Contagion and Efficiency ⋆

Carlos Alós-Ferrer

University of Konstanz, Department of Economics, Box 150, D-78457 Konstanz, Germany. Phone (+49-7531) 88-2340, Fax (+49-7531) 88-4119

Simon Weidenholzer

Department of Economics, University of Vienna, Hohenstaufengasse 9, A-1010 Vienna, Austria. Phone (+43-1) 4277 37424, Fax (+43-1) 4277 37495

⋆ We thank Sanjeev Goyal, Stephen Morris, Larry Samuelson, and an anonymous referee for helpful comments. Financial support from the Austrian Science Fund (FWF) under Project P18141-G09 is gratefully acknowledged.

Email addresses: carlos.alos-ferrer@uni-konstanz.de (Carlos Alós-Ferrer), simon.weidenholzer@univie.ac.at (Simon Weidenholzer).

Preprint submitted to Journal of Economic Theory *23 January 2008*

Abstract

We consider a population of agents, either finite or countably infinite, located on an arbitrary network. Agents interact directly only with their immediate neighbors, but are able to observe the behavior of (some) other agents beyond their interaction neighborhood, and learn from that behavior by imitating successful actions. If interactions are not "too global" but information is fluid enough, we show that the efficient action is the only one which can spread contagiously to the whole population from an initially small, finite subgroup. This result holds even in the presence of an alternative, $\frac{1}{2}$-dominant action.

Keywords: Local Interaction Games, Learning, Imitation
JEL Classification Numbers: C72, D83

1. Introduction

Most day-to-day economic interactions are local in nature, meaning simply that the set of agents who interact with any given economic agent is typically small. It is equally clear, though, that agents routinely collect information which does not originate from their economic partners.

For example, firm managers typically interact with other managers from the same or related branches. Business wisdom, however, can often be readily translated from one branch to another, and successful ideas often spread beyond the branch they originated, with the transmitting agents often being consultants or simply specialized journals. This phenomenon is often referred to as "best practices" or "benchmarking". It is also easy to think of the increased effectiveness of marketing campaigns when a "role group" is targeted first. Further, if we think of university departments, research institutions, or similar departments in different firms, it becomes clear that, while most day-to-day interactions of a given individual are with coworkers, the onset of the internet has enormously increased the importance of information flows beyond the boundaries of such groups.

On a less abstract level, it is clear that agents are often well-informed about "second-level partners", i.e. friends of friends or partners of partners. An economic agent is rarely in complete ignorance about the *other* interactions that one of his or her partners is involved in. This does not mean, though, that every economic agent is perfectly informed about all interactions in the economy; information is still a local matter.

The idea of this paper is to capture these phenomena and study their consequences by making the distinction between local interaction and information explicit; i.e. by modeling not only whom does an agent interact with, but also which agent he receives information from.

Formally, we consider a countable population of agents located on an arbitrary network. Each agent interacts directly only with his or her immediate neighbors, resulting in a *local interaction game*. In contrast to the received literature, we assume that agents are able to observe the behavior of (some) other agents beyond their interaction neighborhood and learn from that behavior.

We also depart from (most of) the literature in two further respects. First, we consider arbitrary networks as in e.g. Morris [30], while most of the literature has concentrated on lattices. Second, whereas most of the literature has concentrated on behavioral rules related to best-reply (see below), we assume here that agents are willing to adopt actions perceived as successful, even though they are not completely aware of why they are successful. That is, we adopt a behavioral rule based on imitation. Shifting the emphasis away from best-reply behavior allows us to make a clear distinction between information and interaction. Further, as observed by Eshel, Samuelson, and Shaked [20], imitation behavior can give rise to interesting dynamics in the presence of local interactions.

We find that, under very general conditions, the fact that the information neighborhood extends (slightly) beyond the interaction neighborhood is enough to allow agents to coordinate on actions leading to Pareto-efficient equilibria, even in the presence of alternative, $\frac{1}{2}$-dominant ones.

The conditions we identify are of independent interest and might shed light on questions like why do inefficient technologies survive. Essentially, efficient actions are selected provided information is "fluid enough", as captured by a connectedness assumption. For finite networks, it is further required that the network is "large enough" relative to the size of interaction neighborhoods. Specifically, the size of the (smallest) interaction neighborhoods in the network should be small relative to the (maximum) number of pairwise disjoint neighborhoods in which the network can be partitioned (in a sense, interaction should not be "too global"). Thus, a transition to an efficient technology, business practice, or social convention might be triggered by engineering small close-knit groups (business parks, role groups, experimental communities, etc), provided they are visible enough.

Related Literature

Our work is of course related to the literature on 'learning from neighbors' (e.g. Ellison and Fudenberg [16,17], Bala and Goyal [5], and Banerjee and Fudenberg [6]). This literature has studied conditions under which efficient actions are eventually adopted by the whole population, if agents receive information only from neighbors. In these models, agents face a common individual problem (a game against nature), thus there is no strategic interaction.

Local interaction games have been examined by Blume [11], Ellison [14,15], Anderlini and Ianni [2], and Morris [30], among many others. In these papers, agents are assumed to adopt (myopic) best-replies to their neighbors' actions. Essentially, the findings of this literature support the selection of $\frac{1}{2}$-dominant actions, [1] in line with standard results from the literature of learning in games (e.g. [26] or [34]).

With the exception of Morris [30], most of this literature has concentrated on lattices. [2] [30] considers (infinite) arbitrary networks, and asks the question of when will a given action be able to take over an infinite network starting from a finite subgroup of users. There are three main differences between his work and ours. First, [30] concentrates on myopic best reply, whereas we work with imitation rules. Second, our main point is the distinction between interaction and information neighborhoods; this point is absent in [30], because such a distinction is hard to motivate under myopic best reply. [3] Third, [30] obtains a number of additional assumptions on the network shape which guarantee the long-run prevalence of the risk-dominant action; in contrast, we find that the efficient action will be selected under much weaker assumptions on the network.

Eshel, Samuelson, and Shaked [20] consider a learning model of imitation on the circular city, and show that imitation might lead to the survival of (strictly dominated) altruistic behavior. We share with [20] an interest on imitation rules. One difference between our work and theirs is of course that we study arbitrary networks; however, we will

[1] An action is $\frac{1}{2}$-dominant if it is a best-response when half of an agent's neighbors adopt it; when there are only two available actions, this corresponds to risk-dominant equilibria.

[2] Blume [10,12] and Young [35, Chapter 6] study coordination games played in general networks, where players adopt best responses perturbed according to the logit choice function. Under appropriate assumptions, the induced *spatial game* has an exact potential, where the profile where all players adopt the risk-dominant action is stochastically stable by virtue of being the only potential maximizer. See also [7] for *asymmetric* networks.

[3] [24] and [22] take a different approach, which is nevertheless conceptually related to ours. They consider models where players only know the number of neighbors in their interaction neighborhood and the distribution of the number of neighbors in the overall network. This incomplete information setup naturally leads to the use of Bayesian Nash equilibrium as a (static) solution concept.

use the circular city as a motivating example. Another difference is the specific behavioral rule. [20] use a rule based on average payoffs, while we concentrate on an imitate-the-best rule which seems to be better founded on empirically observed human behavior. Last, while [20] study the survival of altruism (Prisoner's Dilemma games), we focus on the selection of efficient actions (coordination games and games with an efficient dominant strategy).

Interestingly, there is partial experimental evidence in favor of disentangling information and interaction neighborhoods. Apesteguía, Huck, and Oechssler [4] study imitation in a Cournot oligopoly experiment when subjects face different information structures. In one of the treatments, subjects have both information on their direct competitors and on other subjects they are not interacting with. In this setup, the data shows a fair amount of imitation *outside* the interaction neighborhood.

The paper is organized as follows. Section 2 describes the basic building blocks of our model. Section 3 introduces the idea of a contagious action and shows the implications for finite and infinite networks. In order to develop the intuition, Section 4 discusses a motivating example in detail. Section 5 considers general networks and presents the main result. Section 6 illustrates the results through several additional examples. Section 7 briefly pursues a few extensions and robustness tests.

2. The Model

2.1. *The Base Game*

We are interested in the question of whether an *efficient action* will spread contagiously to a large network. For this purpose, it is enough to consider a symmetric 2×2 coordination game[4]

	A	B
A	a, a	b, c
B	c, b	d, d

where $a > b, c, d$, so that (A, A) is a Pareto-efficient strict Nash equilibrium.

Essentially, this includes two types of games. If $b > d$, A is an efficient dominant strategy. If $b < d$, the game is a symmetric coordination game where both (A, A) and (B, B) are strict Nash equilibria. Then, if $a + b < c + d$, (B, B) is risk dominant in the sense of Harsanyi and Selten [23], while (A, A) is risk dominant if the reverse inequality holds.

Without loss of generality we can normalize this game to obtain

	A	B
A	1	0
B	α	β

(G)

[4] Our results readily extend to a framework with more than two actions, provided there is an efficient one.

where $\alpha = \frac{c-b}{a-b}$ and $\beta = \frac{d-b}{a-b}$. Note that $\alpha, \beta < 1$ (the notation follows [20]). Efficient dominant strategy games correspond to $\beta < 0$, while coordination games are given by $0 < \beta < 1$. In the latter, (B, B) is risk dominant if $\alpha + \beta > 1$, which in turn implies $\alpha > 0$.

2.2. *Interaction*

A *local interaction system* consists of a countable [5] population of agents, such that each of them interacts with a finite subset of the population only. Formally,

Definition 2.1. A local interaction system is a pair $\left(I, (K(i))_{i \in I}\right)$ where I is a countable set of players and $K(i)$ is a subset of I for each $i \in I$ such that
(K1) Irreflexivity: for all $i \in I$, $i \notin K(i)$.
(K2) Symmetry: for all $i, j \in I$, $j \in K(i) \Rightarrow i \in K(j)$.
(K3) Bounded Neighborhoods: there exists $Q > 0$ such that $|K(i)| \leq Q$ for all $i \in I$.

$K(i)$ is called the *interaction neighborhood* of i. If $j \in K(i)$, we say that j is a neighbor of i. In words, irreflexivity requires that no player is his own neighbor. Symmetry states that, if j is a neighbor of i, then i is a neighbor of j. Bounded Neighborhoods is a vacuous condition if the population is finite. For an infinite population, it spells out the requirement that each agent has a relatively small number of neighbors.

Most local interaction systems of interest will additionally be connected, meaning that, for any pair of players, there is some path connecting them; that is, iterating $K(\cdot)$ starting from an arbitrary agent will eventually cover the whole population. Formally,
(K4) Connectedness: there exists $\{i_1, i_2, \ldots, i_L\} \subseteq I$ such that $L \geq 1$, $i_1 = i, i_L = j$ and $i_{l+1} \in K(i_l)$ for each $l = 1, \ldots, L-1$.
We will not, however, assume this additional property. This allows us to encompass *location models*, where agents can interact at a number of alternative, predetermined locations.

Our definition focuses on the sets $K(i)$ for notational convenience. Of course, we could alternatively define a local interaction system as e.g. in [30], i.e. a pair $(I, \sim)$ with I a countable set of agents and $\sim$ a binary relation on I, assumed irreflexive, and symmetric. We obtain then the interaction neighborhood as $K(i) \equiv \{j \in I \mid j \sim i\}$. A further alternative is to focus on the links and treat the local interaction system as a graph; see e.g. [25].

We will consider a dynamical model where, each period, each agent commits to an action and plays the base game (G) against all players in his interaction neighborhood. We assume that, for payoff comparisons, each player is only concerned about (relative) average payoffs, so the size of the interaction neighborhood does not matter. That is, players concentrate on the payoff *per interaction.* [6] So, if $\omega = (s_i)_{i \in I}$ is the profile of strategies adopted by players, the total payoff for player i is

[5] I.e. either finite or countably infinite. Most local interaction models focus on finite populations. [30] considers countably infinite ones.
[6] Since the network is fixed, the number of interactions a player is involved in is not subject to change. Hence, concentrating on per-interaction payoffs seems a reasonable assumption.

$$U(i,\omega) = \frac{1}{|K(i)|} \sum_{j \in K(i)} u(s_i, s_j)$$

where $u(s_i, s_j)$ denotes the payoff of playing s_i against s_j in game (G).

2.3. *Information*

Each agent observes only the actions adopted and the payoffs obtained by all agents in an *information neighborhood* $M(i)$. We explicitly assume $K(i) \bigcup \{i\} \subseteq M(i)$, so that an agent observes at least his own play and the pattern of play within his interaction neighborhood. Formally,

Definition 2.2. An information system for a local interaction system $\left(I, (K(i))_{i \in I}\right)$ is a collection $(M(i))_{i \in I}$ such that, for all $i, j \in I$,
(M1) Observed Play: $K(i) \bigcup \{i\} \subseteq M(i)$.
(M2) Symmetry: $j \in M(i) \Rightarrow i \in M(j)$.
The triple $\left(I, (K(i))_{i \in I}, (M(i))_{i \in I}\right)$ is called a local interaction-information system.

2.4. *Behavior*

Each period $t = 0, 1, 2, \ldots$, after observing the strategies and payoffs in his interaction neighborhood, each agent switches to an action that has earned the highest payoff in the previous period in his information neighborhood. Ties are assumed to be broken randomly. [7] Formally, agent i chooses

$$s_{i,t} = s_{j,t-1} \quad \text{with } j \in \arg\max_{j' \in M(i)} U(j', \omega_{t-1})$$

where ω_{t-1} is the state of the system in the previous period.

We call this behavioral rule "imitate-the-best." It deserves a brief discussion. It prescribes to imitate the strategy that has yielded the highest payoffs in the information neighborhood. This rule can be seen as either naive [8] or optimistic, as agents focus on highest observed payoffs only. Several alternative rules come to mind immediately, for instance imitating the strategy which yields the highest *average* payoffs in the neighborhood (e.g. [20]). The imitate-the-best rule, though, nicely captures some experimentally observed features of actual human behavior, as reported by the psychology literature. These features arise from the well-documented human tendency to focus on *salient* outcomes, e.g. those leading to high payoffs. [9]

For instance, Barron and Erev [8] and Erev and Barron [19] report on a large number of decision-making experiments and identify several interesting effects which lead to deviations from payoff maximization. First, alternatives with the highest recent payoffs are attractive to decision makers, even if they are associated with a low expected return. This points out that imitate-the-best might be more realistic than decision rules based on

[7] For each fixed network, ties are nongeneric in the space of games.
[8] Note that, for large neighborhoods, the imitate-the-best rule poses a minimal computational burden on agents.
[9] In turn, salience constitutes a simple justification for management practices where attention is focused on the industry's top performers.

the imitation of average payoffs. Second, agents show a tendency to switch to strategies with the most attractive (observed) forgone payoffs. This illustrates that observing a high payoff weighs heavily in the human mind. Erev and Barron [19] summarize these observations in a phenomenon called *underweighting of rare events*, which is an observed tendency to imitate actions which lead to large payoffs most of the time even when they are associated with a lower expected payoff.

3. The basic concept

Having spelled out the basic building blocks of our model (interaction, information, and behavior), we now turn to the idea of contagious actions. Intuitively, we want to find out whether the efficient action will be able to spread out from an initially small group of agents to the whole population. Ideally, we would expect that the efficient action has an advantage over the inefficient one, in the sense of being able to spread out from a significantly smaller subset.

3.1. *Contagious Actions*

Denote the state space by Ω. We refer to the state $\overrightarrow{A} = (A, A, A, \ldots)$ as the *efficient convention*; analogously $\overrightarrow{B}$ is the inefficient one. Given $s \in \{A, B\}$, define the basin of attraction of the associated convention $\overrightarrow{s}$, $D(s)$, as the set of states ω such that the dynamics converges to $\overrightarrow{s}$ with probability 1. Then, for each subset $J \subseteq I$, define $\Omega(J, s)$ as the set of states $\omega = (s_i)_{i \in I}$ such that all players $j \in J$ play s. Let

$$\mathcal{J}(s) = \{ J \subseteq I \mid J \text{ finite and } \Omega(J, s) \subseteq D(s) \}$$

be the collection of finite population subgroups J such that s eventually spreads to the whole population if at some point all agents in J adopt s. Denote by $c^*(s) = \min_{J \in \mathcal{J}(s)} |J|$ the size of the smallest such group (if finite).

In order to define contagious action, we focus on two properties. The first requires that it should be able to spread to the whole population from an initially small subgroup. Formally this means that $\mathcal{J}(s)$ should be nonempty (in the infinite case) or $c^*(s)$ should be "small" (in the finite case). The second property requires that the convention should be resilient. That is, once established, it should not be possible to leave it starting with a small group of deviants. In the infinite case, this means that all states where a cofinite set of players adopts s should be in the basin of attraction of $\overrightarrow{s}$. In the finite case, the property translates into the requirement that, the size of such sets of successful deviants should be "large".

Definition 3.1. If I is infinite, we say that action s is contagious if $\mathcal{J}(s) \neq \emptyset$ and for any finite set $J \subseteq I$, $\Omega(I \setminus J, s) \subseteq D(s)$.

If I is finite, we say that s is contagious if $c^*(s) < r^*(s)$, where

$$r^*(s) = \min \{ |J| \mid \Omega(I \setminus J, s) \cap (\Omega \setminus D(s)) \neq \emptyset \}.$$

Alternatively, we could express the conditions for the infinite case as $c^*(s) < \infty$ and $r^*(s) = \infty$. In words, an action is contagious in an infinite local interaction-information system if, first, it can spread to the whole population from a finite subgroup, and second, once it is established, no finite subgroup of deviants will be able to move the population away from the efficient convention. It immediately follows that the contagious action is the only one which can spread to the whole population from a finite subgroup, i.e. if A is contagious then $\mathcal{J}(B) = \emptyset$. This is analogous to the concept studied by [30] and the discussion in [20].

In the finite case, the definition of a contagious action requires that it can spread from a *strictly* smaller subgroup than is needed to leave it once it is established. In particular, the contagious action can spread from a strictly smaller subgroup than any other action.

3.2. *The Finite Case*

For the finite population case, the idea of a contagious action allows an analysis akin to the dynamic models of Ellison [14,15] and others. Essentially, if the efficient action is contagious, it is a simple matter to flesh out specific dynamic models that will select it. As an example, we spell out the result for the well-known *mistakes model* introduced by [26] (see also [31] or [21]).

For finite I, the basic adjustment process gives rise to a finite Markov chain for which standard techniques apply. An *absorbing set* of this chain is a minimal subset of states which, once entered, is never abandoned. An *absorbing state* is an element which forms a singleton absorbing set.

There is of course a multiplicity of absorbing sets. For instance, both conventions $\overrightarrow{A}$ and $\overrightarrow{B}$ are absorbing states. For, if all agents play the same action, only that action will be observed.

In order to select among the absorbing states of the unperturbed learning model, in the mistakes model the dynamics is enriched with a perturbation in the form of experiments (or mistakes) as follows. With an independent probability $\varepsilon > 0$, each agent, in each period, might make a mistake (or "mutate"), and simply pick a strategy at random. [10]

The perturbed process is ergodic, i.e. it has a unique *invariant distribution* $\mu(\varepsilon) \in \Delta(\Omega)$. By the Ergodic Theorem and the Fundamental Theorem of Markov Chains (see e.g. [27]), this distribution summarizes both the long-run behavior of the process and the time-average behavior for any sample path, independently of initial conditions.

The *limit invariant distribution* (as the rate of experimentation tends to zero) $\mu^* = \lim_{\varepsilon\to 0} \mu(\varepsilon)$ exists and its support $\{\omega \in \Omega \mid \mu^*(\omega) > 0\}$ is the union of some absorbing sets of the unperturbed process. The limit invariant distribution singles out a stable prediction of the unperturbed dynamics ($\varepsilon = 0$), in the sense that, for any $\varepsilon > 0$ small enough, the play approximates that described by μ^* in the long run. The states in the support of μ^* are called *stochastically stable states* or *long-run equilibria* (LRE).

Our next result, whose proof is relegated to the appendix, shows that if the efficient action is contagious, the corresponding convention is uniquely selected in the mistakes model.

Theorem 1. *Let I be finite. If the efficient action A is contagious, then it is the unique stochastically stable state in the mistakes model.*

[10] In case of mistakes, strategies are picked up according to some pre-specified probability distribution having full support. The exact distribution does not affect the results, as long as it has full support, and does not depend on ε.

4. A motivating example

As in [14] and [20], we consider a population of N agents, $I = \{1, 2, \ldots, N\}$ arranged around a circle. To cover the case of an infinite population, we also consider the case of the infinite line $\mathbb{Z}$, as in [30]. We refer to this as the (finite or infinite) $2k-$*neighbors model*.

Player i has $i-1$ and $i+1$ as immediate neighbors. Agents play the game (G) against their $2k$ nearest neighbors in discrete time, $t = 1, 2, \ldots$. That is, the *interaction neighborhood* of player i is given by $K(i) = \{i-k, \ldots, i-1, i+1, \ldots, i+k\}$. In the finite case, $i \pm r$ is understood as modulo N; we further assume $k < \frac{N}{2}$.

Each period, each agent selects an action, i.e. a pure strategy of the base game, and plays according to this strategy against all of his neighbors. Thus, if $\omega = (s_1, ..., s_n)$ is the profile of strategies adopted by agents at time t, the (average) payoff for agent i is

$$U(i, \omega) = \frac{1}{2k} \sum_{j=1}^{k} [u(s_i, s_{i-j}) + u(s_i, s_{i+j})].$$

The *information neighborhood* of agent i is assumed to consist of his $2m$ nearest neighbors. i.e. $M(i) = \{i-m, \ldots, i, \ldots, i+m\}$, (modulo N and $m \leq \frac{N}{2}$ on the circle). (M1) then reduces to $m \geq k$, so that an agent at least observes the pattern of play within his interaction neighborhood. According to the imitate-the-best rule, an agent adopts an action that has earned the highest payoff in his information neighborhood in the previous period, i.e.

$$s_{i,t} = s_{j,t-1} \quad \text{where } j \in \arg \max_{j' \in M(i)} U(j', \omega_{t-1})$$

with ties broken randomly, where ω_{t-1} is the state of the system in the previous period.

Assume now that agents always receive information from beyond their interaction neighborhood. So $k < m$ and $K(i) \subset M(i) \setminus \{i\}$. First, consider the two neighbor model where players observe the behavior of their four closest neighbors. So $k = 1$ and $m = 2$. When revising strategies agents do not only consider what is happening within their interaction neighborhood but also take into account relative success of agents who are not direct opponents. The most important feature of this setup is that once efficient outcomes are established somewhere they can spread contagiously. This is similar to the spread of risk-dominant strategies in Ellison's [14] best-reply local interaction model.

The main reason for this result is that any state with three adjacent A-players lies in the basin of attraction of $\overrightarrow{A}$. To see this, consider any state with three adjacent A-players. In the worst case they are surrounded by B-players.

$$\ldots BBBAAABBB \ldots$$

The inner A-player earns a payoff of one. The outer A-players earn $\frac{1}{2}$ and the boundary B-players earn $\frac{\alpha+\beta}{2}$. The A-players and the boundary B-players observe that A earns a payoff of one (which is the maximum payoff). Hence, the A-players will retain their strategy and the boundary B-players will switch to A. In a next step the "new" boundary players will also change to A and so forth. In this manner A will extend to the whole population and we will eventually reach the state $\overrightarrow{A}$ (in the limit for the infinite population case), i.e. $c^*(A) = 3$.

By the observation above, in order to leave the state where everybody plays A we at least have to destabilize every possible A-cluster of size three. In the infinite case, this implies that it is not possible to leave the basin of attraction of the efficient convention through a finite number of deviations. Hence, in the infinite case A is contagious.

In the finite case destabilizing the efficient convention requires at least $\lceil \frac{N}{3} \rceil$ deviations.[11] This implies $r^*(A) \geq \lceil \frac{N}{3} \rceil$. For $N > 9$, $r^*(A) > c^*(A)$ holds. This finding shows that in the finite two-neighbor model with information neighborhood $m = 2$ the efficient convention is contagious.

5. General Networks

We consider now an arbitrary local connected interaction-information structure given by $(I, (K(i))_{i \in I}, (M(i))_{i \in I})$. Recall that $K(i)$ and $M(i) \supseteq K(i) \cup \{i\}$ denote the interaction and the information neighborhood, respectively.

5.1. *Neighbors' Neighbors*

Let us now generalize the above findings to general networks. In the two neighbor model the efficient action can spread contagiously when information extends beyond the narrow bounds of interaction. A natural way to generalize this property to general connected networks is to assume that agents observes the strategies and the payoffs obtained by their neighbors' neighbors: Whenever two agents interact they also exchange information on what is happening in their interaction neighborhoods. Formally,

Assumption 1.

$$M(i) \supseteq \bigcup_{j \in K(i)} K(j).$$

Note that in the $2k$-neighbor model Assumption 1 implies $m = 2k$.

With the help of Assumption 1 we are now able to state the following lemma.

Lemma 1. *Under Assumption 1, in connected networks any state where some player and all of his neighbors play A lies in the basin of attraction of $\overrightarrow{A}$.*

Proof. Assume that agent i and all of his neighbors play A. Agent i receives a payoff of one which is the largest possible payoff. By (M2), all agents in $M(i)$ observe this and hence will switch to A. Now all agents in $K(i)$ receive the maximum payoff of 1. Hence, in the next step, all agents in $M(K(i))$ will switch to A, which includes those in $K(K(i))$ (by (M1)). By Assumption 1, we conclude that the efficient strategy spreads contagiously until we reach $\overrightarrow{A}$. □

[11] B-players must be at most at distance two from each other. Rounding up reflects that the circular nature of the network imposes an additional B-player.

5.2. *An Efficiency Result*

Let $\mathcal{V}$ be the set of all population subsets whose neighborhoods are pairwise disjoint, i.e.

$$\mathcal{V} \equiv \left\{ V \subseteq I \;\middle|\; (K(i) \cup \{i\}) \bigcap (K(j) \cup \{j\}) = \emptyset \; \forall i, j \in V, \; i \neq j \right\}.$$

The maximum number of disjoint neighborhoods of a local interactions system is given by

$$w^* = \max \{ |V| \mid V \in \mathcal{V} \}.$$

Further, let

$$Q_{\min} \equiv \min \{ |K(i)| \mid i \in I \}.$$

be the size of the smallest interaction neighborhood in the network. Then, we can provide the following result for general networks.

Theorem 2. *Under Assumption 1, in any connected local interaction-information system the efficient convention is contagious under the imitate-the-best rule, provided $w^* > Q_{\min} + 1$ holds in the finite case.*

Proof. By Lemma 1, once the efficient strategy is played by a player and all of his neighbors we will move to $\overrightarrow{A}$. Since the number of neighbors will be different across players in general we choose the player with the fewest neighbors. If he and all his neighbors switch to A we will move to $\overrightarrow{A}$. Hence, we have $c^*(A) \leq Q_{\min} + 1$. Note that $Q_{\min} \leq Q$ by (K3), thus $c^*(A)$ is finite.

To move out of the basin of attraction of $\overrightarrow{A}$ we need to destabilize any cluster of A-players that is such that the maximum payoff of one is received by one player. If we have w^* disjoint neighborhoods we at least need one deviation in each of them. Hence, $r^*(A) \geq w^*$. So, if $w^* > Q_{\min} + 1$ efficient strategies are contagious. Note that, in the infinite case, $w^* = \infty$ by (K3). □

5.3. *Contacts*

Consider again the $2k$-neighbors model. Note that the above argument generally applies for $m > k$. Whenever the information neighborhood is larger than the interaction neighborhood the efficient strategy is able to spread contagiously. Formally

Lemma 2. *For $m > k$, any state with $2k + 1$ adjacent A players lies in the basin of attraction of $\overrightarrow{A}$.*

Proof. Consider any state with $2k + 1$ adjacent A-players. The middle A-player earns a payoff of one. The A-players and $2(m - k)$ boundary B-players observe that A is capable of earning this maximum payoff. Hence, the A-players will retain their strategy and the boundary B-players will switch to A. In a next step the "new" boundary agents will also change to A and so forth. Clusters of B-players will disappear and we will reach the state $\overrightarrow{A}$. □

As above we are able to conclude that

Theorem 3. *In the $2k$-neighbor model with information neighborhood $m > k$ the efficient strategy is contagious, provided $N > (2k+1)^2$ in the finite case.*

Proof. By Lemma 2, $c^*(A) = 2k+1$. Unless at least one agent out of each $2k+1$ adjacent agents plays B we will move to $\overrightarrow{A}$. That is $r^*(A) = \lceil \frac{N}{2k+1} \rceil$ on the circle and $r^*(A) = \infty$ on the line. For $N > (2k+1)^2$, $r^*(A) > c^*(A)$ holds. □

There are several features of the result for the $2k$-model which we will attempt to generalize. First, observe that $c^*(A)$ is independent of population size, while the complementary quantity $r^*(A)$ grows with the population. This nicely ties up the finite and the infinite cases. Second, the fundamental condition appears to be that agents do observe other agents beyond their interaction neighborhood ($m > k$).

Note that Assumption 1 requires $m = 2k$. However efficient actions are contagious in the $2k$-neighbors model whenever $m > k$ holds. Hence, there is hope of obtaining a weaker condition on information than Assumption 1.

It is an easy exercise to show to generalize the idea of information extending slightly beyond interaction (i.e. $m > k$) to *Euclidean Networks.*

Example 1. Euclidean Networks. Let $\leftrightarrow$ be a binary relation on a countable set I, assumed symmetric, irreflexive, and connected (i.e. $(I, \leftrightarrow)$ defines a local interaction system). For each $i, j \in I$, let $d(i,j)$ be the length of the shortest path connecting i and j. Clearly, d is a *discrete distance* on I, i.e. a mapping $d : I \times I \mapsto \mathbb{N}$ such that (i) $d(i,j) = 0$ if and only if $i = j$; (ii) $d(i,j) = d(j,i)$ for all $i, j \in I$; and (iii) $d(i_1, i_3) \leq d(i_1, i_2) + d(i_2, i_3)$ for all $i_1, i_2, i_3 \in I$ (the triangle inequality).

The relation $\leftrightarrow$ defines a reference local interaction system on which we now define another local interaction system. Fix $k, m \in \mathbb{N}$, $0 < k < m$, and define for each $i \in I$,

$$K(i) = \{j \in I \mid d(i,j) \leq k,\ j \neq i\}$$

and

$$M(i) = \{j \in I \mid d(i,j) \leq m\}\,.$$

It is then straightforward to check that $(I, (K(i))_{i \in I}, (M(i))_{i \in I})$ is a local interaction-information system, which is called a *Euclidean network.* For example, the $2k$-neighbor model is a Euclidean network where the relation $\leftrightarrow$ correspond to being an immediate neighbor, i.e. $i \leftrightarrow j$ if and only if $|i - j| = 1$. Examples can be built from any connected graph.

Theorem 4. *In Euclidean networks the efficient convention is contagious whenever $m > k$, provided $w^* > Q_{\min} + 1$ holds in the finite case.*

Proof. Consider a player i. If he and all his neighbors play A he will earn the maximum payoff of one. Hence, he and all of his neighbors will retain their strategy. Further, players up to a distance of m observe this and will switch to A. This implies that in a next step players in the set

$$K^*(i) = \{j \in I \mid d(i,j) \leq m - k\}$$

will earn the maximum payoff. In a next step all players in $M(K*(i))$ will switch to A. Since the network is connected A will spread to the entire population. As in Theorem 5, $c^*(A) \leq Q_{\min} + 1$ and $r^*(A) \geq w$ in the finite case. Note that, in the infinite case, $w^* = \infty$ by (K3). □

Note that in Euclidean networks m and k are the same across players. It is however easy to think of examples where this is not the case. To analyze this even more general class of local interactions models we introduce the notion of *contacts*. This allows us to generalize the idea of information slightly extending beyond interaction in an even broader class of local interactions models.

We define the set of *contacts* of agent i as those agents j such that all their interactions are with agents known to i. That is,

$$K^*(i) = \{j \in I \mid K(j) \cup \{j\} \subseteq M(i)\}.$$

With this definition, we immediately obtain

$$M(i) \supseteq \bigcup_{j \in K^*(i)} K(j).$$

That is, each agent observes what is happening in the interaction neighborhoods of all of his contacts. The underlying idea is that contacts correspond to "closest acquaintances." Agents are very likely to exchange information with them and hence will not only know what is happening in their interaction neighborhood, but will at least also have an idea what is going on in their close surrounding. For example, on the $2k$-neighbors model with $m = k + 1$, $K^*(i) = \{i - 1, i + 1\}$.

Notice, though, that the definition of $K^*(i)$ does *not* require $K^*(i) \subseteq K(i) \cup \{i\}$. That is, contacts need not be neighbors. As we will comment below, this allows treatment of examples belonging to the "multiple-location" class.

We introduce the following assumption.

Assumption 2. For each $i, j \in I$, there exists $\{i_1, i_2, \ldots, i_L\} \subseteq I$ such that $L \geq 1$, $i_1 = i, i_L = j$ and $i_{l+1} \in K^*(i_l)$ for each $l = 1, \ldots, L - 1$.

In other words, the relation "to be a contact of" is connected, so that iteration of $K^*(\cdot)$ eventually covers the whole population.

On the $2k$-neighbors model, Assumption 2 reduces to $m > k$. This assumption would e.g. be violated if the network consisted of *informationally* separated components, i.e. subnetworks such that agents in one of them neither interact with nor receive information from agents in other subnetworks.

To motivate both the concept of contacts and Assumption 2, consider again the example of Euclidean networks.

Fix a player i. His contacts are those agents j such that all their neighbors j' are in $M(i)$. That is, $j \in K^*(i)$ if and only if the following implication holds for all $j' \in I$:

$$d(j', j) \leq k \Rightarrow d(j', i) \leq m.$$

Since $d(j', i) \leq d(j', j) + d(j, i)$ by the triangle inequality, it follows that

$$K^*(i) \supseteq \{j \in I \mid d(j, i) \leq m - k\}$$

and, since $m - k > 0$ and $\leftrightarrow$ is connected, we conclude that Assumption 2 holds for any Euclidean network.

We discuss now a different family of examples where contacts are not necessarily neighbors, but Assumption 2 is fulfilled.

Example 2. Multiple Locations. Let I be a finite set and partition it into L different subsets or *locations*, $I_1, \ldots, I_L$. For each $i \in I$, define $K(i) = I_l \setminus \{i\}$, where I_l is the location which contains i, and $M(i) = I$. We refer to this as a *locations model with global information.* That is, interactions are confined to the locations, and are global within each location, but all interactions are publicly observed.[12] Then, $K^*(i) = I$ for all $i \in I$ and Assumption 2 follows. This corresponds to the *locations model* discussed e.g. in Anwar [3] and Ely [18].

With the help of Assumption 2 we are now able to state the following lemma.

Lemma 3. *Under Assumption 2, any state where some player and all of his neighbors play A lies in the basin of attraction of* $\overrightarrow{A}$.

Proof. Assume that agent i and all of his neighbors play A. Agent i receives a payoff of one which is the largest possible payoff. By (M2), all agents in $M(i)$ observe this and hence will switch to A. By construction, now all agents in $K(i)$ receive the maximum payoff of 1. Hence, in the next step, all agents in $M(K^*(i))$ will switch to A, which includes those in $K^*(K^*(i))$ (by (M1)). By Assumption 2, we conclude that the efficient strategy spreads contagiously until we reach $\overrightarrow{A}$. □

This Lemma will be the key for our efficiency result in the next section. In turn, the Lemma hinges upon Assumption 2. This assumption is the natural generalization of the idea that information neighborhoods should extend "at least a bit" beyond the interaction neighborhoods, in a way which allows information to flow throughout the network.

Assumption 2 represents a weak form of connectedness with respect to the information neighborhoods. Such an Assumption is unavoidable, since otherwise the network could consist of completely separated subsocieties.[13]

Although it is clear that Assumption 2 could be relaxed if we constrain ourselves to specific classes of networks, we want to remark that this assumption is already quite weak. For instance, in the local interaction structure of the $2k$-neighbors model, it would be still satisfied if we append information neighborhoods which extend beyond the interaction neighborhoods "on one side".

5.4. *A general efficiency result*

With the help of Lemma 3 we can provide the following result for general networks. The prove merely resembles the proof of Theorem 2 (using Lemma 3 instead of Lemma 1) and is omitted.

Theorem 5. *Under Assumption 2, in any local interaction-information system the efficient convention is contagious under the imitate-the-best rule, provided* $w^* > Q_{\min} + 1$ *holds in the finite case.*

[12] This could be weakened to a form of connectedness among locations for $M(i)$.

[13] In contrast, the selection of risk-dominant strategies in [30] requires connectedness plus two additional properties. The first, δ-*uniformity*, guarantees that the network is not "too lumpy". The second, *low-neighbor growth*, states that the number of players reached in k steps from any initial position does not grow exponentially in k.

This result allows an straightforward application to "regular networks".

Example 3. Interaction on Lattices.
Let $Q \in \mathbb{N}, Q > 0$. We say that a local interaction-information structure $(I, (K(i))_{i\in I}, (M(i))_{i\in I})$ is *Q-regular* if (i) $|K(i)| = Q$ for all $i \in I$; (ii) either I is infinite or $|I| = N$ and the maximum number of pairwise disjoint interaction neighborhoods is $w^* = \lfloor \frac{N}{Q+1} \rfloor$; and (iii) Assumption (A1) holds. Examples include Euclidean Networks on the circle, interactions on a torus, and multidimensional lattice structures.[14]

Theorem 5 then immediately yields

Corollary 1. *In any Q-regular local interaction-information system, the efficient convention is contagious, provided* $N > (Q + 1)^2$ *in the finite case.*

Notice that our result for the $2k$-neighbor model, Theorem 3, agrees with this result taking $Q = 2k$.
A further, straightforward application concerns location models.

Corollary 2. *Consider a location model with global information, and let* $I_1, \ldots, I_L$ *be the locations. The efficient convention is contagious, provided* $L > \min_{l=1,\ldots,L} |I_l|$.

Of course, this latter result is a rather blunt sufficient condition and can be further refined by a more detailed examination of location models.

5.5. *Erdös distance*

Consider a finite local interaction system, and assume connectedness (K4) is fulfilled. As in Example 1, given two players $i, j \in I$ we can define the Erdös distance[15] $d(i, j)$ as the minimum length of a path connecting i and j (in the sense of (K4)). We then define the *diameter* d^* as follows. Then,

$$d^* = \max_{i,j\in I} d(i, j)$$

is the maximum distance between two players in the network. It turns out that the diameter of certain social networks is surprisingly small. This observation is known as the *small worlds phenomenon* (see e.g. [33]).
By construction, every two consecutive players in a minimal path connecting i and j are neighbors, but players two or more positions apart in the path cannot be neighbors. This implies that the path crosses $\lfloor \frac{d^*}{2} \rfloor$ disjoint neighborhoods, hence $w^* \geq \lfloor \frac{d^*}{2} \rfloor$. Theorem 5 then yields a weaker sufficient condition, which might nonetheless be easier to apply.

Corollary 3. *Consider any connected local interaction-information system fulfilling Assumption 2. Let* d^* *be its diameter. The efficient convention is contagious under the imitate-the-best rule whenever* $\lfloor \frac{d^*}{2} \rfloor > Q_{\min} + 1$.

[14] In a lattice structure, (ii) requires the number of agents in the i-th dimension, N_i, to be an exact multiple of $2k + 1$. Dispensing with this requirement, Theorem 5 yields a slightly more complicated condition.
[15] The name comes from the Erdös number, which is the distance to Paul Erdös in the graph of mathematical collaborations.

This sufficient condition is, in general, easier to apply than Theorem 5, specially for irregular networks. The tradeoff is, of course, that the resulting bound is considerably weaker than that obtained from Theorem 5. For instance, in the $2k$-neighbors model, the diameter is $\lceil \frac{N}{2k+1} \rceil$ and $Q_{\min} = 2k$. Ignoring rounding, this yields selection of the efficient action when $N > 2(2k+1)^2$ (approximately), which is a worse bound than that in Theorem 3.

5.6. *The interplay of imitation and information*

The crucial ingredient for our results is the distinction between interaction and information. To be fair, the selection of risk-dominant equilibria in e.g. [14] or [30] hinges on the identification of learning with myopic best-reply, while we are assuming imitation rules. [16] However, imitation alone does not necessarily result in efficiency in a local interactions framework. Contagion of the efficient strategy is not merely a result of the imitate-the-best rule, but rather of the *combination* of imitation and information.

Specifically, Alós-Ferrer and Weidenholzer [1] show that, in the circular city model with $k = m = 1$, under the imitate-the-best rule the risk-dominant equilibrium is uniquely selected. [17] The logic of this result is fundamentally different to the analysis here. With $k = m = 1$, there is a large number of absorbing states, where clusters of A- and B-players alternate. All those states can be connected to each other through chains of single mutations. Risk dominance then implies that, while two mutations suffice for the initial step out of the efficient convention, at least three are required to leave to risk-dominant one. Thus, the risk-dominant convention is the only stochastically stable state. Due to the presence of the intermediate absorbing states (which are not absorbing when $m > k$), though, the risk-dominant action is *not* contagious in the sense of Definition 3.1.

Even though it is not sufficient, the assumption of imitation behavior is of course determinant for our results. To further emphasize this point, consider an arbitrary, connected local interaction system and assume agents play a best reply to observed play within their interaction neighborhood. Let

$$Q_{\max} \equiv \max\{|K(i)| \mid i \in I\}.$$

be the size of the largest interaction neighborhood in the network. A strategy is p-dominant if it is the unique best reply when a fraction p of a player's neighbors adopt it. This is a generalization of $\frac{1}{2}$-dominance (see [29]). Notice that p-dominance implies p'-dominance for $p' > p$.

If a strategy, say B, is $\frac{1}{Q_{max}}$ -dominant, this means that all neighbors of a B-player will adopt B as a best reply independently of other considerations. Hence, if the local interaction system is connected, a single mutation is enough for B to spread contagiously to the whole population, thus B is the unique contagious action.

[16] Kandori, Mailath, and Rob's [26] original model of global interactions, though, can be readily interpreted as an imitation model. See Sandholm [32] for a clarification.

[17] For larger neighborhoods, whether the efficient convention is selected or not depends both on k and the degree of risk-dominance of the other convention. In contrast, the present result holds independently of the exact payoffs of the game.

Note, though, that $\frac{1}{Q_{max}}$ -dominance is unrelated to Pareto-efficiency. Hence, this argument shows that for (nearly) all networks, there exists a 2×2 coordination game where imitation and best reply give different equilibrium predictions.

6. Additional Examples

In this section, we present a few examples of particular networks to illustrate the main result.

Example 4. The Torus. Assume now that $N_1 N_2$ players are situated at the vertices of an $N_1 \times N_2$ lattice on the surface of a torus. We can define the distance separating two players ij and xy as

$$d(ij, xy) = \min\{|i-x|, N_1 - |i-x|\} + \min\{|j-y|, N_2 - |j-y|\}.$$

We consider the Euclidean network where a player is only matched with players at a distance of at most k with $k \le \frac{N_1}{2}$ and $k \le \frac{N_2}{2}$, i. e. a player ij is matched with a player xy if and only if $0 \le d(ij, xy) \le k$. Assume for simplicity that N_1, N_2 are multiples of $2k+1$. Note that a player is not matched with himself. Furthermore, note that within this setup each player has $Q = 2k(1+k)$ neighbors (see Figure 1). Thus, the interaction neighborhood of a player ij is $K(ij) = \{xy : 0 < d(ij, xy) \le k\}$.

Since we consider a Euclidean network, the information neighborhood of a player is made of players up to a distance of $m > k$ from him, $M(ij) = \{xy : d(ij, xy) \le m\}$.

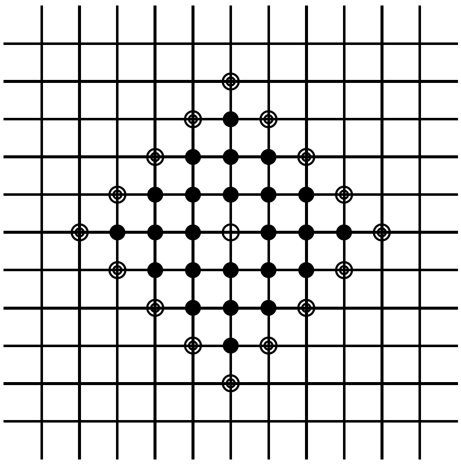

Fig. 1. Interaction neighborhood of size $k = 3$ and information neighborhood of size $m = 4$ on the lattice. A player has $4\sum_{j=1}^{k} j = 2k(1+k) = 24$ neighbors.

Notice that this is a Q-regular local interaction-information structure with $Q = 2k(1+k)$. Hence, we conclude that the efficient convention is contagious provided $N_1 N_2 > (2k(1+k)+1)^2$.

Example 5. Hierarchy. The following example is taken from [30]. The population is arranged in a hierarchy with levels from 0 (top) to L (bottom), where each player except those at level L has S subordinates (S is the *spread*). Figure 2(a) illustrates the case

$L = S = 3$. Formally, $I = \bigcup_{l=0}^{L} I_l$ with $I_l = \{1, \ldots, S\}^l$ for $l \geq 1$ and $I^0 = \{\emptyset\}$. $j \in K(i)$ if and only if $i = (j, n)$ or $j = (i, n)$ for some $n \in \{1, \ldots, S\}$. We further specify that $j \in M(i)$ whenever $j \in K(j')$ and $j' \in K(i)$ for some player j'.

Players at the last level have only one neighbor each, hence $Q_{\min} + 1 = 2$. However, for either $L \geq 3$ or $L \geq 2$ and $S \geq 3$, we have that $w^* \geq 3$. For $L = S = 3$ we already obtain $w^* = 10$ (see Figure 2(a)). Hence Theorem 5 yields selection of the efficient action.

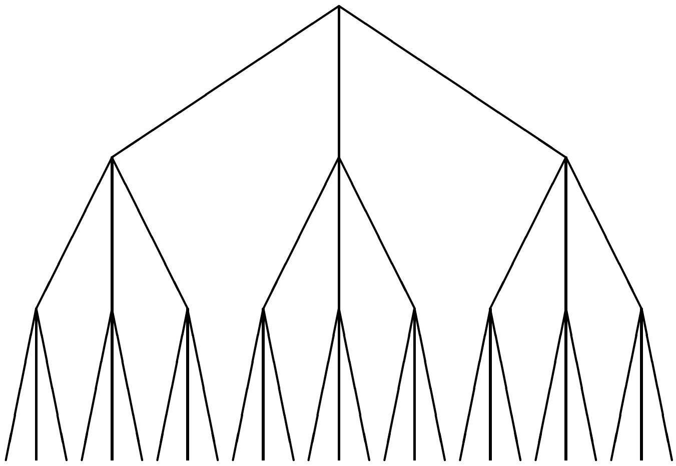

(a) A hierarchical network

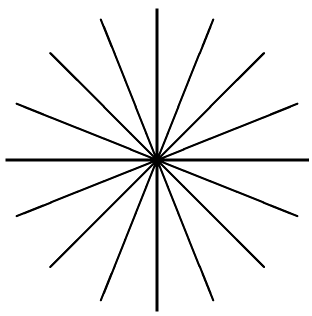

(b) A Star network

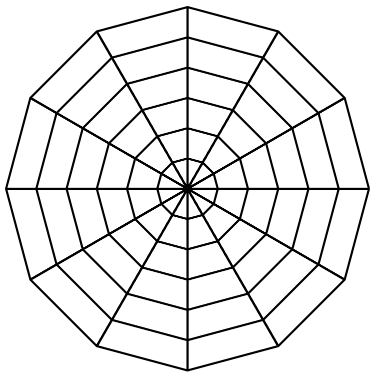

(c) A Spiderweb Network

Fig. 2. Some examples of networks.

Example 6. The Star Counterexample. Let $I = \{0, 1, \ldots, N\}$. Assume $K(0) = I \setminus \{0\}$ and $K(i) = \{0\}$ for each $i \neq 0$. This defines a Star Network (see Figure 2(b)). Let $M(i) = I$ for all i. Here we have that $w^* = 1$ and $Q_{\min} + 1 = 2$, hence Theorem 5 does not apply.

Indeed, if $\alpha > 0$ (which happens e.g. if B is risk-dominant) it is easy to see that a single mutation of player 0 would trigger a transition from $\overrightarrow{A}$ to $\overrightarrow{B}$, while the reverse transition requires two mutations. Thus, the only stochastically stable state is $\overrightarrow{B}$. This is an example where the network is very centralized. In such networks, there will be a low number of disjoint neighborhoods, and hence Theorem 5 fails to apply.

Example 7. Spider Net. Consider the network depicted in Figure 2(c), and assume the $M(i)$ are such that Assumption 2 holds. Taking a player at the external border of the network shows that $Q_{\min} + 1 = 4$. It is unnecessary to count the maximal number of disjoint neighborhoods, since a cursory examination suffices to show that $w^* > 4$. Hence, the efficient action is selected by Theorem 5.

7. Extensions

In this section we explore the robustness of our results with respect to three natural extensions. First, we examine asymmetric information structures. Second, we consider a situation where the information neighborhood is not fixed, but might contain random observations outside the interaction neighborhood. Third, we consider a mutation-free framework as in Lee and Valentinyi [28].

7.1. *Asymmetric Information*

In practice it is very often the case that some agents are more observable than others. There can be information structures which increase the visibility of certain "role groups". As Bala and Goyal [5] observe, this can be the case in situations where agents have access to a public source of information, as consumer magazines, research laboratories, or simply some focal agents (e.g. large farmers in an agricultural community). [18]

In order to take into account the possible presence of role groups in the population it is necessary to slightly generalize our framework to accommodate asymmetric information neighborhoods.

Formally, an *asymmetric information system* is defined as in Definition 2.2, except that the symmetry requirement (M2) is dropped. An asymmetric local interaction-information system is defined accordingly.

We define the set of *observers* of player i as those players j such that all their neighbors observe player i. That is,

$$K^*(i) = \{j \in I \mid i \in M(k) \quad \forall k \in K(j) \cup \{j\}\}$$

Note that $K^*(i)$ coincides with the set of contacts defined above whenever (M2) holds.

Assumption 2 remains formally unchanged, but it now refers to the set of observers. We can then prove the analogue of Lemma 3.

Lemma 4. *For an asymmetric information system fulfilling Assumption 2, any state where some player and all of his neighbors play A lies in the basin of attraction of* $\overrightarrow{A}$.

Proof. Assume that player i and all of his neighbors play A. Player i will receive a payoff of one which is the largest possible payoff. By definition, all players in $K(K^*(i))$ observe this and switch to A. By construction, now all players in $K^*(i)$ receive the

[18] [5] provide a nice example where information is transmitted locally but there is a small role group observed by everybody—a "royal family."

maximum payoff of 1. Hence, in the next step, all players in $K^*(K^*(i))$ will switch to A. By Assumption 2, we conclude that the efficient strategy spreads contagiously until we reach $\overrightarrow{A}$. □

This Lemma allows us to give an immediate analogue of Theorem 5.

Theorem 6. *Under Assumption 2, in any asymmetric local interaction-information system the efficient convention is contagious under the imitate-the-best rule, provided* $w^* > Q_{\min} + 1$ *in the finite case.*

Note that, for a transition to $\overrightarrow{A}$ to occur, it is not necessary that the observers relationship is connected in an undirected way. In fact, it is enough if some (role) group $I^* \subseteq I$ of agents is (iteratively) observable by the whole population. In this case one can modify Assumption 2 accordingly to provide an analogous result to Theorem 5.[19]

7.2. *Spatial Sampling*

In Young's [34] model of adaptive learning, players adopt a best reply to a sample of their opponent's play from the most recent periods. Durieu and Solal [13] give a nice reformulation of this model in terms of local interactions. The idea is that players adopt a best-reply to a sample of their neighbor's play.

This idea can be easily adapted to our context. Players are assumed to observe only a random fraction of the pattern of play in their *information* neighborhood. This can be easily motivated in the presence of information-gathering costs. Alternatively, one could interpret this assumption as another form of bounded rationality as follows. Players lack the the computing capacities to evaluate all information available and hence constrain their information-gathering efforts.

Note that, whereas each agent always has the same interaction neighborhood $K(i)$, the information neighborhood $M(i)$ will now only describe *potential* sources of information. The set of players from which information is received may vary across time.

Formally, assume that, in a given period t, each player i only observes a random subset $\mathcal{M}(i,t) \subseteq M(i)$ of the information available in his information neighborhood. Further, assume that

$$\mathcal{M}(i,t) \supseteq K(i) \cup \{i\} \tag{S1}$$

for all i and all t. That is, each player observes at least himself and always observes the pattern of play in his own interaction neighborhood. The concept of information neighborhood becomes meaningful if we assume that

$$\Pr\left(j \in \mathcal{M}(i,t)\right) > 0 \text{ for all } i, t \text{ and } j \in M(i). \tag{S2}$$

[19] One can define $Q^*_{\min}$ and w^{**} analogously to $Q_{\min}$ and w^* but referring to the role group I^*. If I is finite, we immediately obtain that A is contagious, provided $w^{**} > Q^*_{\min} + 1$. If both I and I^* are infinite, then w^{**} is infinite (by bounded neighborhoods) and hence A is contagious. A difficulty arises, if I is infinite but I^* is finite. Note that if an agent not in the role group and all of his neighbors are playing A, they will never switch to the inefficient action. However, the efficient action might not be able to spread back from this group. Hence, it might be possible to destabilize the efficient convention by strategically "seeding" the role group with deviators and reach, if not the inefficient convention, a mixed state where actions coexist, with the inefficient one dominating the role group. By bounded neighbors, though, the inefficient action can only spread to at most a finite subset of the population.

Hence, agents are informed about play in their interaction neighborhood and also have some idea about play in their information neighborhood.

The set $K^*(i) = \{j \in I \mid K(j) \subseteq M(i)\}$ can now be reinterpreted as the set of *possible contacts* of player i, that is, those players j such that all their interactions are with players that i "sometimes" meets. Assumption 2 has now the same interpretation as before.

As above, we are able to conclude that

Theorem 7. *Under Assumption 2, in any local interaction-information system with spatial sampling satisfying (S1) and (S2), the efficient convention is contagious under the imitate-the-best rule, provided $w^* > Q_{\min} + 1$ in the finite case.*

Proof. First, assume that player i and all of his neighbors play A. Player i will obtain the maximum payoff of one. By (S1), all players in $K(i)$ observe this and hence will retain their strategies. Furthermore, by (M2) and (S2), there is positive probability that each player $j \in M(i)$ observes this and switches to A. Thus, eventually all players in $M(i)$ will play A implying that all players in $K^*(i)$ obtain the maximum payoff. In this manner, due to Assumption 2, the efficient strategy can spread contagiously. Hence, $c^*(A) = B_{min} + 1$.

As in the proof of Theorem 5, we also conclude that $r^*(A) = w^*$ in the finite case, and $r^*(A) = \infty$ in the infinite case. Thus, if $w^* > Q_{\min} + 1$ the efficient strategy is uniquely selected. □

We remark that our result crucially depends on assumption (S1), i.e. that agents always observe their interaction neighborhood. If (S1) is violated, one can construct examples where risk dominant strategies are selected.[20]

7.3. *Contagion without Mutation*

As an alternative to the mistakes model, we follow Blume [11] and Lee and Valentinyi [28] and consider a model where, at the beginning of play, each player chooses each action with positive probability.[21]

Assume that in period $t = 0$ each player plays A with probability $p \in]0, 1[$ and B with probability $1 - p$. Let

$$Q_{\max} \equiv \max\{|K(i)| \mid i \in I\}.$$

be the size of the largest interaction neighborhood in the network. Then, we can prove the following result.

[20] For instance, consider the two neighbor model with $m = 2$ and a coordination game with $\alpha = \frac{2}{3}$ and $\beta = \frac{3}{4}$. Assume $|\mathcal{M}(i,t)| = 3$ and $i \in \mathcal{M}(i,t)$. That is, each player observes three out of the five players in his information neighborhood, himself included. Consider an isolated B-player. He earns a payoff of $\frac{2}{3}$ and the boundary A-players earn a payoff of $\frac{1}{2}$. With positive probability the B-player just sees the boundary A-players and retains his strategy, while each of the boundary A-players just see each other and the B-player and (since $\frac{1}{2} < \frac{2}{3}$) will switch to B. In this manner the risk dominant strategy can spread contagiously. Further, note that one player switching is not enough for the efficient strategy to spread contagiously.

[21] This approach avoids some well-known critiques to the literature on learning through mutations. See e.g. [9].

Theorem 8. *Under Assumption 2 and the Imitate the Best Rule, in any local interaction-information system the probability of converging to the efficient convention is bounded below by*

$$1 - \left(1 - p^{Q_{max}+1}\right)^{w^*}. \tag{1}$$

Proof. By Lemma 3, if some player achieves the maximum payoff we will reach the state $\overrightarrow{A}$. Consider a player i with interaction neighborhood $K(i)$. The probability that he earns the maximum payoff of 1 is equal to the probability that he and all his neighbors initially play A

$$\Pr\left(U(i,\omega_0) = 1\right) = p^{\,|K(i)|+1}.$$

Given a player subset V whose neighborhoods are pairwise disjoint, it follows that the probability that the system will *not* converge to $\overrightarrow{A}$ is smaller than

$$\prod_{i\in V}\left(1 - p^{\,|K(i)|+1}\right) \leq \prod_{i\in V}\left(1 - p^{\,Q_{\max}+1}\right).$$

Since w^* is the largest possible number of disjoint neighborhoods, it follows that the probability of not converging to $\overrightarrow{A}$ is bounded above by $\left(1 - p^{Q_{\max}+1}\right)^{w^*}$. Hence, the probability of converging to the efficient convention is bounded below by (1). □

Note that the bound given in 1 converges to one as $w^* \to \infty$. This corresponds to the infinite case where the probability of converging to the efficient convention is one.

Figure 3 plots the lower bound established in the last result as a function of w^* and $Q_{\max}$. Note that this bound is increasing with w^* and decreasing with $Q_{\max}$, which again reflects part of the intuition behind Theorem 5. A larger number of disjoint neighborhoods favors the eventual emergence of efficient conventions, whereas a relatively large size of the neighborhoods tends to hamper it.

8. Conclusion

This paper considers a local interaction model where players learn through imitation of successful behavior and makes an explicit distinction between interaction and information structures. We find that, for arbitrary networks, in the long run players will learn to coordinate on efficient actions, provided (i) information is able to flow through the network, and (ii) the (minimal) size of neighborhoods is small relative to the maximal number of disjoint neighborhoods.

The first condition is transparent. Information on the success of the efficient action should be able to spread out through the network (Assumption 2), i.e. it should flow beyond the narrow confines of local interactions.

The second condition can be decomposed in two parts. The first (small neighborhood size) points out that efficiency is fostered by the existence of "hubs" where the efficient action can be tried out, and whose size is relatively small, when compared to the population size. This allows the efficient action to achieve the maximum payoff and hence prove its value.

The second part (large number of disjoint neighborhoods) might seem somewhat paradoxical. It essentially requires that, in contrast to information, the nature of the interactions itself should not be "too global." The intuitive reason is that if interaction is

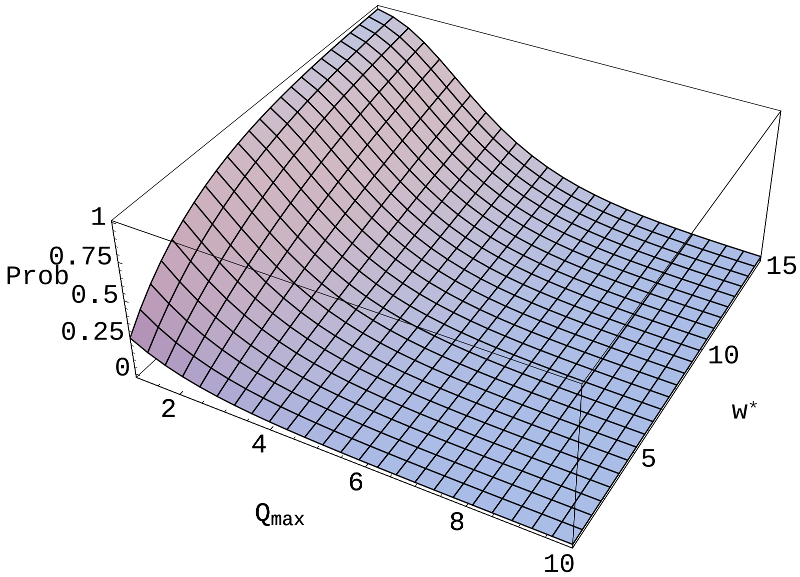

Fig. 3. The lower bound identified in Theorem 8 for $p = 0.5$.

not too global, any attempt to take over the population by an inefficient action would typically leave some clusters of players adopting the efficient one untouched. From such clusters, the efficient action can spread again to the whole population.

It is of course tempting to take this interpretation one step further and conclude e.g. that the onset of the internet, by expanding not only the information but also the interaction neighborhoods, might have favored the survival of inefficient computer-related technologies (software standards, operating systems, etc).

Appendix A. Proof of Theorem 1

We will rely on the characterization of the set of stochastically stable states developed by [26] and [34], and specially on the results of [15]. Detailed overviews can be found e.g. in [21] or [31].

Given two absorbing sets X and Y, let $c(X, Y) > 0$ (referred to as the *transition cost* from X to Y) denote the minimal number of mistakes necessary for a direct transition from X to Y, i.e. a positive probability path starting in an element of X and leading to an element in Y, which does not go through any other absorbing set.

Transitions need not be direct, though. Define a path from X to Y as a finite sequence of absorbing sets $P = \{X = S_0, ..., S_K = Y\}$. Let $S(X, Y)$ be the set of paths from X

to Y. Given a path P, define its length $l(P)$ as the number of elements of the sequence minus 1, so that $P = \{X = S_0, ..., S_{l(P)} = Y\}$. We extend the cost function to paths by $c(P) = \sum_{k=1}^{l(P)} c(S_{k-1}, S_k)$, and define $C(X,Y) = \min_{P \in S(X,Y)} c(P)$ to be the minimal number of mistakes required for a (possibly indirect) transition from X to Y.

We summarize now the results of [15]. The *Radius* of an absorbing set X is defined as

$$R(X) = \min \{C(X,Y) \,|Y \text{ is an absorbing set, } Y \neq X\}$$

i.e. the minimal number of mistakes needed to leave X towards another absorbing set. Intuitively, the radius measures how easy it is to destabilize an absorbing set. To obtain a measure for the accessibility of an absorbing set, we define the coradius of X as

$$CR(X) = \max \{C(Y,X) \,|Y \text{ is an absorbing set, } Y \neq X\}$$

[15] provides a powerful result which states that, if $R(X) > CR(X)$ for a given absorbing set X, then X is the unique stochastically stable set.

Lemma 5. *(Ellison 2000) Let X be an absorbing set. If $R(X) > CR(X)$, the only stochastically stable states are those in X.*

The intuition is clear, for the inequality $R(X) > CR(X)$ simply expresses the idea that X is easier to reach than to leave.

The proof of Theorem 1 is now straightforward.

Proof of Theorem 1. Let $X = \{\overrightarrow{A}\}$. By definition, $c^*(A)$ mutations to A are enough to trigger a transition to $\overrightarrow{A}$ from any other state in an absorbing set. Thus $CR(X) \leq c^*(A)$. Also by definition, in state $\overrightarrow{A}$ at least $r^*(A)$ mutations are required in order to leave the basin of attraction of $\overrightarrow{A}$. Thus $R(X) \geq r^*(A)$. Since A is contagious, $r^*(A) > c^*(A)$ and the conclusion follows from Lemma 5. □

References

[1] C. Alós-Ferrer, S. Weidenholzer, Imitation, Local Interactions, and Efficiency. Econ. Letters 93 (2006), 163–168.

[2] L. Anderlini and A. Ianni, Path Dependence and Learning from Neighbors, Games Econ. Behav. 13 (2006), 141–177.

[3] A. W. Anwar, On the Co-existence of Conventions, J. Econ. Theory 107 (2002), 145–155.

[4] J. Apesteguía, S. Huck, J. Oechssler, Imitation—theory and experimental evidence, J. Econ. Theory 136 (2007), 217–235.

[5] V. Bala, S. Goyal, Learning from Neighbors, Rev. Econ. Stud. 65 (1998), 595–621.

[6] A. Banerjee, D. Fudenberg, Word-of-mouth Learning, Games Econ. Behav. 46 (2004), 1–22.

[7] R. Baron, J. Durieu, H. Haller, P. Solal, Control Costs and Potential Functions for Spatial Games, International Journal of Game Theory 31 (2002), 541–561.

[8] G. Barron, I. Erev, Small Feedback-Based Decisions and their Limited Correspondence to Description-Based Decisions, Journal of Behavioral Decision Making 16 (2003), 215–233.

[9] J. Bergin, B. L. Lipman, Evolution with State-Dependent Mutations, Econometrica 64 (1996), 943–956.

[10] L. Blume, The Statistical Mechanics of Strategic Interaction, Games Econ. Behav. 5 (1993), 387–424.

[11] L. Blume, The Statistical Mechanics of Best-Response Strategy Revision, Games Econ. Behav. 11 (1995), 111–145.

[12] L. Blume, Population Games, in: B. Arthur, S. Durlauf, D. Lane (Eds.), The Economy as an Evolving Complex System II, pages 425–460, Addison-Wesley, Reading, MA, 1997.

[13] J. Durieu,P. Solal, Adaptive Play with Spatial Sampling, Games Econ. Behav., 43 (2003), 189–195.
[14] G. Ellison, Learning, Local Interaction, and Coordination, Econometrica, 61 (1993), 1047–1071.
[15] G. Ellison, Basins of Attraction, Long-Run Stochastic Stability, and the Speed of Step-by-Step Evolution, R. Econ. Stud. 67:17–45, 2000.
[16] G. Ellison and D. Fudenberg, Rules of Thumb for Social Learning, J. Polit. Econ., 101:612–643, 1993.
[17] G. Ellison and D. Fudenberg, Word-of Mouth Communication and Social Learning, Quart. J. Econ., 110:95–126, 1995.
[18] J. C. Ely, Local Conventions, Adv. Theoret. Econ., 2:1–30, 2002.
[19] I. Erev and G. Barron, On Adaptation, Maximization, and Reinforcement Learning among Cognitive Strategies, Psychological Reviews, 112:912–931, 2005.
[20] I. Eshel, L. Samuelson, A. Shaked, Altruists, Egoists, and Hooligans in a Local Interaction Model, Amer. Econ. Rev., 88:157–179, 1998.
[21] D. Fudenberg and D. Levine, The Theory of Learning in Games, The MIT Press, Cambridge, Massachusetts, 1998.
[22] A. Galeotti, S. Goyal, F. Vega-Redondo, M. O. Jackson, and L. Yariv, Network Games, mimeo, Caltech, 2006.
[23] J. Harsanyi and R. Selten, *A General Theory of Equilibrium Selection in Games*, The MIT Press, Cambridge, MA, 1988.
[24] M. O. Jackson, L. Yariv, Diffusion of Behavior and Equilibrium Properties in Network Games, Amer. Econ. Rev. (Papers and Proceedings) 42 (2007), 92–98.
[25] M. O. Jackson, A. Wolinsky, A Strategic Model of Social and Economic Networks, J. Econ. Theory 71 (1996), 44–74.
[26] M. Kandori, G. J. Mailath, R. Rob, Learning, Mutation, and Long Run Equilibria in Games, Econometrica 61 (1993), 29–56.
[27] S. Karlin, H. M. Taylor, A First Course in Stochastic Processes (2nd Ed.), Academic Press, San Diego, 1975.
[28] I. H. Lee, A. Valentinyi, Noisy Contagion without Mutation, Rev. Econ. Stud. 67 (2000), 47–56.
[29] S. Morris, R. Rob, H. Shin, p-dominance and Belief Potential, Econometrica 63 (1995), 145–157.
[30] S. Morris, Contagion, Rev. Econ. Stud. 67 (2000), 57–78.
[31] L. Samuelson, *Evolutionary Games and Equilibrium Selection*, The MIT Press, Cambridge, Massachusetts, 1997.
[32] W. H. Sandholm, Simple and Clever Decision Rules for a Model of Evolution, Econ. Letters, 61 (1998), 165–170.
[33] D. Watts, S. H. Strogatz, Collective Dynamics of ‘Small-World’ Networks, Nature, 393 (1998), 440–442.
[34] P. Young, The Evolution of Conventions, Econometrica 61 (1993) 57–84.
[35] P. Young, Individual Strategy and Social Structure, Princeton University Press, Princeton, New Jersey, 1998.