Structural Inequality in Collaboration Networks

Rafael Ventura

University of Pennsylvania

forthcoming in *Philosophy of Science*

Abstract

Recent models of scientific collaboration show that minorities can end up at a disadvantage in bargaining scenarios. However, these models presuppose the existence of social categories. Here, we present a model of scientific collaboration in which inequality arises in the absence of social categories. We assume that all agents are identical except for the position that they occupy in the collaboration network. We show that inequality arises in the absence of social categories. We also show that this is due to the structure of the collaboration network and that similar patterns arise in two real-world collaboration networks.

Acknowledgements

I thank MindCORE and the Social and Cultural Evolution Working group at Penn for generous funding and a stimulating work environment. I also thank Erol Akçay for helpful discussions, two anonymous referees for their extremely constructive criticism, and especially Alex McAvoy for extensive feedback on earlier drafts.

Contact

112 Leidy Labs, 3740 Hamilton Walk
Philadelphia, 19104 PA
Email: rventura@upenn.edu

1 Introduction

Science is a social enterprise. For the most part, scientists do not work in isolation but collaborate with others when running experiments, analyzing data, or publishing papers. Scientific collaborations have in fact become more common over the past decades throughout academic disciplines (Melin and Persson, 1996; Henriksen, 2016). On the bright side, collaborations can bring about a host of epistemic and practical goods: collaborations seem to increase research output and impact (Beaver, 2004; Lee and Bozeman, 2005), and may even promote the attainment of truth by allowing researchers to pool resources and expertise (Wray, 2002).

But the social dimension of science can also bring about unequal outcomes, as philosophers of science have recently shown. Drawing on results from Bruner (2019) and O'Connor (2017), O'Connor and Bruner (2019) show that minorities can end up at a disadvantage in bargaining models of scientific collaboration merely because of their group size. Similar models suggests that minority disadvantage can hinder progress in epistemic communities (Rubin and O'Connor, 2018), and that intersectionality may aggravate the issue (O'Connor et al., 2019).[1]

Models of inequality in scientific collaboration can be very illuminating: they provide a possible account of how discrimination against minority groups might arise without explicit or implicit bias, or indeed without any difference between groups apart from size. But so far models of inequality in scientific collaboration presuppose the existence of social categories, with agents differing in some arbitrary but visible trait—e.g. race, gender, age, or membership in some other social group. One may therefore be led to conclude that social categories are the main or perhaps the only cause of inequality in epistemic communities. Conversely, it would be a lot more troublesome if inequality could arise in the absence of social categories. Inequality might then persist even if we could somehow erase the divides between distinct social groups.

[1]The social dimension of science can lead to outcomes that are undesirable for epistemic reasons as well. For example, community size and connectivity can restrict how quickly scientists converge on the truth (cf. Rosenstock et al., 2017; Zollman, 2007, 2010). When facing a risk-return trade-off in their work, individual scientists can divide cognitive labor in ways that are suboptimal for the community as a whole (Kummerfeld and Zollman, 2015); see also Kitcher (1990) and Weisberg and Muldoon (2009). Other social aspects of research, such as the influence of funding agencies, can bias epistemic communities and steer scientists away from the truth (Weatherall et al., 2020; Holman and Bruner, 2017).

Here, we present a model of scientific collaboration in which inequality arises in the absence of social categories. Our model represents a collaboration network where scientists must bargain over how much effort to invest in joint projects and how to divide credit for their labor. We then show that some scientists can end up at a disadvantage when all scientists are identical except for the position they occupy in the collaboration network. We also show that this unequal outcome is due to the structure of the collaboration network. Inequality thus emerges in the absence of biases or social categories, although biases and social categories may compound the problem.

The paper proceeds as follows. We begin by reviewing previous results in Section 2. We then describe and justify our model in Section 3. In Section 4, we report results from computer simulations showing that the structure of collaboration networks can lead to inequality in the absence of social categories. We also show that similar patterns arise in two real-world collaboration networks and that different dimensions of inequality can come apart. In Section 5, we discuss how our findings relate to previous work on bargaining models of scientific collaboration. We conclude in Section 6 by considering some limitations of our approach.

2 Previous Models

Recent models of scientific collaboration focus primarily on inequalities that arise due to social categories. There are good reasons for this, as inequality in scientific practice is often linked to social markers. The gender gap is a particularly well-documented case. Female scientists tend to publish fewer papers than male colleagues and are less likely to participate in collaborative research projects (West et al., 2013; Larivière et al., 2013). Female scientists also receive grants less often when funding agencies assess their quality as principal investigators, but not when agencies assess the quality of their research proposals (Witteman et al., 2019). There is further evidence that young female scientists are less likely to be listed as an author in a published paper, despite working more hours in total than male colleagues (Feldon et al., 2017). Similar patterns of discrimination arise with respect to race and ethnicity as well: in many disciplines, members of underrepresented racial and ethnic groups tend to have fewer publications and lower promotion rates (Hopkins et al., 2013; Gabbidon et al., 2004; Abelson et al., 2018).

In an effort to understand inequality of this form, previous models of scientific collaboration consider a simple version of the Nash demand game (Nash, 1950). In this game, two agents decide how to split a resource by demanding a portion of it. If the sum of their demands is equal to or less than the total amount available, each agent gets what they demand. If the sum of their demands exceeds the total amount, each agent gets nothing on the assumption that the negotiation breaks down when they cannot come to an agreement. For simplicity, we assume that agents can only make one of three possible demands: low (*Low*), medium (*Med*), or high (*High*). This is the mini-Nash demand game (Skyrms, 1996), with payoffs shown in Table 1.

Table 1: **Payoffs in the mini-Nash demand game.** In each cell, the first and second entries represent the payoff to the row and column players. Note that $L < M = 0.5 < H$ and $L + H = 1$.

	Low	*Med*	*High*
Low	L, L	L, 0.5	L, H
Med	0.5, L	0.5, 0.5	0, 0
High	H, L	0, 0	0, 0

When agents are perfectly rational, any two demands that sum to 1 is a pure Nash equilibrium of the game. Given any such configuration, neither agent has an incentive to unilaterally demand a different share of the resource. For example, there is an equilibrium where both agents demand *Med* and split the resource evenly. Such equilibria are usually termed “fair”. There are also mixed Nash equilibria in which agents mix two or all three demands with some positive probability. For example, there is an equilibrium in which one agent demands *Low* with probability L/H and the other demands *High* with probability $1 - L/H$. Such equilibria are usually called “unfair”.

Equilibrium results differ when agents are not perfectly rational and instead adjust their strategy via a process of biological or cultural evolution. Using the replicator dynamic as a model of evolution, Skyrms (1996) shows that there are only two equilibria in a population of agents playing the mini-Nash demand game: a symmetric equilibrium with agents who only play *Med*, and a mixed equilibrium with some agents playing *Low* and others playing *High*. Both equilibria are stable. But the equilibrium in which agents play *Low* and *High* is inefficient: when two agents demanding *Low* meet, each gets a positive payoff but a portion of the resource goes to waste.

This inefficient equilibrium can be avoided. If agents differ on the basis of arbitrary but visible group makers, agents can make their strategy conditional on the group membership of others. In this way, agents can coordinate on one of the efficient equilibria (Skyrms and Zollman, 2010). The population then evolves to either the symmetric equilibrium in which everyone plays Med, or the asymmetric equilibrium in which one group demands $High$ and the other group demands Low. The asymmetric equilibrium is known as a "discriminatory norm": a self-reinforcing pattern of behavior that puts some at a disadvantage merely because of group membership (Axtell et al., 2001).

Interesting outcomes are also possible when the population is divided into groups that have different sizes. Although the symmetric equilibrium is still stable in this case, Bruner (2019) and O'Connor (2017) show that the smaller the minority group is, the more likely the population is to evolve to an equilibrium with the minority demanding Low and the majority demanding $High$. Similar results have been observed in experiments where participants play the mini-Nash demand game in groups of different sizes (Mohseni et al., 2019). Under these conditions, the minority is more likely to demand Low because the minority encounters the majority more often than the other way around. As a result, the minority is faster to adapt to the demands of the majority. This outcome is the cultural analogue of the Red King effect: when two populations co-evolve, the population that is slower to adapt gains the evolutionary upper hand (Bergstrom and Lachmann, 2003).

Bargaining games such as the mini-Nash demand game have a long history as models of resource division (Skyrms, 1996; Binmore, 1998). Recently, the mini-Nash demand game has also been used to model the division of resources resulting from scientific collaborations. O'Connor and Bruner (2019), for example, use the mini-Nash demand game to show that members of the minority group can end up at a disadvantage in scientific collaboration simply because of their group size. Rubin and O'Connor (2018) draw on similar models to describe how discrimination can lead to segregation, which decreases the diversity of collaboration networks and is thus likley to hinder epistemic progress in science.

In the next section, we describe a model using the mini-Nash demand game to represent the division of resources resulting from scientific collaboration. But there are no social categories in our model. Yet, we show that inequality can arise because of the structure of the social network.

3 Model Description

The mini-Nash demand game captures important features of scientific collaborations (Rubin and O'Connor, 2018; O'Connor and Bruner, 2019). Scientists must often decide whether or not to enter a collaboration. If they choose to join the project, they must decide how to divvy up the credit for their joint labor. We therefore take a strategy in the mini-Nash demand game to represent a request for a certain amount of credit resulting from the joint project. One example of how a scientist might claim credit is by requesting to be first author. But there are other ways in which a scientist might claim credit. For example, a scientist might claim credit by explicitly describing their role in an author contribution statement, presenting results from the joint project at a conference, or promoting the project in social media. The Low strategy thus corresponds to a case in which a scientist requests a small amount of credit, the Med strategy to a case in which a scientist demands a moderate amount of credit, and the $High$ strategy to a case in which a scientist demands a large amount of credit. We assume throughout that collaborators do enough work to get an output of sufficient quality, thus ensuring that research quality is held constant.

Accordingly, the $Low-Low$ outcome might correspond to a case in which both scientists evince a certain level of timidity, do not promote the project in social media or do not present it at conferences, and therefore claim only a small amount of credit. In this case, both scientists split the credit evenly but claim a small amount of credit in total so each scientist ends up receiving a low payoff. In the $Med - Med$ outcome, both scientists claim a moderate amount of credit—for example, by promoting the project in social media or presenting it at conferences. In this case, scientists again split the credit evenly but each scientist claims a moderate amount of credit and so ends up receiving a moderate payoff. In the $Med - Low$ outcome, the scientist playing Med claims a moderate amount of credit while the scientist playing Low claims a small amount of credit. Thus, the Med scientist gets a moderate payoff and the Low scientist ends up with a small payoff. In the $High-High$ and the $High - Med$ outcomes, both scientists claim too much credit for themselves and conflict erupts between them. As a result, the collaboration breaks down and both are left with a payoff of zero.

In line with this interpretation of the Low, Med, and $High$ strategies, we use the mini-Nash demand game to represent the division of credit in scientific collaborations. In contrast to these models, however, we assume that there are no social categories. We make this assumption because in some cases inequality in science does not appear to be due to social categories, being rather linked to the structure of the social network. A case in point is the "Matthew effect" (Merton, 1968). The Matthew effect describes how more prominent scientists often get more credit than less prominent ones for work of equal worth. Since the mechanism was first proposed, empirical studies have confirmed that the Matthew effect is pervasive in science. For example, early work shows that inequality in publication counts increases as scientists age, suggesting a cumulative effect over time (Allison and Stewart, 1974; Allison et al., 1982). Recent work indicates that citation counts appear to depend in part on how renowned the author already is (Petersen et al., 2014). In fact, the problem seems to be getting worse (Nielsen and Andersen, 2021). A Matthew effect can also be seen in science funding, with recipients of early-career grants being more likely to win further grants than equally qualified peers (Bol et al., 2018).

In light of the evidence that inequality is not always directly due to social categories, we consider how inequality can arise in scientific communities in the absence of social categories. As there are no social categories in our model, we assume that scientists are identical except for the position they occupy in the collaboration network. In particular, we let scientists occupy the N nodes of a graph. Further, we let $e_{ij} = 1$ represent a link between scientists i and j if they collaborate on a joint project and $e_{ij} = 0$ otherwise. Scientist i then plays the mini-Nash demand game with every scientist j such that $e_{ij} = 1$. For simplicity, we assume that every scientist i plays the same strategy with all their collaborators. In each round of interaction, their total payoff is then given by the following expression:

$$\pi_i = \sum_{j}^{N} e_{ij} \cdot r_{ij} \quad , \tag{1}$$

where r_{ij} is the reward that i gets from interacting with j. The total payoff is thus the sum of rewards that a scientist receives from all their collaborators.[2]

As before, we suppose that scientists receive rewards according to Table 1. Since the values of L and H determine how large the gap is between the rewards that *Low* and *High* scientists get, we take these parameters to

[2]We consider the sum, and not the average, of rewards because it is more natural to think of scientists adding the rewards they receive from joint projects instead of averaging them. But results are the same if we instead take the average reward.

represent how "elitist" or "egalitarian" a scientific community is with respect to reward allocation. A large difference between L and H thus represents an elitist community where scientists either get a very low or a very high reward; in contrast, a small difference represents an egalitarian community where scientists mostly get the same reward. Indeed, scientific communities appear to differ in how unequal they are (Han, 2003; Clauset et al., 2015).[3]

To model the structure of the scientific community, we turn to scientometric studies on the topology of collaboration networks. Empirical evidence suggests that collaboration networks often have predictable properties, despite discipline-specific idiosyncrasies. In particular, collaboration networks tend to have a skewed degree distribution (Newman, 2001, 2004). This is to say that the distribution of the number of collaborators per scientist has a long tail, with collaboration networks displaying a hub-and-spoke architecture in which few scientists ("hubs") have many collaborators and many scientists ("spokes") have just a few. More precisely, the degree distribution of collaboration networks has the following form:

$$P(d) \sim d^{-\gamma} \quad , \tag{2}$$

where γ controls the shape of the distribution and d is the degree or the number of collaborators per scientist. Networks with a degree distribution of this form are known as "scale-free". A similar degree distribution is common in other social and biological networks, such as animal societies and gene regulatory networks (Barabási and Oltvai, 2004; Lusseau, 2003).

For this reason, we consider here scale-free networks with a power-law degree distribution. Although there are many models of network formation that result in such a distribution, a simple model that is known to generate a power-law degree distribution is the preferential-attachment model due to Barabási and Albert (1999). In this model of network formation, there is initially a small set of interconnected nodes. Nodes are then added to the network and connected to other nodes with probability proportional to the number of connections that existing nodes already have, giving rise to a

[3] As an anonymous referee points out, some academic communities have a reputation for being especially elitist—e.g. economics. At the same time, economics follows a strict norm of alphabetical author order implying equal contribution in collaborative work. This might be taken to mean that economics is an egalitarian discipline after all. However, it is possible that an alphabetical author order only makes a discipline more elitist: if authors do not disclose their real contribution to a joint project, others must resort to an author's past reputation or institutional affiliation to infer their real contribution.

Matthew effect in network formation. As the network grows, few nodes accumulate many connections and many nodes acquire only a few. In the limit of an infinitely large network, the resulting degree distribution converges on the power law given by equation (2). There are certainly more sophisticated models of network formation, but the preferential-attachment model is a simple and widely used one. For comparison, we consider regular networks in which every node has the same degree d and thus the average degree is also d. In particular, we consider regular networks with $d = 2$ and $d = 5$. These regular networks are not realistic but serve as control cases, as the scale-free networks we analyze have an average degree of about $d = 2$ (see Figure 1).

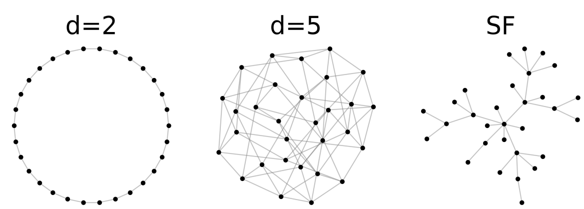

Figure 1: **Network topologies.** *Left*: regular network with $d = 2$. *Center*: regular network with $d = 5$. *Right*: scale-free network given by the preferential-attachment model described in Barabási and Albert (1999) with one initial node. Shown are networks with $N = 30$.

Another important feature of collaboration networks is that they are not static. Scientists sometimes change their behavior, choosing to collaborate when they did not before and vice-versa. There are of course many possible ways to represent this. Following O'Connor (2017), Rubin and O'Connor (2018), and O'Connor et al. (2019), we suppose that scientists update their behavior using a rule known as "myopic best response". This means that, in the first round of interaction, scientists choose a behavior at random. So a third of scientists plays Low, a third plays Med, and a third plays $High$. In each round thereafter, there is a small probability that a scientist updates their behavior. When a scientist updates their behavior, the scientist chooses the strategy that would have been a best response to the set of strategies that they encountered in the previous round. Scientists therefore update their behavior by best responding to previous plays but only keep a record of the most recent interactions.

Given our interest in the emergence of inequality in collaboration networks, we track how unequal the payoff distribution is. To do so, we use the Gini Index (GI). The GI measures the spread in a distribution. Although not entirely free of problems (Langel and Tillé, 2013), the GI is often used in economics to measure income and wealth inequality. It has also been applied to a variety of other contexts, such as in the study of biodiversity and enzyme selectivity (Wittebolle et al., 2009; Graczyk, 2007). The GI is given by:

$$GI = \frac{\sum_{i=1}^{N} \sum_{j=1}^{N} |\pi_i - \pi_j|}{2N \sum_{j=1}^{N} \pi_j} \tag{3}$$

where π_i and π_j are the payoffs that scientists i and j get from their collaborations. The numerator is the mean absolute difference of the payoff distribution and the denominator is twice the mean of the distribution. Since payoffs are always non-negative, the GI ranges from 0 (minimum) to 1 (maximum) depending on the spread of the distribution. The GI thus measures the spread in the payoff distribution.

But we show below that it is possible for different aspects of inequality to come apart. For example, heterogeneity in the distribution of strategies can be low while payoff inequality is high (and vice versa). For this reason, we introduce another measure to track heterogeneity in the distribution of strategies: the Strategy Heterogeneity Index (SI). Since agents get the same payoff when both play Med, we define the SI as the overall frequency of agents who play any of the two extreme strategies (i.e., *Low* and $High$). The SI is therefore given by:

$$SI = f_L + f_H \tag{4}$$

where f_L and f_H give the frequency of agents who play *Low* and $High$, respectively. The SI ranges from 0 (minimum) to 1 (maximum), with 0 indicating that everyone plays Med and 1 that no one plays Med. Unlike the GI, the SI therefore does not track the spread in the payoff distribution; it is instead a simple measure of how far the population deviates from the state in which everyone plays Med.

Having defined the structure of the collaboration network, the strategies that scientists in the collaboration network can adopt, the rule they use to update strategies, their payoffs, as well as two measures of inequality, we report our results in the next section. Pseudo-code, code for simulations, data, and scripts for analyses and figures are available anonymously at: `https://osf.`

io/h6j75/?view_only=479ac3174b8c4fbe8b6e2de1af3e5abe. Pseudo-code is also available in the Appendix.

4 Results

Computer simulations show that collaboration networks reach an equilibrium state in regular and scale-free networks. But regular and scale-free networks arrive at different equilibria. In regular networks with $d = 2$ and $d = 5$, the entire population comes to play Med when $L = 0.1$ (Figure 2, *left*). In scale-free networks, however, only about 70% of the population plays Med at equilibrium. Equilibria also differ when $L = 0.4$ (Figure 2, *right*). While the entire population continues to play Med in regular networks with $d = 5$, about 40% of the population comes to play Med in regular networks with $d = 2$. In scale-free networks, the share of the population playing Med is even smaller: about a third plays Med. The share of the population that plays Med at equilibrium therefore depends on not only network topology, but also average degree and value of L. (Since $L = 1 - H$, it does not matter whether we track L or H; we focus on L when presenting results.)

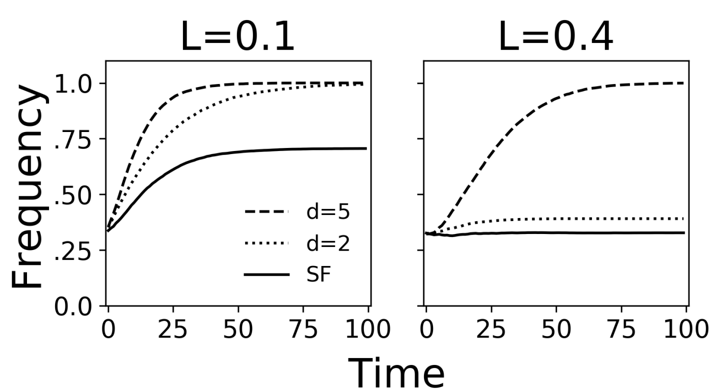

Figure 2: **Frequency of $\boldsymbol{Med}$ over time.** *Left*: when $L = 0.1$, Med takes over regular networks with $d = 2$ (*dotted*) and $d = 5$ (*dashed*); the equilibrium frequency of Med is 0.7 in scale-free networks (*solid*). *Right*: when $L = 0.4$, Med takes over regular networks with $d = 5$ but the frequency of Med is 0.4 in regular networks with $d = 2$ and 0.33 in scale-free networks. Results are average of 100 runs, update probability equal to 0.1, and $N = 100$.

We also find that the equilibrium composition of scale-free networks varies across values of L (Figure 3, *left*). When $L = 0.1$, 72% of the population play Med, while 19% play Low and 9% play $High$. With increasing values of L, the equilibrium frequency of Med goes down while the frequencies of Low and $High$ go up. When $L = 0.4$, the frequency of $High$ is higher than the frequency of Low: 40% of the population play $High$, while 35% play Med and 25% play Low. Depending on L, the population thus goes from having more agents who play Low than $High$ to having more $High$ than Low.

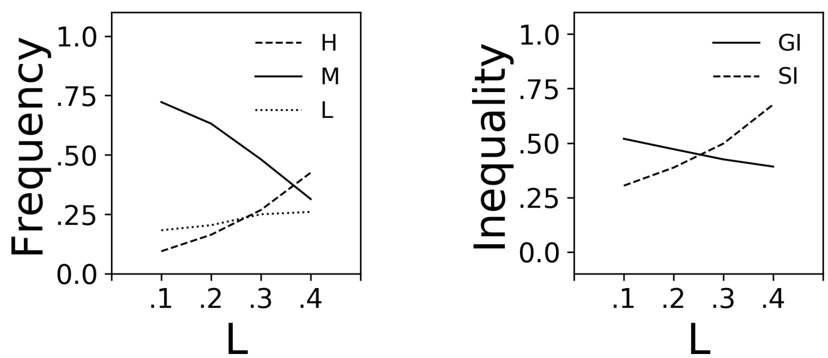

Figure 3: **Equilibrium Composition & Inequality.** *Left*: the equilibrium composition depends on L. *Right*: the Gini Index (GI) decreases with L, while the Strategy Heterogeneity Index (SI) increases with L. Results are average of 100 runs with 100 time steps, update probability equal to 0.1, and $N = 100$.

Next, we find that the payoff distribution becomes less unequal as L goes up (Figure 3, *right*). When $L = 0.1$, GI is about 0.52; when $L = 0.4$, GI is about 0.4. This is not very surprising given that higher (lower) values of L represent more egalitarian (elitist) communities. But the value of L has a very different effect on strategy heterogeneity: SI increases with L, with SI going from 0.3 when $L = 0.1$ to 0.66 when $L = 0.4$. These two measures also differ in that SI is more sensitive than GI to changes in the value of L: SI goes up by 120%, whereas GI goes down by 23%. As L increases, the population thus becomes less unequal with respect to payoff at the same time that it becomes a lot more heterogeneous with respect to its composition. In other words, payoff inequality and strategy heterogeneity come apart.

To better understand what factor(s) could be driving and maintaining payoff inequality and strategy heterogeneity, we consider how an agent's strategy depends on the position that they occupy in the collaboration network. In particular, we compare the degree of agents who play *Low* with those who play *High* (Figure 4, *left*). When $L = 0.1$, agents playing *High* tend to have a higher average degree than agents playing *Low*: the former have about 3.6 collaborators on average, while the latter have about 1.24. But when $L = 0.4$, the pattern is reversed: agents playing *Low* tend to have about 3 collaborators, while agents playing *High* have around 1.36. When L is low, those who play *High* therefore tend to be well-connected agents; when L is high, it is those playing *Low* who are more likely to be well-connected. Inspection of a representative network at equilibrium illustrates this point (Figure 4, *right*). When $L = 0.1$, agents playing *Low* tend to occupy more peripheral nodes than agents playing *High*. Given that agents are identical except for the position that they occupy in the collaboration network, this suggests that it is the structure of the network that drives and maintains inequality in our model.

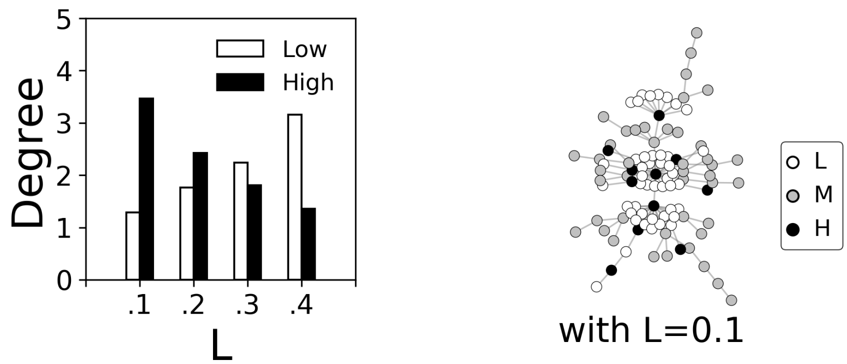

Figure 4: **Degree inequality in model networks.** *Left*: When L is low, the average degree of those playing *High* is higher than the average degree of those playing *Low*; the pattern is reversed when L is high. Results are average of 100 runs with 100 times steps, update probability equal to 0.1, and $N = 100$. *Right*: Population composition after 100 rounds of interactions in a scale-free collaboration network with $L = 0.1$.

But the structure of the collaboration network in our model is simply due to the preferential-attachment model. Although this model of network formation gives rise to a degree distribution that is known to resemble the degree distribution of real-world collaboration networks, it is clearly an idealization. For one, scientists do not always choose who to collaborate with on the basis of how many collaborations potential coworkers already have—among myriad other factors, geographical proximity, institutional affiliation, and personality quirks can also play a role. To examine whether the inequality we observe in our model might arise in the real world, we study the same dynamics of collaboration on two well-known and publicly available collaboration networks: the *GR-QC* and the *Erdos* collaboration network. The *GR-QC* collaboration network includes the authors of papers on general relativity and quantum cosmology posted to the pre-print repository arXiv between 1993 and 2003 (Leskovec et al., 2007). The *Erdos* collaboration network covers all papers written by the extremely prolific mathematician Paul Erdős, his co-authors, and their co-authors (Batagelj and Mrvar, 2000).

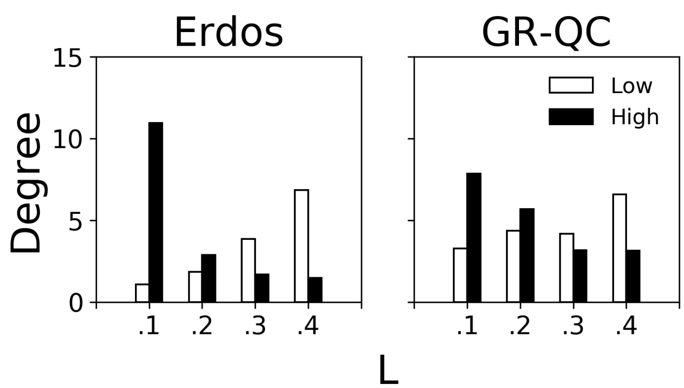

Figure 5: **Degree inequality in real-world networks.** In the *Erdos* ($N = 4,158$; *left*) and the *GR-QC* ($N = 5,094$; *right*) collaboration networks, the average degree of agents who play $High$ is higher than the average degree of agents who play *Low* when L is low; the pattern is reversed when L is high. Results are average of 100 runs with 100 time steps and update probability equal to 0.1.

We obtain similar results from simulations of a population of agents playing the mini-Nash demand game with myopic best response on the *GR-QC* and the *Erdos* collaboration networks (Figure 5). In particular, the average degree is higher for agents playing $High$ than for agents playing *Low* when L is low but the pattern is reversed when L is high. When $L = 0.1$, scientists in *GR-QC* who play *Low* have about 3.1 collaborators on average, while scientists who play $High$ have about 7.9 collaborators. A similar pattern holds in *Erdos*: when $L = 0.1$, scientists playing *Low* have a single collaborator on average but scientists playing $High$ have about 10.9 collaborators. As L goes up, this difference decreases at first and eventually reverses. When $L = 0.4$, scientists in *GR-QC* who play *Low* have about 6.37 collaborators on average, while scientists playing $High$ have about 2.94. Similarly, scientists in *Erdos* who play *Low* have 7 collaborators on average, while scientists playing $High$ have about 1.46. Network structure therefore drives the emergence of inequality in both networks, although the effect is especially pronounced in *Erdos*.

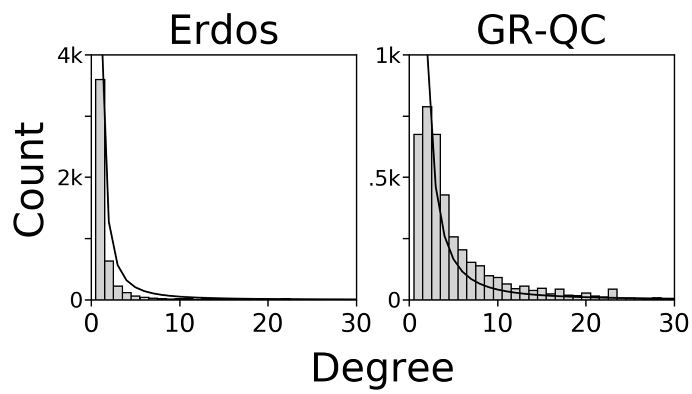

Figure 6: **Degree distribution in model and two real-world networks.** *Left*: the degree distribution given by $P(d) = N \cdot d^{-\gamma}$ with $\gamma = 2$ (*solid line*) approximates the degree distribution in the *Erdos* collaboration network ($N = 4,158$). *Right*: the same expression approximates the observed degree distribution in the *GR-QC* collaboration network ($N = 5,094$). Grey bars show empirical degree distribution.

It is also worth reiterating that the degree distribution of scale-free networks where inequality arises is similar to that of real-world collaboration networks. As already noted, the degree distribution of indefinitely large scale-free networks is given by $P(d) \sim d^{-\gamma}$. Empirical studies find that values of γ for real-world collaboration networks often range between values of 1 and 3, depending on dataset and scientific discipline (Barabási et al., 2002; Albert and Barabási, 2002). Indeed, this expression approximates quite well the degree distribution of both the *Erdos* and the *GR-QC* collaboration networks (Figure 6). Considering that the preferential-attachment model was built to fit the scale-free degree distribution of real-world networks, this is not very surprising. But it serves as a reminder that the inequality we observe in our model is the product of a realistic network structure.

5 Discussion

Our model shows that the structure of collaboration networks can give rise to inequality even in the absence of social categories. In particular, our model shows that inequality in the payoff distribution and heterogeneity in the strategy profile of the population arises and persists in collaboration networks with a heterogeneous degree distribution. Our model also shows that this is so across the full range of values for L—a parameter that controls how elitist or egalitarian the scientific community tends to be. Furthermore, our model highlights that inequality is not a one-dimensional concept: different values of L affect different measures of inequality differently, with inequality in the payoff distribution (GI) being high when heterogeneity in the strategy profile (SI) is low and vice-versa.

These results stand in contrast to previous models showing that population structure can promote an even allocation of resources in the mini-Nash demand game. For example, Alexander and Skyrms (1999) and Alexander (2000) show that spatial structure makes it very likely that a population will converge on the fair equilibrium. But this is due to the fact that spatial organization is a form of population structure where every agents interacts with four neighbors and there is no variation in the degree distribution. When population structure leads many to interact with few and few to interact with many, our model shows that the resulting heterogeneous degree distribution can promote unequal outcomes.

Our model thus adds to a growing body of work showing that a heterogeneous degree distribution can give rise to inequalities in strategic settings. In a network model of the Prisoner's Dilemma, for example, Du et al. (2008), find that a heterogeneous degree distribution favors the spread of cooperation but that it also promotes an unequal payoff distribution. In public goods games, network heterogeneity induces diversity in group size and thus promotes contributions to the public good (Santos et al., 2006, 2008). But network heterogeneity can also lead to unequal outcomes in public good games, as the proliferation of altruistic behaviors ends up harming some individuals (McAvoy et al., 2020).

Our model also reveals two "regimes" in the emergence of inequality in collaboration networks. One regime is when L is low. In this case, poorly connected scientists in the periphery of the collaboration network play Low, while their well-connected collaborators play $High$. The other regime is when L is high. In this case, well-connected scientists play Low, while their poorly connected collaborators play $High$. An analogous pattern is apparent in the way that the Red King/Queen effect leads to inequality in the mini-Nash bargaining game with co-evolving groups of different sizes (Bruner, 2019; O'Connor, 2019; O'Connor, 2017). When L is high, the Red King effect leads the minority to get less than the majority. When L is low, the Red Queen effect kicks in and the minority gets more than the majority.

Despite this superficial similarity, the mechanism driving the emergence of inequality in our model is not the same as in the Red King/Queen. First, the Red King/Queen depends on the minority adapting more quickly to the strategy of the majority. In contrast, the update rule we use is the myopic best response. Strictly speaking, the myopic best response is not an evolutionary update rule because agents do not update their behavior by copying the behavior of others. So it is not a difference in evolutionary tempo that drives inequality in our model. Second, the Red King/Queen relies on there being two groups, groups having different sizes, and individuals conditionalizing their behavior on the group membership of others. In our model, however, the mechanism that gives rise to inequality does not depend on a categorical distinction between groups. In fact, there is no partition of the population into groups at all—let alone groups of different sizes. Third, the Red King/Queen effect causes the minority groups to be at a disadvantage when L is high and thus when payoff inequality is low. But in our model those who are poorly connected end up at a disadvantage when L is low and payoff inequality is high. For all these reasons, the mechanism leading to inequality in our model is not the same as the Red King/Queen.

So what explains the two regimes of inequality that we observe in our model? Since the update rule we use is the myopic best response, to answer this question we follow Rubin and O'Connor's (2018, pp. 386-8) account of how discrimination arises in their model and consider the probability that a strategy is a best response.[4] A strategy is a best response if there is no other strategy that would yield a higher payoff given the strategies that other agents play in the previous round. The probability that a particular strategy is a best response thus depends on the probability with which other agents

[4] We thank an anonymous referee for raising this point.

choose each strategy. For an agent who only interacts with one other agent, the probability that the strategy *Low*, *Med*, or *High* is a best response is just the probability with which the agent encounters another agent who plays *High*, *Med*, or *Low*. Initially, agents choose a strategy at random. The initial probability that each strategy is a best response is thus $\frac{1}{3}$.

In scale-free networks, some agents do interact with only one other agent. But other agents interact with many more. In such cases, the probability that a strategy is a best response can be found in three steps. The first step is to determine what strategy is a best response to every possible combination of strategies that other agents may choose. The second step is to calculate the probability with which each one of these combinations of strategies occurs. The third step is to compute the probability that a strategy is a best response by summing over the probabilities of every combination of strategies to which the strategy in question is a best response. Assuming that agents pick a strategy at random, as they do at first, the probability that *Low*, *Med*, or *High* is a best response is shown in Figure 7.

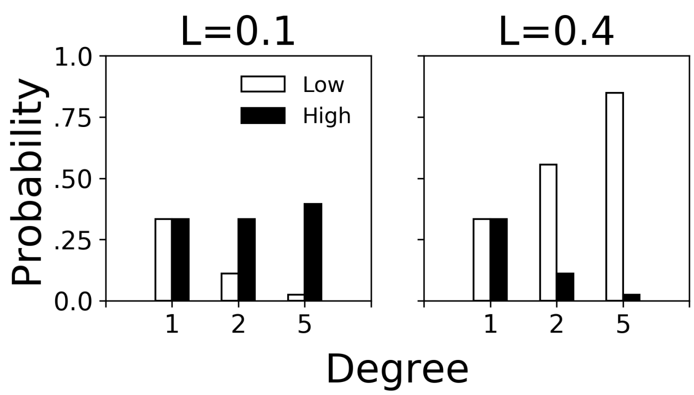

Figure 7: **Initial probabilities that *Low* and *High* is a best response.** *Left*: initial probability that *Low* and *High* are a best response for $d = 1$, $d = 2$, and $d = 5$ when $L = 0.1$. *Right*: initial probability that *Low* and *High* are a best response for $d = 1$, $d = 2$, and $d = 5$ when $L = 0.4$.

Notice that the probability that a strategy is a best response depends on degree. As already noted, each strategy is a best response with probability $\frac{1}{3}$ when an agent interacts with only one other agent—and this is so regardless of L. But when an agent interacts with more than one agent, the probability that a strategy is a best response depends on how many other agents they interact with. When $L = 0.1$, for example, the probability that *Low* is a best response for an agent who interacts with two other agents is about 0.11. But the probability that *Low* is a best response for an agent who interacts with five other agents is only 0.025. When $L = 0.4$, the probability that *Low* is a best response for an agent who interacts with two other agents is about 0.55. But the probability that *Low* is a best response for an agent who interacts with five other agents is about 0.85.

This allows us to gain some insight into the two regimes for the emergence of inequality in our model. Consider two groups of agents: poorly connected agents with $d = 1$, and well-connected agents with $d \geq 5$. When $L = 0.1$, the initial probability that *Low* or $High$ is a best response for poorly connected agents is one third. But for well-connected agents the initial probability that $High$ is a best response is a lot higher than the initial probability that *Low* is a best response. This is because the relative payoff to $High$ is relatively high, so well-connected individuals respond best by "sticking to their guns" and making a $High$ demand that yields a large increase in payoff. For this reason, well-connected agents tend to play $High$ and end up at an advantage when L is low; at the same time, poorly connected agents tend to play *Low* and end up at a disadvantage. When $L = 0.4$, the initial probability that *Low* or $High$ is a best response for poorly connected agents is again one third. For well-connected agents, however, the initial probability that *Low* is a best response is now a lot higher than the initial probability that $High$ is a best response. This is because the relative payoff to *Low* is relatively high, so well-connected individuals respond best by playing it safe and making a *Low* demand instead of holding out for what would be a small increase in payoff. Well-connected individuals therefore tend to play *Low* and end up at a disadvantage when L is high, while poorly connected agents play $High$ and end up at an advantage. The two regimes of inequality we observe in scale-free networks is thus due to differences in the initial probability that a strategy is a best response.[5]

From a social-epistemological perspective, this raises a series of important questions about the structure of collaboration networks. Well-connected

[5]The initial probabilities that either *Low* or $High$ is a best response is higher when $L = 0.4$ than when $L = 0.1$ for $d \geq 2$. This helps explain why a smaller share of the population comes to play Med in scale-free networks and regular networks with $d = 2$ when L is high. In regular networks with $d = 5$, the initial probability that *Low* is a best response is so high that the population quickly becomes saturated with *Low*. This decreases the probability that *Low* is a best response and allows Med to take over.

scientists are more likely to play *Low* and end up at a disadvantage when L is high. This means that well-connected scientists are at a disadvantage in egalitarian communities where payoff inequality is low. Poorly connected scientists, however, are more likely to play *Low* and thus end up at a disadvantage when L is low. Low values of L correspond to elitist communities where payoff inequality is high. Our model therefore raises the specter of a two-fold harm: low values of L put poorly connected scientists at a disadvantage when doing so is particularly harmful.

The two-fold harm of structural inequality is all the more worrisome because members of minority or underrepresented groups are often poorly connected in real-world collaboration networks. Female scientists, for example, have fewer collaborators than their male colleagues (Araujo et al., 2017; Abramo et al., 2009). Black scientists also have fewer collaborators, at least in some disciplines (Del Carmen and Bing, 2000). When payoff inequality is especially high, the two-fold harm is likely to arise and members of these groups might therefore be at a disadvantage. To make matters worse, implicit and explicit biases linked to social categories might only exacerbate the problem: prejudice and discrimination tends to put those groups at a disadvantage who are already vulnerable due to the position that they occupy in the collaboration network. For example, if scientists choose what collaborations to enter on the basis of biases against visible group markers, then biases and social categories might contribute to the formation of collaboration networks where pernicious forms of structural inequality are likely to emerge.

6 Conclusion

Philosophers have long worried that implicit and explicit biases are inevitable in science and that they often contribute to various forms of epistemic injustice (Longino, 1990; Fricker, 2007). In recent years, formal models in philosophy of science have further shown that it is possible for discriminatory norms to lead to an unequal allocation of epistemic credit even when there are no biases (O'Connor and Bruner, 2019; Rubin and O'Connor, 2018; O'Connor et al., 2019). But models proposed so far account for these worrisome patterns in research by positing the existence of social categories. Although biases and social categories remain a source of concern, we show that unequal outcomes are possible even in the absence of social categories:

when scientists bargain with collaborators in a scale-free network, inequality arises simply because of the structure of the collaboration network. We also bring empirical considerations to bear on models of the social organization of science by showing that structural inequality can likewise arise in real-world collaboration networks (cf. Martini and Pinto, 2017).

It is important to keep in mind, however, that our model makes several simplifying assumptions. First, we assume that scientist play the same strategy with all their collaborations. This is unlikely to hold in reality since scientists often negotiate different arrangements with different collaborators. Second, we consider a dynamic population of scientists who change their strategies over time but assume that the structure of the collaboration network is static. This is not the case in the real world where scientists can not only update their behavior, but also adjust their social ties. Third, we assume that all scientists are equally competent. This is again unrealistic because scientists often differ with respect to how productive they are. Fourth, we assume that scientists update their strategy by myopic best response. This is a reasonable assumption but update rules based on imitation are also plausible. While these simplifying assumptions allow us to isolate and better understand an important phenomenon, it would be interesting to relax these assumptions. Future work could therefore consider collaboration networks where scientists pursue different strategies with different collaborators, change who they interact with over time, differ with respect to how productive they are, or update their strategy according to different rules.

Appendix

We use a simple program to simulate the behavior of agents in a network who interact with their neighbors by playing the mini-Nash demand game. In pseudo-code, the program proceeds as follows:

```
FOR each Network Topology, DO:

   FOR each Agent, DO:
      Choose Demand at random from options L, M, and H

      FOR each Time Step, DO:

         FOR each Agent, DO:
            Get Agent's Demand
            Get Demand for each of Agent's neighbors
            Get Agent's payoff based on own Demand and neighbors' Demands

            With probability 0.1, DO:
               Find Agent's Best Response in previous Time Step
               Update Agent's Demand
```

References

Abelson, J. S., N. Z. Wong, M. Symer, G. Eckenrode, A. Watkins, and H. L. Yeo (2018). Racial and ethnic disparities in promotion and retention of academic surgeons. *The American Journal of Surgery 216*(4), 678–682.

Abramo, G., C. D'Angelo, and A. Caprasecca (2009). The contribution of star scientists to overall sex differences in research productivity. *Scientometrics 81*(1), 137–156.

Albert, R. and A.-L. Barabási (2002). Statistical mechanics of complex networks. *Reviews of Modern Physics 74*(1), 47.

Alexander, J. M. (2000). Evolutionary explanations of distributive justice. *Philosophy of Science 67*(3), 490–516.

Alexander, J. M. and B. Skyrms (1999). Bargaining with neighbors: Is justice contagious? *The Journal of Philosophy 96*(11), 588–598.

Allison, P. D., J. S. Long, and T. K. Krauze (1982). Cumulative advantage and inequality in science. *American Sociological Review 47*(5), 615–625.

Allison, P. D. and J. A. Stewart (1974). Productivity differences among scientists: Evidence for accumulative advantage. *American sociological review 39*(4), 596–606.

Araujo, E. B., N. A. Araújo, A. A. Moreira, H. J. Herrmann, and J. S. Andrade Jr (2017). Gender differences in scientific collaborations: Women are more egalitarian than men. *PloS one 12*(5), e0176791.

Axtell, R. L., J. M. Epstein, and H. P. Young (2001). The emergence of classes in a multiagent bargaining model. *Social dynamics 27*, 191–211.

Barabási, A.-L. and R. Albert (1999). Emergence of scaling in random networks. *Science 286*(5439), 509–512.

Barabási, A.-L., H. Jeong, Z. Néda, E. Ravasz, A. Schubert, and T. Vicsek (2002). Evolution of the social network of scientific collaborations. *Physica A: Statistical mechanics and its applications 311*(3-4), 590–614.

Barabási, A.-L. and Z. N. Oltvai (2004). Network biology: understanding the cell's functional organization. *Nature Reviews Genetics 5*(2), 101–113.

Batagelj, V. and A. Mrvar (2000). Some analyses of erdos collaboration graph. *Social Networks 22*(2), 173–186.

Beaver, D. d. (2004). Does collaborative research have greater epistemic authority? *Scientometrics 60*(3), 399–408.

Bergstrom, C. T. and M. Lachmann (2003). The red king effect: when the slowest runner wins the coevolutionary race. *Proceedings of the National Academy of Sciences 100*(2), 593–598.

Binmore, K. G. (1998). *Game theory and the social contract: just playing*, Volume 2. MIT press.

Bol, T., M. de Vaan, and A. van de Rijt (2018). The matthew effect in science funding. *Proceedings of the National Academy of Sciences 115*(19), 4887–4890.

Bruner, J. P. (2019). Minority (dis) advantage in population games. *Synthese 196*(1), 413–427.

Clauset, A., S. Arbesman, and D. B. Larremore (2015). Systematic inequality and hierarchy in faculty hiring networks. *Science advances 1*(1), e1400005.

Del Carmen, A. and R. L. Bing (2000). Academic productivity of african americans in criminology and criminal justice. *Journal of Criminal Justice Education 11*(2), 237–249.

Du, W.-B., H.-R. Zheng, and M.-B. Hu (2008). Evolutionary prisoner's dilemma game on weighted scale-free networks. *Physica A: Statistical Mechanics and its Applications 387*(14), 3796–3800.

Feldon, D. F., J. Peugh, M. A. Maher, J. Roksa, and C. Tofel-Grehl (2017). Time-to-credit gender inequities of first-year phd students in the biological sciences. *CBE—Life Sciences Education 16*(1), ar4.

Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing.* Oxford University Press.

Gabbidon, S. L., H. T. Greene, and K. Wilder (2004). Still excluded? an update on the status of african american scholars in the discipline of criminology and criminal justice. *Journal of Research in Crime and Delinquency 41*(4), 384–406.

Graczyk, P. P. (2007). Gini coefficient: a new way to express selectivity of kinase inhibitors against a family of kinases. *Journal of Medicinal Chemistry 50*(23), 5773–5779.

Han, S.-K. (2003). Tribal regimes in academia: A comparative analysis of market structure across disciplines. *Social networks 25*(3), 251–280.

Henriksen, D. (2016). The rise in co-authorship in the social sciences (1980–2013). *Scientometrics 107*(2), 455–476.

Holman, B. and J. Bruner (2017). Experimentation by industrial selection. *Philosophy of Science 84*(5), 1008–1019.

Hopkins, A. L., J. W. Jawitz, C. McCarty, A. Goldman, and N. B. Basu (2013). Disparities in publication patterns by gender, race and ethnicity based on a survey of a random sample of authors. *Scientometrics 96*(2), 515–534.

Kitcher, P. (1990). The division of cognitive labor. *The Journal of Philosophy 87*(1), 5–22.

Kummerfeld, E. and K. J. Zollman (2015). Conservatism and the scientific state of nature. *The British Journal for the Philosophy of Science 67*(4), 1057–1076.

Langel, M. and Y. Tillé (2013). Variance estimation of the gini index: revisiting a result several times published. *Journal of the Royal Statistical Society: Series A (Statistics in Society) 176*(2), 521–540.

Larivière, V., C. Ni, Y. Gingras, B. Cronin, and C. R. Sugimoto (2013). Bibliometrics: Global gender disparities in science. *Nature News 504*(7479), 211.

Lee, S. and B. Bozeman (2005). The impact of research collaboration on scientific productivity. *Social Studies of Science 35*(5), 673–702.

Leskovec, J., J. Kleinberg, and C. Faloutsos (2007). Graph evolution: Densification and shrinking diameters. *ACM transactions on Knowledge Discovery from Data (TKDD) 1*(1), 2–es.

Longino, H. E. (1990). *Science as social knowledge: Values and objectivity in scientific inquiry*. Princeton University Press.

Lusseau, D. (2003). The emergent properties of a dolphin social network. *Proceedings of the Royal Society of London. Series B: Biological Sciences 270*(suppl_2), S186–S188.

Martini, C. and M. F. Pinto (2017). Modeling the social organization of science. *European Journal for Philosophy of Science 7*(2), 221–238.

McAvoy, A., B. Allen, and M. A. Nowak (2020). Social goods dilemmas in heterogeneous societies. *Nature Human Behaviour 4*(8), 819–831.

Melin, G. and O. Persson (1996). Studying research collaboration using co-authorships. *Scientometrics 36*(3), 363–377.

Merton, R. K. (1968). The matthew effect in science: The reward and communication systems of science are considered. *Science 159*(3810), 56–63.

Mohseni, A., C. O'Connor, and H. Rubin (2019). On the emergence of minority disadvantage: testing the cultural red king hypothesis. *Synthese*, 1–23.

Nash, J. F. J. (1950). The bargaining problem. *Econometrica 18*(2), 155–162.

Newman, M. E. (2001). The structure of scientific collaboration networks. *Proceedings of the National Academy of Sciences 98*(2), 404–409.

Newman, M. E. (2004). Coauthorship networks and patterns of scientific collaboration. *Proceedings of the National Academy of Sciences 101*(suppl 1), 5200–5205.

Nielsen, M. W. and J. P. Andersen (2021). Global citation inequality is on the rise. *Proceedings of the National Academy of Sciences 118*(7), 1–10.

O'Connor, C. (2019). *The origins of unfairness: Social categories and cultural evolution*. Oxford University Press, USA.

O'Connor, C. (2017). The cultural red king effect. *The Journal of Mathematical Sociology 41*(3), 155–171.

O'Connor, C., L. K. Bright, and J. P. Bruner (2019). The emergence of intersectional disadvantage. *Social Epistemology 33*(1), 23–41.

O'Connor, C. and J. Bruner (2019). Dynamics and diversity in epistemic communities. *Erkenntnis 84*(1), 101–119.

Petersen, A. M., S. Fortunato, R. K. Pan, K. Kaski, O. Penner, A. Rungi, M. Riccaboni, H. E. Stanley, and F. Pammolli (2014). Reputation and impact in academic careers. *Proceedings of the National Academy of Sciences 111*(43), 15316–15321.

Rosenstock, S., J. Bruner, and C. O'Connor (2017). In epistemic networks, is less really more? *Philosophy of Science 84*(2), 234–252.

Rubin, H. and C. O'Connor (2018). Discrimination and collaboration in science. *Philosophy of Science 85*(3), 380–402.

Santos, F. C., J. M. Pacheco, and T. Lenaerts (2006). Evolutionary dynamics of social dilemmas in structured heterogeneous populations. *Proceedings of the National Academy of Sciences 103*(9), 3490–3494.

Santos, F. C., M. D. Santos, and J. M. Pacheco (2008). Social diversity promotes the emergence of cooperation in public goods games. *Nature 454*(7201), 213–216.

Skyrms, B. (1996). *Evolution of the social contract*. Cambridge University Press.

Skyrms, B. and K. J. Zollman (2010). Evolutionary considerations in the framing of social norms. *Politics, philosophy & economics 9*(3), 265–273.

Weatherall, J. O., C. O'Connor, and J. P. Bruner (2020). How to beat science and influence people: policymakers and propaganda in epistemic networks. *The British Journal for the Philosophy of Science 71*(4), 1157–1186.

Weisberg, M. and R. Muldoon (2009). Epistemic landscapes and the division of cognitive labor. *Philosophy of Science 76*(2), 225–252.

West, J. D., J. Jacquet, M. M. King, S. J. Correll, and C. T. Bergstrom (2013). The role of gender in scholarly authorship. *PloS one 8*(7), e66212.

Wittebolle, L., M. Marzorati, L. Clement, A. Balloi, D. Daffonchio, K. Heylen, P. De Vos, W. Verstraete, and N. Boon (2009). Initial community evenness favours functionality under selective stress. *Nature 458*(7238), 623–626.

Witteman, H. O., M. Hendricks, S. Straus, and C. Tannenbaum (2019). Are gender gaps due to evaluations of the applicant or the science? a natural experiment at a national funding agency. *The Lancet 393*(10171), 531–540.

Wray, K. B. (2002). The epistemic significance of collaborative research. *Philosophy of Science 69*(1), 150–168.

Zollman, K. J. (2007). The communication structure of epistemic communities. *Philosophy of Science 74*(5), 574–587.

Zollman, K. J. (2010). The epistemic benefit of transient diversity. *Erkenntnis 72*(1), 17.