Discrimination and Collaboration in Science

Hannah Rubin and Cailin O'Connor[*][†]

Abstract

We use game theoretic models to take an in-depth look at the dynamics of discrimination and academic collaboration. We find that in collaboration networks, small minority groups may be more likely to end up being discriminated against while collaborating. We also find that discrimination can lead members of different social groups to mostly collaborate with in-group members, decreasing the effective diversity of the social network. Drawing on previous work, we discuss how decreases in the diversity of scientific collaborations might negatively impact the progress of epistemic communities.

1 Introduction

Philosophers of science have used formal models to argue that the structure of communication and collaboration networks matter in science.[1] One finding from this literature is that diversity of beliefs within an epistemic community is key to ensuring that the group eventually arrives at true beliefs about the world, and that network structure can be crucial to preserving this diversity (Zollman, 2010). At the same time, feminist philosophers of science have argued that personal diversity, i.e., diversity with respect to personal identity markers such as gender, race, and cultural origin, is an important source of such epistemic diversity.[2]

[1]See, for example, Zollman (2007, 2010); Mayo-Wilson et al. (2013); Holman and Bruner (2015); Grim et al. (2015); Rosenstock et al. (2017).

[2]This distinction is similar to that drawn by Fehr (2011) between situational and epistemic diversity. For examples of compelling arguments by philosophers of science for the importance of diversity in science

Given this work, we ask: what factors influence the diversity of epistemic collaboration networks? And, in particular: might discrimination impact these networks? We start by looking at the emergence of discriminatory norms in fixed collaboration networks. Such norms commonly evolve, and in particular we find support for previous work showing that minority status alone can make it more likely for a social group to be disadvantaged by bargaining norms (Bruner, 2017; O'Connor, 2017a). Next, we explore the endogenous emergence of collaboration networks in a population that already has discriminatory norms, finding that such networks tend to become segregated to the point where there are no collaborative interactions across groups. Lastly, we examine the simultaneous co-evolution of discrimination and collaboration, where we see partially segregated networks evolve with some actors upholding the discriminatory norm. Overall these results suggest that discrimination in academia may decrease the personal diversity of collaborative networks. As described above, this may have negative impacts on the ability of epistemic communities to arrive at successful beliefs.

The paper will proceed as follows. In section 2 we describe the Nash demand game, which will be the base model employed here to capture discrimination in academic collaboration. We will justify the use of the model for epistemic communities in particular. In section 3 we present our main body of results. We conclude by discussing the relevance of these results to epistemic communities and to epistemic progress.

2   Academic Bargaining and Discrimination

As stated, our aim is to analyze the dynamics that surround discrimination and collaboration networks in academia. This analysis is inspired, in part, by two sets of empirical results. The first suggests that in epistemic communities, women may get less credit than men for joint work. West et al. (2013) and Sugimoto et al. (2013), for example, find that women are less likely in many disciplines to hold prestigious author positions. Feldon et al. (2017) found that among early PhD students in the biological sciences, women put in significantly more work, but are significantly less likely to be granted authorship on the papers produced. Another set of results suggests that women are less likely to collaborate than men are, and are more likely to collaborate with other women (Ferber and Teiman, 1980; McDowell and Smith, 1992; Boschini and Sjögren, 2007; West et al., 2013). Some findings suggest a similar pattern for black academics, with black criminologists less likely to co-author (Del Carmen and Bing, 2000). Botts et al. (2014) also find that black philosophers tend to cluster in subfields.

Part of our question is: are these sets of results related? Does inequity in academic collaboration lead members of certain groups to self segregate and thus decrease the effec-

---

see Haraway (1989); Longino (1990); Harding (1991). In addition, Page (2008) argues extensively that cognitive diversity is important to creative work and problem solving, and may be generated by some sorts of personal diversity. A large body of research from across the social sciences finds epistemic benefits of personal diversity ranging from financial gains in firms where a significant portion of the leadership are women and members of racial minorities (Dezsö and Ross, 2012; Richard et al., 2007), to improved problem solving in small, racially diverse groups (Phillips et al., 2006), to increased information sharing in racially diverse juries (Sommers, 2006).

		Low	Med	High
	Low	L,L	L,5	L,H
Player 1	Med	5,L	5,5	0,0
	High	H,L	0,0	0,0

Figure 1: Payoff table for a mini Nash demand game.

tive diversity of collaborative teams? It is notoriously difficult to generate empirical data testing cultural evolutionary pathways. To explore these questions, we instead employ game theoretic models. Such models start with a game, or a simplified representation of a strategic interaction. To represent division of labor and credit between academic collaborators we use the *Nash demand game* (Nash, 1950).

This game involves two agents who divide a resource by each demanding some portion of it.[3] If the demands are compatible, each agent gets what she requested. If the demands exceed the total resource, the agents get poor payoffs on the assumption that they cannot peaceably agree on a division. Figure 2 shows a *payoff table* for a 'mini' version of this game where actors have three demands—Low, Med, and High.[4] For simplicity sake, we assume that the total resource is 10, the Med demand is 5, $L < 5 < H$, and $L + H = 10$. This yields, for example, demands of 3, 5, and 7, or 1, 5, and 9. Strategies for player 1 are displayed in the rows of the table, and strategies for player 2 in the columns. Each entry shows payoffs to each player for some combination of demands, with player 1 listed first. The poor payoff when actors over-demand the resource is assumed to be 0.[5]

It will be useful to take a minute to explain why this is a good representation of academic collaboration. Academic collaboration involves joint action which creates a surplus of a credit compared to solo work.[6] However, this joint action necessitates two types of bargaining. First, actors must decide who will do how much work on the project. Second, actors must determine author order as a proxy for credit. The demands in the game, then, are best understood as requests for author position relative to the amount of work done. An actor who does the lion's share of the work and requests first authorship

---

[3]Note that, since we only discuss two player games, our results may be difficult to apply to fields that tend to have larger collaborations.

[4]Mini-games are commonly employed in evolutionary models of bargaining for tractability purposes. See, for example, Skyrms (1994); Alexander and Skyrms (1999). See Sigmund et al. (2001) for a defense of the mini-game approach.

[5]The exact value of these outcomes should not be taken too seriously as these details will not matter much for the outcomes of the model. For instance, all bargaining failures (the project never getting started, failing halfway through, etc.) are given a payoff of zero in this model. Realistically, they should likely be assigned different payoffs, but as long as the payoff for each type of failure is substantially lower than for successful collaboration, this idealization will not significantly affect results.

[6]Collaboration increases academic productivity, and collaborative papers are more likely to be accepted to top journals and to be cited. See Bruner and O'Connor (2018) for an overview of the literature on this topic.

makes a Med demand. One who does more work, and requests second authorship makes a Low demand.[7]

Suppose we have a population with two social groups—women and men, for example, or black and white people. Suppose further that actors can condition their choice of strategy on the group membership of an interactive partner. In a cultural evolutionary scenario, this induces a situation where separate norms can emerge within and between groups.[8] For the Nash demand game, under most reasonable evolutionary dynamics, in-group members will most often evolve to all make fair demands of each other.[9] One of three things will happen between groups. Either the groups will come to demand Med of each other, or else one group will learn to always demand High and the other to always demand Low when meeting out-group members. Axtell et al. (2001) take these two latter sorts of outcomes to represent 'discriminatory norms'—actors treat in- and out-group members differently, to the detriment of one out-group. We follow them in using these outcomes as representations of discriminatory norms of collaboration.

There is one further preliminary we will introduce before describing our models and results. Bruner (2017) outlines the cultural Red King effect—that minority groups can be more likely than majority groups to end up disadvantaged in the emergence of bargaining conventions. This occurs as a result of differential reactivity on the part of the groups. While majority members rarely meet minority types, the reverse is not true, meaning that minority members learn much more quickly how to interact with the majority type. Since it is often best to make safe, low demands in the Nash demand game, these minority types can then end up quickly learning a behavior that ultimately disadvantages them.[10]

In these minority-majority models, it is not always the case that small groups end up disadvantaged as a result of their quick learning. As Bruner (2017) illustrates, under the right conditions bargaining payoffs are such that making high demands is worthwhile despite the risk, meaning that minority groups quickly learn to do so. This complementary effect is the cultural Red Queen. For the mini games just introduced, a cultural Red King is observed when $L > 3$, and a Red Queen otherwise. Below we will draw a connection between our results on the impact of minority status on norms of collaboration and this

---

[7]There are some other, plausible frameworks we might have used. For example, models of the cultural emergence of inequity sometimes focus on a hawk-dove game, or a version of a coordination game with one preferred equilibrium for each player. (See O'Connor (2017b) for an overview and in depth discussion of such models.) We, however, prefer a paradigm where an equitable outcome is available to the players, which is not the case for these games. There are also more complicated versions of bargaining games, where, for instance, actors bargain over many rounds in order to determine outcomes. We choose this simplified version for tractability reasons. From an evolutionary perspective, it also makes more sense to compress the actual process of bargaining and focus on the effects final bargaining outcomes have on dynamics of cultural change.

[8]We follow authors like Young (1993) in labeling emergent patterns of group level behavior in models as 'norms', though this is obviously a thin representation of real world norms.

[9]The other option is to evolve a 'fractious' pattern of bargaining where some actors make High demands and others Low, meaning that miscoordination happens with relative frequency (Axtell et al., 2001; Skyrms, 2014).

[10]O'Connor (2017a) replicates these results in agent based models. O'Connor and Bruner (2017) discuss implications for academic communities in particular. And O'Connor et al. (2017) look at the implications for intersectional populations.

body of literature.

3   Networks and Bargaining

Now we use the framework sketched in the last section to build an explicit model of academic collaboration networks. In our model, agents in the academic community are represented by a collection of *nodes*. The presence of an *edge*, or *link*, between two nodes means that a collaboration exists between the two individuals, whereas the absence of an edge means that they do not collaborate. There are *within* group links, connecting two nodes in the same social identity group, and *between* group links, connecting two nodes in different social identity groups. The set of nodes and edges forms what we call the *collaboration network*.

In what follows, actors belong to one of two social identity groups, where one social identity group comprises a minority of the total population. We will be particularly interested in results where the majority discriminates by demanding High when interacting with minority group members who in turn demand Low. We will also take note of the norm where the minority discriminates, demanding High against majority group members, and the norm of fair division, where both groups demand Med against each other. Our models assume that actors condition their behavior on the social identity of their collaborative partners, so our results will only be relevant to understanding discrimination and lack of diversity when the groups of interest are divided according to a feature real people would actually condition their behavior on, e.g. race or gender but not eye color.[11]

Our exploration of discriminatory norms and collaboration will proceed in three parts. First, we show that when agents are on a network, the minority group can be disadvantaged solely due their relative proportion in the population. Second, we will show that when there are preexisting discriminatory norms in the community, networks tend to become completely segregated. Third, we show how these two parts of the story relate to each other by providing a model where agents' bargaining strategies co-evolve with the structure of the collaboration network. We will see that discriminatory behavior tends to arise between many members of the community. Further, as this behavior arises, the collaboration networks tend to become partially segregated, with agents mostly interacting with others in their own social identity group.

3.1   Part a: the evolution of discrimination on fixed networks

First, we examine the effects of network structure on the evolution of discriminatory norms. Poza et al. (2010) use a similar framework to show that discriminatory norms do commonly arise on networks with agents of different types playing the Nash demand game. In looking at these models with minority/majority statuses for the two groups,

---

[11]Similarly, our epistemic reasons for caring about diversity will only be relevant when there is reason to think that diversity across a particular social identity line produces epistemic benefits, though, of course, one can also care about diversity for non-epistemic reasons.

we find that the minority group can be disadvantaged, in that they are more likely to be discriminated against. Further, we investigate whether *homophily*, the tendency to preferentially form links with members of your own social identity group (Lazarsfeld et al., 1954; Jackson, 2010), exacerbates the effect.

3.1.1 Model

We use multi-type random graphs, networks which are used to model populations with multiple social identity groups (Golub and Jackson, 2012). In this set-up, agents are classified according to which type they are (in this case, whether they are a minority or majority group member). Each agent has some probability of forming a link with an agent of the same type, p_{in}, and some probability of forming a link with an agent of a different type, p_{out}. If $p_{in} > p_{out}$, we say that the agents exhibit homophily.

Once the network is formed, agents update their strategies based on the payoffs they receive by interacting with their collaborators—those they are connected to on the collaboration network. Each agent's strategy consists of two parts: a demand when interacting with an in-group member and a demand when interacting with an out-group member. These strategies are initially randomly assigned. Each round, agents interact with all of their collaborators and, with a small probability, will decide to update their strategy. Strategies are updated using myopic best response: in the next round, the strategy an agent will use is the one that would have gotten them the best payoff in the current round, given the strategies of their collaborators. This captures the fact that agents are trying to choose a strategy that is likely to result in them getting the most out of a successful collaboration, while avoiding the poor payoff from a failed collaboration.

3.1.2 Results

We look at the frequencies with which populations converged to different bargaining norms. Cases where 2 or fewer agents were playing strategies outside the equilibrium expectation were counted as converged since, based on the probabilistic nature of the model, these agents may not have had a chance to update their strategies. The probability of updating a strategy in any given round was .1. Simulations were run for 1,000 rounds over networks ranging from 20 to 100 agents (in intervals of 20), where the high demand (H) ranged from 6 to 9 and the minority group comprised 10% to 50% of the population. While for all simulations the probability of a within group link was held fixed at $p_{in} = .4$, the probability of an out-group link p_{out} ranged from .2 to .8. (That is, we look at cases where the minority was twice as likely to collaborate with in-group members to cases where the minority is twice as likely to collaborate with out-group members.) Each combination of parameter values was run 100 times.

Within each group, populations nearly always evolved to the norm of equal division. Between groups, populations most often evolved to the norm of fair division, but a significant amount of the time they also evolved a discriminatory norm, as in Poza et al. (2010).

Figure 2: Convergence to different norms over size of minority group for $H = 6$.

We now look at whether minority status has any effect on likelihood of being discriminated against. Figure 2 shows results for high demand, $H = 6$. For a small minority group, it is more likely that the majority group will end up demanding High against the minority. As the size of the minority group increases, the fair division becomes more likely and both groups become equally likely to discriminate.

This effect results from an asymmetry in the average number of links for each group. Consider a simple demonstrative example. If there are 10 majority group members, 5 minority group members, and 10 total between group links, on average majority group members will have one between group link and minority group members will have two. If a minority group member has two links, both of these collaborators would have to be demanding 4 in order for their best response to be 6. At least initially, this would happen with only probability 1/9. On the other hand, only one of these collaborators would have to be demanding 6 in order for their best response to be demand 4 and, at least initially, this happens with probability 5/9.[12] In contrast, a majority group member having one link to the minority would have a 1/3 chance of each of the strategies being a best response (they demand the complement of whatever their collaborator demands). So, at the start, a minority group member is much more likely than a majority group member to demand 4 and much less likely to demand 6. This asymmetry drives populations toward

---

[12]To calculate the best response you compare the payoffs of each of possible strategies. So, for instance, when an agent has two collaborators, the first of which demands 5 and a second that demands 6, demanding 4 yields payoff of $4 + 4 = 8$, demanding 5 yields a payoff of $5 + 0 = 5$, and demanding 6 yields a payoff of $0 + 0 = 0$. Therefore, demanding 4 is a best response as it yields the highest payoff. One can calculate the best response for the nine possible configurations of strategies among the collaborations (Low and Low, Low and Med, Med and Low, etc.) to figure out the likelihood of the minority player's strategies being best responses. Note that adding a small probability of error when agents update their strategies does not significantly affect our results.

the outcome where the majority demands 6.

Our findings are somewhat similar to previous results on the cultural Red King/Queen effect described above. However, minority disadvantage does not arise here due to an increased speed of learning in the minority group, but rather due to differences in the best responses of the group. Furthermore, while we observe minority disadvantage which is analogous to the cultural Red King effect when $H = 6$, we do not observe an analogue of the Red Queen for higher levels of H.

To understand why nothing similar to the Red Queen occurs consider the same simplified example as above, but now with $H = 9$. Again, a majority group member having one link to the minority would have a $1/3$ chance of each of the strategies being a best response, as they still only care about demanding the complement their collaborator. Now, though, the minority group members are very unlikely to make the Low demand of 1; it is only a best response with probability $1/9$, when both of their collaborators demand 9. They are just as likely as the majority to demand 9 (doing so with probability $1/3$, whenever both their collaborators demand 1 or one collaborator demands 1 while the other demands 9). More importantly, though, it is most likely that the minority group member's best response will be to demand 5. This will be the best response whenever at least one of their collaborators demands 5, which happens with probability $5/9$. So, the evolution of strategies is again affected by the minority group's initial best responses, but in this case it is likely that the frequency of the Med demand in minority will increase, followed by the majority learning to also make Med demands.

We also found that varying levels of homophily did not influence whether the majority or minority group is more likely to discriminate.[13] Since agents are updating their in-group and out-group strategies separately, how the within group linking probability compares to the out-group linking probability does not have an effect on how often the network converges to a discriminatory norm. So, the existence of majority group advantage does not depend on the existence of homophily.

3.2 Part b: existing discriminatory norms affect network formation

We now examine how the collaboration network will evolve when there is already a discriminatory norm in place. An agent's bargaining strategy is held fixed, while their choice of who to collaborate with evolves over time. O'Connor and Bruner (2017) find that in epistemic communities discriminatory norms can disincentivize collaboration between social groups. They look at a model where there is an option to work alone and actors are not arrayed on a network, and find that minorities will learn to do solo work more often when discriminated against. As will become clear, our results complement this finding.

---

[13]Increasing p_{out} did increase the probability that the collaboration network as a whole converged to one of the possible norms. This is likely because with lower linking probabilities there are often some nodes which do not have any links to the rest of the network (i.e., the network is not connected).

3.2.1 Model

We employ a model similar to Watts (2003) in which agents can choose to form or break links with other agents in the community based on their payoffs from bargaining with those other agents. Each link represents a collaboration and therefore a payoff from the Nash Demand game. Since we are investigating the effect of a pre-existing discriminatory norm, majority group members receive a payoff of 6 from a between group link while minority members receive a payoff of 4.[14] Agents receive a payoff of 5 from within group links, where there is a norm of fair division. Agents have a maximum number of links, capturing the fact that there are a limited number of projects academics can work on. A player can unilaterally sever a link, but both players must consent to a new link being formed. This represents the fact that all the researchers involved in a collaboration must consent to be part of the collaboration. If one person no longer wishes to collaborate, the collaboration fails and the link is broken.[15]

The evolution of the collaboration network proceeds as follows. We begin with an empty network (there are no links between any nodes). At each time-step, two nodes are chosen at random. One of these is an agent who will update their links and the other is a potential or current collaborator of the agent.

If we have chosen a potential collaborator, we determine whether both researchers will consent to form a new link. If neither player has reached the maximum number of links, they will both consent. (Getting some payoff is better than none.) If either, or both, of them already have the maximum number of links, we check whether they will break any of their existing links to form the new link. If a player can increase their payoff by breaking their link with the lowest payoff in order to form this new link, they will consent to the new link. If both players consent to the link forming, this new link will form, and agents will break links to the collaborator with which they receive their lowest payoff.[16] If at least one of them does not consent, no links are formed or broken.

By contrast, if we have chosen a current collaborator, the agent has an option to break the link and form a new one. A potential collaborator is chosen at random from the community. Then, if the agent would get a higher payoff from linking to this new potential collaborator, they try to form a link. If the potential collaborator would also like to form a link with the agent, the link is formed. The agent breaks the link with their old collaborator (and the agent's new collaborator breaks their link with their lowest payoff collaborator if they already had the maximum number of links).

---

[14] The particular values for the high and low demand do not affect the results.

[15] Rather than limiting the number of links, we might have used a model similar to Jackson and Wolinsky (1996), where agents choose how to divide a fixed amount of time among different collaborations (and there is some synergistic benefit from both collaborators spending more time on a collaboration). However, this would add an additional factor into the model, the number of collaborations different groups maintain, so in order to focus the analysis on discrimination and homophily, we use the model of collaboration from Watts (2003).

[16] In the event that an agent has multiple collaborations yielding the same lowest payoff, the link that is broken is chosen among them at random.

Figure 3: An example of the evolution of a network with discriminatory norms, with 100 researchers, the minority group comprising 40% of the population, $H = 6$, and the maximum number of links set to 3.

3.2.2 Results

Across a wide range of parameters, the network reliably converges to the point where researchers only collaborate within their own group.[17] Figure 3 provides an example of how this occurs. Initially, links form steadily within each group and also between groups, as researchers have not yet formed the maximum number of links. Once minority group members have reached the maximum number of links, they begin to break their links with majority group members whenever they have the opportunity to form a new link with another minority group member, which yields them a payoff of 5 rather than a payoff of 4. So, between group links decline and links within the minority group increase.

Simultaneously, links within the majority group increase. This is because, as minority group members break their links with majority group members, the majority group members look to form new links. Minority group members will refuse to form new links with them, but other majority group members (if they also have less than the maximum number of links) will agree to form the link. Note that in figure 3, fewer links exist within the minority than within the majority simply because there are fewer minority group members and so fewer possible links. The important trend is that links between

---

[17]We ran simulations for 10,000 rounds, varying the number of agents in the community (10, 20, 40, 60, 80, or 100), the percent of the community in the minority population (5, 10, 20, 30, 40, or 50), and the maximum number of links agents can form (3, 10, or 20). As long as there were enough minority members so that all the links could be formed within their group (e.g. if there were only 5 minority groups members, but 20 possible links, they could not form all 20 links to other minority group members), simulations show that the collaboration network reliably evolves to a point where at least 95% of the links are formed within social identity groups rather than across.

groups decrease over time until they are essentially non-existent.[18]

To this point, we have interpreted credit in the models here as resulting from purposeful division among collaborators, but it is worthwhile pointing out that academics do not fully control how credit is divvied out. There is evidence, in particular, that members of traditionally disadvantaged groups do not receive as much credit for collaborative work, even when author order is determined fairly. Sarsons et al. (2017) finds that in economics, where author order is alphabetical, coauthoring does not decrease men's chances of getting tenure, but women who coauthor are less likely to be tenured. This effect is not as strong when women coauthor with other women. The implication is that the effects we have been describing might occur even when collaborators attempt a fair division. If the payoff to minority types for between group collaboration is, in fact, lower nonetheless, they still might learn to avoid such collaborations.

3.2.3 Bigger Academics, Bigger Pies

In the model we have just presented, the communal norm determines credit share for all agents, and this predictably impacts the network structure. As we have pointed out, those who are being discriminated against will want to avoid this discrimination. One assumption of this model, though, is that agents of both sorts generate the same amount of credit, which may not apply in certain cases.

There are (at least) two plausible scenarios where members of one social group will produce more credit than those of another. The first is an interpretation of these models where the groups are stages of the academic hierarchy—graduate students, post docs, early professors, established professors, etc. In such a case, established academics have the experience and social standing to publish better papers in better journals, and to expect that their results will be more widely read and cited. The second plausible scenario of this sort is one where members from a traditionally excluded social group are first moving into a discipline. We can expect members from such small minority groups to have less access to the sorts of experience and prestige just described, and so to produce less credit on average.[19]

Importantly, such considerations can change the outcome of the models just discussed. Assume that the different types of collaborative pairings—in-group collaborations among established academics, between group collaborations, and in-group collaborations between less established academics— generate descending amounts of credit for the reasons just described. Perhaps the total credit generated in the three cases is 15, 12, and 10. Demands in the game are then for proportions of these resources, say .4, .5, and .6.

Varying these parameters, maintaining our assumption that discriminatory norms hold, and always assuming that those who produce more credit are the discriminators, we can yield three different sorts of outcomes. The first are those we have described in

---

[18]A few persist for long periods of time just by chance. If the simulation is run for long enough, all of these links disappear.

[19]Minorities might also produce less credit if the type of projects they work on become labeled as low prestige. Thanks to Liam Kofi Bright for this point. Schneider et al. (2017) use similar models to consider this possibility.

section 3.2.2. Second, there are parameters where members of both types are incentivized to engage in only out-group collaboration. Suppose the three pie sizes are 15, 14, and 10. The established group will get 7.5 for within group collaboration, and the less established will get 5. For between-group collaboration, the payoffs will be 8.4 and 5.6 respectively. In other words, the less established group is willing to do more work and get a smaller proportion of the credit in order to share in the larger pie. The more established group is willing to produce a smaller pie as long as they receive a larger portion of it. Such a situation seems to correspond well with what we see when professors and graduate students collaborate, for example.

There is a third outcome where we again see complete segregation, but for a different reason. In this case, the established group produces such an excess of credit that they are unwilling to collaborate with the out-group even when they can take advantage of out-group members. For instance, if our pie sizes are 18, 13, and 10, the in-group payoffs are then 9 and 5, while the between group payoffs are 7.8 and 5.2. Here the less established group would like to collaborate, even if they do more work and receive a smaller proportion of credit. But the established group is unwilling to do so.

3.2.4 Diversity and Credit

In the introduction we discussed the idea that diversity might be important to scientific progress. If so, we might expect that the credit generated by diverse collaborative groups will sometimes be greater than that generated by homogenous ones. Some empirical findings back this up. Freeman and Huang (2015), for example, show that cross-cultural collaborations generate more credit in the form of citations. Campbell et al. (2013) find the same thing for gender diverse collaborations in ecology. And Barjak and Robinson (2008) find that culturally diverse academic teams are more productive.

To capture such a possibility, we can alter the model so that while in-group collaborations generate 10 units of credit, between-group collaborations generate more. Whether or not homophily emerges will then depend on the level of discrimination and the size of the credit boost from diversity. For instance, if between-group collaborations generate a pie of size 12, and the minority group gets 40% of the credit generated, their expected payoffs are 5 with their in-group and 4.8 with the out-group, meaning we should still expect segregation despite the extra payoff diversity generates. If we increase this number to 14, the minority group now gets 5.6, and we expect individuals to form only between-group collaborations. If we suppose more serious discrimination has emerged, so that minority types get only 30% of the credit, say, it will take a between-group pie of 17 to prevent segregation. Still, for any level of discrimination if there is enough of a credit boost for diverse collaborations we should not expect homophily to emerge.

Schneider et al. (2017) use models like those presented here to discuss possibilities for promoting diverse collaborations, despite natural tendencies towards homophily. One such possibility is for funding bodies to give more grants to diverse collaborative teams, increasing their average credit generated. This yields a model like the one just described—between-group collaborations split a bigger pie than within-group collaborations. As they point out, this does decrease homophily in network formation models. There is a

downside, though. In the models we describe in section 3.2.2, while there is complete homophily, there is no active discrimination occurring because no one maintains a collaborative bond where they receive the Low demand. In models where between-group collaboration is incentivized, either by funding agencies, or the natural benefits of diversity, minority groups will be willing to accept inequitable arrangements with out-group members. These arrangements provide a net benefit over in-group collaboration, but provide a much more significant benefit for those in the majority. What we see in such a case is a tension between a social good (diverse collaborations that improve the progress of science) and equity between individuals.[20]

The general take-away here is that the collaboration network we expect to evolve will vary based on both the structural situation for academics (i.e., who is poised to create credit), and on the existing norms of bargaining. In particular, there are multiple, realistic configurations of these features that yield segregated co-authorship networks.

3.3 Part c: the co-evolution of bargaining with networks

So far, we have seen the effects of the collaboration network's structure on the evolution of discriminatory norms and, conversely, the effect of discriminatory norms on the evolution of the collaboration network's structure. In this section we explore what happens when when we allow both agents' strategies and their choices of collaborators to co-evolve.

3.3.1 Model

We start with an empty network, with each agent's strategy randomly determined. In each round, there is some small probability that each agent will take an action. There are two types of possible actions: updating your bargaining strategy and updating your set of collaborators. So, if an agent takes an action in the round, there is chance they will update their set of collaborators and a chance they will update their strategy (agents do not update both at once). Updating sets of collaborators is also done via breaking and reforming links, as described in section 3.2. Strategies are updated via best response, as described in section 3.1.

3.3.2 Results

We look at results for a network of 100 researchers where $H = 6$. We ran simulations for 20,000 rounds, varying the size of the minority group from 10% of the population to 50%, in intervals of 5%, and varying the maximum number of links (either 3 or 9). Each combination of parameter values was run 100 times. In each round, agents take actions with probability 0.1. If they act, there is a 20% chance they will update their set of collaborators and an 80% chance they will update their strategy. (The particular probabilities are not important; similar results can be obtained for a variety of values.)

---

[20]If there is a special benefit to collaborations from diversity, it might be the case that small minority groups are then in a special position in that there are not enough of them to go around. This might improve their bargaining power.

Figure 4: Level of majority discrimination, given by the percent of the majority group whose out-group strategy is to demand High, with different sizes of minority group.

As in section 3.1, within groups, the likely outcome was fair division, so we will focus on between group strategies. We are particularly interested in the likelihood of majority discrimination and the effect it has on the collaboration network. As figure 4 shows, the larger the minority group is, the less likely it is that majority group members will discriminate against them, just as in 3.1.

However, the patterns of discrimination in these models are different from those in previous models. Groups do not end up at consistent norms where all members make the same demand of the out-group. Instead we see outcomes where some discrimination and some fair treatment emerge between groups. Figure 5 shows the evolution of between group strategies for two different runs of the same simulation.[21] In (a) we see an outcome where more majority group members demand High and minority members Low. (b) shows a run where both sides tend to demand Med. Note that in both cases, as mentioned, the network does not fully settle into either norm. To understand why this is the case, let us look at the evolution of the collaboration network.

As in section 3.2; both within and between group links increase initially, and then between group links decrease. In particular, discriminatory between group links tend to be dropped, while fair ones remain. Once an agent has no more links to the other group, they stop updating their between group strategy. This explains why the network never fully settles on one norm or another. Discriminators who have lost their minority group links retain latent discriminatory strategies that are never updated.

Figure 6 shows the network structure of the two runs shown in figure 5. As is easy to see, this process leads to collaboration networks which are homophilic, but not fully segregated. In fact, there is a continuum of possible outcomes, ranging from everyone

---

[21]The minority group comprised 30% of the population, maximum links were set to 9, and the simulation was run for 10,000 rounds.
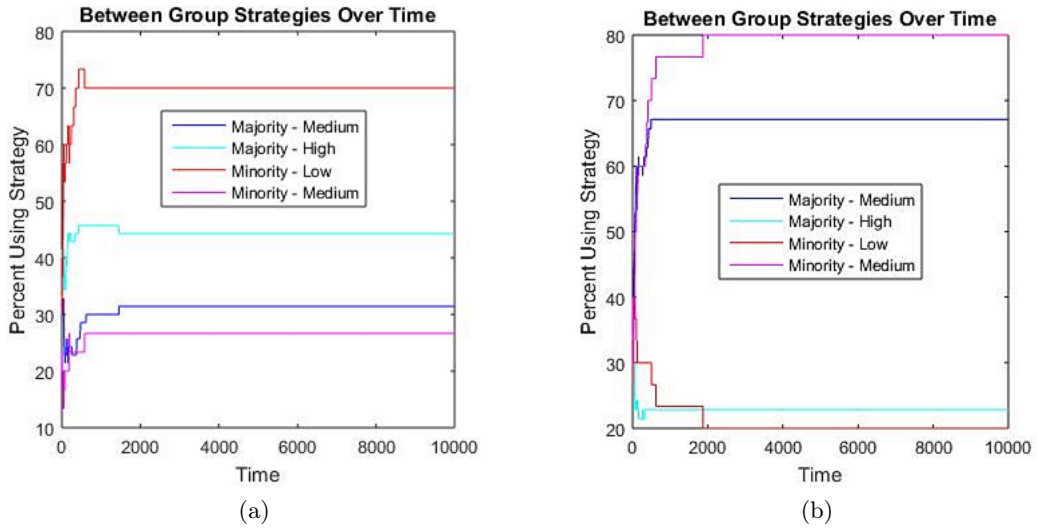
Figure 5: Evolution of between group strategies over time. Data for the majority demanding Low and the minority demanding High are omitted for simplicity. (a) shows an outcome where the high demand quickly spreads in the majority population. (b) shows an outcome where the medium demand quickly spreads in the majority population.


Figure 6: Possible collaboration network outcomes. Pink nodes represent minority group members and blue nodes represent majority group members. (a) shows an outcome where demanding High is common in the majority population and homophily is .4921. (b) shows an outcome where demanding Med is common in the majority population and homophily is .1005.

15

in the network reaching the norm of fair division (with no homophily in the network) to all majority members discriminating (with a totally segregated network having no links between social identity groups). Nearly all outcomes are somewhere between these two extremes: partially segregated networks with some members of the majority group upholding the discriminatory norm.

In order to quantify how the level of discrimination in a network impacts homophily, we use the following measure of *inbreeding homophily* (which we will denote I_i) from Currarini et al. (2009):

$$I_i = \frac{H_i - w_i}{1 - w_i}$$

where H_i is the proportion of a group i's links that are within group links and w_i is the fraction of the population that group comprises. Thus, inbreeding homophily measures how homophilic groups are (by calculating the proportion of their links that are within group) compared to how homophilic they could be (the 1 in the denominator represents 100% of links being within group), while controlling for the possibility that groups could make up different proportions of the total population (by subtracting w_i in both the numerator and denominator).

As figure 7 shows, more discrimination means more homophily. This is a very stochastic process, as one can see from the wide spread of possible levels of homophily for each level of majority discrimination. However, if we look at the regression line we see an overall trend where as more of the majority group discriminates, homophily is stronger.[22]

4 Discussion

We can now return to the empirical results mentioned in the beginning of section 2. As is evident from parts (b) and (c) of the last section, our models suggest a connection between evidence that women receive less credit in collaboration and evidence that women tend to collaborate less and more often with other women. Furthermore, our models provide a potential mechanism for in-group clustering in academia. Those who get less by dint of discriminatory social norms may take steps to protect themselves from discrimination.[23]

What is the upshot for epistemic communities and epistemic progress? As mentioned in the introduction, diversity has been championed as an important feature of successful academic communities both by those in feminist epistemology/philosophy of science, and by those doing formal work in social epistemology. Does the type of homophily we attempt to explain actually impede epistemic progress?

There might be reason to think not. So long as diverse ideas are present somewhere in a community, why should it matter whether collaborations happen between groups

---

[22]One might think that the results in figures 4 and 7 together imply that as minority size increases, homophily decreases. We found that this is true with a few exceptions. These are possibly due to the fact that the measure of homophily does not account for the fact that the maximum number of links may affect groups of different sizes in different ways.

[23]An alternative explanation is that people just feel drawn to those like them. It is entirely possible that both factors are at play in the real world.
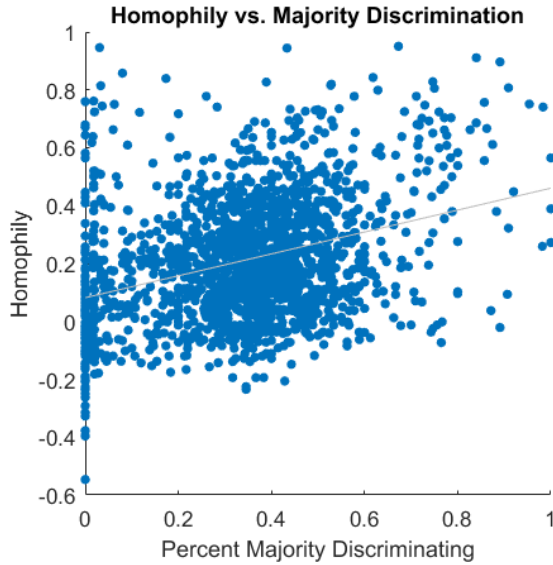
Figure 7: Level of homophily, measured using inbreeding homophily, with changes in majority discrimination, given by the percent of the majority group whose out-group strategy is to demand High.

or not? For example, suppose a researcher from one social identity group is likely to figure out A and a researcher from another group B. If they collaborate, then they might also together conclude C, which follows from A&B. However, the community has another route to knowing C: A and B are published separately, then the community as a whole has access to these ideas and concludes C for themselves.[24] Okruhlik (1994) seems to have a picture somewhat like this in mind—diverse researchers will generate and test diverse hypotheses which will then be assessed by the usual scientific methods.

There are few reasons why we think diversity within collaborations might be important. First, it is possible that A and B independently are not interesting enough to warrant publication and that they are only valuable realizations insofar as they jointly imply C. If this is the case, it is reasonable to think that A and B would never be published on their own. This might be especially likely if members of one group struggle to publish in top journals due to reputational effects. There is reason, as well, to think that the actual process of collaborating with those unlike ourselves might influence deliberation and discovery. For example, Sommers (2006) find that deliberation processes proceed differently on racially diverse juries in that members share more information. And Phillips et al. (2006) find improved problem solving in small, racially diverse groups.

Finally, even if collaborations with diverse authorship do not tend to be of higher quality, homophily alone can impede epistemic progress. Information spreads more slowly through homophilic networks (Golub and Jackson, 2012), so a homophilic epistemic network will be less efficient in that it will take longer for the community to reach various

---

[24]Thanks to Liam Kofi Bright for raising this concern.

17

conclusions.[25] Under certain assumptions, such homophilic networks can also prevent the spreading of new and better scientific practices throughout the community as a whole (Schneider, 2017).

Our models suggest a process by which academic communities can spontaneously un-diversify in the face of discriminatory bargaining norms. Furthermore, they suggest that such norms can emerge in academic communities under many conditions, and are more likely to impact minority groups. As we note above, this may help explain patterns observed in real communities where minority groups receive less credit, and tend to collaborate less. Those interested in intervening on epistemic communities may find these results useful in that they generate a theoretically well grounded hypothesis for why we see such patterns.

Bibliography

Alexander, Jason and Brian Skyrms (1999). "Bargaining with neighbors: Is justice contagious?." *The Journal of philosophy*, 588–598.

Axtell, Robert L, Joshua M Epstein, and H Peyton Young (2001). "The emergence of classes in a multiagent bargaining model." *Social dynamics*, 191–211.

Barjak, Franz and Simon Robinson (2008). "International collaboration, mobility and team diversity in the life sciences: impact on research performance." *Social Geography*, *3*(1), 23.

Boschini, Anne and Anna Sjögren (2007). "Is team formation gender neutral? Evidence from coauthorship patterns." *Journal of Labor Economics*, *25*(2), 325–365.

Botts, Tina Fernandes, Liam Kofi Bright, Myisha Cherry, Guntur Mallarangeng, and Quayshawn Spencer (2014). "What Is the State of Blacks in Philosophy?." *Critical Philosophy of Race*, *2*(2), 224–242.

Bruner, Justin P (2017). "Minority (Dis)advantage in Population Games." *Synthese*, 1–15.

Bruner, Justin P and Cailin O'Connor (2018). "Power, Bargaining, and Collaboration." *Scientific Collaboration and Collective Knowledge*. Ed. Thomas Boyer and Michael Weisberg Conor Mayo-Wilson. Oxford University Press.

Campbell, Lesley G, Siya Mehtani, Mary E Dozier, and Janice Rinehart (2013). "Gender-heterogeneous working groups produce higher quality science." *PloS one*, *8*(10), e79147.

Currarini, Sergio, Matthew O Jackson, and Paolo Pin (2009). "An economic model of friendship: Homophily, minorities, and segregation." *Econometrica*, *77*(4), 1003–1045.

---

[25]This argument assumes that people are more likely to engage with their co-authors, in terms of discussing relevant research and spreading ideas, but we think that this is a reasonable assumption.

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), 237–249.

Dezsö, Cristian L and David Gaddis Ross (2012). "Does female representation in top management improve firm performance? A panel data investigation." *Strategic Management Journal*, *33*(9), 1072–1089.

Fehr, Carla (2011). "What is in it for me? The benefits of diversity in scientific communities." *Feminist epistemology and philosophy of science.* . Springer, 133–155.

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.

Ferber, Marianne A and Michelle Teiman (1980). "Are women economists at a disadvantage in publishing journal articles?." *Eastern Economic Journal*, *6*(3/4), 189–193.

Freeman, Richard B and Wei Huang (2015). "Collaborating with people like me: Ethnic coauthorship within the United States." *Journal of Labor Economics*, *33*(S1), S289–S318.

Golub, Benjamin and Matthew O Jackson (2012). "Network structure and the speed of learning measuring homophily based on its consequences." *Annals of Economics and Statistics/ANNALES D'ÉCONOMIE ET DE STATISTIQUE*, 33–48.

Grim, Patrick, Daniel J Singer, Christopher Reade, and Steven Fisher (2015). "Germs, Genes, and Memes: Function and Fitness Dynamics on Information Networks." *Philosophy of Science*, *82*(2), 219–243.

Haraway, Donna Jeanne (1989). *Primate visions: Gender, race, and nature in the world of modern science.* Psychology Press.

Harding, Sandra G (1991). *Whose science? Whose knowledge?: Thinking from women's lives.* Cornell University Press.

Holman, Bennett and Justin P. Bruner (2015). "The Problem of Intransigently Biased Agents." *Philosophy of Science*, *82*(5), 956–968.

Jackson, Matthew O (2010). *Social and economic networks.* Princeton university press.

Jackson, Matthew O and Asher Wolinsky (1996). "A strategic model of social and economic networks." *Journal of economic theory*, *71*(1), 44–74.

Lazarsfeld, Paul F, Robert K Merton, et al. (1954). "Friendship as a social process: A substantive and methodological analysis." *Freedom and control in modern society*, *18*(1), 18–66.

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

Mayo-Wilson, Conor, Kevin Zollman, and David Danks (2013). "Wisdom of crowds versus groupthink: learning in groups and in isolation." *International Journal of Game Theory, 42*(3), 695–723.

McDowell, John M and Janet Kiholm Smith (1992). "The effect of gender-sorting on propensity to coauthor: Implications for academic promotion." *Economic Inquiry, 30*(1), 68–82.

Nash, John F (1950). "The bargaining problem." *Econometrica: Journal of the Econometric Society*, 155–162.

O'Connor, Cailin (2017a). "The Cultural Red King Effect." *The Journal of Mathematical Sociology.*

O'Connor, Cailin (2017b). "The Dynamics of Inequity.".

O'Connor, Cailin, Liam K. Bright, and Justin P. Bruner (2017). "The Evolution of Intersectional Oppression." *Working Paper*, 2017.

O'Connor, Cailin and Justin Bruner (2017). "Dynamics and Diversity in Epistemic Communities.".

Okruhlik, Kathleen (1994). "Gender and the biological sciences." *Canadian Journal of Philosophy, 24*(sup1), 21–42.

Page, Scott E (2008). *The difference: How the power of diversity creates better groups, firms, schools, and societies.* Princeton University Press.

Phillips, Katherine W, Gregory B Northcraft, and Margaret A Neale (2006). "Surface-level diversity and decision-making in groups: When does deep-level similarity help?." *Group processes & intergroup relations, 9*(4), 467–482.

Poza, David, José Manuel Galán, José Ignacio Santos, and Adolfo López-Paredes (2010). "An Agent Based Model of the Nash Demand Game in Regular Lattices." *Balanced Automation Systems for Future Manufacturing Networks.* . Springer, 243–250.

Richard, Orlando C, BP Murthi, and Kiran Ismail (2007). "The impact of racial diversity on intermediate and long-term performance: The moderating role of environmental context." *Strategic Management Journal, 28*(12), 1213–1233.

Rosenstock, Sarita, Cailin O'Connor, and Justin Bruner (2017). "In Epistemic Networks, Is Less Really More?." *Philosophy of Science, 82*(2), 234–252.

Sarsons, Heather et al. (2017). "Recognition for Group Work: Gender Differences in Academia." *American Economic Review, 107*(5), 141–145.

Schneider, Mike D (2017). "The non-epistemic origins of research strongholds." (Working Paper).

Schneider, Mike D, Hannah Rubin, and Cailin O'Connor (2017). "Promoting Diverse Collaborations." (Working Paper).

Sigmund, Karl, Christoph Hauert, and Martin A Nowak (2001). "Reward and punishment." *Proceedings of the National Academy of Sciences*, *98*(19), 10757–10762.

Skyrms, Brian (1994). "Sex and justice." *The Journal of philosophy*, 305–320.

Skyrms, Brian (2014). *Evolution of the social contract*. Cambridge University Press.

Sommers, Samuel R (2006). "On racial diversity and group decision making: identifying multiple effects of racial composition on jury deliberations.." *Journal of personality and social psychology*, *90*(4), 597.

Sugimoto, Cassidy R, Vincent Lariviere, CQ Ni, Yves Gingras, Blaise Cronin, et al. (2013). "Global gender disparities in science." *Nature*, *504*(7479), 211–213.

Watts, Alison (2003). "A dynamic model of network formation." *Networks and Groups*. Ed. Bhaskar Dutta and Matthew O. Jackson. Springer, 337–345.

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.

Young, H Peyton (1993). "The evolution of conventions." *Econometrica: Journal of the Econometric Society*, 57–84.

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

Zollman, Kevin JS (2010). "The epistemic benefit of transient diversity." *Erkenntnis*, *72*(1), 17–35.