium that optimizes the potential function. Hence, learning algorithms which converge to this efficient Nash equilibrium have proven useful for distributed control.

Log-linear learning is one algorithm that accomplishes this task [5]. Log-linear learning is a perturbed best reply process where agents predominantly select the optimal action given their beliefs about other agents’ behavior; however, the agents

occasionally make mistakes, selecting suboptimal actions with a probability that decays exponentially with respect to the associated payoff loss. As noise levels approach zero, the resulting process has a unique stationary distribution with full support on the efficient Nash equilibria. By designing agents’ objective functions appropriately, log-linear learning can be used to define distributed control laws which converge to optimal steady-state behavior in the long run.

Unfortunately, worst-case convergence rates associated with log-linear learning are exponential in the game size [22]. This stems from inherent tension between desirable asymptotic behavior and convergence rates. The tension arises because small noise levels are necessary to ensure that the mass of the stationary distribution lies primarily on the efficient Nash equilibria; however, small noise levels also make it difficult to exit inefficient Nash equilibria, degrading convergence times.

Positive convergence rate results for log-linear learning and its variants are beginning to emerge for specific game structures [2], [14], [20], [22]. For example, in [20] the authors study the convergence rates of log-linear learning for a class of coordination games played over graphs. They demonstrate that underlying convergence rates are desirable provided that the interaction graph and its subgraphs are sufficiently sparse. Alternatively, in [22] the authors introduce a variant of log-linear learning and show that convergence times grow roughly linearly in the number of players for a special class of congestion games over parallel networks. They also show that convergence times remain linear in the number of players when players are permitted to exit and enter the game. Although these results are encouraging, the restriction to parallel networks is severe and hinders the applicability of such results to distributed engineering systems.

We focus on identifying whether the positive convergence rate results above extend beyond symmetric congestion games over parallel networks to games of a more general structure relevant to distributed engineering systems. Such guarantees are not automatic because there are many simplifying attributes associated with symmetric congestion games that do not extend in general (see Example 2). The main contributions of this paper are as follows:

- We formally define a subclass of potential games, called *semi-anonymous potential games*, which are parameterized by populations of agents where each agent’s objective function can be evaluated using only information regarding the agent’s own decision and the aggregate behavior within each population. Agents within a given population have identical action sets, and their objective functions share the same structural form. The congestion games studied in [22] could be viewed as

This research was supported by AFOSR grants #FA9550-12-1-0359, ONR grant #N00014-12-1-0643, NSF grant #ECCS-1351866, the NASA Aeronautics scholarship program, the Philanthropic Educational Organization, and the Zonta International Amelia Earhart fellowship program.

We wish to acknowledge conversations with Jinwoo Shin regarding the technical content of this paper, and we thank him for his feedback.

H. Borowski is with the Department of Aerospace Engineering Sciences, University of Colorado, 429 UCB, Boulder, CO 80309, 719-213-3254. holly.borowski@colorado.edu.

J. R. Marden is with the Department of Electrical and Computer Engineering, Harold Frank Hall, Rm 5161, University of California, Santa Barbara, 93106, 805-893-2299, jrmarden@ece.ucsb.edu. Corresponding author.

a semi-anonymous potential game with only one population.¹

– We introduce a variant of log-learning learning that extends the algorithm in [22]. In Theorem 1, we prove that the convergence time of this algorithm grows roughly linearly in the number of agents for a fixed number of populations. This analysis explicitly highlights the potential impact of system-wide heterogeneity, i.e., agents with different action sets or objective functions, on the convergence rates. Furthermore, in Example 4 we demonstrate how a given resource allocation problem can be modeled as a semi-anonymous potential game.

– We study the convergence times associated with our modified log-linear learning algorithm when the agents continually enter and exit the game. In Theorem 2, we prove that the convergence time of this algorithm remains roughly linear in the number of agents provided that the agents exit and enter the game at a sufficiently slow rate.

The forthcoming analysis is similar in structure to the analysis presented in [22]. We highlight the explicit differences between the two proof approaches throughout, and directly reference lemmas within [22] when appropriate. The central challenge in adapting and extending the proof in [22] to the setting of semi-anonymous potential games is dealing with the growth of the underlying state space. Note that the state space in [22] is characterized by the aggregate behavior of a single population while the state space in our setting is characterized by the Cartesian product of the aggregate behavior associated with several populations. The challenge arises from the fact that the employed techniques for analyzing the mixing times of this process, i.e., Sobolev constants, rely heavily on the structure of this underlying state space.

II. SEMI-ANONYMOUS POTENTIAL GAMES

Consider a game with agents $N = \{1, 2, \dots, n\}$. Each agent $i \in N$ has a finite action set denoted by \mathcal{A}_i and a utility function $U_i : \mathcal{A} \rightarrow \mathbb{R}$, where $\mathcal{A} = \prod_{i \in N} \mathcal{A}_i$ denotes the set of joint actions. We express an action profile $a \in \mathcal{A}$ as (a_i, a_{-i}) where $a_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$ denotes the actions of all agents other than agent i . We denote a game G by the tuple $G = (N, \{\mathcal{A}_i\}_{i \in N}, \{U_i\}_{i \in N})^2$.

Definition 1. A game G is a semi-anonymous potential game if there exists a partition $\mathcal{N} = (N_1, N_2, \dots, N_m)$ of N such that the following conditions are satisfied:

(i) For any population $N_\ell \in \mathcal{N}$ and agents $i, j \in N_\ell$ we have $\mathcal{A}_i = \mathcal{A}_j$. Accordingly, we say population N_ℓ has action set $\bar{\mathcal{A}}_\ell = \{\bar{a}_\ell^1, \bar{a}_\ell^2, \dots, \bar{a}_\ell^{s_\ell}\}^3$ where s_ℓ denotes the number of actions available to population N_ℓ . For simplicity, let $p(i) \in \{1, \dots, m\}$ denote the index of the population associated with agent i . Then, $\mathcal{A}_i = \bar{\mathcal{A}}_{p(i)}$ for all agents $i \in N$.

(ii) For any population $N_\ell \in \mathcal{N}$, let

$$X_\ell = \left\{ \left(\frac{v_\ell^1}{n}, \frac{v_\ell^2}{n}, \dots, \frac{v_\ell^{s_\ell}}{n} \right) \geq \mathbf{0} : \sum_{k=1}^{s_\ell} v_\ell^k = |N_\ell| \right\} \quad (1)$$

¹Semi-anonymous potential games can be viewed as a cross between a potential game and a finite population game [6].

²For brevity, we refer to G by $G = (N, \{\mathcal{A}_i\}, \{U_i\})$.

³We use the notation $\bar{\mathcal{A}}_\ell$ to represent the action set of the ℓ th population, whereas \mathcal{A}_i represents the action set of the i th agent.

represent all possible aggregate action assignments for the agents within population N_ℓ . Here, the utility function of any agent $i \in N_\ell$ can be expressed as a lower-dimensional function of the form $\bar{U}_i : \bar{\mathcal{A}}_{p(i)} \times X \rightarrow \mathbb{R}$ where $X = X_1 \times \dots \times X_m$. More specifically, the utility associated with agent i for an action profile $a = (a_i, a_{-i}) \in \mathcal{A}$ is of the form

$$U_i(a) = \bar{U}_i(a_i, a|_X)$$

where

$$\begin{aligned} a|_X &= (a|_{X_1}, a|_{X_2}, \dots, a|_{X_m}) \in X, \\ a|_{X_j} &= \frac{1}{n} \left\{ |\{j \in N_\ell : a_j = \bar{a}_\ell^k\}| \right\}_{k=1, \dots, s_\ell}. \end{aligned} \quad (2)$$

The operator $\cdot|_X$ captures each population's aggregate behavior in an action profile \cdot .

(iii) There exists a potential function $\phi : X \rightarrow \mathbb{R}$ such that for any $a \in \mathcal{A}$ and agent $i \in N$ with action $a'_i \in \mathcal{A}_i$,

$$U_i(a) - U_i(a'_i, a_{-i}) = \phi(a|_X) - \phi((a'_i, a_{-i})|_X). \quad (4)$$

If each agent $i \in N$ is alone in its respective partition, the definition of semi-anonymous potential games is equivalent to that of exact potential games in [19].

Example 1 (Congestion Games [4]). Consider a congestion game with players $N = \{1, \dots, n\}$ and roads $R = \{r_1, r_2, \dots, r_k\}$. Each road $r \in R$ is associated with a congestion function $C_r : \mathbb{Z}_+ \rightarrow \mathbb{R}$, where $C_r(k)$ is the congestion on road r with k total users. The action set of each player $i \in N$ represents the set of paths connecting player i 's source and destination, and has the form $\mathcal{A}_i \subseteq 2^R$. The utility function of each player $i \in N$ is given by

$$U_i(a_i, a_{-i}) = - \sum_{r \in a_i} C_r(|a|_r),$$

where $|a|_r = |\{j \in N : r \in a_j\}|$ is the number of players in joint action a whose path contains road r . This game is a potential game with potential function $\phi : \mathcal{X} \rightarrow \mathbb{R}$

$$\phi(a|_X) = - \sum_{r \in R} \sum_{k=1}^{|a|_r} C_r(k). \quad (5)$$

When the players' action sets are symmetric, i.e., $\mathcal{A}_i = \mathcal{A}_j$ for all agents $i, j \in N$, then a congestion game is a semi-anonymous potential game with a single population. Such games, also referred to as anonymous potential games, are the focus of [22]. When the players' action sets are asymmetric, i.e., $\mathcal{A}_i \neq \mathcal{A}_j$ for at least one pair of agents $i, j \in N$, then a congestion game is a semi-anonymous potential game where populations consist of agents with identical path choices. The results in [22] are not proven to hold for such settings.

The following example highlights issues that arise when transitioning from a single population to multiple populations.

Example 2. Consider a resource allocation game with n players and three resources, $R = \{r_1, r_2, r_3\}$. Let n be even and divide players evenly into populations N_1 and N_2 . Suppose that players in N_1 may select exactly one resource from $\{r_1, r_2\}$, and players in N_2 may select exactly one resource from $\{r_2, r_3\}$. The welfare garnered at each resource

depends on how many players have selected that resource; the resource-specific welfare functions are

$$\begin{aligned} W_{r_1}(k) &= 2k, \\ W_{r_2}(k) &= \min \left\{ 3k, \frac{3}{2}n \right\}, \\ W_{r_3}(k) &= k. \end{aligned}$$

where $k \in \{0, 1, \dots, n\}$ represents the number of agents selecting a given resource. The total system welfare is

$$W(a) = \sum_{r \in R} W_r(|a|_r)$$

for any $a \in \mathcal{A}$, where $|a|_r$ represents the number of agents selecting resource r under action profile a . Assign each agent's utility according to its marginal contribution to the system-level welfare: for agent i and action profile a

$$U_i(a) = W(a) - W(\emptyset, a_{-i}) \quad (6)$$

where \emptyset indicates that player i did not select a resource. The marginal contribution utility in (6) ensures that the resulting game is a potential game with potential function W [24].

If the agents had symmetric action sets, i.e., if $\mathcal{A}_i = \{r_1, r_2, r_3\}$ for all $i \in N$, then this game has exactly one Nash equilibrium with $n/2$ players at resource r_1 and $n/2$ players at resource r_2 . This Nash equilibrium corresponds to the optimal allocation.

In contrast, the two population scenario above has many Nash equilibria, two of which are: (i) an optimal Nash equilibrium in which all players from N_1 select resource r_1 and all players from N_2 select resource r_2 , and (ii) a suboptimal Nash equilibrium in which all players from N_1 select resource r_2 and all players from N_2 select resource r_3 . This large number of equilibria will significantly slow any equilibrium selection process, such as log-linear learning and its variants.

III. MAIN RESULTS

Example 2 invites the question: can a small amount of heterogeneity break down the fast convergence results of [22]? In this section, we present a variant of log-linear learning [5] that extends the algorithm for single populations in [22]. In Theorem 1 we prove that for any semi-anonymous potential game our algorithm ensures (i) the potential associated with asymptotic behavior is close to the maximum and (ii) the convergence time grows roughly linearly in the number of agents for a fixed number of populations. In Theorem 2 we show that these guarantees still hold when agents are permitted to enter and exit the game. An algorithm which converges quickly to the potential function maximizer is useful for multi-agent systems because agent objective functions can often be designed so that the potential function is identical to the system objective function as in Example 2.

A. Modified Log-Linear Learning

The following modification of the log-linear learning algorithm extends the algorithm in [22]. Let $a(t) \in \mathcal{A}$ be the joint

action at time $t \geq 0$. Each agent $i \in N$ updates its action upon ticks of a Poisson clock with rate $\alpha n / z_i(t)$, where

$$z_i(t) = |\{k \in N_{p(i)} : a_k(t) = a_i(t)\}|,$$

and $\alpha > 0$ is a design parameter which dictates the expected total update rate. A player's update rate is higher if he is not using a common action within his population. To continually modify his clock rate, each player must know the value of $z_i(t)$, i.e., the number of players within his population sharing his action choice, for all $t \in \mathbb{R}$. In many cases, agents also need this information to evaluate their utilities, e.g., when players' utilities are their marginal contribution to the total welfare, as in Example 2.

When player i 's clock ticks, he chooses action $a_i \in \bar{\mathcal{A}}_{p(i)} = \mathcal{A}_i$ probabilistically according to

$$\begin{aligned} \text{Prob}[a_i(t^+) = a_i \mid a(t)] &= \frac{e^{\beta U_i(a_i, a_{-i}(t))}}{\sum_{a'_i \in \mathcal{A}_i} e^{\beta U_i(a'_i, a_{-i}(t))}} \\ &= \frac{e^{\beta \phi(a(t)|_x)}}{\sum_{a'_i \in \mathcal{A}_i} e^{\beta \phi((a'_i, a_{-i}(t))|_x)}}, \quad (7) \end{aligned}$$

for any $a_i \in \mathcal{A}_i$, where $a_i(t^+)$ indicates the agent's revised action and β is a design parameter that determines how likely an agent is to choose a high payoff action. As $\beta \rightarrow \infty$, payoff maximizing actions are chosen, and as $\beta \rightarrow 0$, agents choose from their action sets with uniform probability. The new joint action is of the form $a(t^+) = (a_i(t^+), a_{-i}(t)) \in \mathcal{A}$, where $t \in \mathbb{R}^+$ is the time immediately before agent i 's update occurs. For a discrete time implementation of this algorithm and a comparison with the algorithm in [22], please see Appendix B.

The expected number of updates per second for the continuous time implementation of our modified log-linear learning algorithm is lower bounded by $m\alpha n$ and upper bounded by $(|\bar{\mathcal{A}}_1| + \dots + |\bar{\mathcal{A}}_m|)\alpha n$. To achieve an expected update rate at least as fast as the standard log-linear learning update rate, i.e., at least n per second, we set $\alpha \geq 1/m$. These dynamics define an ergodic, reversible Markov process for any $\alpha > 0$.

B. Semi-Anonymous Potential Games

Theorem 1 bounds the convergence time for modified log-linear learning in a semi-anonymous potential game and extends the results of [22] to semi-anonymous potential games. For notational simplicity, define $s := |\cup_{j=1}^m \bar{\mathcal{A}}_j|$.

Theorem 1. *Let $G = (N, \{\mathcal{A}_i\}, \{U_i\})$ be a semi-anonymous potential game with aggregate state space X and potential function $\phi : X \rightarrow [0, 1]$. Suppose agents play according to the modified log-linear learning algorithm described above, and the following conditions are met:*

(i) *The potential function is λ -Lipschitz, i.e., there exists $\lambda \geq 0$ such that*

$$|\phi(x) - \phi(y)| \leq \lambda \|x - y\|_1, \quad \forall x, y \in X.$$

(ii) *The number of players within each population is sufficiently large:*

$$\sum_{i=1}^m |N_i|^2 \geq \sum_{i=1}^m |\bar{\mathcal{A}}_i| - m.$$

For any fixed $\varepsilon \in (0, 1)$, if β is sufficiently large, i.e.,

$$\beta \geq \max \left\{ \frac{4m(s-1)}{\varepsilon} \log 2ms, \frac{4m(s-1)}{\varepsilon} \log \frac{8ms\lambda}{\varepsilon} \right\}, \quad (8)$$

then

$$\mathbb{E}[\phi(a(t)|X)] \geq \max_{x \in X} \phi(x) - \varepsilon \quad (9)$$

for all

$$t \geq \frac{2^{2ms} c_1 e^{3\beta} m(m(s-1))!^2 n}{4\alpha} \times \left(\log \log(n+1)^{ms-m} + \log \beta + 2 \log \frac{1}{\varepsilon} \right) \quad (10)$$

where c_1 is a constant that depends only on s .

We prove Theorem 1 in Appendix C. This theorem explicitly highlights the role of system heterogeneity, i.e., $m > 1$ distinct populations, on convergence times of the process. For the case when $m = 1$, Theorem 1 recovers the results of [22]. Observe that for a fixed number of populations, the convergence time grows as $n \log \log n$. Furthermore, note that a small amount of system heterogeneity does not have a catastrophic impact on worst-case convergence times as suggested by Example 2.

It is important to note that our bound is exponential in the number of populations and in the total number of actions. Therefore our results do not guarantee fast convergence with respect to these parameters. However, our convergence rate bounds may be conservative in this regard. Furthermore, as we will show in Section IV, a significantly smaller value of β may often be chosen in order to further speed convergence while still retaining the asymptotic properties guaranteed in (9).

C. Time Varying Semi-Anonymous Potential Games

In this section, we consider a trajectory of semi-anonymous potential games to model the scenario where agents enter and exit the system over time,

$$\mathcal{G} = \{G^t\}_{t \geq 0} = \{N^t, \{\mathcal{A}_i^t\}_{i \in N^t}, \{U_i^t\}_{i \in N^t}\}_{t \geq 0}$$

where, for all $t \in \mathbb{R}^+$, the game G^t is a semi-anonymous potential game, and the set of *active* players, N^t , is a finite subset of \mathbb{N} . We refer to each agent $i \in \mathbb{N} \setminus N^t$ as *inactive*; an inactive agent has action set $\mathcal{A}_i^t = \emptyset$ at time t . Define $\mathcal{X} := \cup_{t \in \mathbb{R}^+} X^t$, where X^t is the finite aggregate state space corresponding to game G^t . At time t , denote the partitioning of players per Definition 1 by $\mathcal{N}^t = \{N_1^t, N_2^t, \dots, N_m^t\}$. We require that there is a fixed number of populations, m , for all time, and that the j -th population's action set is constant, i.e., $\forall j \in \{1, 2, \dots, m\}, \forall t_1, t_2 \in \mathbb{R}^+, \mathcal{A}_j^{t_1} = \mathcal{A}_j^{t_2}$. We write the fixed action set for players in the j -th population as $\bar{\mathcal{A}}_j$.

Theorem 2. *Let \mathcal{G} be a trajectory of semi-anonymous potential games with state space \mathcal{X} and time-invariant potential function $\phi : \mathcal{X} \rightarrow [0, 1]$. Suppose agents play according to the modified log-linear learning algorithm and Conditions (i) and (ii) of Theorem 1 are satisfied. Fix $\varepsilon \in (0, 1)$, assume the parameter β satisfies (8) and the following additional conditions are met:*

(iii) for all $t \in \mathbb{R}^+$, the number of players satisfies:

$$|N^t| \geq \max \left\{ \frac{4\alpha m e^{-3\beta}}{2^{2ms} c_1 m^2 (m(s-1))!^2}, 2\beta\lambda + 1 \right\}, \quad (11)$$

(iv) there exists $k > 0$ such that

$$|N_i^t| \geq |N^t| / k, \quad \forall i \in \{1, 2, \dots, m\}, \forall t \in \mathbb{R}^+, \quad (12)$$

(v) there exists a constant

$$\Lambda \geq 8c_0 \varepsilon^{-2} e^{3\beta} (6\beta\lambda + e^\beta k(s-1)) \quad (13)$$

such that for any t_1, t_2 with $|t_1 - t_2| \leq \Lambda$,

$$|\{i \in N^{t_1} \cup N^{t_2} : \mathcal{A}_i^{t_1} \neq \mathcal{A}_i^{t_2}\}| \leq 1, \quad (14)$$

and, if $i \in N^{t_1} \cap N^{t_2}$, then $i \in N_j^t$ for some $j \in \{1, \dots, m\}$ and for all time $t \in [t_1, t_2]$, i.e., agents may not switch populations over this interval. Here, c_0 and c_1 do not depend on the number of players, and hence the constant Λ does not depend on n .

Then,

$$\mathbb{E}[\phi(a(t)|\mathcal{X})] \geq \max_{x \in X(t)} \phi(x) - \varepsilon \quad (15)$$

for all

$$t \geq |N^0| e^{3\beta} c_0 \left(\frac{(ms-m)! \log(|N^0| + 2) + \beta}{\varepsilon^2} \right). \quad (16)$$

Theorem 2 states that, if player entry and exit rates are sufficiently slow as in Condition (v), then the convergence time of our algorithm is roughly linear in the number of players. However, the established bound grows quickly with the number of populations. Note that selection of parameter β impacts convergence time, as reflected in (16): larger β tends to slow convergence. However, the minimum β necessary to achieve an expected potential near the maximum, as in (15), is independent of the number of players, as given in (8). The proof of Theorem 2 follows a similar structure to the proof of Theorem 4 in [22] and is hence omitted for brevity. The significant technical differences arise due to differences in the size of the state space when $m > 1$. These differences give rise to Condition (iv) in our theorem.

IV. ILLUSTRATIVE EXAMPLES

In this section, we consider resource allocation games with a similar structure to Example 2. In each case, agents' utility functions are defined by their marginal contribution to the system welfare, W , as in (6). Hence, each example is a potential game with potential function W .

Modified log-linear learning defines an ergodic, continuous time Markov chain; we denote its transition kernel by P and its stationary distribution by π . For relevant preliminaries on Markov chains, please refer to Appendix A, and for a precise definition of the transition kernel and stationary distribution associated with modified log-linear learning, please refer to Appendices B and C.

Unless otherwise specified, we consider games with n players distributed evenly into populations N_1 and N_2 . There are three resources, $R = \{r_1, r_2, r_3\}$. Players in population N_1 may choose a single resource from $\{r_1, r_2\}$ and players

in population N_2 may choose a single resource from $\{r_2, r_3\}$. We represent a state by

$$x = (x_1^1, x_2^1, x_2^2, x_3^2), \quad (17)$$

where nx_1^1 and nx_2^1 are the numbers of players from N_1 choosing resources r_1 and r_2 . Likewise, nx_2^2 and nx_3^2 are the numbers of players from N_2 choosing resources r_2 and r_3 respectively. Welfare functions for each resource depend only on the number of players choosing that resource, and are specified in each example. The system welfare for a given state is the sum of the welfare garnered at each resource, i.e.,

$$W(x) = W_{r_1}(nx_1^1) + W_{r_2}(n(x_2^1 + x_2^2)) + W_{r_3}(nx_3^2).$$

Player utilities are their marginal contribution to the total welfare, W , as in (6).

In Example 3, we directly compute convergence times as in Theorem 1:

$$\min\{t : \mathbb{E}_{P^t(y, \cdot)} W(x) \geq \max_{x \in X} W(x) - \varepsilon, \forall y \in X\}, \quad (18)$$

for modified log-linear learning, the variant of [22], and standard log-linear learning. This direct analysis is possible due to the example's relatively small state space.

Example 3. Here, we compare convergence times of our log-linear learning variant, the variant of [22], and standard log-linear learning. The transition kernels for each process are described in detail in Appendix B.

Starting with the setup described above, we add a third population, N_3 . Agents in population N_3 contribute nothing to the system welfare and may only choose resource r_2 . Because the actions of agents in population N_3 are fixed, we represent states by aggregate actions of agents in populations N_1 and N_2 as in (17). The three resources have the following welfare functions for each $x = (x_1^1, x_2^1, x_2^2, x_3^2) \in X$:

$$\begin{aligned} W_{r_1}(x) &= 2nx_1^1, \\ W_{r_2}(x) &= \min\left\{3(nx_1^1 + nx_2^1), \frac{3}{2}(nx_2^1 + nx_2^2)\right\}, \\ W_{r_3}(x) &= nx_3^2. \end{aligned}$$

Our goal in this example is to achieve an expected total welfare that is within 98% of the maximum welfare.

We fix the number of players in populations N_1 and N_2 at $n_1 = n_2 = 7$, and vary the number of players in population n_3 to examine the sensitivity of each algorithm's convergence rate to the size of N_3 .

In our variant of log linear learning, increasing the size of population N_3 does not change the probability that a player from population N_1 or N_2 will update next. However, for standard log-linear learning and for the variant in [22], increasing the size of population N_3 significantly decreases the probability that players from N_1 or N_2 who are currently

choosing resource r_2 will be selected for update.⁴

We select β in all cases so that, as $t \rightarrow \infty$, the expected welfare associated with the resulting stationary distribution is within 98% of its maximum. Then we examine the time it takes to come within $\varepsilon = 0.05$ of this expected welfare. We multiply convergence times by the number of players, n , to analyze the expected number of updates required to reach the desired welfare. These numbers represent the convergence times when the expected total number of updates per unit time is held constant as n increases. Table 3 depicts β values and expected numbers of updates.

For both log-linear learning and our modification, the required β to reach an expected welfare within 98% of the maximum welfare is independent of n_3 and can be computed using the expressions

$$\pi_x^{\text{LLL}} \propto e^{\beta W(x)} \binom{n_1}{nx_1^1, nx_2^1} \binom{n_2}{nx_2^2, nx_3^2}, \quad (19)$$

$$\text{and } \pi_x^{\text{MLLL}} \propto e^{\beta W(x)}. \quad (20)$$

These stationary distributions can be verified using reversibility arguments with the standard and modified log-linear learning probability transition kernels, defined in [22] and Appendix B respectively. Unlike standard log-linear learning and our variant, the required β to reach an expected welfare of 98% of maximum for the log-linear learning variant of [22] does change with n_3 . For each value of n , we use the probability transition matrix to determine the necessary values of β which yield an expected welfare of 98% of its maximum.

Our algorithm converges to the desired expected welfare in fewer updates than both alternate algorithms for all tested values of n_3 , showing that convergence rates for log linear learning and the variant from [22] are both more sensitive to the number of players in population 3 than our algorithm.⁵

We are able to determine convergence times in Example 3 using each algorithm's probability transition matrix, P , because the state space is relatively small. Here, we directly compute the distance of distribution $\mu(t) = \mu(0)P^t$ to the stationary distributions, π^{LLL} and π^{MLLL} for the selected values of β , where P and π . Examples 4 and 6, however, have significantly larger state spaces, making similar computations with the probability transition matrix unrealistic. Thus, instead of computing convergence times as in (18) we repeatedly simulate our algorithm from a worst case initial state and approximate convergence times based on average behavior. This method does not directly give the convergence time of Theorem 1, but the average performance over a sufficiently large number of simulations is expected to reflect expected

⁴Recall that in our log-linear learning variant and the one introduced in [22], an updating player chooses a new action according to (7); the algorithms differ only in agents' update rates. In our algorithm, an agent i in population N_j 's update rate is $\alpha n / z_i^j(t)$, where $z_i^j(t)$ is the number of agents from population j playing the same action as agent i at time t . In the algorithm in [22], agent i 's update rate is $\alpha n / \bar{z}_i(t)$, where $\bar{z}_i(t)$ is the total number of agents playing the same action as agent i .

⁵A high update rate for players in population N_3 was undesirable because they contribute no value. While this example may seem contrived, mild variations would exhibit similar behavior. For example, consider a scenario in which a relatively large population that contributes little to the total welfare may choose from multiple resources.

Algorithm	n_3	β	Expected welfare	Expected # updates
Standard Log Linear Learning	1	3.77	98%	9430
	5	3.77	98%	11947
	50	3.77	98%	40250
	500	3.77	98%	323277
Log Linear Learning Variant from [22]	1	2.39	98%	1325
	5	2.44	98%	1589
	50	2.83	98%	3342
	500	3.72	98%	15550
Our Log Linear Learning Variant	1	1.28	98%	743
	5	1.28	98%	743
	50	1.28	98%	743
	500	1.28	98%	743

TABLE I: This table corresponds to Example 3. Here, there are three populations of agents, N_1 , N_2 , and N_3 , and three resources r_1 , r_2 , and r_3 . Agents in population N_1 may choose from resources r_1 and r_2 , and agents in population N_2 may choose from resources r_2 and r_3 . Agents in population N_3 may only choose resource r_2 . Welfare functions are given in (19); population N_3 contributes nothing to the overall system welfare. Here, we examine the sensitivity of convergence times to the size of N_3 , and keep the sizes of populations N_1 and N_2 fixed at 7. The third column of this table shows the values of β which yield an expected total welfare within 98% of the maximum. These values of β are constant for standard log-linear learning and for our variant, but grow with n for the algorithm in [22]. The final column shows the expected number of updates to achieve the desired near-maximum welfare. This value is constant for our algorithm, but increases with n for the other two. Global update rates are a design parameter dictated by parameter α ; selecting a global update rate of n per second ($\alpha = 1/m$), convergence times would be a factor of n smaller than the number of updates shown.

behavior predicted by the probability transition matrix.

Example 4. In this example we consider a scenario similar to the previous example, without the third population. That is, agents are evenly divided into two populations, N_1 and N_2 ; we allow the total number of agents to vary. Agents in N_1 may choose either resource r_1 or r_2 , and agents in N_2 may choose either resource r_2 or r_3 . We consider welfare functions of the following form:

$$W_{r_1}(x) = \frac{e^{x_1} - 1}{e^2}, \quad (21)$$

$$W_{r_2}(x) = \frac{e^{2x_2 + 2x_3} - 1}{e^2}, \quad (22)$$

$$W_{r_3}(x) = \frac{e^{2.5x_4} - 1}{e^2}. \quad (23)$$

for $x = (x_1^1, x_2^1, x_2^2, x_3^2) \in X$. Here, the global welfare optimizing allocation is $a_i = r_2$ for all $i \in N$, i.e., $x^{\text{opt}} = (0, 1/2, 1/2, 0)$. Similar to Example 2, this example has many Nash equilibria, two of which are x^{opt} and $x^{\text{ne}} = (1/2, 0, 0, 1/2)$.

We simulated our algorithm with $\alpha = 1/4$ starting from the inefficient Nash equilibrium, x^{ne} . Here, β is chosen to yield an expected steady state welfare equal to 90% of the maximum. We examine the time it takes the average welfare to come within $\varepsilon = 0.05$ of this expected welfare.

Simulation results are shown in Figure 1 averaged over 2000 simulations with n ranging from 4 to 100. Average convergence times are bounded below by $2n \log \log n$ for all




Fig. 1: Example 4, number of players vs. average convergence times. Here, there are two equal-sized populations of agents, N_1 and N_2 , and three resources r_1 , r_2 , and r_3 . Agents in population N_1 may choose from resources r_1 and r_2 , and agents in population N_2 may choose from resources r_2 and r_3 . Welfare functions are given in (23).

values of n , and are bounded above by $4n \log \log n$ when $n > 30$. These results support Theorem 1.

Example 5. In this example we investigate convergence times for modified log-linear learning when agents have larger action sets. We consider the situation where n agents are divided into two populations, N_1 and N_2 . Agents in N_1 may choose from resources in $A_1 = \{r_1, r_2, \dots, r_k\}$, and agents in population N_2 may choose from resources in $A_2 = \{r_k, r_{k+1}, \dots, r_{2k-1}\}$. That is, each agent may choose from k different resources, and the two populations share resource




Fig. 2: 5, number of agents vs. average time to reach 90% of the maximum welfare. Agents are separated into two populations, N_1 and N_2 . Agents in N_1 choose from resources r_1, r_2, \dots, r_k , and agents in N_2 choose from resources $r_k, r_{k+1}, \dots, r_{2k-1}$, where k varies from 5 to 15. Resource welfare functions are given by (24), agent utility functions are given by (6), and average convergence times are taken over 200 simulations.

r_k . Suppose resource welfare functions are

$$W_{r_j}(x) = \begin{cases} x/4n & \text{if } j \neq k \\ x^2/n^2 & \text{if } j = k, \end{cases} \quad (24)$$

and suppose agents' utilities are given by their marginal contribution to the total welfare, as in (6). We allow k to vary between 5 and 15, and n to vary between 4 and 50.

The welfare maximizing configuration is for all agents to choose resource r_k ; however, when all agents in populations N_1 and N_2 choose resources r_j and r_ℓ respectively, with $j, \ell \neq k$, this represents an inefficient Nash equilibrium. Along any path from this type of inefficient Nash equilibrium to the optimal configuration, when $n \geq 4$, at least $\lceil (n+4)/8 \rceil$ agents must make a utility-decreasing decision to move to resource r_k . Moreover, the additional resources are all alternative suboptimal choices each agent could make when revising its action; these alternate choices further slow convergence times. Figure 2 shows the average time it takes to reach a configuration whose welfare is 90% of the maximum, starting from an inefficient Nash equilibrium where all agents in N_1 choose resource r_1 and all agents in N_2 choose resource r_{2k-1} . Parameter β is selected so that the expected welfare is at least 90% of the maximum in the limit as $t \rightarrow \infty$. For each value of k , convergence times remain approximately linear in the number of agents, supporting Theorem 1.⁶

In Example 6 we compare convergence times for standard and modified log-linear learning in a sensor-target assignment problem.

⁶In this example, convergence times appear super-linear in the size of populations' action sets. Note that the bound in (10) is exponential in the sum of the sizes of each population's action set. Fast convergence with respect to parameter s warrants future investigation; in particular, convergence rates for our log-linear learning variant may be significantly faster than suggested in (10) under certain mild restrictions on resource welfare functions (e.g., submodularity) or for alternate log-linear learning variants (e.g., binary log-linear learning [3], [16]).

Example 6 (Sensor-Target Assignment). In this example, we assign a collection of mobile sensors to four regions. Each region contains a single target, and the sensor assignment should maximize the probability of detecting the targets, weighted by their values. The targets in regions $R = \{r_1, r_2, r_3, r_4\}$ have values

$$v_1 = 1, \quad v_2 = 2, \quad v_3 = 3, \quad v_4 = 4 \quad (25)$$

respectively. Three types of sensors will be used to detect the targets: strong, moderate, and weak. Detection probabilities of these three sensor types are:

$$p_s = 0.9, \quad p_m = 0.5, \quad p_w = 0.05. \quad (26)$$

The numbers of strong and weak sensors are $n_s = 1$ and $n_w = 5$. We vary the number of moderate sensors, n_m .

The expected welfare for area r_i is the detection probability of the collection of sensors located at r_i weighted by the value of target i :

$$W_{r_i}(k_s, k_m, k_w) = v_i (1 - (1 - p_s)^{k_s} (1 - p_m)^{k_m} (1 - p_w)^{k_w}),$$

where k_s , k_m and k_w represent the number of strong, moderate, and weak sensors located at region r_i . The total expected welfare for configuration a is

$$W(a) = \sum_{r \in R} W_r(|a|_r^s, |a|_r^m, |a|_r^w),$$

where $|a|_r^s$, $|a|_r^m$, and $|a|_r^w$ are the numbers of strong, moderate, and weak sensors choosing region r in a .

We assign agents' utilities according to their marginal contributions to the total welfare, W , as in (6). Our goal is to reach 98% of the maximum welfare. We set the initial state to be a worst-case Nash equilibrium.⁷

To approximate convergence times, we simulate each algorithm with the chosen β value⁸ and compute a running average of the total welfare over 1000 simulations. In Figure 3 we show the average number of iterations necessary to reach 98% of the maximum welfare.

For small values of n_w , standard log-linear learning converges more quickly than our modification, but modified log-linear learning converges faster than the standard version as n_w increases. The difference in convergence times is significant (≈ 1000 iterations) for intermediate values of n_w . As the total

⁷The initial configuration is chosen by assigning weak agents to the highest value targets and then assigning strong agents to lower value targets. In particular, agents are assigned in order of weakest to strongest according to their largest possible marginal contribution. This constitutes an inefficient Nash equilibrium. As a similar example, consider a situation with two sensors with detection probabilities $p_1 = 0.5$ and $p_2 = 1$, and two targets with values $v_1 = 2$ and $v_2 = 1$. The assignment (sensor 1 \rightarrow target 1, sensor 2 \rightarrow target 2) is an inefficient Nash equilibrium, whereas the opposite assignment is optimal. The large state space makes it infeasible to directly compute a stationary distribution, and hence also infeasible to compute values of β that will yield precisely the desired expected welfare. Thus, we use simulations to estimate the β which yields an expected welfare of 98% of the maximum.

⁸To approximate the value of β which yields the desired steady-state welfare of 98% of maximum, we simulated the standard and modified versions of log-linear learning for 1×10^6 iterations for a range of β values. We then selected the β which yields an average welfare closest to the desired welfare during the final 5000 iterations. Note that we could instead set β according to (8) for the modified log-linear learning algorithm; however, in order to compare convergence times of modified and standard log-linear learning, we chose β to achieve approximately the same expected welfare for both algorithms.




Fig. 3: Example 6, number of weak sensors vs. average convergence times. Here, there are three types of sensors which may choose from four resources. Sensor detection probabilities and resource values are given in (26) and (25). We fix the number of strong and moderate sensors and vary the number of weak sensors. This figure shows the average time it takes for the average welfare to reach 98% of maximum. The average is taken over 1000 iterations, and convergence times correspond to a global update rate of 1 per second. Error bars show standard deviations of the convergence times.

number of weak sensors increases, (1) the probabilities of transitions along the paths to the efficient Nash equilibrium begin to increase for both algorithms, and (2) more sensor configurations are close to the maximum welfare. Hence, convergence times for both algorithms decrease as n_w increases.

This sensor-target assignment problem does not display worst-case convergence times with respect to the number of agents for either algorithm. However, it demonstrates a situation where our modification can have an advantage over standard log-linear learning. In log-linear learning, the probability that the strong sensor will update next decreases significantly as the number of agents grows. In modified log-linear learning this probability remains fixed. This property is desirable for this particular sensor-target assignment problem, since the single strong sensor contributes significantly to the total system welfare.

V. CONCLUSION

We have extended the results of [22] to define dynamics for a class of semi-anonymous potential games whose player utility functions may be written as functions of aggregate behavior within each population. For games with a fixed number of actions and a fixed number of populations, the time it takes to come arbitrarily close to a potential function maximizer is linear in the number of players. This convergence time remains linear in the initial number of players even when players are permitted to enter and exit the game, provided they do so at a sufficiently slow rate.

REFERENCES

- [1] C. Alós-Ferrer and N. Netzer. The logit-response dynamics. *Games and Economic Behavior*, 68(2):413–427, 2010.
- [2] I. Arieli and H.P. Young. Fast convergence in population games. 2011.
- [3] G. Arslan, J. R. Marden, and J. S. Shamma. Autonomous vehicle-target assignment: a game theoretical formulation. *ASME Journal of Dynamic Systems, Measurement and Control*, 129(5):584–596, 2007.
- [4] M. Beckmann, C.B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, New Haven, 1956.
- [5] L. Blume. The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 1993.
- [6] L. E. Blume. Population games. *The Economy as a Complex Evolving System II*, pages 425–460, 1996.
- [7] G. Ellison. Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. *The Review of Economic Studies*, pages 17–45, 2000.
- [8] D. P. Foster and H. P. Young. On the Nonconvergence of Fictitious Play in Coordination Games. *Games and Economic Behavior*, 25(1):79–96, October 1998.
- [9] M. Fox and J. Shamma. Communication, convergence, and stochastic stability in self-assembly. In *49th IEEE Conference on Decision and Control (CDC)*, December 2010.
- [10] T. Goto, T. Hatanaka, and M. Fujita. Potential game theoretic attitude coordination on the circle: Synchronization and balanced circular formation. *2010 IEEE International Symposium on Intelligent Control*, pages 2314–2319, September 2010.
- [11] S. Hart and Y. Mansour. How long to equilibrium? The communication complexity of uncoupled equilibrium procedures. *Games and Economic Behavior*, 69:107–126, May 2010.
- [12] S. Hart and A. MasColell. A Simple Adaptive Procedure Leading to Correlated Equilibrium. *Econometrica*, 68(5):1127–1150, 2000.
- [13] M. Kandori, G. Mailath, and R. Rob. Learning, Mutation, and Long Run Equilibria in Games. *Econometrica*, 61(1):29–56, 1993.
- [14] G. Kreindler and H. Young. Fast convergence in evolutionary equilibrium selection. 2011.
- [15] J. Marden. Joint strategy fictitious play with inertia for potential games. *IEEE Transactions on Automatic Control*, 54(2):208–220, 2009.
- [16] J. Marden, G. Arslan, and J. Shamma. Connections between cooperative control and potential games illustrated on the consensus problem. *Proceedings of 2007 the European Control Conference*, 2007.
- [17] J. Marden and A. Wierman. Distributed welfare games. *Operations Research*, 61(1):155–168, 2013.
- [18] D. Monderer and L. Shapley. Fictitious Play Property for Games with Identical Interests. *Journal of Economic Theory*, (68):258–265, 1996.
- [19] D. Monderer and L. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.
- [20] A. Montanari and A. Saberi. The spread of innovations in social networks. *Proceedings of the National Academy of Sciences*, pages 20196–20201, 2010.
- [21] R. Montenegro and P. Tetali. Mathematical Aspects of Mixing Times in Markov Chains. *Foundations and Trends in Theoretical Computer Science*, 1(3):237–354, 2006.
- [22] D. Shah and J. Shin. Dynamics in Congestion Games. In *ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, 2010.
- [23] M. Staudigl. Stochastic stability in asymmetric binary choice coordination games. *Games and Economic Behavior*, 75(1):372–401, May 2012.
- [24] D. Wolpert and K. Tumer. An Introduction To Collective Intelligence. Technical report, 1999.
- [25] M. Zhu and S. Martinez. Distributed coverage games for mobile visual sensors (I): Reaching the set of Nash equilibria. In *Joint 48th IEEE Conference on Decision and Control and 28th Chinese Control Conference*, 2009.

APPENDIX

A. Markov chain preliminaries

A continuous time Markov chain, $\{Z_t\}_{t \geq 0}$, over a finite state space Ω may be written in terms of a corresponding discrete time chain with transition matrix M [21], where the

distribution $\mu(t)$ over Ω evolves as:

$$\mu(t) = \mu(0)e^{t(M-I)} = \mu(0)e^{-t} \sum_{k=0}^{\infty} \frac{t^k M^k}{k!}, \quad t \geq 0 \quad (27)$$

where we refer to M as the kernel of the process Z_t and $\mu(0) \in \Delta(\Omega)$ is the initial distribution. The following definitions and theorems are taken from [21], [22]. Let μ, ν be measures on the finite state space Ω . Total variation distance is defined as

$$\|\mu - \nu\|_{TV} := \frac{1}{2} \sum_{x \in \Omega} |\mu_x - \nu_x|. \quad (28)$$

and

$$D(\mu : \nu) := \sum_{x \in \Omega} \mu_x \log \frac{\mu_x}{\nu_x} \quad (29)$$

is defined to be the relative entropy between μ and ν . The total variation distance between two distributions can be bounded using the relative entropy:

$$\|\mu - \nu\|_{TV} \leq \sqrt{\frac{D(\mu : \nu)}{2}} \quad (30)$$

For a continuous time Markov chain with kernel M and stationary distribution π , the distribution $\mu(t)$ obeys

$$D(\mu(t) : \pi) \leq e^{-4t\rho(M)} D(\mu(0) : \pi), \quad t \geq 0 \quad (31)$$

where $\rho(M)$ is the Sobolev constant of M , defined by

$$\rho(M) := \inf_{\substack{f: \Omega \rightarrow \mathbb{R} \\ \mathcal{L}(f) \neq 0}} \frac{\mathcal{E}(f, f)}{\mathcal{L}(f)} \quad (32)$$

with

$$\mathcal{E}(f, f) := \frac{1}{2} \sum_{x, y \in \Omega} (f(x) - f(y))^2 M(x, y) \pi_x \quad (33)$$

$$\mathcal{L}(f) := \mathbb{E}_\pi \log \frac{f^2}{\mathbb{E}_\pi f^2}. \quad (34)$$

Here \mathbb{E}_π denotes the expectation with respect to stationary distribution π . For a Markov chain with initial distribution $\mu(0)$ and stationary distribution π , the total variation and relative entropy mixing times are

$$T_{TV}(\varepsilon) := \min_t \{\|\mu(t) - \pi\| \leq \varepsilon\} \quad (35)$$

$$T_D(\varepsilon) := \min_t \{D(\mu(t) : \pi) \leq \varepsilon\} \quad (36)$$

respectively. From [21], Corollary 2.6 and Remark 2.11,

$$T_D(\varepsilon) \leq \frac{1}{4\rho(M)} \left(\log \log \frac{1}{\pi_{\min}} + \log \frac{1}{\varepsilon} \right),$$

where $\pi_{\min} := \min_{x \in \Omega} \pi_x$. Applying (30),

$$\begin{aligned} T_{TV}(\varepsilon) &\leq T_D(2\varepsilon^2) \\ &\leq \frac{1}{4\rho(M)} \left(\log \log \frac{1}{\pi_{\min}} + 2 \log \frac{1}{\varepsilon} \right). \end{aligned} \quad (37)$$

Hence, a lower bound on the Sobolev constant yields an upper bound on the mixing time for the Markov chain.

B. Notation and Problem Formulation: Stationary Semi-Anonymous Potential Games

The following Markov chain, M , over state space X is the kernel of the continuous time modified log-linear learning process for stationary semi anonymous potential games. Define $n_j := |N_j|$ to be the size of population j , define $s_j := |\bar{A}_j|$, and let $\sigma := \sum_{j=1}^m s_j$. Let $e_j^k \in \mathbb{R}^{s_j}$ be the k th standard basis vector of length s_j for $k \in \{1, \dots, s_j\}$. Finally, let

$$x = (x_j, \mathbf{x}_{-j}) = (x_1, x_2, \dots, x_m) \in X,$$

where $x_j = (x_j^1, x_j^2, \dots, x_j^{s_j})$ represents the proportion of players choosing each action within population j 's action set. The state transitions according to:

- Choose a population $N_j \in \{N_1, N_2, \dots, N_m\}$ with probability s_j/σ .
- Choose an action $\bar{a}_j^k \in \{\bar{a}_j^1, \bar{a}_j^2, \dots, \bar{a}_j^{s_j}\} = \bar{A}_j$ with probability $1/s_j$.
- If $x_j^k > 0$, i.e., at least one player from population j is playing action \bar{a}_j^k , choose $p \in \{p' \in N_j : \bar{a}_{p'} = \bar{a}_j^k\}$ uniformly at random to update according to (7). That is, transition to $(x_j + \frac{1}{n}(e_j^\ell - e_j^k), \mathbf{x}_{-j})$ with probability

$$\frac{e^{\beta\phi(x_j + \frac{1}{n}(e_j^\ell - e_j^k), \mathbf{x}_{-j})}}{\sum_{t=1}^{s_j} e^{\beta\phi(x_j + \frac{1}{n}(e_j^t - e_j^k), \mathbf{x}_{-j})}}$$

for each $\ell \in \{1, 2, \dots, s_j\}$.⁹

This defines transition probabilities in M for transitions from state $x = (x_j, \mathbf{x}_{-j}) \in X$ to a state of the form $y = (x_j + \frac{1}{n}(e_j^\ell - e_j^k), \mathbf{x}_{-j}) \in X$ in which a player from population N_j updates his action, so that

$$M(x, y) = \frac{e^{\beta\phi(x_j + \frac{1}{n}(e_j^\ell - e_j^k), \mathbf{x}_{-j})}}{\sigma \sum_{t=1}^{s_j} e^{\beta\phi(x_j + \frac{1}{n}(e_j^t - e_j^k), \mathbf{x}_{-j})}} \quad (38)$$

For a transition of any other form, $M(x, y) = 0$. Applying (27) to the chain with kernel M and global clock rate $\alpha\sigma n$, modified log-linear learning evolves as

$$\mu(t) = \mu(0)e^{\alpha\sigma n t(M-I)}. \quad (39)$$

Notation summary for stationary semi-anonymous potential games: Let $G = \{N, \{A_i\}, \{U_i\}\}$ be a stationary semi-anonymous potential game. The following summarizes the notation corresponding to game G .

- X - aggregate state space corresponding to the game G
- $\phi : X \rightarrow \mathbb{R}$ - the potential function corresponding to game G
- M - probability transition kernel for the modified log-linear learning process
- α - design parameter for modified log-linear learning which may be used to adjust the global update rate

⁹Agents' update rates are the only difference between our algorithm, standard log-linear learning, and the log-linear learning variant of [22]. In standard log-linear learning, players have uniform, constant clock rates. In our variant and the variant of [22], agents' update rates vary with the state. For the algorithm in [22], agent i 's update rate is $\alpha n / \bar{z}_i(t)$, where $\bar{z}_i(t)$ is the total number of players selecting the same action as agent i . The discrete time kernel of this process is as follows [22]: (1) Select an action $a_i \in \cup_{i \in N} A_i$ uniformly at random. (2) Select a player who is currently playing action a_i uniformly at random. This player updates its action according to (7). The two algorithms differ when at least two populations have overlapping action sets.

- $\mu(t) = \mu(0)e^{\alpha nt(M-I)}$ - distribution over state space X at time t when beginning with distribution $\mu(0)$ and following the modified log-linear learning process
- N_j - the j th population
- $n_j := |N_j|$ - the size of the j th population
- \mathcal{A}_j - action set for agents belonging to population N_j
- \bar{a}_j^k - the k th action in population N_j 's action set
- $s := |\cup_{j=1}^m \bar{\mathcal{A}}_j|$ - size of the union of all populations' action sets
- $s_j := |\bar{\mathcal{A}}_j|$ - size of population N_j 's action set
- $e_j^k \in \mathbb{R}^{s_j}$ - k th standard basis vector of length s_j
- $\sigma := \sum_{j=1}^m s_j$ - sum of sizes of each population's action set
- π - stationary distribution corresponding to the modified log-linear learning process for game G .
- $(x_j, x_{-j}) = (x_1, x_2, \dots, x_m) \in X$, a state in the aggregate state space, where $x_j = (x_j^1, x_j^2, \dots, x_j^{s_j})$.

C. Proof of Theorem 1

We require two supporting lemmas to prove Theorem 1. The first establishes the stationary distribution for modified log-linear learning as a function of β and characterizes how large β must be so the expected value of the potential function is within $\varepsilon/2$ of maximum. The second upper bounds the mixing time to within $\varepsilon/2$ of the stationary distribution for the modified log-linear learning process.

Lemma 1. *For the stationary semi-anonymous potential game $G = (N, \mathcal{A}_i, U_i)$ with state space X and potential function $\phi : X \rightarrow [0, 1]$, the stationary distribution for modified log-linear learning is*

$$\pi_x \propto e^{\beta \phi(x)}, \quad x \in X \quad (40)$$

Moreover, if condition (i) of Theorem 1 is satisfied and β is sufficiently large as in (8), then

$$\mathbb{E}_\pi[\phi(x)] \geq \max_{x \in X} \phi(x) - \varepsilon/2. \quad (41)$$

Proof: The form of the stationary distribution follows from standard reversibility arguments, using (38) and (40).

For the second part of the proof, define the following:

$$\begin{aligned} C_\beta &:= \sum_{x \in X} e^{\beta \phi(x)}, \\ x^* &:= \arg \max_{x \in X} \phi(x) \\ B(x^*, \delta) &:= \{x \in X : \|x - x^*\|_1 \leq \delta\} \end{aligned}$$

where $\delta \in [0, 1]$ is a constant which we will specify later. Because π is of exponential form with normalization factor C_β , the derivative of $\log C_\beta$ with respect to β is $\mathbb{E}_\pi[\phi(x)]$. Moreover, it follows from (40) that $\mathbb{E}_\pi[\phi(x)]$ is monotonically increasing in β , so we may proceed as follows:

$$\begin{aligned} \mathbb{E}_\pi[\phi(x)] &\geq \frac{1}{\beta} (\log C_\beta - \log C_0) \\ &= \phi(x^*) + \frac{1}{\beta} \log \frac{\sum_{x \in X} e^{\beta(\phi(x) - \phi(x^*))}}{|X|} \\ &\stackrel{(a)}{\geq} \phi(x^*) + \frac{1}{\beta} \log \frac{\sum_{x \in B(x^*, \delta)} e^{-\beta \delta \lambda}}{|X|} \end{aligned}$$

$$\begin{aligned} &= \phi(x^*) + \frac{1}{\beta} \log \frac{|B(x^*, \delta)| e^{-\beta \delta \lambda}}{|X|} \\ &= \phi(x^*) - \delta \lambda + \frac{1}{\beta} \log \left(\frac{|B(x^*, \delta)|}{|X|} \right) \end{aligned}$$

where (a) is from the fact that ϕ is λ -Lipschitz and the definition of $B(x^*, \delta)$. Using intermediate results in the proof of Lemma 6 of [22], $|B(x^*, \delta)|$ and $|X|$ are bounded as:

$$|B(x^*, \delta)| \geq \prod_{i=1}^m \left(\frac{\delta(n_i + 1)}{2ms_i} \right)^{s_i - 1}, \quad \text{and} \quad (42)$$

$$|X| \leq \prod_{i=1}^m (n_i + 1)^{s_i - 1}. \quad (43)$$

Now,

$$\begin{aligned} \mathbb{E}_\pi[\phi(x)] &\geq \phi(x^*) - \delta \lambda + \frac{1}{\beta} \log \left(\frac{|B(x^*, \delta)|}{|X|} \right) \\ &\geq \phi(x^*) - \delta \lambda + \frac{1}{\beta} \log \left(\frac{\prod_{i=1}^m \left(\frac{\delta(n_i + 1)}{2ms_i} \right)^{s_i - 1}}{\prod_{i=1}^m (n_i + 1)^{s_i - 1}} \right) \\ &= \phi(x^*) - \delta \lambda + \frac{1}{\beta} \log \prod_{i=1}^m \left(\frac{\delta}{2ms_i} \right)^{s_i - 1} \\ &\geq \phi(x^*) - \delta \lambda + \frac{m(s-1)}{\beta} \log \left(\frac{\delta}{2ms} \right) \end{aligned}$$

Consider two cases: (i) $\lambda \leq \varepsilon/4$, and (ii) $\lambda > \varepsilon/4$. For case (i), choose $\delta = 1$ and let $\beta \geq \frac{4m(s-1)}{\varepsilon} \log 2ms$. Then,

$$\begin{aligned} \mathbb{E}_\pi[\phi(x)] &\geq \phi(x^*) - \delta \lambda + \frac{m(s-1)}{\beta} \log \left(\frac{\delta}{2ms} \right) \\ &\geq \phi(x^*) - \varepsilon/4 - \frac{m(s-1)}{\beta} \log 2ms \\ &\geq \phi(x^*) - \varepsilon/4 - \frac{\varepsilon m(s-1)}{4m(s-1) \log 2ms} \log 2ms \\ &= \phi(x^*) - \varepsilon/2 \end{aligned}$$

For case (ii), note that $\lambda > \varepsilon/4 \implies \delta = \varepsilon/4\lambda < 1$ so we may choose $\delta = \varepsilon/4\lambda$. Let $\beta \geq \frac{4m(s-1)}{\varepsilon} \log \left(\frac{8\lambda ms}{\varepsilon} \right)$. Then

$$\begin{aligned} \mathbb{E}_\pi[\phi(x)] &\geq \phi(x^*) - \delta \lambda + \frac{m(s-1)}{\beta} \log \left(\frac{\delta}{2ms} \right) \\ &= \phi(x^*) - \varepsilon/4 + \frac{m(s-1)}{\beta} \log \left(\frac{\varepsilon}{8\lambda ms} \right) \\ &= \phi(x^*) - \varepsilon/4 - \frac{m(s-1)}{\beta} \log \left(\frac{8\lambda ms}{\varepsilon} \right) \\ &\geq \phi(x^*) - \varepsilon/4 - \frac{\varepsilon m(s-1)}{4m(s-1) \log \left(\frac{8\lambda ms}{\varepsilon} \right)} \log \left(\frac{8\lambda ms}{\varepsilon} \right) \\ &= \phi(x^*) - \varepsilon/2 \end{aligned}$$

as desired. \square

Lemma 2. *For the Markov chain defined by modified log-linear learning with kernel M and stationary distribution π , if the number of players within each population satisfies condition (ii) of Theorem 1, and t is sufficiently large as in (10), then*

$$\|\mu(t) - \pi\|_{TV} \leq \varepsilon/2. \quad (44)$$

Proof: We begin by establishing a lower bound on the Sobolev constant for the Markov chain, M . We claim that, for the Markov chain M defined in Appendix B, if $\phi : X \rightarrow [0, 1]$ and $m + \sum_{i=1}^m n_i^2 \geq \sigma$, then

$$\rho(M) \geq \frac{e^{-3\beta}}{c_1 m(m(s-1))! 2n^2} \quad (45)$$

for some constant c_1 which depends only on s . Then, from (37), a lower bound on the Sobolev constant yields an upper bound on the mixing time for the chain M .

Using the technique of [22], we compare the Sobolev constants for the chain M and a similar random walk on a convex set. The primary difference is that our proof accounts for dependencies on the number of populations, m , whereas theirs considers only the $m = 1$ case. As a result, our state space is necessarily larger. We accomplish this proof in four steps. In step 1, we define M^* to be the Markov chain M with $\beta = 0$, and establish the bound $\rho(M) \geq e^{-3\beta} \rho(M^*)$. In step 2, we define a third Markov chain, M^\dagger , and establish the bound $\rho(M^*) \geq \frac{1}{s} \rho(M^\dagger)$. Then, in step 3, we establish a lower bound on the Sobolev constant of M^\dagger . Finally, in step 4, we combine the results of the first three steps to establish (45). We now prove each step in detail.

Step 1, M to M^* : Let M^* be the Markov chain M with $\beta = 0$, and let π^* be its stationary distribution. In M^* an updating agent chooses his next action uniformly at random. Per Equation (40) with $\beta = 0$, the stationary distribution π^* of M^* is the uniform distribution. Let $x, y \in X$. We bound π_x / π_x^* and $M(x, y) / M^*(x, y)$ in order to use Corollary 3.15 in [21]:

$$\frac{\pi_x}{\pi_x^*} = \frac{e^{\beta\phi(x)}}{\sum_{y \in X} e^{\beta\phi(y)}} \cdot \frac{\sum_{y \in X} e^0}{e^0} = \frac{|X| e^{\beta\phi(x)}}{\sum_{y \in X} e^{\beta\phi(y)}}$$

Since $\phi(x) \in [0, 1]$ for all $x \in X$, this implies

$$e^{-\beta} \leq \frac{\pi_x}{\pi_x^*} \leq e^\beta \quad (46)$$

Similarly, for $y = (x_j + \frac{1}{n}(e_j^k - e_j^\ell), x_{-j})$,

$$\frac{M(x, y)}{M^*(x, y)} = \frac{s_j e^{\beta\phi(y)}}{\sum_{r=1}^{s_j} e^{\beta\phi(x_j + \frac{1}{n}(e_i^k - e_i^r), x_{-j})}}$$

Since $\phi(x) \in [0, 1]$ for all $x \in X$, for any $x, y \in X$ of the above form,

$$e^{-\beta} \leq \frac{M(x, y)}{M^*(x, y)} \leq e^\beta. \quad (47)$$

For a transition to any y not of the form above, $M(x, y) = M^*(x, y) = 0$. Using this fact and Equations (46) and (47), we apply Corollary 3.15 in [21]:

$$\rho(M) \geq e^{-3\beta} \rho(M^*). \quad (48)$$

Step 2, M^* to M^\dagger : Consider the Markov chain M^\dagger on X , where transitions from state x to y occur as follows:

- Choose a population N_j with probability s_j / σ
- Choose $k \in \{1, \dots, s_j - 1\}$ and choose $\kappa \in \{-1, 1\}$, each uniformly at random.
 - If $\kappa = -1$ and $x_j^k > 0$, then $y = (x_j + \frac{1}{n}(e_j^{s_j} - e_j^k), x_{-k})$.
 - If $\kappa = 1$ and $x_j^{s_j} > 0$, then $y = (x_j + \frac{1}{n}(e_j^k - e_j^{s_j}), x_{-j})$.

Since $M^\dagger(x, y) = M^\dagger(y, x)$ for any $x, y \in X$, M^\dagger is reversible with the uniform distribution over X . Hence the stationary distribution is uniform, and $\pi^\dagger = \pi^*$.

For a transition x to y in which an agent from population N_j changes his action, $M^*(x, y) \geq \frac{1}{s_j} M^\dagger(x, y)$, implying

$$M^*(x, y) \geq \frac{1}{s} M^\dagger(x, y), \quad \forall x, y \in X \quad (49)$$

since $s \geq s_j, \forall i \in \{1, \dots, m\}$. Using (49) and the fact that $\pi^* = \pi^\dagger$, we apply Corollary 3.15 from [21]:

$$\rho(M^*) \geq \frac{1}{s} \rho(M^\dagger) \quad (50)$$

Step 3, M^\dagger to a random walk: The following random walk on

$$C = \left\{ (z_1, \dots, z_m) \in \mathbb{Z}_+^{\sigma-m} : z_j \in \mathbb{Z}_+^{s_j-1}, \sum_{k=1}^{s_j-1} z_j^k \leq n_j, \forall j \right\}$$

is equivalent to M^\dagger . Transition from $x \rightarrow y$ in C as follows:

- Choose $j \in [\sigma - m]$ and $\kappa \in \{-1, 1\}$, each uniformly at random
- $y = \begin{cases} x + \kappa e_j & \text{if } x + \kappa e_j \in C \\ x & \text{otherwise} \end{cases}$.

The stationary distribution of this random walk is uniform. We lower bound the Sobolev constant, $\rho(M^\dagger)$, which, using the above steps, lower bounds $\rho(M)$ and hence upper bounds the mixing time of our algorithm.

Let $g : C \rightarrow \mathbb{R}$ be an arbitrary function. To lower bound $\rho(M^\dagger)$, we will lower bound $\mathcal{E}(g, g)$ and upper bound $\mathcal{L}(g)$. The ratio of these two bounds in turn lower bounds the ratio $\mathcal{E}(g, g) / \mathcal{L}(g)$; since g was chosen arbitrarily this also lower bounds the Sobolev constant. We will use a theorem due to [?] which applies to an extension of a function $g : C \rightarrow \mathbb{R}$ to a function defined over the convex hull of C ; here we define this extension.

Let K be the convex hull of C . Given $g : C \rightarrow \mathbb{R}$, we follow the procedure of [?], [22] to extend g to a function $g_\varepsilon : K \rightarrow \mathbb{R}$. For $x \in C$, let $C(x)$ and $C(x, \varepsilon)$ be the $\sigma - m$ dimensional cubes of center x and sides 1 and $1 - 2\varepsilon$ respectively. For sufficiently small $\varepsilon > 0$ and $z \in C(x)$, define $g_\varepsilon : K \rightarrow \mathbb{R}$ by:

$$g_\varepsilon(z) := \begin{cases} g(x) & \text{if } z \in C(x, \varepsilon) \\ \frac{(1+\eta(z))g(x) + (1-\eta(z))g(y)}{2} & \text{otherwise} \end{cases}$$

where $y \in C$ is a point such that $D := C(x) \cap C(y)$ is the closest face of $C(x)$ to z (if more than one y satisfy this condition, one such point may be chosen arbitrarily), and $\eta := \frac{\text{dist}(z, D)}{\varepsilon} \in [0, 1]$. The dist function represents standard Euclidean distance in $\mathbb{R}^{\sigma-m}$.

Define

$$I_\varepsilon := \int_K |\nabla g_\varepsilon(z)|^2 dz \quad (51)$$

$$J_\varepsilon := \int_K g_\varepsilon(z)^2 \log \frac{g_\varepsilon(z)^2 \text{vol}(K)}{\int_K g_\varepsilon(y)^2 dy} dz. \quad (52)$$

Applying Theorem 2 of [?] for $K \in \mathbb{R}^{\sigma-m}$ with diameter

$$\sqrt{\sum_{i=1}^m n_i^2}, \text{ if } m + \sum_{i=1}^m n_i^2 \geq \sigma, \\ \frac{\varepsilon I_\varepsilon}{J_\varepsilon} \geq \frac{1}{A \sum_{i=1}^m n_i^2}. \quad (53)$$

We lower bound $\mathcal{E}(g, g)$ in terms of $\varepsilon I_\varepsilon$ and then upper bound $\mathcal{L}(g)$ in terms of J_ε to obtain a lower bound on their ratio with Equation (53). The desired lower bound on the Sobolev constant follows.

Using similar techniques to [22], we lower bound $\mathcal{E}(g, g)$ in terms of $\varepsilon I_\varepsilon$ as

$$I_\varepsilon \leq \frac{|C|(\sigma - m)}{\varepsilon} \mathcal{E}(g, g) + O(1).$$

Then, $\varepsilon I_\varepsilon \leq_{\varepsilon \rightarrow 0} |C|(\sigma - m)\mathcal{E}(g, g)$, and hence

$$\mathcal{E}(g, g) \geq_{\varepsilon \rightarrow 0} \frac{\varepsilon I_\varepsilon}{|C|(\sigma - m)}. \quad (54)$$

Again, using similar techniques as [22], we bound J_ε as

$$\frac{J_\varepsilon}{\text{vol}(K)} \geq \frac{|C|}{2^{2(\sigma-m)} \text{vol}(K)(\sigma - m)!^2} \mathcal{L}(f).$$

Then

$$\mathcal{L}(g) \leq_{\varepsilon \rightarrow 0} \frac{2^{2(\sigma-m)}(\sigma - m)!^2}{|C|} J_\varepsilon. \quad (55)$$

Step 4, Combining inequalities: Using inequalities (53), (54), and (55),

$$\frac{\mathcal{E}(f, f)}{\mathcal{L}(f)} \geq \frac{1}{2^{2(\sigma-m)} A(\sigma - m)(\sigma - m)!^2 \sum_{i=1}^m n_i^2}, \quad (56)$$

$\forall f : C \rightarrow \mathbb{R}$. Therefore,

$$\rho(M^\dagger) = \min_{f: C \rightarrow \mathbb{R}} \frac{\mathcal{E}(f, f)}{\mathcal{L}(f)} \\ \geq \frac{1}{2^{2(\sigma-m)} A(\sigma - m)(\sigma - m)!^2 \sum_{i=1}^m n_i^2} \quad (57)$$

Combining equations (48), (50), and (57)

$$\rho(M) \geq \frac{e^{-3\beta}}{2^{2ms} c_1 m^2 (m(s-1))!^2 n^2}$$

where c_1 is a constant depending only on s .

From here, Lemma 2 follows by applying Equation (37) in a similar manner as the proof of Equation (23) in [22]. The main difference is that the size of the state space is bounded as $|X| \leq \prod_{i=1}^m (n_i + 1)^{s_i+1}$ due to the potential for multiple populations. \square

Combining Lemmas 1 and 2 results in a bound on the time it takes for the expected potential to be within ε of the maximum, provided β is sufficiently large. The lemmas and method of proof for Theorem 1 follow the structure of the supporting lemmas and proof for Theorem 3 in [22]. The main differences have arisen due to the facts that i) our analysis considers the multi-population case, so the size of our state space cannot be reduced as significantly as in the single population case of [22], and ii) update rates in our algorithm depend on behavior within each agent's own population, instead of on global behavior.

Proof of Theorem 1:

From Lemma 1, if condition (i) of Theorem 1 is satisfied and β is sufficiently large as in (8), then $\mathbb{E}_\pi[\phi(x)] \geq \max_{x \in X} \phi(x) - \varepsilon/2$. From Lemma 2, if condition (ii) of Theorem 1 is satisfied,

and t is sufficiently large as in (10), then $\|\mu(t) - \pi\|_{TV} \leq \varepsilon/2$. Then

$$\mathbb{E}[\phi(a(t)|_X)] = \mathbb{E}_{\mu(t)}[\phi(x)] \\ \geq \mathbb{E}_\pi[\phi(x)] - \|\mu(t) - \pi\|_{TV} \cdot \max_{x \in X} \phi(x) \\ \stackrel{(a)}{\geq} \max_{x \in X} \phi(x) - \varepsilon$$

where (a) follows from (41), (44), and the fact that $\phi(x) \in [0, 1]$. \square



Holly Borowski is a Ph.D. candidate in the Aerospace Engineering Sciences Department at the University of Colorado, Boulder. Holly received a BS in Mathematics in 2004 at the US Air Force Academy. She received the NASA Aeronautics graduate scholarship (2012), the Zonta International Amelia Earhart scholarship (2014), and the Philanthropic Educational Organization scholarship (2015). Holly's research focuses on the role of information in the control of distributed multi-agent systems.



Jason Marden is an Assistant Professor in the Department of Electrical and Computer Engineering at the University of California, Santa Barbara. Jason received a BS in Mechanical Engineering in 2001 from UCLA, and a PhD in Mechanical Engineering in 2007, also from UCLA, under the supervision of Jeff S. Shamma, where he was awarded the Outstanding Graduating PhD Student in Mechanical Engineering. After graduating from UCLA, he served as a junior fellow in the Social and Information Sciences Laboratory at the California Institute of Technology until 2010 when he joined the University of Colorado. Jason is a recipient of the NSF Career Award (2014), the ONR Young Investigator Award (2015), the AFOSR Young Investigator Award (2012), and the American Automatic Control Council Donald P. Eckman Award (2012). Jason's research interests focus on game theoretic methods for the control of distributed multi-agent systems.