



Priority and privilege in scientific discovery

Hannah Rubin^{a,*}, Mike D. Schneider^{b,1}

^a University of Notre Dame, 100 Malloy Hall, Notre Dame, IN 46556, United states

^b University of Illinois at Chicago, United states



ABSTRACT

The priority rule in science has been interpreted as a behavior regulator for the scientific community, which benefits society by adequately structuring the distribution of intellectual labor across pre-existing research programs. Further, it has been lauded as an intuitively fair way to reward scientists for their contributions, as a special case of society's "grand reward scheme". However, we will argue that the current formal framework utilized to model the priority rule idealizes away important aspects of credit attribution, and does so in a way that impacts the conclusions drawn regarding its function in scientific communities. In particular, we consider the social dynamics of credit attribution in order to show that the priority rule can foster structural disadvantages in socially diverse science, as well as drive the distribution of intellectual labor away from optimal.

1. Introduction

In scientific practice, discoveries of sufficient impact appear to generate certain quantities of prestige, which are bestowed upon particular scientists by the scientific community at large. The priority rule is a broad norm concerning the allocation of that prestige. The inference toward such a norm comes via an observed phenomenon in the history of science whereby, in situations of *multiple discovery*, disputes about who deserves the prestige that comes associated with a discovery are often fought by way of assertions about who was first to make it.

While the phenomenon of multiple discovery in science was recognized by Ogburn and Thomas (1922), who mention in a footnote the common appearance of disputes over priority of discovery in such cases, the investigation of these so-called "priority disputes" was taken up by Merton (1957). Merton sought to explain, by appeal to the institutions and norms of science, how otherwise dignified and reserved scientists would, when involved in instances of multiple discovery, often fight tooth and nail against the possibility that anyone but them bore responsibility for the discovery. Merton makes sense of this state of affairs in terms of the perception amongst scientists of intellectual property rights, for which it is the case that the first person responsible for the invention of a new bit of intellectual property ought to enjoy the profit that comes from it. This feature of the account—that the prestige generated by a particularly impactful discovery is taken by scientists to

properly belong to whoever is truly first to produce it—is what has come to be known as the "priority rule" in science.

Strevens (2003b), following a lead from Kitcher (1990), famously formalizes the priority rule in science as a particular manifestation of a "grand reward scheme" (p. 76) that exists in our society, which allocates prestige associated with new discoveries in such a way that incentivizes scientists to optimally distribute themselves across a variety of programs of research. In other words, Strevens argues that the priority rule, as a norm governing prestige disbursal, plays an epistemically beneficial role. In this account, the power to allocate prestige to an individual is conceptualized as belonging to the scientific community as a whole. By contrast, we argue that the community-level phenomenon of an individual receiving prestige should instead be thought of as driven by the choices of individual scientists; individuals engaged in assigning credit for a discovery in accordance with their beliefs about priority are what, altogether, give rise to the community having disbursed a quantity of prestige for that discovery. Having argued for this alternative conception of prestige disbursal (section 2), our second major goal is to evaluate some important consequences of the way scientific communities disburse prestige, given adherence to the priority rule (section 3). We will do so by means of a simple model, where what we find is that inequities in the underlying social network of the scientists can allow prestige to accumulate in the hands of those historically well-positioned within the scientific community, and meanwhile historically underrepresented or

* Corresponding author.

E-mail address: hannahmrubin@gmail.com (H. Rubin).

¹ We would like to thank more people than it would be prudent to list here, including audiences at the Formal Social Epistemology Workshop at UC Irvine and the Computational Modeling in Philosophy conference at LMU in 2018, as well as in the Social Dynamics seminar at UC Irvine and in the Department of Philosophy at University of Michigan. Special thanks go to The Norms and Networks Cluster at University of Groningen, as well as to Remco Heesen. Finally, we are indebted to our anonymous reviewers.

otherwise marginalized social groups can suffer in the context of multiple discovery. Notably, this disadvantage arises due to facts about the social structure of scientific communities, rather than due to any differences in skill or achievement, or any bias against the minority population. These sorts of phenomena are easily missed in current credit economy models of science, and, as we will argue, can affect conclusions like those Strevens offers, regarding the outcomes of different credit incentives.

In recent literature, there has been increasing attention paid to how formal models can illuminate at least some aspects of structural disadvantage to minorities. Much of this work originates in dissertation work by Justin Bruner, later published as Bruner (2019), in which the ‘cultural red king’ effect allows minorities to become disadvantaged in bargaining contexts, solely due to differences in strategic learning speed generated by differences in group size. This initial work has since engendered two ongoing lines of research: the first establishing that the cultural red king effect is robust across many different modeling assumptions (Bruner & O’Connor, 2015; O’Connor & Bruner, 2019; O’Connor, 2017; O’Connor et al., 2019) and arises among experimental subjects (Mohseni et al., 2019), and the second showing that a similar minority disadvantage can arise via different mechanisms and can have impacts on a community which are hard to repair (O’Connor, 2019; Rubin & O’Connor, 2018; Schneider et al., 2019).

While this previous work shows that minority groups can be disadvantaged in the context of bargaining and collaboration, we show that structural disadvantages can arise through different sorts of mechanisms when it comes to the credit economy.² (As we will elaborate, these mechanisms have to do with the structure of the scientific community.) Furthermore, an upshot of this claim is the suggestion that certain modeling choices regarding the credit economy in science, which have become common in philosophy of science, can mask a relationship between the presence of particular norms in scientific practice—e.g. the priority rule—and the emergence of structural disadvantage in socially diverse scientific communities. This is particularly important for evaluating conclusions, such as Strevens’s, about the value of those norms persisting in practice.

2. From priority disputes to the priority rule and back

The operational notion of profit in priority disputes is not in terms of immediate payment for services rendered. What is at stake is the accumulation of prestige, or what we might think of as ‘wide-scale credit’, in the eyes of the community at large. Merton (1957) lists several examples of how the community confers prestige to individuals, including eponymy (p. 643), honorifics such as the Nobel Prize, introduction into “honorary academies” like the Royal Society and the French Academy of Sciences (p. 644), and posthumous recognition by historians (p. 645). To this list of examples, Strevens (2003b) adds “reputation, a sizable office, the rapt attention of graduate students and the like” (p. 57). Based on these examples then, prestige appears to be a retrospective quantity, conferred on individuals when the community as a whole has come to associate them with the corresponding discoveries.

A key observation is that in most cases of disputes about who ought to receive the prestige that corresponds to a particular discovery, at least one of the parties involved perceives there to be a great injustice afoot: they are being denied access to newfound prestige because another party involved is, for reasons unrelated to the content of the discovery made by the first party, in a better position to win the prestige generated by that discovery instead. Merton proceeds to explain how such morally charged priority disputes emerge as unsurprising artifacts of the institutions and norms that govern scientific practice. In particular, he leans on the notion that scientists perceive discoveries in terms of proprietary intellectual goods, such that scientists responsible for the production of those goods

are entitled to their fruits. A claim to the quantity of prestige by any but the first to the discovery is thus perceived as tantamount to an attempt of theft. So is born the notion of a ‘priority rule’ in science: a norm according to which prestige associated with a discovery rightly belongs to whoever is truly first to produce it. The upshot is that we can expect anyone who adheres to this norm to go about crediting for a discovery exactly the party that they happen to believe has priority (and so would be deserving of the quantity of prestige associated with the discovery).

2.1. Priority, prestige, and plural action

Strevens takes Merton’s work as motivation to formalize the priority rule in such a way that he can assess its impact on scientific inquiry, as a mechanism by which scientists disburse prestige. In particular, he models the priority rule as a reward system characterized by two “parts” (p. 56):

First, rewards to scientists are allocated solely on the basis of actual achievement, rather than, for example, on the basis of effort or talent invested. Second, no discovery of a fact or a procedure but the first counts as an actual achievement.

Strevens then argues that, from the perspective of a central planner, although the priority rule may seem harmful initially (as compared to, e.g., a rule that rewards scientists for hard work and talent), it is actually beneficial for scientific inquiry. In his formal setup, it is by virtue of the priority rule that intellectual labor is efficiently distributed across various research programs with differing odds of success. This will be discussed further in section 3.3, but the basic idea is that scientists balance the odds that a research program will be successful against their expected credit if the research program is successful. Therefore, scientists do not all abandon other lines of research to join the most promising-looking research program. Instead, they distribute themselves among research programs (with more scientists working within more promising-looking programs). Others have since used this same basic modeling framework to discuss, for example, incentives to publish early and frequently (Heesen, 2018) or share intermediate results (Heesen, 2017), disincentives to replicate previous findings (Romero, 2017), and how scientists are motivated by a combination of truth and credit (Zollman, 2018).

Strevens (2003b) and those who follow his modeling framework assume that disbursing prestige is adequately regarded as a power of the scientific community as a whole. What we would like to emphasize is that, while the disbursal of prestige is indeed a community-level phenomenon, it is nonetheless driven by the choices of individual scientists. Contrast the way Strevens (2003b) talks about scientists choosing their research program with how he talks about prestige disbursal. He writes, “... with respect to worker-hours at least, allocation in science is driven to a great extent by certain decisions of individual scientists, namely, their decisions as to what projects to pursue” (p. 65). But when discussing prestige, he states: “That is a power of the scientific community as a whole, to be exercised in accordance with the prevailing reward system ...” (p. 64). Strevens does not explicitly say prestige disbursal is *not* due to individual scientists’ decision making.³ However, his modeling framework idealizes away (or ignores consequences of) prestige disbursal being, ultimately, a community-level phenomenon that emerges from individual scientists’ decisions. While, of course, every model must idealize in some way, we will discuss in section 3 how this idealization masks certain features of the priority rule, which ultimately push against the conclusions Strevens wants to draw. Before that, however, we will provide support for the claim that the community-level phenomenon of prestige disbursal is driven by the choices of individual scientists, in a way not captured by the formal framework Strevens offers.

² For a similar argument focused on citations, a quantifiable part of the credit economy, see Rubin (2021).

³ In fact, here, Strevens is arguing that the power to distribute rewards belongs to the community as a whole, rather than to the individual research programs.

First, one might note that it seems odd that in Strevens's models of the priority rule, there is no possibility of genuine conflict in cases of multiple discovery. After all, it was a litany of such conflicts in Merton's article that gives rise to an inference toward the priority rule in the first place. This gap between Merton's treatment of priority disputes and formal models of the priority rule may, in part, be explained by the shift away from historical cases of priority disputes and toward scenarios in which competing research programs approaching scientific problems are in a winner-takes-all race toward the resolution of those problems. Unlike in Merton's work, "discoveries" are now the resolutions of those problems and prestige is, in effect, *automatically* doled out to whoever wins the race.

Such scenarios certainly occur in science, for instance when it comes to highly anticipated or highly publicized results. However, not all scientific work is like this. Much research is done by individual scientists or labs (rather than an entire research program) and published in academic journals with little attention from the press. In the relevant scientific communities, scientists are generally not aware of every publication, meaning that many novelties are not acknowledged by the community (see, e.g., Hofstra et al., 2020). In these cases, it is plausibly better to start from the perspective that scientists assign credit as they individually become familiar with new discoveries, and that it is something about this individual-level dynamics that ultimately brings about the disbursement of prestige, as considered at the community level.

Second, it may be helpful to note that there is always some ambiguity in plural action claims regarding a scientific community disbursing prestige. When Strevens asks "... why does the scientific community disburse prestige in accordance with the priority rule rather than ... some alternative scheme?" (p. 58), nestled implicitly in the question is a plural action claim:

The scientific community disburses prestige.

Following the terms provided by Ludwig (2016), this claim can either be read as a distributive action sentence or as a collective action sentence. Moreover (following the form of Ludwig's argument, p. 131), we may understand the ambiguity between these two readings to wholly consist in an ambiguity of scope:

On the distributive reading, we mean that each of [the scientists within the community] were separately sole agents of [the disbursement of prestige] in a certain way. On the collective reading we mean that for [the disbursement of prestige] each of [the scientists within the community] (and no one else) was an agent of it in a certain way.

As an example, we may read the sentence "Two people built boats" as claiming either that two people independently built boats (distributive reading) or that two people came together to build boats, e.g., with one sawing and the other hammering pieces together (collective reading). Similarly, scientists could separately take part in the community's disbursement of prestige by each independently engaging in a prestige-relevant task—e.g. individually attributing credit in accordance with the priority rule (distributive reading), or they could come together as a whole to disburse the prestige as a group, again in accordance with the priority rule (collective reading).

Which reading we accept influences how we proceed to idealize prestige disbursement in a community adhering to the priority rule, so as to study its effects therein. Strevens, and others following his way of modeling the credit economy, seem to favor a collective reading. Or, at least, their formal framework can only obviously capture cases where distributive and collective plural action yield the same results: where we need not tend to any complications that may be present at the individual level. But in the next section, we provide explicit evidence for a distributive reading, which motivates our own discussion of the priority rule in section 2.3.

In what follows, to keep our discussion of these two readings as clear as possible, we will use the term 'credit' when referring to individual activities of associating a person with a discovery. We will use the term 'prestige' when talking about the consequences of a community-level association of a person with a discovery (whether that association is

the result of individual scientists' assigning credit or else the result of some community-level activity).

2.2. Evidence for the distributive reading

The scientific community, though in some respects quite hierarchical, is not centrally governed. The awarding of prestige, though associated with individual accomplishments, often happens at no particular moment in time. That individuals come to be awarded prestige at no particular time by no particular decision made on behalf of the community strongly suggests that the disbursement of that prestige is not a single act of which each scientist was an agent. Instead, it seems more plausible to view the disbursement of that prestige, in many cases, as following from the members of the community separately taking actions (of an identical kind to each other) that need not individually resemble the awarding of prestige. That is, individual scientists assign credit upon learning about a discovery in relation to a name, and along the way prestige is conferred by the scientific community. (No one gains the respect of a community because one person gives them credit for a discovery, but if the members of a community all generally associate a particular scientist with a discovery, that scientist will receive the associated reputational benefits, etc.)

To support this alternative view, we will offer historical examples drawn from evolutionary biology, (computational) social science, and fundamental physics. Each of the examples below is a case where the scientific community as a whole "got it wrong", in the sense that the first person to make the discovery was not the one who received the subsequent prestige. However, to be clear, that the scientific community in each of the cases "got it wrong" in this sense is not itself what constitutes evidence for the distributive reading we have advocated.⁴ What we want to emphasize is one feature common to all three cases: *individuals attributing credit as they go about their ordinary affairs in accordance with their personal commitments to the priority rule are what, in aggregate, determines to whom prestige comes to be awarded.*

In evolutionary biology, the distinction between proximate and ultimate causes in biology is attributed to Ernst Mayr. As Laland et al. (2011) point out, this distinction was made well before Mayr wrote about it in the 1960s (they cite an article by J. Baker written in the 1930s). Nonetheless, Mayr's article is what ostensibly led to the distinction's widespread acceptance in evolutionary theory. Moreover, even as Laland et al. (2011) flag the trivia that Mayr was not the first to make the distinction, they consistently refer to it as "Mayr's distinction". This is illuminating because, as an evidently conscious decision of the authors in that article, it demands explanation. One prudent explanation is that of historical usage: since people started citing Mayr when talking about the distinction, it subsequently became known as his distinction. Hence, even when the pre-history of Mayr's work is acknowledged, Mayr's legacy continues to enjoy the prestige.

In the social sciences, Thomas C. Schelling undoubtedly enjoys the prestige that surrounds the basic checkerboard model of segregation, which is equivalently called the "Schelling model" (an instance of eponymy), and is taught in any introductory course on the subject of agent-based computational models in the social sciences. As is now recognized, James M. Sakoda beat Schelling to the public invention of computational models of segregation: the latter's model amounts to a special case of one of the former's, which had its origins in his dissertation work twenty years earlier and was printed in the previous issue of the same journal in which the latter's model appeared. But as (Hegselmann et al 2017, p. 5–6) argues, "No crime happened, no conspiracy was involved ... as to the main actors, nobody did anything wrong."

⁴ The community can also be wrong about who to award prestige under a collective reading. However, the collective reading cannot capture what went wrong in the cases below, where each individual agent assigning credit on account of personal commitments to the priority rule meanwhile results in a conflict between the recognition of the priority of one person and the prestige being bestowed on another.

Instead, one deciding factor it seems was Schelling's subsequent decision to write a book, developing many of the ideas in his paper, accessible for much broader audiences than just those computer scientists who happen to additionally be interested in modeling social dynamics. Those broader audiences were encouraged to try out small, table-top examples of the checkerboard models under scrutiny, whereupon: "They all had experienced how surprisingly fast, right before their eyes, certain unexpected, dramatic macro structures evolved, generated by fairly innocent looking micro-motives— an eye-opening phenomenon *par excellence*" (Hegselmann et al., 2017, p. 87). This got those broader audiences talking about the demonstrative power of these simple computational models, and they were talking about it in the context of Schelling's work. Whatever reasons we give for why Schelling enjoys the prestige, an indispensable part of the story is that he enjoys it because individuals separately began to associate him with the discovery, rather than because the scientific community wholesale decided that it was he who was responsible for the discovery.

A final case worth mentioning is one in fundamental physics, particularly in the history of quantum mechanics. In 1932, von Neumann published an alleged "no-go" proof of the viability of hidden-variables underpinning quantum mechanical behavior in an ultimately deterministic theory. Grete Hermann evidently discovered a flaw in the scope of the proof in 1935, yet this discovery was "not widely known at the time, and her criticism had no impact whatsoever" (Seevinck, 2016, p. 107). In 1964, John Bell happened on the same such discovery, in the aftermath of the development of Bohmian mechanics (whose success as a deterministic, hidden-variables alternative to quantum mechanics clearly stood as proof of the alleged impossible). Bell enjoyed the benefit of offering Bohmian mechanics as a demonstration of his point (whereas in 1935, Hermann could only conclude that the possibility of such a hidden-variables theory was left open). There are many reasons why Hermann's contributions at the time were overlooked. Our point here is that following this neglect, when most scientists were for the first time ready to credit *someone* for identifying the ostensible flaw in von Neumann's proof (that is, following the development of Bohmian mechanics), Bell's independent study of the subject led him to be the recipient of that credit, and, eventually following, prestige.

In each of these historical anecdotes, there are myriad reasons one could give for why a particular person enjoys the prestige. The common feature across each of these histories is that the recipient of the prestige associated with a discovery is the individual who first became largely known to be associated with the discovery, in virtue of a plurality of individuals each attributing credit to them. We take this as evidence from the history of science for the importance of recognizing a distributive reading of the claim "the scientific community disburses prestige": in each of these cases, each of the scientists within the relevant community happened to attach credit for the relevant discovery to some or other scientist. Meanwhile, by virtue of each such scientist acting in this way separately, i.e. as a sole agent, prestige came eventually to be disbursed by the community to the individual that most of the scientists credited.

2.3. Credit attribution contests

Taking a lead from the historical cases just discussed, we offer the following perspective. Scientific developments are not, in general, immediately known by everyone in a scientific community. News of them spreads throughout the community, rather, through an informal social network that spans the community. Insofar as a development proves valuable to the community, the party known to be responsible for that development enjoys a corresponding quantity of prestige. Of course, it need not be the case that *every* individual in the community has assigned credit; the community may just as well bestow prestige when the vast majority has learned about the development (whereupon any stragglers learn about the development as do those outside of the community, as discussed below).

In cases of multiple discovery, just as in any other case, when individuals within the community hear news of pertinent developments, they generally come to associate a development with whomever they first learn is responsible for it. The difference is that in the case of multiple discovery, different individuals may associate the development with different parties. Only when a large majority of the community comes to associate the development with a single party is prestige bestowed on that party (as it is normally), while any other party that independently produced the same development is neglected.

At risk of being denied the prestige that they believe they are due on the basis of their work, any party that does not enjoy the support of the majority will protest that they ought to receive the prestige that in other circumstances would have been awarded to them, for instance had the social network just happened to have been structured differently. As Merton (1957) notes, those in the minority who believe the wrong person received the prestige will often protest, too. On the hypothesis that the conferring of prestige by the community occurs when a sufficient number of individuals within the community associate a particular person with a discovery, the function of such a protest is obvious: to change individuals' associations as to who typically gets credit for a discovery.

Returning to the priority rule, one can grant that scientists ought to disburse prestige on the basis of actual achievement, and, in winner-takes-all cases, none but the first counts as an actual achievement. That is, we agree that there is a norm governing disbursement of prestige, as described by Merton and Strevens. Still, we emphasize that it is individuals who are the subjects that take steps to act in accordance with such a norm. And this leaves open a further question, missed in current models examining the priority rule: how does prestige *get allocated*, when scientists are each acting separately as sole agents assigning credit, consonant with the priority rule?

Based on what we have said above about the nature of priority disputes, we conclude that prestige associated with a particular scientific development is bestowed wholesale upon a particular party as a consequence of a large majority of the surrounding community having all individually come to associate, as a matter of priority, that development with them. And in instances of multiple discovery, we may imagine 'credit attribution contests' occurring among individuals within the community who independently produce similar developments at the same time.⁵

One might object that a large majority is not enough – surely all or nearly all of the community must agree on a discoverer for that person to get the prestige. But there are reasons we think a large majority would suffice. First, in the instance of a priority dispute between the two competing parties, one might imagine that those who enjoy the support of the large majority are more likely to win the priority dispute: along the lines of the observation by Merton just mentioned, those who enjoy the support of the large majority also enjoy a larger collection of possible defenders, ready to fight against an instance of perceived injustice.

Second, even without a priority dispute, a large majority of scientists coming to associate one person with a discovery might be sufficient for prestige to be bestowed on that individual. For instance, when using reference to a name as short-hand for an idea that one's interlocutor will understand (e.g. using "Schelling's segregation model" to refer to a mathematical model whereby minimal conditions for segregation are demonstrated), the most effective name to choose is the name most well-known in connection to that idea (one's interlocutor is less likely to know

⁵ Of course, what is supposed to count as "similar developments" and "at the same time" is highly contextual. As the historical cases suggest, "at the same time" is sometimes better understood as "both occurring before that time at which the content of the discovery becomes a hot topic". This subject merits further scrutiny, particularly in order to get empirical traction on the various models discussed here, with regards to particular instances of priority disputes in real-world scientific communities. But that is a further project, whose ultimate aims differ from those that are relevant here.

of Sakoda's work). One might imagine that there is some tipping point, or threshold, at which it becomes prudent for individuals wishing to communicate in this way to defer to the use of the more well-known name, irrespective of who they individually have associated with it.⁶ For example, Laland et al. (2011) refer to "Mayr's distinction" while noting that Mayr does not deserve priority for that distinction. Similarly, new members of the community, hoping to signal their understanding of the field, will think to provide the name most people within their community associate with the discovery, lest they be thought ignorant. Finally, it is reasonable to expect that the more well-known name would be used in review articles written by members of the community, to communicate the results of their sub-field to a wider audience (and to the few members of the community who may not have already heard of the discovery), and so both the broader scientific community and, eventually, the public at large and history books would come to recognize the more well-known name.⁷

In the next section, we will argue that this picture of how priority interacts with prestige allocation in science does not necessarily align with our ideas of fairly rewarding people for benefits they confer (as in society's grand reward scheme, which Strevens suggests does exactly that). That is, while "it seems plainly fair to reward scientists precisely in proportion to the actual contribution they make to society" (Strevens, 2003b, p. 75), we will see that contributions from some members of the community are more likely to be overlooked. Our sense that the priority rule is fair will (at least partially) disappear, once we turn to focus on consequences of its implementation in the kinds of situations that we have emphasized are common in science.

A discussion about the common implementation of a norm like the priority rule in science is valuable in its own right. In particular, it highlights one plausible way by which the institution of the credit economy can foster structural disadvantages in socially diverse science. But as we go on to elaborate, it also suggests that there are certain cases where the reward system in science does not lead to a (rough) maximization of a social good via optimal distribution of labor. This conclusion stands in contrast with that presented by (Strevens 2011, p. 194) as a summary of his earlier project.

3. Historically underrepresented groups

Here, we will formalize the credit attribution contests just discussed so as to study the influence of network structure on the awarding of prestige in instances of genuine multiple discovery. This model is meant to be descriptive of scientific practices regarding prestige (in accordance with the priority rule), not prescriptive of how a central planner ought to distribute prestige.⁸ As such, we do not compare the efficacy of different mechanisms for disbursing prestige; rather, we take prestige to simply be that which is disbursed in the way we have so far described.

3.1. A basic model

We formalize scientific communities as networks of agents, where nodes of the network represent individual scientists and edges, or links, represent regular information channels between them.⁹ These links are

bidirectional and can be thought of as representing people who talk to each other when working on a new project, or who ask each other if looking to reference a paper on some topic or other.

We model credit attribution contests where information about discoveries spreads throughout these networks. In the model, we start with an instance of multiple discovery. Two scientists each independently make some discovery, and not knowing about the other, they each believe themselves to be the discoverer. In the second time-step, we pick a third node at random. Whoever they are closest to in the network¹⁰ is who they will credit as having priority for the discovery, who they will cite with respect to that idea, etc.¹¹ Then in the next timestep, another random node is chosen. This fourth scientist gets news of that discovery from whoever is closest to them in the network that already has some belief about who made the discovery. (The idea here is that when one goes looking for a paper on some topic x , for instance so as to cite a discussion of that topic, one might ask their friends if they know about a paper on x , or alternatively, one might hear about some new discovery from a trusted collaborator, or so on.¹²) This fourth scientist then has a belief about who deserves credit for the discovery. This process continues until all scientists have attributed credit. If there is a super-majority of 2/3 in favor of one discoverer over another, we say the former wins the credit attribution contest (and so, gets the prestige associated with that discovery). If not, there is a tie.

Since this information spread occurs over networks of scientists, the structure of the network will likely impact who receives credit for a discovery. At first, we will consider basic Barabási-Albert networks (Barabási & Albert, 1999), though in the next section, we will consider alterations to that model, which include social identity types and network change over time. As explained below, these networks capture several important features of scientific communities.

Barabási-Albert networks are formed in the following way. First, we start with a small number of fully connected nodes (nodes with all possible links between them), m_0 . Then, new nodes are added one by one until the network reaches a designated size. Each time a new node enters, it forms a set number of links, m . For the results presented here, $m_0 = m = 4$. New links are formed via preferential attachment. That is, the more links a node already has (i.e., the higher its *degree*), the more likely it is that an entering node will form a link with it. The probability p_i that the new node is connected to node i is:

$$p_i = \frac{d_i}{\sum_j d_j} \quad (1)$$

where d_i is the degree of node i , and $\sum_j d_j$ is the sum of the degrees of all nodes in the network (not including the new entering node). The higher the degree a node already has, the more likely it is the new node will connect to it.

These networks have a couple of important features. First, in these networks, the 'rich get richer': nodes that already have many links are more likely to get new links. This captures a scenario seen in many real-world networks where the oldest members of the community tend to be the most central and well-connected individuals in the network. Moreover, this model takes on a natural interpretation in the context of

⁶ Thanks to The Norms and Networks Cluster at University of Groningen for discussions on this point. In the next section we provide a particular tipping point of 2/3. While this number is arbitrary, the existence of a tipping point is not arbitrary, and the particular number chosen will not significantly affect results. For instance, similar results have been obtained with 3/5.

⁷ Thanks to Jim Joyce for encouraging us to think about the role of review articles.

⁸ We use the modifier 'descriptive' only in contrast to 'prescriptive'; we do not mean to imply that our model is descriptively accurate in that it captures all features of the situation at hand.

⁹ Code for simulations in section 3 can be found at: https://osf.io/b29y7/?view_only=299ebfa8d597497e9314afb69da0907b.

¹⁰ If there is a tie, a node is chosen at random.

¹¹ This can either be interpreted as this third node caring about a particular discovery, and searching for it in that moment, or one of the original discoverers discussing their discovery with someone close to them in the network.

¹² One might think that the nodes along the relevant shortest path should also make attributions as the information must pass between them. While that is a reasonable alternative assumption to explore, we might also imagine cases where those along the path do not make attributions: someone recommending a paper they know about but have not read, someone recommending an author working on a particular problem without knowing about a specific paper, or someone looking into an author's paper because their close collaborators have mentioned a name before.

scientific research communities: as new researchers enter the community (e.g. as graduate students), they often seek out social relationships with the more well established members of the community, the oldest of which are often the most esteemed.

Second, these networks are scale-free, meaning their degree distribution follows a power law. In other words, there are many nodes with a few links and a few nodes with a large number of links; there are a few ‘hubs’ in the network. Many real world networks are (approximately) scale-free, including many types of social networks and collaboration patterns. Among the scale-free networks, Barabási-Albert networks are particularly useful for our purposes. There is evidence that collaboration networks are formed via preferential attachment, similar to the method of preferential attachment used in the formation of Barabási-Albert networks (Barabási et al., 2002; Newman, 2001). Additionally, since in later sections we will be discussing how networks evolve over time, Barabási-Albert networks are useful because they already stipulate what should happen when new nodes enter the network.

For this basic model, we formed a network of 100 people, then ran 1000 contests on each network to estimate the likelihood of each person getting the prestige for their discovery, then performed 100 replications (i.e., we formed 100 different networks of 100 people, and ran 1000 contests on each). Fig. 1 shows the estimated likelihood of winning for each node in the network. That is, of the contests a node was a part of, it shows the percent of contests that node won. In this figure, the lower number a node is, the older it is (nodes 1–4 are all the same age, as the network started with 4 fully connected nodes). As one might expect, older nodes were more likely to win credit attribution contests because they tended to have a higher degree. Intuitively, when there is an instance of multiple discovery, those scientists who are more well connected are more likely to wind up with the prestige associated with their discovery, because the news of their discovery travels faster.

3.2. Types and evolution

This observation points to a possible disadvantage for historically underrepresented groups (HUGs). Since older nodes tend to be members of the historically entrenched group (HEG), the HEG members will tend to be better connected, even when members of a HUG begin to enter the community at an equal rate later on. This means that HEG members tend to receive prestige for their discoveries at a higher rate, even when they make the discovery at the same time as a member of the HUG. Of course, scientific communities also change over time. Older members of the community retire and new scientists enter the community. After a time, if the HUG enters the community at a rate equal to its size in the population, it will eventually achieve proportional representation in the

scientific community. This raises the question: will the HEG advantage over the HUG ever go away? (And, if so, how long will it take?)

In order to address these questions, we introduce types into the basic model: HEG members will be type 1 and HUG members type 2. There is nothing intrinsically important about these types; they are not related to scientific competence or likelihood of producing a scientific discovery. They are, however, socially relevant, in that type 1 enters at a higher rate earlier in time. In particular, we used a logistic growth equation to determine how likely it was a new node was type 2 at each point in time. This represents a case where the HUG finds it hard to enter the scientific community at first, but once there is a sufficient number of them it becomes much easier.¹³ We will consider results at first where, by the end, the HUG enters at roughly the same rate as the HEG.¹⁴ If we run a similar model as in section 3.1, but assign types to the nodes as just described, type 1 individuals win credit attribution contests against type 2 individuals about 42.3% of the time, lose about 16.7% of the time, and tie the rest of the time.

These results are already telling, but we also want to consider how social identity impacts network formation and evolution over time. A pervasive feature of real communities comprised of different social identity groups is *homophily*, or the preference for linking to members of one’s own social identity group. People are homophilic in a variety of contexts (e.g. when forming friendships (Currarini et al., 2009)), and scientific communities are homophilic as well, especially when it comes to co-authorship (Ferber & Teiman, 1980; McDowell & Smith, 1992; Boschini & Sjögren, 2007; del Carmen & Bing, 2000; West et al., 2013; Wang et al., 2019) and citation patterns (Ghiasi et al., 2018; Paris et al., 1998; Wardle, 1995). We include homophily in the model by having agents place some weight, H , on their similarity to a node in addition to their degree. We used the following to determine how much the incoming node k values linking with each of the existing nodes i :

$$v_{ik} = H \times \frac{s_{ik}}{\sum_j s_{jk}} + (1 - H) \times \frac{d_i}{\sum_j d_j} \quad (2)$$

where $s_{ik} = 1$ if nodes i and k are of the same type and 0 otherwise.¹⁵ The probability that the new node links with a particular node, p_{ik} , is then proportional to v_{ik} . The likelihood a node is chosen is thus determined by its value to the new individual, including both homophilic preferences and degree, rather than just its degree.¹⁶

We incorporate network change over time into the model in the following way. There is a maximum network size of 100 scientists, so as a new scientist enters beyond the first 100, the oldest node is removed from the network (it “retires”) along with all its links. When a new node enters, it then forms a set number of links, as in the Barabási-Albert model. Additionally, in order to capture the consequences on the social network structure of a scientific community that trains its younger members,




Fig. 1. Likelihood of winning for each node in the network.

¹³ Results are shown for $P(\text{type } 2) = \frac{w}{1 + 10 \times e^{-0.05t}}$, where w is the fraction of the larger population the HUG comprises and t is the time-step the node enters. However, none of the results we discuss hang on using this particular equation. We also obtained similar results with an equation where the probability of type 2 increases quickly then asymptotes at .5, $P(\text{type } 2) = 1 - \frac{1 + 2e^{-0.05t}}{2 + e^{-0.05t}}$.

¹⁴ That is, we set $w = 0.5$ in the equation in footnote 13.

¹⁵ We incorporate homophily in this way because its influence on the value an incoming node assigns to a new link remains constant as we move node to node, regardless of the degrees of those nodes. Some authors incorporate homophily as a weighting of the degree of a node, and find, similar to our results here, that homophily “makes the rich even richer” (Kim & Altmann, 2017) and that homophily can lead to minority group members occupying less important places in the network (Karimi et al., 2018). Incorporating homophily in this alternative way does not qualitatively affect the results described below.

¹⁶ One might also think that winning a credit attribution contest would make an existing node more “prestigious” and thus would increase its value in the eyes of incoming nodes. This is not included in the model, but would likely intensify the effects reported here.

when the network grows beyond 50 nodes, incoming nodes also choose an ‘advisor’.¹⁷ For each node the advisor is linked to, the new node has a chance of linking with that node as well (a 50% chance for the results below), in addition to its links formed according to Eq. (2). This captures a scientific field that grows to a certain size, becomes established, then begins to adopt practices to train new generations of scientists.¹⁸

To interpret our results, we define *HEG advantage* as the probability a HEG member wins minus the probability a HUG member wins the credit attribution contest in an instance of multiple discovery, where the two discoverers belong to different social identity groups. In order to discuss how long the HEG advantage will persist when the network evolves, we will talk about HEG advantage over *generations* of the scientific community. A generation is defined as the time it takes to have a complete turnover of scientists. Since in each time-step the oldest scientist retires and a new scientist enters, with 100 scientists a generation is 100 time-steps.

Let us first consider a case where the HUG approaches 50% of the population, where by timestep 100 (after 1 generation) they are entering in equal proportions and after 2 generations they have achieved equal representation. For each level of homophily, we formed 250 networks. For each of these networks, we performed 250 credit attribution contests (where a type 1 and type 2 individual were competing for credit) every 25 rounds to get an idea of how likely it was that each social identity group would get prestige for their discovery, and how these chances changed over time. Fig. 2a shows what happens in this case.

We find that the HEG advantage disappears quickly over time. Interestingly, the HUG has an advantage for a short period of time (that is, the HEG advantage goes negative). As homophily increases, this temporary HUG advantage increases. This is likely because, when there are very few members of this group for the initial time period, the one or two that exist serve as focal points for the incoming members. As the HUG starts entering at higher rates, these focal points become highly connected to all the new people such that if a member of their social identity group makes a discovery, they will know about it. Since these focal individuals are so highly connected, they are likely to have at least some connections with new HEG members despite homophily, because people still do care to some extent about forming links with highly connected people. So news of the HUG members’ discoveries can spread.

That the HEG advantage disappears quickly in this case looks somewhat promising—if we can get equal representation, eventually no group is disadvantaged.¹⁹ A situation like this is achievable if we are thinking about men and women, but not if we are thinking about minority groups—like racial minorities, people with disabilities, etc. In these cases, we can talk about what happens if proportional, rather than equal, representation is achieved.

As seen in Fig. 2b, the situation is different when the HUG is a minority group. In this case, even when the HUG achieves proportional representation within two generations, the HEG advantage can be more severe and last longer, especially in the presence of strong homophily. When homophily is low, the HEG advantage disappears around when the HUG achieves proportional representation (with no real period of advantage for the HUG). But as homophily increases, a qualitative change appears: the HEG advantage can persist over time, meaning the HUG would be less likely to get credit for their discoveries indefinitely.




Fig. 2. HEG advantage over time, for different levels of homophily, where (a) the HUG is 50% of the population and (b) the HUG is 10% of the population.

3.3. Research programs

There is also the further question of whether the priority rule can serve the function of promoting a (roughly) optimal division of cognitive labor, in the way Strevens suggests.²⁰ In order for the priority rule to provide this benefit, it must be the case that scientists believe they will receive prestige for their discovery (not merely that they ought to), so that they decide which research program to join based on the likelihood of making a discovery. Merton (1957) already provides some evidence that scientists believe recognition is not automatic after discovery, and

¹⁷ The advisor is the first link formed according to Eq. (2).

¹⁸ Nothing depends on this particular way of doing things. For instance, we obtained similar results when new nodes used a copying mechanism to create new links, modified from Kumar et al. (2000), though homophily had slightly less effect on HEG advantage.

¹⁹ We have argued elsewhere that in other situations, such as when there is a discriminatory bargaining norm in a community of scientists, merely increasing representation of the minority group will not eliminate inequalities (Schneider et al., 2019).

²⁰ Thanks to Remco Heesen for this point. See also (Muldoon & Weisberg, 2011), who explore another way in which switching to an agent-based computational modeling framework can counsel against Strevens's conclusions. Although their models differ considerably from ours, a common theme is that heterogeneous agents who are myopic in their decision-making can be expected to drive community-level patterns, which are ultimately at odds with Strevens's model of science.

(to the contrary) is something that is more likely for some than others. He quotes Norbert Weiner explaining:

I was competitive beyond the run of younger mathematicians, and I knew equally that this was not a very pretty attitude. However, it was not an attitude which I was free to assume or to reject. I was quite aware that I was an out among ins and I would get no shred of recognition that I did not force (p. 649).

If news of a discovery must travel through the network before a credit attribution contest is won and prestige bestowed, scientists likely also choose research programs based on the likelihood they will receive credit for their discoveries within that research program.

There are many ways such considerations could affect scientists' choice of research program. Our purpose in this section is to provide an example of how considerations of network structure and social identity matter to an analysis of the effects of the priority rule on program choice. To this end, we extend the model in section 3.2 to include scientists' decisions regarding which research program to join. We ask: how might the network structure and considerations of social identity pull a network away from the optimal division of labor?

To answer this question, we assume there is an optimal division when the network first reaches its full size, in order to study how quickly the community departs from optimal when scientists choose their research programs based on a combination of likelihood of making a discovery and likelihood of receiving credit for it.²¹ In this model, a scientist chooses a research program, chooses an advisor within that program, then forms links with others in the community. While the advisor must be within the research program chosen, the other links formed may be with any scientist in the community.²²

In the choice of research program, we match Strevens's setup as closely as possible, though there will be some differences, necessitated by our different perspectives. With Strevens, we assume there are two research programs, each with some probability of success, s_1 and s_2 . Both probabilities are functions of the number of people in the program, where it is assumed there are diminishing returns to each new person joining the program. Program 1 is assumed to be the better program, with a higher probability of generating a discovery, i.e. $s_1 > s_2$. The optimal distribution maximizes the overall probability of making a discovery—producing the social good, $s_1 + s_2 - s_1s_2$.²³ In Strevens's setup, a scientist chooses a research program based on the likelihood the program makes a discovery and, in the event that both programs make the discovery, the likelihood their research program does so first, w_i .²⁴ For research program 1, for instance, this would be $s_1 + s_1s_2w_1$. Our slight alteration to this setup is that we interpret w_i as the likelihood a scientist gets credit in the event of multiple discovery, which scientists estimate based on the most recent round of credit attribution contests – how likely was it that a scientist of their type got credit in research program 1 (likewise in

research program 2)? This allows us to see whether, when scientists consider the likelihood of receiving credit in addition to the likelihood of making a discovery in the context of choosing research programs, the priority rule will incentivize the optimal division.

Fig. 3 shows how far the scientific community is pulled away from the optimal division over successive rounds when the HUG is 10% of the population.²⁵ We can see that, even though we start with an optimal division at round 100, the scientific community diverges from this division quite rapidly. While homophily makes a difference to these results, there is substantial deviation for all levels of H .²⁶

The most plausible explanation is that scientists are choosing research programs based on the likelihood someone of their type will get credit for a discovery made within that research program. In fact, we find that differences in w_1 and w_2 for each type predict whether that type will make up a greater proportion of research program 1 or 2. Let p_i refer to the proportion of program i 's members which are HEGs. For the HEGs, the covariance between what we might call their 'credit differential' ($w_1 - w_2$) and 'research program differential' ($p_1 - p_2$) 25 rounds later (after new scientists have joined programs using the estimates of w_1 and w_2) is 0.79. The covariance between similarly defined 'differentials' for the HUGs is the same. Scientists tend to go where they think it is more likely they will get credit for their discoveries, in which case the optimal division of cognitive labor is not incentivized.

Another way to think about this point is that while the priority rule is "well-attuned" to some parameters, e.g. "the degree of correlation between different programs' successes and the speed with which different programs can be expected to succeed" (Strevens, 2003a, pp. 2–3), it is not well-attuned to others. That is, for certain properties of scientific communities, the priority rule will lead scientists to adjust their distribution into research programs both in the right direction and roughly in




Fig. 3. Distance from the optimal distribution over time, for different levels of homophily, where the HUG is 10% of the population.

²¹ To keep the analysis simple, we also now start advising in the model when the network reaches its full size.

²² The scientist's choices are made by the same procedure outlined in section 3.2, i.e. their advisor is chosen according to Eq. (2), their other links are formed by copying their advisor's links and forming their own links, and so on. (Though, as seems realistic, we assume a scientist is twice as likely to form a link with someone in their own research program.)

²³ Our results use the following equations: $s_1 = \frac{n}{n+10}$ and $s_2 = \frac{N-n}{N-n+50}$, where N is the size of the community and n is the number of people in research program 1. With these equations, and a community of size 100, the optimal division is to have 70 people in program 1 and 30 people in program 2.

²⁴ Regarding the role of the scientist's conditional expected prestige, Strevens (2003b) writes: "The further assumption I make ... is that there is a fixed amount of prestige bestowed on each program that is rewarded, to be divided among the program's workers according to their contributions" (p. 71). However, this assumes that the research program gives appropriate credit to all who contribute, which we might think (following the arguments in section 2) ignores how various individual-level decisions determine prestige disbursement.

²⁵ Distance from optimal is just $|\text{actual \# in program 1} - \text{optimal \# in program 1}|$.

²⁶ The HEG advantage remains very similar to what is shown in Fig. 2b. Though, interestingly, if we assume that people are ten times as likely to link with others in their own research program (as opposed to twice as likely, as for the results in Fig. 3), HEG advantage increases notably, e.g. when $H = 0.8$ the HEG advantage is 0.35 by round 500. This suggests that HEG advantage is not a particularly good predictor of how far the community will be from an optimal distribution. This may be of interest, as HEG advantage could plausibly be measured, while distance from optimal is a theoretical quantity, which is more difficult to directly assess.

proportion to what the optimal distribution would be (Strevens, 2003a, p. 2), but for other such properties, it will lead to adjustments in the wrong direction. Specifically, we have shown that the priority rule is not well-attuned to certain parameters relevant to social identity, which reflect salient properties of real-world scientific communities, e.g. the presence of homophily in socially diverse science.

4. Concluding remarks

Strevens (2003b) interprets the priority rule in science as a behavior regulator for the scientific community, part of a grand reward scheme which benefits society by adequately structuring the distribution of intellectual labor across pre-existing research programs. However, prestige does not necessarily go to who deserves it. People are not always rewarded based on the benefit they confer, and so it is difficult to regard the priority rule as part of a grand reward scheme that resonates with our notions of fairness. It may be, as in our models, that the people who are rewarded have done something that could merit a reward. Nonetheless, contributions made by members of certain social identity groups are more likely to be overlooked.

Further, if scientists instead believe that prestige is disbursed in the way we discuss, and consequently that news of a scientist's discovery must spread through the scientific network, Strevens's optimal division of labor is likely not what is incentivized by the priority rule. Section 3.3 gives one example of how social identity and network structure may affect scientists' estimations of how likely they are to receive credit for a discovery, leading to deviations from optimal division of labor, but there are many other factors that scientists may consider. For instance, in the context of our basic model (section 3.1), scientists might be incentivized to pick a research program with fewer well-connected people, so as to decrease the likelihood that the prestige associated with their potential discovery will be "scooped" by another who is in a better position to capture it.

If less-established scientists' contributions are more likely to be recognized in a less promising research program, the scientific community may end up with an inefficient division of cognitive labor, with more scientists working in a less promising program than what would be optimal. This is because not only do less-established scientists consider how likely it is that a research program will lead to a discovery, they also consider how likely it is that their achievements will be recognized and rewarded with the commensurate prestige. In other words, the disincentive of being scooped may often be in tension with a scientist's considerations of the "intrinsic potentials" of the competing programs.²⁷

Absent further study, it is difficult to discern which of these various factors (or others) clearly dominates in the decision-making of scientists,²⁸ and even more difficult to suss out the epistemic consequences of such decision-making; we leave this investigation for future work. Our point is that these considerations are absent in current 'credit economy' models of science, and that the extent to which any such strategies are motivated at all can only be evaluated in a model of the priority rule which takes into account individual scientists' credit attributions. Altogether, the epistemic benefits of the priority rule are far from guaranteed when we take seriously the implications of a distributive reading of the statement: The scientific community disburses prestige.

²⁷ Worse still, it is worth flagging that this state of affairs may also eventually lead to clustering into sub-disciplines according to social identity, as HUG members are often underrepresented amongst those who are well-established. Following the arguments given by Schneider et al. (2019), there is some historical precedent to suggest that such clustering is ultimately detrimental to the general state of our scientific knowledge across disciplines.

²⁸ One might also want to consider subtleties arising from different reference classes valuing contributions differently (Lee, 2020).

Acknowledgements

This material is based on work supported by the National Science Foundation under grant no. 2045007.

References

- Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439), 509–512.
- Barabási, A.-L., Jeong, H., Neda, Z., Ravasz, E., Schubert, A., & Vicsek, T. (2002). Evolution of the social network of scientific collaborations. *Physica A: Statistical Mechanics and Its Applications*, 311(3–4), 590–614.
- Boschini, A., & Sjögren, A. (2007). Is team formation gender neutral? Evidence from coauthorship patterns. *Journal of Labor Economics*, 25(2), 325–365.
- Bruner, J. P. (2019). Minority (dis) advantage in population games. *Synthese*, 196(1), 413–427.
- del Carmen, A., & Bing, R. L. (2000). Academic productivity of African Americans in criminology and criminal justice. *Journal of Criminal Justice Education*, 11(2), 237–249.
- Bruner, J. P., & O'Connor, C. (2015). Power, bargaining, and collaboration. In T. Boyer, C. Mayo-Wilson, & M. Weisberg (Eds.), *Scientific collaboration and collective knowledge*. Oxford University Press.
- Currarini, S., Jackson, M. O., & Pin, P. (2009). An economic model of friendship: Homophily, minorities, and segregation. *Econometrica*, 77(4), 1003–1045.
- Ferber, M. A., & Teiman, M. (1980). Are women economists at a disadvantage in publishing journal articles? *Eastern Economic Journal*, 6(3/4), 189–193.
- Ghiasi, G., Mongeon, P., Sugimoto, C., & Larivière, V. (September 2018). Gender homophily in citations. In *23rd international conference on science and technology indicators (STI 2018)* (pp. 1519–1525), 2018.
- Heesen, R. (2017). Communism and the incentive to share in science. *Philosophy of Science*, 84(4), 698–716.
- Heesen, R. (2018). Why the reward structure of science makes reproducibility problems inevitable. *The Journal of Philosophy*, 115(12), 661–674.
- Hegselmann, R., Schelling, T. C., & Sakoda, J. M. (2017). The intellectual, technical, and social history of a model. *The Journal of Artificial Societies and Social Simulation*, 20(3).
- Hofstra, B., Kulkarni, V. V., Galvez, S. M.-N., He, B., Jurafsky, D., & McFarland, D. A. (2020). The diversity–innovation paradox in science. *Proceedings of the National Academy of Sciences*, 117(17), 9284–9291.
- Karimi, F., Géois, M., Wagner, C., Singer, P., & Strohmaier, M. (2018). Homophily influences ranking of minorities in social networks. *Scientific Reports*, 8.
- Kim, K., & Altmann, J. (2017). Effect of homophily on network formation. *Communications in Nonlinear Science and Numerical Simulation*, 44, 482–494.
- Kitcher, P. (1990). The division of cognitive labor. *Journal of Philosophy*, 87(1), 5–22.
- Kumar, R., Raghavan, P., Rajagopalan, S., Sivakumar, D., Tomkins, A., & Upfal, E. (2000). Stochastic models for the web graph. In *Foundations of computer science, 2000. Proceedings. 41st annual symposium on* (pp. 57–65). IEEE.
- Laland, K. N., Sterelny, K., Odling-Smee, J., Hoppitt, W., & Uller, T. (2011). Cause and effect in biology revisited: Is Mayr's proximate-ultimate dichotomy still useful? *Science*, 334(6062), 1512–1516.
- Lee, C. J. (2020). The reference class problem for credit valuation in science. *Philosophy of Science*, 87(5), 1026–1036. <https://doi.org/10.1086/710615>
- Ludwig, K. (2016). *From individual to plural agency: Collective action I*. Oxford University Press.
- McDowell, J. M., & Smith, J. K. (1992). The effect of gender-sorting on propensity to coauthor: Implications for academic promotion. *Economic Inquiry*, 30(1), 68–82.
- Merton, R. K. (1957). Priorities in scientific discovery: A chapter in the sociology of science. *American Sociological Review*, 22(6), 635–659.
- Mohseni, A., O'Connor, C., & Rubin, H. (2019). On the emergence of minority disadvantage: Testing the cultural red king hypothesis. *Synthese*, 1–23.
- Muldoon, R., & Weisberg, M. (2011). Robustness and idealization in models of cognitive labor. *Synthese*, 183(2), 161–174.
- Newman, M. E. (2001). Clustering and preferential attachment in growing networks. *Physical Review*, 64(2), Article 025102.
- O'Connor, C. (2017). The cultural red king effect. *Journal of Mathematical Sociology*, 41(3), 155–171.
- O'Connor, C. (2019). *The origins of unfairness: Social categories and cultural evolution*. USA: Oxford University Press.
- O'Connor, C., Bright, L. K., & Bruner, J. P. (2019). The emergence of intersectional disadvantage. *Social Epistemology*, 33(1), 23–41.
- O'Connor, C., & Bruner, J. (2019). Dynamics and diversity in epistemic communities. *Erkenntnis*, 84(1), 101–119.
- Ogburn, W. F., & Thomas, D. (1922). Are inventions inevitable? A note on social evolution. *Political Science Quarterly*, 37(1), 83–98.
- Paris, G., De Leo, G., Menozzi, P., & Gatto, M. (1998). Region-based citation bias in science. *Nature*, 396(6708), 210–210.
- Romero, F. (2017). Novelty versus replicability: Virtues and vices in the reward system of science. *Philosophy of Science*, 84(5), 1031–1043.
- Rubin, H. (2021). *Structural causes of citation gaps*. <http://philsci-archive.pitt.edu/18947/>.
- Rubin, H., & O'Connor, C. (2018). Discrimination and collaboration in science. *Philosophy of Science*, 85(3), 380–402.
- Schneider, M. D., Rubin, H., & O'Connor, C. (2021). Promoting diverse collaborations. In G. Ramsey, & A. De Block (Eds.), *The dynamics of science: Computational frontiers in history and philosophy of science*. University of Pittsburgh Press.

- Seevinck, M. (2016). Challenging the gospel: Grete Hermann on von Neumann's no-hidden-variables proof. In *Grete hermann-between physics and philosophy* (pp. 107–117). Springer.
- Strevens, M. (2003a). *Further properties of the priority rule*. manuscript. <http://www.strevens.org/research/scistruc/MorePrior.pdf>.
- Strevens, M. (2003b). The role of the priority rule in science. *Journal of Philosophy*, 100(2), 55–79.
- Strevens, M. (2011). Economic approaches to understanding scientific norms. *Episteme-Edinburgh*, 8(2), 184.
- Wang, Y. S., Lee, C. J., West, J. D., Bergstrom, C. T., & Erosheva, E. A. (2019). *Gender-based homophily in collaborations across a heterogeneous scholarly landscape*. arXiv preprint arXiv:1909.01284.
- Wardle, D. A. (1995). Journal citation impact factors and parochial citation practices. *The Bulletin of the Ecological Society of America*, 76(2), 102–104.
- West, J. D., Jacquet, J., King, M. M., Correll, S. J., & Bergstrom, C. T. (2013). The role of gender in scholarly authorship. *PloS One*, 8(7), Article e66212.
- Zollman, K. J. (2018). The credit economy and the economic rationality of science. *The Journal of Philosophy*, 115(1), 5–33.