

# The Cultural Evolution of Science

Jingyi Wu<sup>1</sup>, Cailin O'Connor<sup>1</sup>, and Paul E. Smaldino<sup>2</sup>

<sup>1</sup>Department of Logic and Philosophy of Science, University of California, Irvine

<sup>2</sup>Department of Cognitive and Information Sciences, University of California, Merced

June 20, 2022

To appear in *The Oxford Handbook of Cultural Evolution*, edited by Jeremy Kendal, Rachel Kendal, and Jamshid Tehrani, Oxford University Press. Expected 2023.

## 1 Introduction

Science is a cultural practice, and cultural practices tend to change over time via processes of cultural selection and social learning. There is a long history of philosophers of science arguing that scientific theories evolve through a “critical” evolutionary process where new hypotheses are criticized, modified, eliminated, or replaced (Popper 1972; Hull 1988). More recent work has suggested that other features of science such as methodologies, beliefs, and norms may develop likewise. Such features of science exhibit key characteristics that make them suitable for evolutionary analysis. They are reliably transmitted via pedagogy and cultural imitation, and produce non-random variation that leads to differential success in subsequent transmission. For this reason, a new body of work has emerged looking at cultural evolutionary processes in science. This research addresses topics ranging from the persistence of poor statistical practices, to conservatism in science, to the ideal communication structure for scientific communities.

In this chapter, we offer a brief review of formal models in which science is treated as a cultural evolutionary system. The field of cultural evolution has a rich theoretical foundation. Cultural evolutionary models help elucidate how patterns of cultural behavior emerge via appeal to mechanisms of cultural transmission and social learning. Here we

highlight learning and transmission processes particular to scientific communities. These include scientists choosing methods for hypothesis testing, competing for jobs and resources, learning from other experts, attempting to accrue credit for their work, and assigning credit for work done. Throughout, we connect our discussion to empirical research on the workings of science. Because this is a very recently developed area of cultural evolutionary science, though, there is relatively little new empirical work derived from these models. We hope they may inspire both future empirical investigations and further theoretical work on the cultural evolution of science.

We divide the models we review into three categories. We first consider “selective” models in which scientific practices are transmitted preferentially by successful individuals. We then consider models that focus on the dynamics of scientific beliefs rather than on the methods used to produce them. These models consider how social network structures shape communication in science, and thus lead to more or less accurate beliefs. Third, we look at models tracking how social identity can impact the norms and structures of science. We conclude with a discussion of the role of cultural evolutionary models in our understanding of science.

## **2 Cultural Evolution of Science Via Selection**

Scientists are ideally motivated by the quest for truth, and should (the story goes) unflinchingly adopt beliefs and practices that directly aid in that quest. In practice, though, scientists are fallible humans that respond to a variety of incentives, particularly those that directly contribute to their ability to work and thrive in their chosen profession. Institutional incentives to publish, secure funding, and obtain prestige can therefore shape scientific practice just as much as the search for truth.

These kinds of incentives impact scientific communities in at least two ways. First, individual scientists may actively respond to incentives by strategically changing their research practices. For instance, a researcher who seeks tenure may choose to work on safe, publishable topics rather than something riskier. There is a large literature on how individual

scientists seek “credit” in this way, and how this impacts scientific communities (Merton 1973).<sup>1</sup> Second, selective dynamics can occur when scientists who accrue institutional credit become more likely to transmit their practices to other researchers. This can occur without any deliberate or strategic adjustment of one’s practices. This might happen, for example, when a famous scientist places many students who use their methods in prominent jobs, or when scientists copy the methods of famous colleagues. The papers reviewed here focus primarily on this second, selective dynamic, as it embodies one form of cultural evolution in science.

Smaldino and McElreath (2016) seek to explain empirical literature showing that low statistical power continues to plague many disciplines in spite of regular calls for change (Sedlmeier and Gigerenzer 1989; Rossi 1990; Smaldino and McElreath 2016; Szucs and Ioannidis 2017). They introduce a model where labs test scientific hypotheses using methods of varying quality.<sup>2</sup> Each lab has a characteristic methodological power, representing the ability to correctly identify a true hypothesis. Increasing power also increases the rate of false positives—the probability that a false hypothesis is incorrectly labeled as true—unless effort is exerted to avoid this. However, exerting additional effort also increases the time between results. The model also assumes that negative results are harder to publish than positive results and, critically, that labs that publish more are more likely to pass along their methods, perhaps through graduate students starting new labs with the same methods. These conditions ensure that there is variability in the quality of heritable methods in the community and a selection pressure (publication) that acts on that variability.

If effort is held constant and methodological power is allowed to evolve, power soon evolves to its maximum, along with the false positive rate. The idea is that if labs do not exert effort to curb the false discovery rate, they produce more positive results. They are then more successful in generating publications and it is more likely their methods are reproduced in subsequent generations. Unsurprisingly then, the most selectively “successful” methods are

---

<sup>1</sup>This includes work on how scientists may adopt questionable research practices in response to incentives (e.g. Higginson and Munafò 2016; Bright 2017; Heesen forthcoming).

<sup>2</sup>The model extends a previous model by McElreath and Smaldino (2015) that analyzed the population dynamics of scientific hypotheses in an idealized community of researchers, but did not consider selection. See Smaldino (2022) for a review of related models.

ones that label every hypothesis to be true. Smaldino and McElreath (2016) then show that when methodological power is held fixed and effort is allowed to evolve, it rapidly declines to its minimum value, so the successful methods are again those that have high false positive rates. Moreover, Smaldino and McElreath (2016) show that replications are not sufficient for curbing the selection of bad methods, even if negative replications carry a hefty punishment for the original lab. This is because the top performing labs are always the ones who are able to cut corners and avoid being caught. These labs have disproportionate influence on future methods adopted by the community.

This result is a powerful demonstration of what is sometimes called Campbell's Law, which states that "the more any quantitative social indicator is used for social decision-making, the more subject it will be to corruption pressures and the more apt it will be to distort and corrupt the social processes it is intended to monitor" (Campbell 1979, p. 49). Here, because there is a mismatch between what is rewarded, namely publication, and what we intend to reward, namely quality, the labs that are able to exploit this discrepancy end up the most success. The model also provides a possible explanation for why poor methods persist in many disciplines and why so many published results have failed to replicate (Collaboration 2015; Nosek and Errington 2017; Camerer et al. 2018; Ritchie 2020; Nosek, Hardwicke, et al. 2022). The model therefore presents a rather bleak picture of science. Is this picture justified? And if so, can cultural evolutionary modeling provide insights into a solution?

Several authors have built models to investigate various interventions that may help. Barnett et al. (2018) consider random audits of scientific labs. This stems the spread of poor methods but requires fairly unrealistic assumptions about the ability of auditors to accurately assess the false positive rates of individual labs, as well as the willingness of researchers to take responsibility for shutting down their corner-cutting colleagues' labs. Smaldino, Turner, et al. (2019) find that the triple threat of reduced publication bias, improved peer review, and modified funding lotteries can impede the natural selection of bad science. Reducing publication bias means that negative results are more likely to be published. Improved peer review means that false positives are more likely to be rejected for publication. (This could be facilitated via measures such as registered reports, double or

triple blind reviews, preprints, etc.) By modified funding lottery, they mean randomly allocating funding to labs that have a false positive rate below a certain threshold, thus shifting incentives. These institutional changes are consistent with recommendations from the open science movement (e.g. Munafò et al. (2017)) and proposals for funding reform (e.g. Fang and Casadevall (2016) and Gross and Bergstrom (2019)). Stewart and Plotkin (2021) take a different approach by allowing effort to be invested in theory development as well as publication. This increases the probability that the hypotheses tested by a lab are correct (Muthukrishna and Henrich 2019; Smaldino 2019). Stewart and Plotkin show that, under a wide range of conditions, such an investment in theory can change the evolutionary outcomes from bad science to good. These models all suggest possible interventions to improve methods that may be worth exploring empirically.

In another use of selective models in science Akerlof and Michailat (2018) consider the tenure process. They model academics who prefer to tenure those who work on similar topics to themselves, reflecting studies of decision-making in science (Mahoney 1977; Travis and Collins 1991; Lamont et al. 2009; Wang et al. 2017). In such a regime, this self-preferential bias shapes selective pressures to favor the status quo, even when it is a bad one. Smaldino and O'Connor (2020) extend this work to explore how interdisciplinarity can facilitate the spread of better methods in such cases. They build evolutionary models where academics are sometimes reviewed and assessed by those in neighboring fields. When this happens, their assessments reflect a different set of biases, meaning that new, better methods can receive credit and spread. Again, this points toward further research looking at the promise of interdisciplinary contact to improve methods.

Selection in science can also act on the propensity of scientists to take risks. Scholars have drawn on evidence to argue that many scientific fields are trending towards increased conservatism (e.g. Stanford 2019; Currie 2019). O'Connor (2019) presents a model to explore why this might be the case. Building on the paradigm of Smaldino and McElreath (2016), she introduces the following variations: some labs are inherently conservative, with a dependable probability of success and associated payoffs, while other labs are inherently risky, with a lower probability of success and higher payoffs when they succeed. If successful labs can stably pass

on their methods to future generations, then the selection process can favor risky science. This is because risky methods are associated with a higher variance of outcomes, and the most successful risky labs will have an out-sized influence on the makeup of the scientific community because their payoffs are higher. However, risky science is often hard to repeat. When it sufficiently difficult to replicate, then high risk, high reward science will be selected against.<sup>3</sup>

All the models described to this point focus on the endogenous credit reward system in science, but other insidious selection forces can also shape epistemic communities. (For well documented examples, see Oreskes and Conway (2011).). Holman and Bruner (2017) and Weatherall et al. (2020) analyze models of industrial impact on science. They use a paradigm where a community of scientists attempts to learn which of two hypotheses is better. The scientists test the options directly via experiment, and also learn from other scientists in their social network. Holman and Bruner (2017) assume that there is a methodological diversity in the epistemic community, so that scientists vary in their ability to track truth. They show that industrial agents who are interested in spreading a false belief can exploit this methodological diversity by selectively funding the scientists who produce results favoring industrial interests. The extra funding can then make these scientists more productive and more likely to pass on their methods to new scientists, eventually threatening the formation of true consensus in the epistemic community. They, and others, point to real cases of this sort of "industrial selection" related to the approval of anti-arrhythmic heart drugs (Holman and Bruner 2017) and research on the link between fat and heart disease (O'Connor and Weatherall 2019).

In the model from Weatherall et al. (2020) an industrial agent cherry picks research from the scientific community that supports their position. Or alternatively they fund new research, but then only publish preferred results. These models are inspired by what Oreskes and Conway (2011) call the "tobacco strategy," used by industry to influence policy makers regarding the health risks of tobacco. Weatherall et al. (2020) argue that both strategies can be extremely effective at influencing the policy makers. What's more, they find that selectively sharing existing research may bias beliefs more effectively than directly funding new research.

---

<sup>3</sup>Kummerfeld and Zollman (2020) build related (though non-evolutionary) models to argue that individual scientists have incentives against exploration at the socially optimal level, creating a free rider problem.

Most of these models consider selection acting only at the level of individuals. However, science is a cultural practice that often involves teams and groups. Like many cultural practices, group-level pressures can have important effects on the evolutionary dynamics of those practices (Henrich 2004; Richerson et al. 2016). Tiokhin et al. (2021) have recently argued that some of problematic trends in scientific communities could be avoided if the “level of selection” in science could shift from individual scientists to groups such as collaborations, labs, or departments. To motivate their argument, they forefront the role of “indirect” contributions that increase the rigor or productivity of other scientists. These include mentorship and high quality peer review. By selecting on the group level, strategies that emphasize indirect rather than direct contributions might have more success. This would have the effect of promoting pro-social behavior in science and improving community productivity.

### **3 Evolving Beliefs in Scientific Networks**

In this section, we turn from the evolution of the methods by which scientific knowledge is generated, to the evolution of the beliefs held by scientists. Although a number of factors impact beliefs in science, social transmission is a central one. Thus, the psychological and communicative features involved in the formation and transmission of scientific beliefs are at the forefront. These models are related to more general models of opinion and belief dynamics (e.g. Flache, Mäs, et al. 2017; Galesic et al. 2021), but unlike most of those models, they assume that actors seek some objective truth in the world, and use evidence and rational updating to reach it.

Social network structure is well known to have a profound influence on the dynamics of social traits (Szabó and Fath 2007; Nowak et al. 2010; Flache and Macy 2011). Thus, we focus our review here on models with explicit network structures. Nodes represent individual scientists and edges represent social connections between those scientists. These connections can thus track the myriad ways scientists influence one another’s beliefs—direct communication, collaboration, publication of research findings, and refereeing.

Zollman (2007, 2010) studies agent-based models in which networked scientists attempt to figure out which of the two available hypotheses are better. This is modeled as a “two-armed bandit” problem—each arm succeeds with a characteristic probability, with one more successful than the other. The goal is to figure out which is better. Agents can learn either by testing the options themselves or by observing evidence from their community. They adjust their beliefs about the actions using Bayesian updating.

Zollman finds that a more sparsely connected network is more likely to converge to a belief in the better hypothesis. Although it may seem that more communication would always be better, this is not the case. In a more tightly connected community, a series of misleading data can lead the community to converge on the false belief prematurely. Sparse connection prevents this outcome by preserving pockets of diverse practice. The tradeoff is that convergence to consensus occurs more slowly in the sparse network.

This result echoes findings from other models of problem solving and innovation on networks, which indicate that, when teams of individuals need to solve difficult problems, more sparsely connected networks can outperform more well connected networks (Lazer and Friedman 2007; Derex and Boyd 2016; Galesic et al. 2021; Gomez and Lazer 2019; Migliano et al. 2020; Boroomand and Smaldino 2021). In these models, when agents rapidly adopt superior solutions by learning from network neighbors, high connectivity leads to a rapid dissolution of diversity that squashes the network’s ability to innovate. In contrast, lower connectivity or temporary isolation facilitates a diversity of possible solutions, and is more likely to preserve the best ones. Empirical studies of problem solving on networks mostly confirm these results (Mason, Jones, et al. 2008; Jönsson et al. 2015; Derex and Boyd 2016). (But see Mason and Watts (2012) for disconfirmatory findings.)

There are other mechanisms that can help a community explore a diversity of options long enough to ensure they reach successful consensus. Zollman (2010), for example, models “stubborn” scientists who are less likely to switch hypotheses when presented with new evidence. This means that competing hypotheses are more accurately assessed before a community settles onto a consensus. Relatedly, Gabriel and O’Connor (2022) show that mild “confirmation bias,” in which scientists preferentially engage with evidence that fits their



prior beliefs, can be beneficial to the evolution of true beliefs in a community. Like stubbornness, confirmation bias makes scientists to stick to their theories for longer. However both Zollman (2010) and Gabriel and O'Connor (2022) find that too much intransigence is bad thing, preventing a community from ever researching consensus.

O'Connor and Weatherall (2018) extend the modeling paradigm introduced in Zollman (2007) to study the effects of belief-based distrust in scientific communities. They assume that agents are less trustful of evidence that comes from scientists who hold very different beliefs.<sup>4</sup> Under these conditions, epistemic communities often become polarized. They divide into two camps where each prefers one of the two hypotheses, and mutual mistrust prevents them from influencing each other further.

Other authors have studied models of epistemic communities with multiple subgroups, and explored how beliefs evolve as a result of interactions between them.<sup>5</sup> Wu (2022) builds a model where one group ignores or devalues evidence from the other. For instance, a dominant social group might ignore contributions from a marginalized one.<sup>6</sup> Wu finds that the marginalized group tends to end up with better beliefs. The dominant group updates on less evidence, and thus tests a variety of hypotheses for longer. However, it is the marginalized group that benefits from this exploration, while the dominant group does not learn from the marginalized group's own exploration.

Fazelpour and Steel (2022) consider two mutually mistrusting subgroups. They find that mild inter-group distrust can actually improve consensus. Similar to the stubbornness and confirmation bias cases mentioned above, with inter-group distrust agents update their beliefs more slowly, making it more likely that they explore an option sufficiently before reaching consensus. Fazelpour and Steel take their models to support the view that demographic diversity can improve group learning, a result also echoed by more general models of collective problem solving (Gomez and Lazer 2019; Boroomand and Smaldino 2021).

---

<sup>4</sup>This is similar to "bounded confidence" models of opinion dynamics (Hegselmann, Krause, et al. 2002). But in these previous models, there is no truth of the matter as to which opinion is best. This means that they are not particularly good at representing scientific communities.

<sup>5</sup>The model by Smaldino and O'Connor (2020) discussed in the previous section is of this type.

<sup>6</sup>This draws on work tracking testimonial injustice (Fricker 2007) and epistemic quieting (Dotson 2011).

## 4 Diversity and the Evolution of Discrimination in Science

The models reviewed to this point have investigated how selection and cultural evolution in science impact learning in scientific communities. All this work has focused on explaining epistemic failures and successes in science (and thereby improving future communities). We turn now to models considering norms of bargaining, communication, and credit giving in science, and especially how social identity can impact the cultural evolution of these practices.

Bruner (2019) uses evolutionary game theoretic models to consider how members of a minority group can be disadvantaged in bargaining scenarios. Differences in the size of communities mean that minorities meet out-group members more often than majority group members. This in turn means they tend to learn how to interact with their out-group more quickly, and thus learn accommodating bargaining strategies more frequently.<sup>7</sup> O'Connor and Bruner (2019) point out that bargaining is a key feature of academic collaboration—whether implicitly or explicitly, collaborators must decide who does what work and who gets what credit. They argue that minority groups may be at extra risk of credit disadvantage as a result of this bargaining process. This may help explain empirical findings on why female academics often receive less credit for collaborative work (Sarsons et al. 2021), and at least one lab based study replicates this pattern (Mohseni et al. 2021).

Rubin and O'Connor (2018) look at how discrimination in science can impact the evolution of collaboration. They model two social identity groups, and assume that one may discriminate against the other when bargaining over collaborative credit. They assume that each individual may choose their collaborative partners, and find that the disadvantaged group can learn to avoid the discriminators as a result. Their results may help explain empirical findings on gender based homophily in scientific collaboration (Ferber and Teiman 1980), and also why academics of color may be less likely to choose collaborative arrangements (Del Carmen and Bing 2000).

---

<sup>7</sup>Ethnographically-inspired modeling indicating that a strategy of adopting majority norms when appropriate ("cross-cultural competence") may actually help minority groups maintain their distinctive cultures (Bunce 2021).

Rubin (2022) applies evolving network models to explore how gender citation gaps can emerge due to structural factors. Specifically, she looks at a model with two subgroups of unequal size and a homophilic network structure, meaning that agents prefer to connect with their in-group members over out-group. On the reasonable assumption that academics prefer to cite those they are more closely connected with, citation gaps between the two groups reliably emerge, where minority members are cited less often than members of majority members. This model can be taken to possibly support empirical findings about gender and racial citation gaps in academia (e.g. Ferber (1986) and Teich et al. (2021)). Similarly Rubin and Schneider (2021) investigate how the "priority rule" can disadvantage minority groups. This rule mandates that in cases of multiple discovery, credit and prestige is awarded to the one who made the discovery first (Merton 1957). In their model, academics from underrepresented groups tend to be less connected because of social identity, and because they tend to have joined the academic network more recently. When multiple discoveries are made, academics give credit to the academic closest in the network to themselves, thus disadvantaging those from underrepresented groups on average. This, again, may help explain identity based citations gaps.

## 5 Conclusion

The dynamic models reviewed in this chapter reveal insights about science that static models (such as those in which rational agents focus on maximizing either accuracy or credit) typically do not. This is because evolutionary thinking is, at its core, *population thinking*. The quality of a strategy, and its associated payoff, depends on the behaviors of others in a population, the structure of social interactions, and the complex ways in which those behaviors and relations change over time. Formal models allow us to examine the consequences of assumptions about the way things are and the way those things change (Smaldino 2017). Cultural evolutionary theories provide a theoretical framework that helps us to ensure that those assumptions are well justified and not merely ad hoc.

Of course, there is no guarantee that the assumptions used in these and other models

are always met by the empirical systems one hopes to represent with them. Models must be tested against empirical data for validation and refinement. But models can also help us to understand which data are most important to gather. Foster and Evans (2019), for example, recently used models to anchor a discussion of how one might use cultural phylogenetics to understand the lineage of technologies in the patent record, with the models indicating the data one would need to test their ideas.

An understanding of cultural evolution and how to model its processes can and should help us do better science. At one level, cultural evolution provides an unparalleled framework for understanding social behavior and cultural dynamics, and so it is valuable to any discipline that tackles those factors. *All* scientists are engaged in the cultural system of science itself. By modeling the dynamic processes by which our norms, practices, and beliefs emerge and evolve, we can discover how to improve them, and in doing so, improve science.

## References

- Akerlof, George A and Pascal Michaillat (2018). "Persistence of false paradigms in low-power sciences". In: *Proceedings of the National Academy of Sciences* 115.52, pp. 13228–13233.
- Barnett, Adrian G, Pauline Zardo, and Nicholas Graves (2018). "Randomly auditing research labs could be an affordable way to improve research quality: A simulation study". In: *PLOS ONE* 13.4, e0195613.
- Boroomand, Amin and Paul E Smaldino (2021). "Hard work, risk-taking, and diversity in a model of collective problem solving". In: *Journal of Artificial Societies and Social Simulation* 24.4.
- Bright, Liam Kofi (2017). "On fraud". In: *Philosophical Studies* 174.2, pp. 291–310.
- Bruner, Justin (2019). "Minority (dis) advantage in population games". In: *Synthese* 196.1, pp. 413–427.
- Bunce, John A (2021). "Cultural diversity in unequal societies sustained through cross-cultural competence and identity valuation". In: *Humanities and Social Sciences Communications* 8.1, pp. 1–9.
- Camerer, Colin F et al. (2018). "Evaluating the replicability of social science experiments in Nature and Science between 2010 and 2015". In: *Nature Human Behaviour* 2.9, pp. 637–644.
- Campbell, Donald T (1979). "Assessing the impact of planned social change". In: *Evaluation and program planning* 2.1, pp. 67–90.
- Collaboration, Open Science (2015). "Estimating the reproducibility of psychological science". In: *Science* 349.6251, aac4716.
- Currie, Adrian (2019). "Existential risk, creativity & well-adapted science". In: *Studies in History and Philosophy of Science Part A* 76, pp. 39–48.
- Del Carmen, Alejandro and Robert L Bing (2000). "Academic productivity of African Americans in criminology and criminal justice". In: *Journal of Criminal Justice Education* 11.2, pp. 237–249.
- Derex, Maxime and Robert Boyd (2016). "Partial connectivity increases cultural accumulation within groups". In: *Proceedings of the National Academy of Sciences* 113.11, pp. 2982–2987.
- Dotson, Kristie (2011). "Tracking Epistemic Violence, Tracking practices of Silencing". In: *Hypatia* 26.2, pp. 236–257.
- Fang, Ferric C and Arturo Casadevall (2016). *Research funding: The case for a modified lottery*.
- Fazelpour, Sina and Daniel Steel (2022). "Diversity, Trust, and Conformity: A Simulation Study". In: *Philosophy of Science* 89.2, pp. 209–231. DOI: 10.1017/psa.2021.25.
- Ferber, Marianne A (1986). "Citations: Are they an objective measure of scholarly merit?" In: *Signs: Journal of Women in Culture and Society* 11.2, pp. 381–389.
- Ferber, Marianne A and Michelle Teiman (1980). "Are women economists at a disadvantage in publishing journal articles?" In: *Eastern Economic Journal* 6.3/4, pp. 189–193.
- Flache, Andreas and Michael W Macy (2011). "Small worlds and cultural polarization". In: *The Journal of Mathematical Sociology* 35.1-3, pp. 146–176.
- Flache, Andreas, Michael Mäs, et al. (2017). "Models of social influence: Towards the next frontiers". In: *Journal of Artificial Societies and Social Simulation* 20.4.
- Foster, Jacob G and James A Evans (2019). "Promiscuous inventions: Modeling cultural evolution with multiple inheritance". In: *Beyond the Meme: Development and Structure*

- in Cultural Evolution—Minnesota Studies in Philosophy of Science*. Ed. by AC Love and W Wimsatt. University of Minnesota Press.
- Fricker, Miranda (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press.
- Gabriel, Nathan and Cailin O'Connor (2022). *Can Confirmation Bias Improve Group Learning?* URL: <http://philsci-archive.pitt.edu/20528/>.
- Galesic, Mirta et al. (2021). "Integrating social and cognitive aspects of belief dynamics: towards a unifying framework". In: *Journal of the Royal Society Interface* 18.176, p. 20200857.
- Gomez, Charles J and David Lazer (2019). "Clustering knowledge and dispersing abilities enhances collective problem solving in a network". In: *Nature Communications* 10, p. 5146.
- Gross, Kevin and Carl T Bergstrom (2019). "Contest models highlight inherent inefficiencies of scientific funding competitions". In: *PLoS biology* 17.1, e3000065.
- Heesen, Remco (forthcoming). "Cumulative Advantage and the Incentive to Commit Fraud in Science". In: *The British Journal for the Philosophy of Science*. ISSN: 1464-3537. DOI: 10.1086/716235. URL: <https://doi.org/10.1086/716235>.
- Hegselmann, Rainer, Ulrich Krause, et al. (2002). "Opinion dynamics and bounded confidence models, analysis, and simulation". In: *Journal of Artificial Societies and Social Simulation* 5.3.
- Henrich, Joseph (2004). "Cultural group selection, coevolutionary processes and large-scale cooperation". In: *Journal of Economic Behavior & Organization* 53.1, pp. 3–35.
- Higginson, Andrew D and Marcus R Munafò (2016). "Current incentives for scientists lead to underpowered studies with erroneous conclusions". In: *PLoS Biology* 14.11, e2000995.
- Holman, Bennett and Justin Bruner (2017). "Experimentation by industrial selection". In: *Philosophy of Science* 84.5, pp. 1008–1019.
- Hull, David L (1988). *Science as a process: an evolutionary account of the social and conceptual development of science*. University of Chicago Press.
- Jönsson, Martin L, Ulrike Hahn, and Erik J Olsson (2015). "The kind of group you want to belong to: Effects of group structure on group accuracy". In: *Cognition* 142, pp. 191–204.
- Kummerfeld, Erich and Kevin JS Zollman (2020). "Conservatism and the scientific state of nature". In: *The British Journal for the Philosophy of Science*.
- Lamont, Michèle et al. (2009). *How professors think*. Harvard University Press.
- Lazer, David and Allan Friedman (2007). "The network structure of exploration and exploitation". In: *Administrative science quarterly* 52.4, pp. 667–694.
- Mahoney, Michael J (1977). "Publication prejudices: An experimental study of confirmatory bias in the peer review system". In: *Cognitive Therapy and Research* 1.2, pp. 161–175.
- Mason, Winter A, Andy Jones, and Robert L Goldstone (2008). "Propagation of innovations in networked groups." In: *Journal of Experimental Psychology: General* 137.3, p. 422.
- Mason, Winter A and Duncan J Watts (2012). "Collaborative learning in networks". In: *Proceedings of the National Academy of Sciences* 109.3, pp. 764–769.
- McElreath, Richard and Paul E Smaldino (2015). "Replication, communication, and the population dynamics of scientific discovery". In: *PloS one* 10.8, e0136088.
- Merton, Robert K (1957). "Priorities in scientific discovery: a chapter in the sociology of science". In: *American sociological review* 22.6, pp. 635–659.
- (1973). *The sociology of science: Theoretical and empirical investigations*. University of Chicago press.
- Migliano, Andrea B et al. (2020). "Hunter-gatherer multilevel sociality accelerates cumulative cultural evolution". In: *Science Advances* 6.9, eaax5913.

- Mohseni, Aydin, Cailin O'Connor, and Hannah Rubin (2021). "On the emergence of minority disadvantage: testing the cultural Red King hypothesis". In: *Synthese* 198.6, pp. 5599–5621.
- Munafò, Marcus R et al. (2017). "A manifesto for reproducible science". In: *Nature human behaviour* 1.1, pp. 1–9.
- Muthukrishna, Michael and Joseph Henrich (2019). "A problem in theory". In: *Nature Human Behaviour* 3.3, pp. 221–229.
- Nosek, Brian A and Timothy M Errington (2017). "Reproducibility in cancer biology: Making sense of replications". In: *eLife* 6, e23383.
- Nosek, Brian A, Tom E Hardwicke, et al. (2022). "Replicability, robustness, and reproducibility in psychological science". In: *Annual Review of Psychology* 73.1, pp. 719–748.
- Nowak, Martin A, Corina E Tarnita, and Tibor Antal (2010). "Evolutionary dynamics in structured populations". In: *Philosophical Transactions of the Royal Society B* 365.1537, pp. 19–30.
- O'Connor, Cailin (2019). "The natural selection of conservative science". In: *Studies in History and Philosophy of Science Part A* 76, pp. 24–29.
- O'Connor, Cailin and Justin Bruner (2019). "Dynamics and diversity in epistemic communities". In: *Erkenntnis* 84.1, pp. 101–119.
- O'Connor, Cailin and James Owen Weatherall (2018). "Scientific polarization". In: *European Journal for Philosophy of Science* 8.3, pp. 855–875.
- (2019). "How Powerful Interests Use Science to Sway Public Opinion". In: *Zocalo Public Square*. URL: <https://www.zocalopublicsquare.org/2019/09/05/how-powerful-interests-use-science-to-sway-public-opinion/ideas/essay/>.
- Oreskes, Naomi and Erik M Conway (2011). *Merchants of doubt: How a handful of scientists obscured the truth on issues from tobacco smoke to global warming*. Bloomsbury Publishing USA.
- Popper, Karl R (1972). *Objective knowledge*. Oxford University Press Oxford.
- Richerson, Peter et al. (2016). "Cultural group selection plays an essential role in explaining human cooperation: A sketch of the evidence". In: *Behavioral and Brain Sciences* 39, e30.
- Ritchie, Stuart (2020). *Science fictions: How fraud, bias, negligence, and hype undermine the search for truth*. Metropolitan Books.
- Rossi, Joseph S (1990). "Statistical power of psychological research: What have we gained in 20 years?" In: *Journal of Consulting and Clinical Psychology* 58.5, pp. 646–656.
- Rubin, Hannah (2022). "Structural causes of citation gaps". In: *Philosophical Studies*, pp. 1–23.
- Rubin, Hannah and Cailin O'Connor (2018). "Discrimination and collaboration in science". In: *Philosophy of Science* 85.3, pp. 380–402.
- Rubin, Hannah and Mike D Schneider (2021). "Priority and privilege in scientific discovery". In: *Studies in History and Philosophy of Science Part A* 89, pp. 202–211.
- Sarsons, Heather et al. (2021). "Gender differences in recognition for group work". In: *Journal of Political Economy* 129.1, pp. 101–147.
- Sedlmeier, Peter and Gerd Gigerenzer (1989). "Do studies of statistical power have an effect on the power of studies?" In: *Psychological Bulletin* 105.2, pp. 309–316.
- Smaldino, Paul E (2017). "Models are stupid, and we need more of them". In: *Computational social psychology*. Ed. by R. R. Vallacher, S. J. Read, and A. Nowak. Routledge, pp. 311–331.
- (2019). "Better methods can't make up for mediocre theory". In: *Nature* 575.7783, p. 9.

- Smaldino, Paul E (2022). “Five models of science, illustrating how selection shapes methods”. In: *The dynamics of science: Computational frontiers in history and philosophy of science*. Ed. by Grant Ramsey and Andreas De Block. University of Pittsburgh Press.
- Smaldino, Paul E and Richard McElreath (2016). “The natural selection of bad science”. In: *Royal Society Open Science* 3.9, p. 160384.
- Smaldino, Paul E and Cailin O’Connor (2020). “Interdisciplinarity can aid the spread of better methods between scientific communities”. In: *MetaArXiv*.
- Smaldino, Paul E, Matthew A Turner, and Pablo A Contreras Kallens (2019). “Open science and modified funding lotteries can impede the natural selection of bad science”. In: *Royal Society Open Science* 6.7, p. 190194.
- Stanford, P Kyle (2019). “Unconceived alternatives and conservatism in science: the impact of professionalization, peer-review, and Big Science”. In: *Synthese* 196.10, pp. 3915–3932.
- Stewart, Alexander J and Joshua B Plotkin (2021). “The natural selection of good science”. In: *Nature Human Behaviour* 5.11, pp. 1510–1518.
- Szabó, György and Gabor Fath (2007). “Evolutionary games on graphs”. In: *Physics Reports* 446.4-6, pp. 97–216.
- Szucs, Denes and John PA Ioannidis (2017). “Empirical assessment of published effect sizes and power in the recent cognitive neuroscience and psychology literature”. In: *PLOS Biology* 15.3, e2000797.
- Teich, Erin G et al. (2021). “Citation inequity and gendered citation practices in contemporary physics”. In: *arXiv preprint arXiv:2112.09047*.
- Tiokhin, Leo et al. (2021). “Shifting the level of selection in science”. In: *MetaArXiv*.
- Travis, G David L and Harry M Collins (1991). “New light on old boys: Cognitive and institutional particularism in the peer review system”. In: *Science, Technology, & Human Values* 16.3, pp. 322–341.
- Wang, Jian, Reinhilde Veugelers, and Paula Stephan (2017). “Bias against novelty in science: A cautionary tale for users of bibliometric indicators”. In: *Research Policy* 46.8, pp. 1416–1436.
- Weatherall, James Owen, Cailin O’Connor, and Justin P Bruner (2020). “How to beat science and influence people: policymakers and propaganda in epistemic networks”. In: *The British Journal for the Philosophy of Science*.
- Wu, Jingyi (2022). “Epistemic Advantage on the Margin”. In: *Philosophy and Phenomenological Research*.
- Zollman, Kevin JS (2007). “The communication structure of epistemic communities”. In: *Philosophy of science* 74.5, pp. 574–587.
- (2010). “The epistemic benefit of transient diversity”. In: *Erkenntnis* 72.1, pp. 17–35.