SAND2013-8667C

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

Kiran Lakkaraju,
Paper ID: 249

ABSTRACT

In this work we study the diffusion of multiple attitudes within a population by using "broad" influence strategies. We define broad influence strategies are ones that have an effect on a large number of people, in contrast to "narrow" strategies which identify and target specific individuals. Mass media ads are the prototypical type of broad influence strategies. Identifying "opinion leaders" in communities and targeting them with free samples etc., is a type of narrow influence strategy.

We develop a new model, the Multi-Agent, Multi-Attitude (MAMA) model that captures social interaction through a social network and the internal dynamics of attitudes predicated on the theory of cognitive consistency. We imbue agents with a "cognitive network" that models the interaction between attitudes. We show that when we have an agenda setting influence strategy (where some entity induces discussion on certain topics) diffusion time can be (statistically) significantly reduced. Which topics are discussed, and with what frequency, has a significant impact on diffusion time.

Categories and Subject Descriptors

I.6 [**Simulation and Modeling**]: Model Development

General Terms

Experimentation

Keywords

social simulation, cognitive consistency

1. INTRODUCTION

We are interested in understanding how populations make decisions that have significant cost and/or long-term impact and how they are influenced by others, internal attitudes and mass media. For instance, purchasing a residential solar panel is capital intensive and also includes assessing non-monetary costs such as maintenance time etc. [25].

A characteristic of these types of behaviors is how they are influenced by a variety of underlying beliefs, values and attitudes. While choosing to retweet or click on a link may be influenced by local factors (I'm bored, let me surf randomly!), longer term, capitol investment decisions will be more thoughtful.

For example, in purchasing a solar panel, a consumer must consider the economics of the purchase, however the attitudes of the consumer can play a role as well. For example, [12] showed how attitudes towards carbon emissions and political ideology can affect energy efficiency product purchasing decisions.

Interpersonal communication and influence still plays a strong role – in fact, [25] shows that peer effects can significantly reduce the decision making time by reducing information uncertainty – they are mediated by these "internal" factors.

In this work we will model the underlying attitudinal change the often precedes behavior change. Specifically, we will model attitudes, which are "general and relatively enduring evaluative responses to objects" where objects can be "a person, a group, an issue or a concept" [40, Page 1]. Attitudes have a long history in social psychology, and have been shown to have an impact on the behaviors of individuals (e.g., voting behavior [15], consumer purchases [13, 10]).

An important fact of attitudes is that there are many of them, on many objects. Attitudes are linked to each other as well, and can influence each other. Theories of cognitive consistency suggest that individuals strive to hold a consistent set of attitudes.

We divide influence mechanisms into two general categories. "Narrow" influence mechanisms focus on identifying specific individuals or small groups that have disproportionate influence on others (also called "opinion leaders" [26]). Many of the techniques derived to solve the influence maximization problem [14, 18, 6, 29] rely on this type of identification of individuals. While powerful, these often require extensive knowledge of the social network of individuals, which may be difficult outside of the online social network realm.

In contrast, we will focus on what we call "broad" influence techniques, such as mass media ads through TV, radio and print media. Somewhat surprisingly, considering the heavy focus on online media, TV is still the "king" of advertising, with a lions share of advertising dollars going to it [1]. The reason is simple, TV ads still work. In fact, in a recent survey of college students, 42% responded that TV ads were the most effective forms of advertising [1].

We believe broad influence techniques merit more attention.

Thus the goal of this work is to develop a better under-

Appears in: *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), Lomuscio, Scerri, Bazzan, Huhns (eds.), May, 5–9, 2014, Paris, France.*

standing of "broad" influence techniques in the context of population wide attitude change. Specifically, in this work we:

1. Develop the Multi-Agent, Multi-Attitude (MAMA) framework that captures social interaction and attitudinal interaction via cognitive consistency.
2. Develop a representation for the "agenda setting" effect that mass media can provide.
3. Show that through "agenda setting" we can (statistically) significantly reduce the time for a particular set of attitudes to diffuse through a population.

1.1 Cognitive consistency

Cognitive consistency is a hypothesized drive for individuals to have attitudes that are "consistent" with each other. Cognitive consistency has long been shown to be an important factor in attitude change. The drive for cognitive consistency is that individuals tend to want to hold a set of attitudes that are consistent with each other [28, 32, 35]. 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 [41].

The surprising, and counter-intuitive, aspect is bidirectional reasoning – under a cognitive consistency model conclusions can affect understanding of premises, rather than the "traditional" means by which premises imply conclusions.

1.2 The role of media

Mass media has been studied extensively to understand its affects on peoples attitudes (and thus behaviors). While there are numerous aspects of media effecting the public, we will discuss two, that of *agenda setting* and *framing*.

Framing refers to how information is presented, and thus how it impacts individuals. This seems to be the most common way of modeling media effects – we can consider it a type of "virtual agent" that interacts with everyone and has a particular set of attitudes.

However, there is another role that media plays, that of agenda setting. Agenda setting is when the media's focus on certain stories increases their importance in the minds of the viewers [42, 20]. Agenda setting and interpersonal communication interact to influence individuals.

2. MODEL

To explore the impact of multiple attitudes we have developed a multi-level agent based model that incorporates two levels. The model contains two levels, the *social* level – which captures interpersonal interaction between agents – and the *cognitive* level which captures the interactions between 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 $(a_i, a_j) \in E_s$ represents a bidirectional influencing relationship between agents i and j. Figure 1 depicts an example social network. Each rectangle is an agent, and each agent has their own cognitive network.

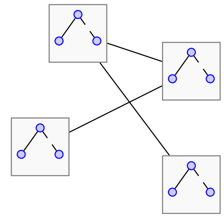

Figure 1: Social network. See text for details.

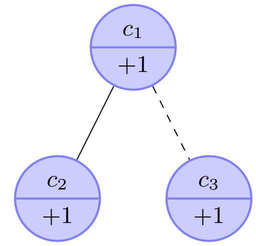

Figure 2: Cognitive network. See text for details.

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

Each cognition represents a concept and the evaluation of the concept. For instance, this could be the concept of "environmentalism" and the evaluation for a particular agent could be positive or negative.

We represent this as a value for each cognition within an agent: $v(i, k)$ is the value of cognition k of agent i. In this work we only consider binary values, either the evaluation can be positive or it can be negative: $v(k) \in \{-1, +1\}$.

Figure 2 depicts a cognitive network. The lines represent a relationship between the cognitions, with the weights indicated. The top of the node is the node name, and the bottom is the currently assigned value.

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

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

Intuitively, if an edge has negative weight, the edge is consistent if the two vertices have *opposing* values. If an edge has a positive weight, the edge is consistent if the two vertices have the *same* value.

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: $s_0 \ldots s_m$. $s_p(k)$ is the value of cognition k in state p.

The consistency of a concept k for agent i is:

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

where l_i is the number of edges incident to concept i.

The consistency of the cognitive network for agent i under state assignment s_p is $\Phi_i(s_p)$ and is defined as:

$$\Phi_i(s_p) = \frac{\sum_{(c_k,c_l)\in E_c} \chi_i(k,l)}{|E_c|} \tag{3}$$

Intuitively, the consistency of a cognitive network is the proportion of edges within the cognitive network that are consistent.

Cognitive networks can be viewed as *bi-valued, binary constraint satisfaction network* [8]. 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.

It is easy to see that the set of edges will determine whether there exists a fully consistent state (where $\Phi_i(s_p) = 1.0$). It may be the case that there are no fully consistent states.

2.1 Cognition change

In this model attitude change occurs as a function of interpersonal communication (social influence) mediated by cognitive aspects. The baseline probability of attitude change is represented by P_{base}. We consider two related effects on concept change – the drive for cognitive consistency and the embeddeness of the concept – drawing from the social psychology literature.

2.1.1 Drive for cognitive consistency

As mentioned before, the drive for cognitive consistency would change the likelihood of changing attitudes.

In Figure 2, concept c_1 is in an inconsistent state with concept c_3 – the link between them is negative, so they should be opposite, however there are positive attitudes towards both concepts. Based on cognitive consistency theory, we should expect c_1 to be more likely to change[1]. On the other hand, concepts that are highly consistent should be less likely to change – since they are consistent with most of their neighbors.

We represent this drive to consistency as a multiplicative weight on the baseline probability.

Let $f_{con}(k,i)$ represent the *inclination* to change concept k of agent i based on the how consistent it is with other concepts (for simplicity we occassionaly drop the reference to the particular agent). Intuitively, the more consistent the concept is with it's neighbors, the less likely to change.

We define $f_{con}(k,i)$ as

$$f_{con}(k) = \epsilon + \frac{2}{1 + e^{-10((1-\phi_i(k))-.5)})} \quad (4)$$

Figure 3 plots this as a function of $1 - \phi\, i(k)$ for $\epsilon = 0.01$. For concepts that have more than 50% of their neighbors in an inconsistent state,$f_{con}(k) > 1.0$, increasing the probability to change (with a max 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).

For example, consider Figure 2; $f_{con}(1) = 0.5$ because concept c_3 is inconsistent with c_1, however c_2 is consistent.

On the other hand, $f_{con}(2) = \epsilon$, since all of its neighbors are consistent with it. Finally, $f_{con}(3) \approx 2.0$, since all of its neighbors are inconsistent with it.

2.1.2 Concept Embeddedness

The second cognitive factor that we capture is the concept of *embeddeness* which captures a resistance to change based on how connected the concept is to other concepts. Consider an individual who initially does not have a positive attitude towards environmentalism. If the individual were to switch,

[1]We are greatly oversimplifying, there are a host of mediating factors, including the type of attitude (implicit or explicit), the persuasion route, etc. See [9, 21]

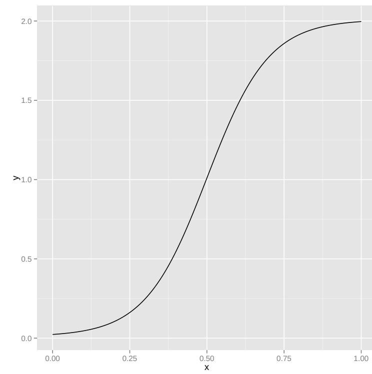

Figure 3: f_{con}(k,i) for $\epsilon = 0.01$ as a function of $1-\phi\, i(k)$

it would have serious consequences to other aspects, such as their attitude towards types of cars, energy production, and energy consumption.

A host of research has shown the importance of embeddedeness on resistance to change (see [9, Chap. 12] for a review).

In Figure 2, concept c_1 is connected to two other concepts vs. c_2, which is only connected to 1 other concept. Thus, we would expect c_1 to have more resistance to change than c_2 or c_3.

We represent this resistance to change as a multiplicative weight on the baseline probability of changing.

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

$$f_{deg}(k,i) = \begin{cases} 1.0 & \text{if } deg(k) < deg(max)/2, \\ 0.5 & \text{else} \end{cases} \quad (5)$$

where $deg(max)$ is the highest degree in the cognitive network. We chose this relative measure because it can apply to a wide variety of cognitive network structures – the resistance to change is relative to the other concepts in your network.

For example, Figure 2; $f_{deg}(1) = 0.5$, since concept c_1 has the highest degree. $f_{deg}(2) = f_{deg}(3) = 1.0$, since they only have one neighbor.

2.1.3 Probability of Change

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

$$P_{change}(k) = P_{baseline} \cdot f_{degree}(k) \cdot f_{con}(k) \quad (6)$$

2.2 Model Dynamics

Algorithm 1 specifies the dynamics of the model. Similar to other work ([14]) 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 concept 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

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 concept 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
```

once per timestep.

Note that these dynamics correspond closely to a voter model , except the probability of switching is varies over the length of the simulation.

2.3 Agenda Setting

The choice of the topic concept will be of significant interest to this work. We define an *agenda* $\pi = [P(c_1), \ldots, P(c_{n_c})]$ as a probability distribution over concepts that can vary over the time of the simulation.

A *time-independent, uniform agenda* is a fixed, uniform probability distribution over the concepts. Every concept has a probability of $1/n_c$ of being chosen.

A *time-independent, non-uniform agenda* is a fixed, non-uniform probability distribution over concepts. For instance, we may choose a particular concept with a higher probability than another.

A *time dependent, non-uniform agenda* (TDNU) is a fixed, non-uniform probability distribution over concepts that vary with time. For instance, from timestep 0 to timestep 1000, the strategy may be $\pi_{1000} = [1/3, 1/3, 1/3]$, but from 1000 onwards, may be $\pi_\infty = [1/9, 1/9, 7/9]$.

3. EXPERIMENTAL STUDIES

In this section we go over three experiments. First, establish that diffusion time is significantly higher in the MAMA model than in the voter model. Secondly, we identify the impact of different time independent strategies on diffusion time. Thirdly, we investigate the impact of time dependent strategies on diffusion time.

3.1 Experimental Setting

We used two social networks. The first was a *k-regular graph* – a graph where each vertex has degree k. The second social network was a small-world network. A small-world network feature a high clustering coefficient and short average path length. Small world networks appear in many real-world social networks [43]. We used the algorithm defined in [43], implemented in [7]. Following [43], 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.

We assumed that all agents had the same type of cognitive network, and only varied in their states. We focus on a single type of cognitive network, depicted in Figure 2. We call this the "3-Fan" network, because it has one central cognition (c_1) and two ancillary nodes 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, represents the interaction between attitudes. Consider the central cognition as an attitude towards a specific decision, such as purchasing an energy efficient lightbulb. The ancillary nodes represent the other cognitions that can impact this decision, say political preference [12]. The bidirectional link between the central node and the ancillary nodes illustrate the bidirectionality of influence; your attitudes can influence your behavior, but your behavior can influence your attitudes (which is in line with the work on cognitive dissonance described earlier).

The goal state in these experiments is $s^* =< +1, +1, -1 >$. 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 say that a population has converged, or a behavior has diffused, when more than 90% of the agents have reached state s^*. The main metric we use is the *mean diffusion time* – the average number of timesteps till the population converges.

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

3.2 Experiment 1: MAMA and the Voter Models

The first question that should be asked is whether adding the cognitive consistency aspects makes any difference to convergence time. 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 [19]

As we mentioned before, the voter model comes closest in terms of update rules to the MAMA model. The 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 [4], convergence time for the small world networks was lower than the regular network.

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

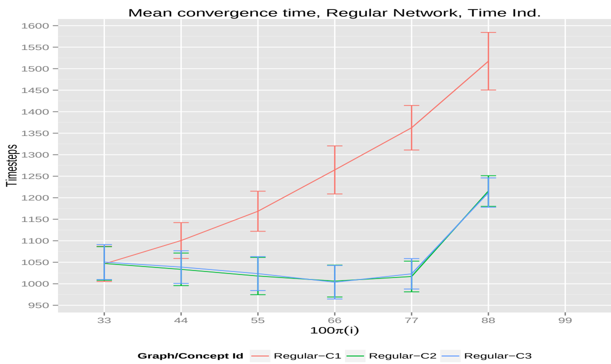

Figure 4: Mean convergence time on a regular network. The x axis indicates the probability of choosing a particular cognition, and each line is a particular cognition (e.g., the red line is the mean convergence time when choosing cognition 1). The bars indicate one standard deviation.

3.3 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 $(1-p)/2$. This strategy will be used through the entire time period of the experiment. We have two parameters that can vary, p and k

Figure 4 5show the mean convergence time (over 100 runs) for the regular graph and small world graph for $k = 1, 2, 3$ and $p \in \{.33, .44, .55, .66, .77, .88, .99\}$. Notice that focusing discussion on the center node doesn't decrease convergence time, and in fact, it increases it as p increases.

On the other hand, a focus on the ancillary cognitions slightly reduces the mean convergence time, although we do see an inflection point p^*, beyond which the convergence time actually increases. Note that none of the runs converged for $p = .99$. The mean convergence time at $p = .66$ is statistically significantly less than the other (two-sided non-paired t-test, $p \approx 0$).

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 discussing cognition 1, the higher the convergence time. Results are not shown here due to space limitations.

We will discuss these results more in Section 3.5

3.4 Experiment 3: Time Dependent Agendas

Is it possible that time dependent agendas can do better? To test this, we considered a pair of agendas. Up until some boundary point b, one agenda would be used, and after b, another would be used. 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 6,7 show the mean convergence time for $p = .3, .6, .89$ on the regular graph (the results were the same for the Small

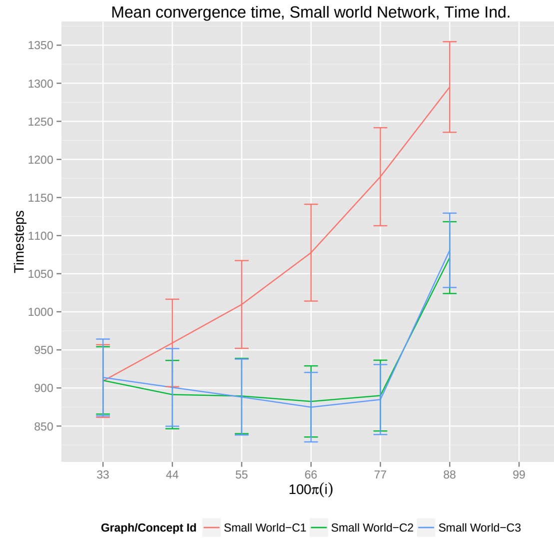

Figure 5: Convergence time on a SW network. The x axis indicates the probability of choosing a particular cognition, and each line is a particular cognition (e.g., the red line is the mean convergence time when choosing cognition 1). The bars indicate one standard deviation.

World network and are omitted for brevity). Results were averaged over 100 runs.

The results show similar patterns to the time independent agendas. Figure 6 shows that a focus on the central cognition results in higher convergence time. The time to convergence is related to the probability and boundary value.

Figure 7 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. We will discuss this in more detail in Section 3.5.

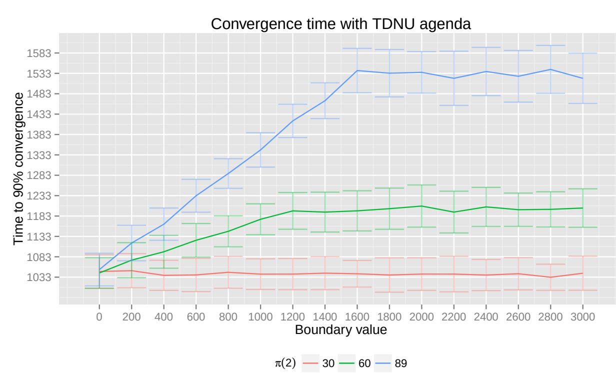

Figure 6: Mean convergence time over 100 runs with a time independent agenda. $k = 1$. Each line represents a different probability p.

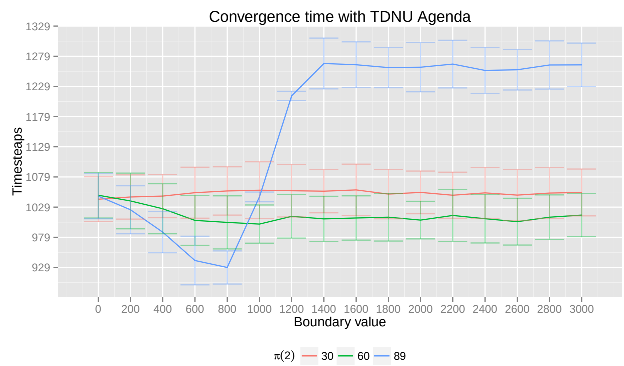

Figure 7: Mean convergence time over 100 runs with a time independent agenda. $k = 1$. Each line represents a different probability p.

3.5 Discussion

The results from Experiment 2 & 3 suggest that time dependent agendas can significantly change the convergence time of the population. However the results prompt three other questions:

1. Why does focusing on the ancillary concepts perform better (quicker convergence) than a uniform agenda?
2. Why does focusing on cognition 1 cause an increase in the convergence time?
3. Why is there an increase in the time to convergence after a certain boundary value (b^*) in the TDNU agenda?

To explore these questions it is useful to consider the changes in state of an agent as a Markov Chain.

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 as they will be the same). 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) [27]. The probability of moving from state s_i to state s_j is denoted by $P(s_i|s_j)p_{ij}$ and is called the *transition probability.* A Markov chain is *absorbing* if there are a set of states from which the process cannot leave ($p_{ij} = 0, \forall j$) [11].

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 8 depicts the Markov chain. The edges are labelled with $P_{change}(i)$ (for clarity, self edges are not shown, but are equal to one minus the sum of the outgoing edges), however the actual transition probability would be:

$$P(s_i|s_j) = P_{change}(k) \cdot P(k = s^*(k)) \cdot \pi(k) \tag{7}$$

where k is the cognition that changes state between $s_i$ands_j. Intuitively, this equation says that the probability of transitioning is the probability of picking as a topic the cognition that needs to change ($\pi(k)$), times the probability of a neighbor's cognition having the correct value ($P(k = s^*(k))$), times the actual probability of changing that cognition within the agent (P_{change}(k)).

Of these elements, P_{change}(k) is fixed by the cognitive network and the state; and $P(k = s^*(k))$ is a function of the distribution of agents in the population. $\pi(k)$ is the only variable that can be changed.

Now we can begin to answer the questions above. First, we want to understand why TDNU agenda's perform better than uniform. It helps to think of the states in the Markov Chain as being in 4 groups, the start states (which is $s_1 = s'$); the states at 1 hamming distance from the start states, $s_2, s_3,$ and s_4 which we call *Level 2*; the states at 2 hamming distance from the start state $s_5, s_6,$ and s_7 which we call *Level 3*; and the ending states $s^* = s_8$.

Note that the probability of leaving the start state is higher when when changing cognitions 2 and 3. This is because they are less resistant to change, being less embedded. However, if one manages to change cognition 1, the probability of going to the end state is relatively high.

With a TDNU strategy that focuses on cognition 2 or 3, many agents will move from the start state to level 2 and level 3. Once there, it will be relatively difficult to get to the end state, since the probability to change is low. On the other hand, in the uniform agenda, fewer agents will leave the start state – all the extra probability mass on choosing cognition 1 as a topic is to waste since few agents will be able to change their cognition. Those who do will very quickly reach the end state, however.

Figure 9 shows the distribution of agents across states comparing the TDNU agenda $\pi_{800,2} = [.89, .055, .055]$ and $\pi_{800,2} = [.3, .35, .35]$ which approximates the uniform agenda. You can see that with $\pi_{800,89}$, agents "pool" in levels 2 and 3, after rapidly leaving the start state. In the uniform case, there are relatively few agents in levels 2 and 3, with most of the agents in the start state or the end state.

The benefit of the TDNU shows when we reach the boundary point at 800. Once we switch to a uniform strategy, all the agents from level 2 and level 3 swiftly make their way to the end state.

So to answer our question from above, the reason why the TDNU agenda helps is that it brings many more agents out of the start state, which allows them to influence others. In contrast, the uniform agenda brings few agents out of the start state, but does get them all the way to the end state.

The second question is why is there a convex pattern to the expected convergence time as we increase the boundary value for the TDNU strategy on C2, with probability 0.89? The reason for this is over saturation. As we move the boundary from 0 to 800, we see more and more agents leaving the start state. Around $b = 800$, under an agenda of $\pi_{b,2} = [.89, .055, .055]$, nearly all the agents have left the start state (see the top graph of Figure 9 which shows only 100 agents remaining in s'). Once all the agents have moved out of the start state, spending time on the non-uniform agenda is a waste since most agents have already been converted. Thus the increase in convergence time past $b \approx 800$. Because we are still using the strategy on cognition 2/3 after all the agents have moved out of the start state, it takes longer to reach convergence.

The final phase, after $b = 1200$ is how long it takes for $\pi_{b,2} = [.89, .055, .055]$ to reach convergence without changing agenda. This is worse than the uniform strategy because the agenda is unduly focused on just one cognition for too long.

Finally the third question – why does focusing on cogni-

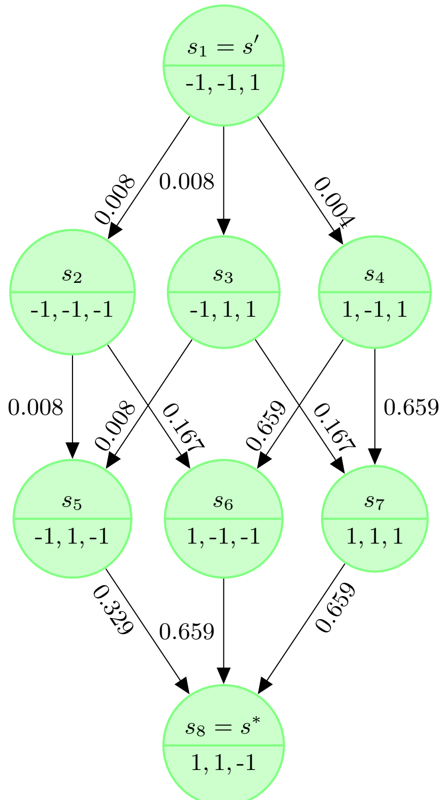

Figure 8: Markov chain of state transitions, with edges labeled by $P_{change}(i)$ Not accounted for: the agenda π and the distribution of agents over states. The bottom of each node are the values for the state.

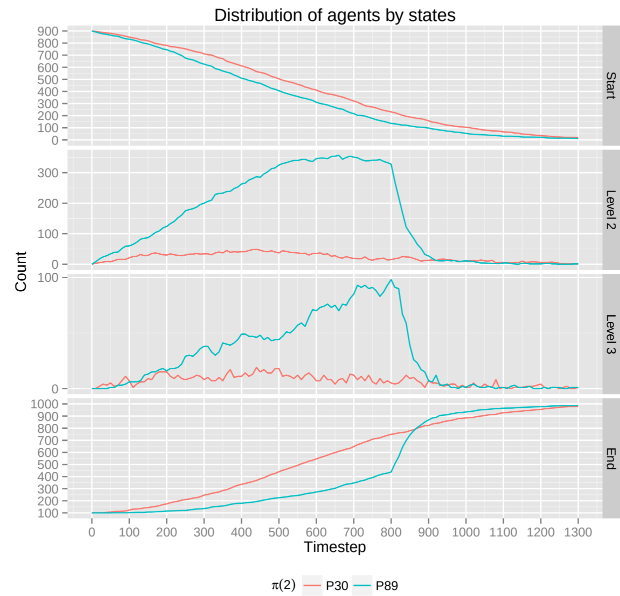

Figure 9: Distribution of agents across states. Level 2 = $s_2 + s_3 + s_4$, Level 3= $s_5 + s_6 + s_7$.

tion 1 hinder efforts? The answer is relatively simple, because the initial transition out of the start state is difficult, agents pool in the start state. This causes a decrease in $P(v(k) = s^*(k))$, and thus generally reducing the probability of transitioning to the end state.

4. RELATED WORK

The communication and interaction dynamics of our model are closely related to the *voter model* a well explored model from the physics domain [36]. In this model there are a set of sites each of which is endowed with a variable that can take either state 0 or state 1. At each timestep a random site changes its state to a randomly chosen neighbors. The voter model has been studied on lattices as well as small-world and heterogenous degree networks [38, 36]. Multiple extensions have been developed that incorporate a majority rule (sites change to the state that is in the majority of their neighborhood) [17, 23, 5]. The main issues with voter models are that they are limited to one concept with only two states; this simplicity is, however, a boon for analysis.

Similar to the voter model, at each timestep we choose one agent which communicates with one of its neighbors. However, our agents are more complicated as they utilize a PCS network.

The conformity-consistency model (CCM) is a model that represents both social and cognitive factors [3, 22]. In the CCM, there are N agents, each endowed with a binary vector of size M that represents their cognitive state. At each time step an agent is chosen and it 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.

The CCM is similar to our model, but differs in two important regards: (1) it does not incorporate modifying the likelihood of changing an concept, and (2) it does not incorporate elements of agenda setting.

Chapter 10 of [39] describes the "consensus = coherence + communication" (CCC) model. In this model each agent has a PCS network where concepts represent hypotheses and evidence, and links represent explanatory relations. The purpose of the model was to explain the diffusion of scientific theories in a population, so pairs of agents could interact (simulating a complete network) and agents could execute "lectures" for many other agents.

The main difference between our model and the CCC model is in the interaction. In the CCC model agents exchange concepts with others, thus changing the structure of the network. We have focused on a different perspective, given that individuals already know the links between concepts, how do persuasive messages between agents lead to attitude change.

The parallel constraint satisfaction model has been explored in a wide variety of contexts, such as impression formation [16], legal inference [33, 31, 34], and as a model of change in attitude to the persian gulf war [37].

Axelrod's cultural diffusion model is a model that contains multiple cultural traits that interact [2]. The model uses a bounded confidence paradigm – so agents that only match on some number of traits are allowed to communicate. However, there is no interaction between traits.

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 [24]. This work also studies the influence effect of media, and not the agenda setting effect.

5. CONCLUSIONS

The MAMA framework captures interactions between attitudes and the influence of agenda setting on diffusion time. We incorporate aspects of the theory of cognitive consistency, drawn from work in social psychology, to capture how attitudes interact and influence each other within an individuals. We show that the diffusion time in the MAMA model is dramatically longer than with more standard models, such as the voter model.

We also showed the impact of agenda setting on diffusion time for agents with the "3-Fan" cognitive network. Through extensive empirical experimentation, we have shown that choosing a non-uniform, time dependent strategy is better than choosing a time-independent uniform strategy – the benefit is a reduction in mean diffusion time by ≈ 100 timesteps, which for the 1000 node network we studies, would be a reduction by $100,000$ agent interactions.

What does this model suggest for advertisers and others wishing to have their product/idea diffuse? We highlight some suggestions. First, it is imperative that one knows the domain of the problem. Some attitudes may be relatively independent, and thus can be considered as independent variables. These may come to convergence through standard convention emergence techniques, dominated by interpersonal interaction.

On the other hand, decisions that are embedded with other attitudes should be considered more carefully. For these, influencers should understand how the set of attitudes are interlinked and why – in order to ascertain the strength of the relationships. Once that is complete, a multi-stage marketing campaign would be a good option, where in each stage one or more of the ancillary attitudes are being focused on. Critically, though, the campaign must be committed to a long term endeavor – since as we saw individuals may pool in a state that is not the end state. However, if a multi-stage campaign can be setup, after you switch to the uniform strategy, there will be a strong swing toward the end state.

If the campaign has a shorter term goal and measures success by individuals in the end state, then perhaps a uniform strategy is the best way to go – in the short term this would outperform the other strategies.

6. FUTURE WORK

There are several paths for future work.

One interesting avenue of exploration is the relative effect of agenda setting as compared to choosing influential agents. Several algorithms exist for identifying influential agents in the influence maximization literature ([14]). We could test the effect of influential agents, agenda setting and the combined interaction.

We have focused our study on regular and small world networks. Another major class of networks are the scale-free graphs. Vertices in a scale free graph have a significantly varied degree distribution, and it would be interesting to see how that effects the effectiveness of agenda setting.

We have focused on the "3-Fan" network, but it's clear there may be many other network structures. In general, identifying these networks of attitudes in a general manner may be difficult. Networks of different attitudes linking together have been created for specific domains, for instance in the health decision making processes[30]. More complex cognitive network topologies may dramatically increase diffusion time, or preclude it (for instance, in the case where there are no consistent states).

We could also consider using analysis techniques derived from epidemic models. Considering the Markov Chain in Figure 8, we can consider agents in state s' as being susceptible, agents in levels 2 and 3 to be "exposed", and agents in state s^* to be infected (and to forever remain so). The goal would be to see how long before the population gets fully infected.

7. ACKNOWLEDGEMENTS

We thank XXXXXXXXXX, XXXXXXX, XXXXXXXX, XXXXXXX for helpful discussion. 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.

8. REFERENCES

[1] Data dive: Us tv ad spend and influence (updated – q2 2013 data), October 2013.

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

[3] J. Bednar, A. Bramson, A. Jones-Rooy, and S. Page. Conformity, consistency, and cultural heterogeneity. Mimeo, University of Michigan, 2006.

[4] C. Castellano, D. Vilone, and A. Vespignani. Incomplete ordering of the voter model on small-world networks. *EPL (Europhysics Letters)*, 63(1):153, July 2003.

[5] P. Chen and S. Redner. Majority rule dynamics in finite dimensions. *Physical Review E*, 71, 2005.

[6] W. Chen, Y. Yuan, and L. Zhang. Scalable influence maximization in social networks under the linear threshold model. In *Proceedings of the 10th IEEE International Conference on Data Mining (ICDM'2010)*, 2010.

[7] G. Csardi and T. Nepusz. The igraph software package for complex network research. *InterJournal*, Complex Systems:1695, 2006.

[8] R. Dechter. From local to global consistency. *Artificial Intelligence*, 55(1):87–107, May 1992.

[9] A. H. Eagly and S. Chaiken. *The Psychology of Attitudes*. Wadsworth Cengage Learning, 1993.

[10] R. H. Fazio and M. A. Olson. Attitudes: Foundations, functions, and consequences. In M. A. Hogg and J. Cooper, editors, *The Sage Handbook of Social Psychology*. Sage, 2003.

[11] C. M. Grinstead and J. L. Snell. *Introduction to Probability*. American Mathematical Soc., Jan. 1997.

[12] D. M. Gromet, H. Kunreuther, and R. P. Larrick. Political ideology affects energy-efficiency attitudes

and choices. *Proceedings of the National Academy of Sciences*, page 201218453, Apr. 2013. PMID: 23630266.

[13] C. R. M. Jones and R. H. Fazio. Associative strength and consumer choice behavior. In C. P. Haugtvedt, P. M. Herr, and F. R. Kardes, editors, *Handbook of Consumer Psychology*, pages 437–459. Psychology Press, 2008.

[14] D. Kempe, J. Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of 2003 SIGKDD*, 2003.

[15] J. A. Krosnick. The role of attitude importance in social evaluation: A study of policy preferences, presidential candidate evaluation, and voting behavior. *Journal of Personality and Social Psychology*, 55(2):196–210, 1988.

[16] Z. Kunda and P. Thagard. Forming impressions from stereotypes, traits, and behaviors: A parallel-constraint-satisfaction theory. *Psychological Review*, 103(2):284–308, 1996.

[17] F. Lima, A. Sousa, and M. Sumour. Majority-vote on directed erdos-renyi random graphs. *Physica A*, 387:3503–3510, 2008.

[18] M. Maghami and G. Sukthankar. Identifying influential agents for advertising in multi-agent markets. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, AAMAS '12, pages 687–694, Richland, SC, 2012. International Foundation for Autonomous Agents and Multiagent Systems.

[19] N. Masuda, N. Gibert, and S. Redner. Heterogeneous voter models. *Physical Review E*, 82(1):010103, July 2010.

[20] M. McCombs. A look at agenda-setting: past, present and future. *Journalism Studies*, 6(4):543–557, 2005.

[21] D. J. O'Keefe. *Persuasion: theory & research*. Sage, Thousand Oaks, Calif. [u.a., 2002.

[22] S. E. Page, L. M. Sander, and C. M. Schneider-Mizell. Conformity and dissonance in generalized voter model. *Journal of Statistical Physics*, 128:1279–1287, 2007.

[23] L. F. C. Pereira and F. G. B. Moreira. Majority-vote model on random graphs. *Physical Review E*, 71(016123), 2005.

[24] L. R. Peres and J. F. Fontanari. The media effect in axelrod's model explained. *EPL (Europhysics Letters)*, 96(3):38004, Nov. 2011.

[25] V. Rai and S. A. Robinson. Effective information channels for reducing costs of environmentally-friendly technologies: evidence from residential PV markets. *Environmental Research Letters*, 8(1):014044, Mar. 2013.

[26] E. M. Rogers. *Diffusion of Innovations*. Free Press, fifth edition, 2003.

[27] S. M. Ross. *Introduction to probability models*. Academic Press, Amsterdam; Boston, 2010.

[28] 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.

[29] P. Shakarian, S. Eyre, and D. Paulo. A scalable heuristic for viral marketing under the tipping model. arXiv e-print 1309.2963, Sept. 2013.

[30] Y. Shoda and W. Mischel. Personality as a stable cognitive-affective activation network: characteristic patterns of behavior variation emerge from a stable personality structure. In S. J. Read and L. C. Miller, editors, *Connectionist Models of Social Reasoning and Social Behavior*. Lawrence Erlbaum, 1998.

[31] D. Simon. A third view of the black box: Cognitive coherence in legal decision making. *University of Chicago Law Review*, 71(511), 2004.

[32] D. Simon and K. J. Holyoak. Structural dynamics of cognition: From consistency theories to constraint satisfaction. *Personality and Social Psychology Review*, 6(6):283–294, 2002.

[33] D. Simon, D. C. Krawczyk, A. Bleicher, and K. J. Holyoak. The transience of constructed preferences. *Journal of Behavioral Decision Making*, 21(1):1–14, 2008.

[34] D. Simon, D. C. Krawczyk, and K. J. Holyoak. Construction of preferences by constraint satisfaction. *Psychological Science*, 15(5):331, 2004.

[35] 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.

[36] V. Sood, T. Antal, and S. Redner. Voter models on heterogenous networks. *Physical Review E*, 77, 2008.

[37] 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.

[38] K. Suchecki, V. M. Eguiluz, and M. S. Miguel. Voter model dynamics in complex networks: Role of dimensionality, disorder and degree distribution. *ArXiv*, 2008. cond-mat/0504482v1.

[39] P. Thagard. *Coherence in Thought and Action*. Life and Mind: Philosophical Issues in Biology and Psychology. MIT Press, 2000.

[40] P. S. Visser and L. M. Clark. Attitudes. In A. Kuper and J. Kuper, editors, *Social Science Encyclopedia*. Routledge, 2003.

[41] P. S. Visser and J. Cooper. Attitude change. In M. Hogg and J. Cooper, editors, *Sage Handbook of Social Psychology*. Sage Publications, 2003.

[42] 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.

[43] D. J. Watts and S. H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, June 1998.