SAND2014-0834C

Reducing diffusion time through agenda setting in a multi-agent multi-attitude model ⋆

Paper 17

No Institute Given

Abstract. We are interested in attitude diffusion – how attitudes change over time in groups of people. Attitudes are important to study as they can predict behavior. Our first contribution is a new model the Multi-Agent, Multi-Attitude (MAMA) model that captures social interaction and the internal dynamics of attitudes predicated on the theory of cognitive consistency. Using the MAMA model, we investigate the impact of mass media based agenda setting – that is where some entity sets the topic of discussion between individuals. Surprisingly, we show through extensive empirical simulations, that merely influencing the conversation through agenda setting can significantly reduce diffusion time.

1 Introduction and Motivating Example

We are interested in attitude diffusion – how attitudes change over time among groups of people. Studying attitudes is important, as attitudes can correlate to behavior (for instance, see theory of reasoned action and theory of planned behavior [4] for models linking attitudes to behaviors). Understanding attitude change can help predict behavior change.

In this work, we will present a novel model of attitude diffusion, the Multi-Agent Multi-Attitude (MAMA) model that incorporates three influences: social (from peer groups), cognitive (attitude interaction), and media (through agenda setting). We show, through extensive empirical experiments, the impact of media influence, through agenda setting in particular, on population wide attitude change.

As a motivating example for the development of the MAMA model and the focus on agenda setting, consider the change in attitudes and behaviors towards immunization that has taken place over the last two decades. First, the rate of parents refusing childhood vaccinations have nearly doubled in the 1991-2004 period [10]. Survey data has observed a correlation between non-vaccination and attitudes toward non-vaccination – showing that attitudes can help in predicting vaccination behavior [12].

Several studies have tried to identify the underlying attitudes and the influences on this decision. First, [6] shows that multiple information sources can influence a parents decision, including information from peers (family, friends),

⋆ Due to the constraints of my employer this paper is only released for reviewing purposes and must not be distributed or copied for any other purpose.

doctors, and government entities. However, another important source of informational influence was from media sources, such as t.v. talk shows, radio, internet and newspapers. This highlights the importance of considering both social and media influences.

Secondly, [15] found that unvaccinated children had parents that were more likely to have a low level of trust in the government (among other factors). This data highlights the interconnected attitudes relevant to this behavior – ones attitude towards the government can influence ones attitude toward vaccinations.

While our example is focused on vaccines, it is important to note that the factors modeled here are relevant to other domains as well. For instance, solar panel adoption has interacting attitudes (towards environmentalism, trust in government, etc) and peer influence [13].

In this work we focus on assessing the impact of select "adopter" agents paired with media influence (through agenda setting). Our goal is to understand how the speed of diffusion changes as a function of media influence. This is similar to the "influence maximization" problem, although here we will be varying media influence rather than which agents are being influences. The problem description is, however, similar.

Cognitive consistency The drive for cognitive consistency is a hypothesized drive for individuals to have attitudes that are "consistent" with each other. For instance, according to these theories, an individual holding a strong positive attitude towards environmentalism should also hold a strong positive attitude towards recycling; if they do not, the attitudes are inconsistent with each other and could cause an uncomfortable feeling (i.e. *cognitive dissonance*) which tends to result in either attitude or behavior change [20].

Cognitive consistency has long been shown to be an important factor in attitude change [14, 16] – thus to understand attitude change in a model, we should consider the impact of this drive.

Media Effects Mass media has been studied extensively to understand its affects on attitudes (and thus behaviors). One of the important characteristics of media is "agenda setting" – where the media's focus on certain stories increases their importance in the minds of the viewers [21, 9]. Several studies have provided evidence towards this claim [9, 8, 21].

We focus on agenda setting because of its generality – it only depends upon the focus of the media, and not necessarily on the content of the arguments. It is the minimal effect of media.

2 MAMA Model

To explore the impact of agenda setting we have developed a multi-level agent based model that contains a *social* level – which captures interpersonal interaction between agents – and a *cognitive* level that captures the interactions of attitudes *within* an agent.

Let $G_s =< V_s, E_s >$ be a undirected graph that represents the social level of the model. Let $a_i \in V_s$ be the set of agents, and $(a_i, a_j) \in E_s$ represent a

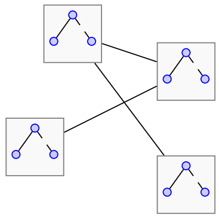

Fig. 1. Social network. See text for details.

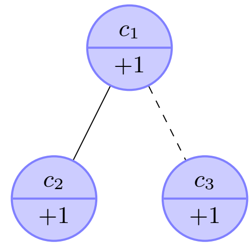

Fig. 2. Cognitive network. See text for details.

bidirectional influencing relationship between agents i and j. Figure 1 depicts an example social network, where each rectangle is an agent

Each agent has a cognitive network associated with it. A cognitive network is a weighted undirected graph, $G_c =< V_c, E_c >$ that represents cognitions and the interactions between them. We use the term cognitions to refer to any entity towards which an individual can have an attitude, such as people, places and things; but also to more abstruse entities like values.

Let $c_k \in V_c$ be the set of cognitions, and $(c_k, c_q) \in E_c$ represents a bidirectional influencing relationship between cognitions k and q. $w(k, q)$ is the weight of edge (c_k, c_q); the weight can either be $+1$, or -1: $w(k, q) \in \{1, -1\}$. The weight represents the relationship between cognitions, as we describe later on. For convenience, we let $n_c = |V_c|$.

An attitude towards a cognition is represented as a real number, called the *value* of the cognition, between -1 and $+1$. The sign of the value represents the *valence* of the attitude; positive values indicate positive attitudes and negative values represent negative attitudes. The size of the value represents how strongly the individuals holds the attitude. So a value of 0.5 would be a mildly positive attitude, whereas a value of -1.0 is a very strong positive attitude. In this work, we limit the values to be either -1 or $+1$. Let $v(i, k)$ be the value of cognition k of agent i.

Figure 2 depicts a cognitive network. The lines represent relationships between cognitions; dashed lines are negative relationships, solid lines are positive relationships. The bottom of each cognition contains the currently assigned value.

Let $\chi_i(k, q)$ be the *consistency* of an edge (c_k, c_q) in the cognitive network of agent a_i. The value of $\chi_i(k, q)$ is:

$$\chi_i(k,l) = \begin{cases} 1 & \text{if } w(k,q)v(i,k)v(i,q) > 0, \\ 0 & \text{Otherwise} \end{cases} \tag{1}$$

Let the *state* of a cognitive network be an assignment of values to its cognitions. There are $m = 2^{n_c}$ states for a cognitive network, labelled: $s_1 \ldots s_m$. $s_p(k)$ is the value of cognition k in state p.

The consistency of a cognition k for agent i is:

$$\phi_i(k) = \frac{\sum_{(c_k,c_q)\in E_c} \chi_i(k,q)}{l_i} \tag{2}$$

where l_i is the number of edges incident to concept i. Intuitively, consistency increases as a cognition has more edges that are consistent.

Cognitive networks can be viewed as *bi-valued, binary constraint satisfaction network* [3]. A significant body of work has been developed around binary constraint satisfaction. The problem there is finding the correct solution; our problem is understanding when a solution diffuses across a network.

2.1 Attitude change

In our model, attitude change is initiated by interpersonal interaction, but mediated by the state of the cognitive network. Given no cognitive influence, the baseline probability that an agents changes their attitude is indicated by P_{base}. In this work we assume that P_{base} is the same for all agents. In the following we describe how the state of the cognitive network modifies P_{base}.

Drive for cognitive consistency Based on the drive for cognitive consistency, we assign a multiplicative weight on the baseline probability of change that varies as a function of how consistent a cognition is with it's neighboring cognitions.

Let $f_{con}(k,i)$ represent the *inclination* to change cognition k of agent i based on its consistency with other cognitions. Intuitively, the more consistent the concept is with its neighbors, the less likely it is to change.

We define $f_{con}(k,i)$ as a sigmoid curve:

$$f_{con}(k,i) = \epsilon + \frac{2}{1 + e^{-10((1-\phi_i(k))-.5)})} \tag{3}$$

For cognitions that have more than 50% of their neighbors in an inconsistent state, $f_{con}(k) > 1.0$, thus increasing the probability they will change (with a maximum multiplicative increase of 2). For those with less than 50% of their neighbors in an inconsistent state, $f_{con}(k) < 1.0$, decreasing the probability to change (with a minimum of ϵ).

While there has been significant work studying cognitive consistency between two cognitions, there has been little or no work done between multiple cognitions. Thus, for our purposes we assume a sigmoid curve with a center at 0.5. As new evidence is uncovered this part may change.

Embeddedness The *embeddedness* of a cognition refers to how well it is connected to other cognitions in the cognitive network. Embeddedness is related to a resistance to change (see [4, Chap. 12] for a review) – a well connected cognition is harder to change because if it did change, the drive for cognitive consistency would cause dissonance with a larger number of other attitudes.

We represent this resistance to change as a multiplicative weight on P_{base}.

Let $f_{deg}(k,i)$ be the *resistance* to change cognition k of agent i based on the cognitions embeddedness, which we measure through its degree $(deg(k))$. Intuitively, we want $f_{deg}(k,i)$ to decrease as we increase the degree of the concept.

$$f_{deg}(k,i) = \begin{cases} 1.0 & \text{if } deg(k) < deg_{max,i}/2, \\ 0.5 & \text{else} \end{cases} \tag{4}$$

where $deg_{max,i}$ is the highest degree in the cognitive network of agent i.

Probability of Change Bringing everything together, let $P_{change}(k,i)$ be the probability of cognition k of agent i changing value, given that i is interacting with another agent with the opposite value for cognition k. Then:

$$P_{change}(k,i) = P_{baseline} \cdot f_{degree}(k,i) \cdot f_{con}(k,i) \tag{5}$$

2.2 Agenda Setting

To model the agenda setting effect of a media source we assume that the media can discuss issues related to the different cognitions. At any point in time the *agenda* determines which cognition is the main focus of the media. We define an *agenda* $\pi = [P(c_1), \dots, P(c_{n_c})]$ as a probability distribution over cognitions. Currently we only consider single cognition agendas, in future work we can extend this to where media can focus on sets of cognitions.

A *Time-Independent Agenda* (TIA) is a fixed probability distribution over the cognitions. A special case is the uniform distribution, where each cognition has a probability of $1/n_c$ of being chosen.

A *Time Varying Agenda* (TVA) is an agenda that changes over time. Essentially, it is some number of agendas which are active at certain times. For instance, for the cognitive network from Figure 2, we can define a time varying agenda by specifying multiple agendas that span the timestep range from $(0, \infty)$. From timestep 0 to timestep 1000, the agenda may be $\pi_{1000} = [1/3, 1/3, 1/3]$, but from 1000 onwards, the agenda may be: $\pi_\infty = [1/9, 1/9, 7/9]$.

2.3 Model Update

Algorithm 1 specifies the update process of the model. Similar to other work ([5]) we study the progressive case. Since we have multiple concepts in our model, we designate a single state s^* as the *goal* state. Once a cognitions switches to the value in the progressive state, it cannot switch back.

Each iteration of t is called a single *timestep*, and within each timestep we randomly and with replacement, sample $N = |V_s|$ agents. Thus on average, every agent is chosen once per timestep.

Note that these dynamics correspond closely to a voter model [18], except the probability of switching varies over the length of the simulation. Voter models have a long history as a simple tool to study diffusion (see Section 4.

Algorithm 1: Model Dynamics

```
for t ← 1 to t_max do
    Choose a strategy π
    for t_s ← 1 to N do
        Choose a random agent a_i
        Choose a_j a random neighbor of a_i
        Choose a topic cognition τ according to π.
        if v(a_i, τ) ≠ s*(τ) and v(a_i, τ) ≠ v(a_j, τ) then
            Set v(a_i, τ) = v(a_j, τ) with probability P_change(τ)
        end
    end
end
```

3 Experimental Studies

Our goal is to study how agenda setting influences the propagation of attitudes in a population. Thus, the metric we will use is the *mean diffusion time* – the number of timesteps the system takes for 90% of the population to reach the goal state, averaged over some number of runs (we also call this time to convergence).

We assume that all agents have the same type of cognitive network, depicted in Figure 2, and only vary in their initial state. We call this the "3-Fan" network, because it has one central cognition (c_1) and two ancillary cognitions that connect to it (c_2, c_3). More generally, a "n-Fan" network would have 1 central cognition and $n-1$ ancillary nodes connected to it.

The fan network, while simple, can effectively represents the interaction between attitudes. In fact, a similar topology was used to model student attitudes to the first Persian Gulf War in [19]. To be more concrete, the central cognition can represent an attitude towards an intention – such as getting a vaccine for your child. The ancillary cognitions would represent relevant attitudes that impact the central cognitions – as mentioned earlier, ones attitude towards the government and ones attitude towards the vaccines etc.

The goal state is set to $s^* =< +1, +1, -1 >$. Initially, a random 10% of the population is assigned the goal state. The rest of the agents are assigned the state of $s^{'} =< -1, -1, +1 >$. Note that these are the only two fully consistent states in the 3-Fan network.

We use two social networks. The first is a *k-regular graph* – a graph where each vertex has degree k. We use a network with 1000 agents, and $k = 4$.

The second social network is a small-world network, which is a network that features a high clustering coefficient and short average path length. Small world networks appear in many real-world domains [22]. We used the algorithm defined in [22], implemented in [2]. We use a network with 1000 agents, and following [22], we set the initial number of neighbors to 10, and the rewiring probability to 0.01. This produces a small world network with clustering coefficient of 0.668519 and average path length of 6.289300.

In all the experiments we set the baseline probability to $P_{base} = 0.333$.

Table 1. Mean convergence time (over 100 runs)

Adoption Mechanism	Regular	Small World
MAMA	1050.700	920.600
Simple Voter	30	20.800
Probabilistic Voter	70.400	61.200

3.1 Experiment 1: MAMA and the Voter Models

Since the MAMA model shares the schema of its update strategy with the voter model, it makes sense to ask if there is any difference in the diffusion time between a MAMA model and the voter model. To address this, we tested three different update algorithms on the regular and small world networks:

Multi-Agent, Multi-Attitude Model The update mechanism outlined in 1, with a time independent uniform agenda.
Simple Voter model Set $P_{change}(\tau) = 1.0, \forall\tau$.
Probabilistic voter model Set $P_{change}(\tau) = P_{base}, \forall\tau$ (similar to the Heterogenous voter model [**?**]

Table 1 shows the results, which make it clear that incorporating interacting attitudes dramatically changes the time to convergence. All pairs of expected convergence time were statistically significant (Wilcox rank sum test, $p \approx 0$)

As expected from existing results [**?**], convergence time for the small world networks was lower than the regular network.

These results reach the same main results as described presented in [**?**] which describes the conformity-consistency model (CCM); an extension to the voter model that represents both social and cognitive factors [**?**]. In the CCM, there are N agents, each with a binary vector of size M that represents their cognitive state. At each time step an agent is chosen and will execute the standard voter model process on of its variables with probability p; with probability $q = 1 - p$ the agent engages in a voter model *with itself*, between the elements of its cognitive state. The voter process within itself can be viewed as reducing dissonance between the elements of its cognitive state – with the extremes of all variables at 0 or all variables at 1 to be the no dissonance situations.

[**?**] shows that adding the dissonance reducing effects within an agent increases the time to convergence – as we show in this experiment.

3.2 Experiment 2: Time Independent Agendas

In this experiment, we identify the best time-independent agenda.

We constructed agendas of the following form. Let π^k be an agenda that sets cognition k to $p \in [0.3, 0.9]$ and the other cognitions to $(1 - p)/2$. We have two parameters that can vary, p and k

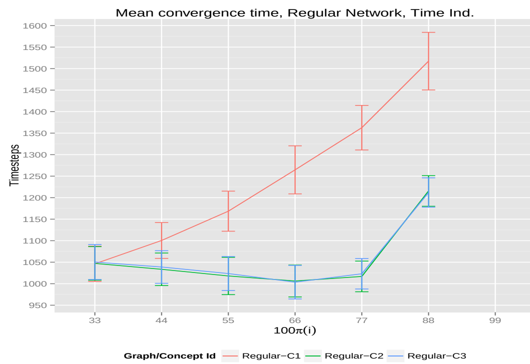

Fig. 3. Mean convergence time on a regular network. The x axis indicates the probability of choosing a particular cognition, and each line is a particular cognition (e.g., the red line is the mean convergence time when choosing cognition 1). The bars indicate one standard deviation.

Figure 3 show the mean convergence time (over 100 runs) for the regular graph (the small world graph had similar results and is not shown) for $k = 1, 2, 3$ and $p \in \{.33, .44, .55, .66, .77, .88, .99\}$. Notice that focusing discussion on the central cognition does not decrease diffusion time, and in fact, it increases as p increases.

On the other hand, a focus on the ancillary cognitions slightly reduces the mean diffusion time, although we do see an inflection point, p^*, beyond which the diffusion time actually increases. Note that none of the runs converged for $p = .99$.

A separate experiment was conducted with agendas of the form $[p_1, p_2, p_3]$, with $p_1, p_2, p_3 \in \{.3, .4, .5, .6, .7, .8, .9\}$ and $p_1 + p_2 + p_3 = 1.0$. These results supported the results shown here – the higher the probability of choosing cognition 1 as a topic, the longer the diffusion time. Results are not shown here due to space limitations.

3.3 Experiment 3: Time Varying Agendas

Is it possible that time varying agendas can reduce diffusion time? To test this, we considered pairs of agendas. At some time b (which we call the *boundary point* we switch agendas. Let agenda π_b^k be an agenda that was used from timestep 0 to timestep b which sets cognition k to $p \in [0.3, 0.9]$ and the other two cognitions to $(1-p)/2$. After timestep b, we set the agenda to $\pi = [1/3, 1/3, 1/3]$. Thus we have three parameters to vary: p, k and b.

Figure 4,5 show the mean convergence time for $p = .3, .6, .89$ on the regular graph (the results were the same for the Small World network and are omitted for brevity), with $k = 1, 2$. Results for $k = 3$ were similar to $k = 2$ and are omitted for space reasons. Results were averaged over 100 runs.

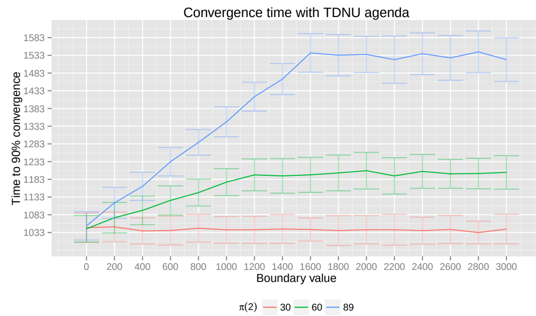

Fig. 4. Mean diffusion time over 100 runs with a time varying agenda. $k = 1$. Each line represents a different probability p.

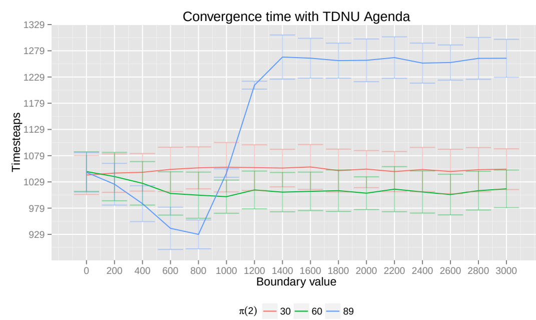

Fig. 5. Mean diffusion time over 100 runs with a time varying agenda. $k = 2$. Each line represents a different probability p.

The results show similar patterns to the time independent agendas. Figure 4 shows that a focus on the central cognition results in higher convergence time.

Figure 5 shows the mean convergence time for agendas in which we change the probability of cognition 2. We see that for low probabilities, the agenda does not seem to make a difference. However, if we increase the probability to .89, we see a dramatic change in the mean convergence time as a function of the boundary value (the mean convergence time at $b = 800$ is less than the mean time for $b = 0$, and is statistically significant, $t - test$, $p \approx 0$)

3.4 Discussion

The results from Experiment 2 & 3 suggest that (1) agenda setting can reduce diffusion time; and (3) time varying agendas have the most impact. However the results prompt three other questions:

1. Why does focusing on the ancillary cognitions reduce diffusion time (in contrast to a uniform agenda)?
2. Why does focusing on cognition 1 cause an increase in diffusion time?
3. Why is there an increase in the time to convergence after a certain boundary value (b^*) in the TVU agenda?

To explore these questions it is useful to consider a Markov Chain representation of state changes. A Markov chain is a stochastic process that varies over time. At each time point t, the process is said to be in a state $s \in s_0, \ldots, s_m$ (we have intentionally used the notation denoting states of the cognitive network) The probability of being in state s_j at time t is a function only of the state the process was in at time $t-1$ (this is the Markov assumption) [**?**]. The probability of moving from state s_q to state s_i is denoted by $P(s_i|s_q) = p_{iq}$ and is called the *transition probability*. A Markov chain is *absorbing* if there are a set of states from which the process cannot leave ($p_{iq} = 0, \forall q$) [**?**].

We can construct an absorbing Markov chain representation of an agent by setting the states of the chain to the states of the cognitive network of the agent (s^* is an absorbing state). Figure 6 depicts the Markov chain. The edges are labelled with $P_{change}(i)$ [1], however the actual transition probability would be: $P(s_i|s_q) = P_{change}(k) \cdot P_{neigh}(v(j,k) = s^*(k)) \cdot \pi(k)$ where k is the cognition that changes value between s_i and s_j and $P_{neigh}(v(j,k) = s^*(k))$ is the probability that the value of cognition k of the chosen neighbor is equal to the goal states value.

We can divide the states in the Markov chain by Hamming distance from the start state; so Level $1 = s_1 = s'$, Level $2 = s_2, s_3,$ and s_4, Level $3 = s_5, s_6,$ and s_7 and Level $4 = s^* = s_8$.

Under a TVA focused on cognitions 2 and 3 first the unequal probability of leaving state 1 causes more agents to reach Level 2 than in the uniform case. Figure 7 illustrates this by showing the distribution of agents across states for a single run of Experiment 3 (per parameter setting). The red line is an approximation of a uniform agenda: $\pi_{800,1} = [.3, .35, .35]$, and the blue line is for a TVA which initially starts with $\pi_{800,2} = [.89, .055, .055]$ and switches to a uniform agenda at timestep 800.

In a TVA agenda, agents pool at Levels 2 and 3, then rapidly get to the end state once the agenda changes; in a uniform agenda fewer agents leave the start state – this is reason TVA agendas focused on ancillary cognitions perform better.

This also answers question 3: an agenda focusing on just the central node will result in fewer agents leaving the start state. The more time spent focused on the central cognition (increasing boundary value in Figure 4), the longer it takes for all the agents to make it to the end state.

The answer to the second question has to do with saturation. Once all the agents have moved out of the start state it is useless to use an agenda that

[1] for clarity, self edges are not shown, but are equal to one minus the sum of the outgoing edges

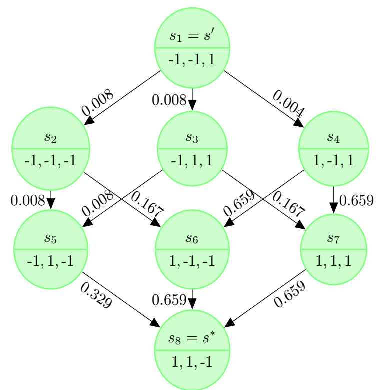

Fig. 6. Markov chain of state transitions, with edges labeled by $P_{change}(i)$ Not accounted for: the agenda π and the distribution of agents over states. The bottom of each node are the values for the state.

focuses on moving agents out of the start state – which is what occurs when $b > 800$

For $b > 1200$ we see the convergence time while using only $\pi_{b,2} = [.055, .89, .055]$. This is worse than the uniform strategy because the agenda is unduly focused on just one cognition for too long.

4 Related Work

To our knowledge, there has been no work that computationally explores the impact of agenda setting on diffusion.

The interpersonal interaction dynamics of our model are closely related to the *voter model* a well explored model from the physics domain [17] in which nodes can take on the values 0 or 1. At each timestep a random nodes takes on the value of one of its neighbors. While the choice of agents is similar, few voter models capture multiple interacting values within an agent.

In 3.1 we described the CCM model, which has a similar structure to the MAMA model. The main difference is that we do not assume that agents can either communicate or modify internal attitudes – both processes can occur simultaneously in the MAMA model.

The constraint satisfaction model of attitude interaction is based on similar work from social psychology – where it is called a "parallel constraint satisfaction" model and has been applied to a variety of domains, such as impression formation [7], legal inference [16], and as a model of change in attitude to the persian gulf war [19].

In Axelrod's model of social dynamics agents have multiple cultural features where each feature can take on a value from a small set of "traits" [1, 11]. Agents can interact only The model uses a bounded confidence in which only agents that match on a certain percentage of features will interact.

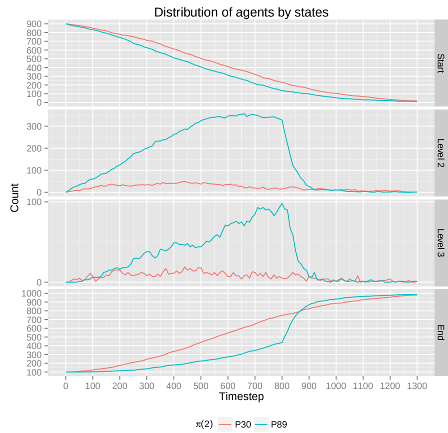

Fig. 7. Distribution of agents across states. Level 2 = $s_2 + s_3 + s_4$, Level 3= $s_5 + s_6 + s_7$.

Several extensions to Axelrod's model have been proposed that incorporate mass media. Often, this is incorporated through a virtual agent that represents the media and which has edges connecting it to all other agents in the population [11]. This work studies the influence effect of media, and not the agenda setting effect.

5 Conclusions

Agent-based simulation is an important tool that allows empirically study of complex interactions, in our case between interpersonal influence and attitudes. In this work, we developed a novel agent based model that captures social and cognitive factors that affect decision making (the MAMA model). We used agent-based simulation to study the impact of agenda setting on diffusion time within the MAMA model. We found that:

1. Agenda-setting can significantly affect diffusion time – both positively and negatively.
2. The best agenda for the "3-Fan" network is a time varying agenda that focuses on an ancillary attitudes first, then switches to a uniform agenda which can reduce mean diffusion time by by $\approx$ 100 timesteps (statistically significant)

Surprisingly, merely setting the topic of discussion (not even influencing attitudes) can have an impact on diffusion time.

References

1. R. Axelrod. The dissemination of culture: A model with local convergence and global polarization. *The Journal of Conflict Resolution*, 41(2), 1997.

2. G. Csardi and T. Nepusz. The igraph software package for complex network research. *InterJournal*, Complex Systems:1695, 2006.
3. R. Dechter. From local to global consistency. *Artificial Intelligence*, 55(1):87–107, May 1992.
4. A. H. Eagly and S. Chaiken. *The Psychology of Attitudes*. Wadsworth Cengage Learning, 1993.
5. D. Kempe, J. Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of 2003 SIGKDD*, 2003.
6. A. Kennedy, M. Basket, and K. Sheedy. Vaccine attitudes, concerns, and information sources reported by parents of young children: Results from the 2009 HealthStyles survey. *PEDIATRICS*, 127(Supplement):S92–S99, Apr. 2011.
7. Z. Kunda and P. Thagard. Forming impressions from stereotypes, traits, and behaviors: A parallel-constraint-satisfaction theory. *Psychological Review*, 103(2):284–308, 1996.
8. V. Larcinese, R. Puglisi, and J. M. Snyder Jr. Partisan bias in economic news: Evidence on the agenda-setting behavior of U.S. newspapers. *Journal of Public Economics*, 95(9–10):1178–1189, Oct. 2011.
9. M. McCombs. A look at agenda-setting: past, present and future. *Journalism Studies*, 6(4):543–557, 2005.
10. Omer SB, Pan WY, Halsey NA, and et al. Nonmedical exemptions to school immunization requirements: Secular trends and association of state policies with pertussis incidence. *JAMA*, 296(14):1757–1763, Oct. 2006.
11. L. R. Peres and J. F. Fontanari. The media effect in axelrod's model explained. *EPL (Europhysics Letters)*, 96(3):38004, Nov. 2011.
12. P. Peretti-Watel, P. Verger, J. Raude, A. Constant, A. Gautier, C. Jestin, and F. Beck. Dramatic change in public attitudes towards vaccination during the 2009 influenza a(H1N1) pandemic in france. *Euro surveillance: bulletin Européen sur les maladies transmissibles = European communicable disease bulletin*, 18(44), 2013. PMID: 24176658.
13. V. Rai and K. McAndrews. Decision-making and behavior change in residential adopters of solar pv. In *Proceedings of the Behavior, Energy and Climate Change Conference*, 2012.
14. J. E. Russo, K. A. Carlson, M. G. Meloy, and K. Yong. The goal of consistency as a cause of information distortion. *Journal of Experimental Psychology: General*, 2008.
15. D. A. Salmon, L. H. Moulton, S. B. Omer, M. P. DeHart, S. Stokley, and N. A. Halsey. Factors associated with refusal of childhood vaccines among parents of school-aged children: a case-control study. *Archives of pediatrics & adolescent medicine*, 159(5):470–476, May 2005. PMID: 15867122.
16. D. Simon, C. J. Snow, and S. J. Read. The redux of cognitive consistency theories: Evidence judgments by constraint satisfaction. *Journal of Personality and Social Psychology*, 86(6):814–837, 2004.
17. V. Sood, T. Antal, and S. Redner. Voter models on heterogenous networks. *Physical Review E*, 77, 2008.
18. V. Sood and S. Redner. Voter models on heterogenous graphs. *Physical Review Letters*, 94, 2005. 178701.
19. B. A. Spellman, J. B. Ullman, and K. J. Holyoak. A coherence model of cognitive consistency: Dynamics of attitude change during the persian gulf war. *Journal of Social Issues*, 49(4):147–165, 1993.
20. P. S. Visser and J. Cooper. Attitude change. In M. Hogg and J. Cooper, editors, *Sage Handbook of Social Psychology*. Sage Publications, 2003.

21. H. N. N. Vu and V. Gehrau. Agenda diffusion: An integrated model of agenda setting and interpersonal communication. *Journalism & Mass Communication Quarterly*, 87(1):100–116, Mar. 2010.
22. D. J. Watts and S. H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, June 1998.