

ARTICLE

Structural Inequality in Collaboration Networks

Rafael Ventura 

University of Pennsylvania, 112 Leidy Labs, 3740 Hamilton Walk, Philadelphia, 19104 PA, United States
E-mail: rventura@upenn.edu

(Received 29 March 2021; revised 09 May 2022; accepted 11 May 2022; first published online 10 June 2022)

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, I present a model of scientific collaboration in which inequality arises in the absence of social categories. I assume that all agents are identical except for the position that they occupy in the collaboration network. I show that inequality arises in the absence of social categories. I also show that this is due to the structure of the collaboration network and that similar patterns arise in two real-world collaboration networks.

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 they 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 suggest that a minority disadvantage can hinder progress in epistemic communities (Rubin and O'Connor 2018) and that intersectionality may aggravate the issue (O'Connor, Bright, and Bruner, 2019).¹

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, Bruner, and O'Connor 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, O'Connor, and Bruner 2020; Holman and Bruner 2017).

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—for example, 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.

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

2. Previous Models

Recent models of scientific collaboration focus primarily on inequalities that arise as a result of 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 articles 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 article, 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

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 $L + H = 1$

	<i>Low</i>	<i>Med</i>	<i>High</i>
<i>Low</i>	L, L	$L, 0.5$	L, H
<i>Med</i>	$0.5, L$	$0.5, 0.5$	$0, 0$
<i>High</i>	H, L	$0, 0$	$0, 0$

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 the payoffs shown in table 1.

When agents are perfectly rational, any two demands that sum to 1 result in 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 markers, 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 their 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, O'Connor, and Rubin 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 coevolve, 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 collaborations 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 likely to hinder epistemic progress in science.

In the next section, I describe a model that uses the mini-Nash demand game to represent the division of resources resulting from scientific collaboration. There are no social categories in my model. Yet, I 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. I therefore consider 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 the 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 through 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. I 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 through 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 through 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, whereas 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 *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, I use the mini-Nash demand game to represent the division of credit in scientific collaborations. In contrast to other models, however, I assume that there are no social categories. I make this assumption because in some cases, inequality in science does not appear to be due to social categories, instead being 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, the model shows how inequality can arise in scientific communities in the absence of social categories. Because there are no social categories in the 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:

$$\pi_i = \sum_j^N e_{ij} \cdot r_{ij}, \quad (1)$$

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. Because 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 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

² We consider the sum of rewards rather than the average 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.

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:

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

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, here we consider 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 attributable 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 a probability proportional to the number of connections that existing nodes already have, giving rise to a 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 will consider regular networks in which every node has the same degree d and thus the average degree is also d . In particular, we will consider regular networks with $d = 2$ and $d = 5$. These regular networks are not realistic but serve as control cases because the scale-free networks discussed here have an average degree of approximately $d = 2$ (see figure 1).

Another important feature of collaboration networks is that they are not static. Scientists sometimes change their behavior, for example, 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 will suppose that scientists update their behavior using

³ As an anonymous referee points out, some academic communities have a reputation for being especially elitist—for example, economics. At the same time, economics follows a strict norm of alphabetical author order implying equal contribution in collaborative works. 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.




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 by Barabási and Albert (1999) with one initial node. Shown are networks with $N = 30$.

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 will play *Low*, a third will play *Med*, and a third will play *High*. In each round thereafter, there is a small probability that a scientist will update 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 keep a record of only the most recent interactions.

Given our interest in the emergence of inequality in collaboration networks, we wish to 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 the following:

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

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. Because payoffs are always nonnegative, 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 I show later in the article 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 need to introduce another measure to track heterogeneity in the distribution of strategies: the strategy heterogeneity index (SI). Because 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 the following:

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

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




Figure 2. Frequency of *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 the average of 100 runs, with update probability equal to 0.1, and $N = 100$.

Med and 1 indicating 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, and two measures of inequality, I report the 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 plays *Med* when $L = 0.1$ (figure 2, left). In scale-free networks, however, only approximately 70% of the population plays *Med* at equilibrium. Equilibria also differ when $L = 0.4$ (figure 2, right). Whereas the entire population continues to play *Med* in regular networks with $d = 5$, only approximately 40% of the population plays *Med* in regular networks with $d = 2$. In scale-free networks, the share of the population playing *Med* is even smaller: approximately a third will play *Med*. The share of the population that plays *Med* at equilibrium therefore depends on not only network topology but also the average degree and value of L . (Because $L = 1 - H$, it does not matter whether we track L or H ; I focus on L when presenting results.)

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 will play *Med*, whereas 19% will play *Low* and 9% will 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




Figure 3. Equilibrium composition and inequality. *Left:* The equilibrium composition depends on L . *Right:* The GI decreases with L , whereas the SI increases with L . Results are the average of 100 runs with 100 time steps, with update probability equal to 0.1, and $N = 100$.

population plays *High*, whereas 35% plays *Med* and 25% plays *Low*. Depending on L , the population thus goes from having more agents who play *Low* than *High* to having more agents who play *High* than *Low*.

Next, we find that the payoff distribution becomes less unequal as L goes up (figure 3, *right*). When $L = 0.1$, GI is approximately 0.52; when $L = 0.4$, GI is approximately 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 the 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 will 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 has approximately 3.6 collaborators on average, whereas the latter has approximately 1.24. But when $L = 0.4$, the pattern is reversed: agents playing *Low* tend to have approximately 3 collaborators, whereas 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 the model.

But the structure of the collaboration network in the 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

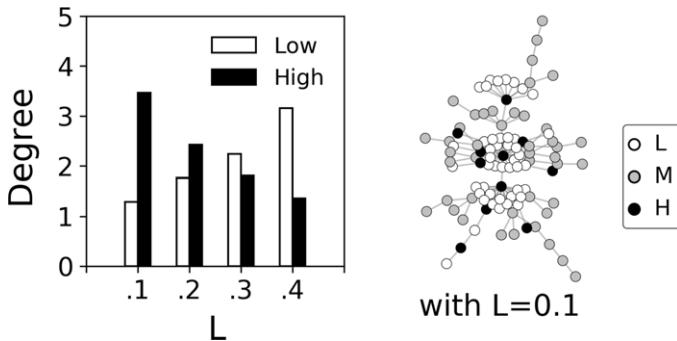


Figure 4. Degree inequality in model networks. *Left:* When L is low, the average degree of those playing *High* is higher than the average degree of those playing *Low*; the pattern is reversed when L is high. Results are the average of 100 runs, with 100 times steps, update probability equal to 0.1,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

not always choose whom 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 the model might arise in the real world, we will consider the same dynamics of collaboration on two well-known and publicly available collaboration networks: the *GR-QC* and *Erdos* collaboration networks. The *GR-QC* collaboration network includes the authors of articles on general relativity and quantum cosmology posted to the preprint repository arXiv between 1993 and 2003 (Leskovec et al. 2007). The *Erdos* collaboration network covers all articles written by the extremely prolific mathematician Paul Erdős, his coauthors, and their coauthors (Batagelj and Mrvar 2000).

Similar results are obtained 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 approximately 3.1 collaborators on average, whereas scientists who play *High* have approximately 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 approximately 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 approximately 6.37 collaborators on average, whereas scientists playing *High* have approximately 2.94. Similarly, scientists in *Erdos* who play *Low* have 7 collaborators on average, whereas scientists playing *High* have approximately 1.46. Network structure therefore drives the emergence of inequality in both networks, although the effect is especially pronounced in *Erdos*.

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 the values of γ for real-world




Figure 5. Degree inequality in real-world networks. In the Erdos ($N = 4,158$; left) and 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 the average of 100 runs, with 100 time steps and update probability equal to 0.1.

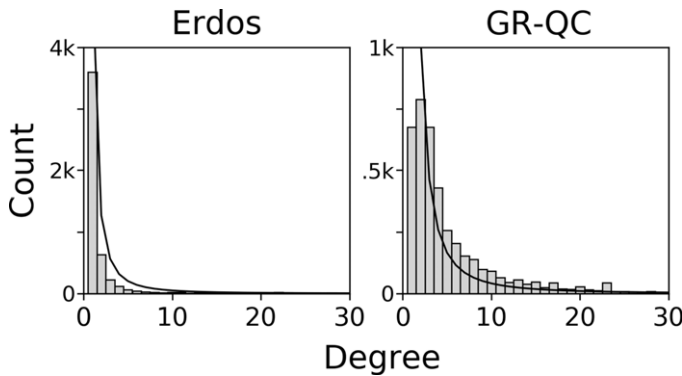


Figure 6. Degree distribution in model and two real-world networks. *Left:* The degree distribution given by $P(d) = N \cdot d^{-\gamma}$ with $\gamma = 2$ (solid line) approximates the degree distribution in the Erdos collaboration network ($N = 4,158$). *Right:* The same expression approximates the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

collaboration networks often range between values of 1 and 3, depending on the data set 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 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 the model is the product of a realistic network structure.

5. Discussion

My model shows that the structure of collaboration networks can give rise to inequality even in the absence of social categories. In particular, the model shows that

inequality in the payoff distribution and heterogeneity in the strategy profile of the population arise and persist in collaboration networks with a heterogeneous degree distribution. The 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, the 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 in which every agent interacts with four neighbors and there is no variation in the degree distribution. When the population structure leads many to interact with few and few to interact with many, my model shows that the resulting heterogeneous degree distribution can promote unequal outcomes.

My 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-good games, network heterogeneity induces diversity in group size and thus promotes contributions to the public good (Santos et al. 2006; Santos et al. 2008). But network heterogeneity can also lead to unequal outcomes in public-good games because the proliferation of altruistic behaviors ends up harming some individuals (McAvoy et al. 2020).

My 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*, whereas their well-connected collaborators play *High*. The other regime is when L is high. In this case, well-connected scientists play *Low*, whereas 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 coevolving groups of different sizes (Bruner 2019; O'Connor 2019, 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 my model is not the same as in the Red King/Queen effect. First, the Red King/Queen effect depends on the minority adapting more quickly to the strategy of the majority. In contrast, the update rule used in my model 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. Thus, it is not a difference in evolutionary tempo that drives inequality in my model. Second, the Red King/Queen effect relies on there being two groups, groups having different sizes, and individuals conditionalizing their behavior on the group membership of others. In my 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




Figure 7. Initial probabilities that *Low* and *High* are 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 $L = 0.4$.

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 my 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 my model is not the same as that in the Red King/Queen effect.

So what explains the two regimes of inequality that we observe in my model? Because the update rule my model uses is the myopic best response, to answer this question, I follow Rubin and O'Connor's (2018, 386–88) 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 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 $1/3$.

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.

Notice that the probability that a strategy is a best response depends on the degree. As already noted, each strategy is a best response with probability $1/3$ when

⁴ I thank an anonymous referee for raising this point.

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 approximately 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 approximately 0.55. But the probability that *Low* is a best response for an agent who interacts with five other agents is approximately 0.85.

This allows us to gain some insight into the two regimes for the emergence of inequality in the 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, whereas poorly connected agents play *High* and end up at an advantage. The two regimes of inequality we observe in scale-free networks are 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 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 the 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 the payoff inequality is high. My model therefore raises the specter of a twofold harm: low values of L put poorly connected scientists at a disadvantage when doing so is particularly harmful.

The twofold 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

⁵ The initial probabilities that either *Low* or *High* is a best response are higher when $L = 0.4$ than when $L = 0.1$ for $d \geq 2$. This helps explain why a smaller share of the population will 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.

collaborators than their male colleagues (Araujo et al. 2017; Abramo, D'Angelo, and Caprasecca 2009). Black scientists also have fewer collaborators, at least in some disciplines (Del Carmen and Bing 2000). When the payoff inequality is especially high, the twofold 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 tend to put those groups that are already vulnerable as a result of the position that they occupy in the collaboration network at a further disadvantage. 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 the 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 the 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, I 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. I 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 my model makes several simplifying assumptions. First, it assumes that scientists play the same strategy with all their collaborators. This is unlikely to hold in reality because scientists often negotiate different arrangements with different collaborators. Second, it considers a dynamic population of scientists who change their strategies over time, but it assumes 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, it assumes that all scientists are equally competent. This is again unrealistic because scientists often differ with respect to how productive they are. Fourth, it assumes that scientists update their strategies by the method of myopic best response. This is a reasonable assumption, but update rules based on imitation are also plausible. Although 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 whom they interact with over time, differ with respect to how productive they are, or update their strategies according to different rules.

Acknowledgments. I thank MindCORE and the Social and Cultural Evolution Working group at the University of Pennsylvania for generous funding and a stimulating work environment. I also thank Erol Akçay for helpful discussions, two anonymous referees for their extremely constructive criticism,

as well as Hannah Read, Justin Bruner, Hannah Rubin, Cailin O'Connor, Dan Singer, and especially Alex McAvoy for extensive feedback on earlier drafts.

References

- Abelson, Jonathan S., Natalie Z. Wong, Matthew Symer, Gregory Eckenrode, Anthony Watkins, and Heather L. Yeo. 2018. "Racial and Ethnic Disparities in Promotion and Retention of Academic Surgeons." *American Journal of Surgery* 216 (4):678–82.
- Abramo, Giovanni, Ciriaco D'Angelo, and Alessandro Caprasecca. 2009. "Gender differences in research productivity: A bibliometric analysis of the Italian academic system." *Scientometrics* 79 (3):517–539.
- Albert, Réka, and Albert-László 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 Brian Skyrms. 1999. "Bargaining with Neighbors: Is Justice Contagious?" *Journal of Philosophy* 96 (11):588–98.
- Allison, Paul D., J. Scott Long, and Tad K. Krauze. 1982. "Cumulative Advantage and Inequality in Science." *American Sociological Review* 47 (5):615–25.
- Allison, Paul D., and John A. Stewart. 1974. "Productivity Differences among Scientists: Evidence for Accumulative Advantage." *American Sociological Review* 39 (4):596–606.
- Araujo, Eduardo B., Nuno A. M., Araújo, André A. Moreira, Hans J. Herrmann, and José S. Andrade, Jr. 2017. "Gender Differences in Scientific Collaborations: Women Are More Egalitarian Than Men." *PLoS One* 12 (5):e0176791.
- Axtell, Robert L., Joshua M. Epstein, and H. Peyton Young. 2001. "The Emergence of Classes in a Multiagent Bargaining Model." *Social Dynamics* 27:191–211.
- Barabási, Albert-László, and Réka Albert. 1999. "Emergence of Scaling in Random Networks." *Science* 286 (5439):509–12.
- Barabási, Albert-László, Hawoong Jeong, Zoltan Nédá, Erzsebet Ravasz, Andras Schubert, and Tamas Vicsek. 2002. "Evolution of the Social Network of Scientific Collaborations." *Physica A: Statistical Mechanics and Its Applications* 311 (3–4):590–614.
- Barabási, Albert-László, and Zoltan N. Oltvai. 2004. "Network Biology: Understanding the Cell's Functional Organization." *Nature Reviews Genetics* 5 (2):101–13.
- Batagelj, Vladimir, and Andrej Mrvar. 2000. "Some Analyses of Erdos Collaboration Graph." *Social Networks* 22 (2):173–86.
- Beaver, Donald deBlasiis. 2004. "Does Collaborative Research Have Greater Epistemic Authority?" *Scientometrics* 60 (3):399–408.
- Bergstrom, Carl T., and Michael Lachmann. 2003. "The Red King Effect: When the Slowest Runner Wins the Coevolutionary Race." *Proceedings of the National Academy of Sciences* 100 (2):593–98.
- Binmore, Kenneth George. 1998. *Game Theory and the Social Contract: Just Playing*, vol. 2. Cambridge, MA: MIT Press.
- Bol, Thijs, Mathijs de Vaan, and Arnout van de Rijt. 2018. "The Matthew Effect in Science Funding." *Proceedings of the National Academy of Sciences* 115 (19):4887–90.
- Bruner, Justin P. 2019. "Minority (Dis)Advantage in Population Games." *Synthese* 196 (1):413–27.
- Clauset, Aaron, Samuel Arbesman, and Daniel B. Larremore. 2015. "Systematic Inequality and Hierarchy in Faculty Hiring Networks." *Science Advances* 1 (1):e1400005.
- Del Carmen, Alejandro, and Robert L. Bing. 2000. "Academic Productivity of African Americans in Criminology and Criminal Justice." *Journal of Criminal Justice Education* 11 (2):237–49.
- Du, Wen-Bo, Hao-Ran Zheng, and Mao-Bin Hu. 2008. "Evolutionary Prisoner's Dilemma Game on Weighted Scale-Free Networks." *Physica A: Statistical Mechanics and Its Applications* 387 (14):3796–800.
- Feldon, David F., James Peugh, Michelle A. Maher, Josipa Roksa, and Colby 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, Miranda. 2007. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.

- Gabbidon, Shaun L., Helen Taylor Greene, and Kideste 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, Piotr P. 2007. "Gini Coefficient: A New Way to Express Selectivity of Kinase Inhibitors against a Family of Kinases." *Journal of Medicinal Chemistry* 50 (23):5773–79.
- Han, Shin-Kap. 2003. "Tribal Regimes in Academia: A Comparative Analysis of Market Structure across Disciplines." *Social Networks* 25 (3):251–80.
- Henriksen, Dorte. 2016. "The Rise in Co-Authorship in the Social Sciences (1980–2013)." *Scientometrics* 107 (2):455–76.
- Holman, Bennett, and Justin Bruner. 2017. "Experimentation by Industrial Selection." *Philosophy of Science* 84 (5):1008–19.
- Hopkins, Allison L., James W. Jawitz, Christopher McCarty, Alex Goldman, and Nandita 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):515–34.
- Kitcher, Philip. 1990. "The Division of Cognitive Labor." *Journal of Philosophy* 87 (1):5–22.
- Kummerfeld, E., and K. J. Zollman. 2015. "Conservatism and the Scientific State of Nature." *British Journal for the Philosophy of Science* 67 (4):1057–76.
- Langel, Matti, and Yves 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):521–40.
- Larivière, Vincent, Chaoqun Ni, Yves Gingras, Blaise Cronin, and Cassidy R. Sugimoto. 2013. "Bibliometrics: Global Gender Disparities in Science." *Nature News* 504 (7479):211.
- Lee, Sooho, and Barry Bozeman. 2005. "The Impact of Research Collaboration on Scientific Productivity." *Social Studies of Science* 35 (5):673–702.
- Leskovec, Jure, Jon Kleinberg, and Christos Faloutsos. 2007. "Graph Evolution: Densification and Shrinking Diameters." *ACM Transactions on Knowledge Discovery from Data* 1 (1):article 2.
- Longino, Helen E. 1990. *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton: Princeton University Press.
- Lusseau, David. 2003. "The Emergent Properties of a Dolphin Social Network." *Proceedings of the Royal Society of London. Series B: Biological Sciences* 270 (suppl. 2):S186–88.
- Martini, Carlo, and Manuela Fernández Pinto. 2017. "Modeling the Social Organization of Science." *European Journal for Philosophy of Science* 7 (2):221–38.
- McAvoy, Alex, Benjamin Allen, and Martin A. Nowak. 2020. "Social Goods Dilemmas in Heterogeneous Societies." *Nature Human Behaviour* 4 (8):819–31.
- Melin, Göran, and Olle Persson. 1996. "Studying Research Collaboration Using Co-Authorships." *Scientometrics* 36 (3):363–77.
- Merton, Robert K. 1968. "The Matthew Effect in Science: The Reward and Communication Systems of Science Are Considered." *Science* 159 (3810):56–63.
- Mohseni, Aydin, Cailin O'Connor, and Hannah Rubin. 2019. "On the Emergence of Minority Disadvantage: Testing the Cultural Red King Hypothesis." *Synthese* 198 (6):5599–621.
- Nash, John F. Jr. 1950. "The Bargaining Problem." *Econometrica* 18 (2):155–62.
- Newman, Mark E. J. 2001. "The Structure of Scientific Collaboration Networks." *Proceedings of the National Academy of Sciences* 98 (2):404–9.
- Newman, Mark E. J. 2004. "Coauthorship Networks and Patterns of Scientific Collaboration." *Proceedings of the National Academy of Sciences* 101 (suppl. 1):5200–5205.
- Nielsen, Mathias Wullum, and Jens Peter Andersen. 2021. "Global Citation Inequality Is on the Rise." *Proceedings of the National Academy of Sciences* 118 (7):1–10.
- O'Connor, Cailin. 2017. "The Cultural Red King Effect." *Journal of Mathematical Sociology* 41 (3):155–71.
- O'Connor, Cailin. 2019. *The Origins of Unfairness: Social Categories and Cultural Evolution*. Oxford: Oxford University Press.
- O'Connor, Cailin, Liam Kofi Bright, and Justin P. Bruner. 2019. "The Emergence of Intersectional Disadvantage." *Social Epistemology* 33 (1):23–41.
- O'Connor, Cailin, and Justin Bruner. 2019. "Dynamics and Diversity in Epistemic Communities." *Erkenntnis* 84 (1):101–19.

- Petersen, Alexander Michael, Santo Fortunato, Raj K. Pan, Kimmo Kaski, Orion Penner, Armando Rungi, Massimo Riccaboni, H. Eugene Stanley, and Fabio Pammolli. 2014. "Reputation and Impact in Academic Careers." *Proceedings of the National Academy of Sciences* 111 (43):15316–21.
- Rosenstock, Sarita, Justin Bruner, and Cailin O'Connor. 2017. "In Epistemic Networks, Is Less Really More?" *Philosophy of Science* 84 (2):234–52.
- Rubin, Hannah, and Cailin O'Connor. 2018. "Discrimination and Collaboration in Science." *Philosophy of Science* 85 (3):380–402.
- Santos, Francisco C., Jorge M. Pacheco, and Tom Lenaerts. 2006. "Evolutionary Dynamics of Social Dilemmas in Structured Heterogeneous Populations." *Proceedings of the National Academy of Sciences* 103 (9):3490–94.
- Santos, Francisco C., Marta D. Santos, and Jorge M. Pacheco. 2008. "Social Diversity Promotes the Emergence of Cooperation in Public Goods Games." *Nature* 454 (7201):213–16.
- Skyrms, Brian. 1996. *Evolution of the Social Contract*. Cambridge: Cambridge University Press.
- Skyrms, Brian, and Kevin J. S. Zollman. 2010. "Evolutionary Considerations in the Framing of Social Norms." *Politics, Philosophy & Economics* 9 (3):265–73.
- Weatherall, James Owen, Cailin O'Connor, and Justin P. Bruner. 2020. "How to Beat Science and Influence People: Policymakers and Propaganda in Epistemic Networks." *British Journal for the Philosophy of Science* 71 (4):1157–86.
- Weisberg, Michael, and Ryan Muldoon. 2009. "Epistemic Landscapes and the Division of Cognitive Labor." *Philosophy of Science* 76 (2):225–52.
- West, Jevin D., Jennifer Jacquet, Molly M. King, Shelley J. Correll, and Carl T. Bergstrom. 2013. "The Role of Gender in Scholarly Authorship." *PloS One* 8(7):e66212.
- Wittebolle, Lieven, Massimo Marzorati, Lieven Clement, Annalisa Balloi, Daniele Daffonchio, Kim Heylen, Paul De Vos, Willy Verstraete, and Nico Boon. 2009. "Initial Community Evenness Favours Functionality under Selective Stress." *Nature* 458 (7238):623–26.
- Wittman, Holly O., Michael Hendricks, Sharon Straus, and Cara 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):531–40.
- Wray, K. Brad. 2002. "The Epistemic Significance of Collaborative Research." *Philosophy of Science* 69 (1):150–68.
- Zollman, Kevin J. S. 2007. "The Communication Structure of Epistemic Communities." *Philosophy of Science* 74 (5):574–87.
- Zollman, Kevin J. S. 2010. "The Epistemic Benefit of Transient Diversity." *Erkenntnis* 72 (1):17.

Appendix

I 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
  
```

Cite this article: Ventura, Rafael. 2023. "Structural Inequality in Collaboration Networks." *Philosophy of Science* 90 (2):336–353. <https://doi.org/10.1017/psa.2022.73>