

Final Project Progress Report # 2

Sara Clifton
Agent Based Modeling
Professor Uri Wilensky
Northwestern University

Introduction

People claiming no religious affiliation constitute the fastest growing religious minority in many countries, including the United States [1]. In fact, the religious “nones” are the only group growing in all 50 US states [2]. Although many scholars attribute the decline of religious affiliation to generational changes, roughly half of the US population changes religious affiliation at some point in their life, often several times [3]. This suggests that religious affiliation shift can be modeled as social group competition, with different religious groups competing for members [4]. Such models predict that religious coexistence is not a stable state; the unaffiliated group will grow until all organized religion has disappeared. Whether or not this comes to fruition, all people have a stake in the outcome.

Dynamical systems model

Abrams et al. split an ideal society into the mutually exclusive religiously affiliated and unaffiliated, with the fraction x belonging to the unaffiliated group and $y = 1 - x$ belonging to affiliated. Assuming that people only switch affiliations based on the fraction of people in each group and the perceived utility of the group, the dynamics of conversion can be modeled by

$$\frac{dx}{dt} = yP_{yx}(x, u_x) - xP_{xy}(x, u_y), \quad (1)$$

where $P_{yx}(x, u_x)$ is the probability per unit time that an individual converts from religious to unaffiliated, $0 \leq u_x \leq 1$ is the perceived utility of being unaffiliated, and $u_y = 1 - u_x$ is the perceived utility of being affiliated. The authors further assume that (1) is symmetric under exchange of x and y and that no individual would switch to a group with no members (i.e. $P_{yx}(0, u_x) = 0$).

If the transition probabilities are smooth and monotonically increasing in both arguments, then there exist at most three fixed points with alternating stability. All available data suggest that the inevitable steady state is $x^* = 1$, or the extinction of religious affiliation. For specificity, the authors chose the power law $P_{yx}(x, u_x) = x^a u_x$; the best fit to data occurs for $a = 1$. The authors extend this model to binary networks of individuals (rather than all-to-all coupling) and allow for a continuous “religiosity” degree (rather than binary in or out of group), but the final state remains the same. As long as the network is not completely disconnected, only a time delay is introduced.

Agent based model

Minimal continuous dynamical systems lend themselves well to rigorous analysis, but important details may be left out for simplicity. Assuming religious affiliation change can be modeled as social group competition, agent based modeling is a natural way to test the robustness of the continuous model results.

Previous version

The previous version verified that an agent based implementation of (1) with an all-to-all network replicates the approximately logarithmic growth of the unaffiliated group and the eventual extinction of the religiously affiliated.

Initially, COMMUNITY-SIZE agents are given a random unaffiliated utility sampled from a normal distribution with mean U-X and standard deviation U-X-VAR (taken to be 0 for initial tests). A proportion INITIAL-X of the population is unaffiliated and the rest are affiliated. To be consistent with all available data, INITIAL-X must be small ($\approx 5\%$).

At each tick, affiliated agents switch to unaffiliated with probability $\text{TIME-SCALE} * u_x * x^a$, where a is initially set to 1 to test agreement with (1). Unaffiliated agents switch to affiliated with probability $\text{TIME-SCALE} * (1 - u_x) * (1 - x)^a$. See Figure 1.

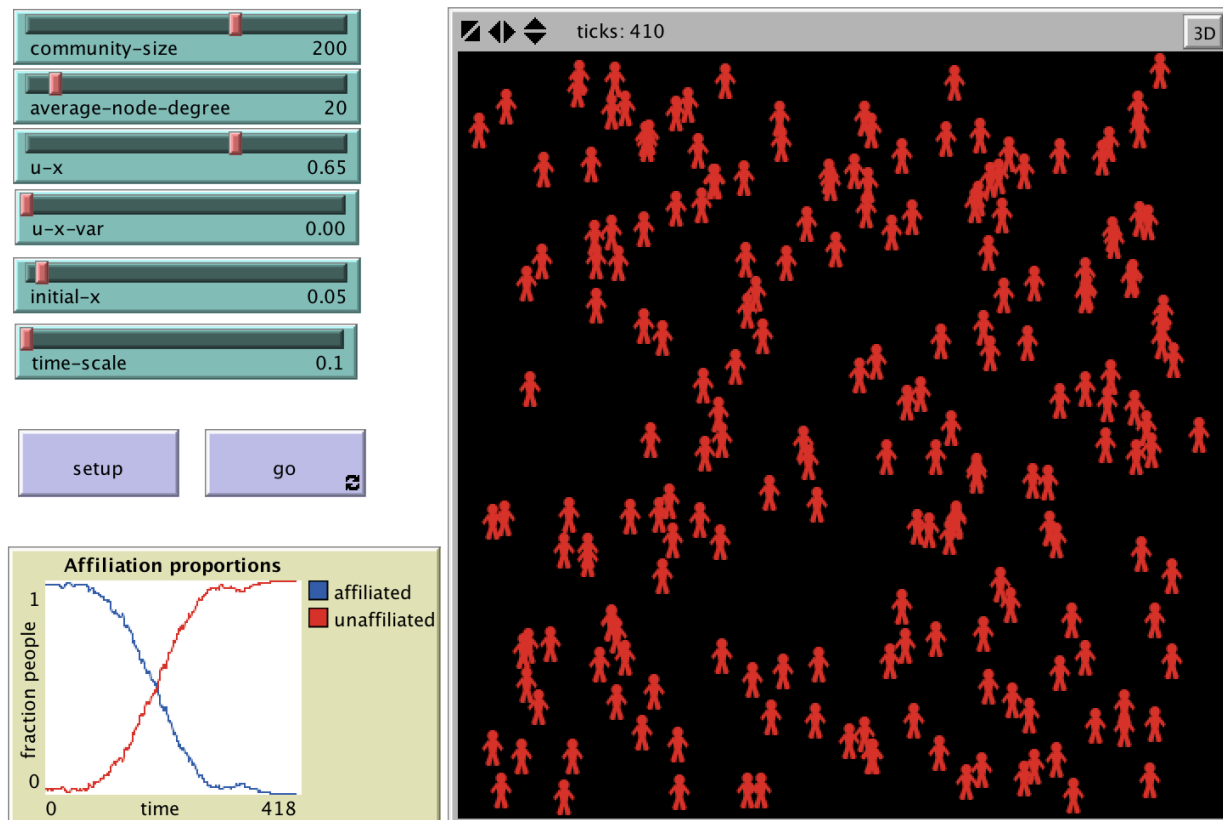


Figure 1: Verification that agent based implementation of (1) replicates the logarithmic growth of an initially small unaffiliated population. For a uniform unaffiliated utility u_x exceeding 0.5, the religiously affiliated group will convert to unaffiliated.

After confirming that the agent based model (with specific settings) produces the same results as the dynamical systems model, I test the robustness of the model results.

Add individual affiliation utility

One of the main assumptions of the continuous model is that the entire population perceives that religious affiliation has the same utility. While this assumption greatly simplifies analysis, it is much more realistic to assume that people perceive a wide range of utility in affiliating themselves with a religious organization. As a first attempt, I assume that unaffiliated utility is normally distributed with mean 0.65, as indicated by world-wide data [4].

The model results are robust to small variation in utility. However, if the standard deviation of utility is large enough that a significant number of people have $u_x < 0.5$, then the system stabilizes with a small fraction of affiliated individuals. In other words, the two groups coexist eventually. See Figure 2.

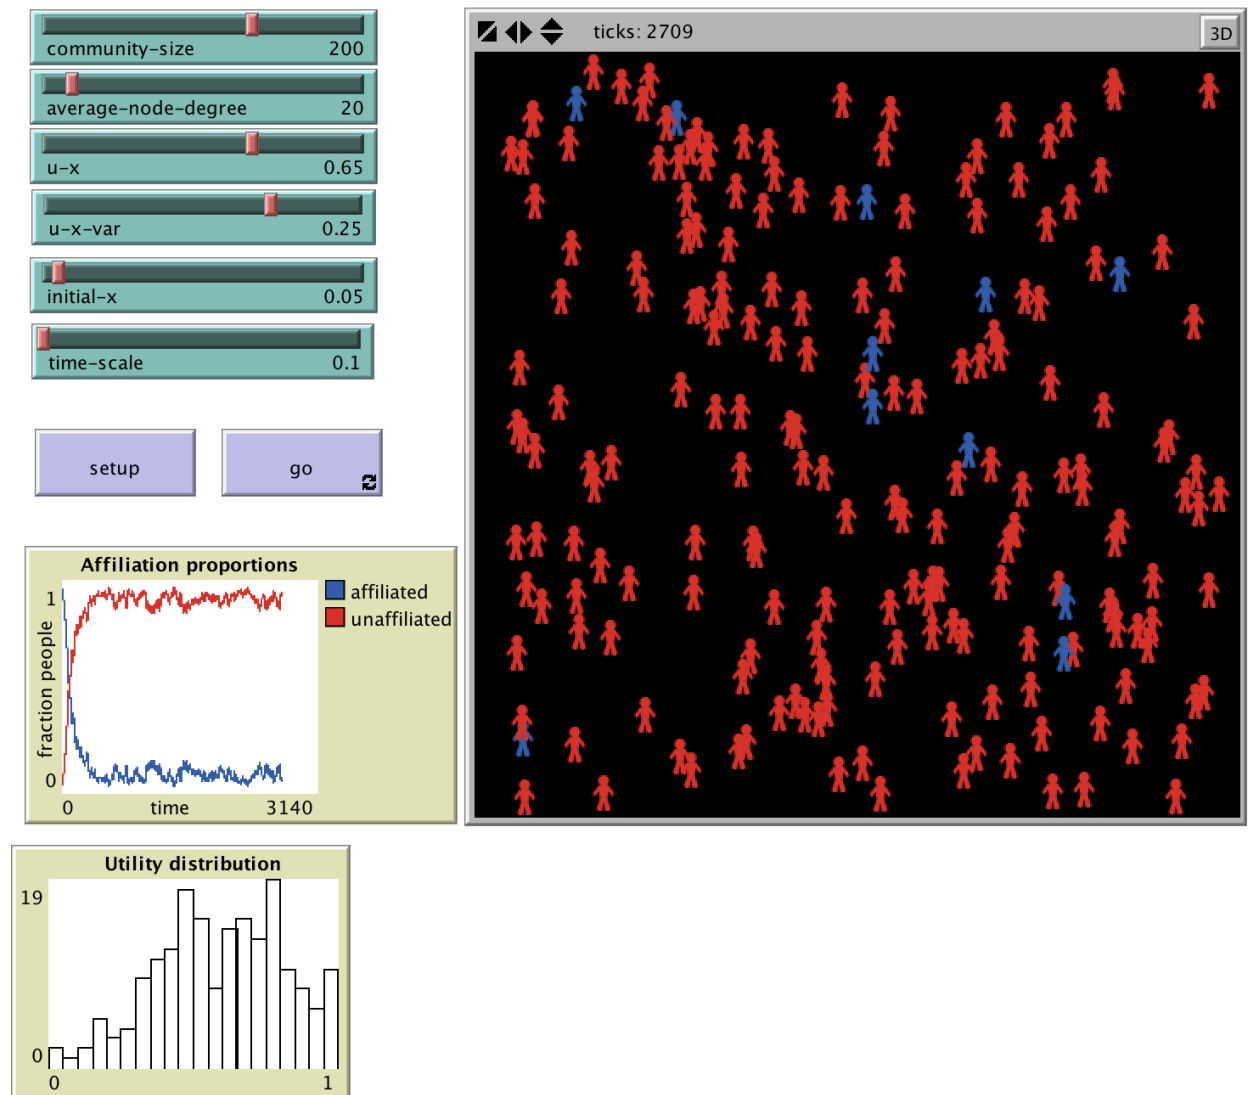


Figure 2: The continuous model results are not robust to large variation in affiliation utility. The population stabilizes with a small but significant proportion of the population still religiously affiliated.

Change exponent of power law

The authors found that the exponent $a = 1$ minimized the model error using real-world data, but they only checked integer values of a for simplicity. While the model is robust to small changes in a , qualitative differences emerge when the exponent changes too much. See Figure 3.

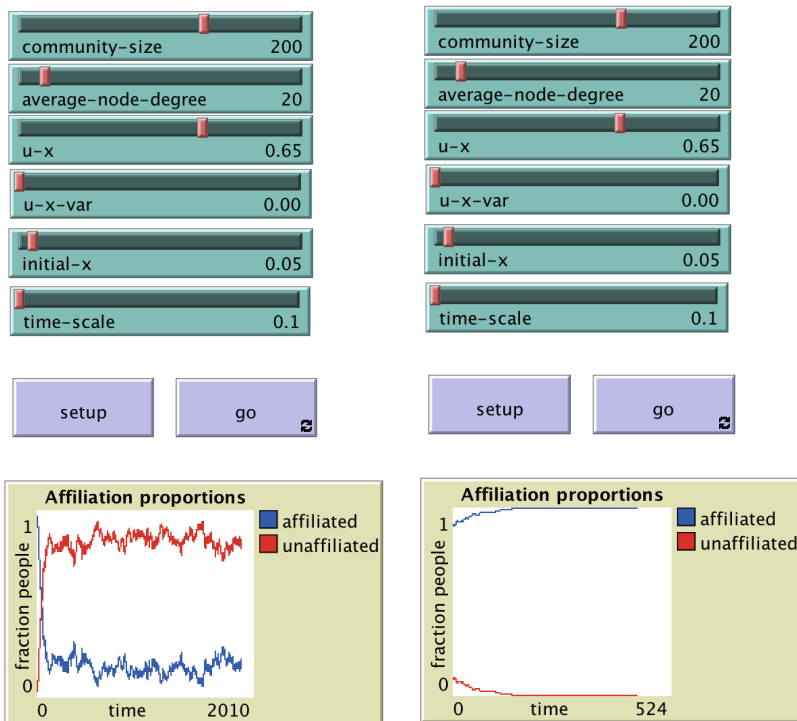


Figure 3: The continuous model results are not robust to changes in the exponent a of the power law switching probability. For $a < 1$ the population can coexist (left, $a = 0.6$). The early dynamics do not contradict the real-world data, so this is a plausible outcome. For $a > 1$ the affiliated can be the sole survivors (right, $a = 1.4$). The early dynamics are inconsistent with real-world data, so this case is not supported by the model.

Current version

The current version connects people in a social network. If NETWORK? is switched on, people change affiliations based only on their friends' affiliations rather than the affiliated proportion of the entire population. The social network also changes as people switch affiliations.

Initially, COMMUNITY-SIZE agents are given a random unaffiliated utility sampled from a normal distribution with mean U-X and standard deviation U-X-VAR (taken to be 0 for initial tests). A proportion INITIAL-X of the population is unaffiliated and the rest are affiliated. If NETWORK? is on, a spatially clustered network with AVERAGE-NODE-DEGREE is created.

At each tick, affiliated agents switch to unaffiliated with probability $\text{TIME-SCALE} * u_x * x^{\text{AFFILIATION-POWER}}$, where x is the unaffiliated proportion of an agent's friends. Unaffiliated agents switch to affiliated with probability $\text{TIME-SCALE} * (1 - u_x) * (1 - x)^{\text{AFFILIATION-POWER}}$. See Figure 4.

As seen in [4], the unaffiliated group grows approximately logarithmically, and the affiliated group

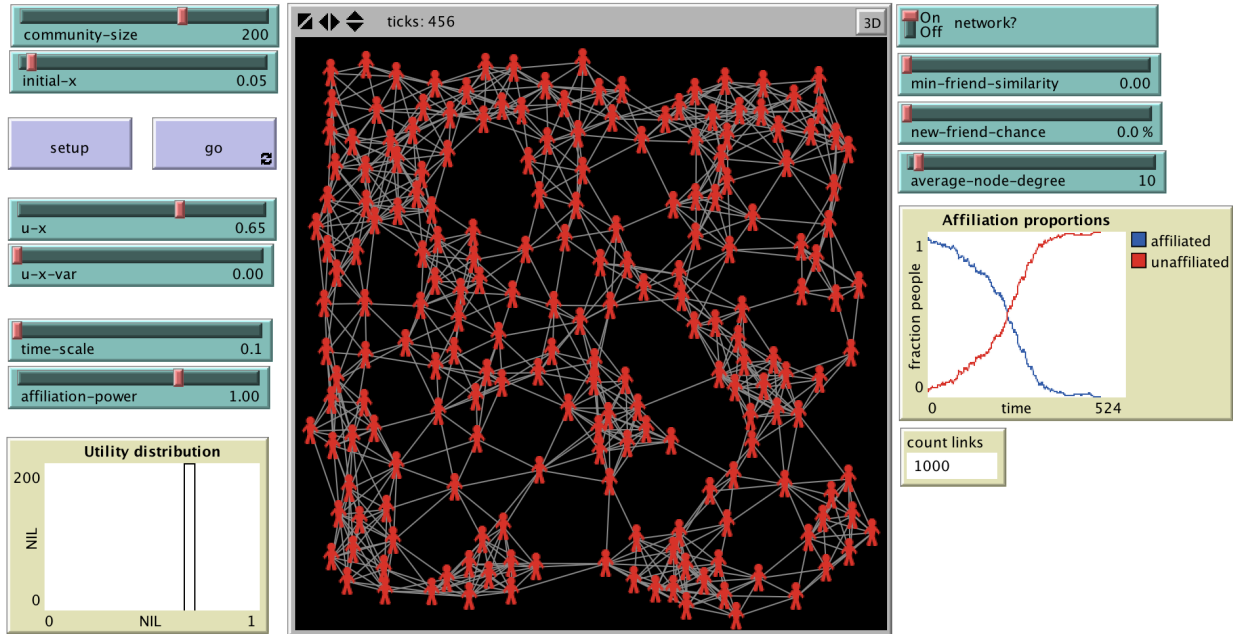


Figure 4: Verification that the agent based implementation of (1) on a network replicates the logarithmic growth of an initially small unaffiliated population. For a uniform unaffiliated utility u_x exceeding 0.5, the religiously affiliated group will convert to unaffiliated, but with a time delay from the all-to-all network.

eventuality disappears. When the network is not all-to-all, the system reaches steady state with a time delay, which is also consistent with the dynamical systems model. As AVERAGE-NODE-DEGREE decreases, the time delay increases. For small average node degree (< 5), the affiliated group no longer goes extinct because small clusters are not connected.

After confirming that this version of the agent based model produces the same results as the dynamical systems model on a network, I continue to test the robustness of the model results.

Add affiliation-dependent network

One of the assumptions of the continuous model (1) is that the network is static, but it's more realistic to assume that people lose and make friends when they switch affiliations. At each time step, people switch affiliations according to the previous transition probabilities. However, if an agent switches from affiliated to unaffiliated (or vice versa) in that time step, it will lose some affiliated friends, but it may gain a new unaffiliated friend.

Each time an agent switches from affiliated to unaffiliated, it calculates what proportion of its friends are unaffiliated. If that proportion is less than the MIN-FRIEND-SIMILARITY desired by all agents, the agent break ties with random affiliated friends until the MIN-FRIEND-SIMILARITY proportion is reached. If agents want most of their friends to have similar affiliations, then the unaffiliated and affiliated groups can coexist. The system stabilizes when the affiliated agents either form their own separate cluster or disconnect from the network entirely. See Figure 5.

In addition to losing friends after switching from affiliated to unaffiliated, each agent might gain a new unaffiliated friend with NEW-FRIEND-CHANCE. As this probability increases, the proportion of unaffiliated agents at equilibrium increases. See Figure 6.

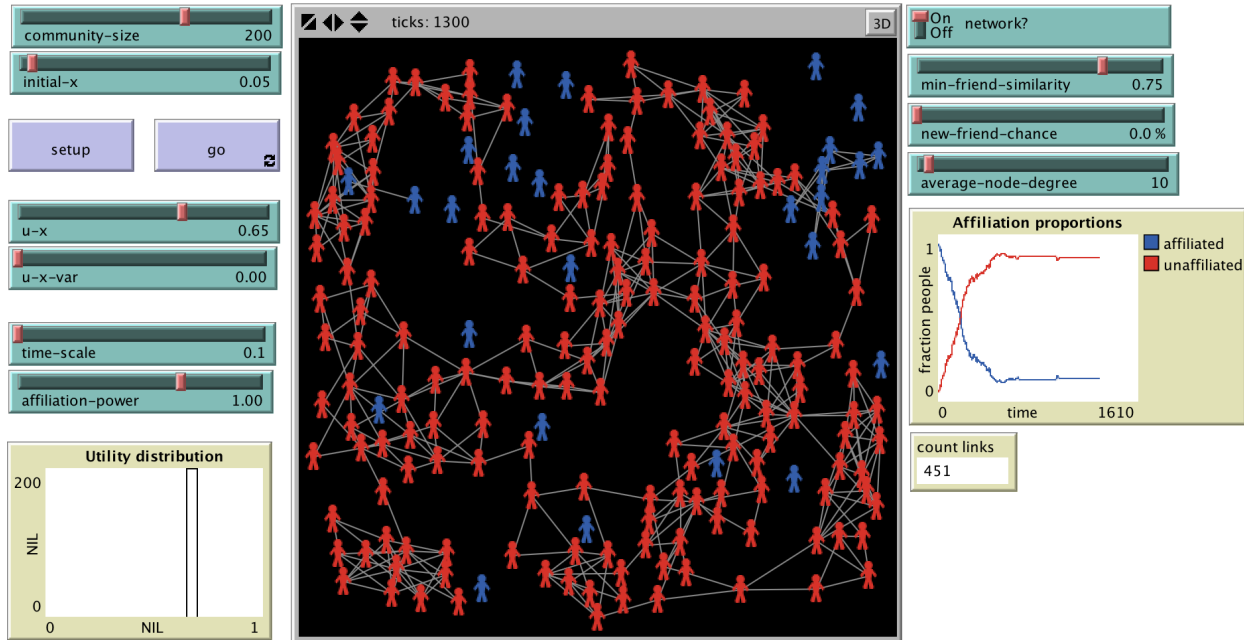


Figure 5: When the agents want to be friends with at least 75% similarly affiliated agents, the two groups can coexist. At equilibrium, affiliated agents have either formed their own separate cluster or disconnected from the network entirely.

Combine all current features

I have thus far tweaked model parameters one at a time. If for instance I add variation to affiliation utility, shift the affiliation power off of exact $a = 1$, set the desired friend similarity to 66%, and the chance of making new friends to 10%, the agent based model results are consistent with real-world data but the equilibrium state disagrees with the continuous model predictions. See Figure 7.

Next steps

The following questions/verifications will be investigated next:

1. The algorithm currently does not perfectly meet agents' desire to have a least MIN-FRIEND-SIMILARITY. This bug will be fixed shortly.
2. What happens when birth, death, and immigration are incorporated into the model? What if affiliations have different birth and death rates? What if immigrants have different religious affiliations from the current population?
3. What happens when there is a cost to switching affiliations? Perhaps a person can only switch affiliations a fixed number of times.
4. What if people make and break friendships all the time (not just when switching affiliations)?

In addition to answering these questions, I will build a generic social group competition model using HubNet where users choose a group to join based on the current membership and utility of the group. It would be

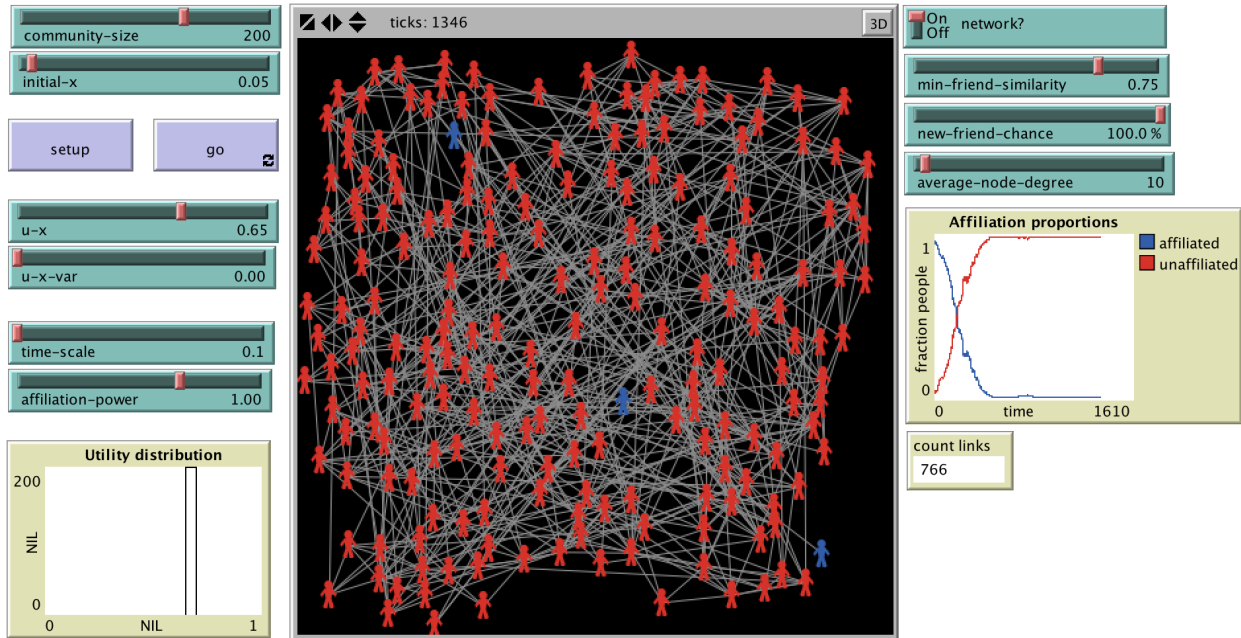


Figure 6: When the agents have a good chance of making friends with similarly affiliated agents after switching groups, the two groups barely coexist. Only a few affiliated agents remain, often unattached to anyone in the network.

interesting to see if the same transition functions emerge from this abstraction of group competition. I imagine that large patches would represent affiliations, and users would move to those patches based on provided utilities (varying across users) and the number of other agents on the same patch. If I want to impose a network structure, non-neighboring agents would be invisible.

Conclusions

The second version of the agent based model of religious affiliation replicates the continuous dynamical systems model on a network under the specified conditions. Specifically, the agent based model predicts logarithmic growth of the unaffiliated group and the eventual extinction of the affiliated group. The early dynamics of the model are consistent with real-world data, but the world is far from religious equilibrium. It is therefore impossible to use real-world data to select which of the models best describes religious shift or what the eventual outcome will be.

I am testing the robustness of the continuous model predictions under various reasonable conditions. After allowing people to perceive different utilities for religious affiliation (while still retaining the average utility implied by the data), coexistence of the affiliated and unaffiliated is possible. Tweaking the exponent of the power law switching probability also allowed coexistence. These realistic additions to the model are still consistent with real-world data, but the final state can be qualitatively different.

Coexistence of the affiliated and unaffiliated groups is also possible when the people are connected in a social network, only looking to their friends to decide their affiliation. When people break ties with differently affiliated people and make friends with similarly affiliated people, the social network breaks into unconnected clusters. Often, the minority affiliated individuals are completely alienated from the rest of the population.

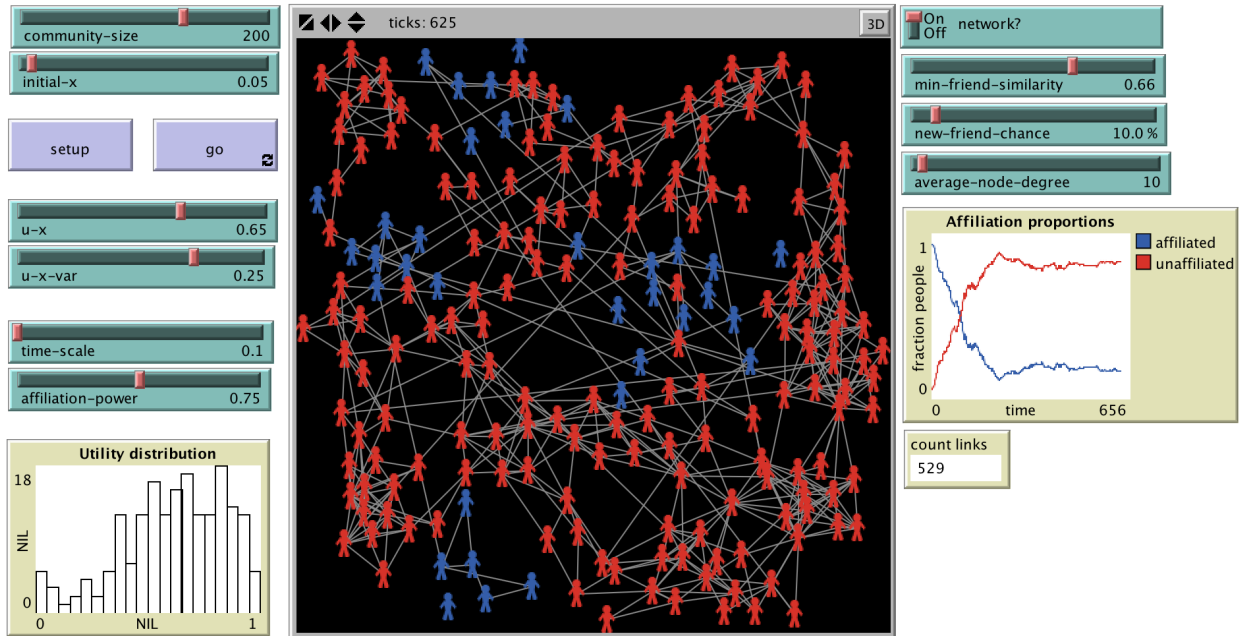


Figure 7: Changing all model parameters at once, I see a consistent pattern. Consistent with real-world data, the number of unaffiliated agents grows until the the unaffiliated group is the clear majority. Then the number of unaffiliated members peaks before settling down into a slightly lower proportion. At equilibrium, the two groups coexist with very few alienated individuals. This seems like a realistic outcome consistent with real-world data.

References

- [1] M. Martin, *The Cambridge companion to atheism*. Cambridge University Press, 2006.
- [2] B. A. Kosmin, A. Keysar, R. T. Cragun, and J. Navarro-Rivera, “American nones: The profile of the no religion population,” *Hartford, CT: Institute for the Study of Secularism in Society and Culture*, 2009.
- [3] L. Lugo, S. Stencel, and J. Green et al., “Faith in flux: Changes in religious affiliation in the us,” 2009.
- [4] D. M. Abrams, H. A. Yapel, and R. J. Wiener, “Dynamics of social group competition: modeling the decline of religious affiliation,” *Physical Review Letters*, vol. 107, no. 8, p. 088701, 2011.