

**The Surprising Effect of Implementation Choices
on the Rate of Convergence of Opinion Dynamics Models**

Stephen Davies

Department of Computer Science
University of Mary Washington
1301 College Avenue
Fredericksburg, VA 22401, USA

Hannah Zontine

Department of Computer Science
University of Mary Washington
1301 College Avenue
Fredericksburg, VA 22401, USA

ABSTRACT

Opinion Dynamics models seek to reproduce the phenomenon of individual agents forming opinions over time via mutually influencing one another. As with all agent-based models (ABMs), researchers seek to faithfully model the real-world process through a selective simplification of the phenomenon.

When creating a simulation to realize such a model, various implementation choices present themselves that at first glance may look arbitrary. In some cases, however, the way in which the details of the simulation are implemented can lead to important differences in its overall behavior; in the extreme, incorrect conclusions may even be drawn about what macro-behavior the abstract model is guaranteed to produce. In this paper, we look at two such choices germane to Opinion Dynamics models: a subtle difference in the way agents are randomly chosen for interaction, and the direction of influence between the two agents in such an encounter. In both cases, the rate of convergence to equilibrium is profoundly affected by what seem to be arbitrary implementation choices. This raises questions about the robustness of certain conclusions that are sometimes drawn from ABMs, and suggests strategies for avoiding pitfalls.

1 INTRODUCTION

In this work, we consider Opinion Dynamics models, which seek to reproduce the phenomenon of individual agents forming opinions over time via mutual influence. The field boasts a large literature, full of claims about the behavior of various such systems, some supported with mathematical proofs, others sustained by empirical evidence from simulation results.

When implementing any agent-based simulation, the designer faces choices that seem inconsequential — “In which order should I update the state variables here?” “Should agents be treated in a consistent order, or should they be shuffled each time?” “Do I use a while loop or a for loop in this function?” “When simultaneously inserting several events into the event queue, which one do I enqueue first?” *Etc.* Yet despite their apparent unimportance, in some cases, system behavior may actually hinge on what choice is made. If this goes undiscovered, there is a risk that broad claims made about a general class of system may in fact be contingent on certain non-obvious implementation concerns.

Among other things, this underscores the importance of the ABM community taking the time to reproduce results from the literature. If several researchers, starting from the same conceptual description of a system, independently build implementations that produce the same macro-level behavior, this increases confidence that the claims made about the system are indeed robust. If not, this may expose the presence of latent implementation-dependent assumptions that should actually be promoted to the model description proper, rather than being omitted and hence left to an implementer. It is then worth considering whether the augmented model, with the new assumptions made explicit, is still a reasonable abstraction of the real-world system being studied.

In this paper, we look at two such simulation variants of possibly the simplest of all Opinion Dynamics models: the original Binary Voter Model (Holley and Liggett 1975, Clifford and Sudbury 1973). We argue that in neither case is one implementation choice obviously preferable to the other, and yet what may seem a matter of indifference actually has a profound effect on the runtime characteristics of the model – in this case, the convergence time to consensus.

2 OPINION DYNAMICS MODELS

Published independently by (Clifford and Sudbury 1973) and (Holley and Liggett 1975) in the 1970s, the Binary Voter Model (BVM) laid the initial foundation from which many other Opinion Dynamics models have been constructed. The BVM, which represents an individual's opinion as a single binary value, was intentionally simplified in several ways for the sake of deriving an analytical solution. One simplification, for instance, is that agents are distributed on a regular lattice rather than an arbitrary graph. Periodically, a randomly chosen agent adopts the opinion of a neighboring agent if the two opinions differ. Importantly, over time the BVM will always reach uniformity of opinion, according to (Aldous and Fill 2002, ch. 14).

Numerous models for the Opinion Dynamics phenomenon have been proposed, varying in a multitude of ways, for instance:

- the way opinions are represented:
 - discrete, with two or more distinct options (Föllmer 1974, Yildiz, Acemoglu, Ozdaglar, Saberi, and Scaglione 2011) or continuous, typically a value between 0 and 1 (Ghaderi and Srikant 2012, Weisbuch, Deffuant, Amblard, and Nadal 2001)
 - single (Weisbuch and Boudjema 1999) or multiple, where opinions on multiple different topics are considered (Deffuant, Neau, Amblard, and Weisbuch 2000, Sirbu, Loreto, Servedio, and Tria 2013)
 - expressed (most models) or latent, in which the true value is possibly hidden from other agents (Friedkin and Johnsen 1990)
- how agents encounter each other:
 - randomly from the whole population (Hegselmann, Krause, and others 2002), or neighbors in a social network (Clifford and Sudbury 1973, Holley and Liggett 1975)
 - pairwise (most models) or in groups (DeGroot 1974)
- how influence takes place:
 - copying another agent's opinion (Holley and Liggett 1975)
 - averaging their opinion with one's own (DeGroot 1974)
 - “disagreement” processes where opinions diverge rather than converge (Sirbu, Loreto, Servedio, and Tria 2013)

Other innovations have been studied as well. (Yildiz, Acemoglu, Ozdaglar, Saberi, and Scaglione 2011) expanded the BVM by adding a binary “stubbornness” attribute to each agent. Stubborn agents never update their opinion and others would often be changing depending upon whom they interacted with. They discovered that the addition of only a few stubborn individuals always results in a graph polarized by opinion (*i.e.*, non-consensus).

(Weisbuch, Deffuant, Amblard, and Nadal 2001) designed a model with continuous opinion values, where agents only adopt the opinion of a neighbor when the difference between their current opinions is below a fixed threshold value. (This is termed “bounded confidence.”) (Ghaderi and Srikant 2012) further expanded this model by considering degrees of stubbornness; agents differ in the degree to which they are biased in favor of their initial opinion. Both of these approaches were motivated by a desire to prevent the model from always reaching uniformity of opinion in equilibrium.

As indicated above, some researchers have explored models where agents have multiple opinion values. (Deffuant, Neau, Amblard, and Weisbuch 2000) gave each agent a vector of discrete (binary, actually)

opinions on different subjects. Agents aggregate opinions through pairwise evaluation, slightly shifting an agent's opinion after looping through all possible pairings within the population. After more than 1000 interactions among the agents, the researchers discovered orthogonalization of opinions, no polarization, and no correlation between the opinion vectors.

Finally, note that while the BVM restricts agents to only interact with their neighbors, other models have simulated interactions where agents consider the opinions of a group of other agents in the graph. (Hegselmann, Krause, and others 2002) (HK) constructed a model where agents encounter others randomly from the whole population. The HK model uses bounded confidence on a continuous opinion scale. Each randomly selected agent aggregates all other agent opinions within his confidence bound and considers the group average for shifting his opinion. In the DeGroot Model (DeGroot 1974), an agent updates his opinion to be a weighted average of his own opinion and the opinions of his neighbors.

3 THE MODEL AND VARIANTS

As explained above, the BVM is deceptively simple. Each node in the graph is initially assigned an opinion (say, 0 or 1) and updates it periodically by copying the opinion of one of its graph neighbors. It has long been known that such a system will reach consensus (uniformity of one opinion or the other) under a wide variety of conditions (see, *e.g.*, (Sood and Redner 2005)). The probability that 0 (as opposed to 1) becomes the dominant opinion as a function of the initial opinion distribution is known, as is the expected number of iterations required to reach consensus for various degree distributions of the graph.

Implementing this model as an agent-based simulation is straightforward. Yet in reproducing these classical results en route to other work, we discovered at least two subtle implementation choices that at first glance would appear unimportant, and yet which impact the convergence time in striking fashion. We present these not so much as important in their own right, but as exemplars of a more general problem: implementation choices that a modeler takes for granted may turn out to be critical to the behavior claimed for that model.

3.1 Simulation variant 1: choosing with or without replacement

When we say “each node periodically updates its opinion,” what exactly does that imply about the *order* in which the nodes are selected for the update process? There are at least two reasonable interpretations:

- For each iteration of the simulation's main loop, a node is chosen uniformly at random, *with replacement*. In this scenario, if there were five nodes in the graph, we might have the following sequence of choices for opinion update: node 3, 5, 5, 5, 2, 3, 2, 5, 4, 2 ... Notice that in this realization, node 5 was selected 40% of the time, node 1 was not selected at all, and node 4 wasn't selected until near the end of the sequence.
- Treat *each* of the nodes once (in random, shuffled order) before treating them all again (in a different, shuffled order), *etc.* Put another way, chose the nodes uniformly at random *without replacement* until the store of nodes is exhausted; then replenish and repeat. In this scenario, the sequence above would never happen; instead we might have something like this: 3, 5, 2, 4, 1, 5, 4, 3, 2, 1, ... In this way, every node is guaranteed to be chosen once in $n = 5$ iterations.

Clearly, based only on the system's English description, above, either of these choices is consistent with the spirit of the model. An implementer might choose either of them, either deliberately or (more probably) unconsciously. They both pass the “select nodes repeatedly in random fashion” test.

For the sake of clarity we adopt the terms “**selection with replacement**” and “**selection without replacement**” as the descriptions for these two alternatives.

Before examining the results, one question we might ask is: which of these two variants is more reflective of the real-world phenomenon? Compelling arguments can be made both ways. In favor of

selection without replacement is the observation that all human beings have 24 hours per day in which to live and interact. If we want to simulate the dynamic behavior of a social system over time, therefore, it is important to ensure that all agents act at a fairly similar pace. After all, in the real world, there is no sequential loop at all, agents are not successively “chosen” to interact, and no agent is “starved” for interaction as a consequence of a peculiar random number sequence.

On the other hand, it is also true that in any social system, some agents will be more active than others. Consider an online social network, such as Facebook or Twitter. Some users post many messages per day, while others only on rare occasions; and some *read* many posts per day, whereas others only glance at their customized news feeds once in a while. This heterogeneity of usage would seem to favor the **selection with replacement** variant, to do justice to the varying rates at which agents interact with one another.

3.2 Simulation variant 2: direction of opinion propagation

The model calls for the implementer to (1) choose a node randomly, (2) choose one of its graph neighbors, and then (3) have one of the nodes copy the opinion value of the other. But which way does the influence go? Is the originally chosen node the one whose opinion is updated, or is the neighbor’s? At first glance, this may seem to be a symmetrical and therefore arbitrary choice, and not be expected to impact the system’s behavior. Indeed it does, however, and in a striking way.

We define “**node influences neighbor**” as the variant in which the originally selected node is the one whose opinion is propagated (to a random one of its graph neighbors) and “**neighbor influences node**” as the alternative, in which the copying goes from the node’s neighbor back to it.

As to the question of which variant 2 choice is more reflective of reality, we again find no definitive answers. In this case, our ambivalence stems from the fact that the algorithm’s influence procedure isn’t itself a very good model of what actually happens in real human interaction. Influence in a real human conversation isn’t (usually) unidirectional, with the “winner” being determined by considering which participant initiated the dialogue. Instead, it is a complex process during which ideas are shared and defended, biases revealed, and opinions adjusted (or not) based on a myriad of factors. In a highly abstract model such as this one, which does not aim to capture such subtleties, how can we resolve this in the simplest and cleanest way? Clearly by choosing from the two alternatives, above. Unfortunately, as we will see, our choice here, too, made arbitrarily in the name of simplification, noticeably changes the system’s macro behavior.

4 EXPERIMENTAL DESIGN

We implemented a factorial design with the above two factors: “*selection method*” (with levels of **with replacement** and **without replacement**) and “*influence direction*” (with levels of **neighbor influences node** and **node influences neighbor**). Our response variable was the *time to convergence*; *i.e.*, the number of iterations (pairwise encounters between nodes) before total uniformity of opinion was reached.

We ran each of the four combinations of factor levels for 200 trials. Each trial began with an Erdős-Rényi random graph (Erdős and Rényi 1959) with 100 nodes and an edge probability of .04, using the R `igraph` package (Csardi and Nepusz 2006). If the random graph generated for a given trial turned out not to be *connected* (*i.e.*, not all nodes were reachable from all others via some walk), that graph was discarded and another generated until a connected graph was produced. (This is because the model does not guarantee convergence to uniformity of opinion for unconnected graphs.)

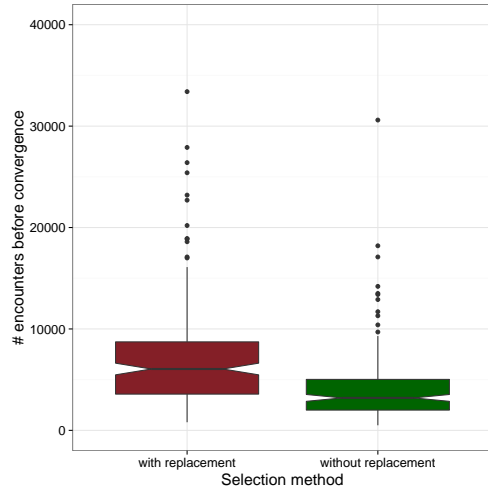
Each node was initially assigned one of the two opinion values at random, with a probability of $p=.5$ for each opinion value.

All of the initial graphs were generated once and then re-used for each of the four treatment combinations, to ensure that every treatment combination was simulated with the identical 200 initial conditions.

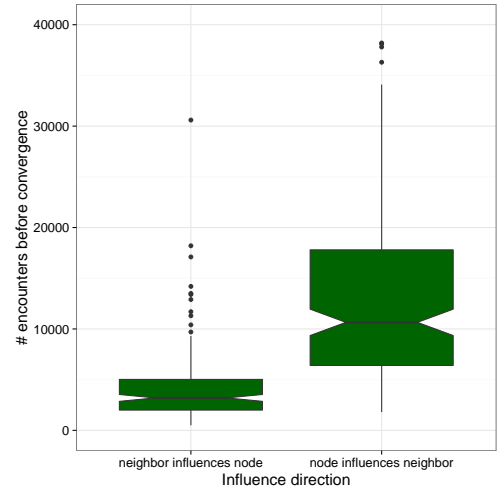
5 RESULTS

5.1 Simulation variant 1 (selection method)

Figure 1a presents a notched boxplot of the time to convergence for the **selection with replacement** vs. **selection without replacement** trials, with the *influence direction* variable held constant at **neighbor influences node**. Clearly the **with replacement** variant ($M=7296$, $SD=5382$) takes significantly longer (nearly twice as long, in fact) to converge than does **without replacement** ($M=4084$, $SD=3493$), as a two-sided t-test for the difference of means confirms ($t(341)=7.080$, $p=8 \times 10^{-12}$).



(a) Comparison of *selection method* levels, with *influence direction* set to **neighbor influences node**.



(b) Comparison of *influence direction* levels, with *selection method* set to **without replacement**.

Figure 1: Time to convergence of uniformity of opinion ($N=400$). (A few of the largest outliers have been omitted for the sake of readability.)

The explanation for this difference would appear to be the following. If the nodes to be influenced are selected *with* replacement, then inevitably some nodes will be relatively unaffected by their neighbors even as other nodes have their opinions updated multiple times. Thus the permeation of the to-be-dominant opinion throughout the system is uneven: the contagion reaches and “converts” some parts of the graph long before the “starved” nodes are influenced. Conversely, if chosen without replacement, every node in the graph regularly has a chance to be influenced, which means no hold-outs can “hide” in the graph.

5.2 Simulation variant 2 (influence direction)

An even greater difference in convergence time exists between the **neighbor influences node** ($M=4084$, $SD=3493$) and **node influences neighbor** ($M=14417$, $SD=11492$) variants, this time holding *selection method* constant at **selection without replacement** (see Figure 1b.) A t-test for the difference between these means yields $t(235)=-12.16$, $p=2 \times 10^{-16}$.

Two explanations of this behavior suggest themselves. The first is a similar “unevenness” argument that we offered for simulation variant 1: if every node is regularly chosen to *be* influenced, this should promote fairly rapid convergence of the graph as a whole. But if every node is simply chosen regularly *to* influence, then there is nothing stopping some of the nodes from receiving very infrequent influence. This, then, would lead to longer convergence times.

A second, more subtle argument relates to the non-uniform degree distribution of an Erdos-Renyi (or indeed, almost any) graph. Some nodes are connected to many neighbors, while others are connected to

very few. This means that while under **neighbor influences node**, a low-degree node will be influenced at the same rate as the others, if **node influences neighbor** is in play, it has a very low probability of ever being influenced. For this to happen, one of its very few neighbors would have to be chosen in an iteration (a rare event), and then of all that neighbors’ neighbors, the low-degree node would have to be the “lucky” one chosen. The necessity of both of these unlikely occurrences happening together makes low-degree neighbors unwittingly stubborn indeed (under **node influences neighbor**, which is perhaps responsible for variant 2 having even larger convergence time differences than variant 1).

5.3 Variable interactions

Finally, we point out a non-trivial interaction between the two factors. As illustrated in Figure 2, the difference in convergence times between the two *selection methods* is only significant under the **neighbor influences node** level. A 2×2 ANOVA for the entire 800-trial design revealed a main effect for *influence direction*, $F(1)=186.00$, $p=2 \times 10^{-16}$, but no significant effect for *selection method*, $F(1)=1.26$, $p=.261$. An interaction between the two effects was significant, $F(1)=20.42$, $p=7 \times 10^{-6}$.

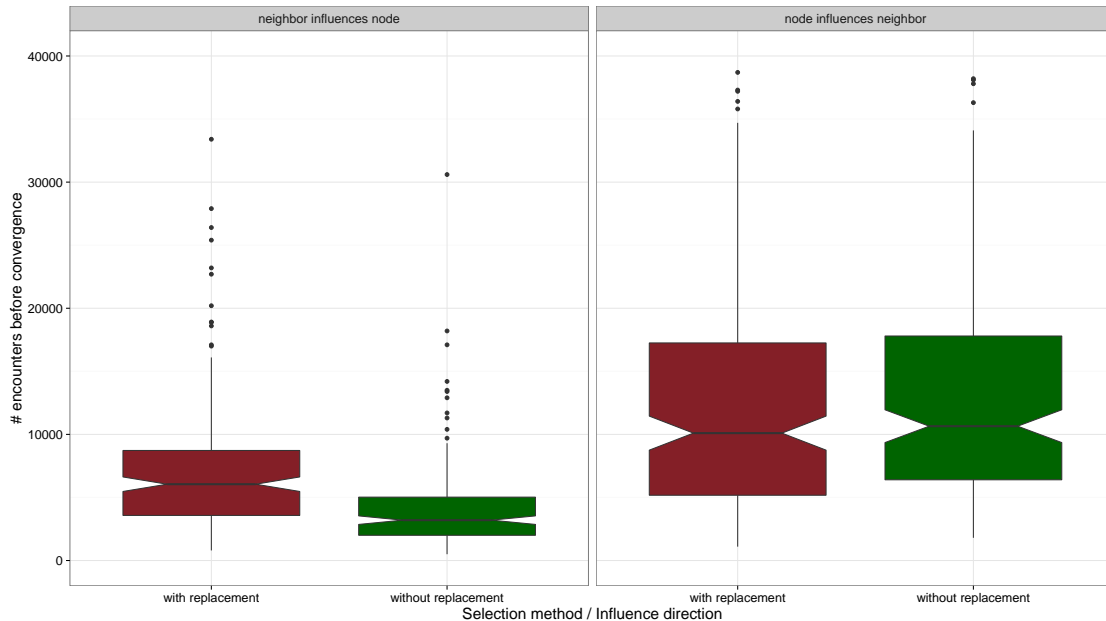


Figure 2: Full factorial design results for all four groups, illustrating interaction between effects (N=800). (A few of the largest outliers have been removed.)

The reason for the interaction is fairly straightforward, stemming directly from the analysis in the sections above. When **neighbor influences node**, it matters very much whether all nodes are chosen on a regular basis, since if some are not, they have no chance to change their opinion, causing an impediment to convergence. If **node influences neighbor**, on the other hand, the only thing at stake with the *selection method* choice is which nodes get to do the most influencing. When some nodes are noticeably less *influential* than others, we can perhaps expect some impact on convergence time (and even then, it’s not clear in which direction), but we will not experience a strong, consistent dampening effect as we will when some nodes are noticeably less *influenced* than others.

6 CONCLUSION

With this brief paper we hope simply to raise an alarm bell in the community of agent-based social simulation. In the midst of the (admirable) efforts to create abstract models reflective of reality, we have

discovered that it is easy for modelers to overlook the impact that presumably irrelevant implementation choices may have on a simulation's behavior. The two examples we raise in this article were ones we encountered in our own work, and which surprised us a great deal. The silver lining was that this led us to redouble our efforts to think through all aspects of our implementation a second time, with an eye to unmasking choices we may have made without realizing that they were important choices.

Only two remedies occur to us. The first is for ABM researchers to develop this kind of scrutinizing mindset as a habit. Moving from the modeling phase to the implementation phase of a project may not always just be a matter of "cranking it out." It would be wise to adopt a cautious attitude and to routinely second-guess every line of code: is the implementation decision I am about to make a natural consequence of what the model demands? Or is this an unwitting subjective choice which may nudge the system into giving a certain outcome and not another, despite the fact that the model itself does not prescribe either?

The second corrective measure is simply for the community to regularly reproduce the work of peers. We hear much about "reproducible research" in the ABM community these days, but perhaps less about actually *reproducing* such research. If models are theoretically reproducible but in practice rarely are, the dangers of unrealized assumptions lurking in computational social science results will have few opportunities to be exposed. For our part, we have resolved to more consistently attempt to replicate the results of publications in our areas of interest in the future, in order to learn more about the details of these models, and to help our community produce even stronger, more transparent science.

ACKNOWLEDGEMENTS

We are grateful to Dr. David Marchette and to the UMW Summer Science Institute for ideas, energy, and financial support.

REFERENCES

- Aldous, David and Fill, James Allen 2002. "Reversible Markov Chains and Random Walks on Graphs". Unfinished monograph, recompiled 2014, available at <http://www.stat.berkeley.edu/~aldous/RWG/book.html>.
- Clifford, P., and A. Sudbury. 1973. "A Model for Spatial Conflict". *Biometrika* 60 (3): 581–588.
- Csardi, G., and T. Nepusz. 2006. "The igraph software package for complex network research". *InterJournal Complex Systems*:1695.
- Deffuant, G., D. Neau, F. Amblard, and G. Weisbuch. 2000. "Mixing beliefs among interacting agents". *Advances in Complex Systems* 3 (01n04): 87–98.
- DeGroot, M. H. 1974. "Reaching a consensus". *Journal of the American Statistical Association* 69 (345): 118–121.
- Erdős, P., and A. Rényi. 1959. "On random graphs, I". *Publicationes Mathematicae (Debrecen)* 6:290–297.
- Föllmer, H. 1974. "Random economies with many interacting agents". *Journal of mathematical economics* 1 (1): 51–62.
- Friedkin, N. E., and E. C. Johnsen. 1990. "Social influence and opinions". *Journal of Mathematical Sociology* 15 (3-4): 193–206.
- Ghaderi, J., and R. Srikant. 2012, August. "Opinion Dynamics in Social Networks: A Local Interaction Game with Stubborn Agents". *arXiv:1208.5076 [cs]*. arXiv: 1208.5076.
- Hegselmann, R., U. Krause, and others. 2002. "Opinion dynamics and bounded confidence models, analysis, and simulation". *Journal of Artificial Societies and Social Simulation* 5 (3).
- Holley, and Liggett. 1975. "Ergodic theorems for weakly interacting systems and the voter model".
- Sirbu, A., V. Loreto, V. Servedio, and F. Tria. 2013, April. "Opinion Dynamics with Disagreement and Modulated Information". *Journal of Statistical Physics* 151 (1/2): 218–237.
- Sood, V., and S. Redner. 2005. "Voter model on heterogeneous graphs". *Physical review letters* 94 (17): 178701.

- Weisbuch, G., and G. Boudjema. 1999, March. “Dynamical Aspects in the Adoption of Agri-Environmental Measures”. *Advances in Complex Systems* 02 (01): 11–36.
- Weisbuch, G., G. Deffuant, F. Amblard, and J. P. Nadal. 2001, November. “Interacting Agents and Continuous Opinions Dynamics”. *arXiv:cond-mat/0111494*. arXiv: cond-mat/0111494.
- Yildiz, E., D. Acemoglu, A. E. Ozdaglar, A. Saberi, and A. Scaglione. 2011, January. “Discrete Opinion Dynamics with Stubborn Agents”. SSRN Scholarly Paper ID 1744113, Social Science Research Network, Rochester, NY.