

Revisiting Log-Linear Learning: Asynchrony, Completeness and Payoff-Based Implementation*

Jason R. Marden[†]

Jeff S. Shamma^{‡§}

September 16, 2008

Abstract

Log-linear learning is a learning algorithm with equilibrium selection properties. In potential games, log-linear learning provides guarantees on the percentage of time that the joint action profile will be at a potential maximizer. The traditional analysis of log-linear learning has centered around explicitly computing the stationary distribution. This analysis relied on a highly structured setting: i) players' utility functions constitute an exact potential game, ii) players update their strategies one at a time, which we refer to as *asynchrony*, iii) at any stage, a player can select any action in the action set, which we refer to as *completeness*, and iv) each player is endowed with the ability to assess the utility he would have received for any alternative action provided that the actions of all other players remain fixed. Since the appeal of log-linear learning is not solely the explicit form of the stationary distribution, we seek to address to what degree one can relax the structural assumptions while maintaining that only potential function maximizers are the stochastically stable action profiles. In this paper, we introduce slight variants of log-linear learning to include both synchronous updates and incomplete action sets. In both settings, we prove that only potential function maximizers are stochastically stable. Furthermore, we introduce a payoff-based version of log-linear learning, in which players are only aware of the utility they received and the action that they played. Note that log-linear learning in its original form is not a payoff-based learning algorithm. In payoff-based log-linear learning, we also prove that only potential maximizers are stochastically stable. The key enabler for these results is to change the focus of the analysis away from deriving the explicit form of the stationary distribution of the learning process towards characterizing the stochastically stable states. The resulting analysis uses the theory of resistance trees for regular perturbed Markov decision processes, thereby allowing a relaxation of the aforementioned structural assumptions.

1 Introduction

The theory of learning in games has sought to understand how and why equilibria emerge in non-cooperative games. Traditionally, social science literature develops *descriptive* game theoretic models for players, analyzes the limiting behavior, and generalizes the results for larger classes of games. Recently, there has been

*Research supported by the Social and Information Sciences Laboratory at the California Institute of Technology, AFOSR grants #FA9550-08-1-0375 and #FA9550-05-1-0321, and NSF grant #ECS-0501394.

[†]J. R. Marden is with the Social and Information Sciences Laboratory, California Institute of Technology, MC 136-93, Pasadena, CA 91107, marden@caltech.edu.

[‡]J. S. Shamma is with the School of Electrical and Computer Engineering, Georgia Institute of Technology, 777 Atlantic Dr NW, Atlanta, GA 30332-0250, shamma@gatech.edu.

[§]Corresponding author.

a significant amount of research seeking to understand these behavioral models not from a descriptive point of view, but rather from a *prescriptive* point of view [2, 7, 8, 17]. The goal is to use these behavioral models as a prescriptive control approach in distributed multi-agent systems where the guaranteed limiting behavior would represent a desirable operating condition.

It turns out that a design of such distributed systems is strongly related to a class of non-cooperative games known as *potential games*, or more generally *weakly acyclic games* [8]. In many multi-agent systems, local agent objectives are appropriately aligned with the global objective. It is precisely this alignment condition that connects the realms of cooperative control and potential games. This connection is important for two main reasons. First, potential games provide a paradigm for designing, analyzing and controlling multi-agent systems. Second, potential games have been studied extensively in the game theory literature and several established learning algorithms with guaranteed asymptotic results could be implemented to control these multi-agent systems.

Most of the learning algorithms for potential games guarantee convergence to a Nash equilibrium. A representative sampling of these algorithms include Fictitious Play [12], Joint Strategy Fictitious Play [10], Adaptive Play [20], and many others [21, 22, 11, 9]. However, in potential games a pure Nash equilibrium may be inefficient with regards to the global objective. In certain settings, such as a congestion game with linear cost functions, it may be possible to bound the efficiency, but in general, such a bound need not be possible [14, 18]. Therefore, from a design perspective, having a learning algorithm that guarantees convergence to the most efficient Nash equilibrium is desirable.

Log-linear learning, originally introduced in [4], is one of the few learning dynamics that embodies this notion of equilibrium selection. In exact potential games, log-linear learning guarantees that only the joint action profiles that maximize the potential function are stochastically stable. The enabler for these results is the introduction of *noise* into the decision making process [20, 4, 5, 6]. In log-linear learning, this noise permits players to occasionally make mistakes, where mistakes represent the selection of suboptimal actions. The structure of the noise in log-linear learning is such that the probability of selecting a suboptimal action is associated with the magnitude of the payoff difference associated with a best response and the suboptimal action. As the noise vanishes, the probability that a player selects a suboptimal action goes to zero.

The traditional analysis of log-linear learning has centered around explicitly computing the stationary distribution. This analysis relies on a highly structured setting:

- (i) Players' utility functions constitute an exact potential game.
- (ii) Players update their strategies one at a time, which we refer to as *asynchrony*.
- (iii) At any stage, a player can select any action in the action set, which we refer to as *completeness*.
- (iv) Each player is endowed with the ability to assess the utility he would have received for any alternative action provided that the actions of all other players remain fixed.

Nonetheless, log-linear learning has received significant research attention [21, 5, 6, 8, 2, 15, 1]. These results range from analyzing convergence rates [15] to the necessity of the structural requirements [1]. In particular, [1] demonstrates that if the structural requirements of (i) and (ii) are relaxed arbitrarily, then the equilibrium selection properties of log-linear learning are no longer guaranteed.

Since the appeal of log-linear learning is not solely the explicit form of the stationary distribution, we seek to address to what degree one can relax the structural assumptions while maintaining that only potential function maximizers are stochastically stable. Our motivation for this relaxation is that in many distributed systems, the structured setting of log-linear learning may inhibit its applicability as a distributed control mechanism.

In this paper, we introduce slight variants of log-linear learning to include both synchronous updates and incomplete action sets. In both settings, we prove that only potential function maximizers are stochastically stable. Furthermore, we introduce a payoff-based version of log-linear learning, in which players are only aware of the utility they received and the action that they played. Note that log-linear learning in its original form is not a payoff-based learning algorithm. In the payoff-based log-linear learning, we also prove that only potential maximizers are stochastically stable. The key enabler for these results is to change the focus of the analysis away from deriving the explicit form of the stationary distribution of the learning process towards characterizing the stochastically stable states. The resulting analysis uses the theory of resistance trees for regular perturbed Markov decision processes [20].

The outline of the paper is as follows. In Section 2 we review the game theoretic concepts that we will use in this paper. In Section 3 we prove that log-linear learning in its original form is a regular perturbed Markov process. Utilizing this connection, we reprove the equilibrium selection properties of log-linear learning using the theory of resistance trees for regular perturbed Markov processes. Furthermore, we analyze the convergence properties of log-linear learning outside the realm of exact potential games. In Section 4, we analyze a variant of log-linear learning that includes synchronous updates. In Section 5, we show that if action sets are not complete, then potential function maximizers may not be stochastically stable in exact potential games. Furthermore, we derive a variant of the traditional log-linear learning that rectifies this problem. In Section 6, we introduce a payoff-based version of log-linear learning that maintains the equilibrium selection properties of log-linear learning. Lastly, in Section 7 we provide some concluding remarks.

2 Setup

We consider a finite strategic-form game with n -player set $\mathcal{I} := \{1, \dots, n\}$ where each player i has an action set \mathcal{A}_i and a utility function $U_i : \mathcal{A} \rightarrow \mathbb{R}$ where $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$.

For an action profile $a = (a_1, a_2, \dots, a_n) \in \mathcal{A}$, let a_{-i} denote the profile of player actions *other than* player i , i.e.,

$$a_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n).$$

With this notation, we will sometimes write a profile a of actions as (a_i, a_{-i}) . Similarly, we may write $U_i(a)$ as $U_i(a_i, a_{-i})$. Let $\mathcal{A}_{-i} = \prod_{j \neq i} \mathcal{A}_j$ denote the set of possible collective actions of all players other than player i . We define player i 's *best response set* for an action profile $a_{-i} \in \mathcal{A}_{-i}$ as

$$B_i(a_{-i}) := \{a_i^* \in \mathcal{A}_i : U_i(a_i^*, a_{-i}) = \max_{a_i \in \mathcal{A}_i} U_i(a_i, a_{-i})\}.$$

In this paper, we will focus on three classes of games: identical interest, potential, and weakly acyclic. In each of these classes of games, a pure Nash equilibrium is guaranteed to exist. Each class of games imposes a restriction on the admissible utility functions.

2.0.1 Identical Interest Games

The most restrictive class of games that we will consider is identical interest games. In such a game, the players' utility functions $\{U_i\}_{i=1}^n$ are the same. That is, for some function $\phi : \mathcal{A} \rightarrow \mathbb{R}$,

$$U_i(a) = \phi(a),$$

for every player $i \in \mathcal{I}$ and for every $a \in \mathcal{A}$. It is easy to verify that all identical interest games have at least one pure Nash equilibrium, namely any action profile a that maximizes $\phi(a)$.

2.0.2 Potential Games

A significant generalization of an identical interest game is a potential game [13]. In a potential game, the change in a player's utility that results from a unilateral change in strategy equals the change in the global utility. Specifically, there is a function $\phi : \mathcal{A} \rightarrow \mathbb{R}$ such that for every player $i \in \mathcal{I}$, for every $a_{-i} \in \mathcal{A}_{-i}$, and for every $a'_i, a''_i \in \mathcal{A}_i$,

$$U_i(a'_i, a_{-i}) - U_i(a''_i, a_{-i}) = \phi(a'_i, a_{-i}) - \phi(a''_i, a_{-i}). \quad (1)$$

When this condition is satisfied, the game is called an (exact) potential game with the potential function ϕ . In potential games, any action profile maximizing the potential function is a pure Nash equilibrium, hence every potential game possesses at least one such equilibrium.

We will also consider a more general class of potential games known as *generalized ordinal potential games*. In generalized ordinal potential games there is a global function $\phi : \mathcal{A} \rightarrow \mathbb{R}$ such that for every player $i \in \mathcal{I}$, for every $a_{-i} \in \mathcal{A}_{-i}$, and for every $a'_i, a''_i \in \mathcal{A}_i$,

$$U_i(a'_i, a_{-i}) - U_i(a''_i, a_{-i}) > 0 \Leftrightarrow \phi(a'_i, a_{-i}) - \phi(a''_i, a_{-i}) > 0.$$

2.0.3 Weakly Acyclic Games

Consider any game with a set \mathcal{A} of action profiles. A *better reply path* is a sequence of action profiles a^1, a^2, \dots, a^L such that, for every $1 \leq \ell \leq L - 1$, there is exactly one player \mathcal{P}_{i_ℓ} such that i) $a_{i_\ell}^\ell \neq a_{i_\ell}^{\ell+1}$, ii) $a_{-i_\ell}^\ell = a_{-i_\ell}^{\ell+1}$, and iii) $U_{i_\ell}(a^\ell) < U_{i_\ell}(a^{\ell+1})$. In other words, one player moves at a time and that player

increases his own utility. A *best reply path* is a better reply path with the additional requirement that each unilateral deviation is the result of a best response, i.e., $a_i^{\ell+1} \in B_i(a_{-i}^\ell)$.

Consider any potential game with potential function ϕ . Starting from an arbitrary action profile $a \in \mathcal{A}$, construct a better reply path $a = a^1, a^2, \dots, a^L$ until it can no longer be extended. Note first that such a path cannot cycle back on itself, because ϕ is strictly increasing along the path. Since \mathcal{A} is finite, the path cannot be extended indefinitely. Hence, the last element in a maximal better reply path from any joint action, a , must be a Nash equilibrium.

This idea may be generalized as follows. A game is *weakly acyclic* if for any $a \in \mathcal{A}$, there exists a better reply path starting at a and ending at some pure Nash equilibrium [21, 22]. A game is *weakly acyclic under best replies* if for any $a \in \mathcal{A}$, there exists a best reply path starting at a and ending at some pure Nash equilibrium [21, 22]. Potential games are special cases of weakly acyclic games.

An equivalent definition of weakly acyclic games utilizing potential functions is given in [8]. A game is weakly acyclic if and only if there exists a potential function $\phi : \mathcal{A} \rightarrow \mathbb{R}$ such that for any action $a \in \mathcal{A}$ that is *not* a Nash equilibrium, there exists a player $i \in \mathcal{I}$ with an action $a'_i \in \mathcal{A}_i$ such that $U_i(a'_i, a_{-i}) > U_i(a_i, a_{-i})$ and $\phi(a'_i, a_{-i}) > \phi(a_i, a_{-i})$.

3 Log Linear Learning

In a repeated game, at each time $t \in \{0, 1, 2, \dots\}$, each player $i \in \mathcal{I}$ simultaneously chooses an action $a_i(t) \in \mathcal{A}_i$ and receives the utility $U_i(a(t))$ where $a(t) := (a_1(t), \dots, a_n(t))$. The action of player i is chosen at time t according to a probability distribution $p_i(t) \in \Delta(\mathcal{A}_i)$ where $p_i^{a_i}(t)$ denotes the probability that player i will select action a_i . We will refer to $p_i(t)$ as the *strategy* of player i at time t . A player's strategy at time t can rely only on information from the games played at times $\{0, 1, 2, \dots, t-1\}$. Different learning algorithms are specified by both the assumptions on available information and the mechanism by which the strategies are updated as information is gathered.

The following learning algorithm is known as log-linear learning [4]. At each time $t > 0$, one player $i \in \mathcal{I}$ is randomly chosen and allowed to alter his current action. All other players must repeat their current action at the ensuing time step, i.e. $a_{-i}(t) = a_{-i}(t-1)$. At time t , player i employs the strategy

$p_i(t) \in \Delta(\mathcal{A}_i)$ where

$$p_i^{a_i}(t) = \frac{e^{\frac{1}{\tau}U_i(a_i, a_{-i}(t-1))}}{\sum_{\bar{a}_i \in \mathcal{A}_i} e^{\frac{1}{\tau}U_i(\bar{a}_i, a_{-i}(t-1))}}, \quad (2)$$

for any action $a_i \in \mathcal{A}_i$ and temperature $\tau > 0$. The temperature τ determines how likely player i is to select a suboptimal action. As $\tau \rightarrow \infty$, player i will select any action $a_i \in \mathcal{A}_i$ with equal probability. As $\tau \rightarrow 0$, player i will select a best response to the action profile $a_{-i}(t-1)$, i.e., $a_i(t) \in B_i(a_{-i}(t-1))$ with arbitrarily high probability. In the case of a non-unique best response, player i will select a best response at random (uniformly).

Consider any exact potential game with potential function $\phi : \mathcal{A} \rightarrow \mathbb{R}$. In the repeated potential game in which all players adhere to log-linear learning, the stationary distribution of the joint action profiles is $\mu \in \Delta(\mathcal{A})$ where [4]

$$\mu(a) = \frac{e^{\frac{1}{\tau}\phi(a)}}{\sum_{\bar{a} \in \mathcal{A}} e^{\frac{1}{\tau}\phi(\bar{a})}}. \quad (3)$$

One can interpret the stationary distribution μ as follows. For sufficiently large times $t > 0$, $\mu(a)$ equals the probability that $a(t) = a$. As one decreases the temperature, $\tau \rightarrow 0$, all the weight of the stationary distribution μ is on the joint actions that maximize the potential function.

The above analysis characterizes the precise stationary distribution as a function of the temperature τ . The importance of the result is not solely the explicit form of the stationary distribution, but rather the recognition that as $\tau \rightarrow 0$, the only stochastically stable states of the process are the joint actions that maximize the potential function. From this point on, rather than looking for the explicit stationary distribution, we will focus on analyzing the stochastically stable states of the process using the theory of resistance trees for regular perturbed Markov decision processes. This relaxation will allow us to modify the traditional log-linear learning to include synchronous updates, incomplete action sets, and a payoff based implementation.

3.1 Proof for Log-Linear Learning Using Theory of Resistance Trees

Before discussing variants of log-linear learning, we will reprove the original result using the theory of resistance trees. The proof approach presented in this section can also be found in [3, 5]. We present

the proof in detail for two reasons: i) the proof approach is similar to the forthcoming arguments for the variations of log-linear learning presented in the subsequent sections and ii) the structure of the proof can be exploited to analyze the limiting behavior of log-linear learning in games outside the realm of exact potential games.

We present an overview of the theory of resistance trees for regular perturbed Markov decision processes in the Appendix. See [20] for a more thorough review.

Before proving that the only stochastically stable states are the potential function maximizers, we will prove two lemmas. The first lemma establishes that log-linear learning induces a regular perturbed Markov decision process.

Lemma 3.1 *Log-linear learning induces a regular perturbed Markov process where the unperturbed Markov process is an asynchronous best reply process and the resistance of any feasible transition $a^0 \rightarrow a^1 = (a_i^1, a_{-i}^0)$ is*

$$R(a^0 \rightarrow a^1) = \max_{a_i^* \in \mathcal{A}_i} U_i(a_i^*, a_{-i}^0) - U_i(a^1). \quad (4)$$

Proof: The unperturbed process, P^0 , is the asynchronous best reply process described as follows. At each time $t > 0$, one player $i \in \mathcal{I}$ is randomly chosen and allowed to alter his current action. All other players must repeat their current action at the ensuing time step, i.e. $a_{-i}(t) = a_{-i}(t-1)$. At time t , player i selects a best response to the action profile $a_{-i}(t-1)$, i.e., $a_i(t) \in B_i(a_{-i}(t-1))$. In the case of multiple best responses, player i selects one at random (uniformly).

Log-linear learning induces a finite, aperiodic, irreducible process over the state space \mathcal{A} . We will analyze this process with respect to $\epsilon := e^{-\frac{1}{\tau}}$ rather than the temperature τ . Let P^ϵ denote the associated transition probability matrix. The probability of transitioning from a^0 to $a^1 := (a_i^1, a_{-i}^0)$ is

$$P_{a^0 \rightarrow a^1}^\epsilon = \frac{1}{n} \frac{\epsilon^{-U_i(a_i^1, a_{-i}^0)}}{\sum_{a_i \in \mathcal{A}_i} \epsilon^{-U_i(a_i, a_{-i}^0)}}. \quad (5)$$

We will assume for convenience that the updating player is selected randomly according to a uniform distri-

bution.¹ Let the maximum utility of player i for any action profile $a_{-i} \in \mathcal{A}_{-i}$ be defined as

$$V_i(a_{-i}) := \max_{a_i \in \mathcal{A}_i} U_i(a_i, a_{-i}).$$

Multiplying the numerator and denominator of (5) by $\epsilon^{V_i(a_{-i}^0)}$, we obtain

$$P_{a^0 \rightarrow a^1}^\epsilon = \frac{1}{n} \frac{\epsilon^{V_i(a_{-i}^0) - U_i(a_i^1, a_{-i}^0)}}{\sum_{a_i \in \mathcal{A}_i} \epsilon^{V_i(a_{-i}^0) - U_i(a_i, a_{-i}^0)}}.$$

Accordingly,

$$\lim_{\epsilon \rightarrow 0^+} \frac{P_{a^0 \rightarrow a^1}^\epsilon}{\epsilon^{V_i(a_{-i}^0) - U_i(a_i^1, a_{-i}^0)}} = \frac{1}{n|B_i(a_{-i}^0)|}.$$

where $|B_i(a_{-i}^0)|$ denotes the size, or number of actions, in player i 's best response set. This implies that the process P^ϵ is a regular perturbed Markov process where the resistance of the transition $a^0 \rightarrow a^1$ is

$$R(a^0 \rightarrow a^1) = V_i(a_{-i}^0) - U_i(a_i^1, a_{-i}^0). \quad (6)$$

Notice that $R(a^0 \rightarrow a^1) \geq 0$.

□

Before stating the second lemma we will introduce some notation. A feasible action path \mathcal{P} is a sequence of joint actions

$$\mathcal{P} = \{a^0 \rightarrow a^1 \rightarrow \dots \rightarrow a^m\}$$

that are the result of unilateral deviations, i.e., for each $k \in \{1, 2, \dots, m\}$, $a^k = (a_i, a_{-i}^{k-1})$ for some player i and action $a_i \in \mathcal{A}_i$. The resistance of a path \mathcal{P} is the sum of the resistance of each edge

$$R(\mathcal{P}) = \sum_{k=1}^m R(a^{k-1} \rightarrow a^k).$$

Lemma 3.2 *Consider any finite n -player exact potential game with potential function $\phi : \mathcal{A} \rightarrow \mathbb{R}$ where all players adhere to log-linear learning. For any feasible action path*

$$\mathcal{P} = \{a^0 \rightarrow a^1 \rightarrow \dots \rightarrow a^m\}$$

and reverse path

$$\mathcal{P}^R = \{a^m \rightarrow a^{m-1} \rightarrow \dots \rightarrow a^0\},$$

¹The player selection process could be relaxed to any probability distribution where the probability of selecting a given player is bounded away from 0.

the difference in the total resistance across the paths is

$$R(\mathcal{P}) - R(\mathcal{P}^R) = \phi(a^0) - \phi(a^m).$$

Proof: We will start by analyzing any edge in the feasible path $a^k \rightarrow a^{k+1}$. Suppose that $a^{k+1} = (a'_i, a^k_{-i})$ for some player i . The resistance across this edge is

$$R(a^k \rightarrow a^{k+1}) = V_i(a^k_{-i}) - U_i(a^{k+1}).$$

The resistance across the reverse edge $a^{k+1} \rightarrow a^k$ is

$$\begin{aligned} R(a^{k+1} \rightarrow a^k) &= V_i(a^{k+1}_{-i}) - U_i(a^k), \\ &= V_i(a^k_{-i}) - U_i(a^k), \end{aligned}$$

where the second equality is because $a^k_{-i} = a^{k+1}_{-i}$. The difference in the resistances across this edge is

$$\begin{aligned} R(a^k \rightarrow a^{k+1}) - R(a^{k+1} \rightarrow a^k) &= U_i(a^k) - U_i(a^{k+1}), \\ &= \phi(a^k) - \phi(a^{k+1}). \end{aligned}$$

Using the above equality, the difference in the resistances across the two paths is

$$\begin{aligned} R(\mathcal{P}) - R(\mathcal{P}^R) &= \sum_{k=0}^{m-1} R(a^k \rightarrow a^{k+1}) - R(a^{k+1} \rightarrow a^k), \\ &= \sum_{k=0}^{m-1} (\phi(a^k) - \phi(a^{k+1})), \\ &= \phi(a^0) - \phi(a^m). \end{aligned}$$

□

Before stating the following theorem, we will introduce the notation of a resistance tree in the context of log-linear learning. A tree, T , rooted at an action profile a , is a set of $|\mathcal{A}| - 1$ directed edges such that, from every action profile a' , there is a unique directed path in the tree to a . The resistance of a rooted tree, T , is the sum of the resistances on the edges

$$R(T) = \sum_{a' \rightarrow a'' \in T} R(a' \rightarrow a'').$$

Let $\mathcal{T}(a)$ be defined as the set of trees rooted at the action profile a . The stochastic potential of the action profile a is defined as

$$\gamma(a) = \min_{T \in \mathcal{T}(a)} R(T).$$

We will refer to a minimum resistance tree as any tree that has minimum stochastic potential, that is, any tree T that satisfies

$$R(T) = \min_{a \in \mathcal{A}} \gamma(a).$$

Proposition 3.1 *Consider any finite n -player exact potential game with potential function $\phi : \mathcal{A} \rightarrow \mathbb{R}$ where all players adhere to log-linear learning. The stochastically stable states are contained in the set of potential maximizers, i.e., $\{a \in \mathcal{A} : \phi(a) = \max_{a^* \in \mathcal{A}} \phi(a^*)\}$.*

Proof: As mentioned earlier, the proposition follows from the known form of the stationary distribution. We now present a proof based on minimum resistance trees in preparation for the forthcoming analysis on variations of log-linear learning. By Lemma 3.1, we know that log-linear learning induces a regular perturbed Markov process. Therefore, an action profile $a \in \mathcal{A}$ is stochastically stable if and only if there exists a minimum resistance tree rooted at a .

Suppose that a minimum resistance tree, T , is rooted at an action profile a that does not maximize the potential function. Let a^* be any action profile that maximizes the potential function. Since T is a rooted tree, there exists a path \mathcal{P} from a^* to a of the form

$$\mathcal{P} = \{a^* \rightarrow a^1 \rightarrow \dots \rightarrow a^m \rightarrow a\}.$$

Notice that \mathcal{P} is a feasible action path. Consider the reverse path \mathcal{P}^R that goes from a to a^* ,

$$\mathcal{P}^R = \{a \rightarrow a^m \rightarrow \dots \rightarrow a^1 \rightarrow a^*\}.$$

Construct a new tree T' rooted at a^* by adding the edges of \mathcal{P}^R to T and removing the redundant edges \mathcal{P} . The new tree will have the following resistance

$$R(T') = R(T) + R(\mathcal{P}^R) - R(\mathcal{P}).$$

By Lemma 3.2, we know that

$$\begin{aligned} R(T') &= R(T) + \phi(a) - \phi(a^*), \\ &< R(T). \end{aligned}$$

We constructed a new tree T' rooted at a^* with strictly less resistance than T . Therefore T cannot be a minimum resistance tree.

□

3.2 Additional Comments for Log-Linear Learning

The insight that log-linear learning is a regular perturbed Markov decision process enables us to analyze the stochastically stable states for any game structure (i.e., not only potential games) using Theorem 4 in [20].

Corollary 3.1 *Consider any finite n -player game where all players adhere to log-linear learning. An action profile $a \in \mathcal{A}$ is stochastically stable if and only if a minimum resistance tree is rooted at a , or equivalently if a has minimum stochastic potential.*

In any regular perturbed Markov process, the stochastically stable states must be contained in the recurrent classes of the unperturbed process. In the case of log-linear learning, the unperturbed process is an asynchronous best reply process. Therefore, regardless of the game structure, if an action profile is stochastically stable, it must be contained in a best reply cycle. A best reply cycle is a best reply path that cycles back on itself, i.e., a best reply path of the form $a^0 \rightarrow \dots \rightarrow a^m \rightarrow a^0$. Notice that a pure Nash equilibrium is a best reply cycle consisting of only one action.

In certain classes of game, analyzing the recurrent classes of the unperturbed process can suitably restrict the possible stochastically stable action profiles. For example, in games that are weakly acyclic under best replies, strict Nash equilibria are the only recurrent classes of the unperturbed process. This leads to the following corollary.

Corollary 3.2 *Consider any finite n -player weakly acyclic game under best replies where all players adhere to log-linear learning. If all Nash equilibria are strict, then the set of stochastically stable states must be*

contained in the set of Nash equilibrium.²

The above analysis also holds true for generalized ordinal potential games that are “close” to being exact potential games as stated in the following corollary.

Corollary 3.3 *Consider any finite n -player generalized ordinal potential game with potential function $\phi : \mathcal{A} \rightarrow \mathbb{R}$ where all players adhere to log-linear learning. Suppose for any player $i \in \mathcal{I}$, actions $a'_i, a''_i \in \mathcal{A}_i$, and joint action $a_{-i} \in \mathcal{A}_{-i}$, players' utility satisfies*

$$|(U_i(a'_i, a_{-i}) - U_i(a''_i, a_{-i})) - (\phi(a'_i, a_{-i}) - \phi(a''_i, a_{-i}))| \leq \delta,$$

for some $\delta > 0$. If δ is sufficiently small, then the stochastically stable states are contained in the set of potential maximizers.

We omit the proof of Corollary 3.3 as it is identical to the resistance tree proof for log-linear learning in the case of an exact potential function. The only difference is the implication of Lemma 3.2 where equality is replaced by a bound on the difference of the resistances across two paths.

4 Revisiting Asynchrony

In this section we explore whether asynchrony is necessary to guarantee the optimality of the stochastically stable states for log-linear learning in potential games. Do the properties of the stochastically stable states change if a limited number of players are allowed to update each period? Or are we bound by asynchrony?

Suppose at each time $t > 0$, a group of players $G \subseteq \mathcal{I}$ is randomly chosen according to a fixed probability distribution $q \in \Delta(2^{\mathcal{I}})$ where $2^{\mathcal{I}}$ denotes the set of subsets of \mathcal{I} and q^G denotes the probability that group G will be chosen. We will refer to q as the *revision process*. At time t , each player $i \in G$ plays a strategy $p_i(t) \in \Delta(\mathcal{A}_i)$ where

$$p_i^{a_i}(t) = \frac{e^{\frac{1}{\tau} U_i(a_i, a_{-i}(t-1))}}{\sum_{\bar{a}_i \in \mathcal{A}_i} e^{\frac{1}{\tau} U_i(\bar{a}_i, a_{-i}(t-1))}},$$

²The strictness condition could be relaxed as follows. If a^* is a Nash equilibrium and there exists a best reply cycle of the form $a^* \rightarrow a^1 \rightarrow \dots \rightarrow a^m \rightarrow a^*$, then the actions $\{a^1, \dots, a^m\}$ are also candidates for the stochastically stable states.

for any action $a_i \in \mathcal{A}_i$. All players not in G must repeat their previous action, i.e., $a_j(t) = a_j(t - 1)$ for all $j \notin G$. We will refer to this learning algorithm as *synchronous log-linear learning* with revision process q . This setting is the focus of [1].

4.1 A Counterexample

Synchronous log-linear learning not only alters the resulting stationary distribution but also affects the set of stochastically stable states. The following example illustrates this phenomenon.

Consider a three player identical interest game with the following payoffs:

	a_2	b_2	
a_1	0	0	
b_1	0	1	
	a_3		

	a_2	b_2
a_1	0	1
b_1	1	1
	b_3	

In this example, there are several Nash equilibria, namely all joint action profiles that yield a utility of 1 in addition to the action profile (a_1, a_2, a_3) that yields a utility of 0. Suppose the revision process q has full support on the player sets $\{1, 2, 3, (1, 2, 3)\}$. An example of such a revision process is

$$\begin{aligned}
 q^1 &= 1/4, \\
 q^2 &= 1/4, \\
 q^3 &= 1/4, \\
 q^{(1,2,3)} &= 1/4.
 \end{aligned}$$

We will see that *all* action profiles are stochastically stable. Before proceeding with the discussion, we first introduce the following notation.

For any action profiles $a, a' \in \mathcal{A}$, the set of groups where the transition $a^0 \rightarrow a^1$ is feasible given the revision process q is defined as

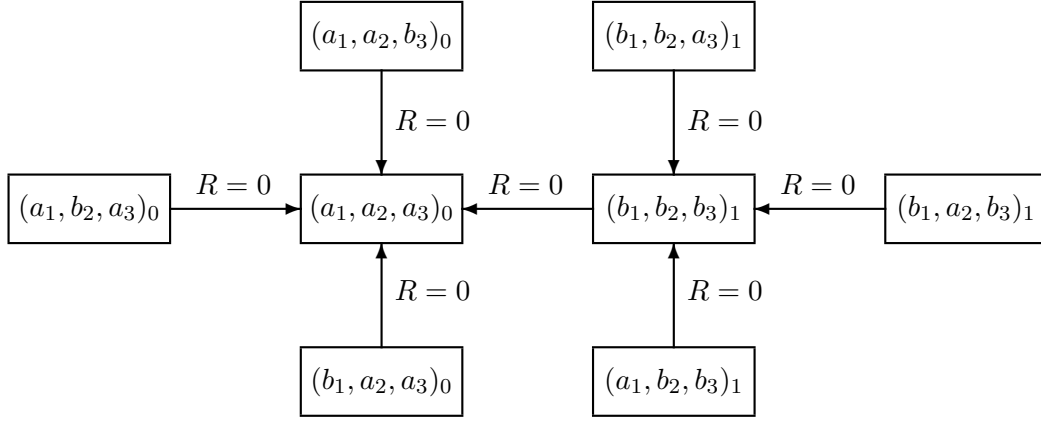
$$\bar{G}(a, a', q) := \{G \subseteq \mathcal{I} : \{i \in \mathcal{I} : a_i \neq a'_i\} \subseteq G, q^G > 0\}.$$

Synchronous log-linear learning with revision process q induces a regular perturbed Markov process, with the unperturbed process being synchronous learning under best replies. In synchronous log-linear learning,

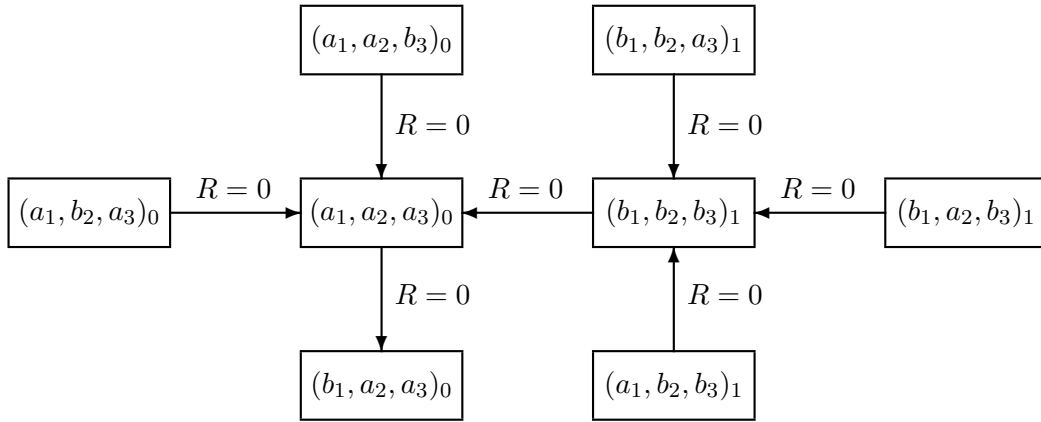
a transition $a^0 \rightarrow a^1$ is feasible if and only if $\bar{G}(a, a', q) \neq \emptyset$. It is straightforward to show that the resistance of a feasible transition $a^0 \rightarrow a^1$ is

$$R(a^0 \rightarrow a^1) = \min_{G \in \bar{G}(a^0, a^1, q)} \sum_{i \in G} \left(\max_{a_i^* \in \mathcal{A}_i} U_i(a_i^*, a_{-i}^0) - U_i(a_i^1, a_{-i}^0) \right). \quad (7)$$

Since synchronous log-linear learning is a regular perturbed Markov process we know that a minimum resistance tree must be rooted at any stochastically stable state. It is straightforward to show that any action profile that yields a utility of 1 has a rooted tree of resistance 0; hence all such action profiles are stochastically stable. Furthermore, there also exists a 0-resistance tree rooted at (a_1, a_2, a_3) of the form



Therefore (a_1, a_2, a_3) is also stochastically stable. Building off the same tree, it is straightforward to construct trees rooted at (b_1, a_2, a_3) , (a_1, b_2, a_3) , and (a_1, a_2, b_3) that also have 0 resistance, e.g.,



Therefore, all action profiles are stochastically stable.

4.2 Regular Revision Processes

Characterizing the stochastically stable states of a synchronous log-linear learning process is highly dependent on the structure of the revision process q in addition to the structure of the game. We start by defining a *regular revision process* (introduced in [1]).

Definition 4.1 (Regular Revision Process) *A regular revision process is a probability distribution $q \in \Delta(2^{\mathcal{I}})$ where for each player $i \in \mathcal{I}$, the probability assigned to the group consisting of only player i is positive, i.e., $q^i > 0$.*

The previous example demonstrates that any action profile may be stochastically stable even if we restrict our attention to regular revision processes and identical interest games.³ However, one can easily characterize the stochastically stable states if the game embodies a special structure. This leads to the following corollary, which extends Theorem 2 in [1] to a larger class of games.

Corollary 4.1 *Consider any finite n -player weakly acyclic game under best replies where all Nash equilibria are strict. If all players adhere to synchronous log-linear learning with regular revision process q , then the stochastically stable states are contained in the set of Nash equilibria.⁴*

Proof: Synchronous log-linear learning with regular revision process q is a regular perturbed Markov process; therefore, the stochastically stable states must be contained in the recurrent class of the unperturbed process. The unperturbed process is as follows: at each time $t > 0$, a group of players $G \subseteq \mathcal{I}$ is randomly chosen according to q . At time t , each player $i \in G$ selects an action from his best reply set, i.e., $a_i(t) \in B_i(a_{-i}(t-1))$. All players not in G must repeat their previous action, i.e., $a_j(t) = a_j(t-1)$ for all $j \notin G$.

For any game that is weakly acyclic game under best replies, if all Nash equilibria are strict, then the recurrent classes of the unperturbed process are precisely the set of Nash equilibria.

□

³See [1] for alternative examples of games and revision processes outside the realm of regular revision processes that exhibit similar behavior.

⁴Any best response potential game is weakly acyclic game under best replies [19]. However, the converse is not true. The strictness condition could be relaxed as stated in the footnote of Corollary 3.2.

4.3 Independent Revision Process

In general, guaranteeing optimality of the stochastically stable states in potential games for an arbitrary revision process and utility interdependence structure is unattainable. In this section, we will focus on particular class of revision processes that we refer to as *independent revision processes*. In such setting, each player independently decides whether to update his strategy by the log-linear learning rule with some probability $\omega > 0$. More precisely, at each time t , each player $i \in \mathcal{I}$ simultaneously plays the following strategy: (i) with probability $(1 - \omega)$, player i selects his previous action, i.e., $a_i(t) = a_i(t - 1)$, or (ii) with probability ω player i plays a strategy $p_i(t) \in \Delta(\mathcal{A}_i)$ where for any action $a_i \in \mathcal{A}_i$

$$p_i^{a_i}(t) = \frac{e^{\frac{1}{\tau}U_i(a_i, a_{-i}(t-1))}}{\sum_{\bar{a}_i \in \mathcal{A}_i} e^{\frac{1}{\tau}U_i(\bar{a}_i, a_{-i}(t-1))}}.$$

Theorem 4.1 *Consider any finite n -player exact potential game with potential function $\phi : \mathcal{A} \rightarrow \mathbb{R}$ where all players adhere to synchronous log-linear learning with an independent revision process with update parameter $\omega > 0$. Set $\omega = (e^{-\frac{1}{\tau}})^m$. For sufficiently large m , the stochastically stable states are contained in the set of potential maximizers.*

Note that Theorem 4.1 considers stochastic stability in terms of the *combined* perturbation of log-linear (vs. maximizing) action selection through τ and the group selection through ω . The perturbations are related by $\omega = (e^{-1/\tau})^m$, with $\tau \downarrow 0^+$ as usual.

Proof: It is straightforward to verify that synchronous log-linear learning with an independent revision process is a regular perturbed Markov process where the resistance of any transition ($a \rightarrow a'$) is

$$R(a \rightarrow a') = m|G| + \sum_{i \in G} \left(\max_{a_i^* \in \mathcal{A}_i} U_i(a_i^*, a_{-i}) - U_i(a'_i, a_{-i}) \right),$$

where $G = \{i \in \mathcal{I} : a_i \neq a'_i\}$. The unperturbed process corresponds to players never experimenting, i.e., $\omega = 0$.

Without loss of generality, we will assume that $\phi(a)$ is bounded between 0 and 1 for all $a \in \mathcal{A}$. Utilizing this normalization, the resistance of the transition ($a \rightarrow a'$) satisfies the following inequality

$$m|G| + n \geq R(a \rightarrow a') \geq m|G|. \quad (8)$$

Let T be a minimum resistance tree rooted at a^* . If for each edge $[a \rightarrow \tilde{a}] \in T$ there is a single deviator, i.e., $|\{i \in \mathcal{I} : a_i \neq \tilde{a}_i\}| = 1$, then arguments from the proof of Proposition 3.1 establish that a^* must be a potential function maximizer.

Suppose there exists an edge $[a \rightarrow \tilde{a}] \in T$ with multiple deviators, i.e., $|\{i \in \mathcal{I} : a_i \neq \tilde{a}_i\}| \geq 2$. Let $G = \{i \in \mathcal{I} : a_i \neq \tilde{a}_i\}$. The resistance of this transition is at least

$$R(a \rightarrow \tilde{a}) \geq m|G|.$$

Consider any path $\mathcal{P} = \{a = a^0 \rightarrow a^1 \rightarrow \dots \rightarrow a^{|G|} = \tilde{a}\}$ where each transition $a^k \rightarrow a^{k+1}$, $k \in \{0, \dots, |G| - 1\}$ reflects a unilateral deviation by some player $i \in G$. According to (8), the resistance of each edge along this path is at most

$$R(a^{k-1} \rightarrow a^k) \leq m + n.$$

Therefore the resistance of the path \mathcal{P} is at most

$$R(\mathcal{P}) \leq |G|(m + n) \leq m|G| + n^2.$$

Construct a new tree T' rooted at a^* by adding the edges of \mathcal{P} to T and removing the redundant edges \mathcal{P}^R . The redundant edges are the set of edges leaving the action profiles $\{a^0, \dots, a^{|G|-1}\}$ in the original tree T . The redundant edges \mathcal{P}^R include the edge $[a \rightarrow \tilde{a}]$ in addition to $(|G| - 1)$ other edges each of which has resistance at least m . The total resistance of the redundant edges is at least

$$R(\mathcal{P}^R) \geq m|G| + m(|G| - 1).$$

The new tree T' has resistance

$$\begin{aligned} R(T') &= R(T) + R(\mathcal{P}) - R(\mathcal{P}^R), \\ &\leq R(T) + m|G| + n^2 - m|G| - m(|G| - 1), \\ &= R(T) + n^2 - m(|G| - 1). \end{aligned}$$

Since $|G| \geq 2$, if $m \geq n^2$ then $R(T') < R(T)$. This implies that any edge in a minimum resistance tree consists of only a single deviator which in turn implies that only potential function maximizers are stochastically stable.

□

4.4 Games Played on Graphs

The independent revision process discussed in the previous section guarantees that the only stochastically stable states are potential function maximizers irrespective of the structure of the utility interdependence. In this section, we will focus on a particular structure of utility interdependence by considering games played on graphs where there are potentially a large number of players and each player's utility is only effected by the actions of a subset of other players. In such settings, we demonstrate that one can exploit this structure to expand the class of revision processes that guarantee optimality of the stochastically stable states.

In games played on graphs, each player's utility is influenced by a subset of other players. Let $N_i \subseteq \mathcal{I}$ represent the set of players that impact player i 's utility. We will refer to N_i as the neighbors of player i . In such a setting, player i 's utility is of the form $U_i : \mathcal{A}^{N_i} \rightarrow \mathbb{R}$ where $\mathcal{A}^{N_i} := \prod_{j \in N_i} \mathcal{A}_j$. To avoid unnecessary notation, we will still express the utility of player i given the action profile a as $U_i(a)$. If player i 's utility depends on the complete action profile, then $N_i = \mathcal{I}$.

We make the following assumption on the revision process q .

Assumption 4.1 *For any player $i \in \mathcal{I}$ there exists a group $G \subseteq \mathcal{I}$ such that $i \in G$ and $q^G > 0$.*

Proposition 4.1 *Consider any finite n -player exact potential game with potential function $\phi : \mathcal{A} \rightarrow \mathbb{R}$. Let all players adhere to synchronous log-linear learning with a revision process satisfying Assumptions 4.1. Assume further that the revision process is such that any group of players G with $q^G > 0$ is conflict free, i.e.,*

$$i, j \in G, i \neq j \Rightarrow i \notin N_j.$$

Then the stochastically stable states are contained in the set of potential maximizers.

Proof: For a given revision process q , a transition $a^0 \rightarrow a^1$ is possible if $\bar{G}(a^0, a^1, q) \neq \emptyset$. According to (7), it is straightforward to exploit the conflict free assumption to verify that

$$R(a^0 \rightarrow a^1) - R(a^1 \rightarrow a^0) = \phi(a^0) - \phi(a^1).$$

The remainder of the proof is identical to the proof for log-linear learning in Section 3.1. We note that this result also could have been proved by analyzing the precise stationary distribution using the detailed balance

condition as set forth in [21].

□

This result is unsatisfying in that it requires a special structure on the revision process. Consider the following relaxation of the revision process. Roughly speaking, suppose that the revision process q usually selects conflict free groups but occasionally selects conflicted groups. To formalize this intuition, we make the following assumption on the revision process.

Assumption 4.2 *Let the revision process $q(\alpha)$ be continuously parameterized by $\alpha \geq 0$, where $q(\alpha)$ satisfies Assumption 4.1 for all α .*

- For any G , $q^G(\alpha) > 0$ for some $\alpha > 0$ implies that there exists a $\rho \geq 0$ such that

$$0 < \lim_{\alpha \rightarrow 0} \frac{q^G(\alpha)}{\alpha^\rho} < \infty;$$

- For any conflicted (i.e., not conflict free) group C such that $q^C(\alpha) > 0$ and any player $i \in C$, there exists a conflict free group $F \subseteq \mathcal{I}$ such that

- $i \in F$;
- $q^F(0) > 0$;
- $q^F(\alpha) \geq \left(\frac{\kappa}{\alpha}\right) q^C(\alpha)$ for some $\kappa > 0$ independent of F and C .

Theorem 4.2 *Consider any finite n -player exact potential game with potential function $\phi : \mathcal{A} \rightarrow \mathbb{R}$ where all players adhere to synchronous log-linear learning with a revision process $q(\alpha)$ that satisfies Assumption 4.1–4.2. Set $\alpha = (e^{-\frac{1}{\tau}})^m$. For sufficiently large m , the stochastically stable states are contained in the set of potential function maximizers.*

As with Theorem 4.1, Theorem 4.2 considers stochastic stability in terms of the *combined* perturbation of log-linear (vs. maximizing) action selection through τ and conflicted (vs. conflict free) group selection through α . The perturbations are related by $\alpha = (e^{-\frac{1}{\tau}})^m$, with $\tau \downarrow 0^+$ as usual.

Proof: Without loss of generality, we will assume that $\phi(a)$ is bounded between 0 and 1 for all $a \in \mathcal{A}$. Synchronous log-linear learning with a revision process $q(\alpha)$ induces a regular perturbed Markov decision process. The unperturbed process consists of conflict free group selection with unilateral best replies.

Let T be a minimum resistance tree rooted at a^* . If for each edge $[a \rightarrow \tilde{a}] \in T$ there exists a conflict free group G such that $q^G > 0$ and $\{i \in \mathcal{I} : a_i \neq \tilde{a}_i\} \subseteq G$ then arguments from the proof of Proposition 4.1 establish that a^* must be a potential function maximizer.

Suppose there exists an edge $[a \rightarrow \tilde{a}]$ where there does not exist a conflict free group G such that $q^G > 0$ and $\{i \in \mathcal{I} : a_i \neq \tilde{a}_i\} \subseteq G$. From Assumption 4.2, the probability of this transition is at most order $\epsilon^{m\rho}$ for some $\rho \geq 1$. Accordingly, the associated resistance satisfies $R(a \rightarrow \tilde{a}) \geq m\rho$.

By Assumption 4.2, there exists a path $\mathcal{P} = \{a = a^0 \rightarrow a^1 \rightarrow \dots \rightarrow a^l = \tilde{a}\}$ of at most length $|\{i \in \mathcal{I} : a_i \neq \tilde{a}_i\}|$, where each transition $a^{k-1} \rightarrow a^k$, $k \in \{1, \dots, l\}$ reflects a unilateral deviation that can be accomplished by a conflict free group. Assumption 4.2 further implies that the resistance of each transition $a^{k-1} \rightarrow a^k$ along the path \mathcal{P} is at most

$$R(a^{k-1} \rightarrow a^k) \leq n.$$

The main idea behind the above inequality is as follows. Resistance computation is based on the product of probability of a group being selected and the probability of an action being selected. There exists a conflict free group with nonvanishing probability that can accomplish the desired transition. The selection of this group contributes zero to the resistance value. The action selection contributes to the resistance by at most n (cf., equation (7) and the assumed normalized bounds of the potential function.) The resistance along the path \mathcal{P} is bounded by

$$R(\mathcal{P}) \leq |\{i \in \mathcal{I} : a_i \neq \tilde{a}_i\}|n \leq n^2.$$

We can conclude that for $m\rho > n^2$,

$$R(\mathcal{P}) < R(a \rightarrow \tilde{a}).$$

Construct a new tree T' still rooted at a^* by adding the edges of \mathcal{P} to T and removing the redundant edges. The new tree T' has strictly less resistance than T , contradicting the assumption that T was a minimum resistance tree. Hence, all minimum resistance trees must consist of only conflict free transitions. This implies that only potential function maximizers are stochastically stable.

□

5 Revisiting Completeness

One role for learning algorithms in distributed control is to guide the decisions of players in real time. That is, the iterations of a learning algorithm correspond to the sequential decisions of players. In such a setting, players may not have the ability to select any particular action in their action set at any given time. One example is multi-vehicle motion control, where an agent’s action set represents discrete spatial locations. Mobility limitations restrict the ability to traverse from one location to another in a given time period.

Standard log-linear learning assumes an agent can access any action at each iteration, which we refer to as “completeness”. These considerations motivate the introduction of constraints between iterations of a learning algorithm.

Let $a(t-1)$ be the joint action at time $t-1$. With constrained action sets, the set of actions available to player i at time t is a function of his action at time $t-1$ and will be denoted as $C_i(a_i(t-1)) \subseteq \mathcal{A}_i$. We will adopt the convention that $a_i \in C_i(a_i)$ for any action $a_i \in \mathcal{A}_i$, i.e., a player is always allowed to stay with his previous action. We will say that a player’s action set is *complete* if $C_i(a_i) = \mathcal{A}_i$ for all actions $a_i \in \mathcal{A}_i$.⁵

We make the following two assumptions on constrained action sets.

Assumption 5.1 For any player $i \in \mathcal{I}$ and any action pair $a_i^0, a_i^m \in \mathcal{A}_i$, there exists a sequence of actions $a_i^0 \rightarrow a_i^1 \rightarrow \dots \rightarrow a_i^m$ satisfying $a_i^k \in C_i(a_i^{k-1})$ for all $k \in \{1, 2, \dots, m\}$.

Assumption 5.2 For any player $i \in \mathcal{I}$ and any action pair $a_i^1, a_i^2 \in \mathcal{A}_i$,

$$a_i^2 \in C_i(a_i^1) \Leftrightarrow a_i^1 \in C_i(a_i^2).$$

5.1 A Counterexample

Log-linear learning, in its original form, requires completeness of players’ action sets. Suppose log-linear learning is employed in a situation with constrained action sets. At time t , player i plays a strategy $p_i(t) \in \Delta(\mathcal{A}_i)$ where

$$p_i^{a_i}(t) = \begin{cases} \frac{e^{\frac{1}{\tau} U_i(a_i, a_{-i}(t-1))}}{\sum_{\bar{a}_i \in C_i(a_i)} e^{\frac{1}{\tau} U_i(\bar{a}_i, a_{-i}(t-1))}} & \text{for any action } a_i \in C_i(a_i), \\ 0 & \text{for any action } a_i \notin C_i(a_i). \end{cases} \quad (9)$$

⁵We note that this scenario could have been formulated as a stochastic game [16] where the state is defined as the previous action profile and the state dependent action sets are defined according to the constrained action sets. We will avoid formally defining the game as a stochastic game in favor of a more direct treatment.

It is easy to see that log-linear learning in this setting induces a regular perturbed Markov process. The resistance of a feasible transition $a^0 \rightarrow a^1$, i.e., a^1 is of the form (a_i^1, a_{-i}^0) for some player i and action $a_i^1 \in C_i(a_i^0)$, is now

$$R(a^0 \rightarrow a^1) = \max_{a_i^* \in C_i(a_i^0)} U_i(a_i^*, a_{-i}^0) - U_i(a^1).$$

The following example demonstrates that log-linear learning applied to constrained action sets need not result in stochastically stable joint actions that are potential function maximizer. Consider a two player identical interest game with payoffs

		Player 2		
		b_1	b_2	b_3
Player 1	a_1	0	0	9
	a_2	10	-10	-10

In this example, there are two Nash equilibria (a_2, b_1) and (a_1, b_3) , but only one potential maximizer (a_2, b_1) . Let the constrained action sets satisfy

$$C_1(a_1) = \{a_1, a_2\},$$


$$C_1(a_2) = \{a_1, a_2\},$$

$$C_2(b_1) = \{b_1, b_2\},$$


$$C_2(b_2) = \{b_1, b_2, b_3\},$$

$$C_2(b_3) = \{b_2, b_3\}.$$

A minimum resistance tree must be rooted at a stochastically stable state. A minimum resistance tree rooted at (a_2, b_1) is



A minimum resistance tree rooted at (a_1, b_3) is



Hence, (a_2, b_1) is not stochastically stable. In fact, one can construct an alternative example which illustrate that the stochastically stable states need not be contained in the set of Nash equilibria for potential games when action sets are incomplete.

The key insight for this example lies in the best reply graph. In log-linear learning with complete action sets, the recurrent classes of the unperturbed process are the Nash equilibria. Constraining action sets changes the structure of the best reply graph. In particular, these changes may induce new *local* Nash equilibria, i.e., action profiles where no player can unilaterally improve his utility by selecting an action in his constrained action set. These local Nash equilibria are now recurrent classes of the unperturbed process and are candidates for the stochastically stable states of the perturbed process.

Furthermore, this example illustrates that Lemma 3.2 no longer holds in the constrained setting because the difference in the resistances across two path is no longer a function of the endpoints.

5.2 Binary Log-Linear Learning

In this subsection, we will introduce a variant of log-linear learning that will rectify the problems caused by constrained action sets.

Consider the following learning algorithm, which we will refer to as *binary log-linear learning* originally proposed in [2, 8]. At each time $t > 0$, one player $i \in \mathcal{I}$ is randomly chosen (uniformly) and allowed to alter his current action. All other players must repeat their current action at the ensuing time step, i.e. $a_{-i}(t) = a_{-i}(t - 1)$. At time t , player i selects one trial action \hat{a}_i (uniformly) from his constrained action

set $C_i(a_i(t-1)) \subset \mathcal{A}_i$. Player i plays a strategy $p_i(t) \in \Delta(\mathcal{A}_i)$ where

$$\begin{aligned} p_i^{a_i(t-1)}(t) &= \frac{e^{\frac{1}{\tau}U_i(a(t-1))}}{e^{\frac{1}{\tau}U_i(a(t-1))} + e^{\frac{1}{\tau}U_i(\hat{a}_i, a_{-i}(t-1))}}, \\ p_i^{\hat{a}_i}(t) &= \frac{e^{\frac{1}{\tau}U_i(\hat{a}_i, a_{-i}(t-1))}}{e^{\frac{1}{\tau}U_i(a(t-1))} + e^{\frac{1}{\tau}U_i(\hat{a}_i, a_{-i}(t-1))}}, \\ p_i^{a_i}(t) &= 0, \quad \forall a_i \neq a_i(t-1), \hat{a}_i \end{aligned}$$

If $a_i(t-1) = \hat{a}_i$, then $p_i^{a_i(t-1)} = 1$.⁶

Reference [8] analyzes the stationary distribution of binary log-linear learning when action sets are constrained. In particular, if the probability of selecting a trial action satisfies a certain specified condition then the stationary distribution of (3) is preserved. This condition involves accommodating for the time-varying cardinality $|C_i(a_i(t-i))|$. We now show that this issue can be resolved by changing the focus of the analysis from the stationary distribution to the stochastically stable states.

Theorem 5.1 *Consider any finite n -player potential game with potential function $\phi : \mathcal{A} \rightarrow \mathbb{R}$ and constrained action sets satisfying Assumptions 5.1 and 5.2. If all players adhere to binary log-linear learning, then the stochastically stable states are contained in the set of potential maximizers.*

Before proving that the only stochastically stable states are the potential maximizers, we will introduce two lemmas. The first lemma establishes that binary log-linear learning is a regular perturbed Markov decision process.

Lemma 5.1 *Binary log-linear learning induces a regular perturbed Markov process where the resistance of any feasible transition, $a^0 \rightarrow a^1 = (a_i^1, a_{-i}^0)$ where $a_i^1 \in C_i(a_i^0)$, is*

$$R(a^0 \rightarrow a^1) = \max_{a_i^* \in \{a_i^0, a_i^1\}} U_i(a_i^*, a_{-i}^0) - U_i(a^1).$$

Proof: The unperturbed process employs a maximizing strategy as opposed to log-linear strategy. Binary log-linear learning with constrained action sets induces a finite, aperiodic, irreducible process over the state space \mathcal{A} . Irreducibility is a consequence of Assumption 5.1. Let $\epsilon := e^{-\frac{1}{\tau}}$. Let P^ϵ denote the associated

⁶The results presented in this section still hold even if the probability distributions are not uniform. In the case of the player selection process, any probability distribution with full support on the player set $\{1, 2, \dots, n\}$ would work. The probability of selecting a trial action can be relaxed in the same fashion.

probability transition matrix. The probability of transitioning from a^0 to $a^1 := (a_i^1, a_{-i}^0)$ where $a_i^1 \in C_i(a_i^0)$ is now of the form

$$P_{a^0 \rightarrow a^1}^\epsilon = \left(\frac{1}{n|C_i(a_i^0)|} \right) \left(\frac{\epsilon^{-U_i(a^1)}}{\epsilon^{-U_i(a^0)} + \epsilon^{-U_i(a^1)}} \right). \quad (10)$$

Redefine the functions:

$$\begin{aligned} V_i(a^0, a^1) &:= \max\{U_i(a^0), U_i(a^1)\}, \\ B_i(a^0, a^1) &:= \{a \in \{a^0, a^1\} : U_i(a) = V_i(a^0, a^1)\}. \end{aligned}$$

Accordingly,

$$\lim_{\epsilon \rightarrow 0^+} \frac{P_{a^0 \rightarrow a^1}^\epsilon}{\epsilon^{V_i(a^0, a^1) - U_i(a^1)}} = \frac{1}{n|C_i(a_i^0)||B_i(a^0, a^1)|}.$$

This implies that the process P^ϵ is a regular perturbed Markov process where the resistance of the transition $a^0 \rightarrow a^1$ is

$$R(a^0 \rightarrow a^1) = V_i(a^0, a^1) - U_i(a^1) \geq 0.$$

□

We state the following lemma without proof as it is almost identical to the proof of Lemma 3.2.

Lemma 5.2 *Consider any finite n -player exact potential game with potential function $\phi : \mathcal{A} \rightarrow \mathbb{R}$, where all players adhere to binary log-linear learning. For any feasible action path*

$$\mathcal{P} = \{a^0 \rightarrow a^1 \rightarrow \dots \rightarrow a^m\}$$

and reverse path

$$\mathcal{P}^R = \{a^m \rightarrow a^{m-1} \rightarrow \dots \rightarrow a^0\},$$

the difference in the total resistance across the paths is

$$R(\mathcal{P}) - R(\mathcal{P}^R) = \phi(a^0) - \phi(a^m).$$

Proof of Theorem 5.1: By Lemma 5.1, we know that binary log-linear learning induces a regular perturbed Markov process. Therefore, an action profile $a \in \mathcal{A}$ is stochastically stable if and only if there exists a minimum resistance tree rooted at a .

Suppose that a minimum resistance tree, T , was rooted at an action profile a that does not maximize the potential function. Let a^* be any action profile that maximizes the potential function. Since T is a rooted tree, there exists a path \mathcal{P} from a^* to a of the form

$$\mathcal{P} = \{a^* \rightarrow a^1 \rightarrow \dots \rightarrow a^m \rightarrow a\}.$$

Notice that \mathcal{P} is a feasible action path. Consider the reverse path \mathcal{P}^R that goes from a to a^* ,

$$\mathcal{P}^R = \{a \rightarrow a^m \rightarrow \dots \rightarrow a^1 \rightarrow a^*\}.$$

Such a path exists because of Assumption 5.2. Construct a new tree T' rooted at a^* by adding the edges of \mathcal{P}^R to T and removing the redundant edges \mathcal{P} . The new tree will have the following resistance


$$R(T') = R(T) + R(\mathcal{P}^R) - R(\mathcal{P})$$

By Lemma 3.2, we know that

$$\begin{aligned} R(T') &= R(T) + \phi(a) - \phi(a^*), \\ &< R(T). \end{aligned}$$

Therefore, we constructed a new tree T' rooted at a^* with strictly less resistance than T . Accordingly, T is cannot be a minimum resistance tree. \square

When revisiting the example in Section 5.1, the minimum resistance trees rooted at (a_2, b_1) for binary log-linear learning is now



The potential maximizer (a_2, b_1) is the only stochastically stable state.

6 Payoff Based Implementation

In any of version of log-linear learning, each player needs to be endowed with the ability to assess the utility that the player would have received had the player selected an alternative action. That is, given any action profile a , each player i is able to compute $U_i(a'_i, a_{-i})$ for all $a'_i \in \mathcal{A}_i$. In this section, we seek to understand what happens if players do not have access to this information. Rather, players have access to *only* (i) the action they played and (ii) the utility they received.

Payoff-based learning algorithms have received significant attention recently [11, 23]. In particular, [11, 23] introduce different forms of payoff-based learning algorithms and show that only Nash equilibria are stochastically stable for certain classes of games.

In this section, we introduce a payoff-based learning algorithm such that only potential function maximizers are stochastically stable in exact potential games. In fact, we will combine two of the algorithms perviously studied, synchronous log-linear learning with an independent revision process in Section 4.3 and binary log-linear learning in Section 5.2, to develop a payoff based version of log-linear learning with the desired properties.

We introduce the following learning algorithm called *payoff based log-linear learning*. For any time $t \in \{1, 2, \dots\}$, let $a(t-1)$ and $a(t)$ be the action profile at time $t-1$ and t , respectively. Define $x_i(t) \in \{0, 1\}$ to be a binary flag that indicates whether player i experimented in time t . At time $t + 1$ each player i simultaneously selects an action $a_i \in \mathcal{A}_i$ according to the following rule:

- If player i did not experiment in period t , i.e., $x_i(t) = 0$, then
 - With probability ω ,
 - * $a_i(t + 1)$ is chosen randomly (uniformly) over \mathcal{A}_i with probability ω ,
 - * $x_i(t + 1) = 1$,
 - With probability $(1 - \omega)$,
 - * $a_i(t + 1) = a_i(t)$,

$$* x_i(t+1) = 0.$$

where $\omega > 0$ is the exploration rate.

- If player i experimented in period t , i.e., $x_i(t) = 1$, then

$$a_i(t+1) = \begin{cases} a_i(t-1) & \text{with probability } \frac{e^{\frac{1}{\tau}U_i(a(t-1))}}{e^{\frac{1}{\tau}U_i(a(t-1))} + e^{\frac{1}{\tau}U_i(a(t))}} \\ a_i(t) & \text{with probability } \frac{e^{\frac{1}{\tau}U_i(a(t))}}{e^{\frac{1}{\tau}U_i(a(t-1))} + e^{\frac{1}{\tau}U_i(a(t))}} \end{cases} \quad (11)$$

and

$$x_i(t+1) = 0.$$

These dynamics can be described as follows. Occasionally players experiment with a new action. In the event that a player does experiment, then the player switches to the new action with a log-linear strategy over the utility received in the previous two time periods.

The above dynamics require that a player knows only his own utility received and action employed in the last two time steps and whether he experimented or not in the previous time step. These dynamics do not require players to have any knowledge regarding the actions, strategies, or utilities of the other players.

Theorem 6.1 *Consider any finite n -player exact potential game with potential function $\phi : \mathcal{A} \rightarrow \mathbb{R}$. where all players adhere to payoff-based log-linear learning. Set $\omega = (e^{-\frac{1}{\tau}})^m$. For sufficiently large m , the stochastically stable states are contained in the set of potential maximizers.*

In a similar manner to Theorem 4.2, Theorem 6.1 considers stochastic stability in terms of the *combined* perturbation of log-linear action selection through τ and experimentation through ω . The perturbations are related by $\omega = (e^{-\frac{1}{\tau}})^m$.

The remainder of this subsection is devoted to the proof of Theorem 6.1. Without loss of generality, we will assume that $\phi(a)$ is bounded between 0 and 1 for all $a \in \mathcal{A}$.

First, define the *state* of the dynamics at time t to be the tuple $z(t) := [a(t-1), a(t), x(t)] \in \mathcal{A} \times \mathcal{A} \times \{0, 1\}^n$.

Claim 6.1 *Payoff-based log-linear learning induces a regular perturbed Markov decision process. Furthermore, the resistance of a transition from state $z^1 = [a^0, a^1, x^1]$ to state $z^2 = [a^1, a^2, x^2]$ is*

$$R(z^1 \rightarrow z^2) = \underbrace{ms(x^2)}_* + \underbrace{\sum_{i:x_i^1=1, a_i^2=a_i^0} (V_i(a^0, a^1) - U_i(a^0)) + \sum_{i:x_i^1=1, a_i^2=a_i^1} (V_i(a^0, a^1) - U_i(a^1))}_{**} \quad (12)$$

where

$$V_i(a^0, a^1) := \max\{U_i(a^0), U_i(a^1)\},$$

$$s(x) := \sum_i x_i.$$

Proof: The unperturbed process corresponds to players never experimenting, i.e., $\omega = 0$. The recurrent classes of the unperturbed process are states of the form $[a, a, \mathbf{0}]$, where boldface $\mathbf{0}$ denotes a zero vector of appropriate dimension.

Valid transitions of the perturbed process are of the following form:

$$[a^0, a^1, x^1] \rightarrow [a^1, a^2, x^2],$$

where

$$x_i^1 = 0 \Rightarrow \begin{cases} x_i^2 = 0, a_i^2 = a_i^1, \text{ or} \\ x_i^2 = 1, a_i^2 \in \mathcal{A}_i, \end{cases}$$

$$x_i^1 = 1 \Rightarrow x_i^2 = 0, a_i^2 \in \{a_i^0, a_i^1\}.$$

Let $z^1 := [a^0, a^1, x^1]$, $z^2 := [a^1, a^2, x^2]$, and $\epsilon := e^{-\frac{1}{\tau}}$. The probability of the transition from $z^1 \rightarrow z^2$ is

$$P_{z^1 \rightarrow z^2}^\epsilon = \left(\prod_{i:x_i^1=0, x_i^2=0} (1 - \omega) \right) \left(\prod_{i:x_i^1=0, x_i^2=1} \frac{\omega}{|\mathcal{A}_i|} \right) \left(\prod_{i:x_i^1=1, a_i^2=a_i^0} \frac{\epsilon^{-U_i(a^0)}}{\epsilon^{-U_i(a^0)} + \epsilon^{-U_i(a^1)}} \right) \left(\prod_{i:x_i^1=1, a_i^2=a_i^1} \frac{\epsilon^{-U_i(a^1)}}{\epsilon^{-U_i(a^0)} + \epsilon^{-U_i(a^1)}} \right)$$

It is straightforward to verify that

$$0 < \lim_{\epsilon \rightarrow 0^+} \frac{P_{z^1 \rightarrow z^2}^\epsilon}{\epsilon R(z^1 \rightarrow z^2)} < \infty,$$

with $R(z^1 \rightarrow z^2)$ as stated in the claim.

□

Note that the resistance of a particular transition is expressed as two terms. The first term in equation (12), denoted (*), captures the number of synchronous deviators. The second term, denoted (**), captures the log-linear response of players.

Partition the state space into the following five sets: X_1 , X_2 , D , E , and E^* :

- X_1 is the set of states $[a, a', x]$ such that $x_i = 1$ for at least one player i .
- X_2 is the set of states $[a, a', \mathbf{0}]$ such that $a \neq a'$.
- D is the set of states $[a, a, \mathbf{0}]$ such that a is not a pure Nash equilibrium.
- E is the set of states $[a, a, \mathbf{0}]$ such that a is a pure Nash equilibrium, but is not a maximizer of ϕ .
- E^* is the set of states $[a, a, \mathbf{0}]$ such that a is a pure Nash equilibrium and is a maximizer of ϕ .

These account for all possible states.

The only candidates for the stochastically stable states are the recurrent classes of the unperturbed process. This eliminates the possibility that states in X_1 and X_2 are stochastically stable. Therefore, from now on we will restrict our attention only to trees rooted at states of D , E , and E^* . For simplicity, we will use shorthand notation to write $[a, a, \mathbf{0}]$ as simply a .

Claim 6.2 *All edges $a^0 \rightarrow a^1$ have a resistance $R(a^0 \rightarrow a^1) \geq m$.*

Proof: Any transition requires at least one player to experiment which happens with at most probability ϵ^m .

□

Note that a single edge between states in D , E , or E^* , will involve transitions between multiple states. Accordingly, the expression $R(a^0 \rightarrow a^1)$ refers to the resistance of the *path* between a^0 and a^1 .

Claim 6.3 *All edges $a^0 \rightarrow a^1$ in a minimum resistance tree must consist of a unilateral deviator; i.e.,*

$$|\{i \in \mathcal{I} : a_i^0 \neq a_i^1\}| = 1.$$

Proof: We will construct resistance trees over vertices in $D \cup E \cup E^*$. Let T be a minimum resistance tree. Suppose there exists an edge $a^0 \rightarrow a^1$ where multiple players switched actions, i.e., there exists a group of players $G \subseteq \mathcal{I}$, $|G| > 1$, such that $a_i^0 \neq a_i^1$ if and only if $i \in G$. We can express the resistance of this transition as

$$R(a^0 \rightarrow a^1) = m|G| + r(a^0 \rightarrow a^1),$$

where $r(a \rightarrow a')$ captures the effect of $(**)$ in equation (12). Because of the assumed normalized potential function,

$$0 \leq r(a^0 \rightarrow a^1) \leq |G|.$$

Construct a new path \mathcal{P} of unilateral deviations

$$\mathcal{P} = \{a^0 = \bar{a}^0 \rightarrow \bar{a}^1 \rightarrow \dots \rightarrow \bar{a}^{|G|} = a^1\},$$

where for all $k \in \{0, 1, \dots, |G| - 1\}$, $\bar{a}^{k+1} = (a_i^1, \bar{a}_{-i}^k)$ for some player $i \in G$. From (12), the resistance of the path \mathcal{P} satisfies

$$m|G| \leq R(\mathcal{P}) \leq m|G| + |G|.$$

Construct a new tree T' by adding the edges of \mathcal{P} to T and removing the redundant edges. The removed resistance of the edge emanating from the originating node is at least $m|G|$. Furthermore, the removed resistance from each of the remaining edges is at least m (by Claim 6.2). Therefore, the overall removed resistance is at least $m|G| + m(|G| - 1)$. The added path \mathcal{P} has resistance of at most $(m+1)|G|$. Accordingly, the new tree has strictly less resistance as long as $|G| \geq 2$. This contradicts the original T being a minimum resistance tree.

□

Claim 6.4 *A minimum resistance tree must be rooted at an action profile that maximizes the potential function.*

Proof: Let T be a minimum resistance tree rooted at $[a^0] \in D$. There exists a better reply path \mathcal{P} from a^0 to a Nash equilibrium, a^* ,

$$\mathcal{P} = \{a^0 \rightarrow a^1 \rightarrow \dots \rightarrow a^k = a^*\}.$$

Since \mathcal{P} is a better reply path, the resistance of the path is precisely $R(\mathcal{P}) = km$. Construct a new tree T' rooted at a^m by adding the edges of \mathcal{P} to T and removing the redundant edges. This results in $R(T) > R(T')$, because the removed edge exiting a^* has resistance $R(a^* \rightarrow a^l) > m$, which is a contradiction. Accordingly, a minimum resistance tree cannot be rooted at a state in D .

Suppose that a minimum resistance tree, T , was rooted at a Nash equilibrium, a^0 , that did not maximize the potential function, i.e., $a^0 \in E$. Let a^* be any action profile that maximizes the potential function. Since T is a rooted tree, there exists a path \mathcal{P} from a^* to a^0 of the form

$$\mathcal{P} = \{a^* = a^k \rightarrow a^{k-1} \rightarrow \dots \rightarrow a^0\},$$

where each edge $a^{l+1} \rightarrow a^l$ is the result of unilateral deviation. Consider the reverse path \mathcal{P}^R

$$\mathcal{P}^R = \{a^0 \rightarrow a^1 \rightarrow \dots \rightarrow a^k = a^*\}.$$

Exploiting the similarities between payoff-based log-linear learning with unilateral updates and binary log-linear learning, we can use Lemma 5.2 to show that the difference in the total resistance across the paths is

$$R(\mathcal{P}) - R(\mathcal{P}^R) = \phi(a^0) - \phi(a^*).$$

Construct a new tree T' rooted at a^* by adding the edges of \mathcal{P}^R to T and removing the redundant edges \mathcal{P} . The new tree will have resistance

$$\begin{aligned} R(T') &= R(T) + R(\mathcal{P}^R) - R(\mathcal{P}), \\ &= R(T) + \phi(a^0) - \phi(a^*), \\ &< R(T). \end{aligned}$$

Since the new tree T' rooted at a^* has strictly less resistance than T , T cannot be a minimum resistance tree.

□

7 Conclusion

The central issue explored in this paper is whether the structural requirements of log-linear learning are necessary to guarantee convergence to the potential function maximizer in exact potential games. In particular,

we have shown that the requirements of asynchrony and completeness are not necessary. The key enabler for these results is to change the focus of the analysis away from deriving the explicit form of the stationary distribution of the learning process towards characterizing the stochastically stable states. Establishing that log-linear learning constitutes a regular perturbed Markov process allows the utilization of theory of resistance trees, as in [20], to analyze the limiting properties of variations on log-linear learning.

References

- [1] C. Alos-Ferrer and N. Netzer. The logit-response dynamics. working paper, 2007.
- [2] 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*, 129:584–596, September 2007.
- [3] A. Beggs. Waiting times and equilibrium selection. *Economic Theory*, 25:599–628, 2005.
- [4] L. Blume. The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5:387–424, 1993.
- [5] L. Blume. Population games. In B. Arthur, S. Durlauf, and D. Lane, editors, *The Economy as an evolving complex system II*, pages 425–460. Addison-Wesley, Reading, MA, 1997.
- [6] L. Blume. How noise matters. *Games and Economic Behavior*, 44:251–271, 2003.
- [7] S. Mannor and J.S. Shamma. Multi-agent learning for engineers. 2007. forthcoming special issue in *Artificial Intelligence*.
- [8] J. R. Marden, G. Arslan, and J. S. Shamma. Connections between cooperative control and potential games illustrated on the consensus problem. In *Proceedings of the 2007 European Control Conference (ECC '07)*, July 2007.
- [9] J. R. Marden, G. Arslan, and J. S. Shamma. Regret based dynamics: Convergence in weakly acyclic games. In *Proceedings of the 2007 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, Honolulu, Hawaii, May 2007.

- [10] J. R. Marden, G. Arslan, and J. S. Shamma. Joint strategy fictitious play with inertia for potential games. *IEEE Transactions on Automatic Control*, 2008. to appear.
- [11] J. R. Marden, H. P. Young, G. Arslan, and J. S. Shamma. Payoff based dynamics for multi-player weakly acyclic games. *SIAM Journal of Control and Optimization*, 2007. submitted to.
- [12] D. Monderer and L. Shapley. Fictitious play property for games with identical interests. *Games and Economic Theory*, 68:258–265, 1996.
- [13] D. Monderer and L.S. Shapley. Potential games. *Games and Economic Behavior*, **14**:124–143, 1996.
- [14] Tim Roughgarden. *Selfish Routing and the Price of Anarchy*. The MIT Press, Cambridge, MA, 2005.
- [15] W. H. Sandholm and M. Benaim. Logit evolution in potential games: Reversibility, rates of convergence, large deviations, and equilibrium selection. working paper, 2007.
- [16] L. S. Shapley. Stochastic games. *Proceedings of the National Academy of Sciences of the United States of America*, 39(10):1095–1100, 1953.
- [17] Y. Shoham, R. Powers, and T. Grenager. If multi-agent learning is the answer, what is the question? forthcoming special issue in *Artificial Intelligence*, 2007.
- [18] A. Vetta. Nash equilibria in competitive societies with applications to facility location, traffic routing, and auctions. In *Proc. of Symp. on Fdns. of Comp. Sci.*, pages 416–425, 2002.
- [19] M. Voorneveld. Best-response potential games. *Economic Letters*, 66:289–295, 2000.
- [20] H. P. Young. The evolution of conventions. *Econometrica*, 61(1):57–84, January 1993.
- [21] H. P. Young. *Individual Strategy and Social Structure*. Princeton University Press, Princeton, NJ, 1998.
- [22] H. P. Young. *Strategic Learning and its Limits*. Oxford University Press, 2005.
- [23] H. P. Young. Learning by trial and error. *Games and economic behavior*, 2008. forthcoming.

8 Appendix: Background on Resistance Trees

The following is a very brief summary of the detailed review presented in [20]. Let P^0 denote the probability transition matrix for a finite state Markov chain over the state space Z . Consider a “perturbed” process such that the size of the perturbations can be indexed by a scalar $\epsilon > 0$, and let P^ϵ be the associated transition probability matrix. The process P^ϵ is called a *regular perturbed Markov process* if P^ϵ is ergodic for all sufficiently small $\epsilon > 0$ and P^ϵ approaches P^0 at an exponentially smooth rate [20]. Specifically, the latter condition means that $\forall z, z' \in Z$,

$$\lim_{\epsilon \rightarrow 0^+} P_{z \rightarrow z'}^\epsilon = P_{z \rightarrow z'}^0,$$

and

$$P_{z \rightarrow z'}^\epsilon > 0 \text{ for some } \epsilon > 0 \Rightarrow 0 < \lim_{\epsilon \rightarrow 0^+} \frac{P_{z \rightarrow z'}^\epsilon}{\epsilon^{R(z \rightarrow z')}} < \infty,$$

for some nonnegative real number $R(z \rightarrow z')$, which is called the *resistance* of the transition $z \rightarrow z'$. (Note in particular that if $P_{z \rightarrow z'}^0 > 0$ then $R(z \rightarrow z') = 0$.)

Let the recurrence classes of P^0 be denoted by E_1, E_2, \dots, E_N . For each pair of distinct recurrence classes E_i and E_j , $i \neq j$, an ij -path is defined to be a sequence of distinct states $\zeta = (z_1 \rightarrow z_2 \rightarrow \dots \rightarrow z_n)$ such that $z_1 \in E_i$ and $z_n \in E_j$. The resistance of this path is the sum of the resistances of its edges, that is, $R(\zeta) = R(z_1 \rightarrow z_2) + R(z_2 \rightarrow z_3) + \dots + R(z_{n-1} \rightarrow z_n)$. Let $\rho_{ij} = \min R(\zeta)$ be the least resistance over all ij -paths ζ . Note that ρ_{ij} must be positive for all distinct i and j , because there exists no path of zero resistance between distinct recurrence classes.

Now construct a complete directed graph with N vertices, one for each recurrence class. The vertex corresponding to class E_j will be called j . The weight on the directed edge $i \rightarrow j$ is ρ_{ij} . A tree, T , rooted at vertex j , or j -tree, is a set of $N - 1$ directed edges such that, from every vertex different from j , there is a unique directed path in the tree to j . The resistance of a rooted tree, T , is the sum of the resistances ρ_{ij} on the $N - 1$ edges that compose it. The *stochastic potential*, γ_j , of the recurrence class E_j is defined to be the minimum resistance over all trees rooted at j . The following theorem gives a simple criterion for determining the stochastically stable states ([20], Theorem 4).

Theorem 8.1 *Let P^ϵ be a regular perturbed Markov process, and for each $\epsilon > 0$ let μ^ϵ be the unique stationary distribution of P^ϵ . Then $\lim_{\epsilon \rightarrow 0} \mu^\epsilon$ exists and the limiting distribution μ^0 is a stationary distribution of P^0 . The stochastically stable states (i.e., the support of μ^0) are precisely those states contained in the recurrence classes with minimum stochastic potential.*