Understanding the Influence of Adversaries in Distributed Systems

Holly P. Borowski and Jason R. Marden

Abstract— **In a multi-agent system, transitioning from a centralized to a distributed decision-making strategy can introduce vulnerability to adversarial manipulation. We study the potential for adversarial manipulation in a class of graphical coordination games where the adversary can pose as a friendly agent in the game, thereby directly influencing the decision-making rules of a subset of agents. The adversary's influence can cascade throughout the system, indirectly influencing other agents' behavior and significantly impacting the emergent collective behavior. The main results in this paper focus on characterizing conditions by which the adversary's local influence can dramatically impact the emergent global behavior, e.g., destabilize efficient equilibria.**

I. Introduction

Engineering and social systems often consist of many agents making decisions based on locally available information. In an engineering system, a distributed decision making strategy can be necessary when communication, computation, or sensing limitations preclude a centralized control strategy. For example, a group of unmanned aircraft performing surveillance in a hostile area may use a distributed control strategy to limit communication and thus remain undetected. Social systems are inherently distributed: individuals typically make decisions based on personal objectives and the behavior of friends and acquaintances. For example, the decision to adopt a recently released technology, such as a new smartphone, may depend both on the quality of the item itself and on friends' choices.

While there are many advantages of distributed decision making, it can create vulnerability to adversarial manipulation. Adversaries may attempt to influence individual agents by corrupting the information available to them, creating a chain of events which could degrade the system's performance. Work in the area of cyber-physical systems has focused on reducing the potential impact of adversarial interventions through detection mechanisms: detection of attacks in power networks [8], estimation and control with corrupt sensor data [1], [4], and monitoring [14]. In contrast to this research, our work focuses on characterizing the impact an adversary may have on distributed system dynamics when no mitigation or detection measures are in place.

We use graphical coordination games, introduced in [3], [17], to study the impact of adversarial manipulation. The foundation of a graphical coordination game is a simple two agent coordination game, where each agent must choose between one of two alternatives, $\{x, y\}$, with payoffs depicted by the following payoff matrix which we denote by $u(\cdot)$:

	x	y
x	$1+\alpha, 1+\alpha$	$0, 0$
y	$0, 0$	$1, 1$

2×2 coordination game, g, with utilities $u(a_i, a_j)$, $a_i, a_j \in \{x, y\}$, and payoff gain $\alpha > 0$

where $\alpha > 0$ defines the relative quality of conventions (x, x) over (y, y). Both agents prefer to agree on a convention, i.e., (x, x) or (y, y), than disagree, i.e., (x, y) or (y, x), with a preference to agreeing on (x, x). The goal of deriving local agent dynamics which lead to the efficient Nash equilibrium (x, x) is challenging because of the existence of the inefficient Nash equilibrium (y, y). Deviating from (y, y) for an individual agent is accompanied by an immediate payoff loss of 1 to 0; hence, myopic agents may be reluctant to deviate, stabilizing the inefficient equilibrium (y, y).

This two player coordination game can be extended to an n-player *graphical coordination game* [9], [13], [19], where the interactions between the agents $N = \{1, 2, \ldots, n\}$ is described by an underlying graph $\mathcal{G} = (N, E)$ where $E \subseteq N \times N$ denotes the interdependence of agents' objectives. More formally, an agent's total payoff is the sum of payoffs it receives in the two player games played with its neighbors $\mathcal{N}_i = \{j \in N : (i, j) \in E\}$, i.e., for a joint decision $a = (a_1, \ldots, a_n) \in \{x, y\}^n$, the utility of agent i is

$$U_i(a_1, \ldots, a_n) = \sum_{j \in \mathcal{N}_i} u(a_i, a_j). \tag{1}$$

Joint actions $\vec{x} := (x, x, \ldots, x)$ and $\vec{y} := (y, y, \ldots, y)$, where either all players choose x or all players choose y, are Nash equilibria of the game; other equilibria may emerge depending on the structure of graph $\mathcal{G}$. In any case, $\vec{x}$ is the unique efficient equilibrium, since it maximizes agents' total payoffs. Graphical coordination games can model both task allocation in engineering systems as well as the evolution of social convention in marketing scenarios.

The goal in this setting is to prescribe a set of decision-making rules that ensures emergent behavior is aligned with the efficient Nash equilibrium $\vec{x}$ irrespective of the underlying graph $\mathcal{G}$ and the choice of α. Any such rule must be accompanied by a degree of noise (or mistakes) as agents must be enticed to deviate from inefficient Nash equilibrium. Log-linear learning [2], [15] is one distributed decision making rule that selects the efficient equilibrium, $\vec{x}$, in the graphical coordination game described above. Although agents predominantly maximize their utilities under

This research was supported by AFOSR grant #FA9550-12-1-0359, ONR grant #N00014-12-1-0643, NSF grant #ECCS-1351866, the NASA Aeronautics Scholarship Program, and the Zonta Amelia Earhart Fellowship Program.

H. P. Borowski is a graduate research assistant with the Department of Aerospace Engineering, University of Colorado, Boulder, holly.borowski@colorado.edu.

J. R. Marden is with the Department of Electrical, Computer, and Energy Engineering, University of Colorado, Boulder, jason.marden@colorado.edu.

log-linear learning, selection of the efficient equilibrium is achieved by allowing agents to choose suboptimally with some small probability that decreases exponentially with respect to the associated payoff loss.

The equilibrium selection properties of log-linear learning extend beyond coordination games to the class of potential games [12], which often can be used to model engineering systems where the efficient Nash equilibrium is aligned with the optimal system behavior [10], [11], [18]. Hence, log-linear learning can be a natural choice for prescribing control laws in many distributed engineering systems [6], [7], [11], [16], [21], as well as for analyzing the emergence of conventions in social systems [15], [20]. This prompts the question: can adversarial manipulation alter the emergent behavior of log-linear learning in the context of graphical coordination games (or more broadly in distributed engineering systems)?

We study this question in the context of the above graphical coordination games. Here, we model the adversary as additional nodes/edges in our graph, where the action selected by these adversaries (which we fix as the inferior convention y) impacts the utility of the neighboring agents and thereby influences the agents' decision-making rule as specified by log-linear learning. We focus on three different models of adversary behavior, referred to as *fixed, intelligent*; *mobile, random*; and *mobile, intelligent.*

- A fixed intelligent adversary aims to influence a fixed set $S \subseteq N$. To these agents the adversary appears to be a neighbor who always selects alternative y. We assume that S is selected based on the graph structure $\mathcal{G}$ and α.
- A mobile, random adversary connects to a random collection of agents $S(t) \subseteq N$ at each time, $t \in \mathbb{N}$ using no information on graph structure, $\mathcal{G}$, or payoff gain, α.
- A mobile, intelligent agent connects to a collection of agents, $S(t) \subseteq N$, at each time, $t \in \mathbb{N}$ using information on graph structure, $\mathcal{G}$, payoff gain α, and the current action profile, $a(t)$.

We will discuss each type of adversary's influence on an arbitrary graph, and then analyze the worst case influence on a set of agents interacting according to a line. We specify the values of payoff gain α for which an adversary can stabilize joint action $\vec{y}$, showing that a mobile, intelligent agent can typically stabilize joint action $\vec{y}$ for larger values of α than a mobile, random agent, and a mobile, random agent can typically stabilize $\vec{y}$ for larger values of α than a fixed, intelligent agent.

A. The model

Suppose agents in N interact according to the graphical coordination game above, with underlying graph $\mathcal{G} = (N, E)$, alternatives $\{x, y\}$ and payoff gain α. We denote the joint action space by $\mathcal{A} = \{x, y\}^n$, and we write

$$(a_i, a_{-i}) = (a_1, a_2, \ldots, a_i, \ldots, a_n) \in \mathcal{A}$$

when considering agent i's action separately from other agents' actions.

Now, suppose agents in N update their actions according to the *log-linear learning* algorithm at times $t = 0, 1, \ldots,$ producing a sequence of joint actions $a(0), a(1), \ldots.$ We assume agents begin with joint action, $a(0) \in \mathcal{A}$, and let $a(t) = (a_i, a_{-i}) \in \mathcal{A}$. At time $t \in \mathbb{N}$, an agent $i \in N$ is selected uniformly at random to update its action for time $t + 1$; all other agents' actions will remain fixed. Agent i chooses its next action probabilistically according to:[1]

$$\begin{aligned}\Pr[a_i(t+1) = x \mid a_{-i}(t) = a_{-i}]& \\ = \frac{\exp\left(\beta \cdot U_i(x, a_{-i})\right)}{\exp\left(\beta \cdot U_i(x, a_{-i}) + \exp\left(\beta \cdot U_i(y, a_{-i})\right)\right.}&. \qquad (2)\end{aligned}$$

Parameter $\beta > 0$ dictates an updating agent's degree of rationality. As $\beta \to \infty$, agent i is increasingly likely to select a utility maximizing action, and as $\beta \to 0$, agent i tends to choose its next action uniformly at random. The joint action at time $t + 1$ is $a(t+1) = (a_i(t+1), a_{-i}(t))$.

Joint action, $a \in \mathcal{A}$ is *strictly stochastically stable* [5] under log-linear learning dynamics if, for any $\varepsilon > 0$, there exist $B < \infty$ and $T < \infty$ such that

$$\Pr[a(t) = a] > 1 - \varepsilon, \quad \text{for all } \beta > B, t > T \qquad (3)$$

where $a(t)$ is the joint action at time $t \in \mathbb{N}$ under log-linear learning dynamics.

Joint action $\vec{x}$ is strictly stochastically stable under log-linear learning dynamics over graphical coordination game G [2]. We will investigate conditions when an adversary can destabilize $\vec{x}$ and stabilize an alternate equilibrium.

Consider the situation where agents in N interact according to the graphical game G, and an adversary seeks to convert as many agents in N to play action y as possible.[2] At each time, $t \in \mathbb{N}$ the adversary attempts to influence a set of agents $S(t) \subseteq N$ by posing as a friendly agent who always plays action y. Agents' utilities, $\tilde{U} : \mathcal{A} \times 2^N \to \mathbb{R}$, are now a function of adversarial and friendly behavior, defined by:

$$\tilde{U}_i((a_i, a_{-i}), S) = \begin{cases} U_i(a_i, a_{-i}) & \text{if } i \notin S \\ U_i(a_i, a_{-i}) & \text{if } a_i = x \\ U_i(a_i, a_{-i}) + 1 & \text{if } i \in S, a_i = y \end{cases} \qquad (4)$$

where $(a_i, a_{-i}) \in \mathcal{A}$ represents friendly agents' joint action, and $S \subseteq N$ represents the set influenced by the adversary. If $i \in S(t)$, agent i receives an additional payoff of 1 for coordinating with the adversary at action y at time $t \in \mathbb{N}$; to agents in $S(t)$ the adversary appears to be a neighbor playing action y. By posing as a player in the game, the adversary has manipulated the utilities of agents belonging to S, providing an extra incentive to choose the inferior alternative, y.

Suppose agents revise their actions according to log-linear learning as in (2), where the utility, U_i defined in (1) is replaced by $\tilde{U}_i$ in (4). An agent $i \in N$ which revises its action at time $t \in \mathbb{N}$ bases its new action choice on the utility $\tilde{U}_i(a(t), S(t))$ if $i \in S(t)$, increasing the probability that agent i updates its action to y. By posing as a player in the

[1] Agent i's update probability is also conditioned on the fact that agent i was selected to revise its action, which occurs with probability $1\,/\,n$. For notational brevity we omit this throughout, and $\Pr[a_i(t+1) = A \,|\, a_{-i}(t) = a_{-i}]$, for example, is understood to mean $\Pr[a_i(t+1) = x \,|\, a_{-i}(t) = a_{-i},\, i$ selected for update$]$.

[2] In this paper we consider a single adversary which may influence multiple agents. Our models can be extended to multiple adversaries whose objectives are either aligned or conflicting.

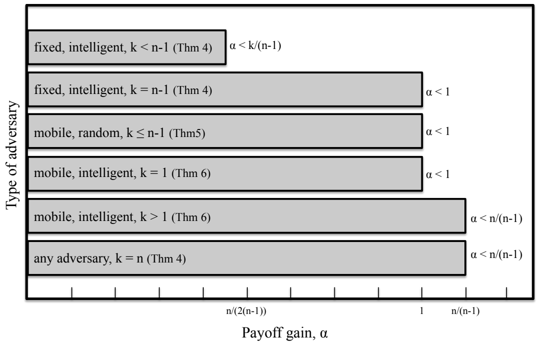

Fig. 1: Values of α for which each type of adversary can stabilize joint action $\vec{y}$ in an n-agent line

coordination game, an adversary manipulates agents' utility functions.thereby modifying their decision making rules.

B. Summary of results

In the following sections, we will precisely define three models of adversarial behavior: fixed, intelligent; mobile, random; and mobile, intelligent. Each type of adversary has a fixed capability, k, i.e., $|S(t)| = k$ for all $t \in \mathbb{N}$. Our analysis of these models will provide insight into an adversary's influence on a general graph, $\mathcal{G}$, and we derive exact bounds on α for adversarial influence on a line. Values of α for which each type of agent can stabilize $\vec{y}$ in the line are summarized below and in Figure 1.

- A fixed, intelligent adversary with capability k can stabilize joint action $\vec{y}$ when $\alpha < k\,/(n-1)$ (Theorem 4).
- A mobile, random adversary with capability $k \leq n-1$ can stabilize joint action $\vec{y}$ when $\alpha < 1$ (Theorem 5).
- A mobile, intelligent adversary with capability $k = 1$ can stabilize joint action $\vec{y}$ when $\alpha < 1$ (Theorem 6).
- A mobile, intelligent adversary with capability $k \geq 2$ can stabilize joint action $\vec{y}$ when $\alpha < n\,/(n-1)$ (Theorem 6).

Note that a mobile, random adversary's influence is the same for any capability k with $1 \leq k \leq n-1$. Similarly, a mobile, intelligent adversary does not increase its influence on agents in a line by increasing its capability above $k = 2$.

II. Universal resilience to an adversary

A graphical coordination game G is universally resilient to an adversary if $\vec{x}$ is strictly stochastically stable for all possible influenced sets $S(t)$, $t \in N$ and adversarial capability, $k \leq n$. The following theorem provides sufficient conditions that ensure G is universally resilient. For sets $S, T \subseteq N$, define

$$d(S,T) := |\{\{i,j\} \in E \,:\, i \in S, j \in T\}|.$$

Theorem 1: Let $\mathcal{G} = (N,E)$, and suppose an adversary influences some set $S(t)$ with $|S(t)| = k$ at each $t \in \mathbb{N}$. If

$$\alpha > \frac{|T| - d(T, N \setminus T)}{d(T,N)}, \quad \forall T \subseteq N \tag{5}$$

Then $\vec{x}$ is strictly stochastically stable. In particular, if $|S(t)| = N$ for all $t \in \mathbb{N}$, (5) is also a necessary condition for strict stochastic stability of $\vec{x}$.

The proof of Theorem 1 follows by using a straightforward adaptation of Proposition 2 in [19] to our adversarial model, included in Appendix B

When α satisfies (5), an adversary cannot influence the game for any $S(t)$. If $\vec{x}$ is strictly stochastically stable when the adversary influences set $S(t) = N$ for all $t \in \mathbb{N}$, then $\vec{x}$ will be strictly stochastically stable for any sequence of influenced sets, $S(t) \subseteq N$. In this case, game G is resilient in the presence of any adversary with capability $k \leq n$.[3]

When (5) is satisfied for some $T \subseteq N$, this means that agents in T have a sufficiently large proportion of neighbors in N. In this case, T can only be influenced by an adversary when the payoff gain, α, is small.

III. Fixed, intelligent adversarial influence

In the fixed, intelligent model of behavior, the adversary knows graph structure, $\mathcal{G}$, and the value of payoff gain, α. Using this information it influences some fixed subset,

$$S(t) = S \subseteq N, |S| = k, \quad \forall t \in \mathbb{N},$$

aiming to maximize the number of agents playing y in a stochastically stable state. Agents in N update their actions according to log-linear learning as in (2) with utilities

$$\tilde{U}_i(a(t), S(t)) = \tilde{U}_i(a(t), S), \quad \forall t \in \mathbb{N}.$$

We begin with two theorems which provide conditions for stochastic stability in an arbitrary graph $\mathcal{G}$ influenced by an adversary, and then we analyze stability conditions in detail for the line.

Theorem 2: Suppose agents in N are influenced by a fixed, intelligent adversary with capability k. Joint action $\vec{x}$ is strictly stochastically stable for any influenced set S with $|S| = k$ if and only if

$$\alpha > \frac{|T \cap S| - d(T, N \setminus T)}{d(T,N)}, \tag{6}$$

$\forall T \subseteq N$, $T \neq \emptyset$ and $\forall S \subseteq N$ with $|S| = k$.

Theorem 3 provides conditions which ensure an adversary can stabilize joint action $\vec{y}$.

Theorem 3: A fixed, intelligent adversary with capability k can stabilize $\vec{y}$ by influencing set $S \subseteq N$ with $|S| = k$ if and only if

$$\alpha < \frac{d(T, N \setminus T) + k - |T \cap S|}{d(N \setminus T, N \setminus T)} \tag{7}$$

for all $T \subseteq N$, $T \neq N$.

The proofs of Theorems 2 and 3 follow similarly to the proof of Theorem 1 and are omitted for brevity.

The line: We now analyze a fixed, intelligent adversary's influence on the line. Let $\mathcal{G} = (N,E)$ with $N = \{1, 2, \ldots, n\}$

[3] Our results can naturally be extended to a multi-agent scenario. The primary differences occur when multiple adversaries can influence a single friendly agent (or, equivalently, when an adversary's influence is weighted by some factor greater than 1). In this scenario, multiple adversaries can more easily overpower the influence of friendly agents on agent i. We will address this in future work.

and $E = \{\{i,j\} : j = i+1\}$, i.e., $\mathcal{G}$ is a line with n nodes. Define

$$[t] := \{1, 2, \ldots, t\} \subseteq N, \text{ and } [i,j] := \{i, i+1, \ldots, j\} \subseteq N.$$

Theorem 4 summarizes stability conditions for the line influenced by a fixed, intelligent adversary.

Theorem 4: Suppose $\mathcal{G}$ is influenced by a fixed, intelligent adversary with capability k. Then:

(a) Joint action $\vec{x}$ is strictly stochastically stable under any influenced set $S \subseteq N$ with $|S| = k$ if and only if

$$\alpha > \max\left\{\frac{k-1}{k}, \frac{k}{n-1}\right\}. \tag{8}$$

(b) If $\alpha < \frac{k}{n-1}$ and the adversary distributes influenced set S as evenly as possible along the line, so that

$$|S \cap [i, i+t]| \leq \left\lceil \frac{kt}{n} \right\rceil$$

for any set of nodes $[i, i+t] \subseteq N$, with $1 \leq i \leq n-t$, $t \leq n$ then $\vec{y}$ is strictly stochastically stable.

(c) Joint action $\vec{y}$ is strictly stochastically stable for all influenced sets S with $|S| = k$ if and only if

$$\alpha < \frac{1+k-t}{n-t-1}, \quad \forall t = 1, \ldots, k. \tag{9}$$

(d) If $\frac{k}{n-1} < \alpha < \frac{k-1}{k}$, the adversary can influence at most

$$t_{\max} = \max\left\{t \,:\, \alpha < \frac{\min\{t,k\}-1}{t}\right\}$$

agents to play $\vec{y}$ in the stochastically stable state by distributing S as evenly as possible along $[t]$, so that

$$|S \cap [i, i+\ell]| \leq \left\lceil \frac{k\ell}{t} \right\rceil \text{ and } S \cap [t+1, n] = \emptyset$$

for any set of nodes $[i, i+\ell] \subset N$ with $1 \leq i \leq t-\ell$, and $\ell < t$.

The proof of Theorem 4 is in Appendix B.

IV. Mobile, random adversarial influence

Now, consider an adversary which influences a randomly chosen set $S(t) \subseteq N$ at each $t \in \mathbb{N}$. The adversary chooses each influenced set, $S(t)$, independently according to a uniform distribution over $\mathcal{S}_k := \{S \in 2^N : |S| = k\}$ An updating agent $i \in N$ revises according to (2), where $i \in S(t)$ with probability $k \,/\, n$.

The line: Suppose a mobile, random adversary attempts to influence a set of agents arranged in a line. Theorem 4 addresses the scenario where $k = n$, since in this case random and fixed agents are equivalent. Hence, Theorem 5 focuses on the case where $1 \leq k \leq n-1$.

Theorem 5: Suppose $\mathcal{G} = (N, E)$ is a line, and agents in N update their actions according to log-linear learning in the presence of a random, mobile adversary with capability k, where $1 \leq k \leq n-1$. Then joint action $\vec{x}$ is strictly stochastically stable if and only if $\alpha > 1$, and joint action $\vec{y}$ is strictly stochastically stable if and only if $\alpha < 1$.

Theorem 5 is proved in Appendix D.

Note that a mobile, random adversary with capability $k = 1$ stabilizes $\vec{y}$ for the same values of α as a mobile, random adversary with any capability $k \leq n-1$. Recall that a fixed, intelligent adversary with capability k could only stabilize $\vec{y}$ when $\alpha < k \,/ (n-1)$. In this sense, a mobile, random adversary with capability $k = 1$ has wider influence than a fixed, intelligent adversary with capability $k \leq n-2$.

V. Mobile, intelligent adversarial influence

Now suppose the adversary chooses $S(t)$ at each $t \in \mathbb{N}$ based on joint action, $a(t)$. We assume a mobile, intelligent adversary with capability k chooses $S(t)$ according to a policy $\mu : \mathcal{A} \to \mathcal{S}_k$ that maximizes the number of agents playing y in a stochastically stable state, given graph structure, $\mathcal{G}$, and payoff gain α. Again, agents in N update their actions according to log-linear learning as in (2), with agent i's utility at time $t \in \mathbb{N}$ given by $\widetilde{U}_i(a(t), \mu(a(t))$. We denote the set of optimal adversarial policies for a given capability k by

$$\mathcal{M}_k = \underset{\mu \in M_k}{\arg\max} \max_{a \text{ stable under } \mu} |\{i \in N \,:\, a_i = y\}| \tag{10}$$

where M_k represents the set of all mappings $\mu : \mathcal{A} \to \mathcal{S}_k$, and "$a$ stable under μ" denotes that joint action $a \in \mathcal{A}$ is strictly stochastically stable under μ. [4]

The line: Theorem 6 establishes conditions for strict stochastic stability of joint actions $\vec{x}$ and $\vec{y}$ in the line influenced by a mobile, intelligent adversary.

Theorem 6: Suppose $\mathcal{G} = (N, E)$ is a line, and agents in N update their actions according to log-linear learning. Further suppose a mobile intelligent adversary influences set $S(t)$ at each $t \in \mathbb{N}$ according to an optimal policy for the line, $\mu^\star \in \mathcal{M}_k$.

(a) If the adversary has capability $k = 1$ then $\vec{x}$ is strictly stochastically stable if and only if $\alpha > 1$, and $\vec{y}$ is strictly stochastically stable if and only if $\alpha < 1$.

In particular, when $k = 1$, the policy $\mu^\star : \mathcal{A} \to \mathcal{S}_1$ with:

$$\mu^\star(a) = \begin{cases} \{1\} & \text{if } a = \vec{x} \\ \{t+1\} & \text{if } a = (\vec{y}_{[t]}, \vec{x}_{[t+1,n]}), \\ & \qquad t \in \{1, 2, \ldots, n-1\} \\ \{1\} & \text{otherwise} \end{cases} \tag{11}$$

is optimal, i.e., $\mu^\star \in \mathcal{M}_1$

(b) If $2 \leq k \leq n$, then $\vec{x}$ is strictly stochastically stable if and only if $\alpha > n \,/ (n-1)$, and $\vec{y}$ is strictly stochastically stable if and only if $\alpha < n \,/ (n-1)$.

If $2 \leq k \leq n-1$, any policy $\mu^\star : \mathcal{A} \to \mathcal{S}_k$ satisfying:

(1) $1 \in \mu^\star(\vec{x})$

(2) $1, n \in \mu^\star(\vec{y})$

(3) For any $a \in \mathcal{A}$, $a \neq \vec{x}, \vec{y}$, there exists $i \in \mu^\star(a)$ such that $a_i = x$ and either $a_{i-1} = y$ or $a_{i+1} = y$

is optimal.

The proof of Theorem 6 is included in Appendix E. Recall that a mobile, random agent with $k \geq 1$ and a fixed, intelligent agent with $k = n-1$ can stabilize $\vec{y}$ any time $\alpha < 1$; an adversary who can intelligently influence a

[4] Note that the optimal set of policies, $\mathcal{M}_k$, depends highly on the structure of graph $\mathcal{G}$, as does the stationary distribution π^μ. In order to maintain notational simplicity, we do not explicitly write this dependence.

different single agent in N each day can stabilize $\vec{y}$ under these same conditions. If the intelligent, mobile adversary has capability $k \geq 2$, it can stabilize $\vec{y}$ when $\alpha < n/(n-1)$, i.e., under the same conditions as an adversary with capability $k = n$.

VI. Summary and future work

We have shown that a mobile, intelligent adversary with capability $k \geq 2$ can stabilize joint action $\vec{y}$ in a line for any $\alpha < n/(n-1)$. Next, an intelligent, mobile adversary with capability $k = 1$ and a random, mobile adversary with capability $k \leq n-1$ can stabilize $\vec{y}$ when $\alpha < 1$. Finally, a fixed, intelligent adversary with capability k can stabilize $\vec{y}$ when $\alpha < k/(n-1)$. Recall that a fixed, intelligent adversary can also stabilize a joint action where some subset of agents play action y; this only occurs when $\alpha < (\min\{t,k\}-1)/t < 1$ for some $t \leq n$.

In future work, we will address the scenario where multiple adversaries aim to influence agents in N.By heavily influencing a single agent, adversaries can cause this agent to choose action y with near certainty. Due to cascading effects, this can allow adversaries to stabilize joint action $\vec{y}$ for significantly larger values of payoff gain, α.

References

[1] C.Z. Bai and V. Gupta. On Kalman filtering in the presence of a compromised sensor: Fundamental performance bounds. *Proceedings of the American Control Conference*, 2014.

[2] L. E. Blume. The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 1993.

[3] R. Cooper. *Coordination Games*. Cambridge University Press, Cambridge, UK, 1999.

[4] H. Fawzi, P. Tabuada, and S. Diggavi. Secure estimation and control for cyber-physical systems under adversarial attacks. *IEEE Transactions on Automatic Control*, 59:1454–1467, 2014.

[5] D. Foster and H.P. Young. Stochastic evolutionary game dynamics. *Theoretical Population Biology*, 38:219–232, 1990.

[6] M. J. Fox and J. S. Shamma. Communication, convergence, and stochastic stability in self-assembly. In *49th IEEE Conference on Decision and Control (CDC)*, December 2010.

[7] 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.

[8] J.M. Hendrickx, K.H Johansson, R.M. Jungers, H. Sandberg, and K.C. Sou. Efficient Computations of a Security Index for False Data Attacks in Power Networks. *IEEE Transactions on Automatic Control*, 59(12):3194–3208, 2014.

[9] M. Kearns, M.L. Littman, and S. Singh. Graphical Models for Game Theory. In *17th Conference in Uncertainty in Artificial Intelligence*, 2001.

[10] J. R. Marden, G. Arslan, and J. S. Shamma. Regret based dynamics: convergence in weakly acyclic games. In *Sixth International Joint Conference on Autonomous Agents and Multi-Agent Systems*, volume 5, 2007.

[11] J. R. Marden and A. Wierman. Distributed welfare games. *Operations Research*, 61(1):155–168, 2013.

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

[13] A. Montanari and A. Saberi. The spread of innovations in social networks. *Proceedings of the National Academy of Sciences*, pages 20196–20201, 2010.

[14] F. Pasqualetti, F. Dörfler, and F. Bullo. Cyber-physical security via geometric control: Distributed monitoring and malicious attacks. *Proceedings of the IEEE Conference on Decision and Control*, (i):3418–3425, 2012.

[15] D. Shah and J. Shin. Dynamics in Congestion Games. In *ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, 2010.

[16] M. Staudigl. Stochastic stability in asymmetric binary choice coordination games. *Games and Economic Behavior*, 75(1):372–401, May 2012.

[17] E. Ullmann-Margalit. *The Emmergence of Norms*. Oxford University Press, 1977.

[18] D. H. Wolpert and K. Tumer. Optimal Payoff Functions for Members of Collectives. *Advances in Complex Systems*, 4:265–279, 2001.

[19] H. P. Young. Colloquium Paper: The dynamics of social innovation. *Proceedings of the National Academy of Sciences*, 108:21285–21291, 2011.

[20] H.P. Young. The Evolution of Conventions. *Econometrica*, 61(1):57–84, 1993.

[21] 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. Log-linear learning and its underlying Markov process

Log-linear learning dynamics define a family of aperiodic, irreducible Markov processes, $\{\tilde{P}_\beta\}_{\beta>0}$, over state space $\mathcal{A} \times \mathcal{S}_k$ with transition probabilities parameterized by β [2]. Under our adversarial model, transition probabilities are

$$\begin{aligned} &P_\beta(((a_i, a_{-i}), S) \to (a_i', a_{-i}), S') \\ &\quad = \frac{1}{n}\Pr[a_i(t+1) = a_i' \mid a_{-i}(t) = a_{-i}, S(t) = S] \end{aligned} \quad (12)$$

for any $i \in N$, $a_i \in \{\vec{x}, \vec{y}\}$, $(a_i, a_{-i}) \in \mathcal{A}$ and $S, S' \in \mathcal{S}_k$. Here S transitions to S' according to the specified adversarial model. If a and $a' \in \mathcal{A}$ differ by more than one agent's action, then $P_\beta(a \to a') = 0$.

For each model of adversarial behavior, it is straightforward to reduce $\tilde{P}_\beta$ to a Markov chain, P_β over state space $\mathcal{A}$. Since P_β is aperiodic and irreducible for any $\beta > 0$, it has a unique stationary distribution, π_β, with $\pi_\beta P_\beta = \pi_\beta$.

As $\beta \to \infty$, the stationary distribution, π_β, associated with log-linear learning converges to a unique distribution, $\pi := \lim_{\beta\to\infty} \pi_\beta$. If $\pi(a) = 1$, then joint action a is *strictly stochastically stable* [5].[5]

As $\beta \to \infty$, transition probabilities $P_\beta(a \to a')$ of log-linear learning converge to the transition probabilities, $P(a \to a')$, of a best response process. Distribution π is one of possibly multiple stationary distributions of a best response process over game G.

B. Stability in the presence of a fixed, intelligent adversary

When a fixed, intelligent adversary influences set S, the corresponding influenced graphical coordination game is a potential game [12] with potential function

$$\Phi^S(a_i, a_{-i}) = \frac{1}{2}\sum_{i\in N}\left(U_i(a_i, a_{-i}) + 2 \cdot \mathbb{1}_{i\in S, a_i = y}\right). \quad (13)$$

This implies that the stationary distribution associated with log-linear learning influenced by a fixed adversary is

$$\pi(a) = \frac{\exp(\beta \cdot \Phi^S(a))}{\sum_{a'\in\mathcal{A}} \exp(\beta \cdot \Phi^S(a'))}, \quad (14)$$

for $a \in \mathcal{A}$ [2]. Hence, $a \in \mathcal{A}$ is strictly stochastically stable if and only if $\Phi^S(a) > \Phi^S(a')$ for all $a' \in \mathcal{A}$, $a' \neq a$.

[5] Note that this definition of strict stochastic stability is equivalent to the definition in the introduction.

Proof of Theorem 1: This proof adapts Proposition 2 in [19] to our adversarial model. Let $\mathcal{G} = (N, E)$ and suppose $S(t) = N$ for all $t \in \mathbb{N}$. Define $(\vec{y}_T, \vec{x}_{N \setminus T})$ to be the joint action $(a_1, \ldots, a_n)$ with $T = \{i : a_i = y\}$. It is straightforward to show that

$$\alpha > \frac{|T| - d(T, N \setminus T)}{d(T, N)}, \quad \forall T \subseteq N$$

if and only if

$$\begin{aligned}\Phi^N(\vec{x}) &= (1+\alpha)d(N, N)\\ &> (1+\alpha)d(N \setminus T, N \setminus T) + d(T,T) + |T|\\ &= \Phi^N(\vec{y}_T, \vec{x}_{N \setminus T}) \end{aligned} \tag{15}$$

for all $T \subseteq N$, $R \neq \emptyset$, implying the desired result. ■

Proof of Theorem 4 part (a): Let $\mathcal{G} = (N, E)$ be a line graph influenced by an adversary with capability k. Joint action $\vec{x}$ is strictly stochastically stable for all $S \subseteq N$ with $|S| = k$ if and only if

$$\begin{gathered}\Phi^S(\vec{x}) > \Phi^S(\vec{y}_T, \vec{x}_{N \setminus T})\\ \Longleftrightarrow\\ (1+\alpha)d(N,N)\\ >\\ (1+\alpha)d(N \setminus T, N \setminus T) + d(T,T) + |S \cap T|. \end{gathered} \tag{16}$$

for all $S \subseteq N$ with $|S| = k$ and all $T \subseteq N$, $T \neq \emptyset$.

Define $t := |T|$, let p denote the number of components in the graph $\mathcal{G}$ restricted to T, and let ℓ denote the number of components in the graph restricted to $N \setminus T$. Since $T \neq \emptyset$, we have $p \geq 1$ and $\ell \in \{p-1, p, p+1\}$.

The case where $T = N$ implies

$$\Phi^S(\vec{x}) = (1+\alpha)(n-1) > n - 1 + k = \Phi^S(\vec{y}),$$

which holds if and only if $\alpha > k\,/(n-1)$.

If $T \subset N$, the graph restricted to $N \setminus T$ has at least one component, i.e., $\ell \geq 1$. Then,

$$\begin{aligned}\Phi^S(\vec{y}_T, \vec{x}_{N \setminus T}) &= (1+\alpha)(n - t - \ell) + t - p + |S \cap T|\\ &\leq (1+\alpha)(n-t-1) + t - 1 + \min\{k, t\}\end{aligned}$$

where the inequality is an equality when $T = [t]$ and $S = [k]$. Then,

$$\begin{aligned}\Phi^S(\vec{y}_T, \vec{x}_{N \setminus T}) &\leq (1+\alpha)(n-t-1) + t - 1 + \min\{k, t\}\\ &< (1+\alpha)(n-1)\\ &= \Phi^S(\vec{x})\end{aligned}$$

for all $T \subset N$ if and only if $\alpha > (k-1)\,/\,k$, as desired. ■

Proof of Theorem 4 part (b): Suppose $\alpha < k\,/(n-1)$. Then

$$\Phi^S(\vec{y}) = n - 1 + k > (1+\alpha)(n-1) = \Phi^S(\vec{x})$$

for any $S \subseteq N$ with $|S| = k$. Then, to show that $\vec{y}$ is stochastically stable for influenced set S satisfying

$$|S \cap [i, i+t]| \leq \left\lceil \frac{kt}{n} \right\rceil,$$

it remains to show that $\Phi^S(\vec{y}) > \Phi^S(\vec{y}_T, \vec{x}_{N \setminus T})$ for any $T \subset N$ with $T \neq \emptyset$ and $T \neq N$. Suppose the graph restricted to set T has p components, where $p \geq 1$. Label these components as $T_1, T_2, \ldots, T_p$ and define $t := |T|$ and $t_i := |T_i|$ Let ℓ represent the number of components in the graph restricted to $N \setminus T$. Since $\mathcal{G}$ is the line graph, we have $\ell \in \{p-1, p, p+1\}$, and since $T \neq N$, $\ell \geq 1$.

For any $T \subset N$ with $T \neq N, T \neq \emptyset$, and $0 < t < n$,

$$\begin{aligned}&\Phi^S(\vec{y}_T, \vec{x}_{N \setminus T})\\ &\quad = (1+\alpha)(n - t - \ell) + \sum_{j=1}^{p} (t_j - 1 + |S \cap T_j|)\\ &\quad < n - 1 + k \qquad (17)\\ &\quad = \Phi^S(\vec{y})\end{aligned}$$

where (17) is straightforward to verify. ■The proofs of parts (c) and (d) follow in a similar manner to parts (a) and (b), by using the potential function Φ^S for stochastic stability analysis.

C. Resistance trees for stochastic stability analysis

When graphical coordination game G is influenced by a mobile adversary, it is no longer a potential game; resistance tree tools defined in this section enable us to determine stochastically stable states.

The Markov process, P_β, defined by log-linear learning dynamics over a normal form game is a *regular perturbation* of a best response process. In particular, log-linear learning is a regular perturbation of the best response process defined in Appendix A, where the size of the perturbation is parameterized by $\varepsilon = e^{-\beta}$. The following definitions and analysis techniques are taken from [20].

Definition 1 ([20]): A Markov process with transition matrix M_ε defined over state space Ω and parameterized by perturbation $\varepsilon \in (0, a]$ for some $a > 0$ is a *regular perturbation* of the process M_0 if it satisfies:

1) M_ε is aperiodic and irreducible for all $\varepsilon \in (0, a]$.
2) $\lim_{\varepsilon \to 0^+} M_\varepsilon(x, y) \to M(x, y)$ for all $x, y \in \Omega$.
3) If $M_\varepsilon(x, y) > 0$ for some $\varepsilon \in (0, a]$ then there exists $r(x, y)$ such that

$$0 < \lim_{\varepsilon \to 0^+} \frac{M_\varepsilon(x, y)}{\varepsilon^{r(x,y)}} < \infty, \tag{18}$$

where $r(x, y)$ is referred to as the *resistance* of transition $x \to y$.

Let Markov process M_ε be a regular perturbation of process M_0 over state space Ω, where perturbations are parameterized by $\varepsilon \in (0, a]$ for some $a > 0$. Define graph $G = (\Omega, E)$ to be the directed graph with $(x, y) \in E \iff M_\varepsilon(x, y) > 0$ for some $\varepsilon \in (0, a]$. Edge $(x, y) \in E$ is weighted by the resistance $r(x, y)$ defined in (18).

Now let $\Omega_1, \Omega_2, \ldots, \Omega_n$ denote the recurrent classes of process M_0. In graph G, these classes satisfy:

1) For all $x \in \Omega$, there is a zero resistance path from x to Ω_i for some $i \in \{1, 2, \ldots, n\}$.
2) For all $i \in \{1, 2, \ldots, n\}$ and all $x, y \in \Omega_i$ there exists a zero resistance path from x to y and from y to x.
3) For all x, y with $x \in \Omega_i$ for some $i \in \{1, 2, \ldots, n\}$, and $y \notin \Omega_i$, $r(x, y) > 0$.

Define a second directed graph, $\mathcal{G} = (V, \mathcal{E})$, where $V = \{1, 2, \dots, n\}$ are the indices of the n recurrent classes in Ω. For this graph, $(i, j) \in \mathcal{E}$ for all $i, j \in \{1, 2, \dots, n\}$, $i \neq j$. Edge (i, j) is weighted by the resistance of the lowest resistance path starting in Ω_i and ending in Ω_j, i.e.,

$$R(i, j) := \min_{i \in \Omega_i, j \in \Omega_j} \min_{p \in \mathcal{P}(i \to j)} r(p), \tag{19}$$

where $\mathcal{P}(i \to j)$ denotes the set of all simple paths in G beginning at i and ending at j. For path $p = (e_1, e_2, \dots, e_k)$,

$$r(p) := \sum_{\ell=1}^{k} r(e_\ell). \tag{20}$$

Let $\mathcal{T}_i$ be the set of all trees in $\mathcal{G}$ rooted at i, and define

$$\gamma_i := \min_{t \in \mathcal{T}_i} R(t) \tag{21}$$

to be the *stochastic potential* of Ω_i, where the resistance of tree t is the sum of the resistances (in $\mathcal{G}$) of its edges,

$$R(t) := \sum_{e \in t} R(e). \tag{22}$$

We use the following theorem due to [20] in our analysis:

Theorem 7 ([20]): State $x \in \Omega$ is stochastically stable if and only if $x \in \Omega_i$ where

$$\gamma_i = \min_{j \in \{1, 2, \dots, n\}} \gamma_j, \tag{23}$$

i.e., x belongs to a recurrent class which minimizes the stochastic potential. Furthermore, x is strictly stochastically stable if and only if $\Omega_i = \{x\}$ and $\gamma_i < \gamma_j, \quad \forall j \neq i$.

D. Stability in the presence of a mobile, random adversary

The following lemma applies to any graphical coordination game in the presence of a mobile, random adversary with capability $k \leq n - 1$. It states that a mobile random adversary decreases the resistance of transitions when an agent in N changes its action from x to y, but does not change the resistance of transitions in the opposite direction.

Lemma 1: Suppose agents in N update their actions according to log-linear learning in the presence of a mobile, random adversary with capability k, where $1 \leq k \leq n - 1$. Then the resistance of a transition where agent $i \in N$ changes its action from x to y is:

$$\begin{aligned} &r((x, a_{-i}) \to (y, a_{-i})) \\ &\quad = \max\{U_i(x, a_{-i}) - U_i(y, a_{-i}) - 1, 0\} \end{aligned} \tag{24}$$

and the resistance of a transition where agent $i \in N$ changes its action from y to x is:

$$\begin{aligned} &r((y, a_{-i}) \to (x, a_{-i})) \\ &\quad = \max\{U_i(y, a_{-i}) - U_i(x, a_{-i}), 0\}. \end{aligned} \tag{25}$$

Here $U_i : \mathcal{A} \to \mathbb{R}$, defined in (1), is the utility function for agent i in the uninfluenced game, G.

Proof: In the presence of a mobile, random agent,

$$\begin{aligned} &P_\beta\left((x, a_{-i}) \to (y, a_{-i})\right) \\ &\quad = \frac{1}{n}\left(\frac{k}{n} \cdot \frac{\exp(\beta(U_i(y, a_{-i}) + 1))}{\exp(\beta(U_i(y, a_{-i}) + 1)) + \exp(\beta U_i(x, a_{-i}))}\right. \\ &\qquad \left. + \frac{n-k}{n} \cdot \frac{\exp(\beta U_i(y, a_{-i}))}{\exp(\beta U_i(y, a_{-i})) + \exp(\beta U_i(x, a_{-i}))}\right) \end{aligned}$$

Define $P_\varepsilon\left((x, a_{-i}) \to (y, a_{-i})\right)$ by substituting $\varepsilon = e^{-\beta}$ into the above equation. It is straightforward to see that

$$0 < \lim_{\varepsilon \to 0^+} \frac{P_\varepsilon\left((x, a_{-i}) \to (y, a_{-i})\right)}{\varepsilon^{U_i(x, a_{-i}) - U_i(y, a_{-i}) - 1}} < \infty,$$

implying

$$\begin{aligned} &r((x, a_{-i}) \to (y, a_{-i})) \\ &\quad = \max\{U_i(x, a_{-i}) - U_i(y, a_{-i}) - 1, 0\}. \end{aligned}$$

Equation (25) may be similarly verified. ■

Proof of Theorem 5: First we show that, for any $\alpha > 0$, $\vec{x}$ and $\vec{y}$ are the only two recurrent classes of the unperturbed process, P, for the line. Then we show that, for the perturbed process, $R(\vec{x}, \vec{y}) < R(\vec{y}, \vec{x}) \iff \alpha > 1$ and $R(\vec{y}, \vec{x}) < R(\vec{x}, \vec{y}) \iff \alpha < 1$. That is, when $\alpha > 1$ and β is large, the lowest resistance path from $\vec{x}$ to $\vec{y}$ occurs with higher probability than the lowest resistance path from $\vec{y}$ to $\vec{x}$ in P_β, and vice versa when $\alpha < 1$. Combining this with Theorem 7 proves Theorem 5.

Recurrent classes of P for the line: Note that, $P(\vec{x}, a) = 0$ for all $a \in \mathcal{A}, a \neq \vec{x}$, and $P(\vec{y}, a) = 0$ for all $a \in \mathcal{A}, a \neq \vec{y}$, implying $\vec{x}$ and $\vec{y}$ are recurrent. To show that no other state is recurrent, we will show that, for any $a \in \mathcal{A} \setminus \{\vec{x}, \vec{y}\}$, there is a sequence of positive probability transitions in P leading from a to $\vec{x}$.

Let $a \in \mathcal{A}$ with $a \neq \vec{x}, \vec{y}$. Without loss of generality, choose $i, i+1$ such that $a_i = y$ and $a_{i+1} = x$. Denote $(a_i, a_{-i}) = a$, and note that:

$$P((y, a_{-i}) \to (x, a_{-i})) = \frac{1}{n} \cdot \frac{n-k}{n} > 0 \tag{26}$$

for any $k \leq n - 1$ and $\alpha > 0$. Since (26) holds for any $a \neq \vec{x}, \vec{y}$, we can construct a sequence of at most $n - 1$ positive probability transitions leading to joint action $\vec{x}$. Therefore a cannot be recurrent in P.

Resistance between recurrent classes $\vec{x}$ and $\vec{y}$: We will show that for all $1 \leq k \leq n - 1$,

$$R(\vec{y}, \vec{x}) = 1, \quad \forall \alpha > 0, \tag{27}$$

$$R(\vec{x}, \vec{y}) \geq \alpha, \quad \forall \alpha > 0, \tag{28}$$

$$\text{and} \quad R(\vec{x}, \vec{y}) = \alpha, \quad \forall \alpha \leq 1. \tag{29}$$

For (27), we have $r(\vec{y}, (x, y, \dots, y)) = 1$, and $r(\vec{y}, a) \geq 1$ for any $a \neq \vec{y}$, implying that $R(\vec{y}, \vec{x}) \geq 1$. Then, since

$$r\left((\vec{x}_{[t]}, \vec{y}_{[t+1,n]}), (\vec{x}_{[t+1]}, \vec{y}_{[t+2,n]})\right) = 0,$$

for any $1 \leq t \leq n - 1$, and

$$r\left((\vec{x}_{[n-1]}, \vec{y}_{[n,n]}), \vec{x}\right) = 0,$$

the path $\vec{y} \to (x, y, \dots, y) \to (x, x, y, \dots, y) \to \cdots \to \vec{y}$ has resistance 1. Since we know $R(\vec{y}, \vec{x}) \geq 1$, this implies that $R(\vec{y}, \vec{x}) = 1$.

Now, for (28), since $r(\vec{x}, a) \geq \alpha$ for any $a \neq \vec{x}$, this implies $R(\vec{x}, \vec{y}) \geq \alpha$. In particular $r(\vec{x} \to (y, x, \ldots, x)) = \alpha$. When $\alpha < 1$,

$$r\left((\vec{y}_{[t]}, \vec{x}_{[t+1,n]}), (\vec{y}_{[t+1]}, \vec{x}_{[t+2,n]})\right) = 0$$

for any $1 \leq t \leq n-1$, and

$$r\left((\vec{y}_{[n-1]}, \vec{x}_{[n,n]}), \vec{y}\right) = 0,$$

implying that the path $\vec{x} \to (y, x, \ldots, x) \to (y, y, \ldots, x) \to \cdots \to \vec{y}$ has resistance α when $\alpha \leq 1$. Hence $R(\vec{x}, \vec{y}) = \alpha$.

Combining (27) - (29) with Theorem 7 establishes Theorem 5. ■

E. Stability in the presence of an intelligent, mobile agent

Define P_β^μ to be the Markov process associated with log-linear learning in the presence of a mobile, intelligent adversary using policy μ.

Proof of Theorem 6 part (a): Let $G = (N, E)$ be the line, influenced by a mobile, intelligent adversary with capability $k = 1$. For any policy $\mu : \mathcal{A} \to \mathcal{S} = N$, if $\alpha \neq 1$, only $\vec{x}$ and $\vec{y}$ are recurrent in the unperturbed process, P^μ. This can be shown via an argument similar to the one used in the proof of Theorem 5.

Define $\mu^\star$ as in (11). We will show that, (1) in $P_\beta^{\mu^\star}$, $\vec{x}$ is stochastically stable if and only if $\alpha > 1$, and $\vec{y}$ is stochastically stable if and only if $\alpha < 1$, and (2) $\mu^\star$ is optimal, i.e., if $\alpha = 1$, $\vec{x}$ is stochastically stable for any $\mu \in M_1$, and if $\alpha > 1$, $\vec{x}$ is strictly stochastically stable for any $\mu \in M_1$.

For policy $\mu \in M_1$, let $r^\mu(a, a')$ denote the single transition resistance from a to $a' \in \mathcal{A}$ in P_β^μ, and let $R^\mu(a, a')$, denote the resistance of the lowest resistance path from a to $a' \in \mathcal{A}$.

For any $\mu \in M_1$, we have $r^\mu(\vec{x}, a) \geq \alpha, \ \forall a \in \mathcal{A}, a \neq \vec{x}$, and $r^\mu(\vec{y}, a) \geq 1, \quad \forall a \in \mathcal{A}, a \neq \vec{y}$. Therefore

$$R^\mu(\vec{x} \to \vec{y}) \geq \alpha, \text{ and } R^\mu(\vec{y}, \vec{x}) \geq 1. \tag{30}$$

If $\alpha < 1$, the path $\vec{x} \to (y, x, \ldots, x) \to (y, y, x, \ldots, x) \to \cdots \to \vec{y}$ in $P_\beta^{\mu^\star}$ has total resistance α. Equation (30) implies that $R^{\mu^\star}(\vec{x}, \vec{y}) = \alpha < 1 \leq R^{\mu^\star}(\vec{y}, \vec{x})$, so by Theorem 7, $\vec{y}$ is strictly stochastically stable in $P^{\mu^\star}$.

If $\alpha = 1$, it is straightforward to show that both $\vec{x}$ and $\vec{y}$ are stochastically stable in $P_\beta^{\mu^\star}$. Moreover, for any $\mu \in \mathcal{M}$, either the resistance of path

$$\vec{y} \to (x, y, \ldots, y) \to (x, x, y, \ldots y) \to \cdots \to \vec{x}$$

or the resistance of path

$$\vec{y} \to (y, \ldots, y, x) \to (y \ldots, y, x, x) \to \cdots \to \vec{x}$$

is 1, and hence it is impossible to find a policy with $R^\mu(\vec{x}, \vec{y}) < R^\mu(\vec{y}, \vec{x})$.

If $\alpha > 1$, similar arguments show that $R^\mu(\vec{y}, \vec{x}) = 1$ for any $\mu \in M_k$. Combining this with (30) implies that $\vec{x}$ is stochastically stable for any P_β^μ, $\mu \in \mathcal{M}$. ■

Proof of (b): Again let $G = (N, E)$ be the line, and suppose the adversary has capability k with $2 \leq k \leq n-1$. We will show that, for a policy $\mu^\star$ which satisfies Conditions 1 - 3 of Theorem 6, $\vec{x}$ is strictly stochastically stable in $P^{\mu^\star}$ if and only if $\alpha > \frac{n}{n-1}$, and $\vec{y}$ is strictly stochastically stable if and only if $\alpha < \frac{n}{n-1}$. Since this is the same bound on α when we have an adversary with capability n, from Theorem 4 part (a), this also proves that policy $\mu^\star$ is optimal, i.e., no other policy can stabilize a state $a \in \mathcal{A}$ with $a_i = \vec{y}$ for some $i \in N$ when $\alpha > \frac{n}{n-1}$.

First note that only $\vec{y}$ is recurrent in $P^{\mu^\star}$ when $\alpha \leq 1$, and hence $\vec{y}$ is strictly stochastically stable in $P_\beta^{\mu^\star}$.

Now assume $\alpha > 1$. Again, it is straightforward to verify that only $\vec{x}$ and $\vec{y}$ are recurrent in $P^{\mu^\star}$. Note that $r(\vec{x} \to a) \geq \alpha, \forall a \neq \vec{x}$, and $r(\vec{y} \to a) = 2, \forall a \neq \vec{y}$. Moreover, the path $\vec{x} \to (y, x, \ldots, x) \to (y, y, x, \ldots, x) \to \cdots \to \vec{y}$ has total resistance $\alpha + (n-2)(\alpha - 1)$ in $P_\beta^{\mu^\star}$.

It is straightforward to verify that this is the least resistance path from $\vec{x}$ to $\vec{y}$ in $P_\beta^{\mu^\star}$, implying $R(\vec{x}, \vec{y}) = \alpha + (n-2)(\alpha - 1)$. The path $\vec{y} \to (x, y, \ldots, y) \to (x, x, y, \ldots, y) \to \cdots \to \vec{x}$ has resistance 2; hence $R(\vec{y} \to \vec{x}) = 2$. ■