Naive learning in social networks: Imitating the most successful neighbor

Tsakas, Nikolas

Universidad Carlos III de Madrid

23 March 2012

Naïve Learning in Social Networks: Imitating the Most Successful Neighbor

Nikolas Tsakas*

Department of Economics, Universidad Carlos III de Madrid

March 23, 2012

Abstract

This paper considers a model of observational learning in social networks. Every period, the agents observe the actions of their neighbors and their realized outcomes, and they imitate the most successful. First, we study the case where the network has finite population and we show that, regardless of the structure, the population converges to a monomorphic steady state, i.e. where every agent chooses the same action. Subsequently, we extend our analysis to infinitely large networks and we differentiate the cases where agents have bounded neighborhoods, with those where they do not. Under bounded neighborhoods, an action is diffused to the whole population if it is the only one initially chosen by infinitely many agents. If there exist more than one such actions, we provide an additional sufficient condition in the payoff structure, which ensures convergence for any network. Without the assumption of bounded neighborhoods, we show that an action can survive even if it is initially chosen by a single agent and also that a network can be in steady state without this being monomorphic.

Keywords: Social Networks, Learning, Diffusion, Imitation.
JEL Classification: D03, D83, D85.

---

*I gratefully acknowledge my advisor Antonio Cabrales for his suggestions, comments and support. I would like to thank as well Friederike Mengel, Pedro Sant'anna, Daniel Garcia, Ignacio Ortuño and seminar participants at Universidad Carlos III de Madrid for their comments. Obviously, all remaining errors are my sole responsibility. Contact details: E-mail: `ntsakas@eco.uc3m.es` ; Universidad Carlos III de Madrid, Department of Economics, Calle Madrid 126, 28903, Getafe, Spain.

1. Introduction

A common characteristic of most economic activities is that they are not organized on a centralized and anonymous way. They rather involve bilateral interactions between agents. Moreover, they tend to be local in nature, meaning that, usually, the agents interact only with a small subset of the rest of the population. The interaction consists of informational exchange, or observation of strategies and outcomes. Indirectly, this interaction affects not only the direct neighbors, but also all the agents which are, somehow, connected via mutual neighbors, or neighbors of neighbors, etc. Agents use the received information to update their choices and improve their future outcomes.

In several dynamic decision problems, the agents seem to behave as if they observe the actions of their neighbors and tend to imitate those that yielded better results in the past, rather than as bayesian utility maximizers. Furthermore, in cases such as the use of a production technology, the choice between substituting products, or even the choice of a route through traffic, we observe that people tend to imitate neighbors that received extremely positive outcomes in the past. Imitation can be explained in several ways, mostly caused by lack of information or bounded rationality. Meaning that, either the agents are not aware of the mechanisms that control the outcome of their choices, or even if they are, they are unable to make the appropriate calculations in order to use the available information.

Learning through imitation of extremely successful behavior is not only observed in real-life examples. Experiments and field studies (see Apesteguia et al., 2007; Conley and Udry, 2010) have shown that, indeed, people tend to mimic the behavior of others that they have observed to be successful in the past, even if they are not aware whether this was a result of a correct choice or because of a lucky draw.

The main contribution of this paper is that we combine the above features, in an environment with purely informational interactions. Namely, we study a problem of dynamic decision making under uncertainty, where the agents interact through a social network. The interactions are only informational, meaning that the decision of an agent does not affect directly the outcome of the others. Agents update their decisions by imitating their neighbor that received the highest outcome in the previous period.

Formally, we consider a countable population forming an arbitrary network. Each agent chooses every period an action from a finite set and receives a payoff, which is drawn from a continuous distribution associated with the chosen action. The agents are not aware of the underlying distributions and the payoffs do not depend on the choices of other agents, i.e. there are no payoff externalities.

However, they observe the actions and realized payoffs of all their neighbors. The draws are independent for every agent, meaning that it is possible that two agents choose the same action and receive different payoffs. After observing actions and payoffs, the agents update their choice, imitating myopically the choice of their neighbor who received the highest payoff in the previous round. We extend the model on cases where only a fraction of the population revises its action every period, as well as on cases where some agents are experimenting by choosing randomly one of the actions they observed.

The model is in line with the literature of observational learning in social networks, capturing an environment with local and purely informational interactions and where the agents are using really few of the available piece of information. The natural question that follows is why this combination is important to deal with. There are many examples supporting the importance of the used context. For example, suppose a farmer who has to choose periodically which product to produce. The quantity she will manage to produce, given spatial and monetary constraints, depends mostly on the product itself, on her skills and some other random events, rather than on the choice of his neighbors. Moreover, a consumer who has to choose which telephone company would give her the highest satisfaction. Also this depends more on factors she may not be aware of, rather than on the choices of the others. This strengthens the idea of using purely informational interactions. Furthermore, the inability of analyzing the importance of each factor enforces the use of imitating mechanisms, instead of bayesian procedures. Finally, behavioral factors indicated by real-life situations, as well as experimental studies, have shown that the observation of surprisingly successful outcomes creates a tendency of imitation, without paying attention to the rest of the available information.

Under this context, we show that, in case the population is finite, the network converges with probability one to a steady state and this steady state has to be monomorphic, meaning that all the agents choose the same action. This action need not be the most efficient. This is because each action is vulnerable to sequential negative shocks, that can lead to its disappearance. The result holds also in the cases of experimentation over other observed actions, as well as if only some of the agents update their choice each period.

The results differ significantly when the population is infinitely large. First of all, without further restrictions we cannot ensure the convergence to a monomorphic steady state. Hence, we need to adjust the conditions under which we can ensure the diffusion of a single action. In case of infinite population, a crucial property is whether or not the agents have bounded neighborhood, meaning that they are observed only by a finite number of agents in the society.

Assuming bounded neighborhoods ensures the diffusion of a single action, if it is the only one that chosen initially by a non-trivial share of the population, i.e. all but a finite number. If this is not the case, then we provide a counter example, where the given network never converges. Nevertheless, we provide a sufficient condition in the payoff structure, that can ensure convergence regardless of the network structure. This condition is more demanding than first order stochastic dominance and for this reason we provide an example where convergence occurs under our sufficient condition, but not under first order stochastic dominance. These observations lead to the conclusion that in a very large network, the diffusion of a single action is very hard and demands either a very large proportion of initial adopters, a special network structure, or an action to be much more efficient compared to the rest. As far as it concerns the use of our sufficient condition, the fact that it disregards completely the importance of the network architecture makes it useful mostly for networks with small upper bound in neighborhoods. The behavior of specific network structure would be a very interesting topic for further research.

Subsequently, if we drop the assumption of bounded neighborhoods, the properties of the network change significantly. Under these circumstances, we provide an example where an action survives although it is chosen by only one agent. This happens because this agent is able to affect the choice of infinitely many others, stressing the role of centrality in social networks. Providing a technology or a product to a massively observed agent can affect seriously the behavior of the population. We provide, also, another example that drops the result of steady state being monomorphic, that was proven to hold for finite population. In this example, we show a network in steady state, where more than one actions are chosen by infinitely many agents.

Concluding, the present paper provides some important results on the study of "imitate-the-best" learning mechanisms, as well as it introduces many important questions, acting as a trigger for future research. The fact that learning is a natural procedure in societies and that the imitation of successful behavior is commonly observable in many aspects of social life and in several economic activities, makes the study of the topic important, promising and, at the same time, interesting and fascinating.

Related Literature

This paper contributes to the literature of social learning in networks. We use a model of observational learning, where the agents observe the behavior and payoffs of their neighbors, and each period they imitate the most successful. Under this framework we study diffusion of a specific behavior in the whole netowrk. In this section, we connect our work with the existing literature in social learning and diffusion in networks, while stressing the different perspectives that arise from our analysis.

Bala and Goyal (1998) study social learning with local interactions, under myopic best-reply. Learning occurs through Bayesian update, using only information about their own neighbors, and not about neighbors of neighbors, etc. They provide sufficient conditions for the convergence of beliefs: Neighborhoods need to be bounded[1] in order to ensure convergence to the efficient action. Existence of a set of agents connected with everybody, often called the "royal family", can be harmful for the society. This is because such a set may enforce the diffusion of their action/opinion, even if it not the optimal for the society.

Bayesian learning and best-response strategies[2] are common in the literature of learning in networks (see Gale and Kariv, 2003; Acemoglu et al., 2011). More specifically, Gale and Kariv (2003) study a Bayesian learning procedure, but using a network and payoff structure quite similar to ours. They show that, given the fact that agent's information is non-decreasing in time, the equilibrium payoff must be also non-decreasing and since they assume it to be bounded, it must converge. Moreover, in equilibrium identical agents gain the same in expectation. In general, in Bayesian procedures the agents accumulate information through time, which is in favor of learning. However, in our setting it is not the case. Hence, it is normal that convergence will be harder to occur, but eventually not impossible.

On the other hand, Golub and Jackson (2010) study an extended version of the standard DeGroot model[3] (see DeGroot, 1974). In their context a crowd is "wise" if the importance of the most influential neighbor vanishes as the society grows. Their idea is along the same lines with our of bounded neighborhood. This is because the existence of agents that are connected with infinitely many agents can prevent from diffusion of a unique action, or even from convergence to a steady state. A crucial difference between the two papers, is that their updating rule allows to the agents to use more information, nevertheless, this is also a type of naïve learning mechanism[4].

As we already mentioned, our paper deals with a procedure of observational learning. Observational learning has been the subject of several studies. The main characteristic of these learning procedures is that agents tend to decide based on the choices made by other agents in the past,

[1] by bounded neighborhood we mean that there exists $K > 0$ such that the number of neighbors of every agent i satisfies $k_i \equiv |N_i| \leq K$, $\forall i \in N$

[2] The agents maximize short-term utility based on the observed behavior of their neighbors in the previous round

[3] In DeGroot's model a set of k individuals may reach consensus starting from a subjective probability distribution for the unknown value of a parameter. Consensus is reached because every agent gets informed about the probability distributions of the other agents and revise her distribution. The revised belief is a linear combination of the distributions of all the agents

[4] Meaning that the agents are using simple rules to update their behavior. Rules that do not demand sophisticated skills.

rather than on calculation of payoffs of any given action. Banerjee (1992) introduces a simple decision model, where agents observe the behavior of the previous decision makers and follow a herd behavior, i.e., they tend to do what others have done, rather than evaluating their own information. Subsequently, Ellison and Fudenberg (1995), and Banerjee and Fudenberg (2004) switch attention to word-of-mouth learning, meaning that agents accumulate information from a random sample of the population, and act based on the average payoffs of each action. Moreover, Smith and Sorensen (2000) observe and study some possible problems in the long-run behavior of observational learning procedures.

Additionally, in Ellison and Fudenberg (1993), agents choose between two technologies, and periodically evaluate their choices based on the average performance of each technology at the previous round. They introduce the concept of an exogenous "window width", which plays a role similar to the network structure in our case.

Even though these papers focus on similar learning procedures, they significantly differ from our work in the role of imitation of the best agent. More specifically, averaging the performance of several agents, reduces the effect of unexpectedly lucky results, in future behaviors. Imitation of very successful results can lead to survival of suboptimal actions through time, in case these results are caused by luckily realized states of nature.

Learning by imitation is also widely-studied in the literature. Various models of evolutionary game theory (Weibull, 1995; Fudenberg and Levine, 1998) are using different types of imitation processes. Josephson and Matros (2004) study an evolutionary process with random shocks, where the agents imitate the most successful behavior. They show that only strictly payoff-dominant strategies survive in the long-run. In this paper, we extend their results — in a slightly different environment — for stochastically dominant strategies. Vega-Redondo (1997) studies the evolution of a Cournot economy, where the agents adjust their behavior by imitating the most successful agent in the previous round. He concludes that, as time proceeds, this adjustment rule leads to spread of the most successful behavior in the market. At first glance, this concept appears to be similar to ours. However, there are still two main differences. First, in our concept, unlike Vega-Redondo (1997), we restrict our attention to payoff functions without externalities. Second, this paper does not include local interaction structure, since it is assumed that every agent's strategy and payoff are publicly observed.

Payoff uncertainty is introduced by Schlag (1998, 1999). In the first paper Schlag (1998), an agent with limited memory is randomly matched with another agent from the population. Then, he studies different adjustment rules, with the property that successful behaviors are imitated, and concludes that the most efficient mechanism is the one where an agent copies probabilistically the other's behavior if it yielded higher payoff with probability proportional to the difference between the realized outcomes. In Schlag (1999), he extends his framework to infinite populations, where each agent is randomly matched with two agents each time. His work differs from ours, because it focuses more on the selection of the most appropriate behavioral mechanism, rather than the properties of a specific mechanism. However, the evidence about efficiency of "imitating-the-best" behavior further strengthens our selected mechanism.

The literature on imitation in networks focuses mainly on coordination games and prisoner's dilemma games (e.g., Eshel et al., 1998; Alós-Ferrer and Weidenholzer, 2008; Fosco and Mengel, 2010). In these models, uncertainty comes from the lack of information about the choices made by each agent's neighbor. However, for each action profile, the payoffs are deterministic. This is not the case in our environment, where payoffs are stochastic. Hence, the interaction between agents is only informational. Alós-Ferrer and Weidenholzer (2008) aim at identifying the conditions for contagion of efficient actions in the long-run, problem similar to the one we study.

Our analysis focuses on the diffusion of behavior, as well as the long-run properties of a network. Large part of the literature in diffusion restricts attention to best-response dynamics, and relates the utility of adopting a certain behavior with the number, or proportion, of neighbors that have already adopted it. Morris (2000) shows that maximal contagion occurs for sufficiently uniform local interaction and low neighbor growth. Along the same line, Lopez-Pintado (2008) and Jackson and Yariv (2007) look into the role of connectivity in diffusion, finding that stochastically dominant degree distributions favor diffusion. The model and the behavior rules are quite different compared to our case. However, once again, the role of connectivity in the network is stressed.

The brief review of the literature makes apparent the importance of learning procedures in the study of social networks, as well as the lack of work dealing with imitation of extremely successful behaviors in a society linked as a network. These facts stress the importance of the results of the present work, as a benchmark for further study of the topic.

The rest of the paper is organized as follows. In Section 2, we explain the model. Section 3, contains some initial remarks on convergence to steady state. Section 4, provides the main results for networks with finite population. While, Section 5 deals with networks with countably infinite population. Finally, Section 6 provides conclusions and possible extensions of the present paper.

2. The Model

2.1. The agents

[5]There is a countable set of agents $N = \{1, ..., n\}$[6], with typical elements i and j. The set N is mentioned as *population* of the network. Each $i \in N$ takes an action a_i^t from a finite set $A = \{\alpha_1, ..., \alpha_M\}$, at every period $t = 1, 2, ...$ Uncertainty is represented by a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where Ω is a compact metric space, $\mathcal{F}$ is the Borel σ-field, and $\mathbb{P}$ is a Borel probability measure. The states of nature are drawn independently every period and for each agent. Hence, it is possible that two agents choose the same strategy at the same period and they receive different payoffs.

There is a common stage payoff function $U : A \times \Omega \to \mathbb{R}$, with $U(a_i^t, \omega_i^t)$ being a bounded and compact function for every action a and continuous in ω_i. Uninteresting cases occur when some action is strictly dominant independently of the state of nature [7]. Hence, we restrict our attention to the case where $\forall a, a' \in A$, $U(a, \Omega)$ is a closed interval in $\mathbb{R}$, $U(a, \Omega) = U(a', \Omega)$ and they have full support.

Denote $A_t = \{\alpha_k \in A : \exists i \in N \text{ such that } a_i^t = \alpha_k\}$ as the set of actions that are chosen in the network at time t. Recall that the action in period t depends on the states of nature in period $t - 1$. Namely, $\Omega_{t-1} = \{\omega_i^{t-1}, \forall i \in N\}$ is the set of realized states of nature in period $t - 1$.

2.2. The Network

A social network is represented by a family of sets $\mathcal{N} := \{N_i \subseteq N \mid i = 1, ..., n\}$, with N_i denoting the set of agents observed by agent i. Throughout the paper N_i is called i's neighborhood, and is assumed to contain i.

The sets $\mathcal{N}$ induce a graph with nodes N, and edges $E = \bigcup_{i=1}^n \{(i, j) : j \in N_i\}$. We focus on undirected graphs: as usual, we say that a network is undirected whenever for all $i, j \in N$, $j \in N_i$ if and only if $i \in N_j$. The network structure describes the flow of information in the network. Namely, each agent $i \in N$ observes the action and the realized payoff of every $j \in N_i$.

A path in a network between nodes i and j is a sequence $i_1, ..., i_K$ such that $i_1 = i$, $i_K = j$ and $i_{k+1} \in N_{i_k}$ for $k = 1, ..., K - 1$. The distance, l_{ij}, between two nodes in the network is the length of the shortest path between them. The diameter of the network, denoted as $d_{\mathcal{N}}$, is the largest distance between any two nodes in the network

[5]Our basic notation follows the one in Gale and Kariv (2003).

[6]In Section 4 we assume N to be finite, whereas in Section 5 is assumed to be countably infinite

[7]If an action is strictly payoff dominant, it will be spread to the network

We say that two nodes are connected if there is a path between them. The network is connected if every pair of nodes is connected. In our analysis, we focus on connected networks. However, the analysis would be identical for connected components[8] of disconnected networks.

2.3. Behavior

At $t = 1$, each agent receives a private signal $\sigma_i(\omega)$, where $\sigma_i : \Omega \to \mathbb{R}$ and based on this she makes her initial choice. At the end of each period, the agents observe the actions and realized payoffs of all their neighbors. Subsequently, at $t = 2, 3, ...$ each agent switches to the action of her neighbor that earned the highest payoff in the previous period. Notice that agents imitate others based on realized, rather than expected payoffs. Tie breaks are broken randomly. Formally, for each $t > 1$,

$$a_i^{t+1} \in \left\{ \, a \in A \ : \ a = a_k^t \text{ where } k = \text{argmax}_{j \in N_i} U(a_j^t, \omega_j^t) \, \right\} \tag{1}$$

where ω_j^t is the actual state at t, for agent j.

The above described behavior is usually called "imitate-the-best" (see also Alós-Ferrer and Weidenholzer, 2008). The important aspect of this myopic behavior is that the agents discard most of the available information. They ignore whatever has happened before the previous round and even from this round they take into account only the piece of information related to the most successful agent. This naïve strategy makes the network vulnerable to extreme shocks, that may be very misleading for the society.

However, there has been empirical and experimental evidence showing that, indeed, agents tend to switch strategies, imitating those with the most attractive realized outcomes (for example see Apesteguia et al., 2007; Conley and Udry, 2010). Moreover, under certain conditions that are explained later, even such a decision rule can lead to efficient outcomes.

Notice that in the present setting there are no payoff externalities, since the outcome of each agent depends only on his action and the current state of nature. This feature differentiates the present work from the extended literature on imitation procedures in several types of games. Nevertheless, there are information externalities, related to the fact that agents adjusted their action based on the information received by their neighbors.

2.4. Steady state and efficiency

State of period t is called the set that contains the action chosen by each agent at this period . A state

[8] A component is a non-empty sub-network N' such that $N' \subset N$, N' is connected and if $i \in N'$ and $(i, j) \in E$ then $j \in N'$ and $(i, j) \in E'$.

is called *monomorphic* if every agent chooses the same action, at period t. Moreover, the network is in *steady state* at period t, if no agent changes her action from this period on. Throughout the paper, it is mentioned the idea of convergence, referring to *convergence with probability 1*. Formally:

Definition 1. *State of period t is called the set $S^t = \{a_i^t, \ldots, a_n^t\}$.*

Definition 2. *A state is monomorphic if there exists $k \in \{1, ..., M\}$, such that $a_i^t = \alpha_k$ for all $i \in N$.*

Definition 3. *Let $t \geq 1$, if for all $i \in N$ holds that $a_i^t = a_i^{t'}$ for all $t' > t$, the network is in steady state.*

We call an action *efficient* if it yields the highest expected payoff. An action is mentioned as *more efficient* than another one if it yields higher expected payoff. Later on, we will focus on actions that can be ordered in terms of stochastic dominance.

3. First Results about Convergence to Steady State

Remark 1. *If the agents were imitating based on the highest expected payoff of each action, then the system would converge to a monomorphic state. At this state every agent would choose the efficient action.*

This first remark is trivial, because being aware of the expected payoff diminishes the uncertainty in the system. This case is the same as having deterministic outcomes for each action, where eventually every agent imitates the efficient action, as soon as she gets aware of it. A more interesting case is extensively studied by Ellison and Fudenberg (1993), where the agents have to choose between two technologies and they adjust their choice based on the average realized payoff of each technology in the previous round.

Remark 2. *Under the behavior rule (1), if a finite population network is* complete*, it will converge with probability 1, to a monomorphic steady state, from the second period on.*

If the network is complete, each agent is able to observe actions and realized payoffs of every agent in the society. Hence, the action chosen by the agent that received the highest payoff in the first period, will be chosen by everyone from the second period on. The probability that two actions will give exactly the same payoff is zero, because we have assumed continuous probability distributions.

In case of network with infinite population, the result is not obvious. If there exist two, or more, actions chosen by infinite players, then for each of them will exist some agent receiving the maximum payoff. However, with probability 1, one of the two action will receive a payoff strictly better than the other. This happens because the agents are countably infinite, whereas the possible outcomes from a continuous distribution are uncountably infinite. Hence, the network converges to a steady state from the second period on.

Moreover, even when the network is not complete, if there exists a set of agents that is observed by everyone, then the initial action chosen by these agents is likely to affect the behavior of the whole system. For finite population, if one of them receives a payoff close to the upper bound, this will lead to the diffusion of the chosen action from the next round on, irrespectively of the being efficient or not (see also "Royal Family" in Bala and Goyal, 1998). Once more, the analysis becomes different in case we assume the population to be infinitely large.

The simple cases analyzed above conclude to the diffusion of a single action to the whole population, in the long run. However, we have not yet shown that this has to be always the case. The following section shows that this is always the case in networks with finite population, whereas section 5 provides appropriate conditions that can ensure convergence to a monomorphic steady state, for infinitely large population.

4. Networks with finite population

In this section, we restrict our attention to networks with finite population. We prove that finite networks always converge to a steady state, regardless of the initial conditions or the network structure. Moreover, every action that is initially chosen in the network can be the one to survive in the long-run (The first proposition shows that in finite networks the steady state has to be monomorphic). Recall, that we have restricted our analysis to cases where there is no action that yields always higher payoffs than all the others.

Proposition 4.1. *Let a connected network, with finite population. If the agents behave under behavioral rule (1), then all possible steady states are necessarily monomorphic.*

Proof. Suppose not. This means that at the steady state at least two different actions are chosen. Given that the network is connected, for any arbitrary structure, there will be at least a pair of neighbors, i, j choosing different actions. Moreover, given that the population is finite, each one of these agents must have a bounded neighborhood. Let $b_{i,1}$ and $b_{i,2}$ the number of neighbors of i choosing actions α_1 and α_2 respectively. If we are in a steady state, this means that agent i will never change her choice. However, because of common support and continuity of payoff functions, as well as boundedness of neighborhoods, it holds always that:[9]

$$Pr[U(j,\omega_j^t) > U(l,\omega_l^t) : \text{ for some } j \in b_{i,2} \text{ and for all } l \in b_{i,1}] = \tag{2}$$
$$= (1 - \prod_{j\in b_{i,2}} Pr[U(j,\omega_j^t) \leq \hat{u}]) \prod_{l\in b_{i,1}} Pr[U(l,\omega_l^t) \leq \hat{u}] = \tag{3}$$
$$= [1 - F_2(\hat{u})^{b_{i,2}}]F_1(\hat{u})^{b_{i,1}} \geq p > 0 \tag{4}$$

Hence, every period there is positive probability p that agent i will change her strategy. The probability that she never changes, given that no other agent in the network does so, is zero, because:

$$\lim_{t\to\infty}\prod_{t=1}^{\infty}(1-p) = \lim_{t\to\infty}(1-p)^t = 0$$

And this leads to contradiction of the initial argument[10]. Concluding, we have shown that it is impossible for a connected network with finite population to be in a steady state, when more than one actions are still present. □

Notice that, the proposition is not valid for networks with infinite agents (see Example 5.4 - two stars). Also, in case the network is not connected, the proposition holds for every connected component. The above proposition is in line with the work of Gale and Kariv (2003), where, although the updating mechanism differs, identical agents end up making identical choices in the long run. Moreover, becomes apparent the advantage of using continuous, rather than discrete, payoff functions, since this discards an unnecessary large amount of ties.

Nevertheless, we have still not ensured the convergence of the system to a steady state. Subsequently, we will focus on identifying necessary and sufficient conditions for convergence. Notice that all monomorphic states are possible steady states (think of the trivial example where initially every agent chooses the same action).

The convergence to steady state is quite intuitive. This is because each strategy has a positive probability of disappearance after a finite number of periods and given that inexistent strategies never reappear, the set of observed actions shrinks over time until it ends up containing a single element (i.e. a unique action chosen by every agent in the network). To prove the result, we need the following three lemmas:

Lemma 4.1. *Let an arbitrary finite network ($n < \infty$), such that more than one action is observed. Each period t, every action $\alpha_k \in A_t$ has positive probability of disappearing after no more than $d_{\mathcal{N}}$ periods.*

[9]For F_1 and F_2 being the cumulative distribution functions of the payoffs of actions α_1 and α_2 respectively
[10]The analysis is identical for more than two actions.

Proof. At period t, there are at least two actions observed in the network. Take all players $i \in N$ such that $a_i^t = \alpha_k$. Let $I_k^t = \{i \in N : a_i^t = \alpha_k\}$ be the set of agents that chose action α_k in period t and $\neg I_k^t = \{i \in N : a_i^t \neq \alpha_k\}$ the set of agents that did not. Take $F_k^t = \{i \in I_k^t : N_i \cap \neg I_k^t \neq \emptyset\}$, i.e. the set of agents in I_k^t that have, at least, one neighbor choosing an action different than α_k. Moreover, take $NF_k^t = \{i \in I_k^t : N_i \cap F_k^t \neq \emptyset \text{ and } i \notin F_k^t\}$, including those agents that have neighbors belonging to F_k^t, but they do not belong themselves.

Think of the following events: $\hat{b}^t = \{\omega_i^t \in \Omega, \text{ for } i \in N_{i,\alpha_k}^t = \{F_k^t \cup NF_k^t\} : U_i^t(\alpha_k, \omega_i^t) \leq \hat{u}^t\}$, meaning that all agents in N_{i,α_k} get payoff lower than a certain threshold and $\hat{B}^t = \{\omega_j^t \in \Omega, \text{ for } j \in N_i \setminus N_{i,\alpha_k}^t, \text{ where } i \in N_{i,\alpha_k}^t \text{ and } k' \neq k : U_j^t(\alpha_{k'}, \omega_j^t) > \hat{u}^t\}$, meaning that all the neighbors of the above mentioned agents that choose different strategy get payoff higher than this threshold.

Given that the states of nature are independent between agents, we can define a lower bound of the event, in which all agents that at time t play α_k, change their strategy in the following period, without making any of their neighbors change to α_k, we denote it as C^t. In fact, $C^t \supset \{\hat{b}^t \cap \hat{B}^t\}$, because this intersection is a specific case of the set C^t. Hence, $Pr(C^t) \geq Pr(\hat{b}^t \cap \hat{B}^t) = Pr(\hat{b}^t)Pr(\hat{B}^t)$. Notice that we impose independence between $\hat{b}^t$ and $\hat{B}^t$, which holds because of the independence between the states of nature.

Remark: $\forall \alpha_k \in A$ and $\omega \in \Omega \Rightarrow Pr(U(\alpha_k, \omega) \leq \hat{u}) \in [b_t, B_t]$ where $b_t, B_t \in (0,1)$. Hence, $Pr(U(\alpha_k, \omega) \leq \hat{u}^t) \geq b_t > 0$ and $Pr(U(\alpha_k, \omega) > \hat{u}^t) \geq 1 - B_t > 0$. This is because of the full support of U and the continuity of Ω.

Using the above remark and the independence of the states of nature we get the result that $Pr(\hat{b}^t) \geq b_t^{|N_{i,\alpha_k}^t|}$ and analogously $Pr(\hat{B}^t) \geq (1 - B_t)^{|N_i \setminus N_{i,\alpha_k}^t|}$[11].Hence,

$$Pr(C^t) \geq Pr(\hat{b}^t \cap \hat{B}^t) = Pr(\hat{b}^t)Pr(\hat{B}^t) \geq b_t^{|N_{i,\alpha_k}^t|}(1 - B_t)^{|N_i \setminus N_{i,\alpha_k}^t|} > 0$$

Independence of the realization of the states of nature yields the result that C^t and C^{t+1} are conditionally independent. The realization or not of C^t does not give any extra information about the realization of C^{t+1}.

One of the possible histories that will lead to disappearance of action α_k is the consecutive realization of the events $C^{t+\tau}$ for $\tau \geq 0$, until all the agents that were using α_k at t have changed their choice. The number of periods needed depends on the structure of the network and more specifically is at most equal to the diameter of the network. For any network with n agents the diameter cannot be greater than $n - 1$.[12]

[11] Recall that $| N_{i,\alpha_k}^t |$ denotes the amount of agents belonging to this set

[12] and this holds only for the "linear" network, which is similar to the 2-neighbor network, except of the fact that two of the agents have only one neighbor.

Hence, we can construct a positive lower bound for the probability of disappearance of an action α_k after no more than $d_{\mathcal{N}}$ periods. Denote this event as $D^t_{\alpha_k}$ and we get:

$$Pr(D^t_{\alpha_k}) \geq Pr(\bigcap_{\tau=0}^{d_{\mathcal{N}}-1} C^{t+\tau}) = [\prod_{\tau=1}^{d_{\mathcal{N}}-1} Pr(C^{t+\tau} \mid C^{t+\tau-1})]Pr(C^t) \geq$$
$$\geq \prod_{\tau=0}^{d_{\mathcal{N}}-1} b_{t+\tau}^{|N^{t+\tau}_{i,\alpha_k}|}(1-B_{t+\tau})^{|N_i \setminus N^{t+\tau}_{i,\alpha_k}|} = \dot{B} > 0. \blacksquare$$

□

Lemma 4.2. *Given that, at time t, $K-1$ actions have disappeared from the network, there is strictly positive probability that after finite number of periods, $\tau < \infty$, exactly K actions will have disappeared.*

Proof. Denote as K_t the following set of histories:

$$K_t = \{\text{at time } t, \text{ exactly } K \text{ actions have disappeared from the network}\}$$

It is enough to show that $Pr(K_{t+\tau} \mid (K-1)_t) > 0$ where $\tau < \infty$ (in fact, at most $\tau = d_{\mathcal{N}}$).

$$Pr[K_{t+\tau} \mid (K-1)_t] = \sum_{\alpha_k \in A_t} \{Pr[D^{t+d_{\mathcal{N}}-1}_{\alpha_k} \mid (K-1)_t]\} - Pr[\bigcup_{m=1}^{M-K-1} (K+m)_{t+\tau} \mid (K-1)_t] \quad (5)$$

where M is the total number of possible actions. Namely, the above expression tells that the probability of exactly one more action disappearing in the next τ periods, equals the sum of probabilities of disappearance of each action, minus the probability that more than one actions will disappear in the given time period. Analogously:

$$Pr[(K+1)_{t+\tau} \mid (K-1)_t] = \sum_{\alpha_k,\alpha_{k'} \in A_t}^{k \neq k'} \{P[D^{t+d_{\mathcal{N}}-1}_{\alpha_k} \bigcap D^{t+d_{\mathcal{N}}-1}_{\alpha_{k'}} \mid (K-1)_t]\} - Pr[\bigcup_{m=2}^{M-K-1} (K+m)_{t+\tau} \mid (K-1)_t]$$

$$Pr[(K+1)_{t+\tau} \mid (K-1)_t] + Pr[\bigcup_{m=2}^{M-K-1} (K+m)_{t+\tau} \mid (K-1)_t] = \sum_{\alpha_k,\alpha_{k'} \in A_t}^{k \neq k'} Pr[D^{t+d_{\mathcal{N}}-1}_{\alpha_k} \bigcap D^{t+d_{\mathcal{N}}-1}_{\alpha_{k'}} \mid (K-1)_t]$$

But:

$$Pr[(K+1)_{t+\tau} \mid (K-1)_t] + Pr[\bigcup_{m=2}^{M-K-1} (K+m)_{t+\tau} \mid (K-1)_t] = Pr[\bigcup_{m=2}^{M-K-1} (K+m)_{t+\tau} \mid (K-1)_t]\}$$

$$\Rightarrow Pr[\bigcup_{m=2}^{M-K-1} (K+m)_{t+\tau} \mid (K-1)_t]\} = \sum_{\alpha_k,\alpha_{k'} \in A_t}^{k \neq k'} Pr[D^{t+d_{\mathcal{N}}-1}_{\alpha_k} \bigcap D^{t+d_{\mathcal{N}}-1}_{\alpha_{k'}} \mid (K-1)_t]$$

Hence, by equation (2):

$$Pr[K_{t+\tau} \mid (K-1)_t] = \sum_{\alpha_k \in A_t} \{Pr[D_{\alpha_k}^{t+d_{\mathcal{N}}-1} \mid (K-1)_t]\} - \sum_{\alpha_k, \alpha_{k'} \in A_t}^{k \neq k'} \{Pr[D_{\alpha_k}^{t+d_{\mathcal{N}}-1} \bigcap D_{\alpha_{k'}}^{t+d_{\mathcal{N}}-1} \mid (K-1)_t]\}$$

By lemma 1, we have shown that the first summation is strictly larger than zero, and we just need to show that it is also strictly larger than the second. Notice that the first summation is trivially weakly larger[13]. If the equality was possible, this would mean that, for every action, its disappearance would necessarily yield the disappearance of another action. However, this cannot be the case for every action, because of the independence of states of nature (The proof is trivial and available upon request).

Hence we have shown that: $Pr[K_{t+\tau} \mid (K-1)_t] \geq \ddot{B} > 0$. □

Lemma 4.3. *At each time t, there is positive probability of convergence to a monomorphic state in the next $T = (M - K + 1)d_{\mathcal{N}}$ periods, when $K - 1$ actions have already disappeared.*

Proof. The event of "convergence to a monomorphic state" is the same as telling that $M-1$ actions will have disappeared after T periods. One possible way of this to happen is if one action disappears every $\tau = d_{\mathcal{N}}$ periods. Namely:

$$\{(M-1)_{t+T}\} \supset \{\{(K)_{t+\tau} \mid (K-1)_t\} \cap \{(K+1)_{t+2\tau} \mid (K)_{t+\tau}\} \cap \dots \cap \{(M-1)_{t+T} \mid (M-2)_{t+T-\tau}\}\}$$

But, notice that the events on the right hand side of the expression are independent, because the states of nature are independent across time. Hence, recalling the result of lemma (2), we get the following expression for the related probabilities.

$$\begin{aligned} P_t[(M-1)_{t+T}] &\geq P_t[\{(K)_{t+\tau} \mid (K-1)_t\} \cap \{(K+1)_{t+2\tau} \mid (K)_{t+\tau}\} \cap \dots \cap \{(M-1)_{t+T} \mid (M-2)_{t+T-\tau}\}] \\ &= P_t[(K)_{t+\tau} \mid (K-1)_t] P_{t+\tau}[(K+1)_{t+2\tau} \mid (K)_{t+\tau}] \dots P_{t+T-\tau}[(M-1)_{t+T} \mid (M-2)_{t+T-\tau}] \\ &\geq \ddot{B}^{(M-K-1)} = \underline{C} > 0 \end{aligned}$$

Meaning that the probability of convergence in finite time is strictly positive. Moreover, notice that, $\forall K < M - 1 \Rightarrow \underline{C} < 1$. This remark is trivial because if $K = M - 1$, the system has already converged, nevertheless we will use it to prove the following theorem. □

Theorem 4.1. *Let an arbitrary finite network ($n < \infty$). Under the behavior rule (1), the network will converge with probability 1 to a monomorphic steady state.*

[13] In general, $Pr[A \mid C] \geq Pr[A \cap B \mid C], \forall A, B, C$

Proof. The previous result shows that the probability that the network will NOT converge to a monomorphic steady state in the next T periods is bounded below 1, given that it has not converged until the current period. Formally:

$$P_t\{[(M-1)_{t+T}]^c \mid [(M-1)_t]^c\} \leq 1 - \underline{C} < 1$$

The event that the network will never converge is just the the intersection of the events were the network is not converging after $t+T$ periods, given that it has not converged until period t.

$$\{\text{The Network never Converges}\} = \{NC\} = \bigcap_{i=0}^{\infty}\{[(M-1)_{t+(i+1)T}]^c \mid [(M-1)_{t+iT}]^c\}$$

But again these expressions are independent. Namely, not converging until $t+2T$ given that it has not converged until $t+T$ is independent of not converging at $t+T$ given that it has not converged until t: $\{[(M-1)_{t+2T}]^c \mid [(M-1)_{t+T}]^c\} \perp \{[(M-1)_{t+T}]^c \mid [(M-1)_t]^c\}$.

Hence, we can transform the above expression in terms of probabilities:

$$\begin{aligned}P[\{NC\}] &= \lim_{s\to\infty}\prod_{i=0}^{s} P\{[(M-1)_{t+(i+1)T}]^c \mid [(M-1)_{t+iT}]^c\} \\ &\leq \lim_{s\to\infty}(1-\underline{C})^s = 0\end{aligned}$$

So, the network will converge with probability 1 to a monomorphic steady state. □

Corollary 4.1. *For a finite network, there is always positive probability of convergence to a sub-optimal action.*

The corollary is apparent from the fact that the efficient action faces a positive probability of disappearance, as long as there are more actions chosen in the network. This result points out a weak point of the present updating mechanism, which is the inability to ensure efficiency. However, if the population is infinitely large, then we provide conditions for the diffusion of the most efficient action (see Section 5).

Probabilistic updating

One of the main assumptions of the model is that all the agents are updating their choices simultaneously and every period. However, this need not be always the case. For this reason we repeat our analysis, allowing for a richer setting. Namely, every period there is positive probability $r > 0$ for each agent of deciding to update her choice. We notice that there is still positive probability that every agent in I_k^t will change her action. The proof is identical to the one of Theorem 4.1, if we multiply the lower bound of Lemma 4.1 by $r^{I_k^t}$. Again, the network converges to a monomorphic steady state with probability 1, although convergence occurs at a slower rate.

Experimentation

The result holds even under certain forms of experimentation. In particular, we transform the updating rule as following: Every period, each agent imitates her most successful neighbor with probability $(1-\epsilon) \in (0,1)$ and with probability ϵ imitates randomly another of the actions she observes, including her current choice. Under this updating rule, the proof remains the same, multiplying again the lower bound of Lemma 4.1 by $(1-\epsilon)^{I_k^t}$. Notice, that here the experimentation is limited to observed actions, which allows us to ensure that once an action dissappears from the network it does not reappear.

5. Networks with countably infinite population

At first glance, one could doubt whether there is a difference between the cases of finitely and infinitely large networks. Throughout this section we show, why the two cases are indeed different. Intuitively, we expect different behavior, since for infinitely large networks there exist actions chosen by infinite agents, where the possibility of disappearance in finite time vanishes to zero. Moreover, the possibility that some agents are connected with infinite number of agents, turn them really important for the long-term behavior of the society.

Even between different networks with infinitely large population, there is a property that can differentiate the appropriate analysis. This is the case where some agents have unbounded neighborhoods, i.e. they are observed by an infinite number of agents. We show how this property affects the behavior of the system and we provide conditions to ensure convergence and efficiency in such networks.

To assist our analysis, we introduce the following assumption. (Keep in mind that the following assumption is used only when it is clearly stated.)

Assumption 1. [Bounded Neighborhood] Exists $K > 0 \in \mathbb{R} : k_i \equiv |N_i| \leq K$, for all $i \in N$[14] $\triangleleft$

In the rest of the section, we compare the cases where Assumption 1 holds or not, while stressing the conditions that make the results of Section 4 to fail.

5.1. Bounded Neighborhoods

In this part we assume that Assumption 1 holds. The main importance of this assumption is that it makes impossible for an agent to affect the behavior of a non-trivial portion of the population, by

[14] k_i is called the *degree* of agent i and means the number of agents that i is connected with.

receiving an extremely good outcome at some period.

An obvious, nevertheless crucial, remark is that for an infinitely large network, there must be at least one action chosen by a non-trivial portion of the population. Obviously, our analysis will be different if there is exactly one such action or there are more.

Proposition 5.1. *Take a network with countably infinite population, where there is only one action, α_k, that is initially chosen by a non-trivial portion of the population. Under behavioral rule (1), and assumption (1), the network will converge with probability 1 to a monomorphic steady state, where every agent will choose α_k.*

Proof. Let $I_k^t = \{i \in N : a_i^t = \alpha_k\}$ be the set of agents that choose α_k at time t and analogously $\neg I_k^t = \{i \in N : a_i^t \neq \alpha_k\}$ the set of agents that do not. By construction of the problem, I_k^t has infinite members, while $\neg I_k^t$ has finite members. For the rest of the notation recall lemma 4.1.

Notice that, every action $k' \neq k$ is chosen by a finite number of agents, so the longest distance, $L_{k'}$, between an agent choosing k' and the closest agent choosing $k'' \neq k'$ must have finite length, $l \leq L_{k'}$. Hence, for all $k' \neq k$, the result of Lemma 4.1 still holds, if we substitute the diameter $d_{\mathcal{N}}$ by the maximum of all these distances, say L.

$$Pr(D_{\alpha_{k'}}^t) \geq Pr(\bigcap_{\tau=0}^{L} C^{t+\tau}) = \prod_{\tau=0}^{L} Pr(C^{t+\tau} \mid C^{t+\tau-1}) \geq \prod_{\tau=0}^{L} b_{t+\tau}^{|N_{i,\alpha_{k'}}^{t+\tau}|}(1 - B_{t+\tau})^{|N_i \setminus N_{i,\alpha_{k'}}^{t+\tau}|} = \dot{B} > 0$$

Notice, as well, the importance of the assumption for bounded neighborhoods. If this did not hold, then we could not ensure that the above product would be strictly positive. Moreover, the expression does not hold for the action α_k. The bounded neighborhood assumption can hold only if the network has infinite diameter. Given that α_k is the only action chosen by infinite number of agents, L_k has to be infinite, giving a trivial lower bound equal to zero.

More intuitively, action α_k faces a zero probability of disappearance in finite time. This is because, each one of the agents choosing a different action can affect the choice only of a finite number of agents, each period. Hence, it is not possible that a non-trivial portion of the population will stop choosing α_k in a finite time period.

Subsequently, the result of Lemma 4.2 still holds with some appropriate modification. Namely $\tau = L$ and $Pr[K_{t+\tau} \mid (K-1)_t, I_k^{t+\tau} \neq \emptyset] \geq \ddot{B} > 0$, meaning that action α_k cannot disappear from the network.

With similar reasoning, we get the modification of Lemma 4.3, which tells that there is positive probability of convergence to a monomorphic steady state in finite time, given that action α_k will not disappear. But, this is equivalent to the case where every agent chooses action α_k, denoted as $\{CA_k\}$. Formally, $P_t[\{CA_k\}] \equiv P_t[(M-1)_{t+T} \mid I_k^{t+T} \neq \emptyset] \geq \ddot{B}^{(M-K-2)} = \underline{C} > 0$.

Finally, as in Theorem 4.1, we get a similar expression, showing that the probability that agents' behavior will not converge to the same action, given as well that action α_k is still present, is bounded below 1. However, this is equivalent to the event of "Not converging to action α_k". This is because action α_k cannot disappear, hence if the system converges to a single action, this has to be α_k. Namely:

$$P[\{NCA_k\}] \equiv P_t\{[(M-1)_{t+T}]^c \mid [(M-1)_t]^c, I_k^{t+T} \neq \emptyset\} \leq 1 - \underline{C} < 1$$

$$P[\{NCA_k\}] = \lim_{s\to\infty} \prod_{i=0}^{s} P\{[(M-1)_{t+(i+1)T}]^c \mid [(M-1)_{t+iT}]^c, I_k^{t+T} \neq \emptyset\}$$

$$\leq \lim_{s\to\infty} (1-\underline{C})^s = 0$$

Which means that the network will converge with probability 1, to a monomorphic steady state, where every agent chooses the action α_k. □

The above proposition covers the case where there is only one action chosen by a non-trivial part of the population. The question that follows naturally is whether the network has the same properties when more than one actions are diffused to infinite agents. For a general network and payoff structure, the answer is negative. The negative result is supported by a counter-example.

Proposition 5.2. *Take an arbitrary network, with countably infinite agents behaving under behavioral rule (1) and satisfying assumption (1). If there are more than one actions chosen by infinite number of agents, we cannot ensure convergence to a monomorphic steady state, without imposing further restrictions in the network or/and payoff structure.*

The argument is proven by the following example.

Example 5.1. *Think of the following infinitely large network (figure 1- linear network).*

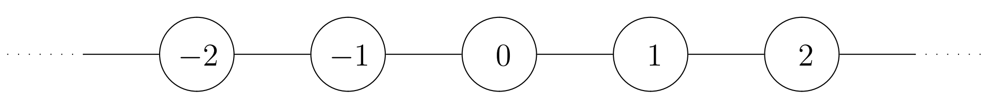

Figure 1: Line...

Notice that, all the agents have bounded neighborhoods, since they are connected with exactly two agents and the diameter of the network is infinite. At time t, there are two actions still present, α_1 and α_2. A line with size equal to half of the population choose each action (i.e. all the agents *left from zero choose α_1 and the rest choose α_2). Both actions have the same support of utilities, $U(\alpha_1, \Omega) = U(\alpha_2, \Omega) \in [0, 1]$. The distributions are as shown in the following figure.*

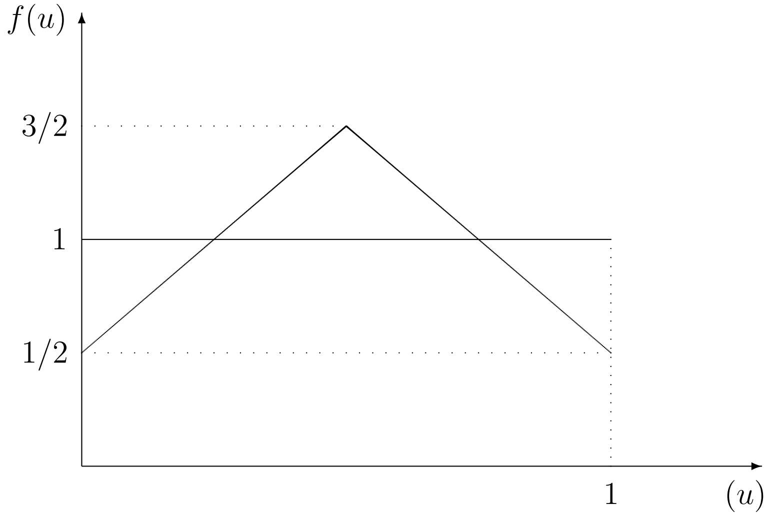

Figure 2: Probability density functions of α_1 and α_2.

For these distributions, an agent choosing action α_1 is, ex-ante, equally likely to get higher or lower utility compared to an agent choosing action α_2, and vice-versa[15]*. Moreover, notice that the only agents who can change are those in the boundary between the groups using each action. This boundary will be moving randomly, but in the form of a random walk without drift. This means that this boundary will never diverge to infinity, not even in the long-run.*

The divergence of the boundary is identical to the disappearance of one of the two actions, i.e. convergence to a monomorphic steady state. Since we have shown that this is not possible, it means that this network will never converge to a monomorphic steady state. In fact it will be fluctuating continuously around zero, without reaching a steady state.

The negative result of the previous example does not allow us to ensure convergence under every structure, when at least two actions are chosen by non-trivial portion of the population. However, there exist sufficient conditions, related to the payoff and network structure, that can ensure convergence to a monomorphic steady state.

In the following proposition, we consider cases where all agents have bounded neighborhoods, as well as, that all the remaining actions are chosen by infinitely many agents. These two facts yield the initial remark that no action will disappear in finite time. Nevertheless, it is shown that it is possible

[15] $P[X_1 > X_2] = \int_0^1 P[X_1 > x_2] f_2(x_2) dx_2 = \int_0^{1/2} (1 - x_2)(\frac{1}{2} + 2x_2) dx_2 + \int_{1/2}^1 (1 - x_2)(\frac{5}{2} - 2x_2) dx_2 = \frac{1}{2}$

for some action to be diffused to a continuously increasing share of the population, capturing the whole of it in the long-run.

Proposition 5.3. *Take a network with countably infinite population, where every agent has bounded neighborhood, behaves under behavioral rule (1) and each of the remaining actions, $\{\alpha_1, ..., \alpha_m\}$ is chosen by a non-trivial share of the population, $\{s_1, ...s_m\}$. If there exists action α_k such that:*
(i) $F_k(u) \leq [F_{k'}(u)]^D$, for all $k' \neq k$[16], and
(ii) s_k is sufficiently large,
then the network will converge, with probability one, to a monomorphic steady state where every agent chooses action α_k.

Proof. Notice, that $F_k(u) \leq [F_{k'}(u)]^D \Leftrightarrow Pr(U(\alpha_k, \Omega) \geq \hat{u}) \geq 1 - Pr(U(\alpha_{k'}, \Omega) \leq \hat{u})^D$, for all $k' \neq k$ and $\hat{u} \in U$. Before convergence occurs, there is at least one pair of agents such that $a_i = \alpha_k$ and $a_j = \alpha_{k'}$, with $k' \neq k$. Condition (i) ensures that for all agents having at least one member of their neighborhood (including themselves) choosing α_k at period t, i.e. $\forall i \in N$ such that, $\exists j \in N_i$: $a_j^t = \alpha_k$, it holds that $P(a_i^{t+1} = \alpha_k) = p_1 > \frac{1}{2}$. Since this holds for each agent we can construct the following variable, which we assume, in the first part of the proof, to be a finite number. Let:

$$r_{t+1} = \{\#i \in N : a_i^{t+1} = \alpha_k, a_i^t = \alpha_{k'}, k' \neq k\} - \{\#j \in N : a_j^{t+1} = \alpha_{k'}, a_j^t = \alpha_k, k' \neq k\}$$

Namely, this represents the difference between the number of agents changing from any action $\alpha_{k'}$ to α_k, minus those that change from α_k to any $\alpha_{k'}$. The expected value of r_t at each round will be, at every time period t:

$$E_n[r_t] = \sum_{i:\exists j \in N_i}^{a_j^t = \alpha_k} [p_1 - (1 - p_1)] = \sum_{i:\exists j \in N_i}^{a_j^t = \alpha_k} [2p_1 - 1] \geq r > 0, \text{ for all } t$$

The last inequality is direct implication of condition (i) and the assumption about finiteness of r_t. Since for every agent is more probable to change to α_k, we expect that more agents will change from other actions, actions to α_k than the opposite.

In order to prove convergence, we need to show that in the long-run and as the population becomes infinite it holds that:

$$\sum_t r_t + s_k n \to n \Leftrightarrow \lim_{n \to \infty} \frac{\lim_{t \to \infty} \sum_t r_t}{n} + s_k \geq 1 \tag{6}$$

By the weak law of large numbers, applied for the population, we get with probability 1:

$$\lim_{t \to \infty} \sum_{\tau=1}^{t} r_\tau = \lim_{t \to \infty} \sum_{\tau=1}^{t} E_n[r_t] \geq \lim_{t \to \infty} rt$$

[16] F_k is the cumulative distribution of payoffs associated with action α_k. D is the upper bound of the neighborhoods and is an exponent of F.

Which makes equation (6) equal to:

$$\lim_{n\to\infty} \frac{\lim_{t\to\infty} \sum_t r_t}{n} + s_k \geq \frac{\lim_{t\to\infty} rt}{\lim_{n\to\infty} n} + s_k = r + s_k \geq 1$$

The equality holds because the limits over t and over n are equivalent. Both variables increase with the same speed. Finally, the last inequality comes from condition (ii) that requires s_k to be sufficiently large. Sufficiency depends on r, which itself depends on the distributions of the payoffs. Obviously, the more efficient is action α_k, the smaller initial population share is needed in order to ensure convergence.

Recall that we assumed r to be a finite number. if r is a share of the population, it is apparent that the result still holds. Even more, the result holds without any restriction on s_k, because the limit we calculated diverges to infinity. □

Our sufficient condition is weaker than first order stochastic dominance, nevertheless we have the advantage of providing a result adequate for every network structure. The important fact in our proof, is that the agents changing to α_k are, in expectation, more than those changing from α_k to some other action. In general, this condition may depend not only on the payoff structure and the initial share, but also on the network structure, which is something we completely disregarded in the present analysis. Hence, if we can construct some condition depending on both network and payoff structure, that yields the same result, then we could again ensure convergence. This result can become the benchmark for future research on stronger conditions for specific structures.

Moreover, it is somehow surprising that a condition as strong as first order stochastic dominance may not be sufficient. This happens either because of the complexity of the possible network structures, or because of insufficient initial share of the analogous action. To clarify more this argument, we construct the following example.

Example 5.2. *Take a linear network, as in example 1. At time $t = 0$, there are two actions present in the network, α_1 and α_2. A line with size equal to half of the population chooses action α_1 and, analogously, the other half chooses α_2. Notice that each agent has a neighborhood of two agents apart from herself. Moreover, every period, there are only two agents, one choosing each action, that face positive probability of changing their chosen action.*

The payoffs are such that $F_1 < F_2$ (i.e. action α_1 First Order Stochastically Dominates action α_2), but $F_1 \geq (F_2)^2$.[17] *Calling r_t the number of agents that change from action α_2 to action α_1 we get the following:*

[17] For example, let $F_1(u) = u$ and $F_2(u) = u^{\frac{2}{3}}$

$$r_t = \begin{cases} 1 & \textit{with probability } p_1 \\ 0 & \textit{with probability } p_2 \\ -1 & \textit{with probability } p_3 \end{cases}$$

FOSD ensures that $p_1 > p_3$. However, the second condition means that $p_1 \leq p_2 + p_3$, or equivalently $p_1 - p_3 < \frac{1}{2}$. The expected value of r_t will be $E[r_t] = (p_1 - p_3)$. So, $\frac{\lim_{\tau \to \infty} \sum_{t=1}^{\tau} r_t + s_1}{lim n \to \infty n} = p_1 - p_3 + s_1$. In order to get diffusion of α_1 we need $p_1 - p_3 + \frac{1}{2} \geq 1$, which here cannot be the case. Hence, although α_1 is FOSD compared to α_2, α_2 will be chosen by infinitely many agents, even in the long run.

Finally, notice that if $F_1 \leq (F_2)^2$, then it can be the case (not necessarily) that $p_1 - p_3 \geq \frac{1}{2}$ and action α_1 is diffused to the whole population.

5.2. Unbounded Neighborhoods

In case neighborhoods are unbounded, even for finite diameter, we cannot ensure convergence to monomorphic steady state. This happens because a single agent can affect a non-trivial portion of the population, meaning that in one period an action can be spread from a finite part of the population to a non-trivial portion of it. This means that we will have more than one actions played by infinite agents, which may lead the network not to converge to a monomorphic state. We clarify the above statement with the following example. Moreover, in case of unbounded neighborhoods, it does not hold the result of proposition 4.1, stating that steady state has to be necessarily monomorphic. We provide an example where a network is in steady state, with more than one actions present.

Example 5.3. *Think of the star network(figure 3 - one star), that satisfies finite diameter and the central agent has unbounded neighbors.*

Suppose there are two actions present in the network, with payoffs same as in example 1. Initially, all the (infinitely many) peripheral agents choose action α_1, while the central agent chooses α_2. It is apparent that the central agent will change her action in the second period, however she will make infinite peripheral agents change to α_2. This is because, for any realized outcome of $U(\alpha_2, \cdot) = u$ it holds that $Pr[U(\alpha_1, \cdot) \leq u] = F_1(u) > 0$.

Given that infinite agents will choose each action, there will be at least one receiving the highest possible outcome, hence the choice of the central agent will change randomly from the second period on. This leads the network to a continuous fluctuation, where infinite number of agents chooses each action. Obviously, the system will never converge to a steady state.

We have shown that we cannot ensure convergence to a steady state in the above example. In the following part we provide an example where the network can be in steady state, without this

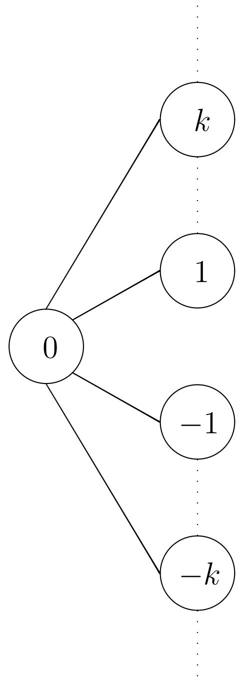

Figure 3: Star...

being monomorphic. This is also a counter-example to proposition 4.1, where we have shown that for finite population, steady states are necessarily monomorphic. For infinite population, this may not be the case.

Example 5.4. *Think of the following network (figure 4 - two stars), that satisfies finite diameter and some of the agents have unbounded neighborhoods.*

Suppose there are two actions present in the network, with payoffs same as in example 1. All the agents on the left star of the figure, including the center choose α_1 *and similarly all the agents on the right star choose* α_2*. The central agents are connected with infinitely many other agents choosing the same action as them and only one choosing differently. Hence, they will continue acting the same, with probability 1*[18]*. The peripheral agents have only one neighbor each, who always chooses the same action as them, so none of them will change her action either. Concluding, this network will be in a steady state where half of the agents choose each action.*

The important difference in this network is the existence of two groups of infinite agents, that are connected only with agents choosing the same and having only a finite number of links with

[18]The extreme case where one of the two centers change action, will happen only if the other center receives at some point the absolute maximum, but by continuity of the distributions, this will happen with probability 0

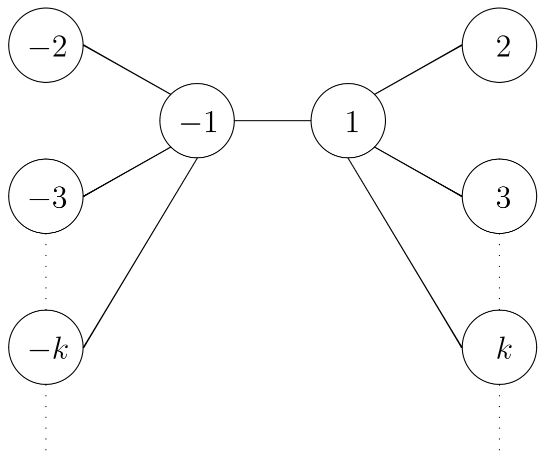

Figure 4: Double star...

agents that choose differently. Moreover, this links are so well connected with agents choosing the same as them, that they never revise their choice. The intuitive conclusion of this observation is that the existence of such groups can ensure the survival of the chosen action in the long-run, since no member of this group is going to change her choice

6. Discussion

This paper considers a model of observational learning in social networks, where agents imitate the behavior of their most successful neighbor. We focus on diffusion of a single action in the whole population and the conditions under which this is possible. Our analysis expresses the different properties between populations of finite and infinite agents.

For finite population, we show that the network necessarily converges to a steady state, and this steady state has to be monomorphic. This happens, mainly, because an action chosen by finitely many agents is vulnerable to a series of sequential negative shocks, that can lead to its disappearance. This fact, combined with the assumption that actions never reappear, leads to the survival of a single action in the long-run. Moreover, it cannot be ensured that the action that is diffused is the socially optimal. In practical terms, this means that small populations can be easily manipulated and even misguided to take sub-optimal decisions. Additionally, for the establishment of a product, or a technology in a society, is crucial that the agents initially choosing it will not receive bad shocks in the first periods, since this can lead to early extinction from the society.

The results differ in case of infinitely large population. In this case, we differentiate between networks where all the agents have bounded neighborhoods, i.e. are observed by a finite number of agents, and networks where there exist agents observed by non-trivial portion of the population.

Under the assumption of bounded neighborhoods, an action is necessarily diffused, regardless of its efficiency, only when it is the only one to be chosen by infinitely many agents. If not, we need some additional conditions in the payoff or/and the network structure to ensure convergence to a monomorphic steady state. This leads to the conclusion that, if there are no agents influencing very large part of the population, for the establishment of a product or a technology in a society, is required to be introduced initially to a sufficiently large portion of the population.

To overcome such problems, we provide a condition in the payoff structure, stricter than first order stochastic dominance, that can ensure convergence regardless of the network architecture. We show an example where first order stochastic dominance fails to guarantee convergence, while our condition does. An important conclusion, is that the diffusion of an action in a very large network is quite hard and requires either very special structure, very large proportion of initial adopters, or the action to be much more efficient compared to the rest. This happens even for the most efficient action.

Also, our sufficient condition is valuable mostly in networks where the maximum neighborhood has quite small size. This increases the importance of studying the role of network architecture, rather than the payoff structure, when applying to network where agents have large neighborhoods. The advanced complexity of this problem, makes hard to deal with its general version. A very interesting and natural extension of the present paper would be to study the behavior of specific network for different payoff structures and vice versa, with importance in applied problems. Our results could work as a benchmark for initial intuitions on the general behavior of such models.

The properties of the network change once again, when the property of bounded neighborhoods is not satisfied. Under these circumstances, we provide an example where an action can survive in the long-run, even if it is initially chosen by a single agent. This happens if this agents behavior is observed by infinite agents. This can lead the network never converging to a steady state. In general, this stresses the importance of central agents in a social network and the influential power they have to the society. Introducing a technology to massively observed agents, such as social media, can become very beneficial for its diffusion.

The last example exposes the possibility of converging to a steady state where more than one actions survive, unlike what we have shown for finite networks. This is because we observe groups of infinite agents that are connected only with other agents choosing identically, except of a finite number of links with agents choosing differently. Those agents will never observe, hence never change their strategy. If, additionally, the agents that observe different actions have infinite links with their own group, then they will never change as well. Concluding, the members of such groups will keep using the same strategy forever.

In the present paper, we have shown some very important properties of an ”imitate-the-best” mechanism in a social network. However, there are still crucial questions to be answered. Specifically, our analysis is referring only to long-run behavior, without mentioning neither the speed of convergence, nor the finite time properties of the society. Different network structures are expected to have much different speed of convergence. For example, networks with lower connectivity, or where there are no groups of agents choosing the same action, may retard the diffusion of an action to the population.

Concluding, learning by imitation, and especially imitation of successful outcomes is a commonly observable behavior in real life societies. This paper provides some explanations of societies whose members act likewise and seek to turn attention into further analysis of similar problems.

References

Acemoglu, D., Dahleh, M.A., Lobel, I. & Ozdaglar, A. (2011). Bayesian Learning in Social Networks. *Review of Economic Studies*, doi: 10.1093/restud/rdr004(1994).

Alós-Ferrer, C. & Weidenholzer, S. (2008). Contagion and efficiency. *Journal of Economic Theory* 143, 251–274.

Apesteguia, J., Huck, S. & Oechssler, J. (2007). Imitation – Theory and experimental evidence. *Journal of Economic Theory* 136, 217–235.

Bala, V. & Goyal, S. (1998). Learning from Neighbours. *Review of Economic Studies* 65, 595–621.

Banerjee, A. (1992). A Simple Model of Herd Behavior. *Quarterly Journal of Economics* 107, 797–817.

Banerjee, A. & Fudenberg, D. (2004). Word-of-mouth learning. *Games and Economic Behavior* 46, 1–22.

Billingsley, P. (1995). *Probability and measure*. John Wiley & Sons, New York.

Conley, T. & Udry, C. (2010). Learning About a New Technology: Pineapple in Ghana. *American Economic Review* 100, 35–69.

DeGroot, M.H. (1974). Reaching a Consensus. *Journal of the American Statistical Association* 69, 118–121.

Ellison, G. & Fudenberg, D. (1993). Rules of Thumb for Social Learning. *Journal of Political Economy* 101, 612–644.

——— (1995). Word-of-mouth communication and social learning. *Quarterly Journal of Economics* 109, 93–125.

Eshel, I., Samuelson, L. & Shaked A. (1998). Altruists, egoists and hooligans in a local interaction model. *American Economic Review* 88, 157–179.

Fosco, C. & Mengel, F. (2010). Cooperation through imitation and exclusion in networks. *Journal of Economic Dynamics and Control*, doi:10.1016/j.jedc.2010.12.002

Fudenberg, D. & Levine, D., (1998). *The Theory of Learning in Games*, MIT Press, Cambridge, MA.

Gale, D. & Kariv, S. (2003) Bayesian Learning in Social Networks. *Games and Economic Behavior* 45, 329–346.

Golub, B. & Jackson, M.O. (2010). Naive Learning in Social Networks and the Wisdom of Crowds. *American Economic Journal: Microeconomics* 2, 112–149.

Jackson, M.O. (2008). *Social and Economic Networks*. Princeton University Press.

Jackson, M.O. & Yariv, L. (2007). Diffusion of Behavior and Equilibrium Properties in Network Games. *American Economic Review, Papers and Proceedings*, 97, 92–98.

Josephson, J. & Matros, A. (2004). Stochastic imitation in finite games. *Games and Economic Behavior* 49, 244–259.

Lopez-Pintado, D. (2008). Contagion in Complex Networks. *Games and Economic Behavior* 62, 573–590.

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

SCHLAG, K. (1998). Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits. *Journal of Economic Theory* 78, 130–156.

——— (1999). Which one should I imitate? *Journal of Mathematical Economics* 31, 493–522.

SMITH, L. & SORENSEN, P. (2000). Pathological Outcomes of Observational Learning. *Econometrica* 68, 371–398.

WEIBULL, J. (1995). *Evolutionary Game Theory*. MIT Press, Cambridge, MA.

VEGA-REDONDO, F. (1997). The evolution of Walrasian Behavior. *Econometrica* 65, 375–384.