The Rumor Mill:
A Preliminary Agent-Based Model of Rumor Spreading

Nate Neligh & Matt L. Miller

Concept

In this project, we model the spread of a rumor within a simple social network. Within this network are agents who are personally concerned with the rumor's spread: some desire that the rumor be accepted as true ("spreaders") and others desiring that the rumor be discounted as false ("quashers"). We begin the modeling process by considering the essential elements of a rumor as compared to other social ideas (political opinions, for example). We consider the following three elements to be the most essential:

1. Initial information: A person who does not know about a rumor cannot act as to spread or quash that rumor.
2. Belief: When a person hears a rumor, they can either accept it as true or discount it as false. This belief can change over time as new information about the rumor is acquired. In the domain of rumors, the source of information is generally social interaction.
3. Activity: A person who hears a rumor may talk about it a great deal, but that person may also simply not care to bring it up again. If people have been discussing a rumor with an individual frequently, that individual may be more likely to bring the rumor up in future conversations.

Model

Individuals within our model have a number of traits. They have an information state: they either know or are ignorant of the rumor depending on whether they have yet had an interaction in which the rumor was brought up. An individual also has a belief level between -1 and 1 that reflects their current stance on the rumor with 1 reflecting complete belief and -1 reflecting complete disbelief; for individual that are ignorant of the rumor, this level reflects the degree of belief the individual would assign to the rumor upon first hearing it. Lastly, the individual has an activity level between 0 and 1 corresponding to the probability that they will bring up the issue in an interaction. Individuals are assigned a regular pattern of eight equally-weighted social connections in a regular pattern; this is modeled by placing the individuals on a discrete, toroidal grid with the nearest neighbors representing each individual's social network.

The system evolves through random pairwise interaction. At each model time step, one individual is selected stochastically; one of that individual's neighbors is then selected randomly as the partner for the interaction. This interaction proceeds as follows. If neither individuals knows the rumor, the interaction does not proceed and both individuals' activity level decays (see below). If one or both of the individuals know the rumor, the probability that the interaction proceeds is equal to the activity level of the first individual (A_1) plus the activity level of their neighbor (A_2) minus the product of their activity levels: $A_1 + A_2 - A_1A_2$. If the interaction does not proceed (modeling the situation in which the rumor is not discussed), the activity level for both individuals decays by a constant fraction, the activity decay rate. If the interaction proceeds, modeling the rumor being discussed, the activity levels of both individuals to a new level. A_1' is the new activity rate of individual 1:

$$A_1' = 1 - \frac{1}{\left(\frac{1}{1-A_1} + k\right)}$$

k is the activity increment. We chose this functional form due to its desired properties of asymptotically approaching 1, concavity, and smoothness, though other functional forms are possible in the absence of ready information on the time course of interest levels for rumors. Note that A_2' is updated in the same way, replacing A_2 for A_1.

When the interaction proceeds (the rumor is discussed), both parties mix their beliefs by updating their belief levels; B_1' being the new belief level for individual 1:

$$B_1' = \lambda B_2 + (1-\lambda)B_1$$

Again, the value for individual 2 is updated symmetrically.

Some empirical evidence suggests that the degree of mixing generally decreases proportional to the difference in belief; many people are less likely to modify their beliefs based on the opinions of thse with whom they strongly disagree. We model this by giving λ the functional form:

$$\lambda = \alpha \frac{1}{1-e^{-\gamma 2}}\left(e^{-\gamma|B_2-B_1|} - e^{-\gamma 2}\right)$$

Thus, the amount of mixing drops off exponentially with the distance between beliefs. Individuals with identical beliefs will have a mixing parameter of α. Individuals with beliefs that differ by 2, the maximum possible difference, will have no influence on one another. As a result, the influence an individual can have on another is single peaked with respect to the difference in beliefs. The rapidness of the decline in influence by difference in prior belief can be modified by changing the parameter γ ("influence steepness") with greater γ increasing the steepness—that is, influence falls off more rapidly with larger values of γ. The influence is symmetric unless one of the individuals is a quasher or a spreader. Quashers have belief level fixed at -1 and spreaders have their belief level fixed at 1, so their partners during the interaction cannot affect their belief level.

At the beginning of the simulation a population is generated on the grid. Each individual has a latent activity level and belief level drawn from a uniform distribution. There are also a number of quashers and spreaders randomly placed on the grid. At initialization, one spreader is given knowledge of the rumor and then the update process begins. By adjusting several parameters, we find a number of interesting results with applications to the real-world spread of information and rumors.

The following figure shows a typical starting point for latent beliefs of individuals regarding the truth of the rumor.

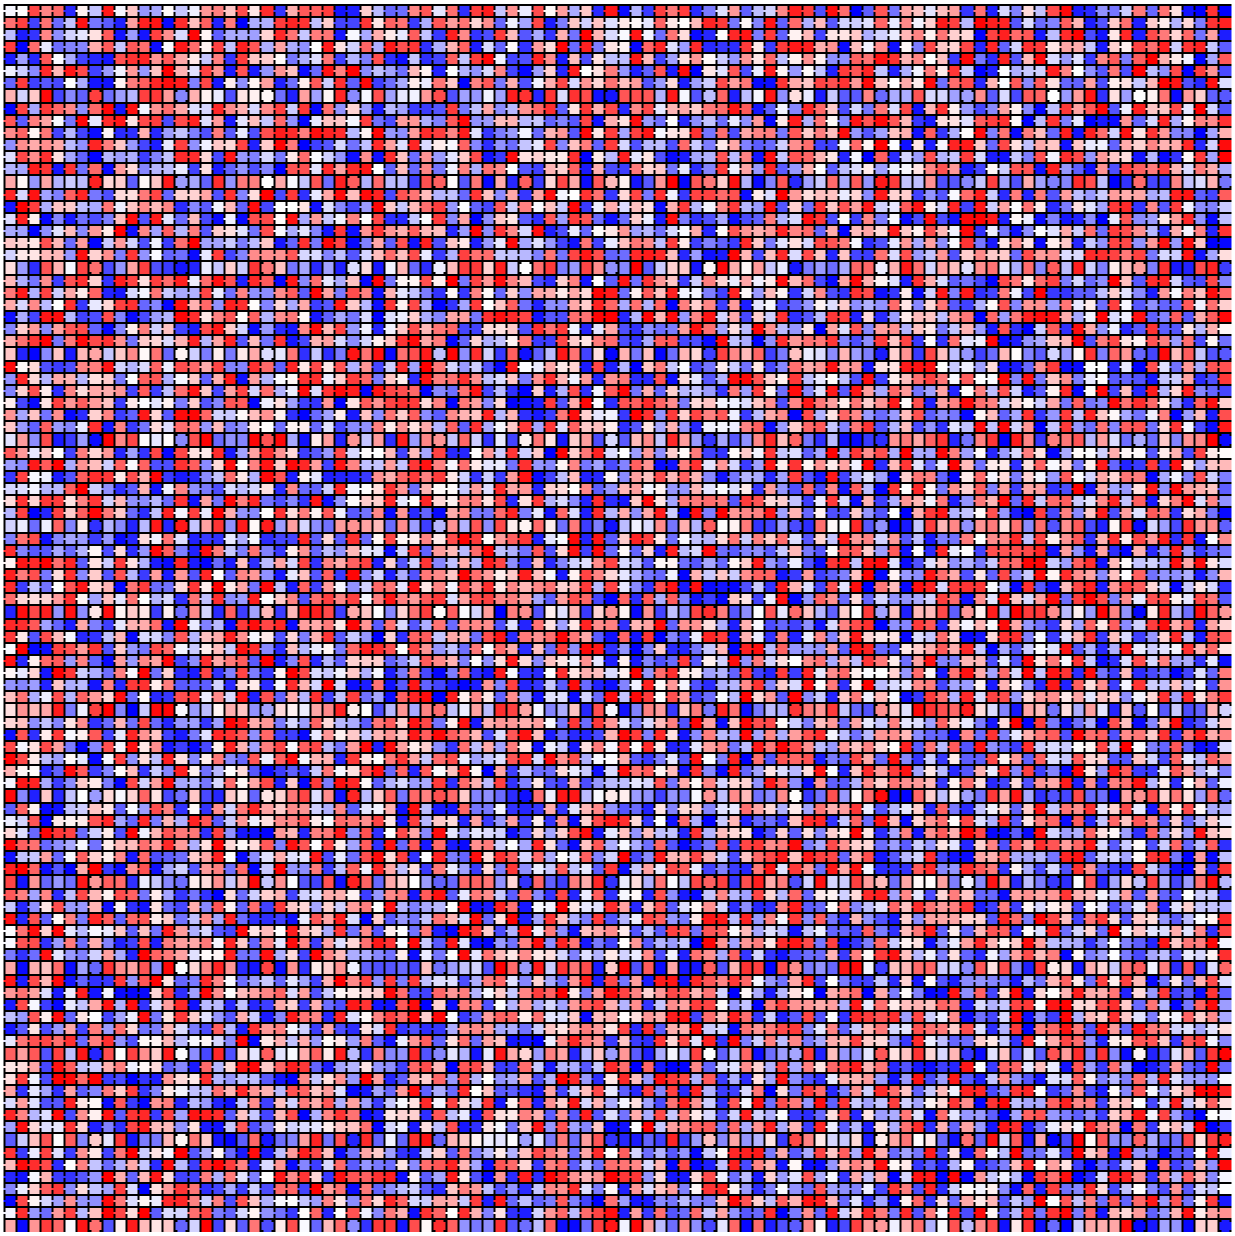

Intense red represents a stronger belief in the rumor. Intense blue points on the grid represent strong disbelief. Lighter colors represent more neutral beliefs

Results

Activity

One interesting result concerns the amount of rumor activity in the model and the number of individuals who have heard the rumors. The primary parameter of interest in the spread of activity is the activity decay level. One would expect that quickly decaying activity might lead to a rumor dying out while higher activity level might lead to sustained activity. These intuitions are supported, but the evolution of the activity level is surprising in a number of ways. To start with let's look at the evolution of the average activity under standard parameters. Standard parameters are:

Parameter	Value
Grid Height	50
Grid Width	50
Activity decay constant (1 – decay rate)	0.95
Activity increment (k)	0.10
Maximum influence (α)	1.00
Influence steepness (γ)	3.00
Initial number of quashers	25
Initial number of spreaders	25

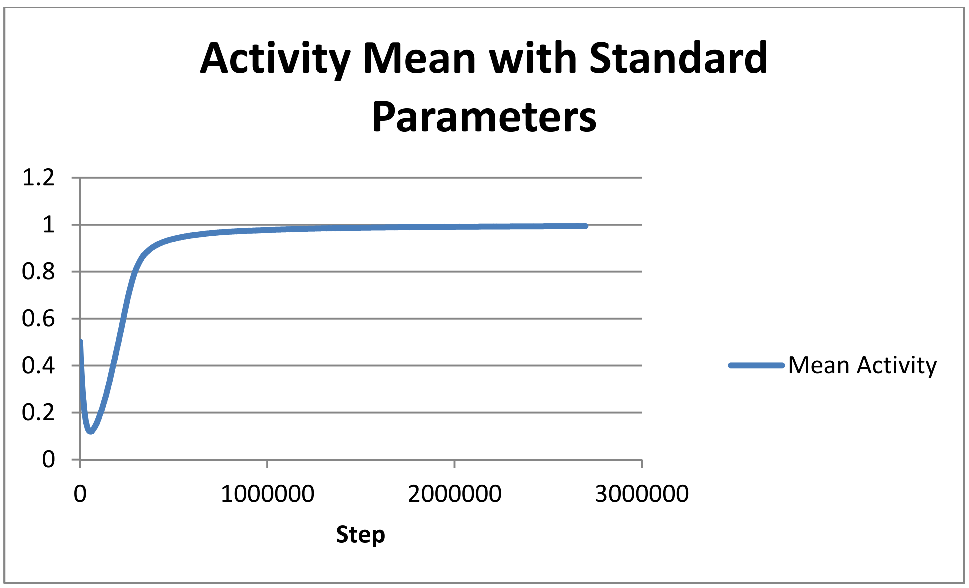

The activity level initially dips and then begins to rise as the activity becomes more spread throughout the population. In the initial phases, the knowledge is very concentrated and very few interactions bring the rumor up, because individuals are mostly not in contact with the few enthusiastic individuals who have heard the rumor. As the information spreads, more individuals becomes active and can spread activity to others; this activity feeds back to those who are enthusiastic.

More rapid activity decay (lower activity decay constant) makes the evolution of the mean activity much noisier. However, in both of the cases shown in the next figure, the knowledge of the rumor eventually reaches 100% of the population.

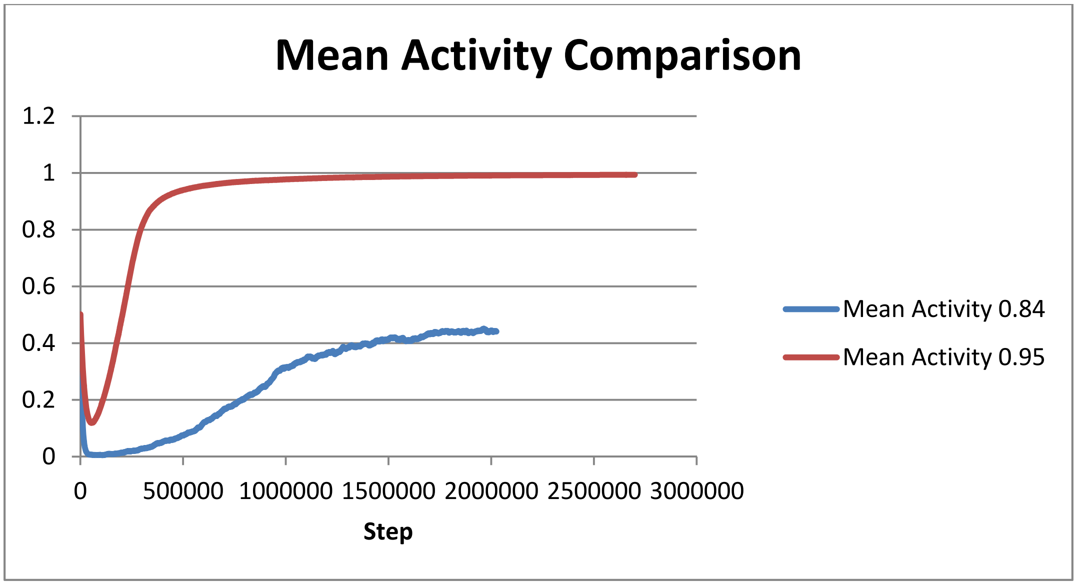

On the other hand, if the decay constant is low enough (indicating a higher decay rate), the activity can die out before all individuals learn the rumor, as occurs when the constant is set to 0.81. The activity level over time is shown in the next figure.

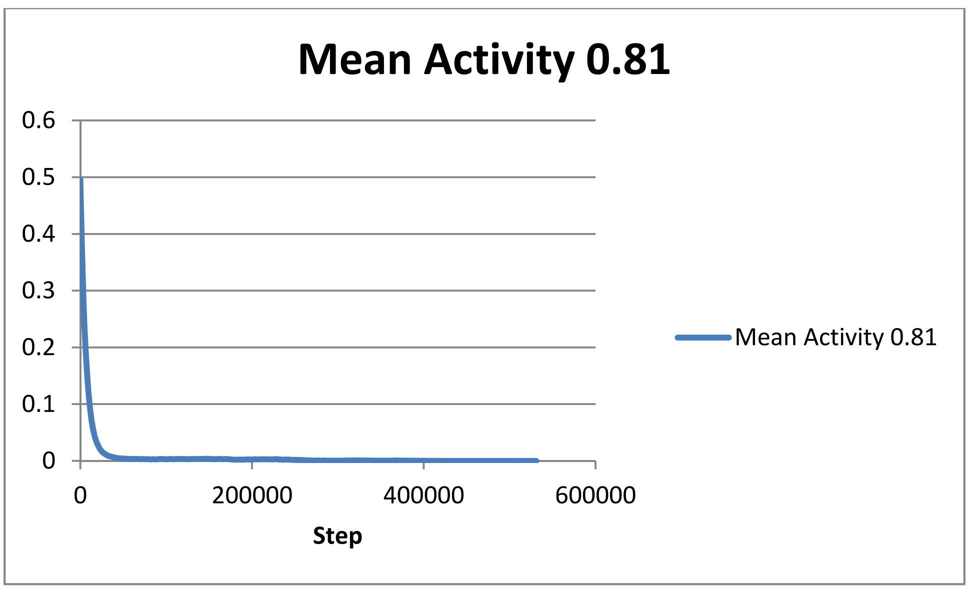

In this simulation, the proportion of individuals who know about the rumor reaches an asymptote at 0.078.

Between these two cases, there is a transition zone where the mean activity neither rebounds to a high level nor moves toward 0. In these cases, stochastic variation plays a large role, with mean activity level appearing to follow a random walk after its initial dip.

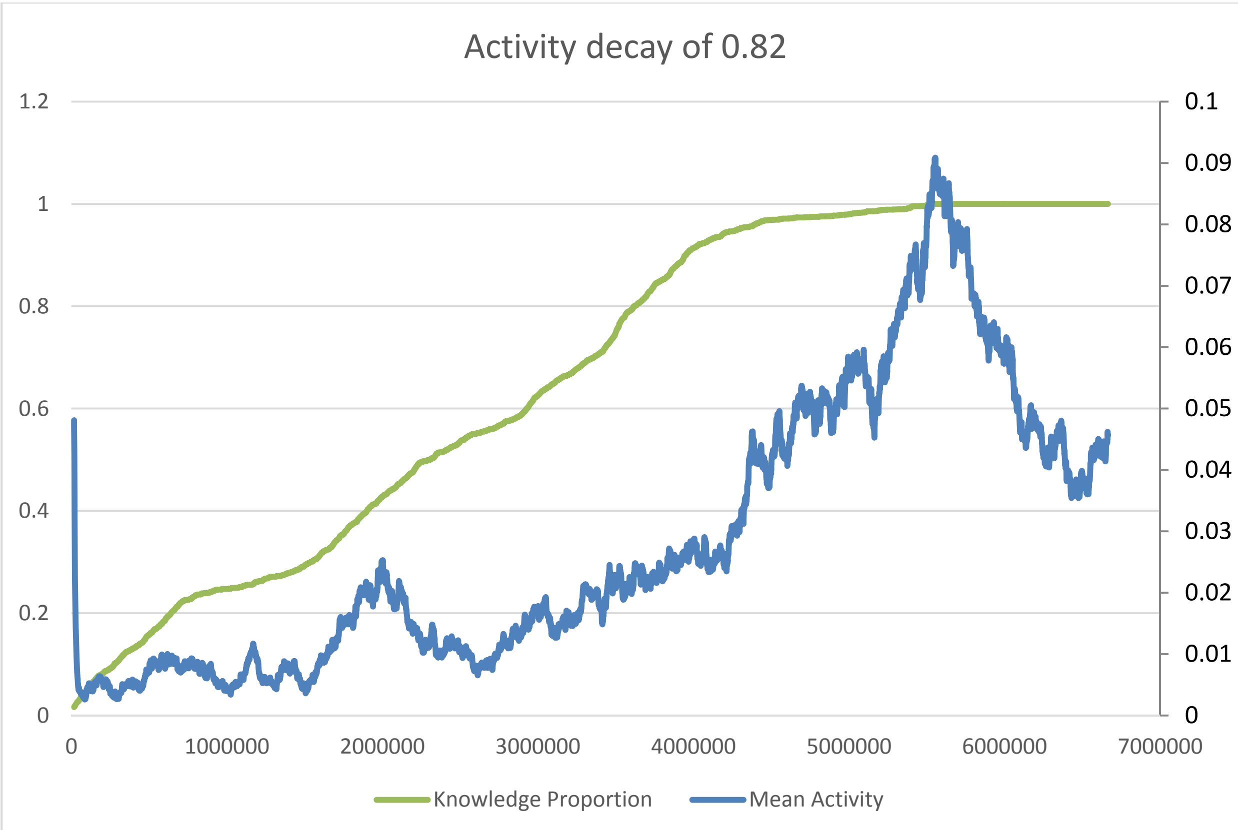

The proportion of individuals that know the rumor is also shown in the above figure (on the left axis, with activity shown on the right axis). Here, we see that the activity level does not reach an asymptote in over 6 million steps. This appears to be a random walk with high momentum, though several unusual features suggest that there may be other underlying processes. One thing is clear: the behavior of the activity level near this activity decay threshold is extremely noisy. It is important to note that when the activity level does not collapse, the knowledgeable proportion does reach 100%. More intriguingly, the activity level seems to peak right at the point where the knowledge level reaches 100%; this could be indicative of some deeper mechanism, an artifact, or simply coincidence emerging from the random walk. More research is required before comment is appropriate. Note that in the other cases, the knowledge proportion behaves as one would expect; to demonstrate this behavior we provide plots showing the activity levels from simulations with activity decay rates away from the threshold including the knowledge proportions.

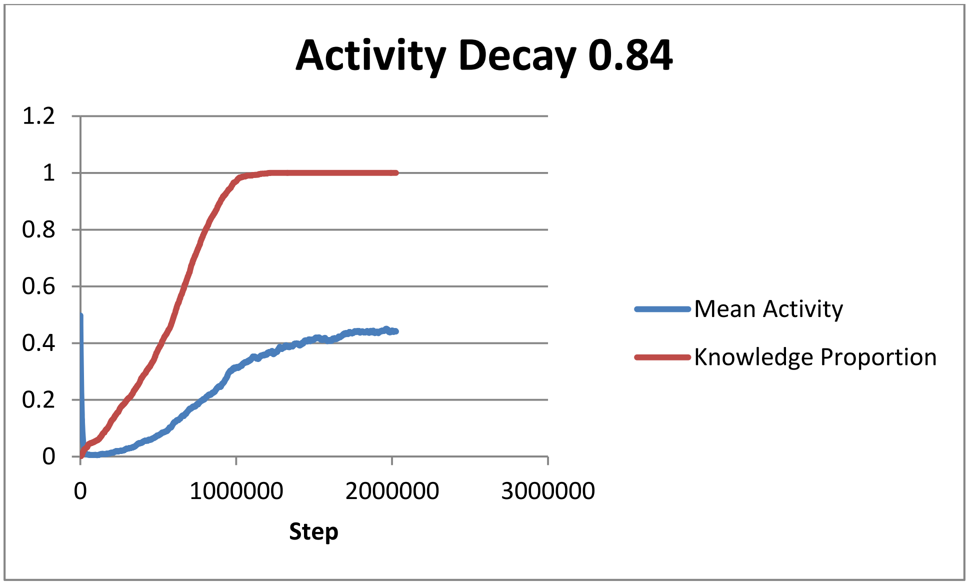

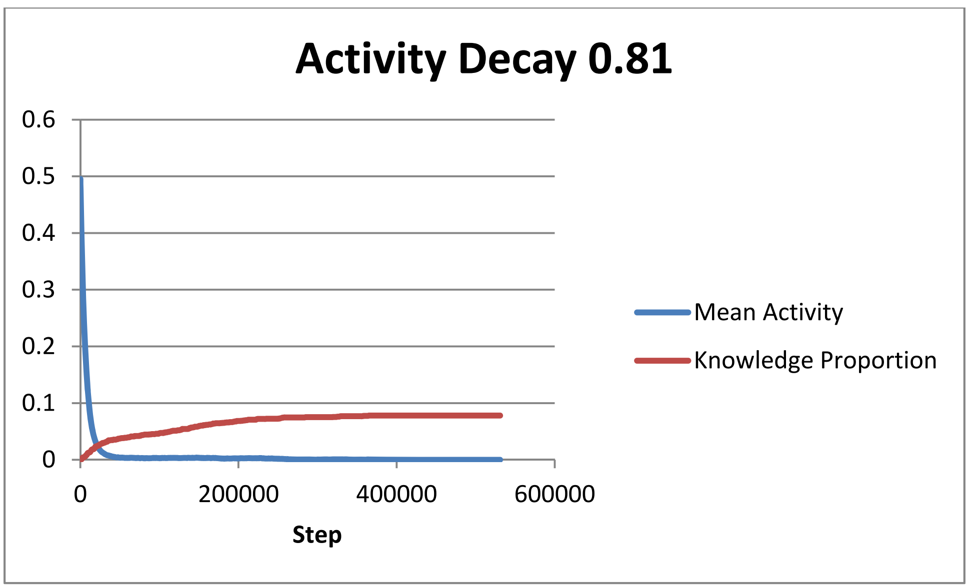

Consensus formation and enclaves

The spreaders and quashers of rumors are generally interested not only in how many individuals hear about the rumor and how active discussion of the rumor is but also about the beliefs people hold about the rumor once they have heard it. Note that though psychological research does suggest that mere exposure to a rumor tends to increase belief at least subconsciously, we do not model that process here, assuming that the only effect on belief level is pre-existing belief and the beliefs of those discussing the rumor with the individual. Because spreaders have access to the earliest influential interactions in the model due to one of them being in sole possession of the rumor, there should be a slight advantage for the spreaders in slow spreading activity as long as knowledge does eventually reach 100%. However, in practice there are almost always enough disbelievers near the start point of the rumor to make this advantage mostly negligible. When one belief does gain an advantage, it is usually due to stochastic variation in the starting conditions and in the ensuing interactions. Spreaders do not seem to have an inherent advantage over quashers. When results emerge, the form they take tends to depend a great deal on influence steepness (γ). Some of the results of this experiment echo results reported in the polarization and segregation literature, but there are also surprises due to the presence of dedicated quashers and spreaders, as well as due to the process of information and activity spread.

Under standard parameters, we eventually see the formation of what looks like a consensus. This is because there is fairly extensive communication influence between the mid-range individuals but less influence between the mid-belief individuals and the extremes. The following image shows a typical evolution of a process on a 100×100 grid with otherwise standard parameters.

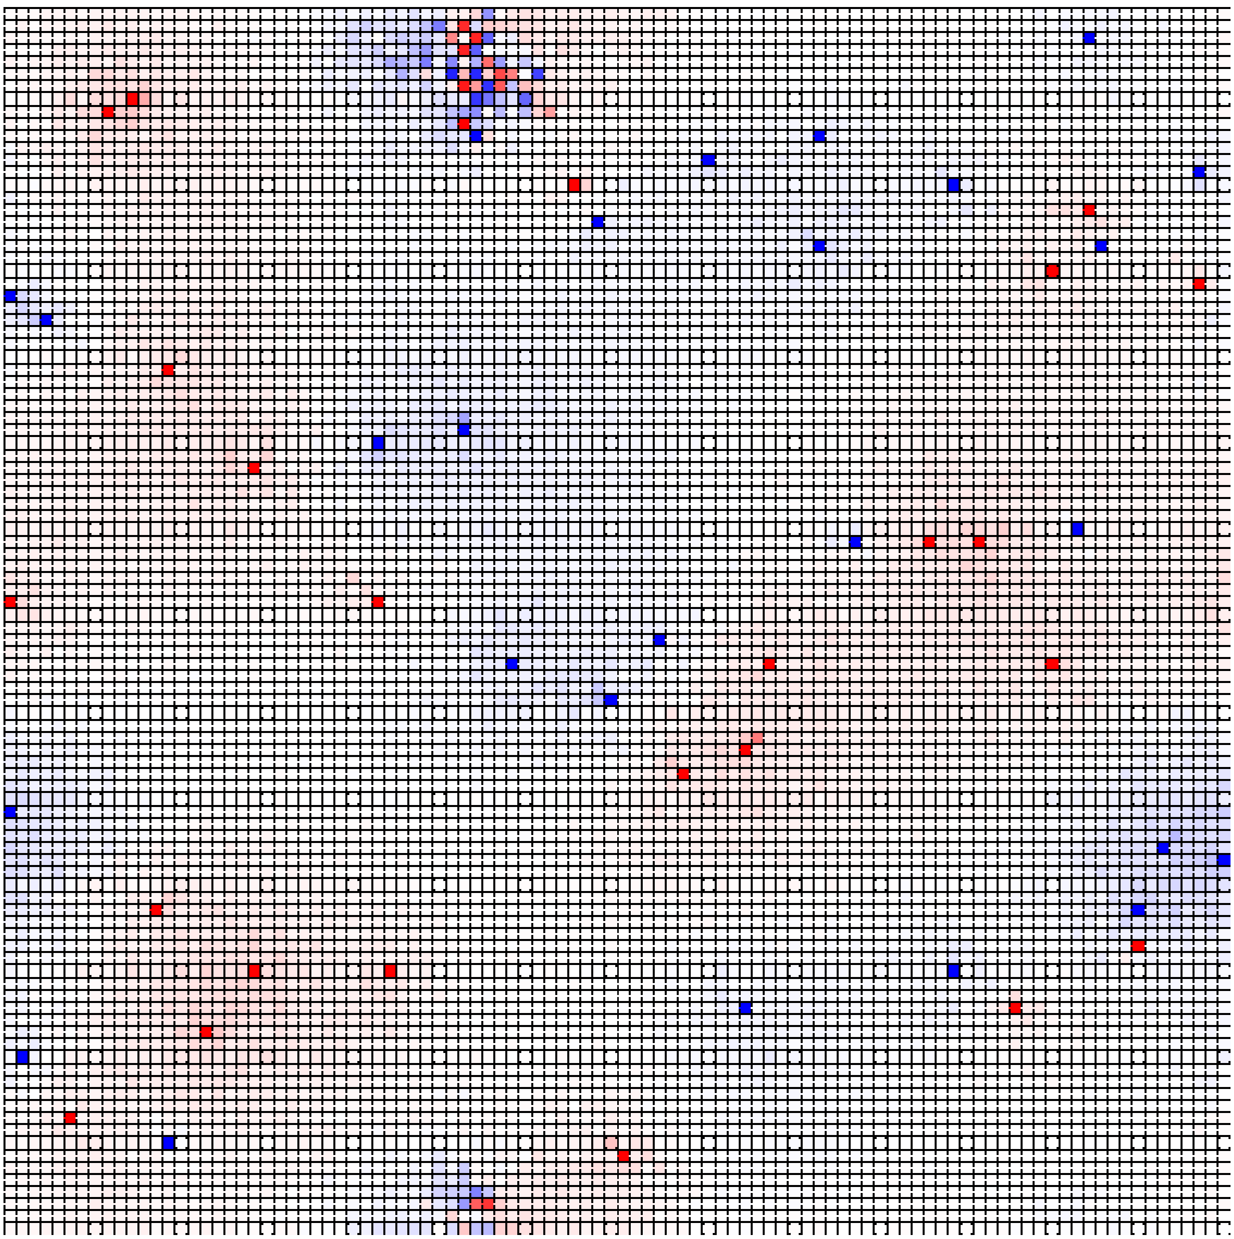

The quashers and spreaders tend to form small areas peaks and troughs of belief, but most of the area is nearly neutral. Since influence decreases as individuals are closer in belief and decreases as they are further apart, hysteresis emerges: after a node has been moved in one direction, they become easier to move farther in that direction.

The influence of a belief level is also rapidly distributed amongst the mid-range belief individuals. As a result, once the median individual is influenced in one direction or another, that side of opinion will tend to take over, although this process takes many millions of steps.

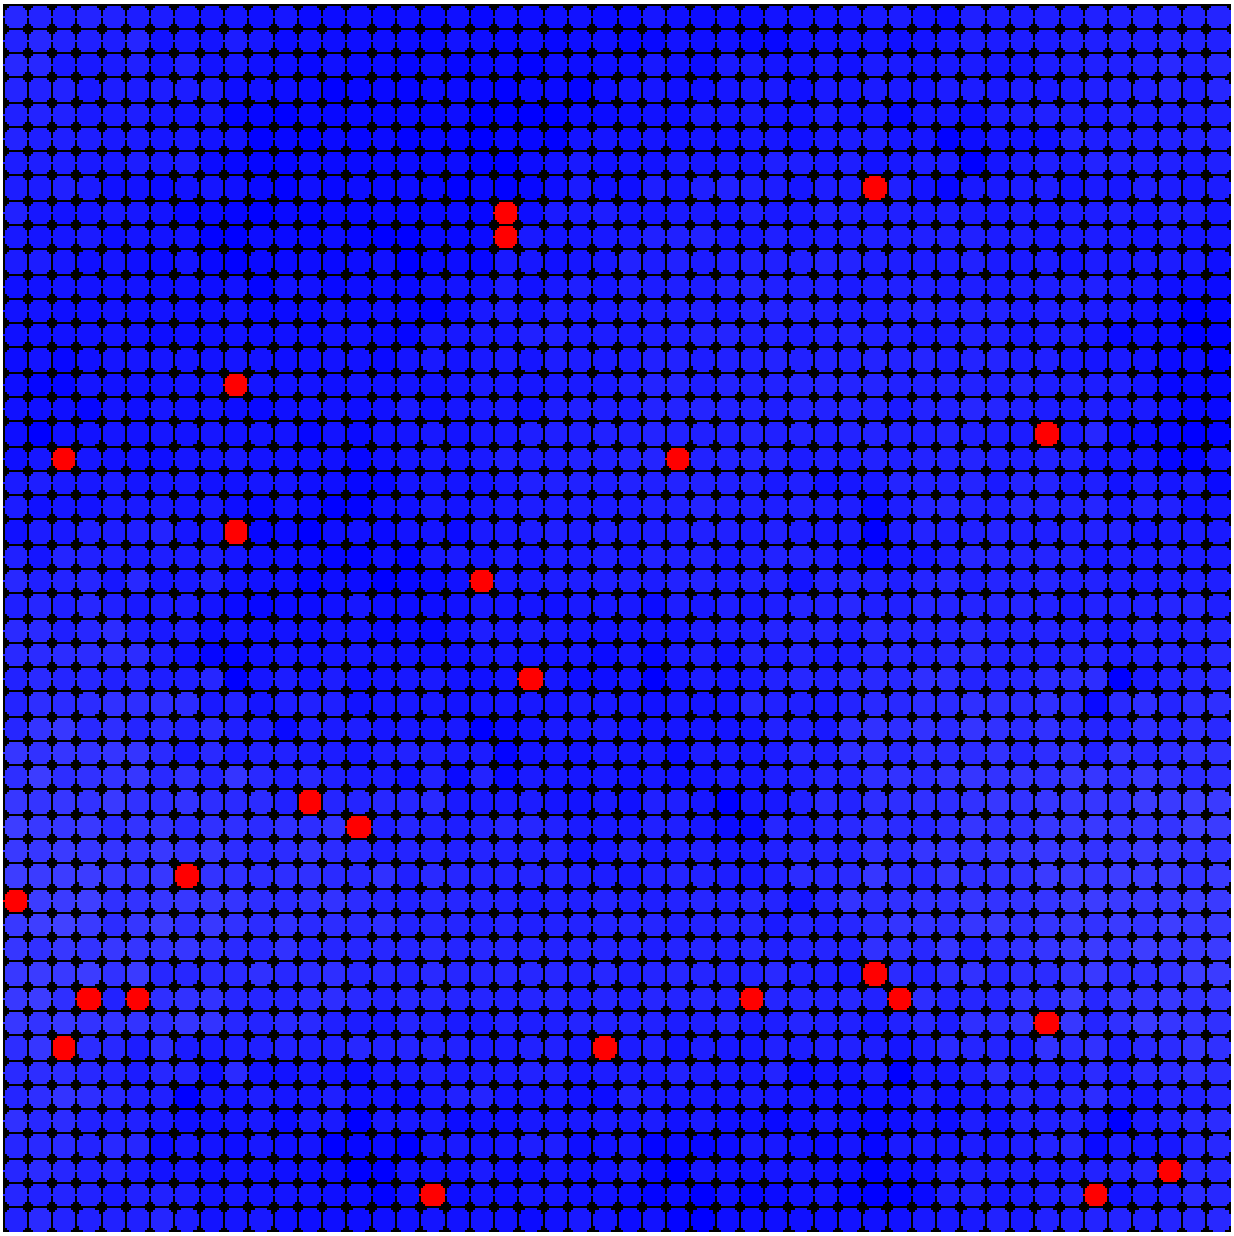

As the figure clearly illustrates, only dedicated spreaders remain away from the consensus. Almost everyone else has been converted strongly to disbelievers.

As we increase the steepness parameter, we see a very interesting type of evolution. Individuals of similar opinion tend to form continuous enclaves. The formation of enclaves seems quite chaotic.

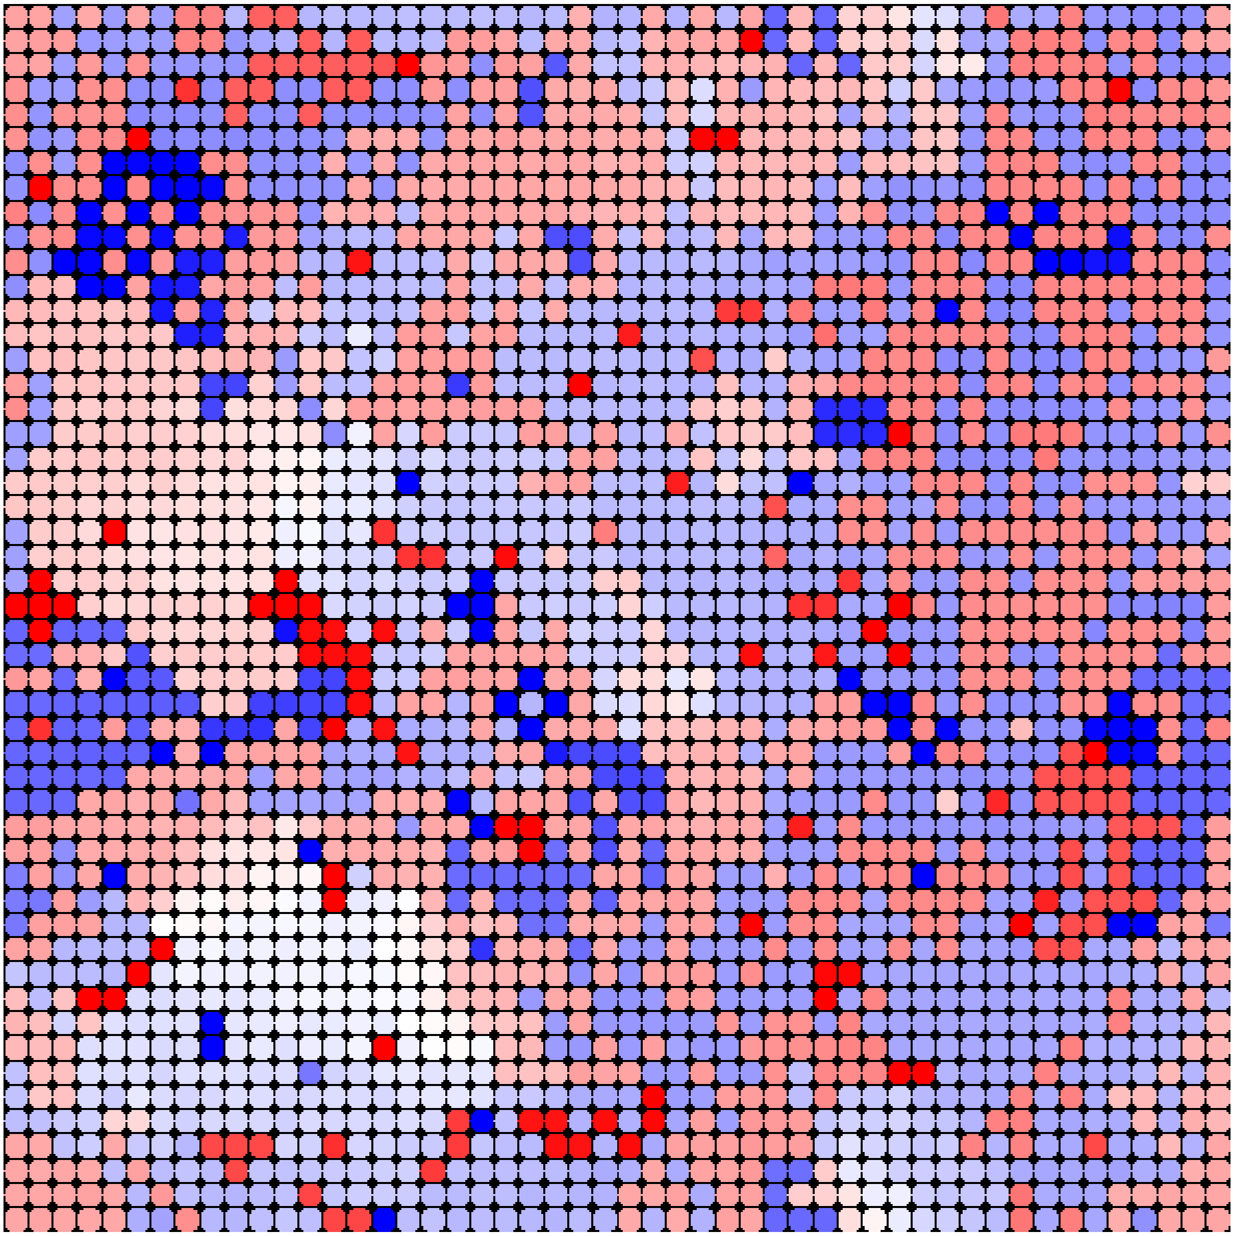

This figure, produced with a steepness of 9.0, illustrates the interesting effects of decreasing the influence of different opinions at an extreme rate. We can see both solid color enclaves and enclaves that have gradients of opinion. The scattering of spreaders and quashers seems to have only a small impact on the layout of these enclaves, demonstrating the powerful impact of the non-persistent initial conditions. It would be quite interesting to see how varying the parameters further changes the number, size, and beliefs of these enclaves, although that research has not yet been conducted.

The beliefs tend to progress towards a trimodal distribution with persistent enclaves of high, medium, and low belief, illustrated below.

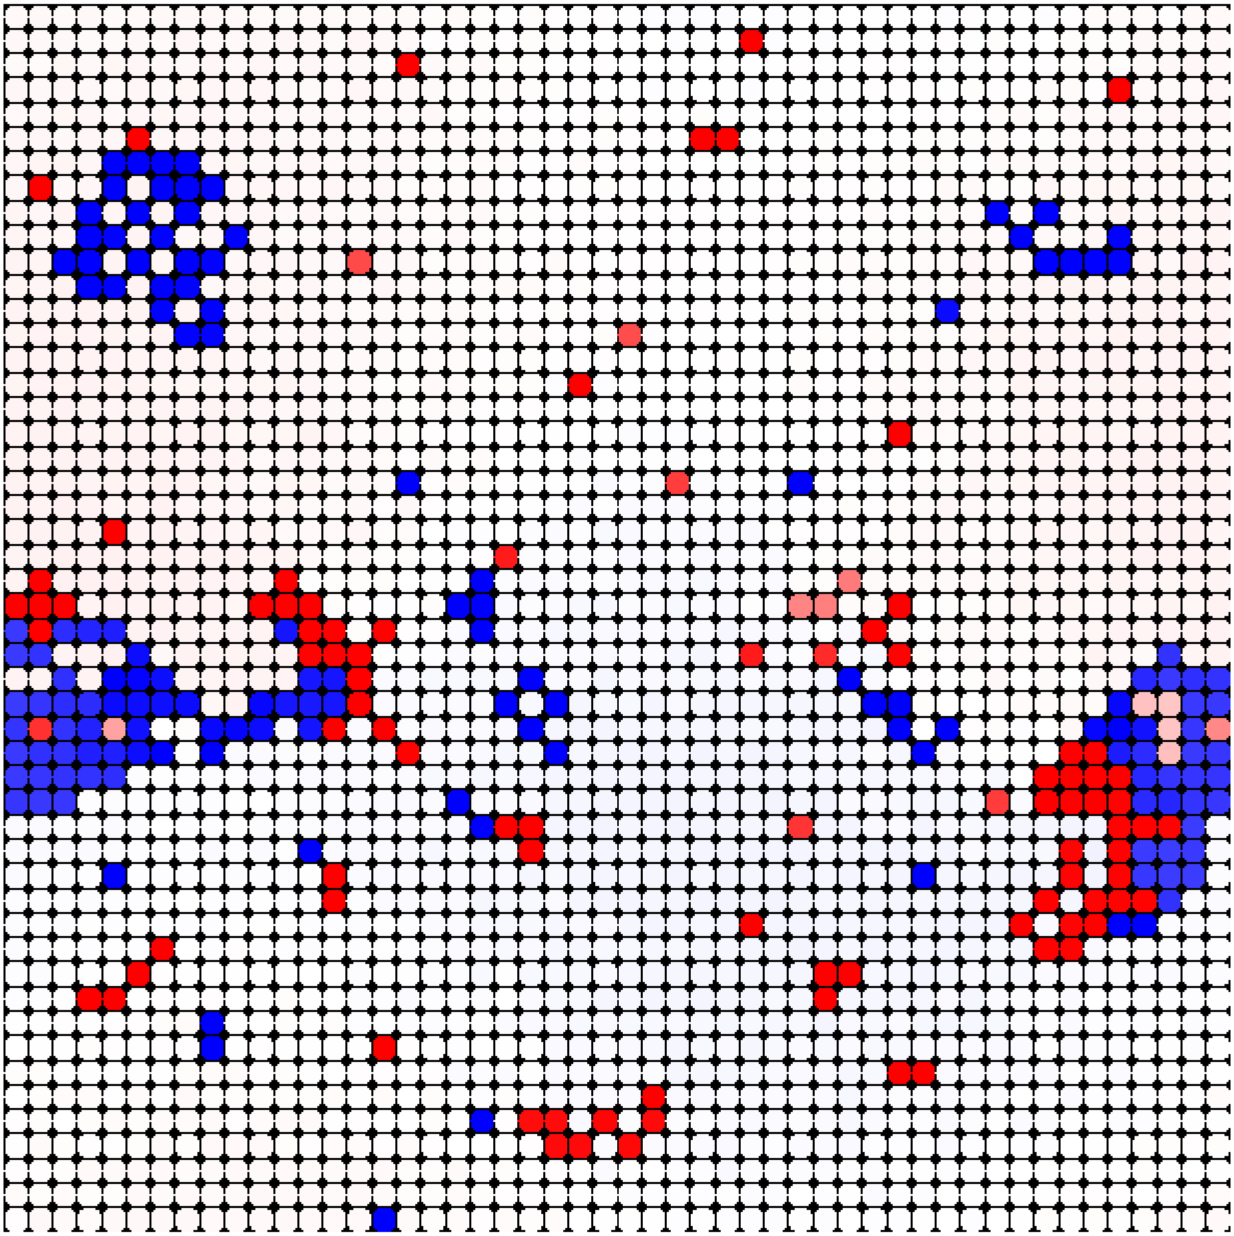

This result differs from the low influence steepness model in the key respect that many of the inherent spreaders and quashers and surrounded by solid enclaves rather than gradients. They are more isolated from the common opinion. It is also interesting that high and low belief enclaves actually tend to cluster together, because the high belief enclave will have longer to convert a slightly believing individual bordering many low belief individuals than a slightly red individual bordering many white nodes. The white nodes will convert the slightly red node quickly, bringing it out of reach of the very red nodes, while the fully blue nodes will have little influence on the slightly red one. Note the similar result that some individuals that are actually quite close to the median position are effectively trapped at some distance from the median by being surrounded by individuals of distant belief ("pink" enclaves within a large low-belief enclave).

In one simulation we increased the proportion of exogenous spreaders and quashers, and as a result we saw more in the way of more polarized enclaves, as one might expect.

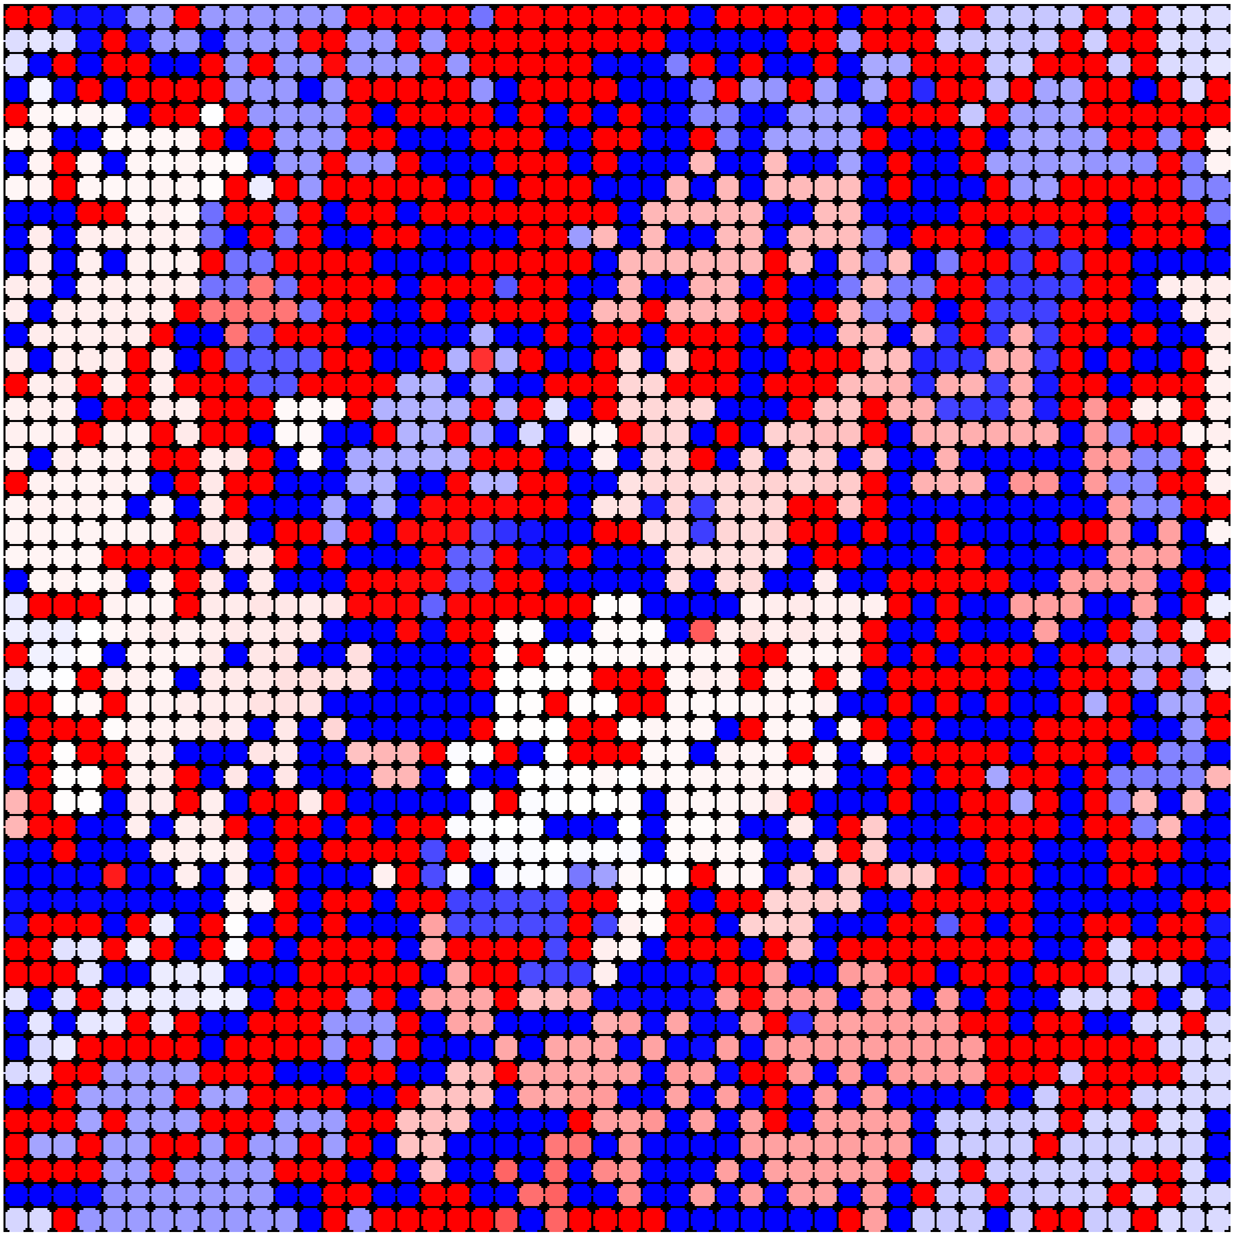

Inactive quashing

One of the most interesting questions about rumor spreading is whether it is a good idea to try to actively quash the rumor by discussing ones disbelief in the rumor or whether simply avoiding discussion of the rumor to prevent its spread is a good idea. To examine this question, we included an option to fix the activity level of the quashers at 0. If asked about the rumor, the quasher will express their beliefs, but they won't bring the topic up on their own; their low activity level also decreases the overall probability that the rumor will be discussed. Our model suggests that under standard parameters, it is unwise to stay silent. The mean activity remains positive, and the influence of the negative beliefs is reduced by their lower activity. Belief in the rumor will almost always take over under these conditions.

If the objective of the quashers is to minimize the number of individuals who know the rumor and have greater than 0 belief, it can be a good idea to avoid bringing the subject up if the decay rate of the rumor is sufficiently high (that is, the rumor does not seem to be generating self-sustaining interest). Consider the boundary case in which the activity decay constant is set at 0.82. In the previous sections, we see that the activity level often wanders around but stays positive enough that the information reaches the entire population. This implies that generally we would expect half of the population to know the rumor and believe it to some degree. By fixing their activity level at 0, quashers can cause the activity level to crash under these conditions.

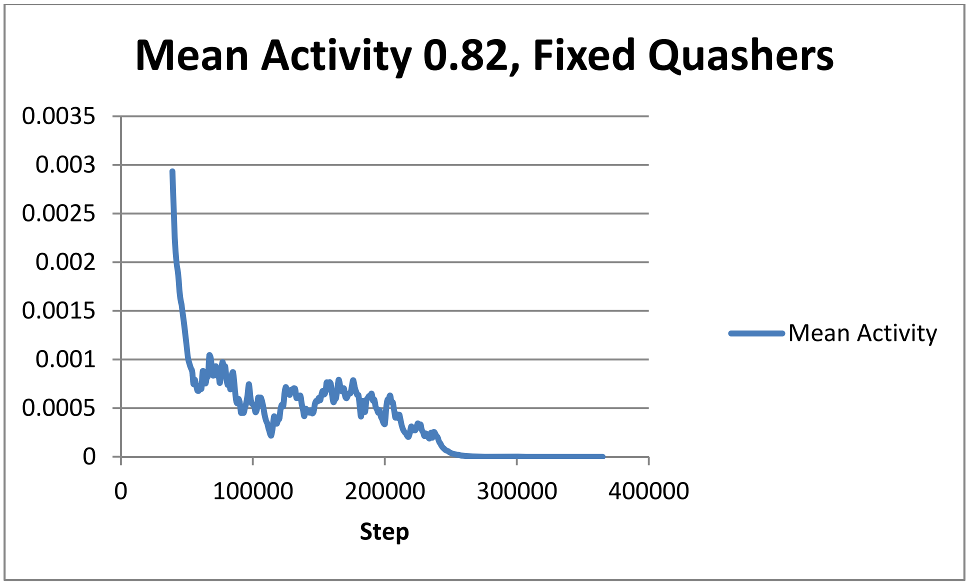

In this example, the proportion of individuals who heard the rumor only reaches 2.12% before leveling off. This could be a very good outcome for quashers. The crash does not occur every time, however, and further research is needed to determine whether the potential benefit of avoiding the topic is worth the potential loss in standing from not actively campaigning. An optimal strategy might involve waiting to see if the rumor goes away on its own for a while and then actively campaigning if the activity takes off (a behavior we see in actual political campaigns to varying success). Another possible avenue for future research might be to determine whether this silence approach is only useful at the phase transition or whether its impact on the asymptotic limit of proportion of knowledgeable individuals might make it a good strategy when activity decays even more quickly.

Further Extensions

A number of important questions could be asked about the benefits and costs of rumor spreading and quashing activities. If individuals could spend resources to increase their interaction rate or find better positions, how might they be able to influence the results? Furthermore, it might be interesting to examine whether is advantageous for the rumor spreader to manage the rate of information spread in order to leverage the advantage of having the source of the rumor be a spreader. Alternatively, perhaps they could try to select an advantageous location on the grid if it is possible to manipulate the social network of interest in such a way.

In our current model, on the initial spreaders and quashers are invested in the rumor's spread such that their beliefs are invariant. Early results, not further reported herein, suggest that if these partisans can convert uncommitted individuals to their position at a low probability when the uncommitted individual's belief becomes extreme enough, even slight early advantages in conversion can result in broad dominance of opinion. This has implications not only for rumors in political campaigns due to their ability to affect partisan standing, but also in the study of the spread of religious and other dogmatic beliefs.

Our model could also be used to examine the spread of new ideas in academia and research, in which belief level might model acceptance of a new idea. Outside of the realm of social science, this model might have isomorphisms with physical processes such as the spread of forest fires or nuclear reactions and could be used to model such processes with the appropriate mapping of analogous variables.