

Equilibria and Efficiency Loss in Games on Networks

Joshua R. Davis* Zachary Goldman† Elizabeth N. Koch*‡ Jacob Hilty*§
David Liben-Nowell* Alexa Sharp¶ Tom Wexler†¶ Emma Zhou*||

Abstract

Social networks are the substrate upon which we make and evaluate many of our daily decisions: our costs and benefits depend on whether—or how many of, or which of—our friends are willing to go to that restaurant, choose that cellular provider, already own that gaming platform. Much of the research on the “diffusion of innovation,” for example, takes a game-theoretic perspective on strategic decisions made by people embedded in a social context. Indeed, multiplayer games played on social networks, where the network’s nodes correspond to the game’s players, have proven to be fruitful models of many natural scenarios involving strategic interaction.

In this paper, we embark on a mathematical and general exploration of the relationship between two-person strategic interactions (a “base game”) and a “networked” version of that same game. We formulate a generic mechanism for superimposing a symmetric two-player base game M on a social network G : each node of G chooses a single strategy from M and simultaneously plays that strategy against each of its neighbors in G , receiving as its payoff the sum of the payoffs from playing M against each neighbor. We denote the *networked game* that results by $M \oplus G$. We are broadly interested in the relationship between properties of M and of $M \oplus G$: how does the character of strategic interaction change when it is embedded in a social network? We focus on two particular properties: the (pure) price of anarchy and the existence of pure Nash equilibria. We show tight results on the relationship between the price of anarchy in M and $M \oplus G$ in coordination games. We also show that, with some exceptions when G is bipartite, the existence or absence of pure Nash equilibria (and even the guaranteed convergence of best-response dynamics) in M and $M \oplus G$ are not entailed in either direction. Taken together, these results suggest that the process of superimposing M on a graph is a nontrivial operation that can have rich, but bounded, effects on the strategic environment.

Appears in *Internet Mathematics*, 7(3):178–205, 2011. An abbreviated preliminary version of this paper appears in the *Proceedings of the IEEE International Conference on Social Computing (SocialCom’09)*, pp. 82–89, August 2009. All authors were supported in part by NSF grant CCF-0728779. This work was also supported by grants from Carleton College, Oberlin College, and Denison University. Thanks to Michael Brooks, Ramesh Johari, Katie Kuksenok, Eric Naeseth, and Charles Noneman for helpful discussions and comments. Please address correspondence to: David Liben-Nowell; Department of Computer Science; Carleton College; 1 N. College St.; Northfield, MN 55057. Phone: +1 507 222 4679. Fax: +1 507 222 4312. Email: dlibenno@carleton.edu. Email addresses of all authors: jdavis@carleton.edu, goldma_z@denison.edu, jacob.a.hilty@gmail.com, elizabeth.koch@alummi.carleton.edu, dlibenno@carleton.edu, alexa.sharp@oberlin.edu, tom.wexler@oberlin.edu, emma@rapleaf.com

*Department of Computer Science, Carleton College, Northfield, MN 55057

†Department of Mathematics and Computer Science, Denison University, Granville, OH 43023

‡Department of Computer Science and Engineering, University of Minnesota, Minneapolis, MN 55455

§School of Philosophy, Psychology and Language Sciences, University of Edinburgh, Edinburgh, Scotland EH8 9AD

¶Department of Computer Science, Oberlin College, Oberlin, OH 44074

||Rapleaf, Inc., San Francisco, CA 94105

1 Introduction

In recent years there has been significant and growing interest in games played on networks, where the game’s players are represented by the nodes of the network. Within this space, there has been particular attention paid to social networks, in which edges connect pairs of people whose actions can directly impact each other. This growing line of research can in part be attributed to a sense that such games are indeed the “right” model for many natural scenarios: many human endeavors can be viewed as games in which a person’s utility is determined by the behavior of those who are in some sense close by—friends, associates, or trading partners, for example. The actions of other players in the network are felt only indirectly; their decisions may or may not cause cascades of local changes that eventually propagate to distant players.

A rich vein of research, and an illustrative example of this game-theoretic style of work on social networks, has been carried out under the “diffusion of innovation” rubric in the literature [6,17–20,30,33–35,37, e.g.]. Imagine, for example, a population of mobile phone users, each of whom must choose whether to subscribe to an unlimited text-messaging plan with her cellular provider, or to use a pay-per-message plan. If many of a person u ’s friends adopt an unlimited-texting plan, then u may receive many (perhaps unsolicited) text messages from her friends, incurring a large number of high pay-per-message charges; conversely, if very few of u ’s friends have adopted the unlimited-message option, then u will likely receive few text messages, and thus the pay-per-use price will likely be a better deal. This scenario—glibly, a strategic multi-agent version of the ski-rental problem—requires people to think not only about the relative costs of “renting” and “buying” but also about the structure of their social network, and the choices that their friends make. In this example, one might expect to observe clusters of adopters in the social network that are relatively isolated from corresponding clusters of nonadopters. In the real world, these clusters will be highly correlated with age, and in part result from the significant correlation in age between people connected by a social tie (see [29]). This “age homophily” illustrates the perhaps-obvious point that in many network-type games the structure of the network matters critically. But, perhaps also obviously, the structure of the game matters too. Sometimes a strategy—adopting a new piece of communication technology, say—becomes more attractive as additional people, particularly one’s friends, choose it; sometimes—ordering a dish at a restaurant where food will be shared, say—it becomes less attractive [23,28, e.g.]. And more complex strategic landscapes are possible too.

In this paper, we are broadly concerned with the way in which local decision making—whether to buy the unlimited text-messaging plan, whether to order the pad thai or the massaman curry, whether to go to the opera or the baseball game—is affected when it is embedded into a social-network context. We formulate a general framework for superimposing a symmetric two-player game, which we call a *base game*, on a social network, creating a networked version of the base game. The nodes of the network correspond to the players, each of whom must choose a single strategy from the set of available strategies in the base game. Each player’s payoff is the sum of its payoffs from playing the base game simultaneously with its neighbors. (In essence, each player simultaneously plays the base game with each of its neighbors in the graph, restricted so that players must act consistently across their multiple games.) Our particular line of inquiry is directed towards exploring the relationship among a base game M , a graph G , and the resulting networked game $M \oplus G$.

Our results. We study the networked games resulting from superimposing arbitrary two-player symmetric games on arbitrary undirected graphs. We are primarily interested in understanding how properties of a base game carry over—or do not carry over—to the networked game. There are many ways to ask about the relationship among M , G , and $M \oplus G$. We focus on two specific instantiations of this question, as initial steps in this line of study: the (pure) price of anarchy and the existence of (pure) Nash equilibria. See Section 3 for definitions. (Throughout, we consider only pure strategies except where mixed strategies are explicitly mentioned; see Sections 5 and 7 for discussion of mixed strategies.)

- **The price of anarchy.** The *(pure) price of anarchy* of a game is the ratio of the social welfare—the sum of the payoffs for all players—in the best outcome (“OPT”) to the worst Nash equilibrium (“WNE”). We are interested in how the price of anarchy of the base game M relates to the price of anarchy of the networked game $M \oplus G$. Here we focus on *coordination games* (also known as *matching games*) as base games. In these games, players receive identical positive payoffs if they choose the same strategy from the set of options, and payoffs of zero if they choose different strategies. (The payoffs the players receive for matching can depend on which strategy they choose.) Coordination games and variants are frequently used to model diffusion of innovation: for example, strategies might represent choices of communication technologies (whether to buy a fax machine, which computer operating system to choose) where there is utility in making the same choice as one’s friends. Let M_{coord} be any coordination game of this type. In Section 5, we give tight bounds on the maximum (taken over all graphs) of the price of anarchy for $M_{\text{coord}} \oplus G$.
- **The existence of pure Nash equilibria.** One of the most basic game-theoretic questions that one can ask is whether a pure Nash equilibrium exists in a given game. The analogous question here is the connection between the existence of pure Nash equilibria in the base game and the existence of pure Nash equilibria in the networked game. In Section 6, we show largely negative results about this connection. If M has a pure Nash equilibrium and G is bipartite, then $M \oplus G$ has a pure Nash equilibrium as well; in all other cases, though, we show that the existence (or absence) of pure Nash equilibria in M does not imply anything about the existence (or absence) of pure Nash equilibria in $M \oplus G$. We give examples of base games M and (bipartite and nonbipartite) graphs G_1 and G_2 such that the properties of M and $M \oplus G_1$ match with respect to the existence of pure Nash equilibria, but M and $M \oplus G_2$ mismatch. Furthermore, we give examples of base games M in which best-response dynamics is guaranteed to converge in M but is not guaranteed to converge in $M \oplus G$, and vice versa.

Our work on the existence of pure Nash equilibria illustrates that the base game and the networked game are qualitatively different. When a base game is “networkified,” pure Nash equilibria can be created or destroyed; the guaranteed convergence of best-response dynamics can be introduced or eliminated. This qualitative difference is largely unsurprising; indeed, networked games are interesting only because they do not precisely replicate their base game. What is more surprising is that there are still nontrivial quantitative similarities between the base game and the networked game, for example in the price of anarchy, that can be proven independent of the size or topology of the graph. We believe that further quantification of the effect of moving a base game to a networked context is a fertile area for study.

2 Related Work

As described in Section 1, there has been appreciable work on networked versions of particular games, usually on particular classes of networks, in modeling the diffusion of innovation [6, 17–20, 30, 33–35, 37]. These models typically capture scenarios in which there is incentive towards assortative behavior; other work has explored models where players have incentive towards disassortative behavior, including *cut games*, and *party affiliation games*, which subsume both cut games and coordination games [2, 5, 12]. (One can think of a cut game as a networked variant of the El Farol Bar problem, in which a player wants to make a decision that is matched by a minority of that player’s neighbors [1].)

A networked game $M \oplus G$ is a special form of a *graphical game* [4, 7, 9–11, 15, 22, 24, 26], an n -player game in which the payoff to a player u is affected only by the strategies of the neighbors of u in an underlying n -node graph. (Other formalisms for games on networks have also been considered; see [14, 16, 21] for some examples.) Graphical games allow u ’s payoff to depend arbitrarily on the strategies chosen by u ’s neighbors; in our networked games u ’s payoff is simply the sum of a payoff on each edge incident to u . Graphical games are interesting when the underlying graph is sparse, as social networks are; otherwise the graphical structure does not impose much limitation. Our networked games, as well as these “sparse” graphical games, form a natural class of *compactly representable games*—games that can be specified in space polynomial in the number of players and strategies. Fully general games are typically of less interest, both practically and theoretically: practically, general games require exponential space to describe and thus are too large to be tractable; and, theoretically, this huge game description trivializes various computational problems—searching for a pure Nash equilibrium in a general game can be solved by brute force in time linear in the input size.

One paper on graphical games is particularly close in spirit to our work here: Ben Zwi and Ronen [4] study the relationship between the (“global”) price of anarchy of a graphical game and what they call the “local price of anarchy.” The latter measures how well any subset S of players in the network responds to choices made by the nodes outside of S (where “how well” is measured in terms of the ratio of the worst Nash equilibrium WNE to the social optimum OPT within the subgame induced by fixing the strategies of all non- S players).

A recent paper by Daskalakis and Papadimitriou [8] studies a special case of our problem: networked games arising from zero-sum (or, more generally, strictly competitive) base games. They demonstrate in this case that the mixed Nash equilibria of the networked game can be computed efficiently, by reducing their computation to that of the mixed Nash equilibria of a particular two-player zero-sum game.

At a more abstract level, one can see a parallel between our work here and recent work in epidemiology. Work in that field can loosely be categorized as falling under the “fully mixed model,” in which one models any two members of a large population as equally likely to interact, or under the “network model,” in which an underlying social/contact network reflects the latent structure in people’s interactions. Network-based models more realistically reflect the ways in which diseases spread, but in most cases the fully mixed model is far more amenable to rigorous mathematical analysis, particularly as the population grows large. In the game-theoretic context, there is a significant body of work on “fully mixed”-type models in which randomly selected pairs of agents from a large population play a particular two-player game, possibly repeatedly—for example, see the classic text of Fudenberg and Levine [13] on learning in games. (Their motivations include giving an account of how players might settle on a Nash equilibrium, and how players might select

a particular equilibrium.) In our setting, we consider a fixed network of interactions; our interests are in the ways that the networked game relates to the base game. This network-based perspective on evolutionary game theory was introduced in the work of Kearns and Suri [25], who were largely interested in extending classic results of evolutionary game theory to the network-based setting.

3 Background and Notation

Games

An n -player game M consists of a set of *players* $\{1, \dots, n\}$, a *strategy set* S_i for each player i , and a *payoff function* $p_i : (S_1 \times \dots \times S_n) \rightarrow \mathbb{Z}$ for each player. Each player i chooses a strategy $s_i \in S_i$; the vector $\mathbf{s} = \langle s_1, \dots, s_n \rangle$ is called a *strategy profile*. The *payoff* to player i under \mathbf{s} is $p_i(\mathbf{s})$, and the *social welfare* of \mathbf{s} is $\sum_i p_i(\mathbf{s})$. The *social optimum*, denoted OPT, is the strategy profile \mathbf{s} that achieves the maximum social welfare. A strategy profile \mathbf{s} is a (*pure*) *Nash equilibrium* if no player can unilaterally deviate from \mathbf{s} to improve her payoff, i.e., if $p_i(\mathbf{s}) \geq p_i(s_1, \dots, s_{i-1}, s'_i, s_{i+1}, \dots, s_n)$ for every i and every $s'_i \in S_i$. We denote the worst Nash equilibrium—that is, the pure Nash equilibrium \mathbf{s} with the lowest social welfare—by WNE, and the best Nash equilibrium by BNE. We also write OPT, WNE, and BNE to denote the social welfare of the corresponding strategy profile OPT, WNE, and BNE. The *price of anarchy* is given by $\text{POA} := \text{OPT}/\text{WNE}$, and the *price of stability* is given by $\text{POS} := \text{OPT}/\text{BNE}$. Note again that all of these quantities refer only to pure Nash equilibria.

In this paper, we are interested in *symmetric* two-player games, where the two players have the same set of strategies $S_1 = S_2 = S$, and $p_1(s, s') = p_2(s', s)$ for any two strategies $s, s' \in S$. Whenever we refer to a “base game,” we implicitly mean it to be a symmetric two-player game.

Graphs

For the purposes of this paper, a *graph* $G = \langle V, E \rangle$ is an undirected graph with a set V of nodes and a set E of edges that has no isolated nodes. For a node u in a graph G , we write $\Gamma_G(u)$ to denote the set of u 's *neighbors*—that is, $\Gamma_G(u) := \{v : \{u, v\} \in E\}$. We omit the subscript when G is obvious from context. We write \mathcal{K}_n to denote the n -node complete graph, and $\mathcal{K}_{n,n}$ to denote the complete bipartite graph with n nodes in each “part” of the graph.

4 Playing Games on Graphs

Let M be an arbitrary symmetric two-player game where each player's set of available strategies is denoted by S , and when player 1 plays s_1 and player 2 plays s_2 then the payoff to player i is given by $p_i(s_1, s_2)$. Let $G = \langle V, E \rangle$ be an arbitrary graph. We define the *networked game* $M \oplus G$ (“game M played on graph G ”) as follows:

- The set of players in $M \oplus G$ is V , the set of nodes in the graph G .
- The set of strategies available to player $v \in V$ is S , the set of strategies for a player in M .
- For a strategy profile \mathbf{s} in which player v plays the strategy $s_v \in S$, the payoff to any player $w \in V$ is given by $p_w(\mathbf{s}) := \sum_{x \in \Gamma(w)} p_1(s_w, s_x)$. That is, the payoff to player w when she plays strategy s_w is the sum of the payoffs that she would receive if she played strategy s_w in the base game with each of her neighbors.

For example, if G is a two-node, one-edge graph (that is, the graph G is \mathcal{K}_2), then $M \oplus G$ is isomorphic to the base game M .

We can see a less trivial example through the classic game of rock–paper–scissors. Define the base game RPS with strategies {rock, paper, scissors} such that

$$p_1(\text{rock, scissors}) = p_1(\text{scissors, paper}) = p_1(\text{paper, rock}) = 1$$

and all other payoffs are zero. This game has no pure Nash equilibria. But in $\text{RPS} \oplus \mathcal{K}_3$, when this game is played on the triangle, we have a three-player game that does have a pure Nash equilibrium—namely, one where each of the three strategies is played by exactly one player.

5 Coordination Games

We begin by analyzing the networked version of coordination games. By way of reminder, our focus in this section (and indeed throughout the paper) is on the mathematical properties of these network games—particularly as those properties relate to their analogues in the base coordination game. In this section, we focus on price-of-anarchy and price-of-stability results.

A two-player *coordination game* (or *matching game*, or *consensus game*) is one in which both players choose from the same set of strategies, and receive positive payoffs only if they make the same selection. As usual, we are interested in symmetric coordination games, which we define as follows:

- Players share a strategy set $S = \{s_1, \dots, s_k\}$.
- The two players both receive identical payoffs of v_i if both choose s_i for any index $i \in \{1, \dots, k\}$, and both receive identical payoffs of 0 if they choose strategies s_i and $s_{j \neq i}$.

We refer to this game as the *k-strategy coordination game with payoffs* $\{v_1, \dots, v_k\}$. Without loss of generality, we always canonically order the strategies so that $v_1 \leq v_2 \leq \dots \leq v_{k-1} \leq v_k$. Throughout, we assume that the v_i 's are integer-valued.

Lemma 1. *Let M be an arbitrary k-strategy coordination game with payoffs $\{v_1, \dots, v_k\}$, and let G be an arbitrary graph. Then $\text{POS}(M) = \text{POS}(M \oplus G) = 1$.*

Proof. The social optimum for the base game M is when both players choose strategy s_k , yielding a payoff of v_k for each. This strategy profile is optimal because v_k is the largest entry anywhere in the payoff matrix for M , and it is a Nash equilibrium for the same reason. Similarly, in $M \oplus G$, the strategy profile in which all players play s_k is again a Nash equilibrium and socially optimal. Thus the socially optimal strategy profile is a Nash equilibrium in both M and $M \oplus G$, and $\text{POS}(M) = \text{POS}(M \oplus G) = 1$. \square

Lemma 2. *Let M be an arbitrary k-strategy coordination game with payoffs $\{v_1, \dots, v_k\}$. Then there exists a graph G such that $\text{POA}(M \oplus G) \geq v_k \cdot \sum_{i=1}^k \frac{1}{v_i}$.*

Proof. To begin, we define a few quantities:

$$p := \prod_{i=1}^k v_i \qquad n_i := \frac{p}{v_i} \qquad n := \sum_{i=1}^k n_i.$$

Note that n_i can also be written as $n_i = \prod_{j \neq i} v_j$; therefore the n_i 's are integral, and thus n is too.

To prove the lemma, we show that M played on the complete bipartite graph $\mathcal{K}_{n,n}$ has the desired price of anarchy; that is, $\text{POA}(M \oplus \mathcal{K}_{n,n}) \geq v_k \cdot \sum_{i=1}^k \frac{1}{v_i}$. We prove two claims:

there is a strategy profile of $M \oplus \mathcal{K}_{n,n}$ that has social welfare $2n^2 v_k$. (1)

there is a Nash equilibrium of $M \oplus \mathcal{K}_{n,n}$ that has social welfare $2pn$. (2)

This establishes the desired result, because

$$\text{POA} = \frac{\text{OPT}}{\text{WNE}} \geq \frac{2n^2 v_k}{2pn} = \frac{nv_k}{p} = \frac{v_k \sum_{i=1}^k \frac{p}{v_i}}{p} = v_k \sum_{i=1}^k \frac{1}{v_i}$$

by (1), (2), and the definitions of n and n_i , respectively.

To prove (1), suppose each player chooses strategy s_k , just as in Lemma 1. Each node in $\mathcal{K}_{n,n}$ thus matches the strategy chosen by all n of its neighbors, and therefore receives a payoff of nv_k . There are $2n$ nodes in the graph, and thus the social welfare under this strategy profile is $2n \cdot nv_k$, as claimed.

To prove (2), consider the strategy profile \mathbf{p} in which precisely n_i of the n nodes in each part of $\mathcal{K}_{n,n}$ play strategy s_i . Note that the payoff to a node u playing s_i is then $v_i n_i$, because u has precisely n_i neighbors matching u , namely the n_i nodes in the other part of the graph who are also playing s_i . Thus

$$\begin{aligned} & \text{social welfare under } \mathbf{p} \\ &= \sum_i 2n_i \cdot v_i n_i \quad (\text{there are } 2n_i \text{ nodes playing } s_i; \text{ each receives } v_i n_i \text{ as above}) \\ &= \sum_i 2n_i \cdot v_i \frac{p}{v_i} = \sum_i 2n_i p = 2pn. \quad (\text{definition of } n_i \text{ and } n) \end{aligned}$$

We further claim that \mathbf{p} is a Nash equilibrium. Indeed, consider any node u . The payoff to u for playing s_i is $v_i n_i = v_i \cdot (p/v_i) = p$, which is a constant independent of i . Thus every node u is indifferent among its strategies, and thus \mathbf{p} is a Nash equilibrium. \square

Lemma 3. *Let M be an arbitrary k -strategy coordination game with payoffs $\{v_1, \dots, v_k\}$, and let G be an arbitrary graph. Then $\text{POA}(M \oplus G) \leq v_k \cdot \sum_{i=1}^k \frac{1}{v_i}$.*

Proof. Fix an arbitrary Nash equilibrium strategy profile \mathbf{p} in $M \oplus G$, and write $\mathbf{p}(w)$ to denote the strategy played by a node w under \mathbf{p} . We write

$$\begin{aligned} \delta(u) &:= |\Gamma(u)| && (u\text{'s degree in } G) \\ \delta^i(u) &:= |\{w \in \Gamma(u) : \mathbf{p}(w) = s_i\}| && (\text{the number of } u\text{'s neighbors playing } s_i \text{ under } \mathbf{p}) \\ m &:= \sum_u \delta(u)/2. && (\text{the total number of edges in the graph } G) \end{aligned}$$

Write p_u to denote u 's payoff under \mathbf{p} . Note that u 's payoff from playing strategy s_i is precisely $v_i \delta^i(u)$. For any node u and any $i \in \{1, \dots, k\}$,

$$p_u \geq v_i \delta^i(u) \tag{3}$$

because \mathbf{p} is a Nash equilibrium. Dividing both sides of (3) by v_i and summing the resulting constraints over all strategies i , we have

$$\sum_i \frac{p_u}{v_i} \geq \sum_i \delta^i(u) \quad \iff \quad p_u \cdot \sum_i \frac{1}{v_i} \geq \delta(u) \quad \iff \quad p_u \geq \frac{\delta(u)}{\sum_i \frac{1}{v_i}}. \quad (4)$$

Therefore

$$\text{social welfare of } \mathbf{p} = \sum_u p_u \geq \sum_u \frac{\delta(u)}{\sum_i \frac{1}{v_i}} = 2m \cdot \frac{1}{\sum_i \frac{1}{v_i}}, \quad (5)$$

where the inequality follows by (4). Note that the social optimum is achieved when all players choose strategy s_k , so we have that the socially optimal outcome achieves the social welfare

$$\text{OPT} = 2m \cdot v_k. \quad (6)$$

Thus

$$\frac{\text{OPT}}{\text{social welfare of } \mathbf{p}} \leq \frac{2m \cdot v_k}{2m \cdot \frac{1}{\sum_i \frac{1}{v_i}}} = v_k \sum_i \frac{1}{v_i}$$

by (5) and (6). Because \mathbf{p} was an arbitrary Nash equilibrium, we have that $\text{OPT}/\text{WNE} \leq v_k \cdot \sum_i \frac{1}{v_i}$ for the worst Nash equilibrium WNE, and thus the price of anarchy of $M \oplus G$ has the same upper bound. \square

These lemmas establish the following. Let M be an arbitrary k -strategy coordination game with payoffs $\{v_1, \dots, v_k\}$. For any graph G , we have $\text{POS}(M \oplus G) = 1$ and $\text{POA}(M \oplus G) \leq v_k \cdot \sum_i \frac{1}{v_i}$; furthermore, there exists a graph G^* for which $\text{POA}(M \oplus G^*) = v_k \cdot \sum_i \frac{1}{v_i}$.

Theorem 4 (Price of anarchy in coordination games). *Let M be any k -strategy coordination game with payoffs $\{v_1, \dots, v_k\}$, for any $k \geq 2$. Then*

$$\text{POA}(M) + k - 1 \leq \max_G \text{POA}(M \oplus G) \leq k \cdot \text{POA}(M).$$

Furthermore, for any k , there are k -strategy coordination games in which both bounds are tight.

Proof. Note that by Lemma 2 and Lemma 3, we have that $\max_G \text{POA}(M \oplus G) = v_k \cdot \sum_i \frac{1}{v_i}$. Let G^* denote the graph achieving this maximum price of anarchy. Also notice that $\text{POA}(M) = v_k/v_1$: in the base game, the only Nash equilibria are when the two players are indeed playing the same strategy, and the worst such equilibrium is when both players play s_1 .

For the lower bound, $\text{POA}(M \oplus G^*) = v_k \cdot \sum_i \frac{1}{v_i} \geq v_k \cdot (\frac{k-1}{v_k} + \frac{1}{v_1}) = k - 1 + \frac{v_k}{v_1} = k - 1 + \text{POA}(M)$, because $1/v_k \leq 1/v_i$ for all i . This bound is tight in the k -strategy coordination game with payoffs $\{1, v, v, \dots, v\}$, where $\text{POA}(M) = v/1 = v$ and $\text{POA}(M \oplus G^*) = k - 1 + v$.

For the upper bound, $\text{POA}(M \oplus G^*) = v_k \cdot \sum_i \frac{1}{v_i} \leq v_k \cdot \frac{k}{v_1} = k \cdot \text{POA}(M)$, because $1/v_i \leq 1/v_1$ for all i . This bound is tight in the k -strategy coordination game M with payoffs $\{1, 1, \dots, 1\}$, where we have $\text{POA}(M) = 1$ while $\text{POA}(M \oplus G^*) = k$. \square

As we stated in Section 1, in this paper we restrict our attention to pure Nash equilibria. In many of the scenarios that motivate our work, the choices participants make (e.g., which operating system to buy, or which cell phone provider to use) are costly to alter, and it is hard to imagine these decisions literally being made by coin tosses by every member of a large population.

That said, it is worth noting that the constructed pure Nash equilibrium on $M \oplus \mathcal{K}_{n,n}$ of Lemma 2 is closely related to a mixed Nash equilibrium in the base game M . In particular, we can associate Player 1 in M with all players on the left side of $\mathcal{K}_{n,n}$ in $M \oplus \mathcal{K}_{n,n}$, and Player 2 with all players on the right side. Given any pure Nash equilibrium in $M \oplus \mathcal{K}_{n,n}$, the corresponding mixed strategy for a player in M is a weighted average of her associated players' strategies. Similarly, any mixed Nash equilibrium on M can be converted into a pure Nash equilibrium in a networked game on a sufficiently large complete bipartite graph (assuming probabilities in the mixed Nash equilibrium are rational). See Lemma 11 for the general, formal version of this statement.

This correspondence implies the following surprising fact: Lemma 2 and Lemma 3 also relate the pure and mixed prices of anarchy in (two-player) coordination games. Of course, Lemma 3 is a more general result, in that it bounds the effect of playing matching games on *arbitrary* networks, not just complete bipartite graphs, and pure Nash equilibria in $M \oplus G$ for general networks G do not correspond (at least not obviously) to mixed Nash equilibria in M . Nevertheless, as part of our future work we intend to further explore the relationship between mixed Nash equilibria in base games and pure Nash equilibria in their networked counterparts.

6 Existence of Pure Nash Equilibria

In this section, we explore the relationship between the existence of pure Nash equilibria in an arbitrary base game M and the existence of pure Nash equilibria in the networked game $M \oplus G$ played on a graph G . Broadly speaking, we show negative results: one cannot infer anything, in either direction, about the existence of pure Nash equilibria in M and the existence of pure Nash equilibria in $M \oplus G$. The limited exception is in the case of a bipartite graph G .

We will highlight our main results first, and then devote the rest of the section to more detailed statements of the results and proofs.

Theorem 5 (bipartite results). *Let G be an arbitrary bipartite graph. Then:*

1. *for a base game M with a pure Nash equilibrium, $M \oplus G$ has a pure Nash equilibrium.*
2. *for a base game M with no pure Nash equilibrium, $M \oplus G$ may or may not have a pure Nash equilibrium. (In particular, for every base game M , there exists a bipartite graph G_M such that $M \oplus G_M$ does have a pure Nash equilibrium.)*

Theorem 6 (nonbipartite results). *Let G be an arbitrary nonbipartite graph. Then, regardless of whether the base game M has a pure Nash equilibrium, $M \oplus G$ may or may not have a pure Nash equilibrium. In particular:*

1. *for any nonbipartite graph G , there is a base game M_G (we can choose whether or not M_G has a pure Nash equilibrium) such that $M_G \oplus G$ does not have a pure Nash equilibrium.*
2. *as in Theorem 5.2, for every base game M (with or without a pure Nash equilibrium), there exists a nonbipartite graph G_M such that $M \oplus G_M$ does have a pure Nash equilibrium.*

Theorem 5.1 is the lone positive result, showing that a pure Nash equilibrium in a base game does translate into a pure Nash equilibrium when that game is played on a bipartite network. The other results are all negative, in that they show that conclusions about the existence of pure Nash equilibria in $M \oplus G$ in terms of the existence of pure Nash equilibria in M are not generally possible.

Of these results, we will specifically highlight the result mentioned in Theorem 5.2 and Theorem 6.2, which is proven in Lemma 11. This result draws an unexpected connection between mixed Nash equilibria in a base game M and pure Nash equilibria in the networked version of M , played on an appropriately chosen complete bipartite or tripartite graph.

Best-Response Dynamics

We also further explore the case when there are pure Nash equilibria in M , by subdividing this case based on whether *best-response dynamics* (*BRD*) always converges. BRD is an algorithm that produces a sequence of strategy profiles by repeatedly allowing a player i to update her strategy to a *best response*, i.e., a strategy in S_i maximizing i 's payoff, holding all other strategies constant. Specifically, BRD is the following algorithm:

1. Start from an arbitrary strategy profile \mathbf{s} .
2. While there is some player who is not playing a best response in \mathbf{s} :
 - (a) Choose such a player i arbitrarily.
 - (b) Update \mathbf{s} by changing s_i to be an arbitrary best response for player i .

If BRD terminates, then it terminates at a pure Nash equilibrium; however, BRD may not terminate even if a pure Nash equilibrium exists. We say that BRD *always converges* if this process terminates regardless of the arbitrary choices (the initial profile \mathbf{s} , which player i updates in each iteration, which best response s_i is chosen if there are more than one), subject to the liveness condition that every player “gets a turn” infinitely often.

Theorem 7 (best-response dynamics results). *For an arbitrary base game M and graph G :*

1. *if BRD always converges in M , then BRD may or may not always converge in $M \oplus G$ (regardless of whether G is bipartite or nonbipartite).*
2. *conversely, if BRD always converges in $M \oplus G$, then*
 - (a) *if G is bipartite, then BRD always converges in M ; but*
 - (b) *if G is nonbipartite, then BRD may or may not always converge in M .*

As before, the lone case in which one can infer anything is Theorem 7.2a: if BRD always converges in the networked game on a bipartite graph, then BRD must always converge on the base game. But all the other results are negative.

These results—and the ones from Theorem 5 and Theorem 6—are summarized in full detail, including complete quantification of each result with respect to graphs and base games, in Figure 1.

Rock–Paper–Scissors

In several of our constructions, we make use of the following k -strategy “rock–paper–scissors” game, which we denote RPS_k . There are $k \geq 2$ strategies $\{s_0, \dots, s_{k-1}\}$, where strategy s_i “beats” the

	no pure Nash equilibrium exists in M	a pure Nash equilibrium exists in M	
		BRD does not converge in M	BRD converges in M
bipartite graph G	$\forall G \exists M :$ $M \oplus G$ has no pure NE (Lemma 8) $\exists G \forall M :$ $M \oplus G$ has no pure NE (bipartite: Lemma 13) (nonbipartite: Lemma 14)	$\forall G \forall M : M \oplus G$ has a pure NE (Lemma 9) $\forall G \forall M :$ BRD does not converge in $M \oplus G$ (Lemma 15)	$\exists M \forall G :$ BRD converges in $M \oplus G$ (Lemma 12) $\exists G \forall M :$ BRD converges in $M \oplus G$ (Lemma 13) $\exists G \exists M :$ BRD does not converge in $M \oplus G$ (Lemma 16)
nonbipartite graph G	$\forall M \exists G :$ $M \oplus G$ has a pure NE (Lemma 11)	$\forall G \exists M : M \oplus G$ has no pure NE (Lemma 10) $\forall M \exists G : M \oplus G$ has a pure NE (Lemma 11) $\forall G \exists M :$ BRD does not converge on $M \oplus G$ (Lemma 10) $\exists G \forall M :$ BRD does not converge on $M \oplus G$ (Lemma 14) $\exists G \exists M :$ BRD converges on $M \oplus G$ (Lemma 17)	$\forall G \exists M :$ BRD does not converge on $M \oplus G$ (Lemma 10) $\exists M \forall G :$ BRD converges in $M \oplus G$ (Lemma 12)

Figure 1: Summary of our results relating the existence of pure Nash equilibria and the convergence of best-response dynamics in base games M and networked games $M \oplus G$. When we write “BRD converges” we mean that best-response dynamics *always* converges, from any starting strategy profile and following any sequence of best-response updates. Similarly, “BRD does not converge” means that BRD *sometimes* fails to converge.

strategy s_{i+1} . Formally, we define

$$\text{payoff to player playing } s_i \text{ against } s_j = \begin{cases} 1 & \text{if } j = (i + 1) \bmod k \\ 0 & \text{otherwise.} \end{cases}$$

For example:

$$\text{RPS}_5 = \begin{array}{c|ccccc} & s_0 & s_1 & s_2 & s_3 & s_4 \\ \hline s_0 & 0,0 & 1,0 & 0,0 & 0,0 & 0,1 \\ s_1 & 0,1 & 0,0 & 1,0 & 0,0 & 0,0 \\ s_2 & 0,0 & 0,1 & 0,0 & 1,0 & 0,0 \\ s_3 & 0,0 & 0,0 & 0,1 & 0,0 & 1,0 \\ s_4 & 1,0 & 0,0 & 0,0 & 0,1 & 0,0 \end{array}.$$

We will write $\alpha \cdot \text{RPS}_k$, for a scalar $\alpha > 0$, to denote the version of this game in which the payoff for playing a strategy s_i against $s_{(i+1) \bmod k}$ is α rather than 1. All other payoffs remain zero.

Note that RPS_k does not have a pure Nash equilibrium for any $k \geq 3$: the only best response to strategy s_i is strategy $s_{(i-1) \bmod k}$, but the only best response to $s_{(i-1) \bmod k}$ is $s_{(i-2) \bmod k}$. Thus no two strategies are mutual best responses because $i \neq (i-2) \bmod k$ for any $k \geq 3$, and thus no pure Nash equilibrium exists in RPS_k . (For $k = 2$, we have $i = (i-2) \bmod k$, and the strategies s_0 and s_1 are mutual best responses. In fact, RPS_2 is a simple “mismatching” game, with pure Nash equilibria $\langle s_0, s_1 \rangle$ and $\langle s_1, s_0 \rangle$.)

Detailed Results and Proofs

Lemma 8. *For every graph G , there exists a base game M with no pure Nash equilibria such that the networked game $M \oplus G$ also does not have a pure Nash equilibrium.*

Proof. It suffices to choose $M = \text{RPS}_{n+1}$, where n is the number of nodes in G .

Suppose for a contradiction that a strategy profile \mathbf{p} is a pure Nash equilibrium for $\text{RPS}_{n+1} \oplus G$. There are $n + 1$ strategies and only n nodes, so there is at least one strategy that is not being played under \mathbf{p} ; consequently, there is an index i and node u such that strategy s_i is played by u under \mathbf{p} , but strategy $s_{(i+1) \bmod (n+1)}$ is not being played by any node under \mathbf{p} .

The structure of RPS_{n+1} implies that the payoff p_u to node u is precisely the number of u 's neighbors in G that play strategy $s_{(i+1) \bmod (n+1)}$ under \mathbf{p} ; thus $p_u = 0$. But the node u has at least one neighbor playing some strategy s_j , and thus u can deviate from s_i to $s_{(j-1) \bmod (n+1)}$ to receive a strictly positive payoff. Thus \mathbf{p} is not a Nash equilibrium, contradicting the assumption. \square

The proof of Lemma 8 uses a base game M with no pure Nash equilibria to show that the networked game $M \oplus G$ need not have a pure Nash equilibrium. If we restrict the base game M to contain a pure Nash equilibrium, then will the same result hold? To answer this question, we have to consider the structure of G . In particular, if G is nonbipartite, there still exists an M such that $M \oplus G$ has no pure Nash equilibrium, even if M itself does; however, if G is bipartite, this is no longer the case.

Lemma 9. *For every bipartite graph G and for every base game M that has a pure Nash equilibrium, the game $M \oplus G$ has a pure Nash equilibrium too.*

Because G is nonbipartite, the graph is not 2-colorable. Interpreting the “strategy types” (a ’s and b ’s) under \mathbf{p} as two colors, we know that there must be two adjacent nodes x and y so that x and y are both “colored” a or both “colored” b by the profile \mathbf{p} . Without loss of generality, suppose that x and y are both colored a . Let A denote the set of a -colored nodes that are connected to x and y in G —that is, let A denote the connected component of the subgraph of G induced by considering only a -colored nodes. Then every node in A plays an a -strategy under \mathbf{p} , and is adjacent to at least one other node in A playing an a -strategy under \mathbf{p} . But $|A| \leq n$; therefore, just as in Lemma 8, by the pigeonhole principle, there must be some node $u \in A$ playing strategy a_i where no node in A plays strategy $a_{(i+1) \bmod (n+1)}$. Therefore $p_u^a = 0$. But that node u is adjacent to some node in S playing a strategy a_j . If u switches to play $a_{(j-1) \bmod (n+1)}$, then u ’s payoff increases: he still receives p_u^b payoff from his b -colored neighbors, and he receives at least $1 - \varepsilon$ from his a_j neighbor. This deviation increases u ’s payoff, and thus contradicts the assumption that \mathbf{p} was a Nash equilibrium.

We have now described a base game M_n such that best-response dynamics always converges in M_n (and thus a pure Nash equilibrium exists in M_n), yet there is no pure Nash equilibrium in $M_n \oplus G$. To prove the second part of the lemma, we must again give a base game M'_n such that M'_n has a pure Nash equilibrium but $M'_n \oplus G$ does not, but this time where BRD does not always converge in M'_n . We construct M'_n by augmenting M_n to block BRD from always converging. We add two strategies c_0 and c_1 , where playing c_i against any non- c strategy has a highly negative payoff, and the c -versus- c sub matrix is a version of the classic Matching Pennies game (a two-player, two-strategy game that has no pure Nash equilibria):

$$M'_n := \begin{array}{c} \begin{array}{ccccc} & s_0 & \dots & b_n & c_0 & c_1 \\ \begin{array}{c} a_0 \\ \vdots \\ b_n \\ c_0 \\ c_1 \end{array} & \begin{array}{|c|} \hline M_n \\ \hline \end{array} & & & \begin{array}{|c|} \hline 0, -n \\ \hline \end{array} & \begin{array}{|c|} \hline 0, -n \\ \hline \end{array} \\ & & & & \begin{array}{|c|} \hline \vdots \\ \hline \end{array} & \begin{array}{|c|} \hline \vdots \\ \hline \end{array} \\ & & & & \begin{array}{|c|} \hline 0, -n \\ \hline \end{array} & \begin{array}{|c|} \hline 0, -n \\ \hline \end{array} \\ \hline & \begin{array}{|c|} \hline -n, 0 \\ \hline \end{array} & \begin{array}{|c|} \hline \dots \\ \hline \end{array} & \begin{array}{|c|} \hline -n, 0 \\ \hline \end{array} & \begin{array}{|c|} \hline 1, -n \\ \hline \end{array} & \begin{array}{|c|} \hline -n, 1 \\ \hline \end{array} \\ \hline & \begin{array}{|c|} \hline -n, 0 \\ \hline \end{array} & \begin{array}{|c|} \hline \dots \\ \hline \end{array} & \begin{array}{|c|} \hline -n, 0 \\ \hline \end{array} & \begin{array}{|c|} \hline -n, 1 \\ \hline \end{array} & \begin{array}{|c|} \hline 1, -n \\ \hline \end{array} \end{array} \end{array}.$$

First observe that, like M_n , the base game M'_n contains a pure Nash equilibrium, namely any strategy profile $\langle a_i, b_j \rangle$ or $\langle b_i, a_j \rangle$. However, BRD does not always converge in M'_n : if the two players begin at the strategy profile $\langle c_0, c_0 \rangle$, then after four best-response updates (player 2 updates to c_1 ; player 1 updates to c_1 ; player 2 updates to c_0 ; player 1 updates to c_0) we have looped back to the original strategy profile.

Now we argue that for the nonbipartite graph G , there is no pure Nash equilibrium in $M'_n \oplus G$. There can be no pure Nash equilibrium in $M'_n \oplus G$ in which any node uses a c -strategy: some c -playing node must receive a payoff of $-n$ along some edge, which means his total payoff must be negative (because he gains at most a payoff of one from every other neighbor), and switching to strategy a_0 , say, guarantees a nonnegative payoff. And there is no pure Nash equilibrium involving only a - and b -strategies, as above. Thus there is no pure Nash equilibrium in $M'_n \oplus G$. \square

In Lemmas 9 and 10 we see how base games with pure Nash equilibria can lead to networked games with and without pure Nash equilibria, depending on the underlying graph structure. Lemma 8 shows a specific base game with no pure Nash equilibria that leads to a networked game

with no pure Nash equilibria. But in Lemma 8, the base game M was chosen with respect to a particular graph G . Here we show that this dependence was crucial: for every base game M , there exists a graph G so that the networked game $M \oplus G$ has a pure Nash equilibrium. Our proof again highlights some of the connections between pure Nash equilibria in networked games and mixed Nash equilibria in base games:

Lemma 11. *For every base game M (independent of whether pure Nash equilibria exist in M or whether best-response dynamics always converges in M), there exist a bipartite graph G and a nonbipartite graph G' such that pure Nash equilibria exist in both $M \oplus G$ and $M \oplus G'$.*

Proof. Let S denote the strategy set of M , and let p denote the payoff function to a player in M . A classic result of Nash [31] says that the symmetric two-player game M must have a symmetric mixed Nash equilibrium—that is, a probability distribution σ over S that is a best response to itself. More formally, the expected payoff to a player for playing any strategy in the support of σ against σ must be maximum:

$$\sigma(s) > 0 \quad \implies \quad \forall s' \in S : \sum_{t \in S} \sigma(t) \cdot p(s, t) \geq \sum_{t \in S} \sigma(t) \cdot p(s', t). \quad (7)$$

Furthermore, because we have a two-player base game, the probability vector σ contains only rational probabilities, so that $\sigma(s) = q_s/n$, where $\sum_s q_s = n$, for a common denominator $n \in \mathbb{Z}^{\geq 1}$ and coefficients $q_s \in \mathbb{Z}^{\geq 1}$. (For a more detailed derivation of the above results, see Papadimitriou [32].)

We can now define the graphs G and G' as the complete bipartite graph $G := \mathcal{K}_{n,n}$ and the complete tripartite graph $G' := \mathcal{K}_{n,n,n}$ with n nodes in each part of each graph. We claim that the following \mathbf{p} is a pure Nash equilibrium in $M \oplus \mathcal{K}_{n,n}$: precisely q_s nodes in each part of the graph play strategy s . Specifically, every node u is playing a best response:

$$\begin{aligned} & u\text{'s payoff from playing } s \\ &= \sum_{w \in \Gamma(u)} p(s, \mathbf{p}(w)) \\ &= \sum_{t \in S} n \cdot \sigma(t) \cdot p(s, t) \quad (u \text{ has precisely } n \text{ neighbors, of which } q_t = n \cdot \sigma(t) \text{ are playing strategy } t) \\ &\geq n \cdot \sum_{t \in S} \sigma(t) \cdot p(s', t) \quad \text{by (7)} \\ &= \sum_{w \in \Gamma(u)} p(s', \mathbf{p}(w)) \\ &= u\text{'s payoff from deviating to } s' \end{aligned}$$

where we can apply (7) because u playing s implies that $\sigma(s) > 0$. Thus \mathbf{p} is a pure Nash equilibrium in $M \oplus \mathcal{K}_{n,n}$. The proof that there is a pure Nash equilibrium in $M \oplus \mathcal{K}_{n,n,n}$ is strictly analogous, where now every node has $2n \cdot \sigma(t)$ neighbors playing strategy t instead of just $n \cdot \sigma(t)$. \square

With these statements about mismatches between base games and networked games proven, we turn to a few simple examples of graphs and base games in which properties of the base game do align with properties of the corresponding networked games.

Lemma 12. *There exists a base game M in which best-response dynamics always converges (and thus M has a pure Nash equilibrium) such that, for every graph G , best-response dynamics always converges in the networked game $M \oplus G$ (which thus also has a pure Nash equilibrium).*

Proof. Let M be a trivial game in which all payoffs are zero. Then every strategy profile is a pure Nash equilibrium in M and in $M \oplus G$, and thus BRD always converges in both. \square

Lemma 13. *There exists a bipartite graph G such that, for every base game M , the existence of pure Nash equilibria and the convergence of best-response dynamics are identical in the games M and $M \oplus G$.*

Proof. Let $G = \mathcal{K}_2$ be the trivial 2-node, 1-edge graph. Then the games M and $M \oplus G$ are in fact identical, and the claim follows immediately. (The same result holds when G is an arbitrary matching, where multiple independent parallel copies of M are played.) \square

Lemma 14. *There exists a nonbipartite graph G such that, for every base game M , a pure Nash equilibrium exists in $M \oplus G$ only if a pure Nash equilibrium exists in M and, furthermore, BRD always converges in $M \oplus G$ only if BRD always converges in M .*

Proof. Let G be a 5-node, 2-component graph containing \mathcal{K}_2 and \mathcal{K}_3 as its two connected components; the graph is not bipartite because of the \mathcal{K}_3 component. If there were a pure Nash equilibrium in $M \oplus G$ then the two nodes in the \mathcal{K}_2 component would have to be playing mutual best responses, and thus their two strategies would form a pure Nash equilibrium in M . Similarly, if there is a sequence σ of best-response updates in M that does not terminate, then BRD does not converge in $M \oplus G$ when σ is the sequence of best-response updates in the \mathcal{K}_2 component. \square

Finally, we conclude this section with three results abouts best-response dynamics. We show that in bipartite graphs, if BRD does not always converge in M , then BRD does not always converge in $M \oplus G$. (This result is analogous to Lemma 9, where we showed the existence of pure Nash equilibria in a base game carried over to networked games on bipartite graphs.) We also give two more examples of mismatches in properties between base games and networked games: a base game in which BRD always converges but for which BRD does not always converge in the networked game on a bipartite graph; and a base game in which BRD does not always converge but for which BRD does always converge in the networked game on a nonbipartite graph.

Lemma 15. *For every bipartite graph G and every base game M in which BRD does not always converge, BRD also does not always converge in the networked game $M \oplus G$ (even if the players' best-response updates are done in round-robin order).*

Proof. Suppose there is a non-converging sequence of best-response strategy choices for the base game M . Denote the sequence of best responses chosen as $\langle s_1, s_2, s_3, \dots \rangle$, where player 1's updates are denoted by odd subscripts and player 2's with even subscripts, where each of these strategies is a best response to the previous strategy.

Fix an arbitrary bipartite graph $G = \langle L \cup R, E \rangle$, and consider the following run of BRD on the networked game $M \oplus G$. We start from the strategy profile \mathbf{p} in which every node in L plays strategy s_1 and every node in R plays strategy s_2 . Players update their strategies in a round-robin order such that all players in L update before all players in R .

Because G is bipartite, when a node $\ell \in L$ updates his strategy, any updates made by other nodes in L do not affect ℓ 's best response. Because strategy s_{i+1} is a best response to strategy s_i in M , by the definition BRD on M , it must also be a best response when playing against many opponents, all of whom are playing s_i . Starting from \mathbf{p} , where L and R nodes are playing s_1 and s_2 , respectively, it is thus a best response for each node $\ell \in L$ to update to s_3 . Now each node in R is adjacent only to nodes playing s_3 , and thus, similarly, it is a best response for each node $r \in R$ to update to s_4 , and so forth. Because there is a non-converging sequence of best responses in M (namely $\langle s_1, s_2, s_3, \dots \rangle$), then there is also a non-converging sequence of best responses in $M \oplus G$ (also $\langle s_1, s_2, s_3, \dots \rangle$), so BRD does not always converge in the networked game. \square

Lemma 16. *There exist a bipartite graph G and a base game M such that best-response dynamics always converges in M but best-response dynamics does not always converge in $M \oplus G$ (even if the players' best-response updates are done in round-robin order).*

Proof. Let the graph $G = \mathcal{K}_{2,2}$ be the complete bipartite graph with nodes $L = \{w, x\}$ and $R = \{y, z\}$ and edges $L \times R$. Define the base game M as follows:

	s_1	s_2	s_3	s_4	s_5	s_6
s_1	3, 3	0, 0	0, 2	0, 2	2, 0	2, 0
s_2	0, 0	3, 3	0, 2	0, 2	2, 0	2, 0
s_3	2, 0	2, 0	3, 3	0, 0	0, 2	0, 2
s_4	2, 0	2, 0	0, 0	3, 3	0, 2	0, 2
s_5	0, 2	0, 2	2, 0	2, 0	3, 3	0, 0
s_6	0, 2	0, 2	2, 0	2, 0	0, 0	3, 3

Intuitively, M is a version of $2 \cdot \text{RPS}_3$ in which there are two strategies of each type (s_1 and s_2 are both “rock”; s_3 and s_4 are both “paper”; s_5 and s_6 are both “scissors”), and players receive a payoff of 3 if they choose the same strategy. The key for BRD not always converging in $\mathcal{K}_{2,2}$ is the following. Suppose that both nodes of each “side” of the graph always play the same of rock, paper, or scissors, but always differ in which of the strategies within that category they play. Then each “losing” node has two neighbors and can gain $2+2$ from winning the two “RPS-type” games against the other side of the graph, which outweighs the 3 from matching one of the other side’s strategies; both nodes will switch to “winning” and the process will continue for the other pair. However, in the non-networked version of M , BRD will always converge, because when there is only one other player then it is better to match that player (for a payoff of 3) than to beat him in RPS (for a payoff of 2).

More formally, we claim that BRD always converges—in zero or one best-response updates, in fact—in the base game M . This is easy to see because in any strategy profile $\langle s_i, s_{j \neq i} \rangle$, either player’s best response is to match the strategy played by the other, yielding a payoff of 3 for both players. Any strategy profile $\langle s_i, s_i \rangle$ is a pure Nash equilibrium, and thus BRD terminates after this step.

Now we claim that BRD does not always converge in $M \oplus G$. Suppose that we start with the strategy profile $\langle s_1, s_2, s_3, s_4 \rangle$ for nodes $\langle w, x, y, z \rangle$, respectively, and updates proceed alphabetically,

in round-robin fashion. Then the strategy profiles after each step of BRD are as follows:

		explanation of updating node u 's update							current strategy			
u	$\Gamma(u)$'s strategies	payoff to u for ...						u 's best response	w	x	y	z
		s_1	s_2	s_3	s_4	s_5	s_6					
<i>initial strategy profile</i>									s_1	s_2	s_3	s_4
w	s_3, s_4	0	0	3	3	4	4	s_5	s_5	s_2	s_3	s_4
x	s_3, s_4	0	0	3	3	4	4	s_6	s_5	s_6	s_3	s_4
y	s_5, s_6	4	4	0	0	3	3	s_1	s_5	s_6	s_1	s_4
z	s_5, s_6	4	4	0	0	3	3	s_2	s_5	s_6	s_1	s_2
w	s_1, s_2	3	3	4	4	0	0	s_3	s_3	s_6	s_1	s_2
x	s_1, s_2	3	3	4	4	0	0	s_4	s_3	s_4	s_1	s_2
y	s_3, s_4	0	0	3	3	4	4	s_5	s_3	s_4	s_5	s_2
z	s_3, s_4	0	0	3	3	4	4	s_6	s_3	s_4	s_5	s_6
w	s_5, s_6	4	4	0	0	3	3	s_1	s_1	s_4	s_5	s_6
x	s_5, s_6	4	4	0	0	3	3	s_2	s_1	s_2	s_5	s_6
y	s_1, s_2	3	3	4	4	0	0	s_3	s_1	s_2	s_3	s_6
z	s_1, s_2	3	3	4	4	0	0	s_4	s_1	s_2	s_3	s_4

Thus we have returned to the same strategy profile as the initial configuration, and BRD is stuck in a loop. Therefore BRD does not always converge in $M \oplus G$. \square

Lemma 17. *There exist a nonbipartite graph G and a base game M with a pure Nash equilibrium but for which best-response dynamics does not always converge in M such that best-response dynamics does always converge on $M \oplus G$.*

Proof. The graph $G = \mathcal{K}_3$ is the triangle, and the base game M is an augmented form of rock–paper–scissors altered with a major tie deterrent and an additional matching-type strategy:

	R	P	S	X
R	-4, -4	0, 1	1, 0	0, 0
P	1, 0	-4, -4	0, 1	0, 0
S	0, 1	1, 0	-4, -4	0, 0
X	0, 0	0, 0	0, 0	4, 4

Note that a pure Nash equilibrium does exist in M , namely the profile $\langle X, X \rangle$. However, best-response dynamics does not always converge in M : if we begin in the strategy profile $\langle R, P \rangle$, then BRD ends up in a cycle (specifically, $RP \rightarrow SP \rightarrow SR \rightarrow PR \rightarrow PS \rightarrow RS \rightarrow RP$). Now we must argue that BRD always converges in $M \oplus G$, which we will do by examining players' best responses depending on the number of the three nodes in the graph who are playing strategy X .

BRD terminates immediately if we start it from the profile $\langle X, X, X \rangle$.

If we start BRD from any profile in which exactly two players have chosen X , then those two players are already playing a best response. As soon as the non- X player has a chance to make an update, her best response is to switch to X as well. BRD then terminates as above.

If we start BRD from any profile in which exactly zero players have chosen X , then the best response for any player is to choose an RPS strategy that does not duplicate any other player's

choice. Therefore, as soon as all three players have gotten a chance to update, BRD will have as its current strategy profile some permutation of $\langle R, P, S \rangle$. This profile is a pure Nash equilibrium, and thus BRD terminates.

Finally, if we start BRD from any profile in which exactly one player has chosen X , then one of two cases holds. One possibility is that the X -player makes the first update: in this case, just as above, his best response is to choose an RPS strategy that does not duplicate any other player's choice. BRD then terminates as in the zero- X -player case. The other possibility is that a RPS-player makes the first update: in this case, her best response is to play X . This update leads to a profile with two players who have chosen X , and, again as above, BRD will always terminate from such a profile. \square

7 Future Directions

Our work points to a substantial number of interesting open questions. We are currently pursuing a number of these directions, which we briefly highlight here.

First, we would like to extend our price-of-anarchy analysis of coordination games to other or more general base games. What properties must a base game possess for similar bounds to hold?

Second, we hope to develop stronger results based on particular structural aspects of the underlying network. A number of our results in this paper distinguish between graphs that are bipartite and those that are not. Could stronger results be found if we restrict ourselves to other classes of graphs, such as regular graphs, trees, or grids? Or, more ambitiously, could we perhaps begin to understand games played on networks that are explicitly intended to represent social structures [3, 27, 36, e.g.]?

Third, in Section 5 and in Lemma 11, we noted a correspondence between mixed Nash equilibria in base games and pure Nash equilibria in their networked counterpart (for a complete k -partite network). Can a more general correspondence be found for arbitrary networks? And could this correspondence further our understanding of mixed Nash equilibria in general?

Fourth, what kind of general algorithmic results are possible in the context of networked games? For example, can we give an efficient algorithm that, given a graph G as input, computes the price of anarchy for a simple coordination game when it is played on the graph G ? Under what circumstances can one find pure Nash equilibria efficiently in networked games? The connection to graphical games may be helpful here, though there are differences. (For example, even for a complete graph G , the game $M \oplus G$ is still compactly representable; graphical games where the underlying graph of direct influence is complete require an exponential-sized description.) Because our $M \oplus G$ networked games are more restrictive than graphical games, we might hope for efficient algorithms in a broader context than what is known for graphical games. An encouraging sign for this research direction is the recent work by Daskalakis and Papadimitriou [8] on reducing a networked game $Z \oplus G$ to a two-player zero-sum game when the base game Z is zero sum (or even strictly competitive).

Finally, the framework we put forth in this work specifies how a large, complex game can be generated via the composition of a simple base game and a network. To what extent might we be able to perform the reverse operation: take a complex game and decompose it (even if only approximately) into a network and one or more simple base games?

References

- [1] W. Brian Arthur. Inductive reasoning and bounded rationality: The El Farol problem. *American Economic Review*, 84:406–411, 1994.
- [2] Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Improved equilibria via public service advertising. In *Proceedings of the Symposium on Discrete Algorithms*, pages 728–737, 2009.
- [3] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.
- [4] Oren Ben-Zwi and Amir Ronen. The local and global price of anarchy of graphical games. In *Symposium on Algorithmic Game Theory*, pages 255–266, 2008.
- [5] George Christodoulou, Vahab Mirrokni, and Anastasios Sidiropoulos. Convergence and approximation in potential games. In *Proceedings of the Symposium on Theoretical Aspects of Computer Science*, pages 349–360, 2006.
- [6] Michael Chwe. Communication and coordination in social networks. *Review of Economic Studies*, 67(1):1–16, January 2000.
- [7] Constantinos Daskalakis and Christos H. Papadimitriou. Computing pure Nash equilibria in graphical games via Markov Random Fields. In *Proceedings of the Annual Conference on Electronic Commerce (EC)*, pages 91–99, 2006.
- [8] Constantinos Daskalakis and Christos H. Papadimitriou. On a network generalization of the minmax theorem. In *International Colloquium on Automata, Languages and Programming (ICALP)*, pages 423–434, 2009.
- [9] Bistra Dilkina, Carla P. Gomes, and Ashish Sabharwal. The impact of network topology on pure Nash equilibria in graphical games. In *AAAI Conference on Artificial Intelligence*, pages 42–49, 2007.
- [10] Edith Elkind, Leslie Ann Goldberg, and Paul W. Goldberg. Nash equilibria in graphical games on trees revisited. In *Proceedings of the Annual Conference on Electronic Commerce (EC)*, pages 100–109, 2006.
- [11] Edith Elkind, Leslie Ann Goldberg, and Paul W. Goldberg. Computing good Nash equilibria in graphical games. In *Proceedings of the Annual Conference on Electronic Commerce (EC)*, pages 162–171, 2007.
- [12] Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar. The complexity of pure Nash equilibria. In *Proceedings of the Annual Symposium on the Theory of Computing (STOC)*, pages 604–612, 2004.
- [13] Drew Fudenberg and David K. Levine. *The Theory of Learning in Games*. MIT Press, 1998.
- [14] Andrea Galeotti, Sanjeev Goyal, Matthew Jackson, Fernando Vega-Redondo, and Leeat Yariv. Network games. *Review of Economic Studies*, 77(1):218–244, 2010.

- [15] Georg Gottlob, Gianluigi Greco, and Francesco Scarcello. Pure Nash equilibria: hard and easy games. *Journal of Artificial Intelligence Research*, 24:357–406, 2005.
- [16] Sanjeev Goyal. *Connections: An Introduction to the Economics of Networks*. Princeton University Press, 2007.
- [17] Nicole Immorlica, Jon Kleinberg, Mohammad Mahdian, and Tom Wexler. The role of compatibility in the diffusion of technologies through social networks. In *Proceedings of the Annual Conference on Electronic Commerce (EC)*, pages 75–83, 2007.
- [18] Matthew Jackson and Leeat Yariv. Diffusion on social networks. *Économie Publique*, 16(1):3–16, 2005.
- [19] Matthew Jackson and Leeat Yariv. Diffusion of behavior and equilibrium properties in network games. *American Economic Review*, 97(2):92–98, May 2007.
- [20] Matthew Jackson and Leeat Yariv. Diffusion, strategic interaction, and social structure. In Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, editors, *Handbook of Social Economics*. Elsevier, 2011. To appear.
- [21] Matthew O. Jackson. *Social and Economic Networks*. Princeton University Press, 2008.
- [22] Sham Kakade, Michael Kearns, John Langford, and Luis Ortiz. Correlated equilibria in graphical games. In *Proceedings of the Annual Conference on Electronic Commerce (EC)*, pages 42–47, 2003.
- [23] Michael L. Katz and Carl Shapiro. Network externalities, competition, and compatibility. *American Economic Review*, 75(3):424–440, 1985.
- [24] Michael Kearns. Graphical games. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay Vazirani, editors, *Algorithmic Game Theory*, pages 159–178. Cambridge University Press, 2007.
- [25] Michael Kearns and Siddharth Suri. Networks preserving evolutionary equilibria and the power of randomization. In *Proceedings of the Annual Conference on Electronic Commerce (EC)*, pages 200–207, 2006.
- [26] Michael J. Kearns, Michael L. Littman, and Satinder P. Singh. Graphical models for game theory. In *Proceedings of the Conference in Uncertainty in Artificial Intelligence (UAI)*, pages 253–260, 2001.
- [27] Jon Kleinberg. The small-world phenomenon: An algorithmic perspective. In *Proceedings of the Annual Symposium on the Theory of Computing (STOC)*, pages 163–170, June 2000.
- [28] S. J. Liebowitz and Stephen E. Margolis. Network externality: An uncommon tragedy. *Journal of Economic Perspectives*, 8(2):133–150, Spring 1994.
- [29] Miller McPherson, Lynn Smith-Lovin, and James M. Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27:415–444, August 2001.
- [30] Stephen Morris. Contagion. *Review of Economic Studies*, 67:57–78, 2000.

- [31] J.F. Nash. Non cooperative games. *Annals of Mathematics*, 54:286–295, 1951.
- [32] Christos Papadimitriou. The complexity of finding Nash equilibria. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay Vazirani, editors, *Algorithmic Game Theory*, pages 29–51. Cambridge Univ. Press, 2007.
- [33] Everett Rogers. *Diffusion of innovations*. Free Press, 4th edition, 1995.
- [34] David Strang and Sarah Soule. Diffusion in organizations and social movements: From hybrid corn to poison pills. *Annual Review of Sociology*, 24:265–290, 1998.
- [35] Thomas Valente. *Network Models of the Diffusion of Innovations*. Hampton Press, 1995.
- [36] Duncan J. Watts and Steven H. Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393:440–442, 1998.
- [37] H. Peyton Young. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton University Press, 1998.