Norms Propagation and their Reinforcement

Francesca Lipari,[1,2,*] Leonhard Horstmeyer,[3,†] Alireza Goudarzi,[4,‡] and Brais Alvarez[5,§]

[1]*Department of Economics, University of Tor Vergata, Rome, Italy*

[2]*Philosophy, Politics and Economics Program,*

University of Pennsylvania, Philadelphia, PA 19104

[3]*Max Planck Institute for Mathematics in the Sciences, 04103 Leipzig, Germany*

[4]*Department of Computer Science, University of New Mexico, Albuquerque, NM 87131*

[5]*Department of Economics, European University Institute, Florence, Italy*

(Dated: September 16, 2014)

Abstract

It is well known that social norms and peer pressure have a big influence on individual decisions and behavior. How do they change and evolve? How can one affect or modify them? And how can an individual break out of the role imposed by them? In this paper we develop two dynamical theoretical models to study these and other related questions. We model a collection of agents forming beliefs about the optimality of different possible choices or states, which will determine their identity. Our first model is based on excitable linear systems. Each agent updates its disposition for each of the different possible states, as a weighted average of the personal beliefs, its neighbors' beliefs, and the influence of an external source of information. In the second model agents follow the same principle, but in this case they observe the actions taken by the whole population, and they modify their dispositions following a Bayesian update. In both models agents' beliefs about the optimality of the different states will determine if they keep their current state or if they move to a new one, changing their identity. These models allow us to study how fast social norms emerge or change, as a function of the strength of the external information and the divergence between agents' own preferences and the established social norms.

* flipari@sas.upenn.edu

† horstmey@mis.mpg.de

‡ alirezag@cs.unm.edu

§ brais.alvarez@eui.eu

I. INTRODUCTION

An individual's image of *self* comprises *personal* and *social* selves. The personal self is a collection of *personal identities*, each of which arises from the individual's group membership. The social self, a collection of *social identities*, on the other hand, consists of individual's perceived membership in various social groups. Each social identity is perceived, by the individual who identifies with it, through a collection of shared values that defines the "us." The values shared within a group impose norms on its members', which influence their behavior and their decision making. What if the imposed norms are harmful to some members, for instance the gender-based repression of decision making? Can an external entity change the group values and norms, or reduce their effects on the group members? We developed two distinct models for social learning and consensus to study this question.

Here, we combine the social identitiy theory and social learning theory to derive two frameworks to model an interacting population of agents under the influence of exeternal information. The agents are bestowed with identities and dispositions to prefer some of the identities over the others. Through interaction with other agents in the population and the external information - e.g public policy campaign- the agents calculate new dispositions and chose their identites according the updated disposition. The first framework builds on the theory of excitable linear systems, in which the agents are nodes with a vector of dispositions. The vector is updated as a linear average of the dispositions of the agents neighbors and the external information. The second framework builds on the theory of Bayesian learning in which the agents start from a prior belief about various dispositions and updates the belief by looking at other agents' actions and estimating a social belief. Combined with some external media, the agents then form a new belief for various identity traits and choose the action leading towards the new identity accordingly. We study the convergence of dispositions and actions in the population as function of two paramters: the trust in the external information λ, and the agents' desire to change their identities τ.

Our preliminary results show that in the presense of external media, the agent's dispositions stablize very quickly and that the speed of convergence is proportional to the strenght of trust on the exeternal media. We find that this result hold for agents with regular grid-like interaction as well as small-world interactions on a preferential attachment network. The linear model shows the convergence in agent's states as τ and λ are reduced. We also found in the linear model the convergence of disposition increases as the strength of external information increases, but, as expected, shows no sensitivity to τ. On the other hand, in the Bayesian model we find that as τ increases we need larger values of λ for the agents to take into account the external information. We will discuss several modifications and extensions to our models, such as heterogeneity in λ and dynamical interaction patterns between agents, which takes into account the similarities of agents' beliefs.

II. RELATED WORK

The present work is connected to several different but intertwined strands of literature. First of all it is directly related to the Social Identity Theory. Social identity refers to that part of an individual's self-concept which derives from his knowledge of belonging to a social group, together with the value and emotional significance attached to that membership [1](Tajfel 1981, p. 255). If we consider the *Self* as collection of self-aspects, each of which is composed of self-relevant attributes that serve significant needs, goals and motives, then the relevance of this membership is determined by the strength of the norms that will shape the potential role of the individual. In other terms, if we denote this sense of belonging as *social identity* - namely the individual's perception of her original social categories - then individuals care about their own self-image, and about the several self-concepts associated to it, according to the one instilled in their society, or circle, that ultimately might affect her preferences and behavior[2–4]. But what if this membership, this sense of belonging, imposes on individuals a harmful norm or expected behavior, like for example gender norms repressing decision making of women throughout the world? Would conformity to this norm be still desirable? Could a policy maker intervene and change the norm, reducing or breaking its impact on the population? These are the open questions we try to address in this paper.

Related to the former, we consider the social norms literature. Social norms are coordination devices to decrease diversity or conflict. They work through empirical and normative expectations on individuals behaviors, implicitly stating what is socially acceptable and expected; thus they constrain individual behavior [2] and decision making.

The third strand of literature is related to the extensive theoretical literature on social learning. For what we have mentioned above, agents are exposed to social learning as receivers but also as senders of information about what is socially acceptable. This literature can be generally divided into two main streams: one that focuses on Bayesian learning and the other that deals with myopic or bounded rational learning. The literature on Bayesian learning originates from the contributions of [5] and [6], who assume an exogenous sequential structure in which each agent, after observing all past actions, optimally updates her belief about an unknown payoff-relevant state of the world and makes a single irreversible choice accordingly. Subsequent papers by [7–11] consider situations where individuals observe only a subset of past actions. These studies differ from the present one in two aspects: first, agents act sequentially and each individual has only one decision node; second, they focus on characterizing the asymptotic properties of different social networks under Bayesian learning. In particular, they study whether individuals converge to payoff-maximizing actions as the size of the social network becomes arbitrarily large. In this context, as the action space is discrete while the signal space is continuous, optimality of convergence is non-trivial under Bayesian learning.

Among the papers belonging to the non-Bayesian learning branch of the literature, those most closely related to ours are [12–15]. These papers study social learning in connected social networks. DeMarzo et al. [14] and Golub and Jackson [15] are particularly relevant for our work as they focus on the properties of consensus beliefs, in settings where individuals converge to the same opinions. These studies provide a characterization of social influence and analyze the likelihood that consensus beliefs will lead to optimal aggregation of information.

III. MODEL

A. General description

The model propose a network of agents who are endowed with individual beliefs or dispositions, $D_i(t)$, about their identity in a social context. Following the social psychology literature on multiple identity traits, each agent $i \in \{1, \ldots, N\}$ possesses a total identity ω_t^i that may change with time t. The traits of identity are multidimensional and social based, in the sense that they are based on the societal belief system that appends on a specific meaning and weight to traits in order to impress on people the notion that individuals in different groups belong to different categories. Members are subjected to different constraints and compel to different behavioral rules. For modeling purposes the total identity will be represented by the presence (=1) or absence (=0) of certain attributes, e.g being a woman, being a mother, having a university degree, etc. Among those attributes the first entry in the vectors includes traits beyond the individual responsibility e.g. race, gender, family background; the second and third entries are represented by attribute for which the DM is fully responsible and it is represented in function of effort disposable or actions taken. An example of this trait could be being an educated girl, coming from a poor family, or highly-gendered society. In our simulations, we choose $K = 2$ attributes to determine the identity of the agent i, which is then given by a K-dimensional binary vector:

$$\omega_t^i = (1, 0, 0) \qquad \sim \qquad \text{(woman=yes, mother=no, university degree=no)}$$

Hence a total of 2^K possible different identities are considered. Formally the identity ω is a vertex in a K-dimensional binary cube Q^K, as depicted in figure 1.

Those attributes are the states of the world relevant to the payoff and the disposition functions. According to the social and cultural characterization of the population, some attributes are more socially acceptable than others, and this assumption limits the possibility of each agent to move freely from one identity to another. We define this transitions as social mobility. Since one cannot change all the attributes within one time step, we constrain the admissible changes in the model in terms of the steps forward that each agent can take. A change is admissible if it differs in at most one attribute, or more formally if its Hamming distance from the current identity is at most one:

$$H(\omega_t^n, \omega_{t+1}^n) \leq 1 \ .$$

Then, in each period each agent has $K + 1$ admissible changes of identity including the case of no change at all. The disposition, or the belief, of agent i at time t over a specific identity is modelled as a probability distribution $\mathbb{D}_t^i$ over $K + 1$ options. Formally it is a probability

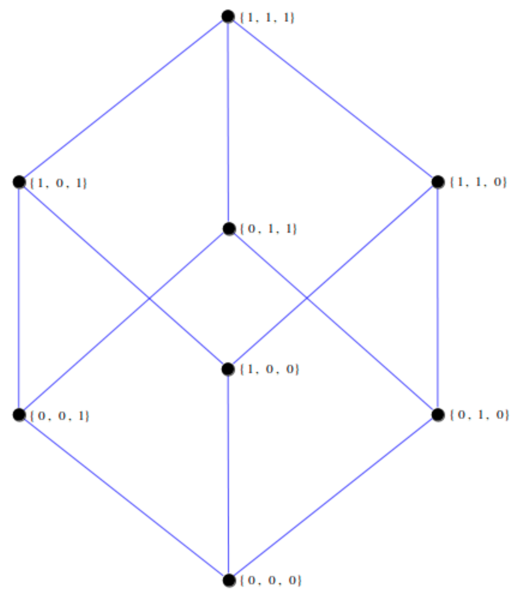

FIG. 1. The Boolean cube

measure $\mathbb{D}_t^i : \omega_t^i \times Q^K \longrightarrow [0,1]$. It can also be viewed as a point in the K-simplex:

$$\begin{aligned} 0 \leq & \mathbb{D}_t^i(\omega_t^i, \omega) \leq 1 \\ 0 = & \mathbb{D}_t^i(\omega_t^i, \omega) && \text{if} \quad H(\omega_t^i, \omega) > 1 \\ 1 = & \sum_{\omega \in Q^K} \mathbb{D}_t^i(\omega_t^i, \omega) \end{aligned}$$

Before choosing to take a given transition, each agent will learn what attribute is socially acceptable by gathering information from her network, and from an external signal coming from the policy maker.

Depending on the model, the piece of information coming from the network is simply given by the influence of different people's attributes, actions, or opinions on the agent. Namely, by seeing what other people do or say, the agent infers the social empirical expectations enforced on each attribute[2]. The higher the number of people in a particular position or a specific attribute, the higher is the evaluation of social acceptance of that position or attribute. These are updated from time to time through the action that each agent decides to take after having weighted the social evaluation with the external influence and her own disposition on the matter. Then we have

$$n_{st}^R = \frac{n_{st}^I}{N} + \frac{\sum n_{s_t}^c - \sum n_{s_t}^l}{N} \ , \tag{1}$$

where for a given period t, n_{st}^R is the revised number of people in state s, n_{st}^I is the initial number, n^c_{st} is the number of newcomers in the state s and finally n^l_{st} are those leaving the state. Such a social norm hypothesis is a probability distribution $\mathbb{N}^i_t : \omega^i_t \times Q^K \longrightarrow [0,1]$, according to which the agent forms beliefs about social mobility.

The external piece of information represents all the policies and campaigns proposed by a "soft-paternalistic" State, whose aim is to raise awareness in agents' choice. Such awareness tool could be represented by legislative interventions, media campaign, educational campaign. External influence is modelled by a stationary distributions $\mathbb{E} : Q^K \times Q^K \longrightarrow [0,1]$ over all admissible identity changes.

In order to understand how all the single parts of the model interact among each other we show the evolution of dispositions over time in figure 2:

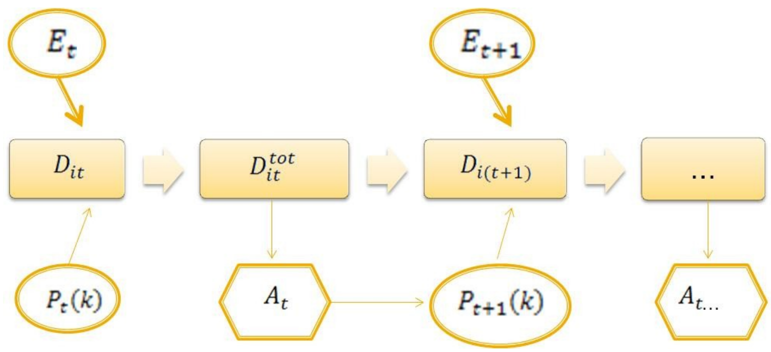

FIG. 2. The flow diagram. $P_t(k)$ represents the number of people being in vertex k at time t

Worthy to note here, is the possibility for each agent to have an individual preference over social mobility that in general might not coincide with the other two pieces of information. In a nutshell, by allowing social expectations[1] over each attribute to influence an agent's "chosen" identity, we are imposing the presence of social norms. The presence of such a norm does not imply that it is always true, and accordingly a particular specification of people beliefs on such norms might be held in an unaware manner with respect to other attributes, and this unawareness can be reinforced within the network. We impose that agents will act under the "veil of ignorance".

The tension between what would be preferred by the agent and what is socially acceptable makes the identity attribute the subject of a learning process, which will determine the agent's action. This act of gathering information will define the agent's disposition $\mathbb{D}^{tot}_i(t)$ towards each possible transition, i.e. she or he allocates preferences over changes of identity from ω^n_t to $\omega^n_{t+1} \in Q^K$. The propensity towards each attribute of the identity, and the willingness to act upon it accordingly is a weighted average of social information $(\sum_j^N W^{ij} D_j(t))$, which includes the own individual perception $W^{ii} D_i(t)$, and the external source of information $(\mathbb{E}(t))$, each receiving a weight of λ_i and $1 - \lambda_i$, respectively. λ_i is a parameter identifying the agent's trust in the external information as a tool to discover the best identity traits.

[1] That for each model takes a different mechanism.

$$\mathbb{D}_i^{tot}(t) = \mathbb{F}\left(\lambda, \mathbb{D}_j, \mathbb{E}(t)\right). \tag{2}$$

Common to both model is the action strategy: compliance decision depends both on the level and direction of social orientation, and on the conviction of agent i about moving to identity s, represented by the threshold τ_{is}[2]. The agent's conviction is an expression of his sense of autonomy[16, 17], and it introduces heterogeneity in the model. It is what we call henceforth social mobility parameter. It is a common practice to assume heterogeneity in learning at the level of the structure of the network or in the information structure within the network. Here we keep the structural heterogeneity but we also introduce the assumption of diversity in preferences(Page, 2007)[18].

We study the problem using two different network structures, in order to study the impact of different connectivites on the results: a 4-regural lattice grid and a preferential attachment model. The followings are common to both. The social network with N agents is modeled by a weighted directed Graph $G = (V, E)$. Each agent is a vertex in V. The weight that agent j places on agent i is $W^{ji} \in [0; 1]$, where W is a non-negative, asymmetric and row-stochastic interaction matrix. Note that since the graph is directed, W^{ji} needs not be the same in general as W^{ij}. A zero in the entry W^{ji} corresponds to the absence of any weight from j to i. The weights represent the trust that each individual places on the different agents.

Finally, we build two different updating mechanism: a myopic, or rule of thumb model, and a Bayesian one. The two differ in the updating mechanism of the agents' dispositions, but the difference is also more substantial.

B. Myopic Model

1. Model Description

We base our work on a traditional model of network influence, the Degroot model[19]. Like in the previous model, and as stated before, the social structure is described by a weighted and possibly directed network. In DeGroot's model agents have beliefs about the probability of some event. At each date, agents communicate with their neighbors in the social network and update their beliefs. The updating process is simple. An agent's new belief is the (weighted) average of his or her neighbors' beliefs from the previous period. Over time, provided the network is strongly connected (so there is a directed path from any agent to any other) and satisfies a weak aperiodicity condition, beliefs converge to a consensus. In our model the belifies are represented by the dispositions towards a particular state of the world implying to take an action to move to some better position on the cube. Dispositions in this scenario represent beliefs on the quality of certain identity traits over the others. At each date, agents communicate with their neighbours to know their dispositions

[2] In the Bayesian model for reasons of tractability of the calculation we imposed that this parameter is not individually based, but rather independent of i and s.

towards the state of the world and at the same time receive external information on the same issue, and then the total disposition is updated. The updating is computed for each position on the cube in order to find the one , e.g. s, that will provide the highest value in terms of disposition. Once it is computed, if such maximum value will be higher than the agent i conviction about moving from the current identity τ_{is}, then the action will be taken. And so on as it is displayed in Figure 2.

Hence, the main equations are:

$$D_i^{tot}(t) = (1-\lambda_i)\sum_j^N W_{i,j} D_j(t) + \lambda_i E(t), \tag{3}$$

$$D_i(t+1) = D_i^{tot}(t), \tag{4}$$

$$A(t) = F(D^{tot}(t)), \tag{5}$$

where $D_i(t)$ is the population disposition matrix whose rows i correspond to each individual and columns d correspond to different dispositions.

Now we find d_i^{max}, the disposition with the highest value for individual i among all the possible positions:

$$d_i^{max} = \arg\max_d D_{i,d}^{tot}(t), \tag{6}$$

and the action is chosen according to:

$$A_i(t) = \begin{cases} A_i(t-1) & , \text{if} D_{i,d_i^{max}}^{tot}(t) < \tau_{is}, \\ d_i^{max} & , \text{if} D_{i,d_i^{max}}^{tot}(t) \geq \tau_{is}. \end{cases} \tag{7}$$

As already mentioned the parameter τ_{is} signals diversity in preferences that each agent has towards a particular state s. Hence, whenever the total disposition is less than the threshold, the agent will not take any action and will stay in the previous position. If the total disposition overcomes τ_{is} then the agent will move towards the position that has the highest disposition value.

C. Bayesian version

1. *Description of the model without weights*

At any time t an agent i observes N_t^i, a vector of occupancies $(n_1^i, \ldots, n_m^i)$ of the states 1 through $m = 2^K$ in the neighborhood of node i. They also form beliefs about the payoff function for a certain identity. This payoff function is a random variable Π_t^i for each period t and agent i. It can take, e.g. the value $\pi = (\pi_1, \ldots, \pi_m)$. So for each individual, each identity gets assigned a payoff. Agents have beliefs about this that are modelled as probability distributions.

Suppose an agent beliefs the payoffs are $\pi = (\pi_1, \ldots, \pi_m)$ and he observes all the people in the network, then what would he expect the probability for n_s number of people in state s to be? We construct this conditional belief as follows:

For any person to go to state s is more likely the higher the payoff for that state. Therefore, this probability is increasing on π_s. However, since π_s can take values in $\{-1, 0, 1\}$ and in particular it can take the value -1, we want to find a function that maps also negative payoffs to positive albeit small probabilities. Choosing the exponential we have:

$$p(s|\vec{\pi}) \sim e^{W\pi_s} \tag{8}$$

where the parameter W regulates how strongly people are affected by the reward associated to each of the possible choices. If $W = 0$ then all choices are equally likely irrespective of their payoff. We now just need to sneak in a normalization constant:

$$p(s|\vec{\pi}) = \frac{e^{W\pi_s}}{\sum_{s'} e^{W\pi'_s}} \ . \tag{9}$$

Technically this normalizing sum has to be only over the attainable states, but since we don't know yet the states of the individuals, we think this is a reasonable approximation for the Bayesian model.

Now, ommitting the agency dependence, the probability to have the output $\vec{n} = (n_1, \ldots, n_m)$ given that we have a payoff $\vec{\pi} = (\pi_1, \ldots, \pi_m)$ is

$$\mathbb{P}(N_t = \vec{n}|\Pi_t = \vec{\pi}) = \binom{n}{n_1 \cdots n_m} p(1|\vec{\pi})^{n_1} \cdots p(m|\vec{\pi})^{n_m} \tag{10}$$

where $n = ||\vec{n}||_1 = \sum_s n_s$ and $\binom{n}{n_1 \cdots n_m}$ are the multinomial coefficients. Written out explicitly[3]:

$$\begin{aligned} \mathbb{P}(N_t = \vec{n}|\Pi_t = \vec{\pi}) &= \frac{1}{(\sum_{s'} e^{\beta\pi'_s})^n} \binom{n}{n_1 \cdots n_m} e^{\beta\pi_1 n_1} \cdots e^{\beta\pi_m n_m}, &(11)\\ &= \binom{n}{n_1 \cdots n_m} \frac{e^{W\langle\vec{\pi},\vec{n}\rangle}}{(\sum_{s'} e^{\beta\pi'_s})^n}, &(12) \end{aligned}$$

where $\langle\vec{\pi}, \vec{n}\rangle$ is the dot product. We can also calculate the probability for the event $N_t = \vec{n}$, which is the sum over all conditional probabilities times the probability for the event $\Pi_t = \vec{\pi}$:

$$\begin{aligned} \mathbb{P}(N_t = \vec{n}) &= \sum_{\vec{\pi}} \mathbb{P}(N_t = \vec{n}|\Pi_t = \vec{\pi}) \cdot \mathbb{P}(\Pi_t = \vec{\pi}), &(13)\\ &= \sum_{\vec{\pi}} \binom{n}{n_1 \cdots n_m} \frac{e^{W\langle\vec{\pi},\vec{n}\rangle}}{(\sum_{s'} e^{W\pi'_s})^n} \mathbb{P}(\Pi_t = \vec{\pi}). &(14) \end{aligned}$$

[3] where is β comes from? what is the meaning?

These things can be put together to form the Bayesian updating rule:

$$\mathbb{P}(\Pi_{t+1} = \vec{\pi}) = \mathbb{P}(\Pi_t = \vec{\pi} | N_t = \vec{n}) = \frac{\mathbb{P}(N_t = \vec{n} | \Pi_t = \vec{\pi})}{\mathbb{P}(N_t = \vec{n})} \mathbb{P}(\Pi_t = \vec{\pi}). \tag{15}$$

On the very right is the prior hypothesis that the payoff vector is $\vec{\pi}$. The left hand side is the posterior distribution for the belief about the payoff vector being in state $\vec{\pi}$. The only thing every agent has to set for himself is an initial prior.

One can also look at the marginal probabilities for the payoffs π_s of a particular identity s, which can take values $\{-1, 0, 1\}$. We can call this marginal probability $p_s^i(\pi_s)$. For instance

$$p_s^i(\pi_s = -1) := \mathbb{P}\Big(\big(\Pi_t^i\big)_s = -1\Big) = \sum_{\pi : \pi_s = -1} \mathbb{P}\Big(\Pi_t^i = \pi\Big). \tag{16}$$

Agents can change their identity, namely their position on the cube. Each agent i is in some identity state ω_t^i and has the option to stay where it is or to change to an adjacent identity s. In formal terms it can choose amongst any states with $H(\omega_t^i, s) \leq 1$. We introduce the threshold parameter $\tau_i \in]0, 1]$, which determines whether the action random variable $A_{s,t}^i$ is zero or one. The default value is zero, unless:

$$\mathbb{E}\big[\pi_s\big] - \mathbb{E}\big[\pi_{\omega_t^i}\big] \geq \tau \qquad \text{and} \quad H(\omega_t^i, s) = 1 \qquad \text{then} \quad A_{s,t}^i = 1, \tag{17}$$

where $\mathbb{E}\big[\pi_s\big] = p_s^i(1) - p_s^i(-1)$ is the expectation value.

$$\omega_{t+1}^i = \begin{cases} \omega_t^i \ \text{ if } \ A_{s,t}^i = 0 \ \ \forall s \\ s' \ \text{ if } \ A_{s',t}^i = 1 \end{cases} \tag{18}$$

In other words, if $A_{s,t}^i = 1$ then at time $t + 1$ there will be one more person in the state s and the process N_{t+1} will adapt accordingly.

An adaptation for this model would be to attribute to a social role s an effort cost that will account for the difficulty to get into s, otherwise high payoff states would attract all the people, even if there is a low saturation point. This could be included by defining $u_s \in [0, 1]$. A modified action condition would be

$$\mathbb{E}\big[\pi_s\big] - \mathbb{E}\big[\pi_{\omega_t^i}\big] \geq \tau \ ,^4 \tag{19}$$

but maybe another mechanism should be found to account for that. Notice also that expectations are linear in their arguments, but in this case we cannot just take the expectation of the difference. The reason is that they are expectations with respect to different distributions, which is hidden, i.e., not included in the notation.

Taken together the assignment for each agent is $(i, t) \mapsto \{\vec{\omega}_t^i, \Pi_t^i, N_t, \mathbb{A}_t^i\}$. The evolution of these quantities is as described above. N_t is a random variable in this model.

The media or policy influence is modelled as a belief process over the payoffs E_t with

[4] Indexed by i? –¿ Each person i has a state ω^i, so it's w.r.t. this state

probability distribution:

$$\mathbb{P}(E_t = \vec{\pi}) \in [0, 1]$$

The updated belief for the payoffs is a convex linear combination with weights λ and $(1-\lambda)$:

$$\mathbb{P}(\Pi_{t+1} = \vec{\pi}) = (1-\lambda)\mathbb{P}(\Pi_t = \vec{\pi}|N_t = \vec{n}) + \lambda\mathbb{P}(E_t = \vec{\pi}). \tag{20}$$

The case $\lambda = 0$ means that there is no external influence.

2. Description of the model including weights

The only social information that we have in the preceding weightless model is the occupation number of states for people that we know, but that are annonymous as far as the model is concerned.

There is a natural way how to include weights into the model. We have to go to a higher resolution than just the occupation numbers and look at each person in the neighborhood separately. We are guided by the idea that a person one does not care about at all should not matter in the updating. So agent i-th belief that agent j enters state s is similar to Equ. 9:

$$p_{ij}(s|\vec{\pi}) = \frac{e^{W_{ij}\pi_{\omega_j}}}{\sum_{s'} e^{W_{ij}\pi_{s'}}}.$$

Note, that a person that i does not care about at all, i.e. $W_{ij} = 0$ will have the same probability of going to any state s irrespectively of the payoff.

The conditional probability for agent i to observe $(\omega_t^1, \ldots, \omega_t^n)$, where n is the number of i-th neighbors, given a payoff allocation $(\vec{\pi})$ is:

$$\mathbb{P}(N_t^i = (\omega^1, \ldots, \omega^n)|\Pi_t = \vec{\pi}) = \prod_{j=1}^{n} p_{ij}(\omega_t^j|\vec{\pi}). \tag{21}$$

It can be easily seen, that this expression will reduce to the unweighted scenario of equation(10). We take $W_{ij} = A_{ij}W$ and sum over all those state configurations that give rise to the occupancy numbers $(n_1^i, \ldots, n_m^i)$. We call this set Ω and calculate :

$$
\begin{aligned}
\mathbb{P}(N_t^i = (n_1, \ldots, n_m) | \Pi_t = \vec{\pi}) &= \sum_{\omega \in \Omega} \mathbb{P}(N_t^i = (\omega^1, \ldots, \omega^n) | \Pi_t = \vec{\pi}), \\
&= \sum_{\omega \in \Omega} \prod_j \frac{e^{A_{ij} W \pi_{\omega_j}}}{\sum_{s'} e^{A_{ij} W \pi_{s'}}}, \\
&= \sum_{\omega \in \Omega} \prod_j \frac{e^{W \pi_{\omega_j}}}{\sum_{s'} e^{W \pi_{s'}}} \qquad \text{since } A_{ij} = 1 \text{ for neighbors}, \\
&= \sum_{\omega \in \Omega} \frac{e^{W \pi_{\omega_1}}}{\sum_{s'} e^{W \pi_{s'}}} \cdot \frac{e^{W \pi_{\omega_2}}}{\sum_{s'} e^{W \pi_{s'}}} \cdots \frac{e^{W \pi_{\omega_n}}}{\sum_{s'} e^{W \pi_{s'}}}, \\
&= \sum_{\omega \in \Omega} \left(\frac{e^{W \pi_1}}{\sum_{s'} e^{W \pi_{s'}}} \right)^{n_1} \cdot \left(\frac{e^{W \pi_2}}{\sum_{s'} e^{W \pi_{s'}}} \right)^{n_2} \cdots \left(\frac{e^{W \pi_m}}{\sum_{s'} e^{W \pi_{s'}}} \right)^{n_m}, \\
&= \binom{n}{n_1 \cdots n_m} \frac{e^{W \langle \vec{\pi}, \vec{n} \rangle}}{\left(\sum_{s'} e^{W \pi_{s'}} \right)^n}. \qquad (22)
\end{aligned}
$$

D. Experimental Setup

The evaluation of both models requires a systematic study of different parameters and initial conditions. In the preliminary study reported here we restrict the study as follows: we consider a network of $N = 100$ agents. In order to compare the implications of different network structures we choose two interaction patterns between the agents. First, a 4-regular lattice grid with fixed boundary condition; and second, the preferential attachment model[20] with connectivity parameter $k = 2$. This is created by a sequential model of network formation in which at each time step a node is added to the network, being connected to each of the k existing nodes with a probability proportional to their degree. To model the heterogeneity of influence between agents we assign weights on the interactions, drawn from an independent and identical uniform distribution on the interval $[0, 1]$. To model beliefs we assume two identity attributes resulting in four possible identity combination and four links to move from one identity to another. The weights on the links are also drawn from a i.i.d uniform distribution on the interval $[0, 1]$. Without lack of generality we normalize the rows of both the beliefs matrix and the interaction matrix so their columns add up to 1.

To simulate the system we set the belief of all agents identically to be dominant towards the state $(0, 0)$ and we fix the external signal to be a belief encouraging a move to $(1, 1)$. We randomly initialize agents state to occupy one of the four possible states with equal probability. We then simulate the system and study agents' beliefs and states as time goes by. We will study the effects of τ and λ on the convergence of states and beliefs on the two different networks.

IV. RESULTS

A. The Myopic Model

We are going to present some of the results of the model.

1. Entropy of beliefs and actions.

In this section we give a representation of the pluralism in a given society at the final state, namely once all the decisions have been made. The two graphs show how the probability of being in each state will change for each combination of the parameters λ and τ. In our simulations λ_t is the same for every agent in each period t; even if it might potentially change over time. τ_{is} is in principle different for every agent.

As τ increases, so does the entropy, and so the difficulty of changing identity increases. If τ is low, everyone goes wherever he prefers, if the parameter is very high, movements to a new position are very unlikely. Hence,in this set-up it will determine the level of heterogeneity across identities. These results are illustrated in Figure 3.

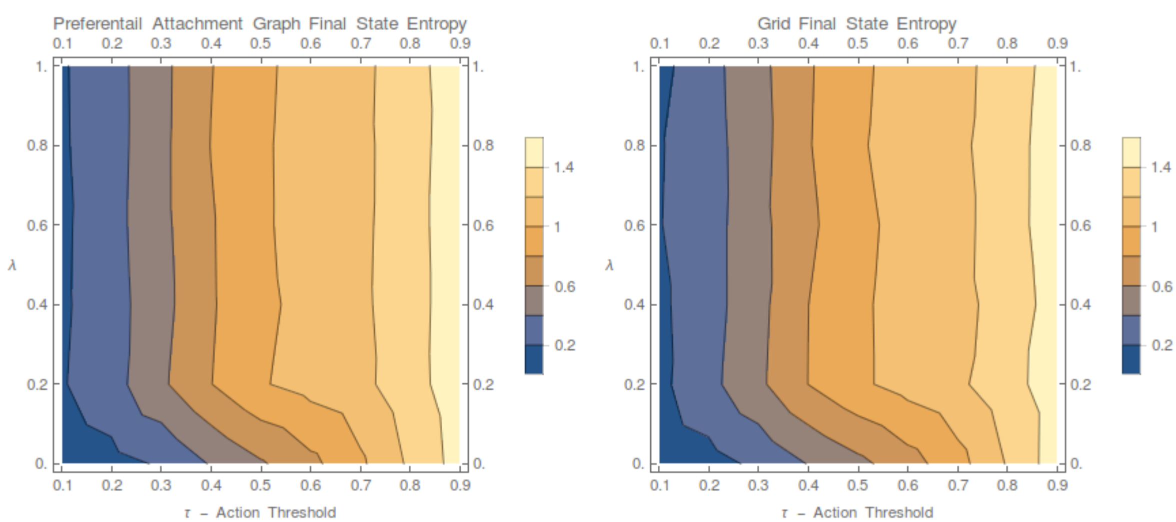

(a) Preferential Attachment.

(b) Grid.

FIG. 3. Final state entropy in the population after 50 time steps for both grid and preferrential attachment networks. For low τ, the agents can easily switch states and go towards the state prefered by average belief or the external signal and therefore we obtain low state entropy. As the τ increases he state entropy increases because the state initialization is random and agents will not be able to change state during the simulation.

2. Relaxation time.

Here we anlayze the *individual belief relaxation time*, i.e., average time for individuals to converge to a stable belief. We meaure this by calculating the mean-square error between the belief of individual i at time t and $t-1$:

$$\Delta mse_i(t) = \frac{1}{m}\sum_{l=1}^{m}(D_{i,l}(t-1) - D_{i,l}(t))^2, \tag{23}$$

where $D_{i,l}(t)$ is the value of the belief l for individual i at time t. We then average $\Delta mse_i(t)$ over the population of agents: $\langle \Delta mse_i(t)\rangle_{i=1}^{N}$. For stable systems $\langle \Delta mse_i(t)\rangle \to 0$ as $t \to \infty$ meaning the agent is converging to its stable belief. Figure 4 shows $\langle \Delta mse_i(t)\rangle$ in the population using both grid interaction and preferrential attachment interraction. The speed of convergence is proportional to λ. For $\lambda = 0$ the convegence occurs very slowly, and as λ increases above zero, individual agents will converge to their final beliefs more quickly. This behavior holds true for both the preferential attachment network and the grid network.

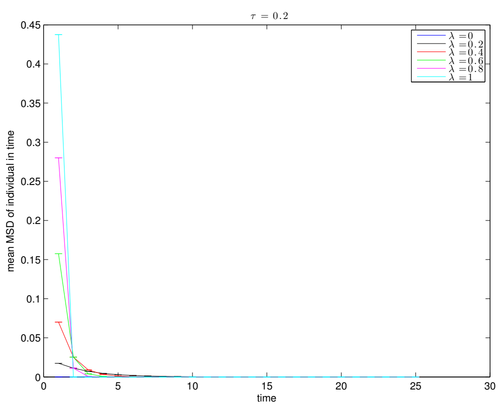

(a) Preferential Attachment.

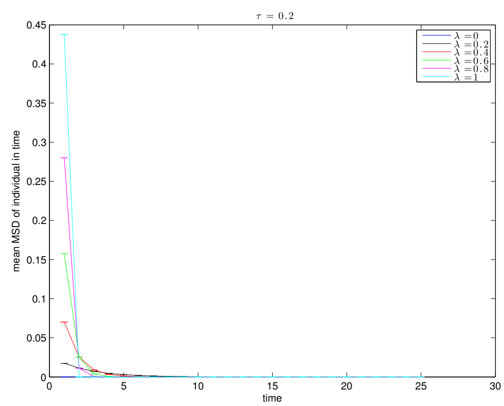

(b) Grid.

FIG. 4. Average individual belief convergence speed. We measure how much each agent changes its belief as the time progresses for different values of λ. Large values of $\lambda > 0.4$ causes very sharp converges for both the grid network and the preferential attachment network.

3. Convergence of population belief.

Next we anlayze the convergence of interacting agents with respect to their beliefs. For this we initial the system and run it for 50 time steps and measure the average local mean-squared error between agents beliefs as follows. Let V_i be the set of agents interacting with agent i. Then we calculate the average local mse as:

$$mse_i^{local}(t) = \frac{1}{|V_i|m} \sum_{j \in V_i} \sum_{l=1}^{m} (D_{i,l}(t) - D_{j,l}(t))^2. \tag{24}$$

We then calculate the average of this measure on the entire population. We see that result of this calculation for both the grid and the preferential attachment networks in Figure 5. The exponential convergence to $mse = 0$ with respect to time depends also on the value of λ that determines the speed of such convergence. When $\lambda = 0$ the convergence of the mse occurs fastly. The average mse increases as λ decreases showing that even if everyone gets the same information, each neighbor averages it differently resulting in unequal D_i.

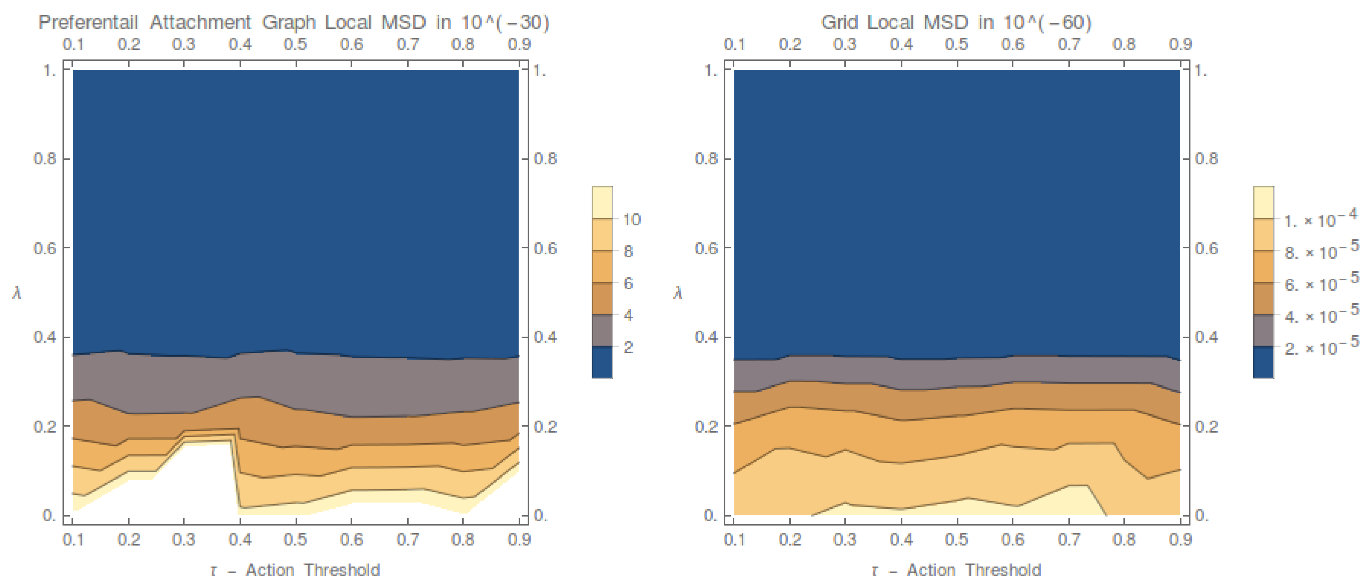

(a) Preferential Attachment.

(b) Grid.

FIG. 5. Average local mean-squared error after 50 time steps for both the grid and the preferrential attachment networks. For both networks the population becomes completely homogeneous for $\lambda > 0.4$, in the presence of significant external influence.

4. Analysis of Convergence

The dynamics of the myopic model (Section III B) follows a linear system model with input. At a descriptive level, these systems are less restrictive than the Markov chains that are typically used to model stochastic processes. Here, we invoke the mathematics of linear systems to anlayze the convergence criteria and the properties of a stable system. As a sufficient requirement we assume that the largest eigenvalue of the connection matrix has *spectral radius* $\gamma < 1$. To analyze the long term convergence properties of the the myopic model we extend the techniques that were initially developed in the context of computing with excitable dynamics[21].

For the rest of this section we assume the following. Let D be a real $N \times m$ matrix whose rows are the belief vectors of each agent and let W be a real $N \times N$ matrix of connectivity between the agents. Moreover, let λ be real in $[0, 1]$, $E(t)$ be a random real number drawn at each time step from identical and independent distributions, and $\bar{e}$ is the vector of ones of length N. Also without loss of generality we assume that $E(t)$ is drawn from zero-mean distribution with variance $\langle E^2 \rangle < \infty$. The time evolution of D is given by:

$$D(t) = (1 - \lambda)W \cdot D(t-1) + \lambda \bar{e} E(t-1) \tag{25}$$

Theorem 1. *The asymptotic state of the system given in Equation 25 depends only on the history of the external signal $E(t)$ as follows:*

$$D(t) = \lim_{t \to \infty} \sum_{i=0}^{t-1} \left((1-\lambda)W\right)^{t-i-1} \lambda \bar{e} E(i). \tag{26}$$

We can see that the effect of the prior belief in the population vanishes as time goes by and is replaced by the history of the external signal. Let us now have a closer look at the mean-squared error between the belief of agent i and agent j at time t, $mse_{i,j}(t)$. But notice that we would like to have an average of this measure over all possible histories of input signal.

Theorem 2. *Assume T identical copies of Equation 25 are running each with its own disdinct history of external signal $E(t)$. If the elements of the agent's belief vector are drawn from identical and independent distributions with mean $\langle d \rangle$ and variance δ, then the ensemble average of mean-squared error between two agents i and j is given by:*

$$\langle mse_{i,j} \rangle(t) = 2\delta. \tag{27}$$

Note that the independence conditions here is met because the elements of the belief vector evolve in time independently of one another. This result implies that the $mse_{i,j}$ only depends on the δ and so we can assume $m = 1$ and calculate the variance of belief at time t in T identical copies of the system running. In fact we can compute an entire covariance of the agents beliefs at time t and denote it $\Delta(t)$.

Theorem 3. *Let $m = 1$, $\langle E^2 \rangle$ the variance of external signal, and I the identity matrix of order N. Also let $W = UVU^{-1}$ the eigenvalue decomposition of W, $\bar{W} = U\bar{V}U^{-1}$ and $\bar{V} = (I - V)^{-1}$. The covariance of the beliefs of all agents with each other at time t represented by $\Delta(t)$ is given by:*

$$\Delta(t) = \langle D(t) D'(t) \rangle = \frac{\langle E^2 \rangle \lambda}{1 - \lambda} \bar{W} \bar{e} \bar{e}' \bar{W}' \tag{28}$$

This shows that the maximum heterogeneity arises in population when $\lambda \to 0$, when the external signal has no influence. As $\lambda \to 1$, the population becomes more and more homogeneous since the external influence dominate the belief dynamics.

B. Bayesian model

Here we analyze the Bayesian model with respect to it's convergence and the heterogeneity of beliefs under the influence of an external information source. We also consider the entorpy of the final state distribution and eventually the success of a policy implementation. Like for the Myopic model, we have used two differnt types of network: a 10×10 grid without periodic boundary conditions and a preferential attachment graph.

The prior beliefs for the payoffs of each state is taken to be independent:

$$\mathbb{P}(\Pi_0 = (\pi_1, \pi_2, \pi_3, \pi_4)) = \mathbb{P}_1(\pi_1)\mathbb{P}_2(\pi_2)\mathbb{P}_3(\pi_3)\mathbb{P}_4(\pi_4)$$

We start with a scenario in which agents belief that the lowest state $\{0,0\}$ has high payoff with probability one, i.e. : $\mathbb{P}_1(\pi_1 = 1) = 1$ and $\mathbb{P}_1(\pi_1 = 0) = \mathbb{P}_1(\pi_1 = -1) = 0$. Likewise we stipulate that agents belief in a low payoff for the highest state $\{1,1\}$ again with probability one, so $\mathbb{P}_4(\pi_4 = 1) = \mathbb{P}_4(\pi_4 = 0) = 0$, but $\mathbb{P}_4(\pi_4 = -1) = 1$. Finally we assume that agents are indifferent about the other social states $\{0,1\}$ and $\{1,0\}$, in other words they believe anything is equally likely: $\mathbb{P}_{2,3}(\pi_{2,3} = 1) = \mathbb{P}_{2,3}(\pi_{2,3} = 0) = \mathbb{P}_{2,3}(\pi_{2,3} = -1) = 1/3$.

On the other hand the external Influence favours the higher state and disfavours the lower, whilst being indifferent about the intermediate states. We also have the independence assumption:

$$\mathbb{P}(E = (\pi_1, \pi_2, \pi_3, \pi_4)) = \mathbb{Q}_1(\pi_1)\mathbb{Q}_2(\pi_2)\mathbb{Q}_3(\pi_3)\mathbb{Q}_4(\pi_4) \tag{29}$$

with $\mathbb{Q}_1(\pi_1 = -1) = 1$ and $\mathbb{Q}_4(\pi_4 = 1) = 1$ and the intermediate ones have beliefs 1/3 in all payoff cases. Figure 6 summarizes these initial data.

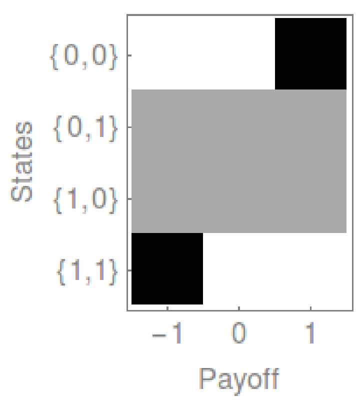

(a) Prior Distribution

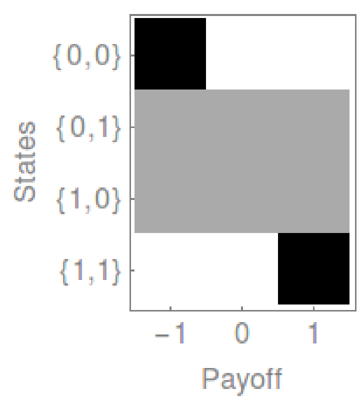

(b) External Incentive

FIG. 6. Depicted are the belief probabilities (black=1, white=0) for the payoffs of the respective states. The prior distribution (a) favours the state $\{0,0\}$ and disfavours $\{1,1\}$, whilst (b) favours $\{1,1\}$ and disfavours $\{0,0\}$.

People are initialized into one of four states $\{0,0\}, \{1,0\}, \{0,1\}$ or $\{1,1\}$ respectively with a uniform probability distribution. So approximately one quarter of the total population is in each state initially. A typical initialization of the 10×10 Grid can be seen in figure 7

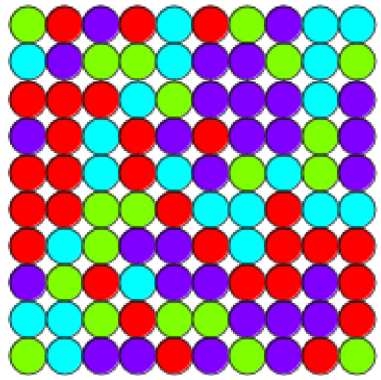

FIG. 7. Typical initialization of the 10×10 grid with states. Decoding: $\{0,0\}$ is red, $\{1,0\}$ is turquoise, $\{0,1\}$ is green and $\{1,1\}$ is purple.

1. *Convergence conditions and the speed of convergence.*

We observe that for general initial conditions on the states we see that the beliefs converge to a stable belief respectively, which of course may differ from one another. Figure 8 shows the results for the Grid and the Preferential Attachment Graph.

For most values of λ and τ the convergence happens very quickly in both networks. Just for values of very low λ there is a slow [5] convergence. This can readily be explained by the fact that there is no belief homogenizing bias that leads to faster equilibration. Similarly for small τ and values of λ around and below 0.3 there is a slow convergene. Since low values of τ signify that states can change quickly if the beliefs are in favour of a change, it is not easy for the beliefs to equilibrate as new states will result which in turn alternate the beliefs. These tendencies can be seen both in the grid (Fig. 8a) and the preferential attachment graph (Fig. 8b).

2. *Heterogeneity of beliefs*

Apart from the convergence of the individual beliefs, we looked at the heterogeneity. This was measured by the average local mean-squared error after 10 time steps [6]. The local mean-squared error measures the average squared Euclidean distance between the belief distributions of adjacent agents [7]. A high value means that there is a high diversity of beliefs represented and neighbors may in general have very different beliefs. A low value on the contrary means that people agree widely with their neighbours, although it is not clear what they agree on.

In figure 9 we see the average local mean squared error for both the grid and the preferential attachment graph. The highest belief diversity occurs for values of big τ and small λ. This observation can be accounted for by the fact that for those values there is also a high diversity of social states. People have a hard time to change their states since the threshold is so high and therefore they will roughly end up in their initial state configuration.

We can also observe another tendency. The higher τ, the threshold for action, the more external influence λ one needs to homogenize the beliefs. This is also in agreement with

[5] maybe include plot of convergence rate

[6] unfortunately different from Ali, because of the time consuming computation

[7] @ Ali, probably it needs to be weighted by the weighting of that node.

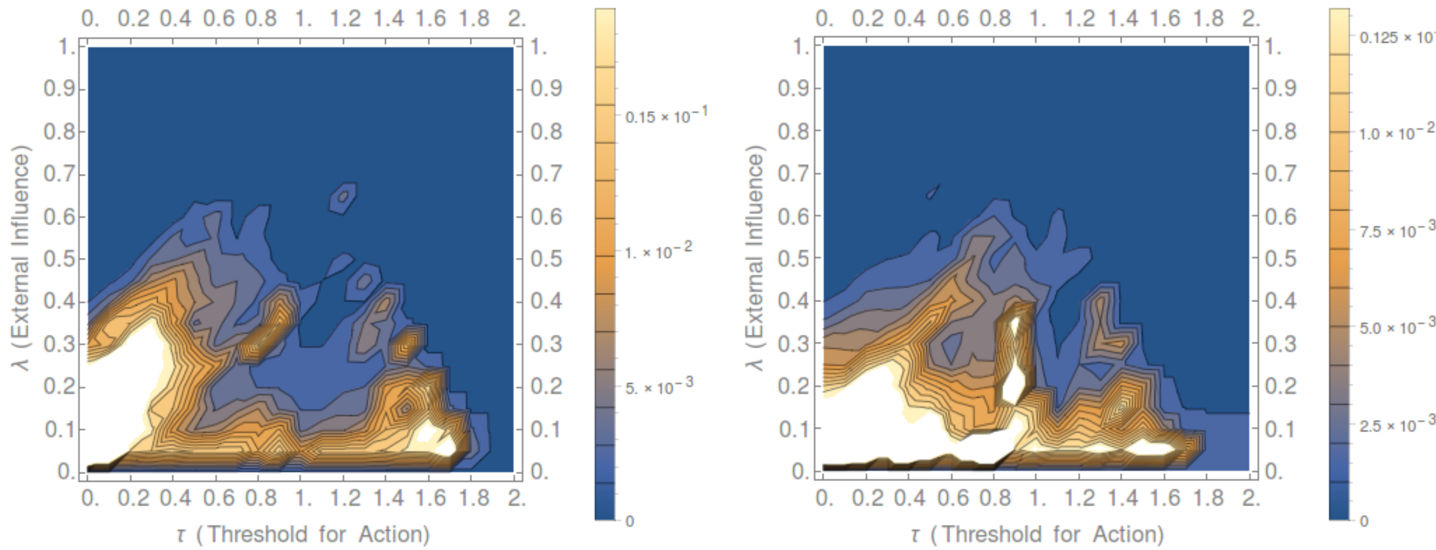

(a) Grid

(b) Preferential Attachment

FIG. 8. Average mean-squared error between successive belief distributions of the agents

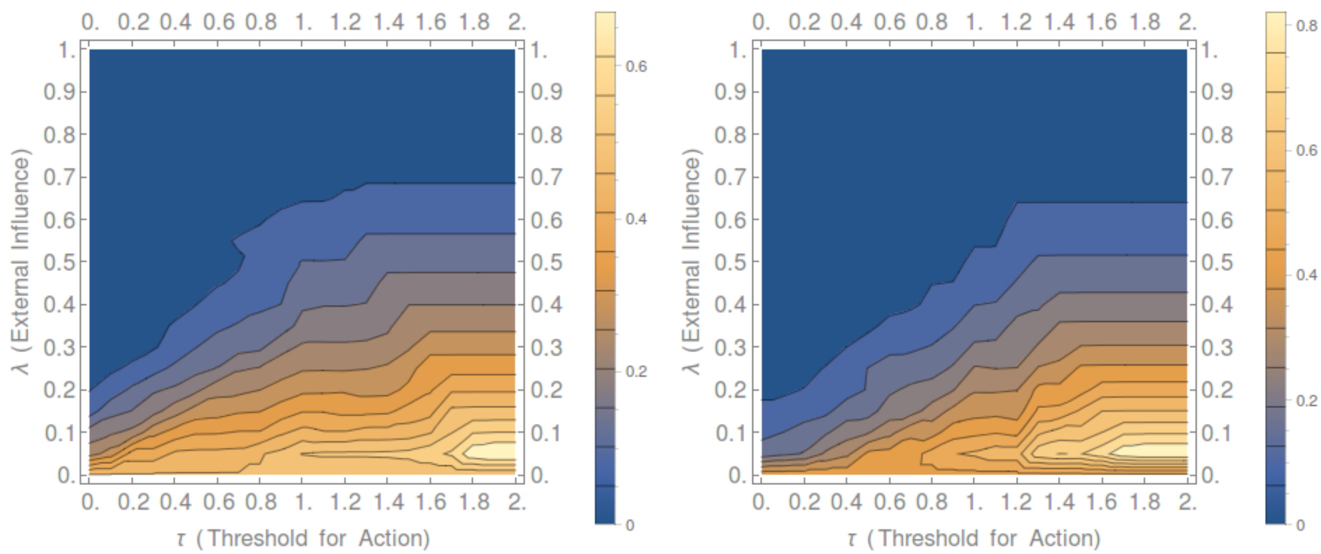

(a) Grid

(b) Preferential Attachment

FIG. 9. Average Local mean-squared distance of the agents belief probability distributions.

the idea that the external influence really needs to pull very hard if people are reluctant to change beliefs and states.

Finally we observer a similar behaviour for the grid (Fig. 9a) as well as the preferential attachment graph (Fig. 9b). The heterogeniyation behaviour seems to be not very sensitive on the type of network. This is also an important result for policy applications, because the exact network structure is often not very well known to the policy makers.

3. Entropy of final state distribution

To study the heterogeneity more in depth, we also looked at the uncertainty for agents to be in a particular final states as it is measured by the entropy of the final state distribution. This measure complements the local mean-squared error on the belief distribution, which only accounts for the agents' dispositions but not their states.

We observe two features in figure 10. Firstly we see that there is a clear cut distinction between state-heterogeneity and state-homogeneity. This delineation follows a slightly linear ascent: The higher the threshold the larger values for λ one needs to lower the entropy, i.e. to have a clear dominance of one particular state. This particular state will of course be the one that is endorsed by the external policy, but this information is not directly contained in the entropy. [8]

We also see that the behaviour of the grid is again very similar to the one of the preferential attachment graph. Therefore the final state entropy is not very sensitive on the type of network. In both cases we see that entropy is also coupled to the belief heterogeneity. The parameter region with low entropy is contained in the parameter region with low belief heterogeneity. If the entropy is low, i.e. the state is fairly certain, then also the beliefs are more homegeneous, i.e. people have similar beliefs. One reason for it is the fact that the absolute values of the entropy are different in the respective network scenarios. That might be due to the fact that beliefs can propagate better in networks with higher vertex degrees, like the preferential attachment graph.

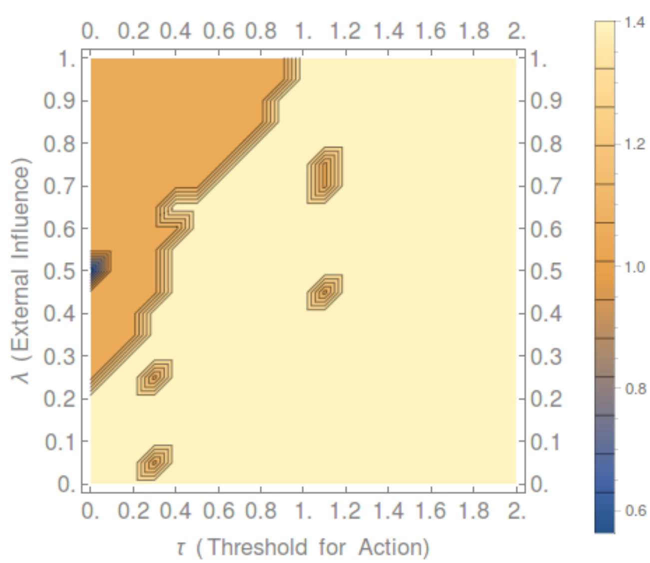

(a) Grid

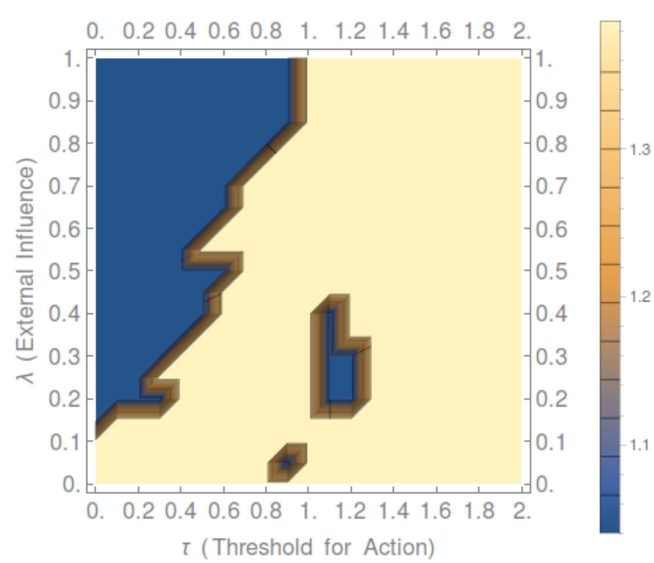

(b) Preferential Attachment

FIG. 10. Entropy of the final state distribution

[8] Secondly we observe that there are small islands of low entropy in the big domain of high entropy. In Figure 10b we see them at $\lambda = 0.3$, and $\tau = 1.1$. This can be explained as follows: People who are initially in a middle state, like $\{0,1\}$ or $\{1,0\}$ will go to the higher state $\{1,1\}$ in the first time step if the threshold permits. This will further increase the attractivity of the high state and lessen the attractivitiy of the middle states, whose occupation numbers drop. As a result the middle will loose their appeal to the ones in the very low state $\{0,0\}$, who will be inert thanks to their high values of τ... don't know how to argue this one well.

4. Success of the policy

The external policy depicted in Figure 6b has been imposed on the agents. We study the success of this policy as we vary both the susceptibility parameter λ and the social mobility parameter τ. The results are shown in figure 11. The fraction of agents in state $\{1,1\}$ ranges between 1/4 and 1 whereas the fraction of agents in state $\{0,0\}$ ranges between 0 and 1/4. This is in agreement with the fact that agents can only move upwards. A successful policy is onewhere a large fraction of people manage to end up in the desired externally stimmulated state and where no large spearation of people in different states occurs.

We see again the similarity between the results for the grid and those for the preferential attachment graph, which is even more similar for the $\{1,1\}$-fraction, arguably the more important one for policy making. This can be viewed again as some degree of insensitivity of the final state fraction to the type of network.

As to the first part of the success criterion, we see that a the parameter space is effectively divided into two regions. At values of τ lower than 1, we can see a high fraction of people in the highest state. This line is very clear for high values of λ but blurs out towards low values of λ and for no external information there is understandibly also nothing to be successfull, so the fraction remains roughly at its lowest value, namely 1/4.

Clearly for very large values of τ the external influence needs to pull very hard to realize its goal as people are reluctant to chnage. If we take $\lambda = 1$ then everyone beliefs exactly what the media imposes (Equ. 29). In order to move into the middle state from the first state one needs a threshold of less than $\mathbb{E}\left[\pi_{\{0,1\}}\right] - \mathbb{E}\left[\pi_{\{0,0\}}\right] = \left(1/3*1+1/3*0+1/3*(-1)\right) - 1*(-1) = 1$ which can be seen in the figures (Fig. 11b and Fig.11d) and similar for moving from a middle state to the top state. These numbers change slightly when the external influence is replaced by social influence, hence the blurring of the line.

With regards to the second success criterion, we can see more intricate behavior. In most of the parameter space people get stuck in the $\{0,0\}$ state if they started out there and thus for those parameters the second success criterion is not met. Only for very low values of τ and values of λ up to about 0.6 do we observe that this criterion is also fulfilled. Especially at $(\tau, \lambda) = (0, 0.5)$ there is a perfect realization point of the policy.

V. CONCLUSIONS AND FUTURE WORK

Much of what influence our behavior and decision making is rooted in our social identity, which consists of our perception of belonging to various social groups. These group memberships make us susceptible to the shared values of each group and thefore their norms. Sometimes such norms may have adverse effects. Here we proposed to frameworks to study the effects of these norms and mechanism by which it could be modifed via external information. In particular, we studied the effects of two paramteres, the social mobility τ and the strength of external information λ, on the convergence of actions and beliefs in a population. We find that as social mobility increases the external information must be more powerful to change the norms, and thefore the actions, in the population. Compared to existing models, the novelty of our models come from multi-dimensional identity attrbitues and inclusion of external information. In future work we extend this framework in several directions. First, we propose a model in which the strength of interaction between agents vary in porportion to the similarity of their beliefs. Moreover, each agent may receive a different external information and with different strength.

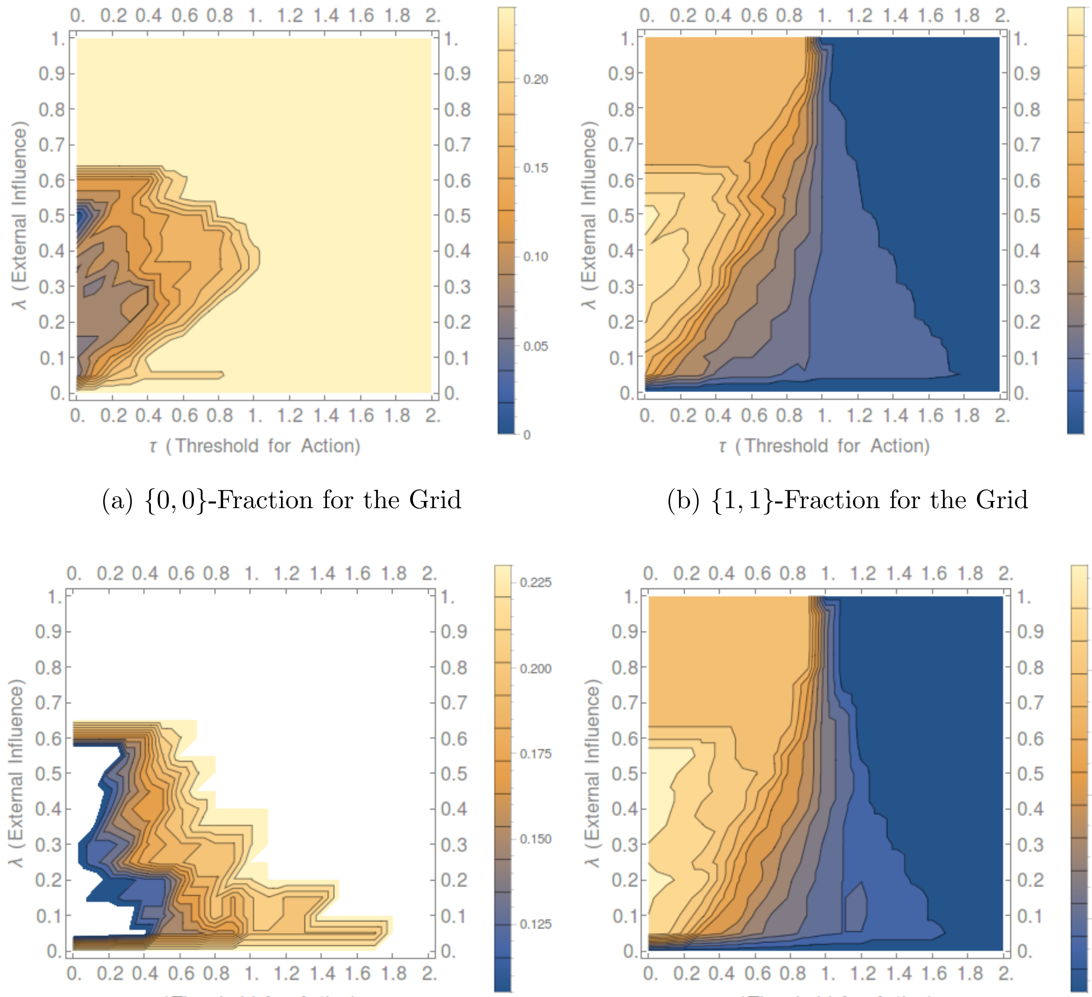

(a) $\{0,0\}$-Fraction for the Grid

(b) $\{1,1\}$-Fraction for the Grid

(c) $\{0,0\}$-Fraction for the Preferential Attachment

(d) $\{1,1\}$-Fraction for the Preferential Attachment

FIG. 11. The fraction of agents who settle to the state $\{0,0\}$ in the Grid (a) or in the PA graph (b), and those that settle to the state $\{1,1\}$ in the grid (c) and the PA grpha (d).

VI. APPENDIX

A. Proof of the Convergence in Linear Model (Equation 1)

In matrix notation we have,

$$D^{tot}(t) = (1-\lambda)WD(t) + \lambda E(t). \tag{30}$$

$$
\begin{aligned}
D^{tot}(1) &= (1-\lambda)WD(0) + \lambda E(0) \\
D^{tot}(2) &= (1-\lambda)WD(1) + \lambda E(1) \\
D^{tot}(3) &= (1-\lambda)WD(2) + \lambda E(2) \\
&\vdots
\end{aligned}
$$

Now let's expand $D(t)$

$$
\begin{aligned}
D^{tot}(1) &= (1-\lambda)WD(0) + \lambda E(0) \\
D^{tot}(2) &= (1-\lambda)W\left((1-\lambda)WD(0) + \lambda H(0)\right) + \lambda E(1) \\
D^{tot}(3) &= (1-\lambda)W\left((1-\lambda)W\left((1-\lambda)WD(0) + \lambda H(0)\right) + \lambda E(1)\right) + \lambda E(2) \\
&\vdots
\end{aligned}
$$

When can be written as:

$$
D(t) = \left((1-\lambda)W\right)^t D(0) + \sum_{i=0}^{t-1} \left((1-\lambda)W\right)^{t-i-1} \lambda E(i). \tag{31}
$$

And since the spectral radius of $(1-\lambda)W < 1$ then:

$$
D(t) = \sum_{i=0}^{t-1} \left((1-\lambda)W\right)^{t-i-1} \lambda E(i). \tag{32}
$$

B. Proof of the $mse_{i,j}$ (Equation 2)

Let us denote that belief vector of agents i and j of length m with $D_i t$ and $D_j(t)$. We would like to know given an connection network between agents and an external signal what is the average mean-squared error between the $D_j(t)$ and $D_i(t)$. To calculate this assume we could run T identical copies of the system but in each one receive its own distinct copy of $E(t)$ from the same distribution. Then at time t we have[9]:

$$
\begin{aligned}
\langle mse_{i,j}(t)\rangle_{m,T} &= \frac{1}{T}\frac{1}{m}\sum_{n=1}^{T}\sum_{l=1}^{m}(D_{n,i,l}(t) - D_{n,j,l}(t))^2 \\
&= \frac{1}{T}\frac{1}{m}\sum_{n=1}^{T}\sum_{l=1}^{m} D^2_{n,i,l}(t) + D^2_{n,j,l}(t) - 2D_{n,i,l}(t)D_{n,j,l}(t) \\
&= \frac{1}{T}\frac{1}{m}\left(\sum_{n=1}^{T}\sum_{l=1}^{m} D^2_{n,i,l}(t) + \sum_{n=1}^{N}\sum_{l=1}^{m} D^2_{n,j,l}(t) - 2\sum_{n=1}^{T}\sum_{l=1}^{m} D_{n,i,l}(t)D_{n,j,l}(t)\right) \\
&= \langle D_i^2(t)\rangle + \langle D_j^2(t)\rangle \\
&= 2\delta
\end{aligned}
$$

[9] We are doing an ensemble average where we calcualte the result over T systems running simultaneously, each with a distinct external input history.

where δ is the variance of the beliefs.

C. Proof of $\Delta(t)$ (Equation 3)

Let $m = 1$, $\langle E^2 \rangle$ the variance of external signal, and I the identity matrix of order N. Also let $W = UVU^{-1}$ the eigenvalue decomposition of W, $\bar{W} = U\bar{V}U^{-1}$ and $\bar{V} = (I - V)^{-1}$. v and $\bar{v}$ are column vectors corresponding to the diagonal elements of V and $\bar{V}$. The covariance of the beliefs of all agents with each other at time t represented by $\Delta(t)$ is given by:

$$\langle \Delta \rangle_T(t) = \langle D(t)D'(t) \rangle_T \tag{33}$$

$$= \sum_{n=0}^{T} \sum_{i=0}^{t-1} \sum_{j=0}^{t-1} ((1-\lambda)W)^{t-i-1} \lambda \bar{e} E_n(i) E_n(j) \bar{e}' \lambda \left(W'(1-\lambda)\right)^{t-j-1}, \tag{34}$$

$$= \sum_{i=0}^{t-1} \sum_{j=0}^{t-1} ((1-\lambda)W)^{t-i-1} \lambda \bar{e} \left(\sum_{n=0}^{T} E_n(i) E_n(j) \right) \bar{e}' \lambda \left(W'(1-\lambda)\right)^{t-j-1}, \tag{35}$$

$$= \lambda^2 \langle E^2 \rangle \sum_{i=0}^{t-1} ((1-\lambda)W)^{t-i-1} \bar{e}\bar{e}' \left(W'(1-\lambda)\right)^{t-i-1}, \tag{36}$$

$$\tag{37}$$

This power sum is not doable. But if we write it using the eigen decomposion of W, we can have:

$$= \frac{\langle E^2 \rangle \lambda}{1-\lambda} \sum_{i=0}^{t-1} UV^{t-i-1}U^{-1}\bar{e}\bar{e}'U'^{-1}V'^{t-i-1}U', \tag{38}$$

$$= \frac{\langle E^2 \rangle \lambda}{1-\lambda} UU^{-1}\bar{e}\bar{e}'U'^{-1} \circ \left(\sum_{i=0}^{t-1} v^{t-i-1} \otimes v'^{t-i-1} \right) U', \tag{39}$$

$$= \frac{\langle E^2 \rangle \lambda}{1-\lambda} UU^{-1}\bar{e}\bar{e}'U'^{-1} \circ \bar{v} \otimes \bar{v}' U', \tag{40}$$

$$= \frac{\langle E^2 \rangle \lambda}{1-\lambda} U\bar{V}U^{-1}\bar{e}\bar{e}'U'^{-1}\bar{V}'U', \tag{41}$$

$$= \frac{\langle E^2 \rangle \lambda}{1-\lambda} \bar{W}\bar{e}\bar{e}'\bar{W}'. \tag{42}$$

ACKLOWLEDGEMENT

F.L. would like to acknowledge the support from the Institue of Humane Studies at the George Mason University under the Hayek Fund for Scholars awards grant of 2014 and from a grant from the John Templeton Foundation.

A.G. would like to acknowledge the support from the National Science Foundation of the United States of America under grant CDI-1028238 and CCF-1318833.

L.H. would like to acknowledge the support from the Max Planck Society under the

IMPRS scholarship.

---

[1] H. Tajfel. *Human groups and social categories*. Camridge: Cambridge University Press, 1981.

[2] Cristina Bicchieri. *The Grammar of Society: the Nature and Dynamics of Social Norms*. New York: Cambridge University Press, 2006.

[3] George A. Akerlof and Rachel E. Kranton. Economics and identity. *The Quarterly Journal of Economics*, 115(3):714–753, 2000.

[4] George A. Akerlof and Rachel E. Kranton. *Identity Economics: How Our Identities Shape Our Work, Wages, and Well-Being*. Princton: Princeton University Press, 2010.

[5] Sushil Bikhchandani, Ivo Welch, and David A. Hirshleifer. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):992–1026, 1992.

[6] Abhijit Banerjee. A simple model of herd behavior. *The Quarterly Journal of Economics*, 107:797817, 1992.

[7] Lones Smith and Peter Sørensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, 2000.

[8] Fudenberg Drew Banerjee, Abhijit. Word-of-mouth learning. *Games and Economic Behavior*, 46:122, 2004.

[9] B. Celen and S. Kariv. Distinguishing informational cascades from herd behavior in the laboratory. *The American Economic Review*, 94(3):484–498, 2004.

[10] B. Celen and S. Kariv. Observational learning under imperfect information. *Games and Economic Behavior*, 47(1):72–86, 2004.

[11] Dahleh Munther Lobel Ilan Ozdaglar Asuman Acemoglu, Daron. Bayesian learning in social networks. *The Review of Economic Studies*, 78:1201–1236, 2011.

[12] Venkatesh Bala and Sanjeev Goyal. Learning from neighbours. *Review of Economic Studies*, 65(3):595–621, 1998.

[13] Venkatesh Bala and Sanjeev Goyal. Conformism and diversity under social learning. *Economic Theory*, 17(1):101–120, 2001.

[14] Vayanos Dimitri Zwiebel Jeffery DeMarzo, Peter M. Persuasion bias, social influence, and unidimensional opinions. *The Quarterly Journal of Economics*, 118:909–968, 1974.

[15] Benjamin Golub and Matthew O. Jackson. Naïve learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, 2(1):112–49, 2010.

[16] Sebastiano Bavetta and Francesco Guala. Autonomy freedom and the right of deliberation. *Journal of Theoretical Politics*, 15:423–43, 2003.

[17] John S. Mill. *On Liberty*. Peterborough, Canada: Broadview Press, 1999.

[18] Scott E. Page. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton, NJ: Princeton University Press, 2007.

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

[20] AL Barabási and R Albert. Emergence of scaling in random networks. *Science*, 286(5439):509, 1999.

[21] A. Goudarzi and D. Stefanovic. Towards a calculus of echo state networks. *to appear in BICA Procedia*, arXiv:1409.0280 [cs.NE], 2014.