Discrimination and Collaboration in Science

Hannah Rubin and Cailin O'Connor

February 27, 2017

In this paper we take an in-depth look at the dynamics of discrimination and academic collaboration. We find that discrimination can lead members of different social groups to only collaborate with in-group members. As we argue, this decrease in the diversity of scientific collaborations may 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 pointed out 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]

Given this work, we ask: where do epistemic collaboration networks come from? And: what factors influence the diversity of these networks? In this paper we take an in-depth look at the dynamics of discrimination and collaboration. 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.[3] 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 collaborations across groups. Lastly, we examine the simultaneous co-evolution of discrimination and collaboration,

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

[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 see Haraway (1989); Longino (1990); Harding (1991).

[3] This echos work on the cultural Red King effect by Bruner (2014); Bruner and O'Connor (2015); O'Connor (2016).

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. This, as described above, 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. In part this analysis is inspired 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. 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 effective 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. 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 peacefully 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

[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.

		Player 2	
	Low	**Med**	**High**
Player 1 **Low**	L,L	L,5	L,H
Med	5,L	5,5	0,0
High	H,L	0,0	0,0

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

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.

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

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.[6] 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.[7] 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. (2000) 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.

[5]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 (2015) for an overview of the literature on this topic.

[6]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.

[7]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., 2000; Skyrms, 2014).

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, we focus on models where actors belong to two social groups, and in particular where one social group is in the minority. We will be particularly interested in results where the majority demands 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.

We will tackle the question of discriminatory norms and collaboration 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 agent's bargaining strategies co-evolve with the structure of the collaboration network. We will see that the discriminatory norm tends to arise between many members of the community. Further, as this norm 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 framework much like ours 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, we find that the minority group can be disadvantaged. Further, we investigate whether *homophily*, the tendency to preferentially form links with members of your own social identity group (Currarini et al., 2009), 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) then 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 at 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 in awhile. 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 an in-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).

First, we look at the effect of minority group size on convergence to the possible bargaining norms between groups. Figure 2 shows results when $H = 6$. 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.

To explain this effect, consider a simple demonstrative example. If there are 10 majority group members and 5 minority group members, and 10 total out-group links, on average majority group member will have one out-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. In contrast, a majority group member having one link to the minority would want to demand 6 if their collaborator demands 4, and this initially happens with probability 1/3. So, it is much more likely that a

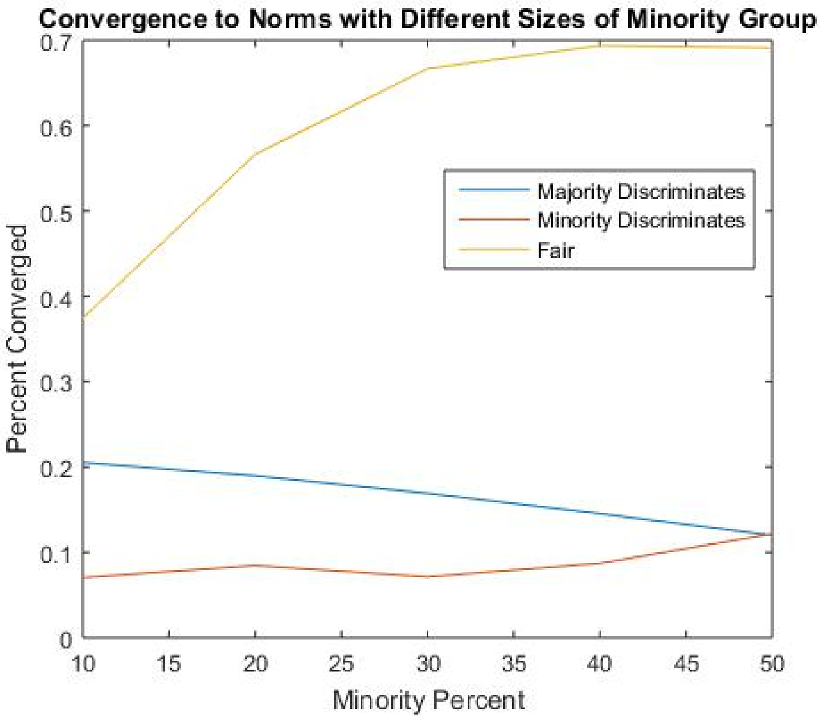

Figure 2: Convergence to different norms over size of minority group.

majority group member's best response is initially to demand 6. By similar reasoning we can see that it is also more likely a minority group member's best response is to demand 4. This asymmetry drives populations toward the outcome where the majority demands High.

Our findings are somewhat similar to previous results on the cultural Red King effect (Bruner, 2014; O'Connor and Bruner, 2016; O'Connor, 2016).[8] Under this effect if one social group is more reactive than another, they may be disadvantaged in that they are less likely to end up at preferable bargaining norms. In particular, the above authors show that minority status can lead to increased reactivity since minority group members are more likely to meet majority group members and so to learn quickly how to interact with them, though the reverse is not true.[9]

Second, we found that varying levels of homophily did not influence whether the majority or minority group is more likely to discriminate.[10] Since agents are updating their in-group and out-group strategies separately, how the in-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; it arises because there are fewer out-group links per majority group member.

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 (2016) find that in epistemic communities discriminatory norms discincentivize collaboration between social groups, and, as will become clear, our results support this finding.

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 preexisting discriminatory norm, majority group members receive a payoff of 6 from a between group link while

[8]The name of this effect comes from evolutionary game theoretic work in biology by Bergstrom and Lachmann (2003) showing that slow evolving species can sometimes gain an advantage in a mutualism.

[9]These authors find that for the Nash mini-game used here, the values of the Low and High demands determine whether the minority group will be advantaged or disadvantaged by their size and corresponding reactivity. When $H < 7$, a Red King is observed, otherwise a weak Red Queen (where the minority group is advantaged by dint of their size) is observed. In these models, we observe an analog to the Red King when $H < 7$, but no Red Queen. For more on the potential for this effect to impact epistemic communities see O'Connor and Bruner (2016).

[10]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 totally connected).

minority members receive a payoff of 4.[11] 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.

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.[12]

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 the lowest payoff if they already had the maximum number of links).

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.[13] Figure 3 provides an example of how this occurs. Initially, links form steadily within each group and also across each group, 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

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

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

[13] 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.

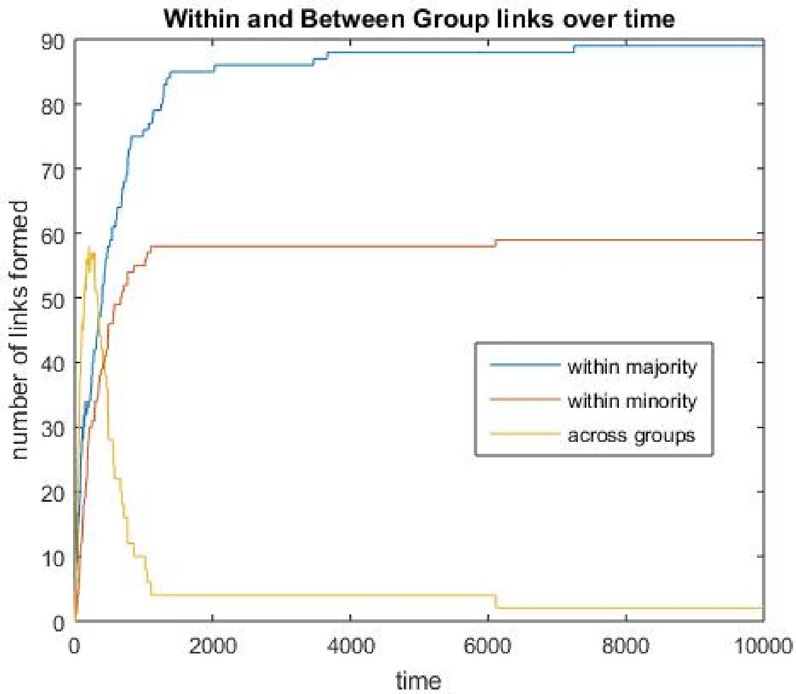

Figure 3: Within and between group links over time, for a network of 100 researchers with the minority group comprising 40% of the population, high demand of 6, low demand of 4, and maximum links set to 3 per person

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, within 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 groups decrease over time until they are essentially non-existent.[14]

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

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

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, each agent takes an action with probability 0.1. There are two types of possible actions: updating your bargaining strategy and updating your set of collaborators. If an agent takes an action in the round, there is a 20% chance they will update their set of collaborators and an 80% chance they will update their strategy. This represents a situation where people are not constantly updating and it is easier to update your bargaining strategy than form a new collaboration. (The particular probabilities are not important; similar results can be obtained for a variety of values.) 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 with the minority group comprising 30% of the population, a high demand of 6 and maximum links set to 10.[15] Simulations were run for 10,000 rounds.

First, we will look at the evolution of strategies over time. As in section 3.1, within groups, the overwhelmingly likely outcome was fair division. Figure 4 shows two possible outcomes for the evolution of between group strategies. Figure 4(a) shows an outcome where at the end of the simulation, more majority group members demand High while minority group members demand Low. Figure 4(b) shows an outcome where both majority and minority members tend to demand Med. Note that in both cases, the evolution of strategies stops before the collaboration network settles into one norm or the other. To understand why this is the case, we look at the evolution of the collaboration network.

The evolution of the collaboration networks is similar to the evolution in section 3.2; both within and between group links increase initially, but between group links decrease later. Here, though, between group links do not disappear altogether. This is because, unlike in section 3.2, usually some majority group members demand Med against the minority. Figure 5(a) and (b) show the end product of this evolution, corresponding to the simulations depicted in Figures 4(a) and (b), respectively. Minority group members will break links with majority group members who demand High in favor of forming links with other minority group members who demand Med. The links that remain between groups are those where both parties demand Med.

Once an agent has no more links to the other group, they will stop updating their between group strategy (all demands are equally best responses, so there is no incentive to change strategies). This explains why the network never fully settles on one norm or

[15]Similar results are found for a variety of parameter values. We chose H = 6 to induce majority advantage.

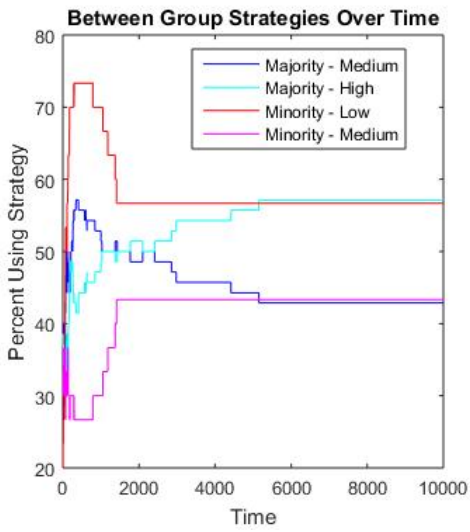

(a)

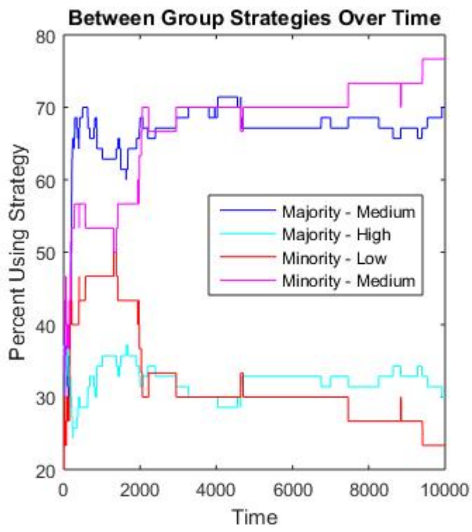

(b)

Figure 4: 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.

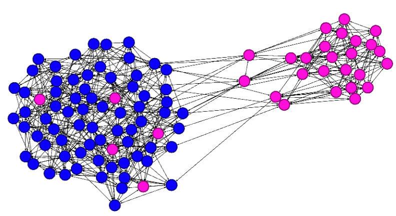

(a)

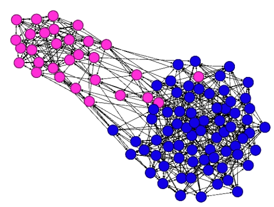

(b)

Figure 5: 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. (b) shows an outcome where demanding Med is common in the majority population.

another. Some majority group members settle on fair division norms with their collaborators in the minority group. Others settle on a discriminatory norm which disadvantages the minority. The minority group members then break links with these discriminating majority group members until the discriminators have no more collaborators in the minority group. This means that these discriminators never update their strategies, and the network persists with some majority group members retaining the discriminatory strategy but without ever having a collaboration in which to discriminate.

As is easy to see from figure 5, this process leads to collaboration networks which are homophilic, with members tending to interact more often with members of their own group. Further, since minority members are willing to keep between group links with majority group members with whom they have settled on a norm of fair division, the more of the majority that demand Med, the more between group links the collaboration network ends up having. There is a continuum of possible outcomes, ranging from everyone 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 will be somewhere between these two extremes: partially segregated networks with some members of the majority group upholding the discriminatory norm.

4 Conclusion

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/work in collaborations 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.

The models in part (a) further suggest that previous results on minority disadvantage in the emergence of bargaining norms replicate in a new context. In academic communities, where it is often the case that women and people of color are in the minority, this is particularly germane.

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. Our models suggest a process by which academic communities will sponaneously un-diversify in the face of discriminatory bargaining norms. Furthermore, they suggest that such norms can spontaneously emerge in academic communities under many conditions, and are more likely to impact minority groups. This is an obvious concern for epistemic progress. To ensure the diversity of epistemic communities and collaborations may take concerted effort to fight the social dynamical forces that divide social groups.

Bibliography

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

Axtell, Robert, Joshua M Epstein, and H Peyton Young. "The emergence of classes in a multi-agent bargaining model.".

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), 593–598.

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 (2014). "Racists and Minorities in Population Games." working paper.

Bruner, Justin P and Cailin O'Connor (2015). "Power, Bargaining, and Collaboration." *Scientific Collaboration and Collective Knowledge*. Ed. Conor Mayo-Wilson Thomas Boyer and Michael Weisberg. Oxford University Press.

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

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.

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.

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.

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.

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 (2016). "The Cultural Red King Effect." working paper.

O'Connor, Cailin and Justin Bruner (2016). "Dynamics and diversity in epistemic communities.".

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.

Rosenstock, Sarita, Cailin O'Connor, and Justin Bruner (pear). "In Epistemic Networks, Is Less Really More?." *Philosophy Compass.*

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.

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.* . 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.