Belief Consolidation and Homophily in Networks

EECS 372 Spring 2013
Jacob Kelter

Big Picture Overview

There are certain beliefs and pieces of common knowledge that seem to be universal in a particular culture. Everyone (almost) in the United States “knows” that the United States won its independence from the British in the Revolutionary War. No one alive today was alive to witness this piece of historical knowledge, but we have historical records and it has been passed down through the generations.

The vast majority of things we believe come from sources other than direct experience. Knowledge and beliefs spread through populations and are passed down through generations. For any event witnessed by multiple people, there will be multiple accounts of what happened; yet somehow a common knowledge of what “actually” occurred eventually seems to be consolidated in the population.

Often times, there are multiple differing beliefs on a topic within a population. Under some conditions, these beliefs will be consolidated into a single belief shared by the whole populations, and under other conditions the beliefs will remain distinct. When multiple beliefs exist within a population, such as different political or religious beliefs, people tend to have many more friends that are similar to them in these areas than different.[1] This tendency to have connections to similar people is termed homophily.

This model simulates how communal knowledge and beliefs about an event emerge from the differing accounts of multiple individuals, how multiple beliefs are either maintained in a population or converge to single group belief and how people with similar beliefs end up maintaining bonds with each much longer than people with divergent beliefs.

What Can Be Learned

The intention of this model is not to fully explain the complex cultural phenomenon of how communal knowledge forms. Rather, the intention of the model is to serve as a focal point for discussion and debate about how collective knowledge might form and evolve and about what mechanisms cause homophily in a network.

That being said, there are some specific questions this model attempts to answer.

1. What effect does noisy transmission of beliefs have on communal beliefs?
2. Do individuals or groups transmit beliefs to large population more accurately?
3. Under what conditions can multiple beliefs be maintained in a population, and under what conditions do beliefs get consolidated into a single group belief?
4. Under what conditions can a small group convince a larger population of a new belief?
5. Can the rules of this model create homophilous networks?

A section of analysis will be devoted to each of these questions.

Motivation and Rationale

Although this model could be used to simulate the spread and consolidation of many types of beliefs, the main motivation is to model how religious beliefs might have spread in relatively small populations in the distant past.

A network structure for agent interactions is the only reasonable way to model the dynamics of collective beliefs, because we know how influential social networks are in influencing beliefs and behaviors.

Modeling communal belief dynamics in a computational agent-based environment could potentially lead to new insights, because dynamics of the interactions between heterogeneous individuals can be observed. Also, the assumptions of how people transfer beliefs will be made explicit, allowing for discussion around how these processes really happen. This model is clearly very simplified, but as an object to think with it could be useful in developing theories about how beliefs spread in populations.

Model Implementation

The model consists of a network in which nodes represent people and links represent friendships (or any other relationship). Individuals have a belief represented by a list of three numbers in the range 1-255 that is displayed as a color, a stubbornness in the range 0 to 1 and a tolerance in the range 0 to 1. Figure 1 shows a sample visualization of the model.

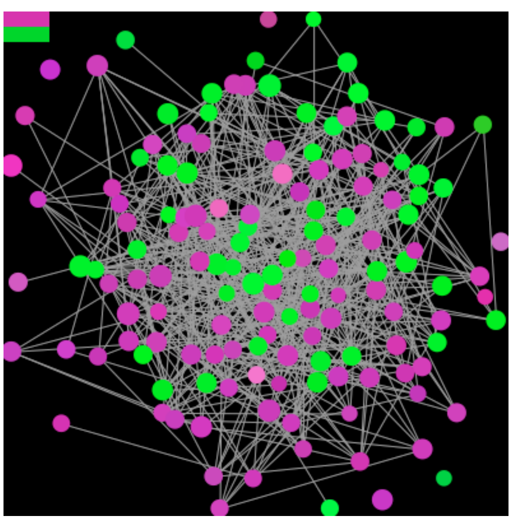

Figure 1: A sample visualization from the model. This population has two distinct beliefs.

Stubbornness represents how willing an agent is to change beliefs. A stubbornness of 1 means the person never changes beliefs (at least not knowingly), and a stubbornness of 0 means that the person fully adopts the belief of anyone she interacts with; a stubbornness of .5 means the person averages her belief with anyone she interacts with. Stubbornness is displayed by the size of the circle representing the person; a larger circle means more stubborn. People with maximal stubbornness (i.e. 1) are visualized as a square to set them off visually, because they often play a very important role in the model.

Tolerance represents how willing the person is to exchange beliefs with someone in the first place. A person with a tolerance of 1 will exchange beliefs with anybody and a person of tolerance of 0 will only "exchange" beliefs with someone that holds exactly the same belief. When exchanging beliefs, each person still weights her own belief with her stubbornness and the other belief with (1- her stubbornness).

Friendships have a friendship strength property. If friendship strength ever falls below 0, the friendship ends.

The rules agents follow at each tick are given in 1st person:

- I choose someone I'm linked with to interact with.
- We see how different our beliefs are. This is calculated as a fraction for each of the three numbers in the belief, and then these fractions are averaged. So if our beliefs are maximally different (all 1s and all 255s), the difference between our beliefs will be 1 and if they are exactly the same, it will be 0.

- If both our tolerances are greater than the differences between our beliefs, we share beliefs. If not, we stop interacting.
    - To share our beliefs, we each use a weighted average according to our stubbornness. For example, if my stubborness is 0.8, I will multiply each component of my belief by 0.8 and my friend's belief by 0.2 and add them component wise. Noise is added either to my friend's belief before I average them, or to the total after I average them.
    - After we share our beliefs, our friendship is strengthened by a friendship gain factor.
- Next I make random friends.
    - With a certain probability, I make a new friend with a random other person.
    - With a probability three times higher than this, I make a new friend with someone who is a friend of one of my friends.
    - If I have no friends, I always try to make a random new friend.

After all the individuals have gone through the above procedures, all friendships are weakened by a friendship decay factor. If the strength of a friendship reaches 0, the friendship ends.

The Effect of Noise on Communal Beliefs

Noise is added to the model in a few different ways. The model is initialized with "The Event", and a certain number of people have knowledge of this event. However their own beliefs about the "The Event" have noise added to them to represent subjective experience, so they all have somewhat different beliefs about “The Event.” Then whenever people pass their belief on to someone else who doesn’t yet have a belief, noise is added in the transmission. This is to reflect imperfect transmission of beliefs.

When two people who already have beliefs share them, there are two different ways noise is added. If I am one of the people, either the noise is added to the other person’s belief and then I average her belief with mine, or we average our beliefs and then noise is added. This seemingly trivial difference has a large effect on the model’s output for two reasons. First, the noise is stronger when it is added after the beliefs are averaged, because it is not multiplied by the averaging factor. Second, and more importantly, completely stubborn people never change their beliefs if noise is added first to the other person’s belief and then averaged (because the other person’s belief gets multiplied by 0). But, if noise is added after the beliefs are averaged, then stubborn people’s beliefs still change. Either of these scenarios could make sense depending on the situation being modeled. Even very stubborn people will change beliefs over time, even if they are never influenced by others (if nothing else senility might change their beliefs eventually). On the other hand, for certain beliefs it would make sense for stubborn people to never change. For the remainder of this paper, the method of adding noise is “other belief” (noise added to the other person’s belief and then the two beliefs are averaged) unless otherwise noted.

Noise in a community with a single belief widens the distribution of beliefs. Figure 2 below shows the distribution of beliefs for each belief component when noise is set to 15 and when noise is set to 5 for the same world. The average belief for each component is the same for each situation, but the belief distribution has a higher standard deviation when the noise is higher, i.e. there is a wider spread of beliefs.

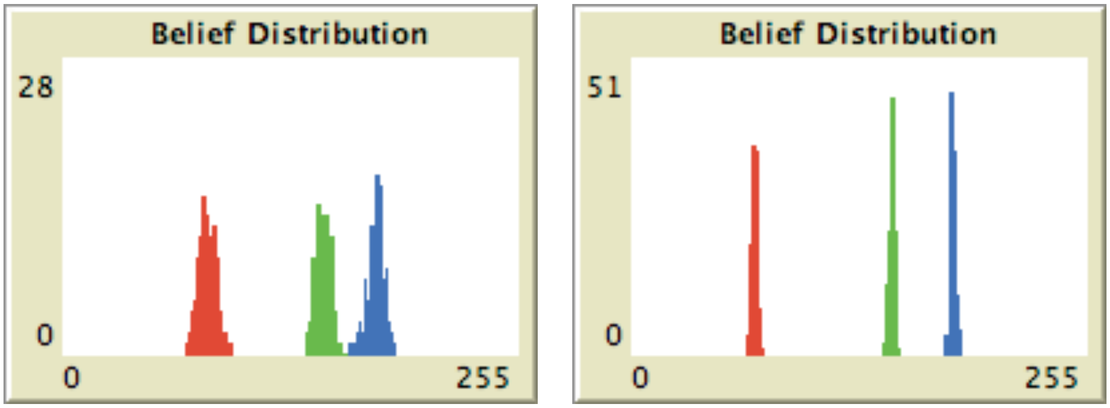

Figure 2: Belief distribution when noise is added to "other belief." (left) Noise=15 (right) Noise = 5.

Figure 3 shows the average of the standard deviations of the distributions of each component of the belief for different amounts of noise. The left figure is for noise being added to the "average belief" and the right figure for noise being added to the "other belief" as defined above. In both cases, the standard deviation increases fairly linearly, but the noise has a greater effect when added to the "average belief", because it doesn't get multiplied by the averaging factor (1-stubbornness). For the same reason, the stubbornness of the population doesn't change the effect of noise when it is added to the "averaged belief." However, when noise is added to the "other belief," increased stubbornness leads to a decrease in the standard deviation of the beliefs. This is because with a higher stubbornness, people listen less to the noisy beliefs of others. An average stubbornness of 0.7 and an average tolerance of 0.3 were used for all runs in Figure 3.

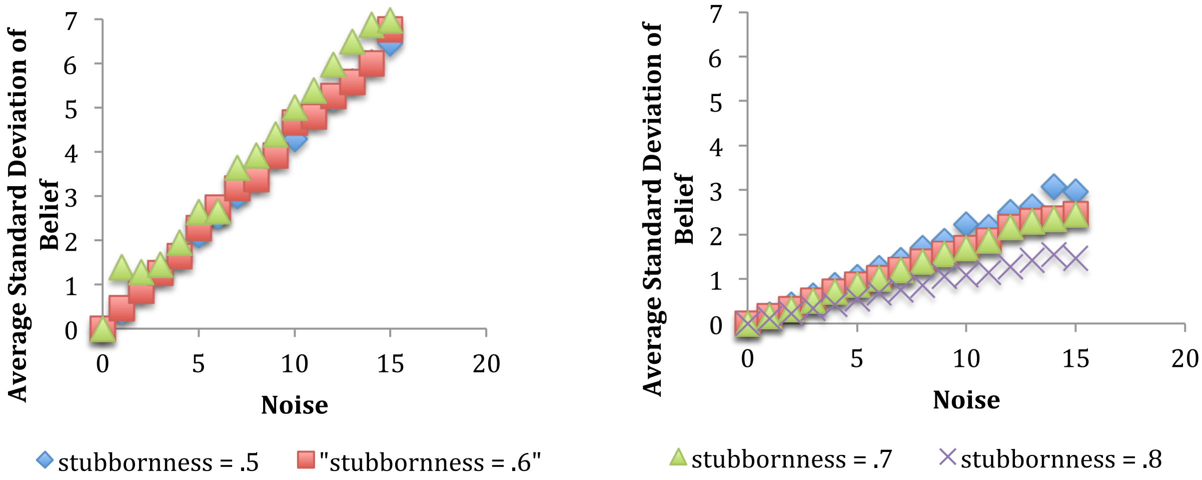

Figure 3: The average standard deviation of the components of the communal belief with increasing noise. The population was only initialized with one belief. (Left) Noise was added to the "averaged belief". (Right) Noise was added to the "other belief."

Analysis

Individuals vs. Small Groups Transmitting Beliefs to a Large Population

Common sense suggests that if a large group of people witnesses an event, they will remember it better than a single person alone. But, is a larger group always better? Figure 4 shows the percent disagreement between the average belief of a population and the initial event after 100 ticks, given a certain number of people originally witnessing the event.

The original witnesses of the event are still subject to noise, so a group of people initially witnessing the event will all "remember" it slightly differently. It is clear from Figure 4, that in the model, when less than 10 people witness the event, the percent disagreement increases, and also becomes more erratic. When around 10 or more people witness an event, the percent disagreement seems to level off; this probably means that with more than 10 people witnessing an event, their different initial experiences average out as they spread across the population, leading to a more accurate memory of the real initial event. As expected, the percent disagreement with the initial event increases with increasing noise levels. An average stubbornness of 0.7 and an average tolerance of 0.3 were used for all runs with a population of 150.

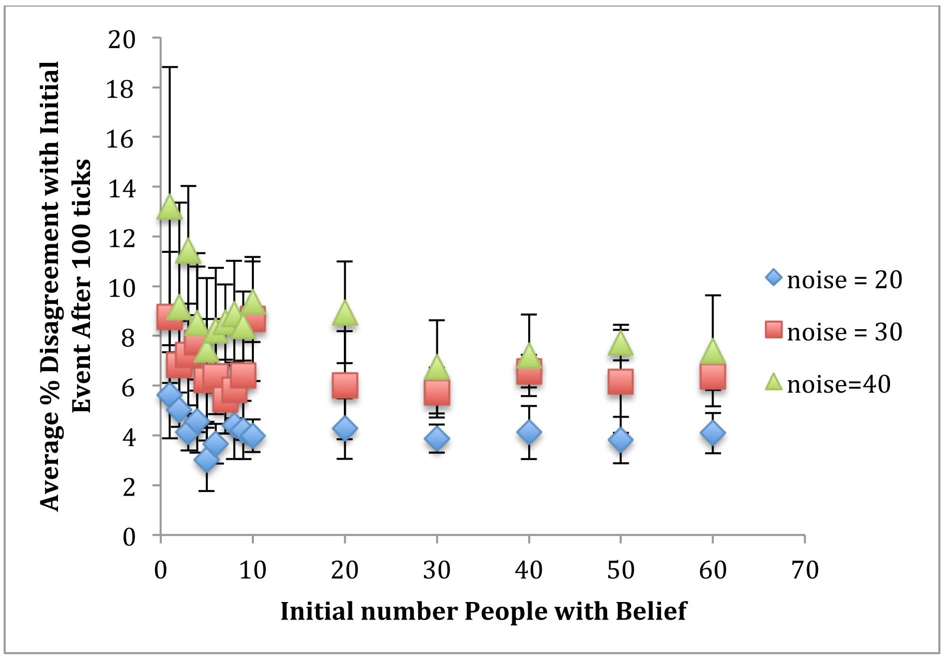

Figure 4: The average disagreement of the population with the group after 100 ticks given a certain number of people starting with the belief and spreading it. The group population was 150.

Conditions for Multiple Beliefs

The tolerance of the population is the main factor determining whether the population maintains multiple beliefs or converges to a single belief. If there are two beliefs in the population, and the members of each belief group are tolerant enough to share beliefs, eventually their beliefs will average and the group will converge to a single belief. If, however, at least one of the populations is not tolerant enough to share beliefs with the other group, both beliefs will be maintained in the population. Figure 5a shows the visualization of a network with two beliefs each held by ten people. The average stubbornness of the entire population was set to 0.7 and the noise was set to 15. In Figure 5b, the average tolerance was set to 0.15 and the network was allowed to evolve for 2000 ticks; the difference between the beliefs was great enough, that they were both maintained in the network. The belief distribution beneath the network in Figure 5b clearly shows two different peak values for each belief component. In Figure 5c, the same network was evolved, but with an average tolerance of 0.23; after around 1000 ticks, the network had converged to a single belief somewhere between the original two. The belief distribution beneath the network in Figure 5c clearly shows a single peak value for each belief component.

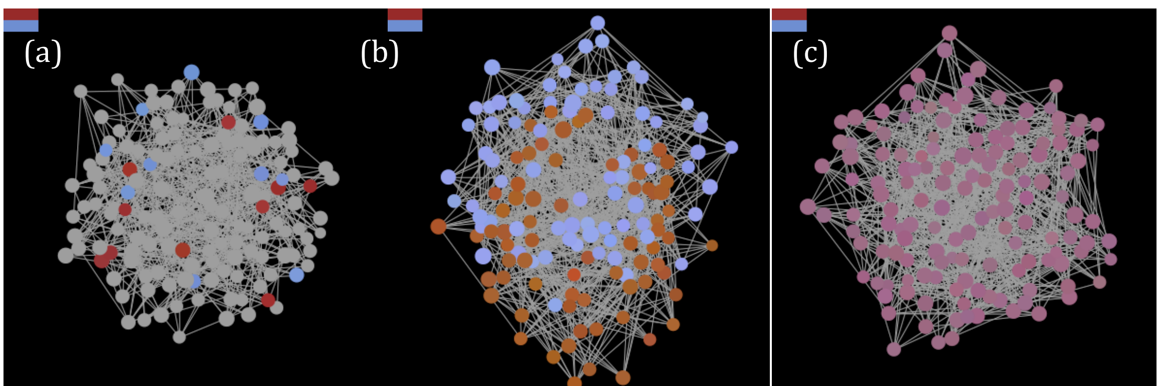

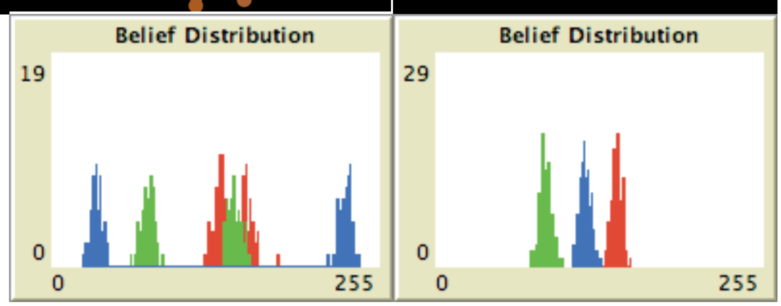

Figure 5: The same initial network (a) was evolved forward with:
(b) An average tolerance of 0.15 resulting in 2 beliefs
(c) An average tolerance of 0.23 resulting in 1 belief.

Small Group Convincing a Large Group

As the previous section showed, a population with multiple beliefs will converge to one belief if the average tolerance is high enough. This means that for a small group to convince a large group of a new belief, both groups will have to be tolerant enough to interact with one another. In addition, the smaller group will have to be much more stubborn than the large group on average, otherwise they will end up adopting the majority opinion instead of convincing the majority of a new belief. Figure 6 shows a network with a majority group of 140 turtles that have an average stubbornness of 0.55 and an average tolerance of 0.5 along with 10 minority turtles that have an average stubbornness of .9 and an average tolerance of 0.5. Since both groups are tolerant enough to interact, but the minority group is so much more stubborn, the whole population eventually converges to the minority belief.

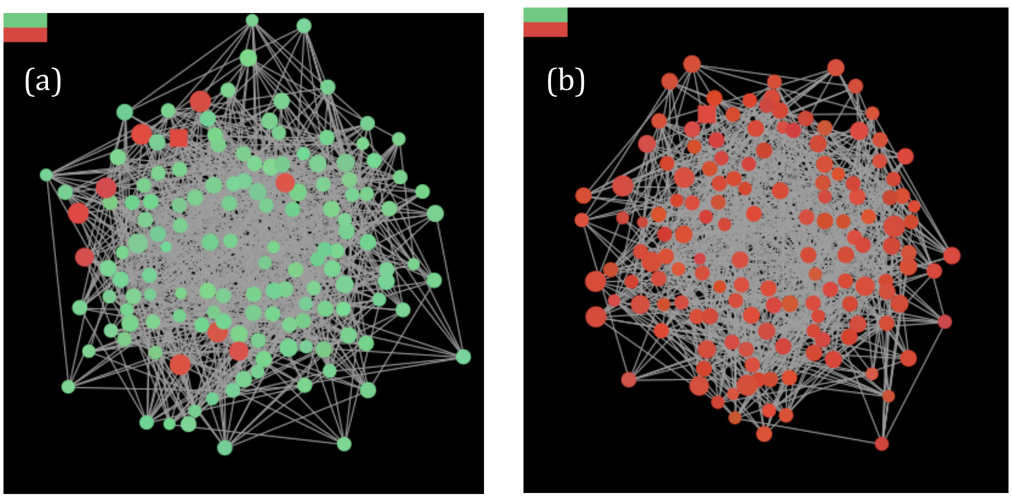

Figure 6: (a) A network with 140 people holding one belief and 10 people holding another. All the people have a tolerance of 0.50, but the majority nodes only have an average stubbornness of 0.55, while the minority nodes have an average stubbornness of 0.9. (b) The same network after being evolved for around 1000 ticks. Because both groups were tolerant enough to interact and the minority group was so stubborn, the entire population adopted the minority belief.

Network Structure

A detailed analysis of the network structure resulting from the model will not be undertaken here. However, it is worth noting two things. First, homophilous networks are definitely possible in the model, and second the network has a normal degree distribution. By changing the friendship gain and friendship decay parameters and the probability of making new friends on each tick, this degree distribution could be altered. It is also possible that a broader distribution could be attained by giving individuals their own personal friendship gain and decay parameters. Figure 7 shows a visualization of a network with two beliefs and largely homophilous friendships, along with the degree distribution for this network.

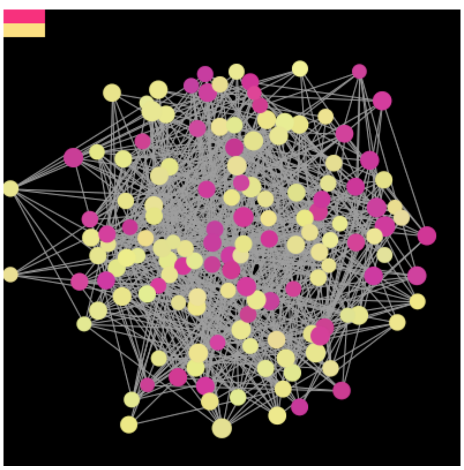

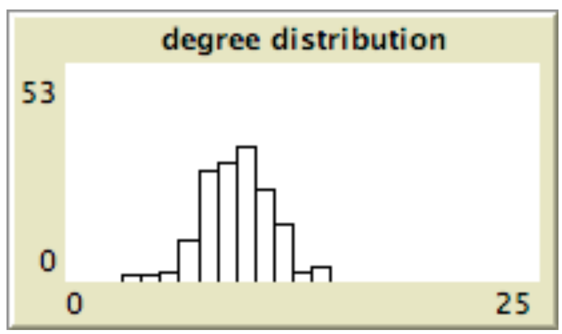

Figure 7: A homophilous network generated by the model and its degree distribution

Conclusion

Using some very basic rules of belief sharing, link formation and strengthening and link decay, this model gives some insight into how beliefs might spread on a network, how multiple beliefs can be maintained in a network, how multiple beliefs converge to single a belief and one mechanism of how homophily might occur. Future work could include trying to incorporate both written and oral transmission of beliefs. The shift from purely oral historical traditions to written historical traditions was surely an extremely important change in the way humans passed on beliefs and knowledge and it would be fascinating to try to make an agent based model of this phenomenon.

References

1) McPherson M., Smith-Lovin L., Cook, J. 2001. Birds of a Feather: Homophily in Social Networks. *Anu. Rev. Socio.*