Information Aggregation in Networks

MSc Thesis

Submitted by: Mattan Lebovich

Supervisor: Dr. Kobi Gal, Dr. Anat Lerner

Open University Ra'anana, Israel

Abstract

This thesis studies information aggregation in networks in which agents' utilities depend on an unknown event. Agents initially receive a noisy signal about the event and take actions repeatedly while observing the actions of their neighbors in the network at each round. Such settings characterize many distributed systems such as sensor networks for intrusion detection and routing systems for internet traffic. The thesis formally defines notions of convergence in action and in knowledge for agents in general networks, including rules of behavior for a general class. Our theoretical results show that (1) agents converge in action and in knowledge for a general class of decision making rules and for all network structures; (2) all networks converge to playing the same action regardless of the network structure; and (3) for particular network configurations, agents can converge to the correct action when using a well defined class of myopic decision rules. These results are supported by a new empirical test-bed that allows researchers to simulate agent behavior over time for different types of networks and behavior rules. These results have implications for network designers in showing

Contents

List of Figures

List of Tables

1 Introduction

In many networked decision-making settings information about the world is distributed across multiple agents, and agents' success depends on their ability to convey this information to others over the network. Examples abound, and include sensor networks for intrusion detection, routing systems for facilitating traffic, and consensus tasks to determine tag images on the web or classify galaxy images taken by the Hubble space telescope [9, 22, 24].

In such settings, the structure of the network and the way agents convey information in the network are key determinants of how agents aggregate information over time. Thus, there is a growing need to understand the way different rules of behavior affect whether and how agents converge in knowledge and in action in different types of network structures. This thesis presents a formal model of information aggregation in networks in which information is distributed among different agents and agents make decisions over time. In our setting, agents receive a noisy signal about the result of a certain process which the agents do not observe. At each round, the agents take actions (e.g., decide whether an intrusion occurred, declare their preferred sports team to win a match) while observing their neighbors' actions in the previous round. For example, an agent in a sensor network can revise its belief about whether an intrusion occurred based on its own signal and the information it receives about the signal from its neighbors' actions. Over time, the agent uses its observations to try to infer the signal that its neighbors received, as well as the signal that its neighbors' neighbors received, and so on.

Our study focuses on notions of convergence in action and in knowledge for agents in the network. We show that for a large class of decision-making rules, each agent is guaranteed to converge in action and in knowledge for any network structure. In this class of rules, agents choose the action that is most likely shared by other agents in the worlds that they consider possible. We show that when all agents use this rule we can guarantee that agents converge to the same action in particular network configurations. We also provide the conditions under which this convergence is correct, in that the action is optimal for agents given the initial signals they receive.

Our theoretical results are supported by a new open-source empirical test-bed for facilitating the study of information aggregation in general networks. The software is flexible, in that it allows to vary the complexity of the decision-making over a variety of dimensions such as the network structure, agents' decision-making rules, and the number of agents. It allows to empirically study the effects of complex network structures on agents' behavior, show how agents' knowledge about the world changes over time, and provide explanations for why agents choose certain actions given their observations about other agents. This system can be of use to network designers that wish to measure the effects of different network structures on agents' decision-making.

The contribution of this thesis is fourfold. First, it formalizes a novel framework for information aggregation in general networks. Second, it uses the formalism to define the conditions under which agents converge in action in general networks. Third, it shows that for particular network structures, agents' behavior is guaranteed to converge to the correct action when using a large class of myopic decision-rules. Fourth, it presents a novel experimental toolbox for information aggregation analysis and for explaining agents' decisions in the networks.

2 Related Work

Information aggregation and group consensus has been studied from economical, computational and behavioral perspectives. We expand on each of these aspects in turn.

Work in economics has its roots in DeGroot's work [11] which constructed a model that describes how a group can reach consenses on a common probability distribution. Each individual has its own opinion about the distribution for the unknown value of some parameter. Agents communicate their subjective beliefs about the signal to all agents in the form of distributions. It is shown that convergence can be reached by a linear weighting of their beliefs. This work put down the foundations to many other similar models [12, 3, 17].

Rational learning requires that all agents in every period consider the set of possible information sets of all other agents and how their choices impact the information sets of their neighbors in the subsequent period. This is an increasingly complex tasks in large networks. Relevant to our work is the model by Meuller-Frank [20] which models rational learning in networks leads to the following outcomes. Like our own work, the notion of learning refers to the refinement of agents' information sets, which are the samllest subset of the state space which the agent knows to contain the true state of the world. First, this work shows that rational learning can lead to heterogenous choices by agents. Specifically, any two agents that receive the same observations disagree about their final actions only if they are indifferent between them. Second, he shows that the speed of convergence is a function of the diameter of the networks. Lastly, under some conditions, incomplete networks (where agents' histories of actions are not common knowledge) may actually dominate complete networks in terms of the speed of convergence. Our own work is novel in showing that a general class of myopic strategies can be optimal for many types of networks. Further, we provide a theoretical guarantee of convergence in our model.

Several works have emphasized the types of dysfunctional outcomes that can ensue when public information eclipses the agents' private knowledge. Bikhchandani et al. [7] provide two types of models for the spread of information in a group of agents. In the observable-signals scenario, the information signals form a pool of public information as they arrive. Because all past signals are publicly observed, information keeps accumulating so that individuals, all of whom have the same payoffs from taking the same action, eventually settle on the correct action and behave alike. In contrast, in the observed-actions scenario, information may not accumulate when the pool of public information becomes more informative than the private signal of agents. They show that when communicating actions, the agents may reach an information cascade, in which public information stops accumulating and agents ignore their signals.

Banerjee and Fundenberg [6] compute a Bayesian Nash equilibrium which is a mapping from history (signals and others actions) to actions for each agent. This equilibrium uses information contained in the decisions made by others participants, which makes each persons own decisions less responsive to her own information and hence less informative to others. Thus agents may ignore their own information and join the herd (this phenomena is called herding). Other game theoretic approaches have focussed on coordinated strategies, defining a coordinated equilibrium (with multiple equilibria) over agents actions given their private signals [2]. Other works in economics has defined conditions under which agents' actions converge to equilibrium given that agents are fully rational, Bayesian decision makers [1, 23].

Blume et al. [8] studied the spread of cascading failure in networks in which each node in the network has a threshold that is generated stochasitcally, and the node fails if the number of failed neighbors exceed the threshold. Dodds and Watts [13] suggested a contagion model which incorporate memory of past interactions. Kempe et al. [19] study the problem of choosing the best individuals given a social network for seeding a contagion process. This problem has important applications like viral marketing in which a few influential members of the network are targetted to trigger a cascade of influence by which friends will recommend the product to other friends, and many individuals will ultimately try it. They provide the first provable approximation for the NP-hard problem.

In Acemoglu et al. [2], a large number of agents sequentially choose between two actions, In addition to his own signal, each agent observes the choices of others in a network setting that is unknown to the agents and is generated stochastically. In line with our work, each individual knows the identity of the agents in his neighbourhood. But he does not observe their private signal or knows what information these agents had access to when making their own decisions. The focus in the paper is on the notion of excessively influential agents in a network and how they influence convergence.

We briefly mention prior work in the multi-agent systems literature that has modeled agents' decision-making in networks using a limited set of rules and did not analyze the evolution of their behavior over time [14]. Other work has modeled information propagation in networks over time, but assumed agents' beliefs depend only on their neighbors' previous actions rather than on the entire history [15], and did not provide theoretical guarantees of convergence. Our work is also related to possible world semantics for propositional logics of reasoning about knowledge [16, 21]. In these works agents' actions are descriptive, that is they are defined within the logic. Our formalism is distinct from these works in that agents' behavior in the network is inferred over time.

Lastly, we turn to a multitude of lab studies that have examined how people relay information in networks and have compared their behavior to theoretical models. Both the Baysian based economic models and DeGroot's learning were also verified in different laboratory experiments [10] and through large-scale data analysis [4]. Kearns et al. [18] have empirically studied the way humans coordinate and convey information in different types of graphs. Banerjee et al. [5] showed the information about a new microfinance program was injected in several important nodes (with high eigenvector centrality) of a 43 villages social network in South India, and the information was successfully diffused around the network.

3 The Model

Our model is built around a network of agents. Each agent is connected to at least one other agent, which makes them neighbors. Initially (at time 0), each agent in the network receives a signal about an unknown outcome, such as the winner of an election, whether an intrusion had occur, or the result of a sporting event. The signal represents subjective information about the event that is available to the agents and is probabilistically related to the actual outcome. Following the initial signal, at each future round, all of the agents take actions simultaneously after observing the actions taken by their neighbors in the previous round. The real outcome of the event itself is never revealed to the agents, and we assume agents' utilities fully depend on each others' signals as well as their actions.

We use a running example to illustrate the model in which agents' signals represent their vote in an election. In our example there are two candidate agents, agent Red and agent Green. The signal at time 0 for each agent is simply their own vote, because that is all of the information about the election that they have at time 0. The real results of the election is not revealed. Therefore, agents' actions at each round are a declaration of who they believe won the election. This example reflects a real-world event, the elections of Iran in 2009, which highlighted the role of social networks in the public's reluctance to accept Ahmadinijad as the declared winner of the election.

3.1 The Network and Agents

Definition A *System* is a quintuple $\Sigma = (A, S, C, N, W)$ where (1) A is a set of agents; (2) S is a set of signals; (3) C is a set of actions; (4) N is a reflexive binary relation on A. If $(i, j) \in N$, we say that j is a neighbor of i. (N need not be symmetric); (5) W is a set of possible worlds, where each world $w \in W$ consists of a signal $s \in S$ for each of the agents.

There are $|S|^{|A|}$ possible worlds in a system. Also, agent's i signal in some possible world w will be denoted as $w_i \in W$. An action $c \in C$ at round t is a declaration of an agent about who won the election. The function $N(i)$ returns the neighbors of agent i.

To describe the behavior of the agents in the system, including the signals and the actions played at every round, we introduce the concepts of a behavior rule.

Definition A *behavior rule* is a function B that assigns an action $c \in C$ to an agent $i \in A$ in a possible world $w \in W$. Formally, $B : w \times i \to c$.

For example, an intuitive behavior rule for an agent at a given round is to declare the winner to be the candidate who received the majority of declarations from itself and its neighbors in the previous round. Once we have a behavior rule, we can inductively define the realization of the system that results from applying the rule to an original set of signals.

Definition A *realization* of a system Σ is a function R from a possible world $w \in W$, a behavioral rule B, and round $t \in T$ that returns a set of action assignments F for every agent in the system $i \in A$ in every round from 1 up to t. Formally, $R : w \times t \times B \to F$.

Furthermore, for ease of presentation, we will use F_i^t to denote the action of agent i in round t in the set of action assignments F. We also sometimes add the behavioral role and possible world for the realization that created F to the notation as follows $F_i^t(B, w)$. Also, if $I \subseteq A$ is a set of agents, the actions played by them in round t can be denoted F_I^t (and $F_I^t(B, w)$ when needed). A realization of our election example will include the agents' votes in the election (the signals), as well as their declaration of the winner at each round (the actions). Note that the above definition implicitly assumes that all the agents use the same behavioral rule.

We now define the observation history of agents in the system, including their original signals, using the notion of a context.

Definition Given a realization R, the *context* for agent i in round t is a subset defined as $K_i^t(w) = \{w_i, F_{N(i)}^1, ..., F_{N(i)}^{t-1}\}$, which includes the signal for i and its own actions as well as its neighbors' actions through round $t-1$. The context captures all of the agent's observations through round t.

In our example, the context for an agent at round t includes his own vote in the election (his signal, w_i) and his own and its neighbors' declarations of the winner at each round up to round $t-1$. The set of all possible contexts, for all agents and for all possible worlds, is denoted as $K_A^T(W)$.

Now, what sort of behavior rules should we consider? A natural possibility is to envision the agents reasoning about which worlds are possible given their observations. In this approach, an agent's decision depends only on what it knows about the current set of possible worlds and ignores how it arrived at that knowledge. This intuition is captured by the following definitions.

Definition Two worlds w', w'' are *indistinguishable* to agent i at round t if $K_i^t(w') = K_i^t(w'')$, that is the context for agent i in t is the same for both worlds w', w''.

Indistinguishable worlds describe a relationship between two possible worlds, in which the agent cannot tell the different between them, because at the current round (t) they suggest the same behavior, and in all previous rounds $(1, ..., t-1)$ its neighbors behaved exactly the same. With that in mind, going back to the definition of a behavior rule, a possible world w can be now a set ("cluster") of indistinguishable worlds.

Definition Given a system $\Sigma = (A, S, C, N, W)$, a world $w \in W$ and a context $K_i^t(w)$ of an agent i at round t, the set of worlds the agent considers possible at w is denoted as the *knowledge* of agent i, $E(K_i^t(w))$.

The knowledge of an agent at some point in time describes the set of possible worlds that are still valid from the agent's point of view. This set is monotonically decreasing as more observations about the behavior of its neighbors get accumulated as the rounds pass. The notion is further illustrated in Figure 1.

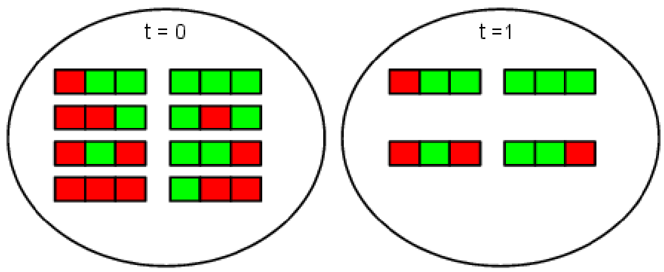

Figure 1: An example of the concept of knowledge. In this world we have 3 agents (the squares) and two possible actions (Red and Green). Looking at the world from the first agent's perspective (left square) at $t = 0$ we can see that there are 8 (2^3) possible worlds. At $t = 1$, after observing that its neighbor's signal (middle square) is Green, the set of possible worlds was reduced to 4 worlds.

3.2 Convergence properties

Following the definition of the system and its agents, our goal is to study what happens in our setting for various kinds of networks. We are particularly interested in convergence properties. Do the agents converge to the same action or to the same beliefs? The following definitions formalize convergence properties and conditions in our setting.

Definition Given a system $\Sigma = (A, S, C, N, W)$ and a behavior rule B, we say that agent i *converges in action* to action c in world w by round T if for all $t \geq T$, $F_i^t = \{c\}$.

The above definition states that an agent converges in action if his behavior remains the same after some round. In other words, if in some world w there exist some round T in which in all rounds that follows the agent always issues the same action, c, we can say that agent i converged to c. In our voting example, an agent converges in action at a round if it does not change its declared winner after the round.

Definition Given a system $\Sigma = (A, S, C, N, W)$ and a behavior rule B, we say that agent i *converges in knowledge* to set of worlds $V \subseteq W$ in world w by round T if for all $t \geq T$, $E(K_i^t(w)) = V$.

An agent converges in knowledge at a round if it does not learn any more information regarding agents' signals (their votes) after the round. In other words, convergence in knowledge occur when the information available to the agent through observation on its neighbors does not change. It is important to note that convergence in knowledge is to a set of possible worlds (V), which can include several worlds. Thus it might be that after convergence in knowledge the agent does not know exactly which world is the true world, but it does know that the true world is one within the set V. When the set V contains *exactly one world*, we can say that the agent has *complete knowledge* as defined below.

Definition We say that an agent has complete knowledge at round T if $E(K_i^T(w)) = \{w\}$.

In our example, the correct action for agents is to declare the candidate that received the most votes as the winner. An agent has complete knowledge when it knows all the other agents' votes. With the following definitions in mind, we can show that when an agent coverages in knowledge, it will also converge in action.

Proposition 1 (convergence in knowledge $\Rightarrow$ convergence in action) *For any system Σ and behavior rule B, convergence in knowledge to a set of worlds V by round T implies convergence in action to some action $c \in C$. Formally,* $(\forall t \geq T, \exists c \in C)(E(K_i^t(w)) = V) \Rightarrow (F_i^t(B, w) = \{c\})$

Proof *The proposition follows from the fact that convergence in knowledge results in a set of indistinguishable worlds, V. When the world converged in knowledge to a set V, by definition the agent will not be able to separate any of the worlds in V. As such, by the definition of behavioral rule, it will result in the same action, c in our case, for any world in the set V. Each of the agents in the system sees the same signals, which consequently results in taking the same action (the behavior rule is deterministic). Note that the other side of the implication, that is convergence in action does not necessarily implies the convergence in knowledge.*

Definition We denote the action that should have been taken if the information was fully available to the agents as the **correct action**. Formally, c is the correct action for agent i if $E(K_i^T(w)) = \{w\} \Rightarrow B(w,i) = c$.

In the voting example, if the agent would have known the real polls counting, he would have been able to declare the correct winner (or in other words take the correct action). The following proposition states that when an agent has a complete knowledge, it knows the exact world (w), and as such it will choose the correct action in future rounds.

Proposition 2 (complete knowledge $\Rightarrow$ converges to the correct action) *If agent i converges in knowledge to a single possible world, than the agent knows the actual world, and consequently, the correct action to take. Formally, $(\forall t \geq T, \exists c \in C)(E(K_i^t(w)) = \{w\}) \Rightarrow (F_i^t = \{c\})$, and c is denoted as the* correct action *to take.*

Proof *The proof follows directly from the definitions of convergence in knowledge and correct action.*

We can now show our first result, stating that all finite systems will eventually converge both in knowledge and consequently in action (though, not necessarily to the *correct* action).

Theorem 1 (All systems converge in knowledge and action) *Let $\Sigma = (A, S, C, N, W)$ be a system in which A and S are finite, Let B be a behavioral rule, and let $w \in W$ be an initial world. Then there exists a round T such that Σ converges in knowledge and action by T.*

Proof *We first show that if all agents have the same knowledge at two successive rounds T and $T+1$, then Σ will have converged in knowledge and action by T. Assume, for each agent i, that $E(K_i^{T+1}(w)) = E(K_i^T(w))$. Since knowledge at T fully determines action at $T+1$ we get that*

$$F_i^{T+1}(B, w) = F_i^T(B, w) \tag{1}$$

That is, agent i in world w behaves the same in both T and $T+1$ (using Prop. 1). Now, suppose $E(K_i^{T+2}(w)) \neq E(K_i^{T+1}(w))$. Then there must be a world $w^{'}$ which was indistinguishable to i from w in round $T+1$ which became distinguishable in round $T+2$. For this to happen some neighbor j would have played a different action in $w^{'}$ from what it actually played in w at time $T+1$, but acted the same at time T. Therefore, we get that

$$\begin{aligned} F_j^{T+1}(B, w^{'}) &\neq F_j^{T+1}(B, w) \\ F_j^{T}(B, w^{'}) &= F_j^{T}(B, w) \end{aligned} \tag{2}$$

Equation 1 is stated without loss of generality, so we can plug $w^{'}$ and j and get that

$$F_j^{T+1}(B, w^{'}) = F_j^{T}(B, w^{'}) \tag{3}$$

From Equation 4 and 3 we get that

$$F_j^{T+1}(B, w^{'}) = F_j^{T}(B, w) \tag{4}$$

We plug in j in Equation 1 and get

$$F_j^{T}(B, w) = F_j^{T+1}(B, w) \tag{5}$$

From equation 5 and 6 we get $F_j^{T+1}(B, w^{'}) = F_j^{T+1}(B, w)$, but this contradicts Equation 1. Therefore $E(K_i^{T+2}(w)) = E(K_i^{T+1}(w))$. Now, since A and S are finite, W must also be finite. Since the set of worlds an agent considers possible is non-increasing, there must be two successive rounds at which the agents have the same knowledge. Therefore the theorem holds.

It's important to note that this theorem states that all systems converges in knowledge and consequently in action. However, it does not state that it will necessarily reach a state of having a *complete knowledge* (which according to Proposition 2 will result in taking the correct action). The convergence might be to a larger set of worlds, which might result in an action that might be incorrect had the agent had complete knowledge.

4 Line Network

Given that Theorem 1 has shown that every system converges in knowledge, and consequently each agent in a system converges in action, but not necessarily the same action. In This section we study whether we can show a stronger criteria, mainly that all agents converge to the *same* action, as we now define:

Definition The system (A, S, C, N, W) *uniformly converges* in action under behavior rule B in world w by round T if all agents converge to the same action by round T. We say a system uniformly converges to the *correct* action if all agents converge in action to the correct action.

The rest of this section defines special classes of systems that include line configurations for which we can guarantee uniform convergence to the correct action. To facilitate this analysis we use a constrained system in which both actions and signals are binary. All agents use the same behavior rule, in which agents choose the action that is most common in all the worlds they consider possible. This intuitive decision rule is inspired by consensus tasks, in which agents' actions correspond to choosing the most likely candidate to win given their observations, with ties being consistently broken in favor of the action f (the default action). Such a rule is optimal if we assume that the agents have a uniform distribution over worlds and their utility of playing an action in a world is proportional to the number of agents who have a signal that is the same as the action. Therefore it is in the agents' interest to make decisions that correspond to the majority signal, as in our election example.

Definition Let $\#_w(S)$ be the number of times signal s appears in world w, and $\#_W(S) = \Sigma_{w \in W} \#_w(S)$. The *majority behavioral rule* is $B_{maj} = \arg\max_w \#_W(S)$

The following definitions relate to the positions of agents and their neighbors within a network. An *Edge agent* has only one neighbor; agent i is a *Border agent* in world w at time t if there is a path from i to an edge agent in which each agent $j \neq i$ takes the same action $F_j^t(B, w)$, and $F_j^t(B, w) \neq F_i^t(B, w)$. Agent i is a *minority agent* in world w at time t given a set of agents G if i's action is played by a strict minority of agents in G, that is, $\#(F_i^t(B, w)) < \frac{|G|}{2}$. An *agent cluster* is a set of agents that take the same action. Formally, an agent cluster of size n at time t in world w is a set of agents $a_i, \ldots, a_n$ such that $F_{a_1}^t(B, w) = F_{a_2}^t(B, w) = \ldots = F_{a_n}^t(B, w)$. Lastly, we define a *Consecutive line* as a set of c agent clusters of sizes $n_1, n_2, ..., n_c$ such that the agents in each adjacent cluster take a different action.

For example, the following configuration is a consecutive line at round t with 2 clusters of size $n_1 = 2, n_2 = 3$. The agent a_2 (in bold) at round t is both a *minority* and a *border* agent (termed "minority border" agent). Agent a_3 is also a *border* agents while agents a_1 and a_5 are *edge* agents.

$$T_1 - \mathbf{T_2} - F_3 - F_4 - F_5$$

The Boolean function $cons_{n_1, n_2, \ldots, n_c}^t(w)$ returns "true" if a network configuration in world w at time t is a consecutive line of c clusters. We can now state convergence for a simple graph which includes consecutive lines of two clusters.

We begin our analysis with networks forming a simple line of 3 agents. We assume the signal for agents can be G or R, and use notation R_i to represent the fact that the signal for agent a_i was R (and similarly for G_i). For example, a line configuration of 3 agents can be $G_1 - R_2 - G_3$.

Proposition 3 (Line of 3 agents) *Let Σ be a system comprising a consecutive line of 3 agents, binary signals and actions. Then the system uniformly converges to the correct action in 3 rounds.*

Proof At round 0, the context for each agent includes solely its own signal. Therefore, the best agents can do in this round given our decision rule, is to repeat their signal. At round 1 the middle agent a_2 observes the actions of both of its neighbors, and therefore its context includes the actions representing the signals of all of the agents. As a result it has full knowledge at round 2, and by Theorem 1, it has converged to the correct action at round 2. Both agents a_1 and a_3 will necessary copy this action at round 3 onwards, because there does not exist a world in which this agent chooses the wrong action at this round.

However it is not the case that all of the agents in a 3-agent configuration converge to full knowledge. For example, take the two line configuration $G_1 - G_2 - R_3$ and $G_1 - G_2 - G_3$. In both of these configurations, agent a_2 will take action G at round 2. Therefore, agent a_1 will not be able to distinguish between the two worlds corresponding to these configurations.

Moreover, in general, line configurations do not converge to the correct action. Consider the following example of 4 agents who observe the following signals: $F_1 - T_2 - T_3 - F_4$. The context of agent a_3 at round 1 includes the signals $w_1 = F, w_2 = T$ and $w_3 = T$ at round 1, and it will thus play T at round 2. This action will reveal to agent a_1 that $w_3 = T$, or else agent a_2 would have played F at round 2. Therefore a_1 changes to T in time step 3. Now, agent a_3 will play T at round 2 for both cases $w_4 = T, w4 = F$. Therefore agent i_2 does not discover the value of w_4 at round 3 and will continue to play T at round 3. i_2 leads i_1, then i_1 will switch to T at round 3. In addition the leader i_2 will not change its action at round 3 and will continue to play T. This is because i_3 will play T at round 3 regardless of whether i_4's signal is T or F. In both of these cases, it will have observed more T's then F's. A symmetric argument can be made to show that i_4, which leads i_3, will play T at round 3. Therefore both agents i_1, i_4 are led to play T at round 4. It is easy to see that no agent will change its behavior after round 3, thus this configuration converges to the wrong action T.

We next look at class of systems in which the network includes consecutive lines of two clusters.

Theorem 2 (Convergence of two clusters consecutive line) — *Let Σ be a system comprising a line of agents, with binary actions and signals, and in which the possible worlds W that agents* *consider correspond to a consecutive line of* 2 *clusters with sizes* n_1 *and* n_2*. That is* $\forall w \in W$: $cons^0_{n_1,n_2}(w)$*. Then the system uniformly converges to the correct action in* w *after* n_1 *rounds when* $n_1 < n_2$.[1]

We will demonstrate the process of convergence using a running example of a configuration consisting of a consecutive line of two clusters $n_1 = 4, n_2 = 5$:

$T_1 - T_2 - T_3 - T_4 - F_5 - F_6 - F_7 - F_8 - F_9$

The basis of the proof relies on the Lemma below. Note that the condition $cons^0_{n_1,n_2}(w)$ means that agents know they are on a line with two clusters (but not their size).

Lemma 3 (Convergence: line of two clusters) *The following statements specify the convergence process for a line of two clusters.*

- *At round* $t \geq n_1$ *in world* w*, it holds that* $cons^t_{n_1-t,n_2+t}(w)$*. The minority border agent at round t in world w is* a_{n_1-t}
- *At round* $t+1$*, we have* $F^t_j(B, w) = F^{t+1}_j(B, w)$ *for any agent* a_j *that is not a minority border agent at round* t*.*
- *The minority border agent at round* t *has complete knowledge at round* $t+1$*.*
- *At round* $t+1$ *the minority border agent at round t takes the correct action.*

Statement 1 guarantees that at each round t, there is a consecutive line with 2 clusters with sizes $n_1 - t, n_2 + t$. Applying this statement to the configuration shown above shows that at round $t = 2$, for example, we have the following line configuration with two clusters $n_1^{'} = 2$, $n_2^{'} = 7$:

$T_1 - \mathbf{T_2} - F_3 - F_4 - F_5 - F_6 - F_7 - F_8 - F_9$

[1]When $n_1 = n_2$ it can be shown that the configuration converges to the default action.

The minority border agent at round 2 (in bold) is a_2. Statement 2 states that at round $t+1$, all agents that are not the minority border agent at round t do not change their action. In our example configuration, this means that the only agent to change action at round 3 is the minority border agent at round 2, which is a_2. Statement 3 states that the minority border agent learns the world w at time $t+1$. According to Statement 4, it will choose to change to the correct majority action at round 3, which is F (the default action in case of ties). This complete proof with the inductive process can be found in the Appendix.

4.1 Simulation Environment

We demonstrate this process using an empirical test-bed called ONetwork that is designed to study information aggregation in networks using different rules of behavior.[2]

Figure 2(a) shows the progression of agents' behavior in a consecutive line of two clusters of size $n_1 = 7, n_2 = 8$. Agents are aligned along the x-axis, time is aligned along the y-axis. Bright green nodes represent agents that take action G (Green), while bright red nodes represent agents that take action R (Red). Agents' behavior at round 0 corresponds to their signals. As shown by the Figure, the system converges uniformly after n_1 rounds, as stated by Theorem 2.[3]

ONetwork can also be queried to show an agent's subjective knowledge about others' actions given a context. In this case the colors represent the way all agents act in the worlds that the agent considers possible. Figure 2(b) describes the knowledge of agent a_4 at round 3 given its context shown in dashed outline, which includes agent a_5 changing from G to R at round 3. As a result agent a_4 learns all agents' signals. This is shown in that agents' actions at time 0 appear in bright red and green colors (solid outline). In contrast, Figure 2(c) describes the knowledge of agent a_4 at round 3 given a context in which agent a_5 stays G at round 3. As a result, agent a_4 does not

[2] ONetwork is free software and will be made available for download under the GNU public license.

[3] The numbers appearing inside a G square represent the number of possible worlds in which an agent acts G (and similarly for R). The string "CONSTRAIN" denotes a function in the software that allows to constrain an agent to behaving in a certain way. We do not relate to these in the analysis.

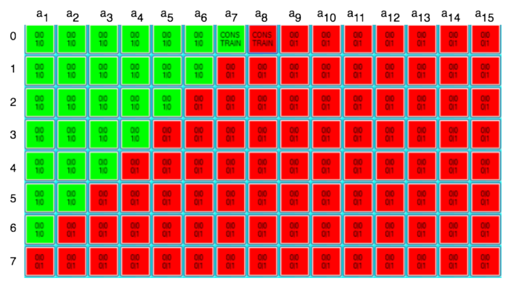

(a)

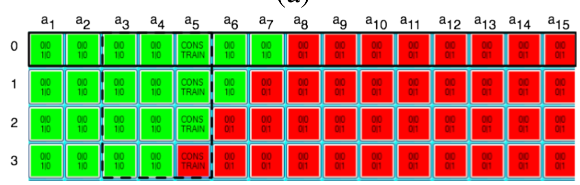

(b)

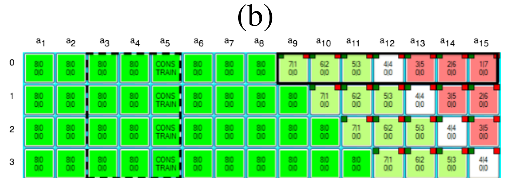

(c)

Figure 2: (a) Progression of behavior in a consecutive line of two clusters. In (b) agent a_4 has full knowledge after 4 rounds. In (c) agent a_4 never has full knowledge.

learn the signals of agents $a_{9,\ldots,13}$, because a_5 will behave the same regardless of their signals. This uncertainty is shown in the figure by using pale or white colored squares to describe the actions of $a_{9,\ldots,13}$ at round 0 (solid outline).[4] In fact a_4 will never learn the signals for these agents. However, it will still converge to the correct action, following Theorem 2.

4.2 Consecutive Lines of Three Clusters

In this section we extend Theorem 2 to consecutive lines of three clusters.

Theorem 4 (Convergence of three clusters consecutive line) *— Let Σ be a system comprising a line of agents, with binary actions and signals, and in which the possible worlds W correspond to a consecutive line of 3 clusters with size n_1, n_2 and n_3, that is $\forall w \in W$, we have that $cons^0_{n_1,n_2,n_3}(w)$. Then the system uniformly converges to the correct action after $\lceil \frac{n_2}{2} \rceil$ rounds given that $n_2 < \lceil \frac{n_1}{2} \rceil + \lceil \frac{n_3}{2} \rceil$.*

Under the conditions specified above, the action represented in the middle cluster is in the minority, and the configuration uniformly converges to the majority action represented by the first and third cluster.

Lemma 5 *The following statements specify the convergence process of a consecutive line of 3 clusters.*

1. *At round $t \leq \frac{n_2}{2}$ in world w, it holds that $cons^t_{n_1+t,n_2-2t,n_3+t}(w)$. There are two border agents in the minority at round t, namely a_{n_1+t} and $a_{n_1+n_2-t}$.*

2. *Let I be the group of border agents in the minority at round t in world w. At round $t+1$, for each agent j ($j \notin I$) it holds that $F^t_j(w) = F^{t+1}_j(w)$.*

[4] Pale green color means that the agent is more likely to choose G in the worlds agent a_4 considers possible (and similarly for pale red color), white denotes tie.

(a)

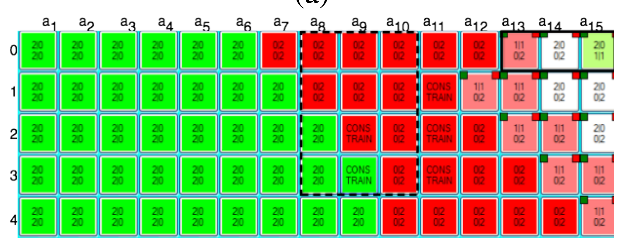

(b)

Figure 3: Progression of behavior in a line of three clusters. In (a) the system is shown to uniformly converge to the incorrect action. In (b) the knowledge of agent a_{10} is shown in round 4.

3. *Let I be the group of border agents in the minority at round t in world w. All agents in I will have full knowledge at round $t+1$ and behave correctly at round $t+1$.*

Statement 1 states that at time t there is a consecutive line with 3 clusters with size n_1+t, n_2-2t, n_3+t. Statement 2 states that all agents other than the border agents in the minority will not change their action, and Statement 3 states that the border agents will behave correctly.

Interestingly, when the size of the middle cluster no longer satisfies $n_2 < \lceil \frac{n_1}{2} \rceil + \lceil \frac{n_3}{2} \rceil$, the line configuration may not converge to the correct action. An example using ONetwork is shown in Figure 3(a), in which a consecutive line with 3 clusters of sizes $n_1 = 5, n_2 = 6$ and $n_3 = 3$ uniformly converges to the wrong action represented in the middle cluster. Here, A "fan effect" occurs in which the action represented by the first cluster propagates to the middle cluster and back. We provide an intuitive explanation for this progression. Consider for example the minority border agent a_8 at round 2. This agent observed its neighbor a_7 change to G at the previous round 1. Because agents only consider line configurations of 3 clusters, agent a_8 learns that the actions for agents $a_1, \ldots, a_6$ at round 0 was G. Also, a_8 observes that the agent a_9 did not change its action at round 2. Thus it learns that the action of agent a_{10} at round 0 was R. Now, there are more G than R signals in the worlds that a_8 considers possible and therefore it changes to G at round 2. In a similar fashion, agent a_9 changes to G at round 3. This "left-to-right" propagation terminates when agent a_{10} learns at round 4 there are more Rs than Gs in the worlds it considers possible. The knowledge of agent a_{10} at round 4 is shown in Figure 3(b). As shown by the Figure, the agent knows the signals for agents $a_1, \ldots, a_{12}$ (the brightly colored squares in round 0), and is uncertain about the actions of agents $a_{13}, \ldots, a_{15}$ (the pale, outlined, colored squares in round 0). Because there is an equal amount of Gs and Rs in the world it considers possible, agent a_{10} chooses the default action and continues to play R at round 4. This process reveals the actions of agents a_{11} and a_{12} at round 0 to agent a_9, which causes it to choose the action R at round 5. This "right-to-left" pass propagates this knowledge to agents $a_{1,\ldots,a_8}$ and they all choose R, without ever discovering the actions at round 0 of agents $a_{13}, \ldots, a_{15}$. A symmetric "fan effect" occurs with agents $a_{11}, \ldots, a_{15}$.

5 Convergence in Clique-based Network

In this section we show that it is possible to guarantee convergence for a general class of network structures characterizing many distributed decision-making settings. We will prove convergence in a network constructed of n-cliques and a constraint class of behavior rules, and through that continue to prove that in a fully connected network, all agents will always converge to the **same** action.

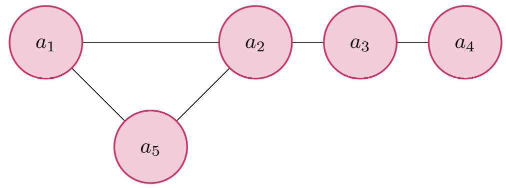

Figure 4: A network for illustrating leadership

5.1 Additional definitions

We start with few new definitions; The first being the notion of "leadership" in a network.

leader — Agent i is a leader of agent j if the following hold: (1) Agent i is a neighbor of agent j; (2) Any neighbor of j's neighbor (other than i) is also a neighbor of i. Formally, $(\forall y, x \in A) N(j, x) \wedge N(x, y) \rightarrow N(i, y)$.

To illustrate, we observe that in the network shown in Figure 4, agent a_2 is a leader of agent a_1. This is because the sole neighbor of agent a_1 (other than a_2) is agent a_5, which is a neighbor of agent a_2. Similarly, it is easy to see that a_2 is also a leader of a_5. However, agent a_2 is not a leader of agent a_3. This is because agent a_4 is a neighbor of agent a_3 but not a neighbor of agent a_2.[5] The following proposition states that all agents "follow their leader" in a general network.

Proposition 4 ("follow the leader") *If i is a leader of j, then $F_j^{t+1}(B, w) = F_i^t(B, w)$ for any behavioral rule B, any world w, and any round $t > 1$.*

Proof *By definition of a leader, if i is a leader of j, all new information available to j will be given solely by observing i. Therefore, if j issue a new action in time $t + 1$, it will only be due to new observation on i at time t.*

A leader can observe at a current round t all of the actions of the observations that the agents it leads will see at the next round $t + 1$. For example, in the network shown in Figure 1, agent a_2

[5] Leadership is not a symmetric property, as exemplified by the fact that agent a_5 is not a leader of agent a_2.

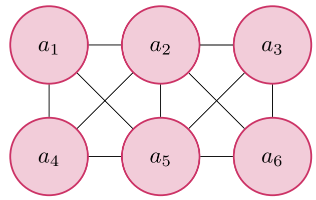

Figure 5: A 2-clique network. Agents a_2 and a_5 are the common members.

will be able to observe the signal from agent a_3 before making an action that is observed by agent a_5. In particular, all information from agents a_3 and a_4, both of which are not directly connected to a_5 is conveyed to a_5 by a_2. In addition, all agents that are directly connected to a_5 (i.e., agent a_1 and a_2 itself) are also connected to a_2. Consequently, any information that is observed by agent a_5 at round t was already observed by agent a_2 in the previous round $t-1$, Therefore, the best action agent a_5 can do under any behavior rule is to follow its leader a_2.

We can use the notion of leadership to characterize agents' behavior in a family of networks that comprise unions of (possible intersecting) complete graphs. Each of these complete graphs corresponds to a set of agents who can observe each others' actions. These may represent sensor networks deployed at an installation, a group of friends that coordinate their activities together on a social network, or work peers at an institution who share information. As agents may be members of many groups, these complete graphs may overlap. Formally, we make the following definitions.

n-clique network — Let $C_1, ..., C_n$ be a set of complete graphs. An n-clique network as a union of n complete graphs $C_1, ..., C_n$. The common members of an n-clique network are defined as $C_1 \cap ... \cap C_n$.

For example, a 2-clique network can describe sensors arranged on two sides of a border or two groups of friends with a (possibly empty) set of common members. The common members of the network may represent those sensors that can observe all of the other sensors, or those friends that are common to both groups. An example of a 2-clique network with common members is shown in Figure 5.

5.2 Convergence of 1-clique network

To formalize convergence in this class of networks, we begin with a system comprising a single complete graph (a 1-clique network). In such a configuration, all agents can observe each other's behavior. The following Proposition states that this system uniformly converges to the correct action in 2 rounds. The Proposition holds for a general class of rules of behavior that only requires agents to repeat their signals in the first round of interaction. We denote the rules that belong to this class as "first-action-signaling" rules of behavior. This is a reasonable assumption to make as the the context for each agent only includes its own signal. Note that this class does not restrict agents' strategies for any future rounds.

Proposition 5 (convergence of 1-clique network) *Let Σ be a system comprising a 1-clique. Then for any rule of behavior of the first-action-signaling class, the system will uniformly converge to the correct action at round 2.*

Proof At round 0, all agents will repeat their signal because they use a first-action-signaling rule. At round 1, each agent observes the signals of all other agents. Thus, at round 2, all agents have complete knowledge of each other's signals. knowledge. Therefore each agent will take the *same* correct action at round 2. Now, because the agents are configured in a clique, there is an edge from a_i to all other agents, and in particular, to the neighbors of a_i. By Definition 5.1, any agent a_i in the system is a leader of a_i itself. By Proposition 4, it follows that all agents will take the same correct action in all rounds following round 2, and thus the system will uniformly converge to the correct action.

5.3 Convergence of n-clique network

The following theorem claim that *all* n-cliques networks necessarily converge to the correct action.

Theorem 6 (n-clique networks convergence to the correct action)
Let Σ be a system in which the graph G is a n-clique network $C_1, ..., C_n$ in which the set of common

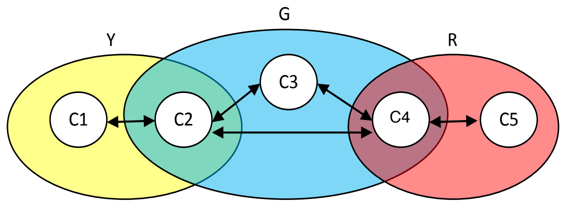

Figure 6: A 2-clique with no common members

members is not empty. Then, for any behavior rule of the first-action-signaling class, the system will uniformly converge to the correct at round 2.

Proof *As before, all agents will repeat their signal at round 0 because they use a first-action-signaling rule. Let $A = C_1 \cup C_2 \cup \ldots\ C_n$ be the set of common nodes in the network. By construction, there is an edge from agent that is a common member to all other agents. Thus, by Definition 5.1, any common agent is a leader of all agents in the network. At round 2, all common agents will necessarily have full knowledge. Therefore they will take the correct action at round 2. By Proposition 4, it follows that all agents will take the same correct action in all rounds following round 2, and thus the system will uniformly converge to the correct action.*

5.4 Convergence of cliques without common members

Lastly, we introduce networks comprising cliques that do not have a common members. An abstraction of such a network is shown in Figure 6. There are three clique groups in the Figure, represented by the yellow (Y), green (G), and red circles (R). Each node in the network represents a complete graph of agents. As shown in the Figure, the set of agents C_2 are common members to the 2-clique network containing G and Y, while the set of agents C_4 are common members to the 2-clique network containing R and G. All agents in C_2 are leaders of all agents in Y and C_3. However, there is no agent in the graph that is a leader of all other agents, and the convergence properties of the network cannot be determined in advance using Proposition 6. In these types of settings it is not possible to prove convergence to the correct action.

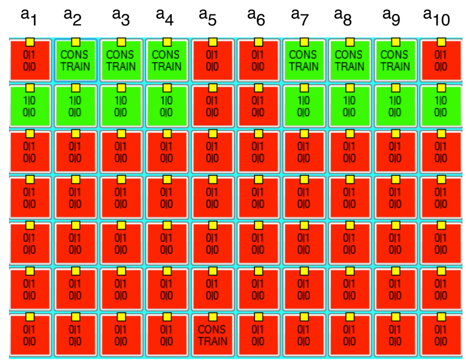

(a)

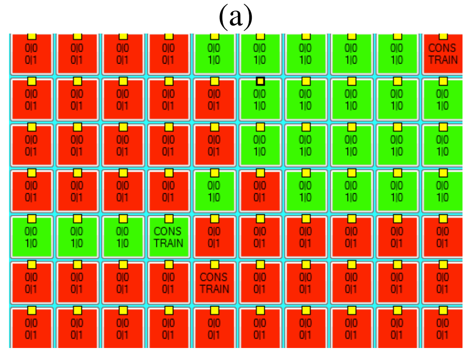

(b)

Figure 7: The two rooms example, convergence to the correct action (a) and incorrect action (b)

An example configuration of such networks are agents that are situated in two rooms separated by a wall with a window. The first room corresponds to the Y clique group and contains the agents in C_1 and C_2. The second room corresponds to the R clique and contains the agents in C_4 and C_5. All agents in C_2 are leaders of the Y clique group and all agents in C_4 are leaders of the R clique group. The leaders in C_2 and C_4 are the only agents that can observe each others action through the window. (In this example the set of agents C_3 is the empty-set).

Although convergence cannot be proven when the network has no common members, we can simulate agents' actions in the two-rooms example using the ONetwork simulation. Figure 7(a) shows 10 agents. Agent $a_1, ..., a_5$ are in the first room, and led by agent a_5. Agents $a_6, ..., a_{10}$ are in the second room, and led by a_6. Bright green nodes represent agents that take action G (Green), while bright red nodes represent agents that take action R (Red). Agents' behavior at round 0 (the first line) corresponds to their signals. At time step 1 all agents repeat their signal. At time step 2 the leader agents stay red. Therefore at time 3 all the agents turn to red and the system converges to the incorrect action (as there is a green majority in agents' signals at round 0). Figure 7(b) shows a different configuration of the two-room example. Here the leader a_5 turns to red in time step 1 so all agents it leads stay red in time 2. At time step 4 this agent turns back to green (after observing that the other leader a_6 has not turned to red in 2 time steps). Consequently, all of the agents led by a_5 turn to green in the next time step. This process continues until agents converge to the correct default action given an equal number of red and green signals.

5.5 Connected network — All agents will Act the Same

In the following subsection we will present a theorem stating that when we have a connected graph of agents, that is, there is a path from each pair of agents, after converging in knowledge, all agents will issue the same action. This will not necessarily be the *correct* action, but all will act the same. We will start with additional definitions and built our proof on them.

The following relation connects every pair of indistinguishable worlds by some agent i at time t.

Indistinguishable worlds relation (Cluster)
$Ind(i,t) = \{(w,w^{'}) \in Ind(i,t) | K_i^t(w) = K_i^t(w^{'})\}$

The set of all indistinguishable worlds relation for some agent i at time t will be denoted as $IND(i,t)$ and individual relations in it will be denoted as $Ind_n(i,t) \in IND(i,t)$.

The following proposition states that the indistinguishable worlds relation is an equivalence relation. In other words, each world can be in only a single cluster and any intersection of clusters is empty.

Proposition 6 (Equivalence relation) *$Ind(i,t)$ is an equivalence relation*

Proof *We need to show that the relation is reflexive, symmetrical, and transitive. (1) Reflexivity - identical worlds are in the relation by setting* $w^{'} = w$ *in the definition. Formally,* $(w, w) \in Ind(i, t) \Leftrightarrow K_i^t(w) = K_i^t(w)$*. (2) Symmetry - if two worlds are in the relation, their symmetric pair is also in the relation by setting* $w^{'} = w$ *and* $w = w^{'}$*. Formally,* $(w, w^{'}) \in Ind(i, t) \Leftrightarrow (w^{'}, w) \in Ind(i, t)$ *as by definition* $K_i^t(w) = K_i^t(w^{'}) \Leftrightarrow K_i^t(w^{'}) = K_i^t(w)$*. (3) Transitivity -* $(w, w^{'}) \in Ind(i, t), (w^{'}, w^{''}) \in Ind(i, t) \Rightarrow (w, w^{''}) \in Ind(i, t)$*. This is also true by definition of* K.

The theorem we present will be true under *any* behavior rule as long as it adhere to the following additive criteria.

Additive behavior rule — An *additive behavior rule* is a behavior rule which provide the same behavior under union operation of set of worlds (clusters). In other words, if the rule results in similar behavior in two clusters of worlds, it will show *the same* behavior in union of these clusters. Formally, B is an additive behavior rule, when the following holds:

$(\forall i \in A, t \in T, Ind_n(i, t), Ind_m(i, t) \in IND(i, t), n \neq m)$

$F_i^t(B, Ind_n(i, t)) = F_i^t(B, Ind_m(i, t)) = F_i^t(B, (Ind_n(i, t) \cup Ind_m(i, t)))$

In our voting example throughout the paper we have used the majority selection voting rule as our behavior rule. It is easy to see intuitively that if, for example, in two clusters of worlds the selected action is Red (as there are more possible worlds in which the Red is the winner), the union of these two clusters will still result in Red as the majority selected action. We now formally show in the following property that this behavior rule is indeed additive.

Proposition 7 (The Majority behavior rule is additive)

$(\forall i \in A, t \in T, Ind_n(i, t), Ind_m(i, t) \in IND(i, t), n \neq m)$

$F_i^t(B, Ind_n(i, t)) = F_i^t(B, Ind_m(i, t)) = F_i^t(B, (Ind_n(i, t) \cup Ind_m(i, t)))$

Proof *Without loss of generality, let* s *be the selected action according to the majority behavioral*

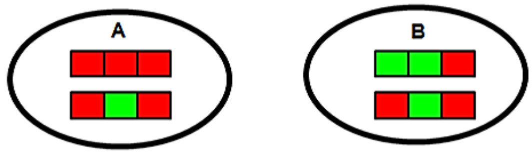

(a)

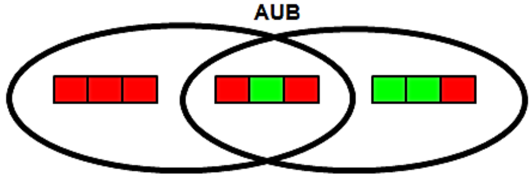

(b)

Figure 8: In the upper figure (a) we can see two clusters. In the left hand one we see two worlds that will result in a Red vote (the bottom world is a tie which will be broken towards Red as the default action). In the right hand side we see two worlds, one which will result in a Green vote (upper), and another in the Red vote (lower). So in both clusters the Red vote will prevail. In figure (b) we can see the union operation on the clusters, which will also result in a Red vote from the joint cluster (as two worlds will be Red and one will be Green).

rule (while the other action will be denoted as $\bar{s}$).

$$F_i^t(B, Ind_n(i,t)) = F_i^t(B, Ind_m(i,t)) = s$$

$$F_i^t(B, Ind_n(i,t)) = s \Leftrightarrow \Sigma\#_{w \in Ind_n(i,t)}(s) - \Sigma\#_{w \in Ind_n(i,t)}(\bar{s}) > 0$$

$$F_i^t(B, Ind_m(i,t)) = s \Leftrightarrow \Sigma\#_{w \in Ind_m(i,t)}(s) - \Sigma\#_{w \in Ind_m(i,t)}(\bar{s}) > 0$$

$$\begin{aligned}
&F_i^t(B, (Ind_n(i,t) \cup Ind_m(i,t))) = s \\
&\Leftrightarrow \Sigma\#_{w \in (Ind_n(i,t) \cup Ind_m(i,t))}(s) - \Sigma\#_{w \in (Ind_n(i,t) \cup Ind_m(i,t))}(\bar{s}) > 0 \\
&\Leftrightarrow (\Sigma\#_{w \in Ind_n(i,t)}(s) - \Sigma\#_{w \in Ind_n(i,t)}(\bar{s})) + (\Sigma\#_{w \in Ind_m(i,t)}(s) - \Sigma\#_{w \in Ind_n(i,t)}(\bar{s})) - \\
&\qquad\qquad (\Sigma\#_{w \in (Ind_n(i,t) \cap Ind_m(i,t))}(s) - \Sigma\#_{w \in (Ind_n(i,t) \cap Ind_m(i,t))}(\bar{s}))
\end{aligned}$$

$$\Leftrightarrow F_i^t(B, (Ind_n(i,t) \cup Ind_m(i,t))) = 0^+ + 0^+ - 0 \tag{6}$$

$$\Leftrightarrow F_i^t(B, (Ind_n(i,t) \cup Ind_m(i,t))) = s$$

The above proposition uses in Equation 7 the fact that when uniting two clusters relations, the disjunction of them will be empty (they are an equivalence relation), thus the additional information will not change the behavior. Figure 8 presents an example of the additivity concept.

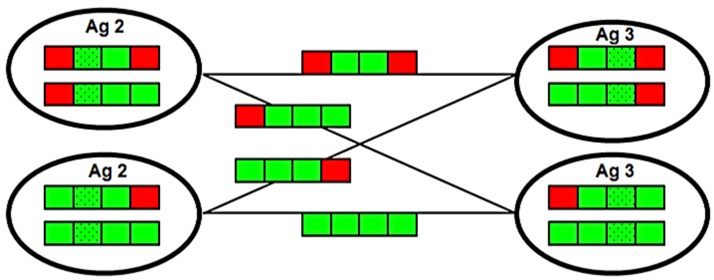

Figure 9: A shared knowledge graph between agents 2 and 3. Each agent has 4 possible worlds, that are divided to two clusters in which the worlds withing them are indistinguishable. The worlds on the edges are ones which exist as possible worlds for both agents.

Theorem 7 (Converging in knowledge $\Rightarrow$ agents converge to the same action) *Let* $\Sigma = (A, S, C, N, W)$ *be a system, and* B *a behavioral rule. If* $\exists t \in T$ *s.t.* $(\forall i \in A) E(k_i^{t+1}(w)) = E(k_i^t(w))$*, then* $\forall(w \in W, i, j \in A) F_i^t(B, w) = F_j^t(B, w)$

Proof *We construct a "shared knowledge graph" of time* t *between both agents* $(i, j) \in N$ *as follows. For each agent, we construct a node for each cluster in its indistinguishable worlds relation. That is, the shared knowledge graph* $SKG_{(i,j)}^t = (V, E)$*, where each node* $v \in V$ *is a cluster from* $IND(i, t)$ *or* $IND(j, t)$ *(which is also an equivalence class as shown above). Let us denote nodes* v, u *as the following clusters,* $v = Ind_v(i, t), u = Ind_u(j, t)$*. Every edge* $e \in E$ *is* *constructed between two nodes $v, u \in V$ if and only if the clusters of these nodes are* not *mutually exclusive. Formally, $\{u, v\} \in E \Leftrightarrow Ind_v(i,t) \cap Ind_u(j,t) \neq \phi$.*

An example of a shared knowledge is presented in Figure 9. Next, we present a lemma that shows that in a shared knowledge graph, at any time t, the neighbors of a node will provide the same behavior at $t-1$.

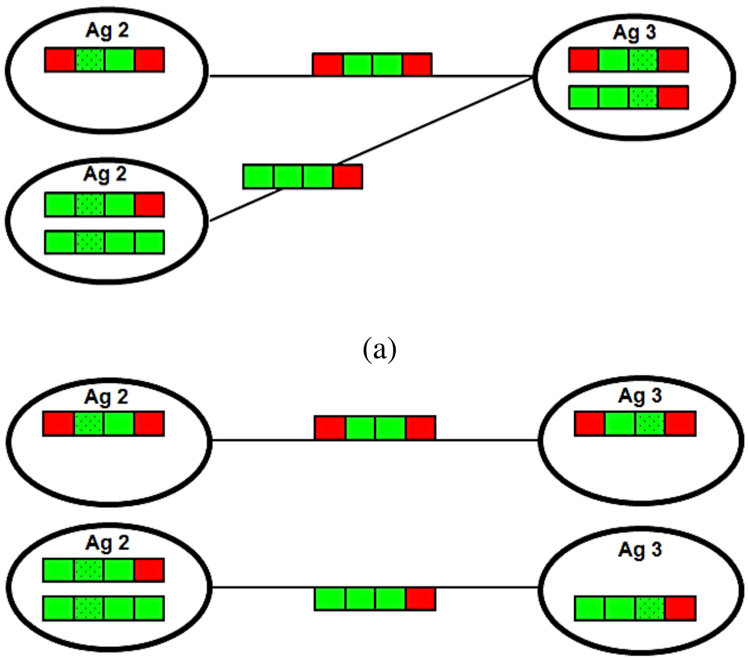

(a)

(b)

Figure 10: The figure shows a snapshot of the Shared Knowledge Graph between agents 2, and 3 (middle agents), in two consecutive time steps. In figure (a) we can see that Agent 3 cannot distinguish between the two worlds in his cluster. We can also see that Agent 2 has 2 clusters: in the upper one he will vote Red (there is a tie and Red is the default action), while in the lower cluster he will vote Green in both worlds. In figure (b), after observing agent 2's signal, agent 3 can now distinguish between the worlds, and the cluster divides to two clusters.

Lemma 8 *Given a shared knowledge graph, $SKG^t_{(i,j)}$. The neighboring nodes of every node will* *behave the same at the previous time step. Formally,* $\forall(v = Ind_v(i,t) \in SKG^t_{(i,j)}, \forall u, q \in SKG^t_{(i,j)}, \{u,v\}, \{q,v\} \in E), F_i^{t-1}(B,u) = F_i^{t-1}(B,q)$

Proof *By definition of a shared knowledge graph,* $SKG^t_{(i,j)}$*, at time* t *the worlds in node* v *are indistinguishable by agent* i*. If at time* $t-1$ *agent* i *could have distinguish between the behavior of its current neighboring nodes, he would have divide his own cluster and they would not be neighbors anymore.*

In Figure 10 we can see the cluster of agent 3 divides are learning the signal of agent 2. The following corollary, derived directly from Lemma 8, states that if the system had converged in knowledge, we will see the same behavior of the neighboring nodes in the shared knowledge graph. This is due to the fact that the graph remains static after convergence.

Corollary 9 *Given a shared knowledge graph,* $SKG^t_{(i,j)}$*. If the system had converged in knowledge in round* T*, the neighboring nodes of every node will behave the same at all time steps after* T*.*

We continue by considering a single connected component of the shared knowledge graph. Denote C_i *to be the set of all worlds in all nodes belonging to agent* i*, and* C_j *to be the set of all worlds in all nodes belonging to agent* j*. We now show that* $C_i = C_j$*.*

Lemma 10 *Given a shared knowledge graph,* $SKG^t_{(i,j)}$*. After convergence in knowledge, in any connected component the union of nodes for each agent contains the same set of worlds. Formally,* $C_i = C_j$*.*

Proof *By definition of a shared knowledge graph, an edge exist only if there is a non empty intersection between two nodes (and each node belongs to each of the agents). As such, assume node* v *belongs to agent* i *and node* u *belongs to agent* j*. Given a world* w *found on the* $\{u,v\}$ *edge. When constructing the union of nodes for each agent, by definition of the the shared knowledge graph,* w *will be included in both. This will be true for each world found in* C_i *or* C_j *as the nodes are part of a connected graph.*

After showing that $C_i = C_j$, the behavior function will return the same behavior in each of the joint clusters. Moreover, according to the additivity constraint on the same behavior will also be taken in the union of all the clusters in C_i or C_j. Formally, $F_I^t(B, C_i)) = F_I^t(B, C_j))$. Finally, moving back from the shared knowledge graph to the graph of the agent themselves (recall, a shared knowledge graph is constructed only between two neighboring agents). Assuming a connected agent graph will allow us, by induction, to conclude that all agents will behave the same after convergence.

6 Conclusions

Our study focused on notions of convergence in action and in knowledge for networks of decision making agents. In our setting the information available to each agent is given as signals, or observations of its neighbors' behavior. Thus, information is being propagated and aggregated over time.

We started by presenting a formal theoretical model of our setting. We then defined convergence properties on that model and differentiated between convergence of knowledge, and action, and showed that the first entails the second. In our first theorem we showed that *any* system, regardless of its structure, will eventually converge in both knowledge and action.

We then proceeded to inquire about a stronger convergence criteria, uniform convergence, in which all agents in the system converge to the *same* action. And an even stronger one, correct uniform convergence, in which all agents in the systems converge to the *correct* action. As those criteria are strong, we started by exploring different network structures. We started by showing that different types of *line networks* converge to the correct action. We also verified our theoretic results with a ONetwork simulation software that was developed for that matter.

The next step was to move to more complex structures that describes certain scenarios in real life. We showed that clique-based network (with various properties) that have common members will uniformly converge to the correct action. Lastly, we prove a general claim that in any form of connected network of agents, after converging in knowledge, all agents will issue the same action (though not necessarily the correct one).

While our work is theoretic in nature, it bares several practical implications, as real life situation that resembles our model can be analyzed and solved in our framework.

References

[1] D. Acemoglu, M. Dahleh, I. Lobel, and A.E. Ozdaglar. Bayesian learning in social networks. *MIT LIDs Working Paper*, 2008.

[2] D Acemoglu, M A Dahleh, I Lobel, and A Ozdaglar. Bayesian Learning in Social Networks. *The Review of Economic Studies*, 78(4):1201–1236, October 2011.

[3] Daron Acemoglu and Asuman Ozdaglar. Opinion dynamics and learning in social networks. *Dynamic Games and Applications*, 1(1):3–49, 2011.

[4] Vivi Alatas, Abhijit Banerjee, Arun G. Chandrasekhar, Rema Hanna, and Benjamin A. Olken. Network structure and the aggregation of information: Theory and evidence from indonesia. Working Paper 18351, National Bureau of Economic Research, August 2012.

[5] Abhijit Banerjee, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson. The diffusion of microfinance. Working Paper 17743, National Bureau of Economic Research, January 2012.

[6] Abhijit Banerjee and Drew Fudenberg. Word-of-mouth learning. *Games and Economic Behavior*, 46(1):1–22, January 2004.

[7] Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. Learning from the behavior of others: Conformity, fads, and informational cascades. *The Journal of Economic Perspectives*, pages 151–170, 1998.

[8] L. Blume, D. Easley, J. Kleinberg, R. Kleinberg, and E. Tardos. Which networks are least susceptible to cascading failures? In *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*, pages 393–402, Oct 2011.

[9] J.-F. Chamberland and V.V. Veeravalli. Decentralized detection in sensor networks. *IEEE Transactions on Signal Processing*, 51(2):407–416, 2003.

[10] Arun G. Chandrasekhar, Horacio Larreguy, Juan, and Pablo Xandri. Testing models of social learning on networks: Evidence from a framed field experiment. Technical report, MIT, 2012.

[11] Morris H DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.

[12] Peter DeMarzo, Dimitri Vayanos, and Jeffrey Zwiebel. Persuasion bias, social influence, and unidimensional opinions. *The Quarterly Journal of Economics*, 118(3):909–968, 2003.

[13] Peter Sheridan Dodds and Duncan J. Watts. Universal behavior in a generalized model of contagion. *Phys. Rev. Lett.*, 92:218701, May 2004.

[14] Y. Gal, R. Kasturirangan, A. Pfeffer, and W. Richards. A Model of Tacit Knowledge and Action. In *AAAI Fall Symposium*, 2008.

[15] Robin Glinton, Paul Scerri, and Katia Sycara. Exploiting scale invariant dynamics for efficient information propagation in large teams. In *AAMAS*, 2010.

[16] J.Y. Halpern and Y. Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial intelligence*, 54(3):319–379, 1992.

[17] Ali Jadbabaie, Pooya Molavi, and Alireza Tahbaz-Salehi. Information heterogeneity and the speed of learning in social networks. *Columbia Business School Research Paper*, (13-28), 2013.

[18] M. Kearns, S. Suri, and N. Montfort. An experimental study of the coloring problem on human subject networks. *Science*, 313(5788):824, 2006.

[19] David Kempe, Jon M. Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In Lise Getoor, Ted E. Senator, Pedro M. Domingos, and Christos Faloutsos, editors, *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, August 24 - 27, 2003*, pages 137–146. ACM, 2003.

[20] Manuel Mueller-Frank. A general framework for rational learning in social networks. *Theoretical Economics*, 8(1):1–40, January 2013.

[21] E. Pacuit and S. Salame. Majority logic. In *Principles of Knowledge Representation and Reasoning, KR*, volume 4, pages 598–605, 2004.

[22] M.J. Raddick, G. Bracey, P.L. Gay, C.J. Lintott, P. Murray, K. Schawinski, A.S. Szalay, and J. Vandenberg. Galaxy Zoo: exploring the motivations of citizen science volunteers. *Astronomy Education Review*, (9), 2010.

[23] L. Smith and P. Sørensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, 2000.

[24] L. Von Ahn and L. Dabbish. Designing games with a purpose. *Communications of the ACM*, 51(8):58–67, 2008.