

Local Coordination and the Geometry of Social Networks

Tom Hutchcroft

Olga Rospuskova

Omer Tamuz*

February 4, 2026

Abstract

We study agents playing a pure coordination game on a large social network. Agents are restricted to coordinate locally, without access to a global communication device, and so different regions of the network will converge to different actions, precluding perfect coordination. We show that the extent of this inefficiency depends on the network geometry: on some networks, near-perfect efficiency is achievable, while on others welfare is strictly bounded away from the optimum. We provide a geometric condition on the network structure that characterizes when near-efficiency is attainable. On networks in which it is unattainable, our results more generally preclude high correlations between outcomes in a large spectrum of dynamic games.

1 Introduction

Pure coordination problems involve a decision between alternatives that are *a priori* equivalent, and where the only motivation is to match the choices of others. For example, the choice of a weekly day of rest has obvious coordination incentives, and one can plausibly assume that *a priori* it does not matter which day of the week it is, assuming everyone chooses the same one. Other examples include the choice of outlet plug standard, the side of the street to drive on, and the term used to describe a new concept or object.

We consider agents who have to coordinate with their social network neighbors. When the network is large and coordination is local, one cannot hope that the entire population will make the same choice, and so inefficiencies must arise. The question we ask is: on which social networks is it possible to achieve low inefficiency? In particular, what geometric property of the network determines the range of possible outcomes?

*Caltech. Tom Hutchcroft was supported by a National Science Foundation award (DMS-2246494) and a Packard Fellowship for Science and Engineering. Olga Rospuskova was supported by a Citadel fellowship. Omer Tamuz was supported by a National Science Foundation CAREER award (DMS-1944153) and by a MURI award (N000142412742). We would like to thank Mira Frick, Ryota Iijima and Gábor Pete for stimulating discussions.

We show that low inefficiency is very closely tied to the *amenability* of the social network graph.¹ This is a well-studied graph property, which roughly means that the graph can be divided into distinct communities, each of which has relatively few edges to other communities. We show that on amenable graphs agents can coordinate well in equilibrium (Theorems 1, 2 and 3). On non-amenable graphs, we show that near-efficient coordination is impossible (Theorem 4). Indeed, we show that it is impossible not just in equilibrium, but in any strategy profile.

We consider a setting in which agents can communicate locally: there is some radius of communication r , and agents can exchange messages with their neighbors, neighbors of neighbors, etc, up to distance r . This can be viewed as a reduced form model that can capture the end state of a more detailed process, perhaps involving multiple rounds of communication with direct neighbors, which allows messages to travel some distance.

Messages can be either (locally) public, in which case they are observed by all agents up to distance r (e.g., posting a yard sign) or private. On amenable graphs, this suffices for the existence of equilibria with low inefficiency (Theorem 1). Eliminating private communication still allows for low inefficiency, but with some loss (Theorem 2).

In non-amenable graphs, inefficiency is high, even when private communication is allowed, and indeed even when incentives are not a constraint. Our general result in this setting (Theorem 5) is thus also a contribution to the study of stochastic processes on non-amenable graphs, which is an active area of research. This result also has a distributed computing interpretation: no distributed algorithm running on the nodes of a graph using local information can yield highly correlated outputs when the graph is non-amenable.

The low-inefficiency equilibria we construct on amenable graphs have a simple structure: agents are divided into communities, each community has a leader, and each leader messages the community with instructions determining the action they coordinate on. We call such an equilibrium a leader equilibrium. When communities are insular there are no incentive issues. However, when agents have neighbors who are members of a different community one has to make sure they have no incentive to deviate and follow the other leader's instructions. This issue can be solved by private communication between leaders and their communities (Theorem 1). When private communication is not allowed, these issues can be overcome by a careful construction of communities that are sufficiently insular (Theorem 2). These results show that on amenable graphs there are efficient leader equilibria. The converse is also true: the existence of an efficient leader equilibrium implies that the graph is amenable.

We show that inefficiency must be high on non-amenable graphs, and not just in leader

¹The term *hyperfiniteness* is often used to describe this property in the math literature; see more in the related literature discussion.

equilibria. To this end, we demonstrate that given any low-inefficiency strategy profile one can construct a leader equilibrium that also has low inefficiency. This shows that this simple coordination device is (almost) as powerful as any other.

The main technical tool we use to prove this result is what we call the *Shapley influence distribution* (Owen, 2014; Owen and Prieur, 2017): a measure of how much a variable affects the outcome of a function that depends on several independent inputs. This distribution is given by the Shapley values of a cooperative game played by the input variables that determine the function’s output. We show that the Shapley influence distribution has a strong contraction property: when applied to two similar functions, it yields similar distributions.

In any equilibrium (or indeed, in any strategy profile) we can view a player’s action as a function of the independent inputs given by the messages of the players within the radius of communication. We apply the Shapley influence distribution to this function to yield, for that player, a distribution over the players around them. This will be the distribution used by this player to choose a leader. Intuitively, this means that in the constructed leader equilibrium players are more likely to choose as leaders those in their neighborhood that are more influential to their decisions. The contraction property we prove ensures that agents tend to choose the same leaders as their neighbors, ensuring low inefficiency.

Related literature. The study of coordination has a long history in economics and game theory, including in both cooperative and non-cooperative game theory; a complete survey is beyond the scope of this paper. Coordination on social networks has also been the subject of a number of paper. Perhaps the closest to ours is Morris (2000), who studies contagion in a two-action, repeated local interaction model. Agents live on an infinite graph and repeatedly best-respond to their neighbors’ actions, myopically maximizing payoffs in a coordination game. The question is: starting from a finite region in which all agents play one action, when will best responses eventually make that action spread to the entire graph? A key message is that this depends on the graph’s geometry: slow neighborhood growth supports contagion, reflecting the amenability intuition that communities should have relatively small boundaries for efficient coordination.

Earlier, Ellison (1993) studied a similar two-action coordination problem, but in a large finite population under a noisy best-response dynamics. He highlights that coordination outcomes depend on the interaction structure. When coordination is local, the agents converge quickly to the same action, whereas under uniform interactions convergence is extremely slow. A number of other papers follow a similar path; see Weidenholzer (2010) for a survey of best-response type models of coordination dynamics. This literature is closely tied to the study of interacting particle systems in probability, which originated with the classical Ising model (Ising, 1925) and is very active to this day (see, e.g., a survey

by [Durrett \(2025\)](#)). Both [Morris \(2000\)](#) and [Ellison \(1993\)](#) focus on the long-term outcomes, after repeated interactions have percolated throughout the network, rather than the local interactions we are interested in.

More recent contributions to the coordination on networks literature include [Sadler \(2020\)](#), who analyzes diffusion from a seed, where global spread depends on whether adoption opportunities percolate through a giant connected cluster, captured by a branching-factor condition. [Chwe \(2000\)](#) shows that in a collective action setting, coordination hinges on the network's ability to generate sufficiently strong higher-order beliefs, via a hierarchical communication structure. In his geometric threshold model, dense local overlap facilitates coordination in low-dimensional networks, while weak ties become important mainly once the network is very dense. [Pesk \(2025\)](#) studies random-threshold coordination and shows that in large networks equilibria admit a simple cutoff structure. Its existence is guaranteed under a small influence condition, where each agent's payoff is not driven by any single neighbor.

Beyond coordination, there is a very large literature that studies how the geometry of the social network affects outcomes in various strategic settings. A very partial list of examples includes [Jackson and Yariv \(2007\)](#), who study how network structure affects the diffusion of behavior; [Ballester et al. \(2006\)](#) and [Bramoullé et al. \(2014\)](#), who consider quadratic games with linear best-responses, and show that spectral properties of the adjacency matrix determine equilibrium behavior; and [Golub and Jackson \(2012\)](#), [Acemoglu et al. \(2011\)](#) and [Mossel et al. \(2015\)](#), who study how network geometry affects social learning outcomes.

Several inequivalent notions of amenability have appeared in the mathematics literature. The notion we use is equivalent to so-called *hyperfiniteness* as introduced by [Elek \(2006\)](#); see also [Schramm \(2008\)](#). This notion of (non)amenability could also reasonably be referred to as *statistical (non)amenability* as it concerns a property of graphs determined by the *average* geometry at a random vertex and is robust to the presence of rare outlying regions with atypical geometry. This notion of amenability has been studied extensively by mathematicians in the setting of *unimodular random rooted graphs* ([Aldous and Lyons, 2007](#)) and *measured equivalence relations* ([Kechris and Miller, 2004](#)), the latter being a central topic in modern descriptive set theory.

The concept of amenability was first introduced by [von Neumann \(1929\)](#), who originally defined it as a property of groups, with the goal of understanding the origins of the Banach-Tarski paradox. Classically, a graph is said to be amenable if it contains finite sets of vertices with arbitrarily small surface-to-volume ratio; for groups, being amenable in the sense of von Neumann is equivalent to having Cayley graphs that are amenable in this sense. Finite graphs are always amenable in this sense since the entire vertex set has no boundary. The most classical non-vacuous notion of non-amenability in the finite context is that of

being an *expander graph*: Given $\varepsilon > 0$, a finite graph is said to be an ε -expander if every set containing at most half the vertices has surface-to-volume ratio at least ε . Such finite expander graphs have many applications in computer science and algorithm design (see, e.g., [Hoory et al., 2006](#)). Being an expander graph might reasonably be referred to as *uniform non-amenable* in contrast to our statistical notion, and is a strictly stronger property than being non-amenable in our sense. For example, the giant cluster of a random graph on n vertices with average degree $d > 1$ is not an expander with high probability when n is large due to the presence of rare “bad regions”, but is non-amenable in our statistical sense. The same is true for many “small-world” models of social networks ([Watts and Strogatz, 1998](#)), making our (statistical) notion of non-amenable the more relevant concept in many applications. The two notions are equivalent for transitive graphs ([Benjamini et al., 1999](#)), that is, graphs in which any two vertices are related by a symmetry. Analogues of our main probabilistic theorem (Theorem 5) for transitive graphs have been proven using spectral methods in e.g. [Backhausz and Virág \(2017\)](#) and [Hutchcroft \(2023\)](#), but these proof methods do not apply under our notion of (statistical) non-amenable. [Csóka, Harangi, and Virág \(2020\)](#) prove a version of this theorem for general (non-transitive) graphs, under a uniform non-amenable assumption (see Remark 2).

2 Model

Let N be a finite set of agents. The agents are connected by a social network whose graph is $G = (N, E)$, where $(i, j) \in E \subseteq N \times N$ signifies that i is a neighbor of j . We assume that the graph is undirected, i.e., $(i, j) \in E$ if and only if $(j, i) \in E$. By convention, we set $(i, i) \notin E$. We denote the neighborhood of $i \in N$ by $N_i = \{j : (i, j) \in E\}$. The degree of agent i is the size of N_i , and we denote by d_{\max} the maximum degree of all agents. We think of d_{\max} as being small: each agent has many fewer neighbors than there are agents in the network.

We say that i is connected to j if there is a sequence of agents starting in i and ending in j such that each subsequent pair are in each other’s neighborhood. The connected component of i is the set of agents to whom i is connected. We denote by $B_r(i) \subseteq N$ the set of agents who are at distance at most r from i in the social network. I.e., the agents to whom i is connected by a sequence of agents of length at most r . We refer to this set as the i ’s r -neighborhood.

Each agent has to choose an action in $A = \{-1, +1\}$. The utility of player i for an action profile $a \in A^N$ is given by

$$u_i(a) = - \sum_{j \in N_i} |a_i - a_j|. \quad (1)$$

This is a pure coordination game in which each agent wants to match its neighbors, receiving a payoff of 0 for each match and a payoff of -2 for each mismatch. All of our results extend to similar coordination games with more actions; we opt to focus on the two action case for simplicity of notation and exposition.

We would like to capture the idea that in order to coordinate, agents can communicate locally and pass information among themselves, but that information cannot travel too far in the network. To this end, we allow for two rounds of local communication, after which agents choose their actions.

In the first round, each agent i can choose a message m_i to broadcast to all agents in $B_r(i)$. We call this message agent i 's *locally public* message, as the same message is observed by all the agents in the i 's r -neighborhood. This element of our model can capture a public action, such as posting a sign on one's yard. It can also be a reduced form approach to modeling a more complicated processes, e.g., one in which some message from agent i is communicated to the agents in $B_r(i)$ through chains of exchanges between direct neighbors.

For some of our results we will require an additional, private (but still local) communication stage. After receiving the locally public messages of the agents in $B_r(i)$, agent i can send each agent $j \in B_r(i)$ an additional private message $\hat{m}_{i,j}$. As with locally public messages, private messages sent directly to agents within distance r can succinctly model a more complicated process.

Messages, both locally public and private, take values in a set of messages \mathcal{M} . The set \mathcal{M} will be large (uncountable), placing no restriction on the complexity of the message. For reasons that will become clear soon, we put a particular structure on \mathcal{M} . Let \mathcal{A} be a finite alphabet that includes the two actions $\{-1, +1\}$ as letters, as well as the letter \emptyset . Let $\mathcal{M} = \mathcal{A}^{\mathbb{N}}$, so that a message is an infinite string written using the alphabet \mathcal{A} . Given a letter $\alpha \in \mathcal{A}$, we identify the string $(\alpha, \emptyset, \emptyset, \dots)$ with α itself. In particular, $(-1, \emptyset, \emptyset, \dots)$ is the message -1 , and likewise $(+1, \emptyset, \emptyset, \dots)$ is the message $+1$. The message $(\emptyset, \emptyset, \dots)$ is the message \emptyset , which we call the empty message.

After sending and receiving the private messages, agent i chooses an action $a_i \in \{-1, +1\}$ and receives a payoff according to (1). Formally, there are five periods, in which each agent i

1. sends a locally public message m_i ;
2. learns the locally public messages $m_{B_r(i)} = (m_j)_{j \in B_r(i)}$;
3. sends a private message $\hat{m}_{i,j}$ to each $j \in B_r(i)$;
4. learns the private messages $\hat{m}_{j,i}$ for all $j \in B_r(i)$;

5. chooses an action a_i .

This defines an extensive form game. The information available to agent i at the end of stage 2 consists of the locally public messages of the agents in $B_r(i)$. At the end of stage 4, agents also know the private messages they received and sent. Thus, a pure strategy of an agent is a tuple $\sigma_i = (m_i, n_i, o_i)$, where m_i is the locally public message, n_i is a measurable map that assigns to every realization of $m_{B_r(i)}$ the private messages $(\hat{m}_{i,j})_{j \in B_r(i)}$, and o_i is a measurable map that assigns to every realization of $(m_{B_r(i)}, (\hat{m}_{i,j})_{j \in B_r(i)}, (\hat{m}_{j,i})_{j \in B_r(i)})$ an action in $\{-1, +1\}$. We use the same notation to denote mixed profiles, in which case $(m_i, n_i, o_i)_{i \in N}$ will be independent random variables, all defined on the same probability space. Within this probability space, we will denote by a_i the (random) action taken by i .

We will be interested in mixed equilibria of this game. In particular, we will be interested in the average welfare of this game, i.e., in

$$W = -\frac{1}{|E|} \sum_{(i,j) \in E} |a_i - a_j|.$$

This is the payoff averaged over all ordered pairs (i, j) of agents who are neighbors. It will be more convenient to study the average inefficiency

$$I = -W = \frac{2}{|E|} \sum_{(i,j) \in E} \mathbb{1}_{a_i \neq a_j},$$

that is, twice the fraction of pairs of neighbors that mismatch. This is equal to 0 if all agents coordinate, and to 2 if all agents miscoordinate.

For any graph, there are two equilibria that achieve first-best: the one in which agents ignore all information and all choose the action $+1$, and another in which they all choose -1 . We are interested in how agents might coordinate, rather than in the fact that these are stable courses of action; assuming that agents play one of these equilibria relegates the question of coordination on actions to that of coordination on equilibria. Instead, we would like to model a setting in which the agents need to interact and exchange information in order to coordinate.

To this end, we will define *action-symmetric* strategy profiles, which are invariant to renaming the actions. Playing an action-symmetric strategy captures an a priori indifference to the two actions, making coordination an interesting problem. We will not restrict players to play action-symmetric strategies, but instead study equilibria in which players have no incentive to deviate away from action-symmetry.

In a one-shot game without communication in which agents only choose an action, the only action-symmetric equilibrium is the one in which all agents choose each action with probability one half. Indeed, this is the only action-symmetric strategy profile. In

the extensive form game with communication, agents can use messages to break ties and coordinate without a prior preference to either action.

Formally, we define the negation operator on the message alphabet \mathcal{A} by having it map -1 to $+1$ and vice versa, and mapping each other letter to itself. This operator is extended in the obvious way to the set of messages $\mathcal{M} = \mathcal{A}^{\mathbb{N}}$, so that the negation of $m = (\alpha_1, \alpha_2, \dots) \in \mathcal{M}$ is $-m = (-\alpha_1, -\alpha_2, \dots) \in \mathcal{M}$. A tuple of messages (e.g., a message received from each neighbor) is negated in the same way.

Let \mathcal{X}, \mathcal{Y} be spaces on which negation is defined (e.g., $\{-1, +1\}, \mathcal{A}, \mathcal{M}, B_r(i)^{\mathcal{M}}$). Let $f: X \rightarrow Y$ be a function (e.g., n_i or o_i). We define the function $\iota(f): X \rightarrow Y$ by $[\iota(f)](x) = -f(-x)$. For example, suppose $X = Y = \{-1, +1, \emptyset\}$, in which case we can think of f as a map that chooses a message to be sent as a function of a received message. If f is the constant $+1$, then $\iota(f)$ is the constant -1 . More generally, if $f(\emptyset) = +1$ then $[\iota(f)](\emptyset) = -1$. If $f(\alpha) = \alpha$ then $\iota(f) = f$, and likewise if $f(\alpha) = -\alpha$ then $\iota(f) = f$.

Given a pure strategy $\sigma_i = (m_i, n_i, o_i)$, we define $\iota(\sigma_i) = (-m_i, \iota(n_i), \iota(o_i))$. We say that i 's mixed strategy is action-symmetric if it is ι -invariant: σ_i has the same distribution as $\iota(\sigma_i)$. Equivalently, if we denote by A_i the set of strategies of i , then ι is a bijection from A_i to A_i , and a probability measure μ_i over A_i (i.e., a mixed strategy for player i) is action-symmetric if for any measurable $E \subseteq A_i$ it holds that $\mu_i(E) = \mu_i(\iota(E))$. Intuitively, action-symmetry is invariance to the renaming of the two actions, capturing the notion that players are agnostic to the actions' names.

Note that if all players use action-symmetric strategies, then the probability that $a_i = 1$ is one-half for all players i . In fact, this would suffice as an assumption for all our results: namely, we can replace the action-symmetry assumption with the assumption that strategy profiles are such that each agent takes each action with probability one half.

Coordination is possible even when all players use action-symmetric strategies. For a simple example, suppose that r is equal to the diameter of the graph (i.e., $j \in B_r(i)$ for all $i, j \in N$), so that every player can communicate with every player. Then there is an equilibrium in which players coordinate perfectly: fix a player $\ell \in N$ whom we call the “leader”. The leader chooses a locally public message m_ℓ uniformly from $\{-1, +1\}$. The players then choose to follow the leader ℓ , by taking the action $a_i = m_\ell$. Thus, our communication structure allows for perfect coordination when r is large enough. The question becomes more interesting when r is much smaller than the extent of the graph.

Regardless of the graph, there is always an equilibrium in which the agents ignore all information and choose a_i uniformly at random. In this case, the expected inefficiency will be high (it will, in fact, be equal to one). The main question of this paper is the following: Given $\varepsilon > 0$ and a radius of communication r , for which graphs is there an equilibrium that achieves expected inefficiency at most ε ?

2.1 Amenable graphs

The answer to this question turns out to be closely related to a property of graphs called *amenability*. In this section we initially discuss amenability of infinite graphs, where it is simpler to define, before returning our focus to finite graphs.

To explain this idea, we first consider our game on the line graph, where $N = \mathbb{Z}$ and $j \in N_i$ if $|i - j| = 1$ (see Figure 1). We fix some large radius of communication r . Notice that it is impossible to achieve perfect coordination: if i and j are more than $2r$ apart, their actions must be statistically independent of each other in any action-symmetric profile, and so it is impossible that all agents take the same action. Nevertheless, it is possible to achieve low inefficiency.

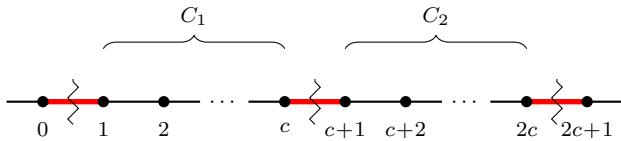


Figure 1: \mathbb{Z} , amenable. Communities have length $c = 2r + 1$. Miscoordination occurs only along the red edges, which form a small fraction of the edges.

To this end, we divide the line into communities of length $c = 2r + 1$: $C_1 = \{1, \dots, c\}$, $C_2 = \{c + 1, \dots, 2c\}$, etc. The leader of C_k will be $\ell_k = k(2r + 1) - r$, the agent in the middle of the interval who is within the distance r to all agents in C_k . This agent will send the same private message to the entire community and will take the action equal to this message. The agents who are not leaders will send empty messages and will choose the action equal to their leader's message.

Agents who are within the same community will choose the same action, and so miscoordination will only occur on the boundary of the communities. This is an equilibrium: agents who are not on the boundary will lose payoff if they deviate, as both of their neighbors will take the same recommended action. Agents who are on the boundary of a community likewise will face a lower expected payoff from deviating, since they know which action their community neighbor will take but do not know which action will be taken by their neighbor across the community border.

Since miscoordination occurs only on the community borders, and since there it occurs with probability one half, the inefficiency will be $2/(2c) = 1/(2r + 1)$. In particular, when r is large, the inefficiency is small. Clearly, on a finite graph that within radius r looks like this graph (i.e., a cycle of length much longer than r) the same conclusion applies, with low inefficiency for large r .

A similar equilibrium can be constructed on the two dimensional grid (see Figure 2). Here, the communities will be $(r + 1) \times (r + 1)$ squares, and so having radius r (the radius

of a connected component is the minimum r such that the entire component is contained in the ball of radius r of some member of the component), and again inefficiency will be small when r is large (this time $1/(r+1)$, so again of order $1/r$). As with the line, the same conclusion applies to finite graphs that locally look like the grid.

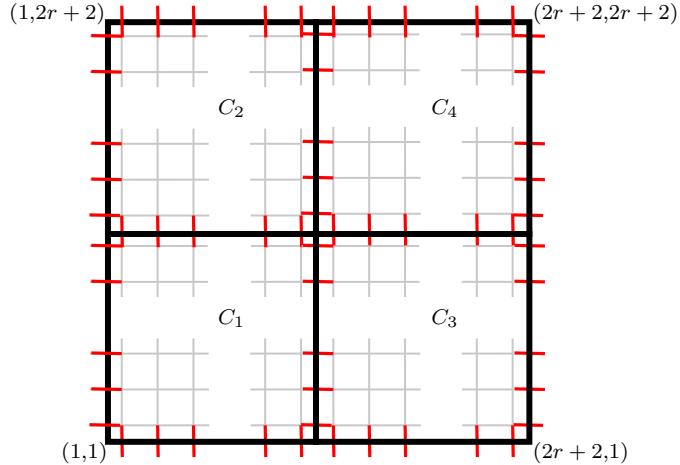


Figure 2: \mathbb{Z}^2 , amenable. Communities are $(r+1) \times (r+1)$ squares. Miscoordination occurs only along the red edges, which, as on the line, form a small fraction of the edges.

Both of these constructions rely on a geometric property of these graphs: namely, that they can be partitioned into communities in such a way that each community has a small boundary as compared to the number of agents it contains. This is not possible on every graph. For example, consider the (infinite) binary tree (see Figure 3). Here, if we choose as a community, say, the first ten levels, then more than half of the agents will be on the boundary of the community, and so inefficiency will be high. This holds if we choose the community to consist of any number of levels. In fact, this holds for any finite community on this graph: the number of edges between members of the community and members outside the community will be at least the size of the community. Thus, on a finite graph that locally looks like this tree, there cannot be efficient leader equilibria, regardless of how large r is.²

In infinite graphs, the existence of finite communities with small boundaries is captured by the definition of an *amenable graph*. Formally, let $G = (N, E)$ be a graph, and assume

²An example of a finite graph that locally looks like a tree is an $(n, d/n)$ -Erdős-Rényi graph for $d > 0$, which, for large n , with high probability looks like a tree within an r -neighborhood of almost all vertices. More formally, if $G_n = (N, E)$ is chosen from this measure over graphs of n vertices, and if i is a uniformly chosen vertex, then the probability that i is in a cycle of length at most r tends to zero as n tends to infinity. The graph will locally look like a Galton-Watson tree, rather than a binary tree.

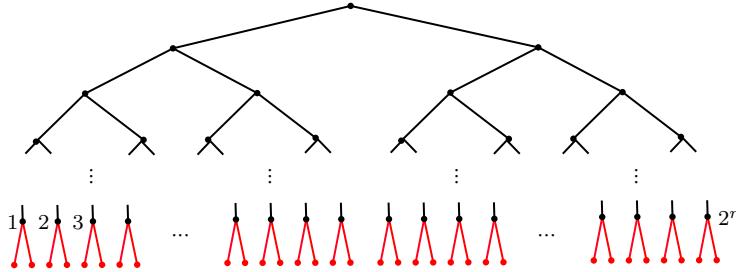


Figure 3: Binary tree, non-amenable.

that N is countable and each neighborhood N_i is finite. We write $\mathcal{P}_f(N)$ for the family of finite subsets of N . For $F \in \mathcal{P}_f(N)$ define its *boundary*

$$\partial F = \{(i, j) \in E : i \in F, j \notin F\},$$

i.e., the edges that cross from F to outside of F . The *surface-to-volume ratio* of F is $\frac{|\partial F|}{|F|}$. The graph G is said to be *amenable* if it has finite sets with an arbitrarily small surface-to-volume ratio:

$$\inf_{F \in \mathcal{P}_f(N)} \frac{|\partial F|}{|F|} = 0.$$

The infinite line and infinite grid are amenable, whereas the infinite binary tree is not.

In our model graphs are finite, and so we will need a finer definition, as all finite graphs are amenable.³ Rather than just requiring the existence of communities with low surface-to-volume ratio, this definition will require a partition of the graph into such sets, so that almost all agents must be members of such communities. We will also need these sets to have radius at most r , so that agents can coordinate actions within their communities; we say that $C \subseteq N$ has radius r if there is an agent $i \in C$ such that C is a subset of $B_r(i)$.

Definition 1. A finite graph $G = (N, E)$ is (ε, r) -amenable if there exists a partition of N into connected components of radius at most r such that the total number of edges connecting nodes in different components is at most $\varepsilon|E|$.

Equivalently, there is a set of edges $D \subseteq E$ such that $|D| \leq \varepsilon|E|$, and for each $i \in N$ the connected component of i in $G' = (N, E \setminus D)$ has radius at most r .⁴ That is, one can remove a small number of connections so that the remaining graph is separated into small communities. Note that while some of these communities may have a large surface-to-volume ratio, it is impossible for these communities to contain many agents, since there

³Let $F = V$. Then ∂F is the empty set, and so $|\partial F|/|F| = 0$.

⁴An alternative, equivalent definition is to partition the graph into subsets that are not necessarily connected.

are at most $\varepsilon|E|$ edges in the union of all boundaries. Indeed, the average surface-to-volume ratio (weighted by population size) is precisely $|D|/|N|$, which is at most $\varepsilon|E|/|N|$. Assuming each agent is connected to a small number of others, this means that when ε is small, almost all agents will be in connected components that have small surface-to-volume ratios.⁵ In the math literature, this property is often called *hyperfiniteness* rather than amenability. A perhaps more appropriate name is *statistical amenability*, as it requires that almost all nodes be the members of a set with a small surface-to-volume ratio, rather than just the existence of such sets, as in the original definition of amenability for infinite graphs.

It is easy to see that for any $\varepsilon > 0$, every cycle graph will be (ε, r) -amenable, assuming r is large enough. The same holds for an n -by- n two dimensional grid. An interesting class of finite graphs which are amenable for non-obvious reasons are the bounded degree finite planar graphs⁶, which are amenable by a theorem of [Lipton and Tarjan \(1977, 1979\)](#) in the following sense: for every maximal degree d_{\max} and $\varepsilon > 0$ there is an r such that every planar graph is (ε, r) -amenable. This holds more generally for bounded degree graphs with any excluded minor ([Alon, Seymour, and Thomas, 1990](#)). An example of a non-amenable graph is an Erdö-Rényi graph, in which each pair of agents are connected by an edge at random, independently and with some probability p . In particular, fixing the expected degree $d = np > 1$, there is an $\varepsilon > 0$ such that for any r the probability that the graph is (ε, r) -amenable tends to zero as the number of agents n tends to infinity. These graphs locally look like trees, which is impossible for a finite planar graph.

3 Coordination on amenable graphs

Our first main result shows that when a graph is (ε, r) -amenable, we can find an equilibrium in which agents coordinate well. These equilibria will be *leader equilibria*, which we informally described above.

Fix $G = (N, E)$ and a partition of N into communities $\{C_1, C_2, \dots\}$ of radius at most r . In each community C_k we choose a leader ℓ_k such that C_k is contained in $B_r(\ell_k)$; the existence of ℓ_k follows immediately from the fact that C_k has radius at most r .

In a leader equilibrium, agents do not use public messages, and so we set $m_i = \emptyset$ for all i . Each leader ℓ_k of C_k sends an empty message $\hat{m}_{\ell_k, j}$ to each $j \in B_r(\ell_k)$ who is not in their community. To the members of their community, leaders send the same private message, chosen uniformly from $\{-1, +1\}$. I.e., if $i, j \in C_k$, then $\hat{m}_{\ell_k, i} = \hat{m}_{\ell_k, j} \in \{-1, +1\}$. Non-leaders send empty messages. Finally, upon observing $\hat{m}_{\ell_k, i}$, each agent $i \in C_k$ chooses

⁵Formally, if G is (ε, r) -amenable, then at most a $\sqrt{2\varepsilon d_{\max}}$ -fraction of the agents will be in connected components with surface-to-volume ratio greater than $\sqrt{2\varepsilon d_{\max}}$.

⁶A graph is *planar* if it can be drawn on the plane without any of its edges intersecting.

$$a_i = \hat{m}_{\ell_k, i}.$$

Off path, any private message that does not come from one's leader or is not in $\{-1, +1\}$ is ignored. If an agent does not receive a private message in $\{-1, +1\}$ from their leader, they choose an action uniformly at random.⁷

This strategy profile is action-symmetric. It is straightforward to verify that it is an equilibrium. First, a leader has no incentive to deviate: since they are a member of their own community, adhering to the equilibrium guarantees coordination with all neighbors inside that community. Even for a leader with neighbors in another community, sending non-empty messages to neighbors outside the community cannot improve payoffs because such messages are ignored by those neighbors under the prescribed behavior. Moreover, since leaders do not observe the private messages other leaders sent to their communities, they cannot hope to coordinate with neighbors from outside their community with probability greater than one half, which they already achieve in this equilibrium. The same holds for non-leaders: following the leader guarantees coordinating with members of their community, while coordinating with members of other communities is impossible to achieve with probability more than one half, since leaders only send messages to their own communities. Finally, non-leaders have no incentive to send non-empty messages, since these are ignored.

Note that implementing these equilibria with public messages would not be possible in general graphs. Indeed, consider the strategy profile in which each leader ℓ_k chooses uniformly at random a public message $m_{\ell_k} \in \{-1, +1\}$, and all agents in $i \in C_k$ take the action $a_i = m_{\ell_k}$. Suppose also that the graph is such that there is an agent i in a community C that has more neighbors in a different community C' than in C , and who is also within the radius of communication of ℓ' , the leader of C' . Then i would have a profitable deviation of choosing $a_i = m_{\ell'}$. Since all we know about the graph is that it is (ε, r) -amenable, we cannot exclude in general the existence of such an agent i , and so we revert to private communication to maintain incentives.

Note also that implementing a leader equilibrium does not require that the communities C_k be connected. Nevertheless, since our communication model is a reduced form of a more detailed process in which only direct neighbors communicate, it is nice to have each C_k connected. As we note above, the definition of (ε, r) -amenability is equivalent to one in which we do not require the communities to be connected.

We use leader equilibria to show our first main result.

Theorem 1. *Suppose G is (ε, r) -amenable. Then there exists an action-symmetric leader equilibrium with radius of communication r and expected inefficiency at most ε .*

⁷We do not address subgame perfection in this paper. Nevertheless, we believe that, under appropriate definitions, these equilibria can be made sequential.

Thus, if we think of ε as very small, we have that on graphs that are very amenable, agents can coordinate very well. Note, that the converse of the theorem statement also holds: if G admits a leader equilibrium with inefficiency at most ε , the set of edges D that connects members of distinct communities witnesses the (ε, r) -amenability of G . Of course, this result does not preclude the existence of equilibria that are more efficient than leader equilibria.

Proof of Theorem 1. Suppose $G = (N, E)$ is (ε, r) -amenable. Then there is a set of edges $D \subseteq E$ such that $|D| \leq \varepsilon|E|$ and each connected component of $(N, E \setminus D)$ is of radius at most r . Denote by $\{C_1, C_2, \dots\}$ the connected components of the graph, which form a partition of N .

Consider a leader equilibrium associated with this partition. In this equilibrium, members of the same community coordinate perfectly, and members of different communities coordinate with probability one half. Since neighboring members of different communities must be connected by an edge in D , the total expected inefficiency is at most $|D|$, and thus the average expected inefficiency is at most $|D|/|E| \leq \varepsilon$. □

3.1 Equilibria without private communication

Theorem 1 shows that on amenable graphs there are leader equilibria that achieve low inefficiency. These equilibria make critical use of the private communication phase and, in particular, hinge on having the leader discriminately communicate to their community only. In light of this, one may wonder what could nevertheless be achieved without private communication, using only locally public messages $(m_i)_i$ to coordinate. We say that a strategy profile is *no-private-communication* if $\hat{m}_{i,j} = \emptyset$ for all i, j .

Our main result of this section is that without private communication there still exist low-inefficiency leader equilibria on amenable graphs, although with slightly worse guarantees.

Theorem 2. *Suppose G is (ε, r) -amenable. Then there exists a no-private-communication, action-symmetric leader equilibrium with radius of communication r and expected inefficiency at most εd_{\max} .*

The main idea behind this proof is the notion of stable communities. We say that a community $C \subseteq N$ is *stable* if for all $i \in C$ it holds that $|N_i \cap C| \geq |N_i \setminus C|$. That is, each agent in C has at least as many neighbors in C as outside of C . The advantage of a stable community is that an agent in a stable community does not have an incentive to deviate from the community's coordinated action, even if they know what their neighbors from outside the community will do.

Suppose that (C_1, C_2, \dots, C_K) are disjoint stable communities, and let $U = N \setminus \cup_k C_k$ be the agents who are not members of any of these communities, so that $(C_1, C_2, \dots, C_K, U)$ is a partition of N . The leader ℓ_k of each stable community will send a public message m_{ℓ_k} chosen uniformly from $\{-1, +1\}$. Non-leaders, as well as members of U , will send an empty public message. Members of stable communities will follow their leader's message. We do not specify explicitly what members of U do, and instead show in the proof that there exist action-symmetric strategies for these agents that make this an equilibrium. Its inefficiency will be at most the size of the set of edges

$$D = (\cup_k \partial C_k) \bigcup \{(i, j) : i \in U, j \in N_i\},$$

since miscoordination can only occur on community boundaries, or outside the communities on edges adjacent to U .

Given this, to prove Theorem 2, it remains to show that if G is (ε, r) -amenable, then it admits a set of edges D of size at most $\varepsilon d_{\max}|E|$ whose removal leaves stable communities of radius at most r , and some isolated vertices.

Proposition 1. *Suppose that a finite graph $G = (N, E)$ is (ε, r) -amenable. Then there exists a set of edges $D \subseteq E$ such that $|D| \leq \varepsilon d_{\max}|E|$, and for each $i \in N$ the connected component of i in $G' = (N, E \setminus D)$ has radius at most r and is either singleton or stable.*

The idea behind the proof of this proposition is simple: start with communities separated by a small set of edges, whose existence is given by the fact that G is (ε, r) -amenable. If a community is not stable, it has a member with more neighbors outside the community than inside it. Banish this member to its own singleton community, and repeat until all communities are stable or singletons. The crux is to show that this ends after a small number of steps.

Proof of Proposition 1. Let (C_1, \dots, C_K) witness (ε, r) -amenability, so each C_k has radius at most r and $\sum_k |\partial C_k| \leq \varepsilon|E|$. If C_k is stable, set $C'_k := C_k$. Otherwise, there exists $i \in C_k$ such that $|N_i \cap C_k| < |N_i \setminus C_k|$. Let $S := C_k \setminus \{i\}$. Removing i decreases the boundary by at least 1: it removes $|N_i \setminus C_k|$ boundary edges and creates $|N_i \cap C_k|$ new boundary edges, so $|\partial S| \leq |\partial C_k| - 1$. If S is still unstable, repeat the same operation. Since $|\partial(\cdot)|$ is a nonnegative integer that drops by at least 1 each time, this procedure terminates after at most $|\partial C_k|$ deletions, yielding a set $C'_k \subseteq C_k$ that is either empty or stable.

Let $U := N \setminus \bigcup_k C'_k$ be the set of vertices not contained in the resulting stable sets. By the bound above, the peeling procedure deletes at most $|\partial C_k|$ vertices from C_k , hence $|U| \leq \sum_k |\partial C_k|$. Now let D be the union of $\bigcup_k \partial C'_k$ and all edges incident to the vertices

in U . Then,

$$|D| \leq \sum_k |\partial C'_k| + d_{\max}|U|.$$

Since the number of boundary edges decreased by at least the size of U ,

$$\sum_k |\partial C'_k| + d_{\max}|U| \leq (\sum_k |\partial C_k| - |U|) + d_{\max}|U| \leq d_{\max} \sum_k |\partial C_k| \leq \varepsilon d_{\max}|E|.$$

In $G' = (N, E \setminus D)$, every vertex in U is isolated, hence a singleton component. Any other vertex lies in some $C'_k \subseteq C_k$, which is stable. Since C_k has radius at most r , so does C'_k . \square

Proof of Theorem 2. Since $G = (N, E)$ is (ε, r) -amenable, by Proposition 1 there exists a set of edges $D \subseteq E$ such that $|D| \leq \varepsilon d_{\max}|E|$, $N = C \cup U$, where C is the union of the stable sets (C_1, C_2, \dots) , each C_k is a connected component of the graph $(N, E \setminus D)$ of radius at most r , and each $i \in U$ is a singleton component in this graph. Assign to each community C_k a leader $\ell_k \in C_k$ such that $C_k \subseteq B_r(\ell_k)$.

Let $C = \cup_k C_k$ so that N is the disjoint union of C , the agents who are members of communities, and U , the agents who are not. We will construct our action-symmetric equilibrium by fixing the behavior of the agents in C , and relying on an existence result to extend the strategy profile to the agents in U . For both groups, we assume that all private messages are empty, and any private messages are ignored, so that there is no incentive to send private messages.

We start by specifying the behavior of the agents in C . The strategy of each leader $\ell_k \in C_k$ is as follows:

1. choose a public message m_{ℓ_k} uniformly at random from $\{-1, +1\}$;
2. take action $a_{\ell_k} = m_{\ell_k}$, regardless of any messages received;
3. in the (off-path) event that m_{ℓ_k} is not in $\{-1, +1\}$, choose the action a_{ℓ_k} uniformly from $\{-1, +1\}$.

Fix the strategy of any non-leader $i \in C_k$ to

1. send an empty public message;
2. take action $a_i = m_{\ell_k}$ if the leader's message is in $\{-1, +1\}$;
3. in the (off-path) event that m_{ℓ_k} is not in $\{-1, +1\}$, choose the action a_i uniformly from $\{-1, +1\}$.

Thus, members of communities follow a leader equilibrium using the leader's public messages (as opposed to the private ones used in Theorem 1), and ignore any off-path messages. Note that these strategies are action-symmetric.

Given that the agents in C follow these strategies, it is not immediately obvious what the best-responses are of the agents in U . The issue is that depending on the graph structure and the set D , it could be that an agent in U receives the public messages of many leaders, and that their neighbors, who are also in U , likewise see the same messages. In this case, they could, for example, coordinate to follow the message of one particular leader, or alternatively follow another one. To address this problem, we do not construct their strategies explicitly, but instead apply an existence result.

We define a finite normal form game played by the agents in U . For each $i \in U$, let L_i be the set of leaders who are in $B_r(i)$, i.e., $L_i = \{\ell_k : \ell_k \in B_r(i)\}$. A pure strategy of i is a map $s_i : \{-1, +1\}^{L_i} \rightarrow \{-1, +1\}$. This will be interpreted as the action i will take in the original game when observing a given tuple of public messages from the leaders in their radius of communication. Given a pure strategy profile s , the expected utility of the agents is the expectation of (1), their utility in the original game, where the agents in C are dummy players playing as described above, and the agents in U use their strategies in s to choose their actions as a function of the public messages sent by the leaders.

The map ι defining action-symmetry is what [Nash \(1951\)](#) calls an *automorphism* or *symmetry* of this game, as it preserves payoffs: $u_i(\iota(s_1), \dots, \iota(s_n)) = u_i(s_1, \dots, s_n)$. He shows (Theorem 2, p. 289) that given an automorphism, there exists a mixed equilibrium that is invariant to the automorphism, which in our setting precisely translates to action-symmetry. Thus, we have an action-symmetric equilibrium s^* of this restricted game of the agents in U . For each $i \in U$ we define a strategy in the original game as follows:

1. send an empty public message;
2. if the public messages of the leaders in L_i are all in $\{-1, +1\}$, take the action specified by applying s_i^* to these messages;
3. if at least one leader in L_i sends a public message not in $\{-1, +1\}$, choose an action uniformly at random.

This is the obvious lifting of the strategy s_i^* to the original game. It follows immediately from the equilibrium property of s^* that agents in U are best responding on path. It thus remains to be shown that agents in C are best responding.

Non-leaders are best-responding because of the stability property of the communities: they have at least as many neighbors in their own community, and these will take the action recommended by the leader, and so no improvement is possible. Deviations other

than not following the leader—i.e., sending non-empty public or private messages—are easily ruled out, as all such messages are ignored.

To verify that leaders are best-responding, we need to check that they have no incentive to deviate from choosing their public messages uniformly from $\{-1, +1\}$ and instead choosing (say) $+1$. All other deviations (including sending a public message not in $\{-1, +1\}$ and taking an action other than the message) are again easy to rule out.

Fix a leader ℓ . If ℓ has no neighbors outside its community then clearly there is no profitable deviation, since ℓ is guaranteed to coordinate with all its neighbors. Otherwise, let j be a neighbor of ℓ . If $j \in C$ then j ignores ℓ 's messages and there's no incentive to deviate. Suppose $j \in U$. The action of this agent is a (random) function s_j^* of the messages of the leaders in $B_r(j)$, that is $a_j = s_j^*(m_\ell, m_{B_r(j) \setminus \{\ell\}})$. Thus, the probability that j chooses $a_j = m_\ell$ is

$$\mathbb{P}[a_j = m_\ell] = \mathbb{P}[s_j^*(m_\ell, m_{B_r(j) \setminus \{\ell\}}) = m_\ell].$$

Conditioned on sending the message $m_\ell = +1$, this probability is

$$\begin{aligned} \mathbb{P}[a_j = m_\ell | m_\ell = +1] &= \mathbb{P}[s_j^*(+1, m_{B_r(j) \setminus \{\ell\}}) = +1 | m_\ell = +1] \\ &= \mathbb{P}[s_j^*(+1, m_{B_r(j) \setminus \{\ell\}}) = +1], \end{aligned}$$

where the second equality is a consequence of the fact that m_ℓ , $m_{B_r(j) \setminus \{\ell\}}$ and s_j^* are all independent, as the randomization in the choice of mixed strategy is independent of what occurs in the game.

Now, s_j^* and $\hat{s}_j^* = \iota(s_j^*)$ have the same distribution by action-symmetry. Thus, and since s_j^* is independent of the public messages,

$$\mathbb{P}[a_j = m_\ell | m_\ell = +1] = \mathbb{P}[\hat{s}_j^*(+1, m_{B_r(j) \setminus \{\ell\}}) = +1].$$

It follows from the definition of ι that

$$\mathbb{P}[a_j = m_\ell | m_\ell = +1] = \mathbb{P}[s_j^*(-1, -m_{B_r(j) \setminus \{\ell\}}) = -1].$$

Since $m_{B_r(j) \setminus \{\ell\}}$ and $-m_{B_r(j) \setminus \{\ell\}}$ have the same distribution (and are independent of s_j^*) it follows that

$$\mathbb{P}[a_j = m_\ell | m_\ell = +1] = \mathbb{P}[s_j^*(-1, m_{B_r(j) \setminus \{\ell\}}) = -1],$$

which is equal to the probability that $a_j = m_\ell$ conditioned on $m_\ell = -1$. Thus, the probability that j decides to follow the leader's message is independent of the message, and mixing is a best response for the leader. \square

3.2 Local strategy profiles

The proof of Theorem 1 relies on a given partition of the social network into connected components with small boundaries. One may object that agreeing on this partition requires some global coordination, and that it would be interesting to know whether coordination on actions can be achieved without first coordinating on communities.

Leader equilibria also require a choice of leader, which again can be viewed as a coordination problem. As we show below, this is a relatively easy problem to solve: given a partition, we can choose a leader locally at random using public messages, and then proceed with the leader equilibrium, coordinating on actions using private messages. The more challenging problem is the local choice of partition.

In this section, we consider graphs that “look the same” around each vertex, up to radius r . An illustrative example is the $n \times n$ torus: a two dimensional grid, in which nodes at the boundaries are connected to the corresponding nodes on the opposite side. In this graph the r -neighborhood of all agents are identical. We consider strategy profiles that are local, in the sense that they do not require a globally-coordinated partition of the graph into communities, and also do not require a priori coordination on leaders. In this setting, we show that amenability remains a sufficient condition for coordination, but with weaker guarantees. Specifically, we show that for any graph that is (ε, r) -amenable there exists a local, action-symmetric equilibrium with expected inefficiency at most $\varepsilon \left(1 + \log \frac{1}{\varepsilon}\right)$. Thus, inefficiency is still small when ε is small, but larger than what we can guarantee without the locality restriction.

To coordinate locally, agents will use both public and private communication. First, agents will use public messages to construct communities, and to choose their leaders. They will then play a leader equilibrium, implemented as before using private communication.

To construct communities using public communication, the agents will decide locally which edges of the original graph to remove, with the communities being the connected components of the remaining graph. Fix a connected subgraph F of radius r ; on the torus we take F to be a $(r + 1) \times (r + 1)$ squares, whose surface-to-volume ratio is $\varepsilon = 4/(r + 1)$. In general, we will want F to have the lowest surface-to-volume ratio possible among all subgraphs of radius r .

Let F' denote a subgraph that is isomorphic to F ; on the torus, these subgraphs are precisely the translated copies of the square. Using public communication, the agents will decide whether or not to select each F' . This decision will be made independently for each such F' , including overlapping ones. The probability of selecting F' will be chosen so that each agent will, with high probability, be a member of some selected F' . The edges removed are the boundary edges of all the selected F' s, with the selection probability being low enough so that the total number of edges removed is unlikely to be too large. The

edges adjacent to agents who are not members of any selected F' are also removed, making them singleton communities. The connected components of the remaining graph form the communities, and finally the public messages are also used to select the leader of each community. Figure 4 illustrates this rule.

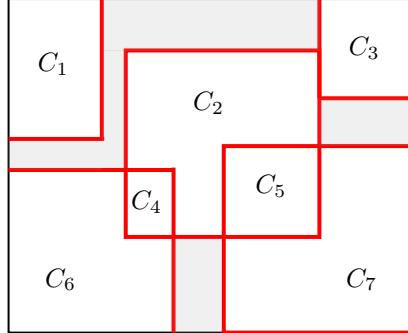


Figure 4: In the equilibrium constructed in the proof of Theorem 3, agents choose at random which tiles should be selected. Communities are the intersections of the selected tiles. Agents outside the selected tiles form singleton communities.

Note that since agents prefer not to be at the boundaries of communities, care needs to be taken to ensure that the local choice of communities is incentivized. Our main tool is the computer-science technique of “secret-sharing” (Shamir, 1979): when a group of agents F' needs to choose at random with probability p whether to select their subgraph, we have each send a public message m_i chosen uniformly from $[0, 1]$, add these messages, and select F' if the fractional part of their sum is in $[0, p]$. It is easy to verify that as long as F' is not a singleton, agents are indifferent between sending any message in $[0, 1]$. Since we have to repeat this for many subgraphs F' , the public message m_i will in fact have to be a tuple of many independent messages, each distributed uniformly on $[0, 1]$.

We formalize the notion of graphs that “look the same” around each vertex as follows. We say that $G = (N, E)$ is *r-locally-transitive* if for any two vertices $i, j \in N$ there exists a *rooted graph isomorphism* between $B_r(i)$ and $B_r(j)$: a bijection $\varphi : B_r(i) \rightarrow B_r(j)$ such that (1) $\varphi(i) = j$, and (2) the bijection preserves vertex adjacency: $(i, j) \in E \iff (\varphi(i), \varphi(j)) \in E$. A strategy profile σ is *network-symmetric* if for all $i, j \in N$ for every rooted graph isomorphism $\varphi : B_r(i) \rightarrow B_r(j)$ it holds that $\varphi(\sigma_i) = \sigma_j$, where we extend φ to a map $\varphi : A_i \rightarrow A_j$ in the obvious way. This notion captures the idea that all agents follow the same strategy, rather than having different roles in the strategy profile, avoiding the conceptual need for a central planner.

Theorem 3. *Suppose G is (ε, r) -amenable and $(2r+1)$ -locally transitive. Then there exists a network-symmetric, action-symmetric leader equilibrium with radius of communication $2r$ and expected inefficiency at most $\varepsilon \left(1 + \log \frac{1}{\varepsilon}\right)$.*

This theorem shows that efficient coordination is possible on an (ε, r) -amenable graph even without reliance on a global partition. Nevertheless, in our construction communication is less efficient, by a factor of $1 + \log \frac{1}{\varepsilon}$, and requires a larger radius of communication.

4 Near-optimality of leader equilibria on cycle graphs

Theorem 1 shows that on an (ε, r) -amenable graph there is an equilibrium with radius of communication r that achieves inefficiency ε . A natural question is: are there equilibria that are more efficient than leader equilibria? In general, we do not know the answer to this question. In the next section, we give a lower bound, showing that if the best leader equilibrium has inefficiency ε , then no strategy profile can achieve inefficiency less than $\varepsilon^2/8$. In this section, we show that for the particularly simple example of cycle graphs leader equilibria are optimal. In fact, we show that no action-symmetric strategy profile can achieve lower inefficiency.

In a cycle graph, the set of agents is $N = \{0, 1, \dots, n-1\}$, and i is connected to j if $i = j \pm 1 \pmod n$. We assume that n is much larger than the radius of communication ($n \geq 5r$ suffices), and, moreover, for simplicity, that the number of agents n is a multiple of $2r+1$, so that we can divide the circle into segments of length $2r+1$ (and hence radius r) as in Figure 1. The average inefficiency of the corresponding leader equilibrium is $1/(2r+1)$, as we explain in § 2.1 above.

Consider an action-symmetric strategy profile that achieves average inefficiency x . We think of the actions $(a_i)_{i \in N}$ as random variables defined on the same probability space \mathbb{P} . These random variables take values in $\{-1, +1\}$, and all have expectation zero, by action-symmetry. Thus, they are vectors in L_0^1 , the space of zero-mean random variables. This is a vector space, and indeed a Banach space under the metric $d(X, Y) = \mathbb{E}[|X - Y|]$. Under this metric, each a_i has unit norm.

Since the average inefficiency is x , there must be a sequence of agents $i, i+1, \dots, i+2r+1 \pmod n$ with average inefficiency $\frac{1}{2r+1} \sum_{j=i}^{i+2r} \mathbb{E}[|a_j - a_{j+1}|] \leq x$. Hence $\sum_{j=i}^{i+2r} d(a_j, a_{j+1}) \leq (2r+1)x$. From the triangle inequality it follows that $d(a_i, a_{i+2r+1}) \leq (2r+1)x$. Since the radius of communication is r , a_i and a_{i+2r+1} are independent, and so $d(a_i, a_{i+2r+1}) = 1$. We have thus shown that $x \geq 1/(2r+1)$, and so no strategy profile can achieve inefficiency lower than the leader equilibrium.

5 Impossibility of coordination on non-amenable graphs

Theorems 1 and 2 show that efficient coordination is possible on an amenable graph. We now turn to the other direction and show that non-amenable graphs do not admit efficient coordination. On such graphs every finite set has a large surface-to-volume ratio. As

a result, any attempt to partition the graph into communities necessarily creates many agents on the boundary, who face a high probability of mismatch with their neighbors. This precludes the existence of a leader equilibrium with low inefficiency.

The impossibility is in fact stronger. On a non-amenable graph no equilibrium of any kind can achieve low inefficiency. In fact, even if we ignore incentives and consider general strategy profiles, the resulting actions will still generate a large amount of mismatch.

Thus, amenability is not only a sufficient, but also a necessary condition for the existence of efficient coordination. The next theorem makes this precise.

Theorem 4. *Suppose G admits an action-symmetric strategy profile with a radius of communication r and average expected inefficiency ε . Then G is $(\sqrt{8\varepsilon}, r)$ -amenable.*

This result provides a partial converse to Theorem 1. It is partial since $\sqrt{8\varepsilon}$ is much larger than ε (for small ε) and so we cannot preclude that for intermediate ranges there are non-leader equilibria that achieve low inefficiency. Nevertheless, these results together qualitatively characterize amenability as a necessary and sufficient condition for low inefficiency. Since a graph is (ε, r) -amenability if and only if it admits a leader equilibrium with inefficiency ε , this theorem immediately implies (by a transformation of ε) that if the best leader equilibrium has inefficiency ε , then no equilibrium can achieve inefficiency less than $\varepsilon^2/8$.

As noted above, incentives do not play a role here: this result applies to any strategy profile. We accordingly rephrase it in a more general, probabilistic framework. The advantage will be that this version is easily applied to a large class of related economic models, as we explain in §6.

Theorem 5. *Let $G = (N, E)$ be a finite graph, and let $(Z_i)_{i \in N}$ be independent random variables associated to each vertex. Let $(X_i)_{i \in N}$ also be random variables associated to each vertex, with each X_i measurable with respect to the sigma-algebra generated by $(Z_j)_{j \in B_r(i)}$. Suppose that $\mathbb{E}[X_i] = 0$ and $\mathbb{E}[X_i^2] = 1$ for all i . If*

$$\frac{1}{|E|} \sum_{(i,j) \in E} \frac{1}{2} \mathbb{E}[(X_i - X_j)^2] \leq \varepsilon,$$

then G is $(\sqrt{8\varepsilon}, r)$ -amenable.

Theorem 4 is an immediate corollary, as each action a_i (corresponding to X_i) is measurable with respect to the sigma-algebra generated by the mixed strategies (corresponding to Z_j) of the agents in $B_r(i)$.

To prove Theorem 5, we start with random variables $(X_i)_{i \in N}$ satisfying the hypothesis of the theorem. These correspond to an action profile on G with small average expected inefficiency ε .

In the first step, for each agent we construct a probability distribution over the agents in her r -neighborhood that measures how much each agent influences her action. Low inefficiency implies that neighboring agents rarely take different actions. We show that this, in turn, forces their influence distributions to be very close to each other in terms of total variation distance.

In the second step, we use these influence distributions to associate to each agent a leader. Since neighboring agents have similar distributions, the probability that they end up choosing the same leader is high. Agents that select the same leader form a community and choose the same action, so only a small fraction of agents lie on the boundaries between communities. This gives a partition of the graph into small communities with low surface-to-volume ratio, which characterizes amenable graphs. We explain the proof ideas in more detail below.

5.1 Shapley influence distributions

Let $Z = (Z_i)_{i \in N}$ be independent (but not necessarily identically distributed) random variables in a standard Borel probability space, indexed by a finite or countable set N . For a finite $S \subseteq N$ denote $Z_S = (Z_i)_{i \in S}$.

Let X be a random variable with mean zero and unit variance, measurable with respect to $\sigma(Z)$, i.e., $X = f(Z)$ for some measurable f . We would like to calculate an influence distribution: a probability measure μ_X over N that captures how much each Z_j influences X . In our setting, this will measure how much each agent's action was influenced by each other agent in its r -neighborhood.

Formally, let $B(Z)$ be the collection of real random variables X that are measurable with respect to the sigma-algebra generated by Z and such that $\mathbb{E}[X] = 0$, $\mathbb{E}[X^2] = 1$. An influence distribution map $X \mapsto \mu_X$ assigns to each $X \in B(Z)$ a probability measure μ_X over N . We require μ_X to be supported on those j 's that determine X . That is, for any $S \subseteq N$ such that X is measurable with respect to $\sigma(Z_S)$, it holds that $\mu_X(S) = 1$.

A desirable property of an influence distribution map is that if $X, Y \in B(Z)$ are close, then the influence distributions μ_X and μ_Y should be close. In particular, for our purposes we will need that the total variation distance between μ_X and μ_Y should be small whenever X and Y are close in L^2 . Recall that the total variation distance between two probability measures $\mu, \nu \in \Delta(N)$ is given by

$$d_{\text{TV}}(\mu, \nu) = \sup_{A \subseteq N} |\mu(A) - \nu(A)| = \sum_{i \in N} \frac{1}{2} |\mu(i) - \nu(i)|.$$

To this end, given random variable $X \in B(Z)$, define a cooperative game v_X for the

set of player N via

$$v_X(S) = \text{Var}(\mathbb{E}[X|Z_S]). \quad (2)$$

That is, v_X assigns to each non-empty finite coalition $S \subseteq N$ the variance of the random variable $\mathbb{E}[X|Z_S]$. This is a measure for how much the collection of random variables Z_S influences X . To see this, it is helpful to take a geometric perspective: $\mathbb{E}[X|Z_S]$ is the projection, in the Hilbert space of square-integrable random variables, of the random variable X to the subspace of $\sigma(Z_S)$ -measurable random variables. The variance is the square of the norm, and when this is high Z_S contains much information about X .

Let $\mu_X: N \rightarrow \mathbb{R}$ be the Shapley values⁸ of the game v_X . Since $v_X \geq 0$, v_X is monotone, and $v_X(N) = \text{Var}(X) = 1$, μ_X is indeed a probability measure on N .

We call μ_X the Shapley influence distribution of X with respect to $(Z_i)_i$. It is easy to see that if X is measurable with respect to $\sigma(Z_S)$, and if $j \notin S$, then $v_X(T \cup \{j\}) = v_X(T)$ for all T —i.e., j is a null player—and thus $\mu_X(j) = 0$. Hence the Shapley influence distribution is indeed an influence distribution.

The next proposition is the main technical contribution of this paper, showing that the Shapley influence distribution map has a strong contraction property.

Proposition 2. *Let $Z = (Z_j)_{j \in N}$ be independent random variables indexed by a finite or countable set N . Then for all $X, Y \in B(Z)$ the Shapley influence distribution map μ satisfies $d_{\text{TV}}(\mu_X, \mu_Y) \leq \sqrt{\mathbb{E}[(X - Y)^2]}$.*

Conceptually, this contraction result shows that the Shapley influence distribution map is stable to small perturbations.⁹ This will be useful to us, as we will apply it to $X = a_i$ and $Y = a_j$ for a pair of highly correlated neighbors $(i, j) \in E$, and will be able to conclude that their Shapley influence distributions are close. The remainder of this section is devoted to the proof of this proposition.

Denote by \mathcal{L} the set of $\sigma(Z)$ -measurable random variables X such that $\mathbb{E}[X] = 0$ and $\mathbb{E}[X^2] < \infty$. The set \mathcal{L} is a separable real Hilbert space, when equipped with the inner product $(X, Y) = \mathbb{E}[X \cdot Y]$. Denote by \mathcal{L}_i the subspace of $\sigma(Z_i)$ -measurable elements of \mathcal{L} . We observe that \mathcal{L}_i is a closed subspace of \mathcal{L} . For each i , let \mathcal{U}_i be a (finite or countable) orthonormal basis of \mathcal{L}_i . An orthonormal basis of \mathcal{L} is then the collection \mathcal{U} of random variables of the form $U_S = \prod_{i \in S} U_i$ where S is a finite subset of N and each U_i is in \mathcal{U}_i . Given a finite, non-empty $S \subseteq N$, let \mathcal{U}_S be the collection of all such U_S , so that $\mathcal{U} = \cup_S \mathcal{U}_S$.

⁸When N is infinite the usual definition of the Shapley value cannot be applied, as one cannot draw uniformly a permutation of N . Instead, we use the Harsanyi dividend definition of the Shapley value, which is still well-defined.

⁹In Appendix B we discuss the stability of classical Shapley values.

We can thus write any $X \in \mathcal{L}$ as a linear combination of elements of \mathcal{U} :

$$X = \sum_S \sum_{U_S \in \mathcal{U}_S} \hat{x}_{U_S} U_S,$$

where $\hat{x}_{U_S} = (X, U_S) \in \mathbb{R}$ is the coefficient of U_S .

By Parseval, we have that for $X, Y \in \mathcal{L}$

$$(X, Y) = \sum_S \sum_{U_S \in \mathcal{U}_S} \hat{x}_{U_S} \cdot \hat{y}_{U_S}.$$

Using this, for a finite $S \subseteq N$ we can write

$$\mathbb{E}[X|Z_S] = \mathbb{E} \left[\sum_T \sum_{U_T \in \mathcal{U}_T} \hat{x}_{U_T} U_T \middle| Z_S \right] = \sum_T \sum_{U_T \in \mathcal{U}_T} \hat{x}_{U_T} \mathbb{E}[U_T|Z_S].$$

Recall that $U_T = \prod_{i \in T} U_i$, where each U_i is $\sigma(Z_i)$ -measurable, and thus this is a product of independent random variables. Note that

$$\mathbb{E}[U_i|Z_S] = \begin{cases} U_i & \text{if } i \in S \\ 0 & \text{if } i \notin S. \end{cases}$$

Hence

$$\mathbb{E}[U_T|Z_S] = \begin{cases} U_T & \text{if } T \subseteq S \\ 0 & \text{if } T \not\subseteq S. \end{cases}$$

Thus,

$$\mathbb{E}[X|Z_S] = \sum_{T \subseteq S} \sum_{U_T \in \mathcal{U}_T} \hat{x}_{U_T} U_T.$$

Since \mathcal{U} is an orthonormal basis, we get by Parseval that

$$v_X(S) = \text{Var} \left(\sum_{T \subseteq S} \sum_{U_T \in \mathcal{U}_T} \hat{x}_{U_T} U_T \right) = \sum_{T \subseteq S} \sum_{U_T \in \mathcal{U}_T} \hat{x}_{U_T}^2.$$

Denote

$$\hat{x}_T^2 = \sum_{U_T \in \mathcal{U}_T} \hat{x}_{U_T}^2, \tag{3}$$

so that

$$v_X(S) = \sum_{T \subseteq S} \hat{x}_T^2.$$

This is exactly the Harsanyi dividend representation of a cooperative game (Harsanyi, 1982), with dividends \hat{x}_T^2 . Therefore, by the standard Shapley value formula in terms of Harsanyi dividends,

$$\varphi_k(v_X) = \sum_{S \ni k} \frac{1}{|S|} \hat{x}_S^2.$$

Fix $X, Y \in \mathcal{L}$ with $\mathbb{E}[X^2] = \mathbb{E}[Y^2] = 1$. By the definition of μ ,

$$d_{\text{TV}}(\mu_X, \mu_Y) = \frac{1}{2} \sum_i |\mu_X(i) - \mu_Y(i)| = \frac{1}{2} \sum_i \left| \sum_{S \ni i} \frac{1}{|S|} \hat{x}_S^2 - \sum_{S \ni i} \frac{1}{|S|} \hat{y}_S^2 \right|.$$

By the triangle inequality we have

$$d_{\text{TV}}(\mu_X, \mu_Y) = \frac{1}{2} \sum_i \left| \sum_{S \ni i} \frac{\hat{x}_S^2 - \hat{y}_S^2}{|S|} \right| \leq \frac{1}{2} \sum_S |\hat{x}_S^2 - \hat{y}_S^2|.$$

Recalling the definition of \hat{x}_S^2 in (3),

$$d_{\text{TV}}(\mu_X, \mu_Y) \leq \frac{1}{2} \sum_S \left| \sum_{U_S \in \mathcal{U}_S} \hat{x}_{U_S}^2 - \sum_{U_S \in \mathcal{U}_S} \hat{y}_{U_S}^2 \right|,$$

and by another application of the triangle inequality

$$d_{\text{TV}}(\mu_X, \mu_Y) \leq \frac{1}{2} \sum_{U \in \mathcal{U}} \left| \hat{x}_U^2 - \hat{y}_U^2 \right|.$$

By Cauchy-Schwarz,

$$\sum_{U \in \mathcal{U}} |\hat{x}_U^2 - \hat{y}_U^2| = \sum_{U \in \mathcal{U}} |\hat{x}_U - \hat{y}_U| \cdot |\hat{x}_U + \hat{y}_U| \leq \sqrt{\sum_{U \in \mathcal{U}} (\hat{x}_U - \hat{y}_U)^2} \cdot \sqrt{\sum_{U \in \mathcal{U}} (\hat{x}_U + \hat{y}_U)^2}.$$

By Parseval, $\sum_U (\hat{x}_U - \hat{y}_U)^2 = \mathbb{E}[(X - Y)^2]$ and $\sum_U (\hat{x}_U + \hat{y}_U)^2 = \mathbb{E}[(X + Y)^2]$. The latter is at most 4, and so

$$d_{\text{TV}}(\mu_X, \mu_Y) \leq \sqrt{\mathbb{E}[(X - Y)^2]}.$$

This completes the proof of Proposition 2.

Remark 1. Proposition 2 is related to (Galicza and Pete, 2024, Lemma 2.9), which establishes a looser relationship between covariances and the “clue”, a way of measuring the dependence of X on the input variables $(Z_j)_{j \in N}$ using Fourier analysis on the hypercube. The proof of Proposition 2 shows that this “clue” is in fact equal to the Shapley influence distribution when the Z_j are i.i.d. unbiased coinflips, so that our Proposition 2 strengthens and clarifies their Lemma 2.9.

5.2 Coupling

In addition to the influence distributions of Proposition 2, to prove Theorem 5 we will make use of what we call grand couplings.

Let μ_1, μ_2 be probability measures on a finite or countable set N . A coupling of these measures is a pair of random variables L_1, L_2 defined on the same probability space and such that L_i has a distribution μ_i . We will be interested in couplings that maximize the probability that $L_1 = L_2$. A classical result (due to Doeblin, see [Lindvall, 1991](#)) is that the minimum of $\mathbb{P}[L_1 \neq L_2]$ over all couplings of μ_1, μ_2 is achieved and is exactly equal to the total variation distance $d_{\text{TV}}(\mu_1, \mu_2)$.

One can ask the same question for three measures, $\mu_1, \mu_2, \mu_3 \in \Delta(N)$. In this case it turns out that one can always find a coupling (or perhaps a throupling) L_1, L_2, L_3 such that for each $i, j \in \{1, 2, 3\}$ it holds that $\mathbb{P}[L_i \neq L_j] \leq 2d_{\text{TV}}(\mu_i, \mu_j)$. Remarkably, the same holds for coupling of more than three measures. In fact, the following proposition, which is due to [Kleinberg and Tardos \(2002\)](#) (see also [Broder, 1997](#), for a closely related idea), shows that one can couple all the probability measures on N and achieve the same guarantee. A particularly good exposition is given by [Angel and Spinka \(2019\)](#).

Proposition 3 ([Kleinberg and Tardos \(2002\)](#); [Angel and Spinka \(2019\)](#)). *Let N be a finite or countable set, and let $(\mu_i)_i$ be the set of all probability measures on N . Then there exists a probability space with random variables $(L_i)_i$ such that L_i has a distribution μ_i , and for each $i, j \in N$ it holds that $\mathbb{P}[L_i \neq L_j] \leq 2d_{\text{TV}}(\mu_i, \mu_j)$.*

We describe a simple construction that achieves this, given by [Angel and Spinka \(2019\)](#); we refer the reader to the proof there. Let $(E_k)_{k \in N}$ be i.i.d. random variables with the unit exponential distribution (i.e., $\mathbb{P}[E_k \leq x] = 1 - e^{-x}$ for $x \geq 0$) defined on a common probability space. For each $\mu_i \in \Delta(N)$, let

$$L_i = \arg \min_{k \in N} \frac{E_k}{\mu_i(k)}.$$

Then $\mathbb{P}[L_i = k] = \mu_i(k)$ and

$$\mathbb{P}[L_i \neq L_j] \leq \frac{2d_{\text{TV}}(\mu_i, \mu_j)}{1 + d_{\text{TV}}(\mu_i, \mu_j)} \leq 2d_{\text{TV}}(\mu_i, \mu_j).$$

The relevance of Proposition 3 to amenability of graphs was first noted by [Hutchcroft and Lopez \(2024\)](#). We follow a similar approach, namely, to use the grand coupling to construct random partitions of the vertex set according to common outputs, so that the average surface-to-volume ratio of cells is bounded by twice the average total variation distance between neighbors.

5.3 Proof of Theorem 5

Given Proposition 2 and Proposition 3, we are ready to prove the main result of this section.

Proof of Theorem 5. Since each X_i is measurable with respect to the sigma-algebra generated by $(Z_j)_{j \in B_r(i)}$, by Proposition 2 we can assign to each X_i a probability measure $\mu_i = \mu_{X_i}$ on N such that $\mu_i(B_r(i)) = 1$ and such that for any pair of vertices $i, j \in N$, the total variation distance $d_{\text{TV}}(\mu_i, \mu_j)$ is at most $\sqrt{\mathbb{E}[(X_i - X_j)^2]}$. Using the grand coupling of Proposition 3 we can couple all these measures in a process $(L_i)_{i \in N}$ taking values in N such that L_i has law μ_i and

$$\mathbb{P}[L_i \neq L_j] \leq 2d_{\text{TV}}(\mu_i, \mu_j) \leq 2\sqrt{\mathbb{E}[(X_i - X_j)^2]}.$$

Let D be the (random) subset of E given by

$$D = \{(i, j) \in E : L_i \neq L_j\}.$$

Then L_i is constant on each connected component of $G' = (N, E \setminus D)$. Since $\mu_i(B_r(i)) = 1$ it follows that $L_i \in B_r(i)$, and so the connected component of i is a subset of $B_r(L_i)$.

Finally,

$$\mathbb{E}[|D|] = \sum_{(i,j) \in E} \mathbb{P}[L_i \neq L_j] \leq 2 \sum_{(i,j) \in E} \sqrt{\mathbb{E}[(X_i - X_j)^2]} \leq \sqrt{8\epsilon}|E|,$$

where the last inequality follows from Jensen's inequality. \square

Remark 2. Both the results and proof of Theorem 5 is reminiscent of the methods of [Csóka, Harangi, and Virág \(2020\)](#), who worked under a “uniform for small sets” notion of non-amenable rather than our weaker “statistical” notion of non-amenable (i.e., non-hyperfiniteness). The authors of that paper work primarily with the entropy rather than the variance, but discuss entropy versions of their inequalities in Section 4.2. Their method is somewhat different than ours and does not identify the contraction property of Proposition 2.

6 Conclusion

In this paper, we identify amenability as a geometric condition that captures the possibility of local coordination on a social network. Theorems 1, 2 and 3 show that on amenable graphs it is possible to achieve low inefficiency, while Theorem 4 shows that low inefficiency implies that a graph is amenable.

Theorem 5 is a general result stating that if a finite graph admits local random variables that are highly correlated, then it must be amenable. Given the generality of this result,

we expect it to have further applications. For example, if we consider a repeated game version of our setting, then agents' actions at time t would depend on their neighbors in a ball of radius t , and so on non-amenable graphs high correlation is impossible. This still holds in even more general settings that go beyond pure coordination games. For example, agents could have idiosyncratic preferences over the actions, i.e., independent private types that influence their payoffs, making them prefer to coordinate on a particular action.

A question that we leave unanswered is whether leader equilibria are always optimal. Equivalently, whether (ε, r) -amenability coincides with having ε inefficiency in equilibria with a radius of communication r . For cycle graphs, we show that this holds in §4, but the technique used there does not easily generalize. A possible avenue for tackling this problem could be through applying Shapley values to the L^1 norm, rather than the L^2 norm that we use. That is, letting the influence distribution be given by the Shapley values of the game $v'_X(S) = \|\mathbb{E}[X|Z_S]\|_1$, rather than our current definition (2) of $v_X(S) = \|\mathbb{E}[X|Z_S]\|_2^2$. For these Shapley values we do not know how to prove a contraction theorem corresponding to Proposition 2.

References

D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar. Bayesian learning in social networks. *The Review of Economic Studies*, 78(4):1201–1236, 2011.

D. Aldous and R. Lyons. Processes on unimodular random networks. *Electron. J. Probab.*, 12:no. 54, 1454–1508, 2007. ISSN 1083-6489. doi: 10.1214/EJP.v12-463. URL <https://doi.org/10.1214/EJP.v12-463>.

N. Alon, P. Seymour, and R. Thomas. A separator theorem for graphs with an excluded minor and its applications. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing*, pages 293–299, 1990.

O. Angel and Y. Spinka. Pairwise optimal coupling of multiple random variables. *arXiv preprint arXiv:1903.00632*, 2019.

Á. Backhausz and B. Virág. Spectral measures of factor of iid processes on vertex-transitive graphs. 2017.

C. Ballester, A. Calvó-Armengol, and Y. Zenou. Who's who in networks. wanted: The key player. *Econometrica*, 74(5):1403–1417, 2006.

I. Benjamini, R. Lyons, Y. Peres, and O. Schramm. Group-invariant percolation on graphs. *Geometric & Functional Analysis GAFA*, 9(1):29–66, 1999.

I. Benjamini, O. Schramm, and A. Shapira. Every minor-closed property of sparse graphs is testable. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 393–402, 2008.

Y. Bramoullé, R. Kranton, and M. D’amours. Strategic interaction and networks. *American Economic Review*, 104(3):898–930, 2014.

A. Z. Broder. On the resemblance and containment of documents. In *Proceedings. Compression and Complexity of SEQUENCES 1997 (Cat. No. 97TB100171)*, pages 21–29. IEEE, 1997.

M. S.-Y. Chwe. Communication and coordination in social networks. *The Review of Economic Studies*, 67(1):1–16, 2000.

E. Csóka, V. Harangi, and B. Virág. Entropy and expansion. *Annales de l’Institut Henri Poincaré-Probabilités et Statistiques*, 56(4):2428–2444, 2020.

R. Durrett. Interacting particle systems: Ideas, techniques, applications. 2025.

G. Elek. The combinatorial cost. *arXiv preprint math/0608474*, 2006.

G. Ellison. Learning, local interaction, and coordination. *Econometrica: Journal of the Econometric Society*, pages 1047–1071, 1993.

P. Galicza and G. Pete. Sparse reconstruction in spin systems. I: iid spins. *Israel Journal of Mathematics*, 262(1):43–96, 2024.

B. Golub and M. O. Jackson. How homophily affects the speed of learning and best-response dynamics. *The Quarterly Journal of Economics*, 127(3):1287–1338, 2012.

J. C. Harsanyi. A simplified bargaining model for the n-person cooperative game. In *Papers in game theory*, pages 44–70. Springer, 1982.

S. Hoory, N. Linial, and A. Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006.

T. Hutchcroft. Continuity of the ising phase transition on nonamenable groups. *Communications in Mathematical Physics*, 404(1):227–286, 2023.

T. Hutchcroft and I. M. Lopez. A relation between isoperimetry and total variation decay with applications to graphs of non-negative ollivier-ricci curvature. *arXiv preprint arXiv:2411.04988*, 2024.

E. Ising. Beitrag zur theorie des ferromagnetismus. *Zeitschrift für Physik*, 31(1):253–258, 1925.

M. O. Jackson and L. Yariv. Diffusion of behavior and equilibrium properties in network games. *American Economic Review*, 97(2):92–98, 2007.

A. S. Kechris and B. D. Miller. *Topics in orbit equivalence*. Number 1852. Springer Science & Business Media, 2004.

J. Kleinberg and E. Tardos. Approximation algorithms for classification problems with pairwise relationships: Metric labeling and markov random fields. *Journal of the ACM (JACM)*, 49(5):616–639, 2002.

S. Kumabe and Y. Yoshida. Lipschitz Continuous Allocations for Optimization Games. In K. Bringmann, M. Grohe, G. Puppis, and O. Svensson, editors, *51st International Colloquium on Automata, Languages, and Programming (ICALP 2024)*, volume 297 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 102:1–102:16, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. ISBN 978-3-95977-322-5. doi: 10.4230/LIPIcs.ICALP.2024.102. URL <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ICALP.2024.102>.

T. Lindvall. W. Doeblin, 1915-1940. *The Annals of Probability*, 19(3):929–934, 1991.

R. J. Lipton and R. E. Tarjan. Applications of a planar separator theorem. In *18th Annual Symposium on Foundations of Computer Science (sfcs 1977)*, pages 162–170. IEEE, 1977.

R. J. Lipton and R. E. Tarjan. A separator theorem for planar graphs. *SIAM Journal on Applied Mathematics*, 36(2):177–189, 1979.

S. Morris. Contagion. *The Review of Economic Studies*, 67(1):57–78, 2000.

E. Mossel, A. Sly, and O. Tamuz. Strategic learning and the topology of social networks. *Econometrica*, 83(5):1755–1794, 2015.

J. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951.

A. B. Owen. Sobol’ indices and Shapley value. *SIAM/ASA Journal on Uncertainty Quantification*, 2(1):245–251, 2014.

A. B. Owen and C. Prieur. On Shapley value for measuring importance of dependent inputs. *SIAM/ASA Journal on Uncertainty Quantification*, 5(1):986–1002, 2017.

M. Pęski. Random utility coordination games on networks. *Theoretical Economics*, 20(2):583–622, 2025.

E. Sadler. Diffusion games. *American Economic Review*, 110(1):225–270, 2020.

O. Schramm. Hyperfinite graph limits. *Electron. Res. Announc. Math. Sci.*, 15:17–23, 2008. ISSN 1935-9179.

A. Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979.

J. von Neumann. Zur allgemeinen theorie des masses. *Fundamenta Mathematicae*, 13: 73–116, 1929.

D. J. Watts and S. H. Strogatz. Collective dynamics of ‘small-world’ networks. *nature*, 393 (6684):440–442, 1998.

S. Weidenholzer. Coordination games and local interactions: a survey of the game theoretic literature. *Games*, 1(4):551–585, 2010.

A Proof of Theorem 3

An important tool in the proof of Theorem 3 is *secret-sharing*, which allows us to use public messages to choose for each subgraph F' of radius at most r an independent, uniform $[0, 1]$ random variable, in an incentive compatible way.

Let ν be the probability measure on the message space $\mathcal{M} = \mathcal{A}^{\mathbb{N}}$ given by choosing each letter independently, from the distribution that assigns probability $1/2$ to \emptyset and probability $1/4$ to both $+1$ and -1 . Given a message $m_i = (\alpha_1, \alpha_2, \dots)$ chosen from ν , we can map it to a message distributed uniformly on $[0, 1]$ by $f(m_i) = \sum_{n=1}^{\infty} 2^{-n}v(\alpha_n)$, where $v(\emptyset) = 0$, $v(-1) = v(+1) = 1$. This mapping is invariant to negation, i.e., $f(-m_i) = f(m_i)$. By a similar construction, we can in fact map m_i to a sequence (m_i^1, m_i^2, \dots) of i.i.d. uniform $[0, 1]$ messages. For our purposes we will not need an infinite sequence, and instead will map m_i to a tuple $(m_i^0, (m_i^{F'})_{F'})$, where F' ranges over all radius most r connected subgraphs of G that contain i and at least one other agent. That is, we interpret a message as consisting of an initial number m_i^0 and another number $m_i^{F'}$ for each non-trivial subgraph i participates in.

Lemma 1. *Suppose that each agent i sends a public message $m_i = (m_i^0, (m_i^{F'})_{F'})$ as above. For each connected subgraph $F' = (V', E')$ of radius at most r and having at least two vertices, let $Z_{F'} = \sum_{i \in V'} m_i^{F'} \bmod 1$. Then the collection $Z_{F'}$ is i.i.d. uniform on $[0, 1]$, and furthermore it is i.i.d. uniform on $[0, 1]$ even when conditioned on any message m_i .*

The proof is immediate, relying on the fact that if X is uniform on $[0, 1]$, then so is $X + a \bmod 1$ for any $a \in \mathbb{R}$.

Given this lemma, we can define our equilibrium strategies. We fix a connected subgraph F of radius at most r , to be determined later. Given a subgraph F' , we write \bar{F}' for the

subgraph F' , together with the edges in $\partial F'$, and the nodes contained in these edges. I.e., \bar{F}' is F' , expanded to include the nodes connected to it. We call it the expanded graph of F' .

All agents will send public messages $m_i = (m_i^0, (m_i^{F'})_{F'})$ as above. As in Lemma 1, let $Z_{F'} = \sum_{i \in V'} m_i^{F'} \bmod 1$. We fix a probability $p \in [0, 1]$, also to be determined later.

Given this, we define a random set of edges $D = D_1 \cup D_2$, as follows. Let F' be a subgraph isomorphic to F and such that \bar{F}' is isomorphic to \bar{F} . We say that F' is *selected* if $Z_{F'} \leq p$. Lemma 1 thus implies that the event that various F' 's are selected are independent, have probability p , and moreover are still independent and have probability p conditioned on any messages sent by an agent i .

The set D_1 is the union of all the boundaries of the selected subgraphs. The set D_2 is the set of all edges of all agents i that do not belong to any selected F' . Let C_1, C_2, \dots be the connected components of the social network graph, with the set of edges D removed: $(V, E \setminus D)$. We say that $\ell_k \in C_k$ is the leader of C_k if $m_{\ell_k}^0$ maximizes $\{m_i^0 : i \in C_k, C_k \subseteq B_r(i)\}$. I.e., the leader of each community is the member whose first-coordinate message is highest.

Note that each agent, after observing the public messages, knows which community C_k they are in, and who its leader is, since the radius of observation is $2r$. For the same reason, leaders can send private messages to all the members of their communities. Indeed, in the next stage, they send uniform $\{-1, +1\}$ private messages to their community members, and in the last stage, members follow these recommendations when choosing their actions. As in previous constructions, all private messages not in $\{-1, +1\}$ are ignored and lead to actions being chosen uniformly at random. Public messages in an unexpected format are likewise ignored and result in the same choice of actions.

Regardless of the choice of F and p , this strategy profile is an equilibrium. That there is no incentive to deviate from the public message distribution follows from Lemma 1. The remaining elements of the profile are as in previously discussed leader equilibria, and hence again admit no profitable deviations. We also note that this strategy profile is action-symmetric. It is also network-symmetric, as its definition does not involve the identity of an agent, or their location in the graph.

The efficiency of this equilibrium depends on the choice of F and p . To complete the proof of Theorem 3, we show that we can choose F and p to get the required efficiency.

Since the network is an (ε, r) -amenable graph, it can be (deterministically) partitioned into connected components $\{C_1, \dots, C_K\}$ of radius at most r , and so that the set of edges E' connecting different components is of size $|E'| \leq \varepsilon|E|$. Then,

$$\frac{|E'|}{|N|} = \frac{1}{|N|} \sum_k |\partial C_k| = \frac{1}{|N|} \sum_k \frac{|\partial C_k|}{|C_k|} |C_k|,$$

which is the average surface-to-volume ratio weighted by component size. Thus, there exists at least one component C such that

$$\frac{|\partial C|}{|C|} \leq \frac{|E'|}{|N|} \leq \frac{\varepsilon|E|}{|N|} = \varepsilon d_{max}.$$

We let F be this connected component.

Let \mathcal{K} denote the family of subgraphs F' of G that are isomorphic to F , and such that \bar{F}' is isomorphic to \bar{F} . Let M_i be the number of subgraphs in \mathcal{K} that contain i . Let M be the number of subgraphs in \mathcal{K} that contain i . Since G is $(2r+1)$ -locally transitive, every $F' \in \mathcal{K}$ that contains i is contained in the ball of radius $2r+1$ around i . Hence this number is the same for every vertex; denote it by $M := M_i$. Note that $|\mathcal{K}| = M|N|/|F|$.

Let

$$p = \min\left\{\frac{1}{M} \log \frac{1}{\varepsilon}, 1\right\}$$

We end the proof of Theorem 3 by showing that with this choice of F and p we have the required inefficiency.

Let \mathcal{K}' be a random subgraph of \mathcal{K} obtained by selecting each $K \in \mathcal{K}$ independently with probability p . The expected number of boundary edges of the selected subgraphs, $D_1 = \bigcup_{K \in \mathcal{K}'} \partial K$, is bounded above by the expected number of selected sets times the number of boundary edges in each, that is

$$\mathbb{E}[|D_1|] \leq p \cdot |\mathcal{K}| \cdot |\partial F| = p \cdot \frac{|E| \cdot M}{d_{max} \cdot |F|} \cdot |\partial F| \leq p \cdot M \cdot \varepsilon \cdot |E|.$$

Next, the probability that a vertex i is not contained in any selected $K \in \mathcal{K}'$ equals

$$\mathbb{P}[i \notin K, \forall K \in \mathcal{K}'] = (1-p)^M.$$

Given that each vertex has degree d_{max} , the expected number of edges whose both endpoints do not belong to any selected K – denoted by D_2 – is bounded by

$$\mathbb{E}[|D_2|] \leq (1-p)^M \cdot d_{max} \cdot |N| = (1-p)^M \cdot |E|.$$

Therefore, the expected total number of deleted edges, $D = D_1 \cup D_2$, satisfies

$$\mathbb{E}[|D|] \leq p \cdot M \cdot \varepsilon \cdot |E| + (1-p)^M \cdot |E|.$$

Since $p = \min\{\frac{1}{M} \log \frac{1}{\varepsilon}, 1\}$, we obtain the desired bound on the number of deleted edges. Particularly, if $\frac{1}{M} \log \frac{1}{\varepsilon} \leq 1$, then

$$(1-p)^M = \left(1 - \frac{1}{M} \log \frac{1}{\varepsilon}\right)^M \leq e^{-\log(1/\varepsilon)} = \varepsilon.$$

Otherwise, $p = 1$, in which case $(1-p)^M = 0 \leq \varepsilon$. This completes the proof of Theorem 3.

Remark 3. *Benjamini, Schramm, and Shapira (2008) apply a similar technique to general graphs (rather than just to r -transitive ones, as we do) with the goal of showing that (ε, r) -amenability—i.e., the existence global, deterministic partition into radius at most r connected communities using an ε fraction of the edges—implies the existence of a local randomized rule that generates the same kind of partition, but (as in our case) using more edges. Their rule is somewhat involved, making it less plausible as a mechanism for coordination.*

B Stability of the Shapley value

Proposition 2 shows a contraction property of the Shapley influence distribution map. This motivates the question of the contraction properties of the Shapley value.

Let N be a finite set of players, and let $\mathbb{R}^{2^N \setminus \emptyset}$ be the collection of cooperative games for these players. The Shapley value is a map $\varphi: \mathbb{R}^{2^N \setminus \emptyset} \rightarrow \mathbb{R}^N$. Contraction properties of the Shapley value capture how stable it is to changes (or measurement errors) in the game. Let v, w be games, and suppose $|v - w|_\infty = \max_S |v(S) - w(S)|$ is small. How different are the Shapley values of v and w ?

There can be a number of ways to quantify the difference between the Shapley values. One is via the sup-norm on the Shapley values: $|\varphi(v) - \varphi(w)|_\infty = \max_i |\varphi_i(v) - \varphi_i(w)|$. Another is by the total variation norm

$$d_{\text{TV}}(\varphi(w), \varphi(v)) = \frac{1}{2} |\varphi(w) - \varphi(v)|_1.$$

When the former is small, the error in the value of each player is small. When the latter is small, the error in each group of players is small, which is a stronger notion. Proposition 2 yields the latter type estimate, although in a restricted setting that does not apply to all games.

Since φ is a linear map, understanding the deviations around a game v is the same as understanding the deviations around zero. We thus define the operator norms

$$\|\varphi\|_{\infty \rightarrow \infty} = \max\{|\varphi(v)|_\infty : |v|_\infty \leq 1\}.$$

and

$$\|\varphi\|_{\infty \rightarrow 1} = \max\{|\varphi(v)|_1 : |v|_\infty \leq 1\}.$$

It is easy to see that $\|\varphi\|_{\infty \rightarrow \infty} = 1$. The upper bound $\|\varphi\|_{\infty \rightarrow \infty} \leq 1$ follows immediately from the definition of the Shapley value, and the matching lower bound is achieved by the game in which $v(S)$ is one if $i \in S$ and zero otherwise, for some fixed player i .

One can interpret this fact as a demonstration of the stability of the Shapley value in the sup-norm sense: if we change v by at most ε in each coordinate, the Shapley value of

any player will not change by more than ε . In contrast, Kumabe and Yoshida (2024) show that the Shapley value is far from stable in the total-variation sense: a small change in the game can cause a group of players to have a very large change in their total value—when the number of players is large. Specifically, they show that there’s a universal constant $C > 0$ such that $\|\varphi\|_{\infty \rightarrow 1} \geq C \log |N|$, and in particular $\|\varphi\|_{\infty \rightarrow 1}$ goes to infinity as the set of players becomes larger. This result also contrasts Proposition 2, in which the bound is independent of $|N|$.