Game Couplings: Learning Dynamics and Applications

Maria-Florina Balcan NINAMF@CC.GATECH.EDU
College of Computing, Georgia Institute of Technology, Atlanta, GA 30332

Florin Constantin* FLORINC@POST.HARVARD.EDU

Georgios Piliouras GEORGIOS@GATECH.EDU
School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332

Jeff S. Shamma SHAMMA@GATECH.EDU
School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332

Abstract

Games encode complex strategic interactions involving numerous agents. Specific game properties, such as the existence of a potential function (Monderer & Shapley, 1996) or (λ, μ)-smoothness (Roughgarden, 2009), allow for making predictions about the resulting agent behavior. In this paper, we examine under which conditions the coupling between two (or more) games that exhibit such an advantageous property leads to a combined game that preserves this property. To study this question formally, we introduce a novel game-theoretic construct that we call game-coupling.

Game coupling intuitively allows us to stitch together the payoff structures of two or more games into a new game. We establish sufficient and necessary conditions for the coupling of two potential games to result into a new potential game. Similar questions are explored for the case of weakly acyclic games. Furthermore, we present settings that allow for particularly advantageous couplings that can enhance desirable properties of the original games, such as convergence of best response dynamics and low price of anarchy. Furthermore, we extend the price of anarchy framework in this setting, to account both for the social welfare within each subgame as well as that of the coupled game. Such concerns give rise to a new notion of equilibrium, as well as a new learning paradigm. We provide welfare guarantees for both individual subsystems as well as for the global system, using generalizations of the (λ, μ)-smoothness framework.

1. Introduction

Many large networked systems, either by evolution or by design, exhibit subsystems of increased internal homogeneity. The prototypical example in the field of technological networks is that of the Internet. The internet consists of multiple independently operated autonomous systems (AS), each implementing a clearly defined routing policy. It has no centralized governance in either technological implementation or policies for usage. As a result, each subnetwork sets its own standards. The coupling of these individual subnetworks is achieved via the centralized maintenance of a few principal name spaces and the standardization of the core protocols (IPv4 and IPv6).

*. Most of this work was completed while the author was with the College of Computing, Georgia Institute of Technology, Atlanta, GA 30332

Hierarchically organized systems also exist in socioeconomic environments. The European Union is a political and economic union of a number of European countries. The member states are internally regulated by locally elected governments with highly diverse economies. The European Union on the other hand enables a closer coupling of these economies via a monetary union and a number of policies that enable the free movement of people, goods and capital.

A unifying characteristic of the above systems is that they consist of pockets of internally homogeneous subsystems. The individual subsystems are endowed with different abilities and local (possibly conflicting) goals, and may have little or no knowledge of the implementation details of other subsystems. Such systems are ubiquitous both in technological as well as socioeconomic networks, however, very little is formally known about their properties and global dynamics.

The heterogeneity of these environments have given risen to much speculation about the long term stability of such environments (Hall, Anderson, Clayton, Ouzounis, & Trimintzios, 2001). The hope, of course, is that these coupled systems do not come at a significant cost for the involved parties. Ideally, we would like to be able to argue that the dynamics of the emerging system do not merely lead to stable states but furthermore that the emergent system is, in some sense, more efficient that the sum of its parts.

We investigate these questions by introducing a novel game-theoretic construct that we call game-coupling. Game coupling intuitively allows us to stitch together the payoff structures of two or more games into a new game. Specifically, we consider populations with J disjoint groups $N^1 \dots N^J$ of agents such that agents in group N^j take part in a game with a desirable property $\mathcal{P}$ for any fixed behavior of all agents from other groups. We define such settings as couplings of $\mathcal{P}$-games. We examine the effect of different classes of properties $\mathcal{P}$ on the performance of each group. We also look for necessary or sufficient conditions for couplings to preserve (to some degree) $\mathcal{P}$.

Ever since the seminal work of Monderer and Shapley (Monderer & Shapley, 1996) the study of dynamics in economic environments has been closely connected with a structural property of games known as potential (and generalizations thereof). Informally, a potential function maps states of the games to the real line, in such a way such that it captures at any state the deviation incentives of all agents. Local optima of the potential correspond to pure Nash equilibria and several learning dynamics are known to converge to them (Monderer & Shapley, 1996; Kleinberg, Piliouras, & Tardos, 2009; Marden, Arslan, & Shamma, 2007). Although the set of potential games is sparse within the set of all games, the existence of a potential function is a highly desired property, and a lot of recent work has been focused on design paradigms that infuse such properties in distributed systems (Marden, Arslan, & Shamma, 2009a; Li & Marden, 2011).

The above approaches, however, are based on the assumptions that there exists a single designer for the whole system or that the individual subsystems exist in isolation from each other. Focusing on the reality of coupled systems, we present conditions for couplings to preserve the existence of potential functions (and generalizations thereof). Our necessary and sufficient condition for preserving an exact potential leverages a condition in (Monderer & Shapley, 1996). We also provide conditions for game couplings to preserve the existence of a weak potential (equivalently, weak acyclicity), a more general notion guaranteeing weaker convergence than an exact potential.

We continue by presenting distributed settings where coupling between heterogeneous systems is provably not merely sufficient but also necessary to achieve convergence to globally efficient outcomes. Aggregation games (Mol, Vattani, & Voulgaris, 2011) model populations aiming for high internal connectivity given an underlying (social) network. Agent type can either be of the form of opinion-leaders, favoring being a node with high degree, or of opinion-follower favoring being in the neighborhood of many other agents. Prior work has established that homogeneous groups can exhibit highly inefficient equilibria, whereas well mixed populations do not. We complete the understanding of these settings by solving in the affirmative an open question regarding the convergence of Nash dynamics of heterogeneous populations in aggregation games.

The efficiency of equilibrium states is captured by the notion of price of anarchy. A game is said to have a low price of anarchy (Koutsoupias & Papadimitriou, 1999) if the performance of the worst (Nash) equilibrium state is comparable to that of the socially optimal state. Recent work (Roughgarden, 2009) has established a structural property known as (λ, μ)-smoothness that provably implies tight bounds on the inefficiency of standard equilibrium concepts in several game classes (Nadav & Roughgarden, 2010; Roughgarden, 2009; Roughgarden & Schoppmann, 2011).

In the case of coupled subsystems the notion of efficiency is more intricate as both the collective efficiency of the individual subgroups as well as that of the merged system are of interest. For example, in the case of socioeconomic systems, such as the European Union, a successful coupling would improve the overall system welfare without imposing significant penalties to the welfare of any individual state. We show how to achieve tight guarantees on group performance when each subgame N^j exhibits a local version of the (λ^j, μ^j)-smoothness. These performance guarantees carry over to a wide set of equilibrium notions as well as for no-regret learning algorithms. Using LP-duality arguments, we identify new equilibrium notions, which are applicable to any coupled systems, and for which these bounds extend automatically to individual groups.

The possibility of subnetwork designers with broadcasting power to their respective subgroups also opens up the possibility of novel learning paradigms. We introduce a novel learning framework modeling interactions of competing groups/institutions (e.g. Internet autonomous systems). In our framework each group has a center that provides public advertisement (Balcan, Blum, & Mansour, 2009, 2010; Balcan, Krehbiel, Piliouras, & Shin, 2012) of (possibly different) strategies to each agent in the group. We analyze centers whose advertised behaviors exhibit vanishing average regret with hindsight. This is a rather natural benchmark, since several simple learning algorithms offer such guarantees (Cesa-Bianchi & Lugosi, 2006). On the side of the agents, we make a similarly weak assumption. We assume that the average performance of each agent will eventually be roughly as high as that of his advertised strategy. We provide welfare guarantees for each group in this framework as well as global guarantees for the system as a whole.

2. Related work

Structural properties of games often follow from analysis of their subgames. Sandholm (Sandholm, 2010) considers subgames defined by all subsets of agents (as opposed to a specific partition as we do here) and shows that a game has an exact potential if and only if all active agents in a sub-game have identical utility functions. Fabrikant et al. (Fabrikant, Jaggard, & Schapira, 2013) show that the uniqueness of a Nash equilibrium in any sub-game is sufficient, but not necessary, for a game to be weakly acyclic. Candogan et al. (Candogan, Menache, Ozdaglar, & Parrilo, 2011) introduce a decomposition of any game via its payoffs matrix rather than subsets of players. They show that any game can be projected orthogonally onto three subspaces, corresponding to potential games, harmonic games and non-strategic games.

Monderer (Monderer, 2007) defines the classes of J-potential games and J-congestion games for $J \in \mathbb{N}$ and shows they are isomorphic. In a J-congestion game, each agent's' delay functions can belong in one of J classes. The case $J = 1$ is treated in Monderer and Shapley's seminal paper (Monderer & Shapley, 1996).

Moving towards the goal of improving system security and robustness for distributed services, several models have been developed that allow for different classes of buggy, selfish, or even malicious nodes. Byzantine-Altruistic-Rational (BAR) (Aiyer, Alvisi, Clement, Dahlin, Martin, & Porth, 2005) behavioral models are used to model such environments and the main goals of this research agenda is to reduce the cost of such Byzantine fault tolerant systems as well as to broaden their applicability. Another model of such malicious behavior is due to Gairing (Gairing, 2008), who studies atomic congestion games with players that are either purely malicious or purely selfish (with probabilities p and $1 - p$). Even in singleton linear congestion games, the resulting games may not have a pure Bayesian Nash equilibrium and deciding the existence of such an equilibrium is NP-complete. In our work, we also introduce a model of malicious agents and show how the question of the existence of Nash equilibria can be addressed via our game coupling tools.

The view of a population divided into groups is often adopted in distributed adaptive control, where a center can only control a local group of agents, e.g. in the collective intelligence (Tumer & Wolpert, 2004) and probability collectives (Wolpert & Bieniawski, 2004) frameworks.

One of the main advantageous structural properties of a distributed environment is the existence of stable states with high performance. A standard measure of such distributed inefficiency is the price of anarchy (PoA) (Koutsoupias & Papadimitriou, 1999), defined as the ratio of the social cost of the worst pure Nash equilibrium to the optimum. Recent work has shown that most positive results in this area can be derived via a canonical property of games known as (λ, μ)-smoothness (Roughgarden, 2009; Nadav & Roughgarden, 2010). This analysis allows to extent such guarantees in a black-box manner to much wider equilibrium sets (e.g. correlated equilibria, no-regret learning (Blum, Hajiaghayi, Ligett, & Roth, 2008)). In the setting of coupled systems we will analyze both the collective performance of the system as well as the performance of each subsystem.

We also study and provide performance guarantees for learning dynamics both from the perspective of each group as well as from a global perspective. Work in the area of learning dynamics has provided strong guarantees in the terms of convergence and performance of learning dynamics (*e.g.* congestion games (Kleinberg et al., 2009; Kleinberg, Piliouras, & Tardos, 2011)), in some games matching the theoretically optimal (centralized) guarantees both in terms of speed as well as approximation guarantees (Piliouras, Valla, & Végh, 2012). Recent work has started incorporating learning in structured/hierarchical environments, *e.g.* language games with coevolutionary dependence between languages and linguistic communities (Fox, Piliouras, & Shamma, 2012), oligopolistic markets with evolving coalition structures (Nadav & Piliouras, 2010; Immorlica, Markakis, & Piliouras, 2010). Interestingly, in some cases persistent disequilibrium behavior can provide performance guarantees that are significantly stronger than that of the best equilibrium outcome (Kleinberg, Ligett, Piliouras, & Tardos, 2011; Ligett & Piliouras, 2011).

The structure of the paper is as follows: In section 3 we start off by presenting some basic facts and definitions about games, price of anarchy and potential games. We continue in section 4 by providing the formal definition of game couplings. Section 5 presents conditions under which the coupling of two exact potential games gives rise to another exact potential game. This question is also addressed in terms of the coupling of weak potential games. Furthermore, in section 6 we study the case of aggregation games and resolve the open question about convergence of dynamics in the case of heterogeneous populations. Finally, in section 7 we argue about the efficiency of general coupled systems. This discussion gives rise to novel notions of equilibria and learning paradigms whose performance we analyze utilizing a generalization of the (λ, μ)-smoothness framework.

3. Preliminaries

We model interactions within a population as games G with simultaneous moves. For a game G we denote by $\{1, \dots, n\}$ the agents, by Σ_i agent i's set of (pure) strategies, and by $\Delta(\Sigma)$ all distributions over outcomes $\Sigma = \times_{i=1}^{n} \Sigma_i$. Any agent i aims to minimize his *cost*[1] $C_i(\sigma_1, \dots, \sigma_n)$ with $\sigma_h \in \Sigma_h$ chosen by agent $h = 1 \dots n$. Individual costs are aggregated by the *social cost* $C(\sigma) = \sum_i C_i(\sigma)$. We denote by $v_{-i} = (v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_r)$ a vector $v = (v_1, \dots, v_r)$ without its ith entry, for $1 \le i \le r$.

A pure Nash equilibrium is a stable outcome in which for any agent, its strategy minimizes its cost given others' strategies. Formally, a strategy vector $(\sigma_1, \dots, \sigma_n) \in \Sigma$ is a *pure Nash equilibrium* (PNE) if any agent i minimizes its cost by playing σ_i, i.e. $C_i(\sigma_i, \sigma_{-i}) \le C_i(\sigma_i', \sigma_{-i}), \forall i, \forall \sigma_i' \in \Sigma_i$.

PNE may not be efficient from a social perspective. A standard measure of distributed inefficiency is the *price of anarchy* (PoA) (Koutsoupias & Papadimitriou, 1999), defined as the ratio of the social cost of the worst PNE to the optimum:

$$\text{PoA} = \frac{\max_{\sigma \in \text{PNE}} C(\sigma)}{\min_{\sigma^* \in \Sigma} C(\sigma^*)}$$

Potential games: Game G exhibits an *exact potential* (Monderer & Shapley, 1996) $\Phi : \Sigma \to \mathbb{R}$ if $C_i(\sigma_i, \sigma_{-i}) - C_i(\sigma_i', \sigma_{-i}) = \Phi(\sigma_i, \sigma_{-i}) - \Phi(\sigma_i', \sigma_{-i})$ $\forall i$, $\forall \sigma_i, \sigma_i' \in \Sigma_i, \sigma_{-i} \in \Sigma_{-i}$. Game G has an *ordinal potential* $\Phi : \Sigma \to \mathbb{R}$ if $C_i(\sigma_i, \sigma_{-i}) < C_i(\sigma_i', \sigma_{-i}) \Leftrightarrow \Phi(\sigma_i, \sigma_{-i}) < \Phi(\sigma_i', \sigma_{-i})$

G has a *weak potential* $\Phi : \Sigma \to \mathbb{R}$ (Marden et al., 2009a) if and only if at any strategy vector $(\sigma_1, \dots, \sigma_n)$ that is not a (pure) Nash equilibrium, there exists an agent i that can simultaneously lower both his cost and Φ by switching to some strategy σ_i': $C_i(\sigma_i, \sigma_{-i}) > C_i(\sigma_i', \sigma_{-i})$ and $\Phi(\sigma_i, \sigma_{-i}) > \Phi(\sigma_i', \sigma_{-i})$.

If G has an exact potential (respectively ordinal potential, or weak potential) then G is called an *exact potential* (respectively *ordinal potential*, or *weakly acyclic*) game.

1. In Section 6 we use utility maximization instead of cost minimization.

The relations between the different potential notions have as follows:

$$\text{exact potentials} \subseteq \text{ordinal potentials} \subseteq \text{weak potentials}$$

An ordinal potential Φ is also a weak one, since it suffices that *one* agent increases Φ upon improving his utility. The existence of an ordinal potential is equivalent (Monderer & Shapley, 1996) to the convergence of *better response dynamics*, i.e. asynchronous updates by each agent to a better strategy given others' current strategies. Any weakly acyclic game has at least one PNE, for example the global optimum of the weak potential.

4. Game Couplings

In the preliminaries section, we went over a number of beneficial system level properties such as a low price of anarchy and the existence of a pure Nash equilibrium. We denote any game that exhibits such a property $\mathcal{P}$ as a $\mathcal{P}$-*game*. Our goal is to understand under which conditions such a property $\mathcal{P}$ extends from individual subgames to the coupled system. To this goal we introduce and study the concept of game *coupling*.

In terms of notation, we will be using subscripts to refer to agents, while superscripts refer to sets of agents, i.e., groups. In other words, we partition the set N of agents to subgroups N^j where $j \in \{1, \dots J\}$. Let $\Sigma^j = \times_{i \in N^j} \Sigma_i$ and $\Sigma^{-j} = \times_{i \notin N^j} \Sigma_i$. For any j and any fixed vector $\sigma^{-j} \in \Sigma^{-j}$ of agents in N^{-j}, we denote as $G|_{N^{-j} \leftarrow \sigma^{-j}}$ the sub-game of G (played by N^j) that arises when the strategies of the agents outside group j are fixed according to σ^{-j}. In other words, $G|_{N^{-j} \leftarrow \sigma^{-j}}$ is a game of $|N^j|$ agents where each agent i has a set Σ_i of pure strategies and her cost at any outcome $\sigma^j_0 \in \Sigma^{-j}$ is equal to $C_i(\sigma^j_0) = C_i(\sigma^{-j}, \sigma^j_0)$. Next, we formally define a game coupling between $\mathcal{P}$-games:

Definition 1. *G is a $(N^1, \dots, N^J)$*-coupling of $\mathcal{P}$-games *if*

- *$N^1, \dots, N^J$ are* groups *partitioning $\{1 \dots n\}$, i.e. $N^j \cap N^{j'} = \emptyset, \forall 1 \le j < j' \le J$, and $N^1 \cup \dots \cup N^J = \{1 \dots n\}$. The groups are fixed: no agent can choose its group.*
- *For any j and any fixed vector $\sigma^{-j} \in \Sigma^{-j}$ of agents in N^{-j}, the sub-game $G|_{N^{-j} \leftarrow \sigma^{-j}}$ (played by N^j) has property $\mathcal{P}$.*

Coupling Example. In a load balancing game, each agent (job) chooses a machine. Each machine e has a specific cost function c_e which depends only on e's load, i.e., number of jobs on it. PNE always exist in such games (Monderer & Shapley, 1996). Jobs are rarely this homogeneous; instead, there are often groups of jobs, e.g., computation-intensive or memory-intensive. In this case, each machine has a cost function c^j_e for each type j of jobs. When fixing the strategies of other jobs, the game experienced by jobs of any type j is a standard load balancing game and hence admits a PNE. Thus, a load balancing game with heterogeneous jobs is a coupling of games that admit PNE. The question is then, when does the coupled (global/heterogeneous) game also admit a PNE?

5. Couplings and potentials

In this section we consider the effect of game couplings on two closely connected $\mathcal{P}$-properties, the existence of an exact potential and weak acyclicity.

5.1 Exact potential games

The existence of an exact potential is a highly desirable game property, since it implies, among others, the convergence of better response dynamics. In multi-agent systems, exact potentials arise for example in the "wonderful life utility" scheme (Tumer & Wolpert, 2004), by which a planner can ensure that individual agents will act in accordance to the common welfare.

We provide sufficient and necessary conditions for a coupling of exact potential games to also have an exact potential. Our analysis leverages the following well-known characterization of (exact) optential games (Monderer & Shapley, 1996).

Lemma 1. *(Monderer & Shapley, 1996) Game G has an exact potential if and only if the changes in payoff for the deviating agents along any closed path of length 4 sum up to zero. Equivalently, for any agents i, k, any strategies σ_i, σ_i' of i and σ_k, σ_k' of k and for any σ_{-ik} of the other agents we have $d_{\sigma_i\sigma_i'\sigma_k\sigma_k'}(\sigma_{-ik}) = 0$ where*

$$d_{\sigma_i\sigma_i'\sigma_k\sigma_k'}(\sigma_{-ik}) := \Delta_{\sigma_i\sigma_k}(\sigma_{-ik}) - \Delta_{\sigma_i\sigma_k'}(\sigma_{-ik}) - \Delta_{\sigma_i'\sigma_k}(\sigma_{-ik}) + \Delta_{\sigma_i'\sigma_k'}(\sigma_{-ik}) \tag{1}$$

and $\Delta_{\hat\sigma_i\hat\sigma_k}(\sigma_{-ik}) := C_i(\hat\sigma_i, \hat\sigma_k, \sigma_{-ik}) - C_k(\hat\sigma_i, \hat\sigma_k, \sigma_{-ik}), \forall$ strategies $\hat\sigma_i, \hat\sigma_k$.

The forward direction is trivial, since for any closed path in any potential game the changes in payoff correspond to changes in potential and the total change of the potential along any closed path is equal to zero. However, this lemma further states that checking all 4-cycles suffices to verify that the game has indeed an exact potential.

Our condition on couplings abstracts away the individuals in a group and considers an auxiliary game using group potentials.

Theorem 1. *Let G be a (N^1, N^2)-coupling of exact potential games: $\forall \sigma^j \in \Sigma^j$, the subgame $G|_{N^j \leftarrow \sigma^j}$ induced by group N^j playing strategy vector σ^j has exact potential $\Phi_{\sigma^j}(\cdot)$, for $j = 1, 2$. Define a game Γ with agents $\{1, 2\}$ in which agent j's strategy space is Σ^j and the utilities from playing (σ^1, σ^2) are $(\Phi_{\sigma^2}(\sigma^1), \Phi_{\sigma^1}(\sigma^2))$, then G is an exact potential game if and only if Γ is an exact potential game.*

Proof. " $\Rightarrow$ " Assume first that G has an exact potential $\Phi(\cdot, \cdot)$. We show that Γ is an exact potential game as well.

We note that $\hat\Phi^2_{\sigma^1}(\sigma^2) = \Phi(\sigma^1, \sigma^2)$ is a potential for the game $G|_{N^1 \leftarrow \sigma^1}$ as Φ is a potential for G. Defining analogously exact potentials $\hat\Phi^2_{\bar\sigma^1}(\cdot)$ of $G|_{N^1 \leftarrow \bar\sigma^1}$, $\hat\Phi^1_{\sigma^2}(\cdot)$ of $G|_{N^2 \leftarrow \sigma^2}$ and $\hat\Phi^1_{\bar\sigma^2}(\cdot)$ of $G|_{N^2 \leftarrow \bar\sigma^2}$, we get that $d_{\sigma^1\bar\sigma^1\sigma^2\bar\sigma^2}(\hat\Phi) = 0$ clearly holds for the game Γ induced by $\hat\Phi$ (in particular this game is one of identical interest, i.e. any Δ as defined in Eq. (1) equals 0). Let $\sigma^1, \bar\sigma^1 \in \Sigma^1$ and $\sigma^2, \bar\sigma^2 \in \Sigma^2$ and exact potentials $\Psi^2_{\sigma^1}(\cdot)$ of $G|_{N^1 \leftarrow \sigma^1}$, $\Psi^2_{\bar\sigma^1}(\cdot)$ of $G|_{N^1 \leftarrow \bar\sigma^1}$, $\Psi^1_{\sigma^2}(\cdot)$ of $G|_{N^2 \leftarrow \sigma^2}$ and $\Psi^1_{\bar\sigma^2}(\cdot)$ of $G|_{N^2 \leftarrow \bar\sigma^2}$. For any other potential function $\Psi^2_{\sigma^1}(\cdot)$ of $G|_{N^1 \leftarrow \sigma^1}$ there exists (Monderer & Shapley, 1996) $d_{\sigma^1} \in \mathbb{R}$ with

$$\Psi^2_{\sigma^1}(\sigma^2) = \hat\Phi^2_{\sigma^1}(\sigma^2) + d_{\sigma^1} = \Phi(\sigma^1, \sigma^2) + d_{\sigma^1}, \forall \sigma^2 \tag{2}$$

From analogous conditions to Eq. (2) we easily get $d_{\sigma^1\bar\sigma^1\sigma^2\bar\sigma^2}(\Psi) = d_{\sigma^1\bar\sigma^1\sigma^2\bar\sigma^2}(\hat\Phi) = 0$.

"$\Leftarrow$" We establish Lemma 1's condition (Eq. (1)) for G given that it holds for Γ.

$$\begin{aligned}
0 &= (\Phi^2_{\sigma^1}(\sigma^2) - \Phi^1_{\sigma^2}(\sigma^1)) - (\Phi^2_{\sigma^1}(\bar{\sigma}^2) - \Phi^1_{\bar{\sigma}^2}(\sigma^1)) \\
&\quad - (\Phi^2_{\bar{\sigma}^1}(\sigma^2) - \Phi^1_{\sigma^2}(\bar{\sigma}^1)) + (\Phi^2_{\bar{\sigma}^1}(\bar{\sigma}^2) - \Phi^1_{\bar{\sigma}^2}(\bar{\sigma}^1)) \Leftrightarrow \qquad (3) \\
0 &= (\Phi^2_{\sigma^1}(\sigma^2) - \Phi^2_{\sigma^1}(\bar{\sigma}^2)) - (\Phi^1_{\sigma^2}(\sigma^1) - \Phi^1_{\sigma^2}(\bar{\sigma}^1)) \\
&\quad + (\Phi^1_{\bar{\sigma}^2}(\sigma^1) - \Phi^1_{\bar{\sigma}^2}(\bar{\sigma}^1)) - (\Phi^2_{\bar{\sigma}^1}(\sigma^2) - \Phi^2_{\bar{\sigma}^1}(\bar{\sigma}^2)) \qquad (4)
\end{aligned}$$

For G, consider two agents i and k, fixed (pure) strategies σ_{-ik} of the other agents (that we will drop for ease of notation), and (pure) strategies $\sigma_i, \bar{\sigma}_i$ of i and $\sigma_k, \bar{\sigma}_k$ of k. If i and k are part of the sub-game N^l (where $l \in \{1, 2\}$) then Eq. (1) holds since N^l is a potential game. By symmetry, we can focus on the case $i \in N^1, k \in N^2$.

Let $\sigma^1 = (\sigma_i, \sigma^1_{-i}), \bar{\sigma}^1 = (\bar{\sigma}_i, \sigma^1_{-i}), \sigma^2 = (\sigma_k, \sigma^2_{-k})$ and $\bar{\sigma}^2 = (\bar{\sigma}_k, \sigma^2_{-k})$. Using the exact potential property in each subgame, we can rewrite Eq. (4) as

$$\begin{aligned}
0 &= (C_k(\sigma_i, \sigma_k) - C_k(\sigma_i, \bar{\sigma}_k)) - (C_i(\sigma_i, \sigma_k) - C_k(\bar{\sigma}_i, \sigma_k)) \\
&\quad + (C_i(\sigma_i, \bar{\sigma}_k) - C_i(\bar{\sigma}_i, \bar{\sigma}_k)) - (C_k(\bar{\sigma}_i, \sigma_k) - C_k(\bar{\sigma}_i, \bar{\sigma}_k))
\end{aligned}$$

which, after rearranging, becomes Eq. (1). □

We note that this result is tight in the following sense: There exists a game G that does not admit an exact potential, with the following property: For any partition of agents (N^1, N^2), there *exist* fixed strategy vectors σ^1, σ^2 such that the game $G|_{N^j \leftarrow \sigma^j}$ has an exact potential. Furthermore, the game Γ induced by potentials of G's sub-games is an exact potential game. An example of such a game is an arbitrary two-agent game, where we pick the unique (non-trivial) partition of the agents in two sets and respectively assign to each of the agents a specific strategy.

5.2 Weakly acyclic games

We further extend the study of the effects of game coupling on the existence of a potential by focusing on the case of weakly acyclic games. Weak acyclicity is defined as the existence of a better-response path from any strategy vector to a PNE. Similarly to the stronger property of the existence of an exact potential, the existence of a weak potential allows for the provable convergence of some families of distributed dynamics (Marden et al., 2009a; Marden, Arslan, & Shamma, 2009b; Young, 1998, 2004).

In the following proposition we present a sufficient condition for the coupling of two weakly potential games to give rise to a coupled system that also exhibits a weak potential.

Proposition 1. *If a coupling G of weakly acyclic games satisfies*

- $G|_{N^1 \leftarrow \sigma^1}$ *has weak potential* $\Phi^2_{\sigma^1}$ *for* $\forall \sigma^1 \in \Sigma^1$,
- $G|_{N^2 \leftarrow \sigma^2}$ *has weak potential* $\Phi^1_{\sigma^2}$ *for* $\forall \sigma^2 \in \Sigma^2$,
- *if* σ^2 *is PNE in* $G|_{N^1 \leftarrow \sigma^1}$ *then any better-response* $\bar{\sigma}^1_i$ *to* σ^1_{-i} *(and* σ^2*) by any* i *in sub-game* $G|_{N^2 \leftarrow \sigma^2}$ *does not increase the weak potential:* $\Phi^2_{\bar{\sigma}^1_i, \sigma^1_{-i}}(\sigma^2) \leq \Phi^2_{\sigma^1}(\sigma^2)$

then G is a weakly acyclic game.

Proof. We will show that G has a weak potential of the form $\Phi(\sigma^1, \sigma^2) = C \cdot \Phi^2_{\sigma^1}(\sigma^2) + \Phi^1_{\sigma^2}(\sigma^1)$ for some large enough $C > 0$. It suffices to choose $C > 0$ such that for all $\sigma^1 \in \Sigma^1, \sigma^2_i, \bar{\sigma}^2_i \in \Sigma^2_i, \sigma^2_{-i} \in \Sigma^2_{-i}$ with $\Phi^2_{\sigma^1}(\sigma^2) - \Phi^2_{\sigma^1}(\bar{\sigma}^2_i, \sigma^2_{-i}) > 0$ it holds that $C \cdot (\Phi^2_{\sigma^1}(\sigma^2) - \Phi^2_{\sigma^1}(\bar{\sigma}^2_i, \sigma^2_{-i})) > |\Phi^1_{(\bar{\sigma}^2_i, \sigma^2_{-i})}(\sigma^1) - \Phi^1_{\sigma^2}(\sigma^1)|$. Such a choice is always possible in any finite game.

Consider a non-PNE strategy vector $\sigma \in \Sigma$: there exists a sub-game $j \in \{1, 2\}$ and an agent $i \in N^j$ that can (strictly) decrease his cost in G by switching to strategy $\bar{\sigma}_i$. If choice $j = 2$ is viable, always choose $j = 2$. In this case, i equally decreases his cost in sub-game $G|_{N^1 \leftarrow \sigma^1}$. The choice of C implies that Φ decreases as well. Otherwise, we have that σ^2 is a PNE in $G|_{N^1 \leftarrow \sigma^1}$ and $j = 1$, i.e. $\exists i \in N^1$ that decreases Φ^1, by weak acyclicity. By assumption, i's switch to $\bar{\sigma}_i$ cannot increase Φ^2 and thus Φ decreases. □

We showcase the power of the Proposition 1 by applying it to a establish the existence of PNE for classes of games with malicious agents.

5.3 Congestion-seeking malice

We apply Proposition 1 to the well-studied setting of *congestion* games (Monderer & Shapley, 1996; Rosenthal, 1973). These games arise in many settings with joint usage of resources and are isomorphic to exact potential games. They are non-cooperative games in which the utility of each player depends only on the player's strategy and the number of other players that either choose the same strategy, or some strategy that "overlaps" with it.

Formally, a *congestion game* is defined by the tuple $(N; E; (\Sigma_i)_{i \in N}; (c_e)_{e \in E})$ where N is the set of *players*, E is a set of *facilities* (also known as *edges* or *bins*), and each player i has a set Σ_i of subsets of E ($\Sigma_i \subseteq 2^E$). Each pure strategy $\sigma_i \in \Sigma_i$ is a set of edges (a *path*), and c_e is a cost (negative utility) function associated with facility e. Given a pure strategy profile $\sigma = (\sigma_1, \sigma_2, \ldots, \sigma_N)$, the cost of player i is given by $C_i(\sigma) = \sum_{e \in \sigma_i} c_e(k_e(\sigma))$, where $k_e(\sigma)$ is the number of players using e in σ. For any strategy profile σ, we define the support of σ, $Supp(\sigma)$, as the set of congested elements in use by some agents, i.e. $Supp(\sigma) = \{e \in E : k_e(\sigma) \geq 1\}$. Congestion games admit the following potential function: $\Phi(\sigma) = \sum_{e \in E} \sum_{j=1}^{k_e(\sigma)} c_e(j)$.

In several classes of congestion games, especially the ones modeling routing applications through transportation networks, agents typically experience higher costs for higher congestion. Such agents are naturally vulnerable to malicious agents that seek to increase the system congestion. Leveraging Proposition 1 with N^2 as the set of malicious congestion-seeking agents, we can establish that a model of malice preserves some of the structure of the original congestion game, unlike other models with a similar scope (Babaioff, Kleinberg, & Papadimitriou, 2009).

Definition 2. *Given a base congestion game $G_{base} = (N; E; (\Sigma_i)_{i \in N}; (c_e)_{e \in E})$ with non-decreasing latency functions c_e and a set of allowable malicious strategies $\Sigma_{mal} \subset 2^E$, we define as (N^1, N^2) malicious congestion-seeking coupling G any game exhibiting the following properties:*

- *For every $\sigma^2 \in \Sigma^2 = \times_{i \in N^2} \Sigma_{mal}$, $G|_{N^2 \leftarrow \sigma^2}$ defines the following congestion subgame $(N_1; E; (\mathcal{S}_i)_{i \in N}; (c^{\sigma^2}_e)_{e \in E})$ with $c^{\sigma^2}_e(x) = c_e(x + k_e(\sigma^2))$ for all $x \in \mathbb{N}$.*

- *For every $\sigma^1 \in \Sigma^1$, $G|_{N^1 \leftarrow \sigma^1}$ defines a "malicious" subgame $G|_{N^1 \rightarrow \sigma^1}$ with the property that for any $i_2 \in N^2$ if $Supp(\sigma) \setminus \sigma_{i_2} \neq \emptyset$ then i_2 can improve his cost by also using some $r \in Supp(\sigma) \setminus \sigma_{i_2}$, i.e. $\sigma_{i_2} := \sigma_{i_2} \cup \{r\}$.*

Next, we will apply Proposition 1 to derive the following characterization of malicious congestion-seeking couplings.

Corollary 1. *Any (N^1, N^2) malicious congestion-seeking coupling with $|N^2| \geq 2$ is a weak acyclic game.*

Proof. Let Φ^1 be a (weak) potential for $G|_{N^2 \leftarrow \sigma^2}$, let m be the total number of resources in the game and

$$\Phi^2(\sigma) = -2m|Supp(\sigma)| - \textstyle\sum_{i \in N^2} |\{e \in \sigma_i : \exists i' \in N^1 \cup N^2 \setminus \{i\} \text{ with } e \in \sigma_{i'}\}|$$

i.e. the negation of $2m$ times the total number of resources used plus, for each N^2 agent, the number of resources he uses together with at least one other agent (in $N^1 \cup N^2$).

We establish that game $G|_{N^1 \leftarrow \sigma^1}$ is weakly acyclic with weak potential Φ^2. Consider a non-NE profile σ. If there exists $r \in Supp(\sigma)$ that is not used by some $i_2 \in N^2$ then, by the malicious subgame definition, i_2 can improve his utility by adding r. Clearly, this implies that Φ^2 decreases.

Suppose now that all resources in $Supp(\sigma)$ are used by all agents in N^2. We claim that an agent $i \in N^2$ can better respond only by using (at least) a new resource $e \notin Supp(\sigma)$ while possibly dropping some resources in σ_i. Indeed, i cannot improve by simply dropping a single resource e_1, i.e. $\sigma_i := \sigma_i \setminus \{e_1\}$: otherwise, were i to play $\sigma_i \setminus \{e_1\}$, he would not benefit from using e_1 despite being the only one in N^2 not to do so, contradicting the definition of the malicious subgames.

We prove by induction on k that i cannot improve by simply dropping $k \leq |\sigma_i|$ resources, i.e. by letting $\sigma_i := \sigma_i \setminus \{e_1, \ldots, e_k\}$. We have just established the base case $k = 1$. For the inductive step, since i is the only agent in N^2 not to use $e_1 \ldots e_k$, by the definition of the malicious subgame there exists some j (assume wlog $j = 1$) such that i improves over $\sigma_i \setminus \{e_1, \ldots, e_k\}$ by $\sigma_i \setminus \{e_2, \ldots, e_k\}$, i.e. using e_j. But then i can profit from σ_i by dropping the $k - 1$ resources $\{e_2, \ldots, e_k\}$, contradicting the inductive hypothesis.

For each new resource e that i starts to use, adding it to $Supp(\sigma)$, Φ^2 decreases by $2m$ from $2m|Supp(\sigma)|$. For each resource e_j dropped, Φ^2 increases by at most 2 – only when exactly one other agent i', with $i' \in N^2$, was using e_j in σ. Since we multiply $|Supp(\sigma)|$ by $2m$, Φ^2 is guaranteed to decrease.

Weak acyclicity now follows from Proposition 1: at any PNE σ in $G|_{N^1 \leftarrow \sigma^1}$, each resource in $Supp(\sigma)$ is used by at least two N^2 agents. Thus Proposition 1's last assumption, that better responses in $G|_{N^2 \leftarrow \sigma^2}$ does not increase Φ^2, holds. □

Applications of malicious congestion-seeking agents: In this section, we briefly discuss settings where the results about the congestion-seeking malicious agents apply.

Load balancing games. There are m machines accepting jobs, each with a distinct latency function. Each player in N^1 owns a job (of load 1) and chooses a machine to execute his job so as to minimize its completion time. On the other hand malicious agents aim to "jam" the machines with spurious jobs in a denial-of-service attack.

For specific instances of this model, Prop. 2's statement can be shown to be tight in the following sense: If we allow only a single malicious agent then we can create a coupled system with no PNE. For example, a game between a single rational agent and a single malicious agent who can choose between two congested elements of linear latencies has the same best response dynamics as a matching pennies games and therefore no PNE exists.

Market-sharing games. In a market-sharing game, there is a set of markets (resources) $r_1, \ldots, r_m$, where market r_j has total available revenue w_j. Each player must choose one market out of a subset available to him and aims to maximize his revenue. Any player joining market r_j has a cost C_j. If n_j players choose market r_j then each one has revenue w_j/n_j and profit $w_j/n_j - C_j$. A market-sharing game is an exact potential game and thus weakly acyclic. Any malicious player can reduce the profit of all other players on a market by jamming it (we can assume that a malicious player's cost is low enough such that jamming is profitable).

We note that a similar result holds for malicious players in facility location games (Vetta, 2002), a game-theoretic distributed counterpart to the classical centralized optimization problem of facility location.

6. Aggregation via heterogeneity

Convergence guarantees for dynamics are made more relevant by quantitative statements about the quality of equilibria (or equilibria that such dynamics can reach)[2]. We will consider the interplay between such issues and coupling of heterogeneous systems in the class of aggregation games and show that heterogeneity is not only sufficient but also necessary condition for combining stability and efficiency in such settings.

Aggregation games (Mol et al., 2011) model populations aiming for high internal connectivity. Specifically, consider an undirected graph $Gr = (\{1, \ldots, N\}, E)$ without self-loops. There exist $n \leq N$ agents. Each must choose a different vertex in $1..N$; denote by H the set of all n agents' vertices: $|H| = n$. Each agent i has a parameter $\beta_i \in [0, 1]$, inducing the utility function[3] u_{β_i} it aims to maximize, where if v_i is i's vertex

$$u_{\beta_i}(v_i, H \setminus \{v_i\}) = E_{v_i,H} + \beta_i E_{v_i,\{1,\ldots,N\}\setminus H} \tag{5}$$

and $E_{I_1,I_2} = |\{e \in E : e = (i_1, i_2), i_1 \in I_1, i_2 \in I_2\}|$ denotes the number of edges between vertex sets I_1 and I_2. We denote by $G(Gr, \beta_1, \ldots, \beta_n)$ the resulting aggregation game G.

Agents with $\beta = 0$ are called *followers* as they maximize $u_0 = E_{v,H}$ i.e. the number of edges to H (the other agents' vertices). In contrast, agents with $\beta = 1$ are called *leaders* because they maximize $u_1 = E_{v,\{1,\ldots,N\}}$ i.e. the degree of v in the hope that other agents, in particular followers, will be drawn to the adjacent vertices. An agent with a general β is called a β-leader; note $u_\beta(\cdot) = \beta u_1(\cdot) + (1 - \beta)u_0(\cdot)$. We call an aggregation game $G(Gr, \beta_1, \ldots, \beta_n)$ *homogeneous* if $\beta_i = \beta, \forall i$.

2. Profiles during learning dynamics, even non-convergent ones, may however be much better than any PNE (Kleinberg et al., 2011).
3. Utilities are more natural than costs for evaluating connectivity.

The *social welfare* (the counterpart of social cost) is the number $E_H := \frac{1}{2}\sum_i E_{v_i,H}$ of internal edges, *for any* $\beta_1,\ldots,\beta_n$. Price of anarchy is defined as:

$$\mathrm{PoA} = \frac{\max_{H^*} E_{H^*}}{\min_{H\,\mathrm{PNE}} E_H}.$$

We now consider issues regarding PNE, PoA and convergence of dynamics in aggregation games. For $n = \Theta(N)$ agents there exist (Mol et al., 2011) graphs Gr for which any uniform β leads to high PoA $= \Theta(N)$. A balanced mix of β-leaders and followers has constant PoA for constant β, but existence of PNE had only been established for $\beta = 1$.

Theorem 2. *(Mol et al., 2011) There exist connected graphs Gr such that for any β and homogeneous aggregation game $G = G(Gr, \beta, \ldots, \beta)$, $PoA(G) \geq n$.*

For any graph Gr and aggregation game $G(Gr, 0, \ldots, 0, \beta, \ldots, \beta)$ with λn β-leaders ($\beta \geq \frac{1}{n}$) and $(1-\lambda)n$ followers, we have $PoA = O(\frac{1}{1-\lambda}\min(n, \frac{1}{\beta\lambda}))$. Hence, PoA is constant for constant λ (i.e. a balanced mix) and constant β.

Any homogeneous aggregation game has an exact potential (implicitly shown in (Mol et al., 2011)). The form of the potential implies that Nash dynamics converge to PNE in polynomial time.

Theorem 3. *A homogeneous aggregation game G, i.e. $\beta_i = \beta \in [0,1]\,\forall i$ has exact potential $\Phi_\beta(H) = (1+\beta)E_H + \beta E_{H,\{1,\ldots,N\}\setminus H}$*

In contrast, the only known structural result for a heterogeneous population is that when all β_i are either 0 or 1, i.e. a mix of leaders and followers, the game has an ordinal potential. Structural results are however critical to Theorem 2 since it bounds the quality of PNE without proving that they exist.

We significantly generalize these results by using a coupling argument that proves the existence of a weighted potential function. This is always an ordinal potential and it is an exact potential if and only if all weights are 1. Formally, a game has a *weighted potential* function (Monderer & Shapley, 1996) $\Phi : \times_{i=1}^{n}\Sigma_i \to \mathbb{R}$ with (positive) weights $w_1, \ldots, w_n$ if $u_i(\sigma_i, \sigma_{-i}) - u_i(\sigma_i', \sigma_{-i}) = w_i \cdot (\Phi(\sigma_i, \sigma_{-i}) - \Phi(\sigma_i', \sigma_{-i}))$ for any agent i and any strategies $\sigma_i, \sigma_i' \in \Sigma_i, \sigma_{-i} \in \Sigma_{-i}$.

Next, we show this section's main result: any set of agents, with arbitrary $\beta_i < 1$ parameters, leads to a weighted potential function. The weighted potential is an explicit mapping of potentials in each sub-game. An analogous result when some β's equal 1 follows easily. Thus Nash dynamics converge (and PNE exist) in any aggregation game.

Theorem 4. *Fix an aggregation game G with $H = H^1 \cup \ldots \cup H^J$ where H^j are vertices occupied by all β^j-leaders, then G exhibits a weighted potential (with weights $1-\beta^j > 0$ for each agent $i \in H^j$)*

$$\Phi(H) = E_H + \sum_{j=1}^{J} \frac{\Phi_{\beta^j}(H^j, H \setminus H^j) - E_H}{1-\beta^j}$$

where $\Phi_{\beta^j}(H^j, H\setminus H^j) = (1+\beta^j)E_H + \beta^j E_{H^j,\{1,\ldots,N\}\setminus H}$ is an exact potential of the group over the (homogeneous) H^j given fixed vertices of others (in $H \setminus H^j$).

Proof. By expanding Φ in theorem (4), we get

$$\Phi(H) = \sum\nolimits_{1 \le i < i' \le n} (1 + \frac{\beta_i}{1-\beta_i} + \frac{\beta_{i'}}{1-\beta_{i'}}) E_{v_i v_{i'}} + \sum\nolimits_{1 \le i \le n} \frac{\beta_i}{1-\beta_i} E_{v_i, \{1,\ldots,N\} \setminus H} \quad (6)$$

Consider agent i that updates its vertex from v_i to v_i'. We consider the resulting state $H' = \{v_1, \ldots, v_{i-1}, v_i', v_{i+1}, \ldots, v_t\}$. We have that:

$$\Phi(v_i', v_{-i}) = \sum\nolimits_{j' : (v_i', v_{j'}) \in E} (1 + \frac{\beta_i}{1-\beta_i} + \frac{\beta_{j'}}{1-\beta_{j'}}) + \sum_{1 \le p < p' \le n, i \notin \{p,p'\}} (1 + \frac{\beta_p}{1-\beta_p} + \frac{\beta_{p'}}{1-\beta_{p'}}) E_{v_p v_{p'}}$$
$$+ \frac{\beta_i}{1-\beta_i} E_{v_i', \{1,\ldots,N\} \setminus H'} + \sum\nolimits_{p \ne i} \frac{\beta_p}{1-\beta_p} (E_{v_p, v_i} + E_{v_p, \{1,\ldots,N\} \setminus (H' \cup \{v_i\})}) \text{ and}$$

$$\Phi(v_i, v_{-i}) = \sum\nolimits_{j : (v_i, v_j) \in E} (1 + \frac{\beta_i}{1-\beta_i} + \frac{\beta_j}{1-\beta_j}) + \sum_{1 \le p < p' \le n, i \notin \{p,p'\}} (1 + \frac{\beta_p}{1-\beta_p} + \frac{\beta_{p'}}{1-\beta_{p'}}) E_{v_p v_{p'}}$$
$$+ \frac{\beta_i}{1-\beta_i} E_{v_i, \{1,\ldots,N\} \setminus H} + \sum\nolimits_{p \ne i} \frac{\beta_p}{1-\beta_p} (E_{v_p, v_i'} + E_{v_p, \{1,\ldots,N\} \setminus (H \cup \{v_i'\})})$$

where we emphasized v_i as not in H' and v_i' as not in H in the sums in the second and fourth line. After canceling common terms and using $E_{v,\Sigma} = \sum_{v' \in \Sigma} E_{v,v'}$ we get

$$\begin{aligned}
\Phi(v_i', v_{-i}) - \Phi(v_i, v_{-i}) &= \frac{1}{1-\beta_i} E_{v_i', H'} + \frac{\beta_i}{1-\beta_i} E_{v_i', \{1,\ldots,N\} \setminus H'} - \\
&\quad - (\frac{1}{1-\beta_i} E_{v_i, H} + \frac{\beta_i}{1-\beta_i} E_{v_i, \{1,\ldots,N\} \setminus H}) \\
&= (u_{\beta_i}(v_i', v_{-i}) - u_{\beta_i}(v_i, v_{-i})) \frac{1}{1-\beta_i}
\end{aligned}$$

i.e. Φ is a weighted potential with weights $1 - \beta_i, \forall i$. □

In a homogeneous aggregation game ($J = 1$, i.e. same β for all), this weighted potential reduces to the exact one in Theorem 3.

The only aggregation games not covered by Theorem 4 are ones containing leaders ($\beta = 1$). Dealing with all leaders separately, one can easily identify an ordinal potential[4]. Hence, we have that:

Corollary 2. *Each aggregation game has an ordinal potential and thus a PNE.*

7. Price of Anarchy within Groups

In this section, we will argue more generally about efficiency in coupled systems. A game is said to have a low price of anarchy (Koutsoupias & Papadimitriou, 1999) if the performance of the worst (Nash) equilibrium state is comparable to that of the socially optimal

4. The utility of a leader depends only on her actions alone. So, given any improving move of any leader their social utility will increase. As a result, a weighted sum of the social cost of the leaders and of the weighted potential function identified suffices to give an ordinal potential for any aggregation game.

state. In the case of coupled games/systems the notion of efficiency is more involved and nuanced. The success of any coupling between technological or socioeconomic systems, such as the Internet or the European Union, depends both on the performance of the system as a whole as well as on the performance of each individual subsystem. If high overall performance is achieved at the expense of some subsystem then system stability will eventually be compromised. The challenge here is to design systems that behave almost optimally at different levels of granularity. We will address such questions by generalizing the notion of (λ, μ)-smoothness (Roughgarden, 2009; Nadav & Roughgarden, 2010). A cost-minimization game with minimum-cost outcome σ' is (λ, μ)-smooth (Nadav & Roughgarden, 2010) if for all $\sigma \in \Sigma$

$$\sum_{i=1}^{n} C_i(\sigma_i', \sigma_{-i}) \leq \lambda \cdot C(\sigma') + \mu \cdot C(\sigma)$$

If G is (λ, μ)-smooth (with $\lambda \geq 0$ and $\mu \in (0,1)$), then (Roughgarden, 2009) each of G's PNE has cost at most $\lambda/(1-\mu)$ times that of a socially optimal outcome, i.e. PoA $\leq \lambda/(1-\mu)$.

Price of anarchy bounds based on (λ, μ)-smoothness extend (Roughgarden, 2009) to three other standard equilibrium concepts that we review now. A mixed Nash equilibrium (MNE) is a product probability distribution in $\Delta(\Sigma)$ in which each agent minimizes its (expected) cost given others' strategies. For any correlated equilibrium (CE) $\pi \in \Delta(\Sigma)$, if a mediator draws σ from a publicly known distribution π and reveals to each agent i only its strategy σ_i then i minimizes her expected cost by playing σ_i, assuming others also follow σ_{-i}. A coarse correlated equilibrium (CCE or equivalently Hannan-consistent strategy (Cesa-Bianchi & Lugosi, 2006)) is more general than a CE. A CCE, like a Nash equilibrium, is a probability distribution over outcomes such that no agent can improve her expected cost by deviating to a fixed strategy. Unlike Nash equilibria, a CCE can not generally be expressed as a product of agent (mixed) strategies. Average coarse correlated equilibria with respect to a socially optimal $\sigma^* \in \Sigma$ (ACCE* (Nadav & Roughgarden, 2010)) comprise the class of distributions for which the (expected) social cost is lower than the sum of costs when each agent i unilaterally deviated to σ_i^*. The best PoA bounds derived via (λ, μ)-smoothness arguments are tight for ACCE* in every game.

Formally, a *correlated equilibrium* (CE) $\pi \in \Delta(\Sigma)$ is a distribution such that $\forall i$, $\forall \sigma_i, \sigma_i' \in \Sigma_i$,

$$\sum_{\sigma_{-i} \in \Sigma_{-i}} C_i(\sigma_i, \sigma_{-i}) \pi(\sigma_i, \sigma_{-i}) \leq \sum_{\sigma_{-i} \in \Sigma_{-i}} C_i(\sigma_i', \sigma_{-i}) \pi(\sigma_i, \sigma_{-i}).$$

At a *coarse correlated equilibrium* (CCE) $\pi \in \Delta(\Sigma)$, $\forall i$, $\forall \sigma_i' \in \Sigma_i$,

$$\sum_{\sigma \in \Sigma} C_i(\sigma_i, \sigma_{-i}) \pi(\sigma_i, \sigma_{-i}) \leq \sum_{\sigma_{-i} \in \Sigma_{-i}} C_i(\sigma_i', \sigma_{-i}) \pi_i(\sigma_{-i}),$$

where $\pi_i(\sigma_{-i}) = \sum_{\tau_i \in \Sigma_i} \pi(\tau_i, \sigma_{-i})$ is the marginal probability that vector $\sigma_{-i} \in \Sigma_{-i}$ will be played. At an $ACCE^*$ π for some socially optimal σ^* (i.e. $C(\sigma^*) \leq C(\sigma), \forall \sigma \in \Sigma$) we have

$$\sum_i \sum_{\sigma \in \Sigma} C_i(\sigma_i, \sigma_{-i}) \pi(\sigma_i, \sigma_{-i}) \leq \sum_i \sum_{\sigma_{-i} \in \Sigma_{-i}} C_i(\sigma_i^*, \sigma_{-i}) \pi_i(\sigma_{-i})$$

We presented these equilibrium notions in increasing order of generality.

$$PNE \subseteq MNE \subseteq CE \subseteq CCE \subseteq ACCE^*$$

We denote the ratio of the social cost[5] of the worst equilibrium in class $\mathcal{C} \subseteq \Delta(\Sigma)$ to the optimum $\sigma^* \in \Sigma$ by[6] $\mathrm{PoA}_{\mathcal{C}} = \dfrac{\sup_{\sigma\in\mathcal{C}} C(\sigma)}{\min_{\sigma^*\in\Sigma} C(\sigma^*)}$.

For any cost-minimization game G[7], as shown in (Roughgarden, 2009; Nadav & Roughgarden, 2010)), $\mathrm{PoA} := \mathrm{PoA}_{\mathrm{PNE}} \leq \mathrm{PoA}_{\mathrm{MNE}} \leq \mathrm{PoA}_{\mathrm{CE}} \leq \mathrm{PoA}_{\mathrm{CCE}}$ [8] $\leq \mathrm{PoA}_{\mathrm{ACCE}^*} = \inf\{\frac{\lambda}{1-\mu} : G \text{ is } (\lambda,\mu)\text{-smooth}\}$.

We present localized (λ,μ)-smoothness arguments for game couplings. Analogously to the social cost, we define group j's total cost at $\sigma \in \Sigma$ by $C^j(\sigma) = \sum_{i\in N^j} C_i(\sigma)$.

Definition 3. *[group (λ^j,μ^j)-smoothness] A coupling is (λ^j,μ^j)-smooth with respect to $\sigma' \in \Sigma^j$ if for group N^j, for all $\sigma^j \in \Sigma^j$ and for all $\sigma^{-j} \in \Sigma^{-j}$:*

$$\sum\nolimits_{i\in N^j} C_i(\sigma'_i, \sigma^j_{-i}, \sigma^{-j}) \leq \lambda^j \cdot C^j(\sigma', \sigma^{-j}) + \mu^j \cdot C^j(\sigma^j, \sigma^{-j})$$

We define the local PoA for a class of probability distributions $\mathcal{C}$ by comparing the worst possible group j's cost at some $s \in \mathcal{C}$ against its cost at a benchmark state $\sigma' \in \Sigma^j$.

Definition 4. *The local price of anarchy of group j in G for a given equilibrium concept $\mathcal{C} \subseteq \Delta(\Sigma)$ with respect to benchmark $\sigma' \in \Sigma^j$ is $PoA^j_{\mathcal{C}}(\sigma') = \sup_{s\in\mathcal{C}} \frac{C^j(s)}{C^j(\sigma', s^{-j})}$.*

$\mathrm{PoA}^j_{\mathcal{C}}(\sigma')$ tests the worst-case performance of an equilibrium concept $\mathcal{C}$ (e.g. NE) from the perspective of group N^j against that of a benchmark group action $\sigma' \in \Sigma^j$. For a single group ($J=1$), if we denote $\arg\min_\sigma \sum_i C_i(\sigma)$ as OPT, the classic notion of price of anarchy corresponds to $\mathrm{PoA}^1_{\mathcal{C}}(\mathrm{OPT})$.

7.1 Dual Equilibrium Notions

We use LP duality to characterize the distributions for which $\mathrm{PoA}^j_{\mathcal{C}}(\sigma')$ bounds derived via local smoothness arguments are tight. For some target group-j benchmark action profile $\sigma' \in \Sigma^j$ we can express bounds on $\mathrm{PoA}^j_{\mathcal{C}}(\sigma')$ using linear fractional problem representations:

$$\begin{aligned} &\text{Minimize } \lambda^j \ / \ (1-\mu^j) \\ &\text{s.t.} \sum_{i\in N^j} C_i(\sigma'_i, \sigma_{-i}) \leq \lambda^j C^j(\sigma', \sigma^{-j}) + \mu^j C^j(\sigma), \forall \sigma \in \Sigma \\ &\quad \mu^j < 1 \end{aligned} \tag{7}$$

Introducing $p^j = \frac{\lambda^j}{1-\mu^j}$ and $z^j = \frac{1}{1-\mu^j}$ yields the linear program (LP)

$$\begin{aligned} &\text{Minimize } p^j \\ &\text{s.t. } p^j C^j(\sigma', \sigma^{-j}) + z^j\Big(C^j(\sigma) - \sum_{i\in N^j} C_i(\sigma'_i, \sigma_{-i})\Big) \geq C^j(\sigma), \ \forall \sigma \in \Sigma \\ &\quad z^j > 0 \end{aligned}$$

5. For randomized outcomes, i.e. distributions π, we abuse notation $C(\pi)$ to express the expected cost $\mathbb{E}_{\sigma\sim\pi}[C(\sigma)]$.
6. Linearity of expectations implies that it suffices to examine only pure (i.e. deterministic) strategy outcomes.
7. Analogous results can be shown for utility-maximization games
8. Blum et al. (Blum et al., 2008) call $\mathrm{PoA}_{\mathrm{CCE}}$ the price of total anarchy of game G.

the corresponding dual to which is as follows:

$$\begin{aligned}
&\text{Maximize } \textstyle\sum_{\sigma\in\Sigma} s_\sigma C^j(\sigma)\\
&\text{s.t. } \textstyle\sum_{\sigma\in\Sigma} s_\sigma C^j(\sigma',\sigma^{-j}) \le 1 \ \text{ and } \ s_\sigma \ge 0,\ \forall\sigma\in\Sigma\\
&\qquad \textstyle\sum_{\sigma\in\Sigma} s_\sigma\big(\sum_{i\in N^j} C_i(\sigma'_i,\sigma_{-i}) - C^j(\sigma)\big) \ge 0
\end{aligned}$$

Since the social costs are positive, we can replace the first inequality with an equality. Furthermore, since this quantity is a constant (and equal to 1), we can divide the objective by it without having any effects on the system:

$$\begin{aligned}
&\text{Maximize } \textstyle\sum_{\sigma\in\Sigma} s_\sigma C^j(\sigma) \ / \ \sum_{\sigma\in\Sigma} s_\sigma C^j(\sigma',\sigma^{-j})\\
&\text{s.t. } \textstyle\sum_{\sigma\in\Sigma} s_\sigma C^j(\sigma',\sigma^{-j}) = 1 \ \text{ and } \ s_\sigma \ge 0,\ \forall\sigma\in\Sigma\\
&\qquad \textstyle\sum_{\sigma\in\Sigma} s_\sigma\big(\sum_{i\in N^j} C_i(\sigma'_i,\sigma_{-i}) - C^j(\sigma)\big) \ge 0
\end{aligned}$$

Finally, due to scaling invariance the normalization $\sum_{\sigma\in\Sigma} s_\sigma C^j(\sigma',\sigma^{-j}) = 1$ can be replaced by $\sum_{\sigma\in\Sigma} s_\sigma = 1$, leading to:

$$\begin{aligned}
&\text{Maximize } \textstyle\sum_{\sigma\in\Sigma} s_\sigma C^j(\sigma) \ / \ \sum_{\sigma\in\Sigma} s_\sigma C^j(\sigma',\sigma^{-j})\\
&\text{s.t. } \textstyle\sum_{\sigma\in\Sigma} s_\sigma = 1 \ \text{ and } \ s_\sigma \ge 0,\ \forall\sigma\in\Sigma\\
&\qquad \textstyle\sum_{\sigma\in\Sigma} s_\sigma\big(\sum_{i\in N^j} C_i(\sigma'_i,\sigma_{-i}) - C^j(\sigma)\big) \ge 0
\end{aligned} \tag{8}$$

We wish to identify classes of distributions for which the worst case performance guarantees (for group j) match the upper bounds derived from smoothness arguments. The optimizing function of (8) has a form similar to a price of anarchy definition, hence, by defining a class of distributions based on the feasibility constraints we can achieve our goal of having equilibrium classes for which smoothness arguments are tight. Specifically, we can define group j's *average coarse correlated equilibria with respect to* $\sigma' \in \Sigma^j$ ($\text{ACCEL}^j(\sigma')$) as distributions for which local PoA^j bounds via (λ^j,μ^j)-smoothness are always tight.

Definition 5. $ACCEL^j(\sigma') = \{s\in\Delta(\Sigma) \colon \sum_{\sigma\in\Sigma} C^j(\sigma)s(\sigma) \le \sum_{\sigma\in\Sigma}\sum_{i\in N^j} C_i(\sigma'_i,\sigma_{-i})s(\sigma)\}$.

Clearly, $\text{ACCEL}^j(\sigma')$ is nonempty for any (coupled) game G and any $\sigma' \in \Sigma^j$ since it includes all of its coarse correlated equilibria (i.e. $\text{CCE} \subseteq \text{ACCEL}^j(\sigma')$). It also captures exactly the feasibility constraints of (8). Tightness of PoA^j bounds follows:

Theorem 5. *For* $\mathcal{C} = ACCEL^j(\sigma')$,
$PoA^j_{\mathcal{C}}(\sigma') = \inf\{\frac{\lambda^j}{1-\mu^j} \colon$ *G is* (λ^j,μ^j)*-locally smooth for group* N^j *with respect to* $\sigma' \in \Sigma^j\}$.

Proof. By definition, the local price of anarchy of group j in G for any set of (randomized) outcomes $\mathcal{C} \subseteq \Delta(\Sigma)$ is $\text{PoA}^j_{\mathcal{C}}(\sigma') = \sup_{\sigma\in\mathcal{C}} \dfrac{C^j(\sigma)}{C^j(\sigma',\sigma^{-j})}$. Given any such set of (randomized) outcomes $\mathcal{C} \subseteq \Delta(\Sigma)$, we can find a sequence of distributions $\sigma_\tau \in \mathcal{C}$, such that $\text{PoA}^j_{\mathcal{C}}(\sigma') = \lim_{\tau\to\infty} \frac{C^j(\sigma_\tau)}{C^j(\sigma',\sigma_\tau^{-j})}$. We have that $\text{PoA}^j_{\text{ACCEL}^j(\sigma')}(\sigma')$ can be expressed as a limit of solutions of linear programs of the form (8). LP duality implies that any $\frac{\lambda^j}{1-\mu^j}$ for which G is locally (λ^j, μ^j)-smooth with respect to σ' is an upper bound on $\mathrm{PoA}^j_{\mathrm{ACCEL}^j(\sigma')}(\sigma')$ as a valid solution to the corresponding dual linear fractional program (7). The theorem follows immediately, since by strong LP duality both $\mathrm{PoA}^j_{\mathrm{ACCEL}^j(\sigma')}(\sigma')$ and $\inf\{\frac{\lambda^j}{1-\mu^j}$: G is $(\lambda^j, \mu^j)-$ locally smooth with respect to $\sigma'\}$ capture the limit $\lim_{\tau\to\infty} \frac{C^j(\sigma_\tau)}{C^j(\sigma', \sigma_\tau^{-j})}$. $\square$

7.2 Coupled Games and Learning via Public Advertising

So far, we have mainly been focusing on different classes of equilibrium behavior, either establishing their existence and the convergence of simple (*e.g* best-response) dynamics to them, or arguing about their efficiency. We will finish our exposition of coupled game-theoretic environments by considering families of online learning dynamics in this setting.

A coupled game essentially establishes a hierarchical environment. A group defines a higher level of organization that consists of several self-interested entities/units. At the same time, as we have argued in the previous section, an organization is also a goal driven entity that strives for (approximate) optimality of its collective performance. Along those lines, it makes sense to consider online learning paradigms that infuse sophisticated learning behavior both at the level of each group as well as at the level of individual agents. At the core of our approach lie ideas from regret-minimizing online learning (Appendix A, (Cesa-Bianchi & Lugosi, 2006; Young, 2004) for more detailed presentation). These are practical algorithms that provide strong performance guarantees and at the same time enjoy connections to equilibrium notions such as coarse correlated equilibria. Specifically, when each agent updates her strategy according to a regret-minimizing dynamic then the long term average of action profiles converges weakly to the set of coarse correlated equilibria (CCE) (Young, 2004). We will explore analogous connections between learning dynamics in coupled games and our ACCEL/local smoothness arguments to provide insights about the possible outcomes of sophisticated learning behavior in the setting of coupled games.

We introduce a novel learning procedure that incorporates public advertising and which allows for provable welfare guarantees both on the level of groups as well as globally. Intuitively, the setting is as follows: Within each group j there exists a broadcasting center that can broadcast to all agents in the group. On each day $t = 1 \dots T$, the center of group j computes a strategy vector $\text{ADV}^j(t)$ for the group and advertises to each agent i his respective strategy $\text{ADV}^j_i(t)$. There exist two high level issues in any such model: first, how does the center decide on which vector to advertise and second, how do the individual agents respond to the recommendations?

In terms of center actions, prior public advertising models (Balcan et al., 2009, 2010, 2012) assumed that there exists a single center with full information over the whole game that is able to broadcast to all agents. In such settings the center can easily broadcast a global optimum solution or the best Nash equilibrium. In contrast, we are moving towards a more restricted and realistic model where each center only controls a *local* neighborhood of agents. Many real life settings share this structure (e.g. competing Internet autonomous systems, or more generally competing institutions/organizations). In such settings, the managing centers have a high incentive in employing sophisticated online algorithms in order to effectively calibrate their predictions. Here, we will analyze centers whose advertised behaviors exhibits vanishing average regret with hindsight. This is a rather natural benchmark, since several simple learning algorithms can offer such guarantees[9].

On the side of the individual agents, we make a similarly weak assumption. We assume that the average performance of each agent i (of group j) will eventually be (almost) as high as that of his advertised strategy $\text{ADV}_i^j(t)$. Any dummy agent can meet this benchmark merely by following the recommended strategy. A more realistic agent could still achieve such guarantees by interpolating between his innate learning strategy and the provided advice. We will show that advertising-guided learning offers guarantees analogous to those of the ACCEL framework. We start by bounding the possible negative effects of agents' experimentation. We will use a slightly stronger local smoothness property:

Definition 6 (*strong*$(\vec{\lambda},\vec{\mu})$-smoothness)**.** *A coupling is $(\vec{\lambda},\vec{\mu})$-smooth if for each group N^j, for all $\sigma' \in \Sigma^j$, for all $\sigma^j \in \Sigma^j$ and for all $\sigma^{-j} \in \Sigma^{-j}$:*

$$\sum_{i\in N^j} C_i(\sigma'_i, \sigma^j_{-i}, \sigma^{-j}) \leq \lambda^j \cdot C^j(\sigma', \sigma^{-j}) + \mu^j \cdot C^j(\sigma^j, \sigma^{-j})$$

This definition has a stronger flavor than the one in Definition 3, since here the deviating condition must hold for all possible group responses and not only for target benchmarks. This definition is more reminiscent of the original (λ,μ)-smoothness definition in (Roughgarden, 2009), which is slightly different than the one in (Nadav & Roughgarden, 2010). When we are referring to $(\vec{\lambda},\vec{\mu})$-smoothness in this section, we will be working with this slightly stronger definition.

Lemma 2. *If in group j the time-average[10] cost of each i is (almost) as low as that of i's advertised strategies ADV_i^j,*

$$\tfrac{1}{T}\sum_t C_i\big(\sigma(t)\big) \leq \tfrac{1}{T}\sum_t C_i\big(\text{ADV}_i^j(t), \sigma_{-i}(t)\big) + o(1)$$

then advertising-guided learning only incurs a $\lambda^j/(1-\mu^j)$ overhead when compared to the advertised strategy.

$$\tfrac{1}{T}\sum_t C^j(\sigma(t)) \leq \tfrac{\lambda^j}{1-\mu^j}\tfrac{1}{T}\sum_t C^j\big(\text{ADV}^j(t), \sigma^{-j}(t)\big) + o(1)$$

Proof.

$$\begin{aligned}
&\tfrac{1}{T}\sum_t C^j\big(\sigma(t)\big) = \tfrac{1}{T}\sum_t \sum_{i\in N^j} C_i\big(\sigma(t)\big)\\
&\leq \tfrac{1}{T}\sum_t \sum_{i\in N^j} \Big(C_i\big(\text{ADV}_i^j(t), \sigma_{-i}(t)\big) + o(1)\Big)\\
&\leq \tfrac{1}{T}\sum_t \big(\lambda^j C^j\left(\text{ADV}^j(t), \sigma^{-j}(t)\right) + \mu^j C^j\left(\sigma(t)\right)\big) + o(1)
\end{aligned}$$

Rearranging terms leads to the desired form. □

9. A policy (sequence of strategies) satisfies the no-regret property if its average payoff is almost as good as that of the best fixed (time-invariant) strategy given the history of play. CCE are limit points of time-averages of no-regret policies. Generally, no-regret algorithms offer guarantees in expectation over their randomized strategies. For ease of notation, we consider pure strategy outcomes. The analysis trivially extends to the case of randomized strategies.

10. In this section, whenever we write $\sum_t$, we mean $\sum_{t=1}^T$.

Given a history of play $\sigma(1), \ldots, \sigma(T)$, we denote the best group response of group j with hindsight as $\textsc{opt}^j(T)$:

$$\textsc{opt}^j(T) = \text{argmin}_{s^j \in \Sigma^j} \tfrac{1}{T} \sum_t C^j\big(s^j, \sigma^{-j}(t)\big)$$

Given a game coupling $(N^1, N^2, \ldots, N^J)$, we define its *super-game* as follows: it is a game with J agents, the available strategies to each super-agent j correspond to strategy tuples for all agents in group j, i.e. $\sigma^j \in \times_{i \in N^j} \Sigma_i$. Finally, the cost of the super-agent j is the group cost for all agents in group j, i.e. $C^j(\sigma) = \sum_{i \in N^j} C_i(\sigma)$. We also assume that super-game is $(\lambda^{sup}, \mu^{sup})$-smooth[11]. Finally, we define a socially optimal strategy vector as $global\_\textsc{opt} \in \text{argmin}_\sigma C(\sigma)$.

We now prove cost bounds for advertising-guided learning.

Theorem 6. *If in each group j the time-average cost of each i is (almost) as low as that of i's advertised strategies $\textsc{adv}_i^j$,*

$$\tfrac{1}{T} \sum_t C_i\big(\sigma(t)\big) \le \tfrac{1}{T} \sum_t C_i\big(\textsc{adv}_i^j(t), \sigma^{-i}(t)\big) + o(1),$$

and the advertised strategy for each group j has vanishing time-average regret,

$$\tfrac{1}{T}\sum_t C^j\big(\textsc{adv}^j(t), \sigma^{-j}(t)\big) \le \tfrac{1}{T}\sum_t C^j\big(\textsc{opt}^j(T), \sigma^{-j}(t)\big) + o(1)$$

then for advertising-guided learning, the group cost satisfies

$$\frac{1/T \sum_t C^j\big(\sigma(t)\big)}{1/T \sum_t C^j\big(\textsc{opt}^j(T), \sigma^{-j}(t)\big)} \le \frac{\lambda^j}{1-\mu^j} + o(1)$$

and for $\min_j \frac{1-\mu^j}{\lambda^j} > \mu^{sup}$*, the social cost satisfies*

$$\frac{1/T \sum_t C\big(\sigma(t)\big)}{\min_{\sigma'} C(\sigma')} \le \frac{\lambda^{sup}}{\min_j \frac{1-\mu^j}{\lambda^j} - \mu^{sup}} + o(1)$$

Proof. By applying Lemma 2, we have for each group j:

$$\begin{aligned} \tfrac{1}{T} \sum_t C^j(\sigma(t)) &\le \frac{\lambda^j}{1-\mu^j} \tfrac{1}{T} \sum_t C^j\big(\textsc{adv}^j(t), \sigma^{-j}(t)\big) + o(1) \\ &\le \frac{\lambda^j}{1-\mu^j} \tfrac{1}{T} \sum_t C^j\big(\textsc{opt}^j(T), \sigma^{-j}(t)\big) + o(1) \qquad (9) \\ &\le \frac{\lambda^j}{1-\mu^j} \tfrac{1}{T} \sum_t C^j\big(global\_\textsc{opt}^j, \sigma^{-j}(t)\big) + o(1) \qquad (10) \end{aligned}$$

From Eq. (9), we have for each group j:

$$\frac{1/T \sum_t C^j\big(\sigma(t)\big)}{1/T \sum_t C^j\big(\textsc{opt}^j(T), \sigma^{-j}(t)\big)} \le \frac{\lambda^j}{1-\mu^j} + o(1)$$

11. This is the standard notion of (non-coupled) game smoothness

We know that the (super)-game is $(\lambda^{sup}, \mu^{sup})$-smooth, therefore:

$$\frac{1}{T}\sum_t \sum_j C^j\big(global\_\text{OPT}^j, \sigma^{-j}(t)\big) \leq \\ \frac{1}{T}\sum_t \lambda^{sup} C\big(global\_\text{OPT}\big) + \frac{1}{T}\sum_t \mu^{sup} C\big(\sigma\big)$$

Combining this with Eq. (10), we have for $\min_j \frac{1-\mu^j}{\lambda^j} > \mu^{sup}$:

$$\frac{1}{T}\sum_t C\big(\sigma(t)\big) \leq \frac{\lambda^{sup}}{\min_j \frac{1-\mu^j}{\lambda^j} - \mu^{sup}} \min_{\sigma'} C(\sigma') + o(1)$$

□

8. Concluding remarks

Modern engineering systems (such as the Internet), as well as social systems, exhibit internal hierarchical structure. Indeed, such systems can intuitively be viewed as a coupling between subsystems of increased internal homogeneity. In fact, for some applications, such a design paradigm is highly desirable since it allows for maximum flexibility and scalability. Naturally, however, questions about the stability and the long-term performance of such systems arise. Nevertheless, despite the ubiquitous nature of such systems and inquiries, few formal tools have been developed that allow us to argue about the degradation (or possible enhancement) of their local properties as we move towards larger coupled and increasingly heterogeneous systems.

In this work, we have introduced and studied game couplings, a concept that encapsulates globally heterogeneous populations exhibiting local homogeneity. We gave several applications of this framework in terms of learning in games, quality of equilibria (PoA) and structural properties. Specifically, we have shed light on the nature of local properties that allow us to argue about the evolution and performance of the global system. Furthermore, we have studied the stresses between local subsystem optimality and global system performance. Finally, we have identified specific settings where the coupling of heterogeneous systems is not only an elegant design solution but actually necessary for global system optimality.

An interesting direction for future work would be the principled study of heterogeneity itself. Specifically, it would be rather enticing to identify novel measures of heterogeneity of structure that capture the reality of real world coupled systems and which are theoretically tractable. Finally, our work here focuses mainly on systems with two levels of hierarchy (individual subsystems and coupled global system). A natural extension of these models would include more intricate hierarchical structures such as rooted trees, or directed acyclic graphs.

This work was supported by NSF grants CCF-0953192 and CCF-1101215, by ONR grant N00014-09-1-0751, by AFOSR grant FA9550-09-1-0538, and by a Microsoft Research Faculty Fellowship.

References

Aiyer, A. S., Alvisi, L., Clement, A., Dahlin, M., Martin, J.-P., & Porth, C. (2005). Bar fault tolerance for cooperative services. *SIGOPS Oper. Syst. Rev.*, *39*(5), 45–58.

Babaioff, M., Kleinberg, R., & Papadimitriou, C. H. (2009). Congestion games with malicious players. *Games and Economic Behavior*, *67*(1), 22–35.

Balcan, M., Krehbiel, S., Piliouras, G., & Shin, J. (2012). Minimally invasive mechanism design: Distributed covering with carefully chosen advice. In *51st IEEE Annual Conference on Decision and Control (CDC)*, pp. 2690–2695.

Balcan, M.-F., Blum, A., & Mansour, Y. (2009). Improved equilibria via public service advertising. In *Proceedings of Symposium on Discrete Algorithms (SODA)*, pp. 728–737.

Balcan, M.-F., Blum, A., & Mansour, Y. (2010). Circumventing. In *Proc. of Innovations in Computer Science*, pp. 200–213.

Blum, A., Hajiaghayi, M., Ligett, K., & Roth, A. (2008). Regret minimization and the price of total anarchy. In *Symposium on Theory of Computing (STOC)*, pp. 373–382.

Candogan, O., Menache, I., Ozdaglar, A., & Parrilo, P. A. (2011). Flows and decompositions of games: Harmonic and potential games. *Mathematics of Operations Research*, *36*(3), 474–503.

Cesa-Bianchi, N., & Lugosi, G. (2006). *Prediction, Learning, and Games*. Cambridge University Press.

Fabrikant, A., Jaggard, A., & Schapira, M. (2013). On the structure of weakly acyclic games. *Theory of Computing Systems*, *53*(1), 107–122.

Fox, M., Piliouras, G., & Shamma, J. (2012). Medium and long-run properties of linguistic community evolution. In *9th International Conference on the Evolution of Language (Evolang IX)*, pp. 110–118.

Gairing, M. (2008). Malicious Bayesian congestion games. In *Workshop on Approximation and Online Algorithms (WAOA)*, pp. 119 – 132.

Hall, C., Anderson, R., Clayton, R., Ouzounis, E., & Trimintzios, P. (2001). *Resilience of the Internet Interconnection Ecosystem*. European Network and Information Security Agency.

Immorlica, N., Markakis, E., & Piliouras, G. (2010). Coalition formation and price of anarchy in Cournot oligopolies. In *Proceedings of the 6th international conference on Internet and network economics*, WINE'10, pp. 270–281, Berlin, Heidelberg. Springer-Verlag.

Kleinberg, R., Piliouras, G., & Tardos, É. (2011). Load balancing without regret in the bulletin board model. *Distributed Computing*, *24*(1), 21–29.

Kleinberg, R., Piliouras, G., & Tardos, É. (2009). Multiplicative updates outperform generic no-regret learning in congestion games: extended abstract. In *Proceedings of the 41st annual ACM symposium on Theory of computing*, STOC '09, pp. 533–542, New York, NY, USA. ACM.

Kleinberg, R. D., Ligett, K., Piliouras, G., & Tardos, É. (2011). Beyond the Nash equilibrium barrier. In *Innovations in Computer Science (ICS)*, pp. 125–140.

Koutsoupias, E., & Papadimitriou, C. H. (1999). Worst-case equilibria. In *Symposium on Theoretical Aspects of Computer Science (STACS)*, pp. 404–413.

Li, N., & Marden, J. R. (2011). Designing games for distributed optimization. In *Proceedings of 50th IEEE Conference on Decision and Control*, pp. 2434–2440.

Ligett, K., & Piliouras, G. (2011). Beating the best Nash without regret. *SIGecom Exch.*, *10*(1), 23–26.

Marden, J., Arslan, G., & Shamma, J. S. (2007). Autonomous vehicle-target assignment: A game theoretical formulation. *ASME Journal of Dynamic Systems, Measurement and Control*, *129*(5), 584–596.

Marden, J., Arslan, G., & Shamma, J. S. (2009a). Cooperative control and potential games. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, *39*(6), 1393–1407.

Marden, J., Arslan, G., & Shamma, J. S. (2009b). Payoff-based dynamics for multi-player weakly acyclic games. *SIAM Journal on Control and Optimization, special issue on Control and Optimization in Cooperative Networks*, *48*(1), 373–396.

Mol, P., Vattani, A., & Voulgaris, P. (2011). The effects of diversity in aggregation games. In *Innovations in Computer Science*, pp. 509–521.

Monderer, D., & Shapley, L. (1996). Potential games. *Games and Econ. Behavior*, *14*, 124–143.

Monderer, D. (2007). Multipotential games. In *Proceedings of the 20th international joint conference on Artificial Intelligence (IJCAI)*, pp. 1422–1427.

Nadav, U., & Roughgarden, T. (2010). The limits of smoothness: Primal-dual framework for price of anarchy bounds. In *WINE*, pp. 319–326.

Nadav, U., & Piliouras, G. (2010). No regret learning in oligopolies: Cournot vs. Bertrand. In *Proceedings of the third international conference on Algorithmic game theory*, SAGT'10, pp. 300–311, Berlin, Heidelberg. Springer-Verlag.

Piliouras, G., Valla, T., & Végh, L. A. (2012). LP-based covering games with low price of anarchy. In *Proceedings of the 8th international conference on Internet and Network Economics*, WINE'12, pp. 184–197, Berlin, Heidelberg. Springer-Verlag.

Rosenthal, R. (1973). A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, *2*, 65–67.

Roughgarden, T. (2009). Intrinsic robustness of the price of anarchy. In *Proc. of Symposium on Theory of Computing (STOC)*, pp. 513–522.

Roughgarden, T., & Schoppmann, F. (2011). Local smoothness and the price of anarchy in atomic splittable congestion games. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 255–267. SIAM.

Sandholm, W. (2010). Decompositions and potentials for normal form games. *Games and Econ. Behavior*, *70*(2), 446–456.

Tumer, K., & Wolpert, D. (2004). A survey of collectives. In *Collectives and the Design of Complex Systems*, pp. 1–42. Springer.

Vetta, A. (2002). Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *Foundations of Computer Science, 2002. Proceedings. The 43rd Annual IEEE Symposium on*, pp. 416–425. IEEE.

Wolpert, D., & Bieniawski, S. (2004). Distributed control by Lagrangian steepest descent. In *Proc. 43rd IEEE Conf. on Decision and Control*, pp. 1562–1567.

Young, H. P. (1998). *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton, NJ: Princeton University Press.

Young, H. P. (2004). *Strategic Learning and Its Limits*. Oxford UK: Oxford University Press.

Appendix A. Regret minimization

We will give a formal definition of having no-regret in an online sequential problem.

Definition 7. *An online sequential problem consists of a feasible set* $F \in \mathbb{R}^m$*, and an infinite sequence of functions* $\{f^1, f^2, \dots,\}$ *where* $f^t : \mathbb{R}^m \to \mathbb{R}$.

At each time step t, an online algorithm selects a vector $x^t \in \mathbb{R}^m$. After the vector is selected, the algorithm receives f^t, and collects a payoff of $f^t(x^t)$. All decisions must be made *online*, in the sense that an algorithm does not know f^t before selecting x^t, i.e., at each time t, a (possibly randomized) algorithm can be thought of as a mapping from a history of functions up to time t, $f^1, \dots, f^{t-1}$ to the set F.

Given an algorithm A and an online sequential problem $(F, \{f^1, f^2, \dots\})$, if $\{x^1, x^2, \dots\}$ are the vectors selected by A, then the payoff of A until time T is $\sum_{t=1}^{T} f^t(x^t)$. The payoff of a static feasible vector $x \in F$, is $\sum_{t=1}^{T} f^t(x)$. Regret compares the performance of an algorithm with the best static action in hindsight:

Definition 8. *The external regret of algorithm* A*, at time* T *is defined as*

$$R(T) = \max_{x \in F} \sum_{t=1}^{T} f^t(x) - \sum_{t=1}^{T} f^t(x^t)$$

An algorithm is said to have no-external regret, if for every online sequential problem, its regret at time T *is* $o(T)$*.*

The definition of regret minimization in the case of cost optimization (negative payoffs) is completely analogous. More detailed definitions and examples can be found here (Cesa-Bianchi & Lugosi, 2006; Young, 2004).