

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Industry Influencing Collective Scientific Reasoning: A Bayesian, Agent-based Exploration

Permalink

<https://escholarship.org/uc/item/06t6b7v0>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 47(0)

Author

Schöppel, Klee

Publication Date

2025

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Industry Influencing Collective Scientific Reasoning: A Bayesian, Agent-based Exploration

Klee Schöppel (l.u.schoppl@rug.nl)

Department of Theoretical Philosophy, University of Groningen
Groningen, 9712 GL NL

Abstract

Recent work in Bayesian, agent-based modelling of scientific communities has employed the Bala-Goyal framework to study the mechanisms involved when industry influence applies the so-called ‘Tobacco Strategy’ to undermine collective inquiry. Motivated by limitations of these models, we propose an alternative based on a recently introduced framework for normative argument exchange across networks. We implement representations of two distinct types of industry influence: ‘Obfuscating’ influence directs inquiry to experiments with low expected value of information. ‘Misleading’ influence filters private research and only communicates misleading signals from the world. We explored the impacts of both strategies on the polarization & mean error of, and flow of information through, social networks of scientists via computer simulations. We conclude that even against highly optimistic background assumptions, and in a less simplified model of inquiry and argumentation, industry influence poses a plausible threat to collective deliberation.

Keywords: Argumentation; Agent-based modelling; Communication; Normative reasoning; Polarization;

Introduction

In “Merchants of doubt” (2011) Oreskes & Conway document how political and societal consensus on the perils of human-made climate change has decreased since the 1990s, as a direct result of intentional, ideologically motivated propping up of the notion of scientific uncertainty on the issue. They identify these strategies as adapted from corporate PR efforts to undermine certainty in the harms of smoking tobacco.

Since then, multiple publications in computational philosophy of science (Holman & Bruner 2017, O’Connor & Weatherall 2019, Weatherall et al. 2020) have started using agent-based two-armed bandit models (see Bala & Goyal 1998, Zollman 2007, Zollman 2010) to further explicate the social-epistemological processes involved when the ‘Tobacco Strategy’ so disrupts evidence-based policymaking. However, while helpful, these models are limited by the constraints of the two-armed bandit framework, which encodes an overly flat representation of scientific inquiry, both with respect to the questions investigated and the arguments exchanged by agents; and features reasoners making strategic choices that are difficult to justify outside of modelling particular disciplinary contexts (such as drug trials).

We propose an alternative model of industry influence in the Bayesian, agent-based framework ‘NormAN’ (Assaad et




Figure 1: The ‘Asia network’ (Lauritzen & Spiegelhalter 1988).

al. 2023), short for normative argument exchange across networks, which allows us to study industry influence on science in a model of argumentation¹, on the basis of complex causal relationships between evidence and the hypothesis.

It has recently been used to study the effects of different communication rules on polarization (Assaad & Hahn 2024), the effects of self-censoring in discourse participation (Schöppel & Hahn 2024), and to evaluate the merits of opinion averaging as a means of group deliberation (Hahn et al. 2024).

In this paper, we newly extend NormAN to feature agents that combine strategic communication with strategic inquiry, distinguishing two types of industry influence in the process. We then study the epistemic effectiveness of mixed populations, which consist partly of default NormAN agents, and partly of agents behaving according to either ‘obfuscating’ or ‘misleading’ industry strategies. We argue that our model represents these strategies’ effects in a highly optimistic context, with respect to the rationality of agents, the difficulty of the epistemic situations, and the tools available to industry actors. As such, our results can only really err on the side of optimism, and any negative effects we may find give us additional reason for concern.

The Model

Just like the base-model of NormAN, our extension² is implemented in *NetLogo* (Wilensky 1999), and makes use of its *R*-extension (Thiele & Grimm 2010) and the *R*-packages *gRain* (Højsgaard 2012) and *bnLearn* (Scutari 2009).

¹See Assaad et al. (in press) for a detailed discussion of what this entails, both generally, as well as applied to this model.

²Accessible on OSF under https://osf.io/bx248/?view_only=be1fbl673f34f3a9aef9c06f1fa3f2f.




Figure 2: Sample ‘Watts-Strogatz small-world’ network of 30 agents with $k = 2$ and $p = 0.2$ Watts & Strogatz (1998). While future work may explore in detail the effects of network structure on the impacts of industry influence, Watts-Strogatz can provide a helpful baseline for initial simulations, as they tend to avoid creating highly clustered networks, as well as those featuring overly central agents.

Being a model of argument exchange, NormAN models require both a source and a topic for said arguments. To that end, they feature a ground-truth evidence distribution generated by a causal structure represented in a Bayes’ net (see Pearl 1988 and Bovens & Hartmann 2004), a directed acyclical graph with a conditional probability distribution, of which one node becomes the hypothesis variable under consideration, and a set of remaining nodes become evidence variables. In each run of a simulation, the hypothesis is initialized as true or false with a probability matching its base rate. Then, the Bayes’ net is conditionalized on the actual truth-value of the hypothesis, and evidence nodes are initialized as true or false with probabilities matching the posterior marginal distribution. Take, as an example, the simple Bayes’ net displayed in Fig. 1: When the ground-truth in a NormAN simulation is generated via the so-called ‘Asia network’ (Lauritzen & Spiegelhalter 1988), then the model selects the question whether a particular individual suffers from lung cancer (L) as the hypothesis node. In accordance with its base rate, the L -node will then stochastically be initialized as ‘true’ in about 7% of model runs, and as ‘false’ otherwise. Consequently, the marginal probabilities of other, causally related variables will change, such as the probability that the individual is a smoker (S), or shows abnormal results when subjected to a lung X-ray (X). In runs in which the individual in question actually does suffer from lung cancer, it is highly likely (but not guaranteed), that they are also a smoker and show abnormal X-ray results. These evidence variables are thus stochastically initialized as true or false using the updated probabilities, to generate a distribution of evidence pieces that are causally relevant to the truth value of the hypothesis.

NormAN worlds are inhabited by agents, Bayesian reasoners aware of the causal structure governing their world. They may attempt to access these generated, actual values of the evidence nodes through their inquiry, the effectiveness

of which may be adjusted via settings in the user interface. For instance, the modeller may choose how often and with what success rate agents initially inquire (pre-draws, initial-draws-chance), how likely they are to inquire at each time step (curiosity), and at which point they will stop inquiring (max-draws). When learning the truth value of a piece of evidence, agents use Bayes’ rule and their knowledge of the ground-truth Bayes’ net to update their beliefs about the hypothesis, as well as their expectations for the probabilities of yet unknown evidence pieces. Agents are arranged in a social network, with links determining the interlocutors with whom they may communicate. See Fig. 2 for an example of the kinds of networks used in this paper. During communication, agents may truthfully communicate evidence pieces as arguments pro-/ or contra the hypothesis to their interlocutors, at a throughput of one argument per time-step per agent. A variety of communication rules have been studied in the NormAN framework, but for this short paper, we will consider as a baseline only agents that assert known arguments at random. Upon receiving evidence as an argument from an interlocutor, NormAN agents again update their beliefs using Bayesian conditionalization. In doing so, they will always correctly assess the actual value of the evidence-node in question, as well as how that piece of evidence relates to all other variables encoded in the ground-truth Bayes’ net.

Our extensions

From there, we implemented representations of two distinct aspects of industry influence in science:

Firstly, ‘obfuscating’ industry influence pushes scientists to explore the least diagnostic roads of inquiry, tests, experiments, or research projects, i.e., to prioritize those evidence-variables with the lowest expected informational value, from the ex ante perspective prior to observing their actual values.

We calculate this in the following way: First, we determine the veritistic value ($|p(HYP)|$ or closeness to truth, see Goldman 1999) of the initial hypothesis belief, to set a baseline. Then, for each evidence variable, we determine the veritistic values of the posterior degrees of hypothesis-belief one would reach by conditionalizing only on its truth (or falsity, respectively) in isolation. We then simply weighted these so-called ‘singular posteriors’ by the respective marginal probabilities of that evidence-node being true/false, given the actual value of the hypothesis. This returns an ex ante expected veritistic value of updating on the value of this particular evidence node, which we can then compare to the baseline to arrive at its expected informational value.

Returning to our running example, in the ‘Asia network’, learning only whether or not the individual in question is a smoker will, in expectation, sway the community’s mean hypothesis belief towards the actual value, which incentivizes industry-sponsored research against pursuing this avenue. On the other hand, learning only whether the patient has recently visited Asia (A) has neutral expected informational value. While this information would be informative with respect to the question of whether the individual suffers from tubercu-




Figure 3: The ‘Big net’ from base NormAN, featuring hypothesis A, and evidence nodes 1 through 9, of which the first three correlate with the hypothesis, the last three anti-correlate, and the middle three are non-diagnostic.

losis (T), and so further downstream could shine light on the meaning of potentially abnormal X-ray results, the A node alone will not provide evidence for or against the hypothesis until further evidence pieces are ‘filled in’. As such, in the Asia network, both A and T are prioritized by obfuscating industry influence.

Secondly, ‘misleading’ industry research privately casts a wide net, but publishes only those (factual) results that stochastically happen to be misleading (similar to Weatherall et al. 2020). In each simulation, being initialized stochastically, some (finite) number of evidence values may happen to mislead. Industry-sponsored research will only publish and re-share those misleading results. Just as previously, pieces of evidence are categorized as misleading or helpful depending on calculations of the veritistic values of the resulting singular posteriors. As an example in the context of the Asia network from Fig. 1, let us assume that a patient is indeed suffering from lung cancer, and industry actors wish to mislead the scientific community as to that fact. If in this scenario, the patient misleadingly does not happen to suffer from dyspnea (D), then asserting this will propagate through the Bayes’ net and decrease the marginal expectations of the patient suffering from Bronchitis (B), being a smoker (S) and, indeed, suffering from lung cancer (L), in turn.

Because the ‘meanings’ of arguments in NormAN are determined by the causal relationships encoded in the underlying Bayes’ nets, they may fulfil complex, strategic roles in the scientific discourse. As such, agents using obfuscating or misleading strategies, if they evaluate arguments in appropriate depth, will automatically employ ‘distracting causes’ (or ‘explaining away’), one of the main mechanisms of industry influence in science, as identified in (Freeborn & O’Connor 2024). For example, in the Asia network, the tuberculosis- and lung cancer nodes form a collider (or ‘common effect’) substructure together with the node which represents having either disease (E), and which in turn is a direct parent of the node representing (ab-)normal X-ray results. As such, in cases where industry influence seeks to mislead about the reality of the patient suffering from lung cancer, it may do so by asserting the ‘distracting cause’ tuberculosis to ‘explain away’ such X-ray results. The possibility of this exact example will, of course depend on the stochastic initialization of the involved nodes such that L, T, X are each true. We will




Figure 4: Effects of 25% obfuscating industry agents on the polarization and mean error of a community of reasoners over time. These agents still perform inquiry as usual, but their efforts are directed at nodes with low expected informational value.

further explore these and other complex group-dialectic dynamics involved in industry influence in future work, but limit the scope of this paper to the initial exploration of the dynamics of our extension.

It should also be noted that, despite their names and the intentions behind them, we model these strategies as labouring under extremely optimistic constraints: Agents may point to misleading evidence or prioritize less helpful research, but will never report falsified results. Instead, because the scientific community in our model labours to exhaust a shared finite space of possible experiments (unlike, for instance, in O’Connor & Weatherall 2019, Weatherall et al. 2020, or Holman & Bruner 2017, where agents may keep generating misleading evidence), even industry-influenced inquiry ultimately works to fill in missing evidence for the overall community. We also forewent the implementation of mechanisms for Matthew effects (see Heesen & Romeijn 2019, unlike Holman & Bruner 2017), which might otherwise exacerbate the impacts of industry influence through self-reinforcing dynamics. And lastly, nodes tend to be encoded into Bayes’ nets by virtue of their being causally relevant, so opportunities for obfuscation are rather tempered in the model, compared to the real world. Taken together, these optimistic features make our model that of a ‘best case scenario’ of industry influence from the perspective of the pursuit of truth; and as such, so are any of our results about their impacts on epistemic performance. Still, while we expect industry influence to have more pronounced effects in the real world, where these idealizations are not present to reign it in, the primary aim of this paper is not to give a quantitative estimation of effect sizes, but rather to study the plausibility of and processes generated by industry influence in a context of argument exchange.

We combine these extensions with the implicitly typed agents introduced to NormAN in (Schöpl & Hahn 2024), which allows us to study mixed populations consisting of default agents and those influenced by industry interests.

Simulation experiments

Setup

Using the thus extended model, we performed explorative computational experiments to study the epistemic performance and discourse behaviour of a community of scientists, contrasting the presence and absence of both obfuscating and misleading industry influence. To that end, we measured primarily two values, namely polarization, as the average distance of agents' beliefs $p(HYP)$ from the population mean (see Angere & Olsson 2017),

$$\sqrt{\frac{1}{n} \sum (p_n(HYP) - p_{mean}(HYP))^2}$$

as well as mean error, understood as the average distance of agents' beliefs from the 'optimal posterior' ($p_{opt}(HYP)$), i.e., the belief reached by agents upon conditionalizing on all the pieces of evidence available in the world of a particular simulation run.

$$\sqrt{\frac{1}{n} \sum (p_n(HYP) - p_{opt}(HYP))^2}$$

To make transparent the 'rise and fall of arguments' that results from our interventions, we track the relative frequencies of assertions of pieces of evidence over time, to observe in detail the (effects on the) flow of information through the social network.

For the simulations used in the results presented in this paper, we selected a setup in which the modelled scientific community is facing an appropriately challenging epistemic situation, such that figuring out the true state of the world is neither trivial nor exceptionally hard. This requires creating a balance between the amount of information available to the community (in terms of, e.g., its size and ease of inquiry) and the amounts required to reach the optimal posterior (in terms of, e.g., the total number and diagnosticity of evidence pieces). Our simulations thus feature 30 random-sharing agents arranged in randomly generated Watts-Strogatz networks (Watts & Strogatz 1998), with a mean degree of 2 and a rewiring probability of 0.2. Default agents have an initial draw chance of 10%, a curiosity of 1%, up to max-draws of 1. Ground-truth evidence in our simulations is generated by the 'Big net' DAG included in the base-model (see Fig. 3).

Note that because we are particularly interested in cases where industry agents are such by virtue not of intentional fraud or malice, but by more subtle selection mechanisms, we include the posteriors of all agents equally in our analysis of convergence speed, verisimilitude, and polarization. At 25% industry agents, we selected a number that is significant, while still highly unlikely to fully separate groups of non-industry agents merely by virtue of their random distribution throughout the social network.

Results

We will first consider the effects of the milder intervention, obfuscating industry influence. Fig. 4 contrasts the performance of the baseline population to one in which a quarter of agents instead prioritize low expected informational value nodes in their inquiry, and later communication. Both populations start in very similar positions, the former with slightly higher polarization, and the latter with slightly higher mean error. From there, the baseline population mean error quickly descends, and polarization, too, declines steadily. This behaviour is what we expected when modelling random-sharing populations in the NormAN-framework: While the population initially does not have access to all relevant pieces of evidence, and the ones already unearthed are yet to be widely shared, both will reliably be worked out over the first 100 time steps. Compared to that, the introduction of obfuscating industry agents will, starting from around step 10, increase both mean error and polarization, neither of which will ever quite be resolved over the course of the observed time span.

The mechanism behind this becomes clear when one takes a look at the rise and fall of arguments as displayed in the sample run of Fig. 5: When a Bayes' net features nodes that are, taken individually, not diagnostic with respect to the hypothesis under consideration, obfuscating industry influence shifts the inquiry and communicative resources of their agents towards an inefficient path. The resulting overrepresentation of least diagnostic arguments then clogs up communication channels and, at least initially, serves to derail collective inquiry. By virtue of their being randomly placed throughout a network where agents have two communication partners on average, obfuscating industry agents will block off the further distribution of highly informative parcels, even as they receive them from their interlocutors. Given the optimistic assumptions encoded in our implementation of this strategy, these effects will be resolved in the medium term, as non-industry agents, undeterred, will continue to inquire into the missing pieces of evidence and eventually distribute them through the community.

Next, let us take a look at the effects of the stronger, misleading industry influence. Mirroring the previous case study, Fig. 6 displays the results of introducing 25% misleading industry actors to the same baseline population. Where the mean error of the baseline population still starts at around $\frac{1}{3}$ and rapidly declines, misleading industry influence adds to that throughout the entirety of the observed 100 time steps, roughly doubling it between steps 10 and 50. After that, the relative difference increase becomes starker, as the baseline population tends to converge to the optimal posterior around 60 to 70 steps. Mean polarization, too, is consistently worsened by misleading industry influence, starting at more than double the baseline, and remaining elevated throughout the observed timespan. These results are tempered still by the fact that they include those runs in which no evidence happens to be initialized as misleading, so that industry agents can only stand idly by. In the simulations considered here,




Figure 5: Sample run of 25% obfuscating industry actors being introduced to the standard population, and observed for 50 time steps. As displayed in the monitors on the bottom right, this run features a false hypothesis, and a correspondingly low optimal posterior in the hypothesis of around 4%. Of the 9 evidence nodes in ‘Big net’, nodes ‘four’, ‘five’, and ‘six’ are characterized as providing no expected informational value. It is thus these three arguments that obfuscating agents will prioritize in their inquiry and communication, which in this run caused arguments ‘four’ and ‘six’ (represented by the brown and darker green graphs in the upper plot) to be overrepresented with respect to the frequency with which they were asserted across the social network, especially in the first third of the runtime. This diffusion of undiagnostic arguments served to lead the population astray, causing an initial rise in mean error as observable in the bottom left plot.




Figure 6: Effects of 25% misleading industry agents on the polarization and mean error of a community of reasoners over time. This includes those 20% of runs in which no misleading evidence was available in the world.

this was the case in roughly 20% of runs.

To showcase the underlying mechanism, we included another sample run, displayed in Fig. 7. As evidence is stochastically generated using the marginal probabilities resulting from conditionalizing the Bayes’ net on the actual truth value of the hypothesis, each diagnostic piece will tend to be epistemically helpful in expectation. However, in each particular run, any number of evidence pieces may be outliers to this trend and be initialized as misleading. It is that set of evidence that misleading industry actors will have immediate access to and will exclusively share with the rest of the population, while non-industry actors still work to slowly unearth the remaining evidence. This causes an initial flood of misleading arguments across the social network, which may push all agents to uniformly erroneous hypothesis beliefs in the medium run. If the mean beliefs of even non-industry scientists were to be sampled by outside agents at any of these points in time, industry interests rather than the pursuit of truth would be served. Even so, in our optimistic implementation, communicating misleading results still eventually contributes to the collective checking off all missing pieces




Figure 7: Sample run of 25% misleading industry actors being introduced to the standard population, and observed for 50 time steps. As displayed in the monitors on the bottom right, this run features a true hypothesis, and a correspondingly high optimal posterior in the hypothesis of around 89%. Of the 9 evidence nodes in ‘Big net’, nodes ‘two’ and ‘three’ are misleadingly initialized as false. Misleading agents will thus exclusively propagate the corresponding parcels (colored orange and red, respectively, in the upper plot), causing them to be heavily overrepresented in the frequency with which each argument is asserted over time. This quick diffusion of misleading arguments through the network works to initially lead all of the population astray, causing the high mean error and low polarization displayed in the bottom left.

of evidence, and will thus help the population approach the optimal posterior in the long run. This is because misleading industry influence in our model does not increase the amount of misleading evidence in the world; it merely directs initial attention towards it.

Conclusion

We extended the Bayesian, agent-based modelling framework NormAN to represent mixed populations of scientists, partly consisting of agents whose inquiry and communication are governed by industry interests. We considered obfuscating industry agents, which direct their efforts towards experiments with ex ante low expected informational value for the scientific community, as well as misleading agents, who direct their efforts towards the diffusion of experimental results that, ex post, have happened to come out as misleading.

In preliminary computational experiments, we observed both types of industry influence to worsen the epistemic performance of the population, with misleading influence having a stronger impact. These results obtain despite the fact that agents in our model are perfectly Bayesian reasoners aware of the causal structure which generates their evidence, and despite the fact that even misleading agents in our model never

falsify facts, never increase the number of non-diagnostic or misleading evidence otherwise available in the world, do not cause other agents to adapt their strategies, and instead, in the long run, even help the population unearth the entirety of evidence.

This leads us to the preliminary conclusion that industry influence remains a plausible threat to the epistemic performance of science. In particular, we are concerned with the repercussions this implies for questions of competition in science, since a crucial avenue for industry influence to enter the picture is via the competitive edge that industry funding may promise researchers harbouring under high-pressure resource allocation mechanisms (see also Holman & Bruner 2017).

Finally, we expect that moving agent-based models of the tobacco strategy into a framework in which agents exchange complex arguments will provide a stepping stone for further study of the phenomena at hand, also when eventually considering de-idealizations that deviate from the ‘best-case’ scenarios studied here, as well as more extensive simulation studies.

Acknowledgements

Klee Schöppel was supported by the research program Sustainable Cooperation – Roadmaps to Resilient Societies (SCOOP). They are grateful for funding from the Netherlands Organization for Scientific Research (NWO) and the Dutch Ministry of Education, Culture and Science (OCW) in the context of its 2017 Gravitation Program (grant number 024.003.025).

References

- Angere, S., & Olsson, E. J. (2017). Publish late, publish rarely!: Network density and group performance in scientific communication. In *Scientific collaboration and collective knowledge* (pp. 34–62). Oxford University Press.
- Assaad, L., Fuchs, R., Jalalimanesh, A., Phillips, K., Schöppel, K., & Hahn, U. (2023). A bayesian agent-based framework for argument exchange across networks. *arXiv preprint arXiv:2311.09254*.
- Assaad, L., Fuchs, R., Phillips, K., Schöppel, K., & Hahn, U. (in press). Capturing argument in agent-based models. *TOPOI*.
- Assaad, L., & Hahn, U. (2024). Rational polarization: Sharing only one's best evidence can lead to group polarization. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 46).
- Bala, V., & Goyal, S. (1998). Learning from neighbours. *The review of economic studies*, 65(3), 595–621.
- Bovens, L., & Hartmann, S. (2004). *Bayesian epistemology*. OUP Oxford.
- Freeborn, D., & O'Connor, C. (2024). Industrial distraction.
- Goldman, A. I. (1999). *Knowledge in a social world*. Oxford University Press.
- Hahn, U., Assaad, L., & Burton, J. W. (2024). Opinion averaging versus argument exchange. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 46).
- Heesen, R., & Romeijn, J.-W. (2019). Epistemic diversity and editor decisions: A statistical matthew effect. *Philosophers' Imprint*, 19(39), 1–20.
- Højsgaard, S. (2012). Graphical independence networks with the grain package for r. *Journal of Statistical Software*, 46, 1–26.
- Holman, B., & Bruner, J. (2017). Experimentation by industrial selection. *Philosophy of Science*, 84(5), 1008–1019.
- Lauritzen, S. L., & Spiegelhalter, D. J. (1988). Local computation with probabilities on graphical structures and their application to expert systems (with discussion). *Journal of the Royal Statistical Society. Series B (Methodological)*, 50(2), 157–224. Retrieved from <https://www.jstor.org/stable/2345762>
- O'Connor, C., & Weatherall, J. O. (2019). *The misinformation age: How false beliefs spread*. Yale University Press.
- Oreskes, N., & Conway, E. M. (2011). *Merchants of doubt: How a handful of scientists obscured the truth on issues from tobacco smoke to global warming*. Bloomsbury Publishing USA.
- Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan kaufmann.
- Schöppel, K., & Hahn, U. (2024). Exploring effects of self-censoring through agent-based simulation. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 46).
- Scutari, M. (2009). Learning bayesian networks with the bnlearn r package. *arXiv preprint arXiv:0908.3817*.
- Thiele, J. C., & Grimm, V. (2010). Netlogo meets r: Linking agent-based models with a toolbox for their analysis. *Environmental Modelling & Software*, 25(8), 972–974.
- Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *nature*, 393(6684), 440–442.
- Weatherall, J. O., O'Connor, C., & Bruner, J. P. (2020). How to beat science and influence people: Policymakers and propaganda in epistemic networks. *The British Journal for the Philosophy of Science*.
- Wilensky, U. (1999). Netlogo. *Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.* Retrieved from <https://ccl.northwestern.edu/netlogo/>
- Zollman, K. J. (2007). The communication structure of epistemic communities. *Philosophy of science*, 74(5), 574–587.
- Zollman, K. J. (2010). The epistemic benefit of transient diversity. *Erkenntnis*, 72(1), 17–35.