



Structural causes of citation gaps

Hannah Rubin¹ 

Accepted: 26 November 2021

© The Author(s), under exclusive licence to Springer Nature B.V. 2021

Abstract The social identity of a researcher can affect their position in a community, as well as the uptake of their ideas. In many fields, members of under-represented or minority groups are less likely to be cited, leading to *citation gaps*. Though this empirical phenomenon has been well-studied, empirical work generally does not provide insight into the causes of citation gaps. I will argue, using mathematical models, that citation gaps are likely due in part to the structure of academic communities. The existence of these ‘structural causes’ has implications for attempts to lessen citation gaps, and for proposals to make academic communities more efficient (e.g. by eliminating pre-publication peer review). These proposals have the potential to create feedback loops, amplifying current structural inequities.

Keywords Philosophy of science · Structural inequity · Citation gaps · Formal models of science

1 Introduction

How do ideas spread throughout academic communities? One important factor to consider is that the social identity of a researcher can affect their position in a community, as well as the uptake of their ideas. Recognition of this fact has been significant for philosophers of science, who attempt to understand the way scientific communities function, as well as social epistemologists, social scientists, and others who attempt to understand knowledge production more generally.

There are various ways to operationalize ‘position in a community’ and ‘uptake of ideas’ in order to more concretely explore the impact of social identity on the way

✉ Hannah Rubin
hannahmrubin@gmail.com

¹ Department of Philosophy, University of Notre Dame, South Bend, IN, USA

these communities function. It is common to use networks (which summarize the various connections among a community's members) along with measures of 'connectedness' or 'centrality' of researchers, to capture a researcher's place in a community. We will see below how these networks capture the 'structure' of a community, in terms who is connected to whom, and how social identity can importantly affect this structure. One common way to measure the uptake of ideas by a research community is to look at citations accumulating to published papers, perhaps along with some measures of an author's impact that are based on citations, like the h-index or i10-index.

Though there are other ways to operationalize these concepts, relying on these common measures will allow this paper to illuminate under-appreciated causal processes that lead to work by marginalized groups being overlooked. Further, I will argue that these processes have important consequences for our understanding of how research communities (ought to) function.

Before advancing these arguments, I will discuss evidence that members of underrepresented and minority groups are often less likely to be cited; they are overlooked in favor of members of a majority group, even when the minority group member published the same results or arguments first or around the same time. These inequities in citations, or *citation gaps*, have been the subject of a number of empirical studies, which will be reviewed in Sect. 2. One limitation of empirical work on citation patterns, however, is that it generally does not determine the causes of citation gaps, but merely determines whether or not they exist.

While there are a variety of proposed explanations of this phenomenon, I will argue that citation gaps are likely due, at least in part, to the structure of academic communities (Sect. 3). This is, of course, not to deny the existence of other factors (e.g. psychological biases), as Sects. 3.3 and 3.4 will discuss. Given the difficulty of using empirical data to determine causes of citation gaps, this paper will present mathematical models to show that the way academic communities are structured, in terms of who is connected to whom, can influence how likely it is that people have heard of certain papers, and therefore how likely it is that they will cite papers from certain social identity groups.

This situation is concerning, not only because we ought to care about equitable opportunities, but also because our epistemic goals are likely to be hindered if good work is repeatedly overlooked. Furthermore, as I will argue in Sect. 4, the existence of these structural causes of citation gaps means that certain attempts to make academic communities more efficient (e.g. by eliminating pre-publication peer review) have the potential to create feedback loops, where initial inequities in citation practices feed back into greater and greater inequities over time. As Sect. 5 will discuss, identifying the causes of citation gaps is also crucial to understanding the impact of possible interventions aimed to ameliorate them. This will, of course, influence our thinking as we attempt to both promote equitable opportunities and further our epistemic goals.

2 Citation gaps and structural causes

Citation gaps according to gender have been found in such fields as economics (Ferber, 1988; Ferber & Brün, 2011), ecology (Cameron et al., 2016), political science (Dion & Mitchell, 2012; Maliniak et al., 2013; Mitchell et al., 2013; Dion et al., 2018), library and information sciences (Håkanson, 2005), linguistics and sociology (Leahey et al., 2008), health and natural sciences (Aksnes et al., 2011; Beaudry & Larivière, 2016), social psychology (Nosek et al., 2010) and neuroscience (Dworkin et al., 2020). Some of these studies also find that women are less likely to be cited in sub-disciplines which are more male dominated (e.g. Dion et al., 2018)—so, plausibly, representation matters as subfields with greater gender balance tend to have smaller citation gaps. While there are far fewer studies regarding citation rates and race or ethnicity, citation gaps according to race have been found, for instance, in law (Merritt, 2000) and social psychology (Nosek et al., 2010).

There are also studies where gender citation gaps have not been observed (e.g., in public administration (Corley & Sabharwal, 2010), social problems (Ward et al., 1992), international relations (Østby et al., 2013), economic history (Di Vaio et al., 2012), forestry and geography (Slyder et al., 2011), and criminal justice (Stack, 2002)) or where women have been found to receive more citations than men (biochemistry (Long, 1992) and construction studies (Powell et al., 2009)).¹ In addition to the studies already mentioned, that focus on individual fields and often focus on particular countries, some studies find that gender citation gaps exist when we look across many different disciplines and many different countries (Larivière et al., 2013; Huang et al., 2020).

There are many ways to dive into the details of these studies, but, for the purposes of this paper, I will mention one more specific empirical result that will be important for understanding the models of structural causes in Sect. 3. As Dion et al. (2018) note, studies focusing only on overall citation counts “miss the sociological aspects of how scholars recognize the work of their peers” (p. 314), e.g. whether social factors like gender affect citation practices. For this reason, at least some of the studies on citation gaps also look to see whether the gender of the author(s) of a paper predicts the gender composition of authors in its references section. That is, they look for what we might call a *citation ratio gap*, capturing how often men are citing men versus citing women compared to how often women are citing men versus citing women:

$$\frac{\text{men citing men}}{\text{men citing women}} - \frac{\text{women citing men}}{\text{women citing women}} > 0$$

These citation ratio gaps have been found in a number of studies (Ferber, 1986, 1988; Ferber & Brün, 2011; Dion & Mitchell, 2012; Mitchell et al., 2013; Dion et al., 2018), even studies that did not find a citation gap (Ward et al., 1992).

¹ There are a variety of reasons why we might not expect gender citation gaps in every field. See Sect. 3.4 for further discussion the factors relevant to the existence of citation gaps.

Of course, as will be discussed at greater length in Sect. 3.4, there are many different potential factors interacting to produce citation patterns. So, when we look at studies on citation gaps, we should not expect, e.g. the correlation between representation and whether or not there is a gap to be perfect, or to find citation gaps in every field. It would be good to know something more about the causes of these gaps. Some proposed explanations of citation gaps (and citation ratio gaps) are implicit and explicit cognitive biases, with the emphasis generally on implicit bias as a possible cause. While there are some studies showing that bias could be at work (e.g. Knobloch-Westerwick et al., 2013), interpreting the evidence is not necessarily straightforward (Lee et al., 2013; Lee, 2016). The productivity gap, where men are found to publish more papers than women, combined with the fact that men have been found to self-cite at a greater rate could explain some part of observed gender gaps. Different lengths of careers are also plausibly relevant to explaining differential accumulation of citations. Women's research careers tend to be shorter than men's; for whatever reason there is a 'dropout gap' or 'leaky pipeline' where women leave positions at a greater rate.

The possible explanation I will focus on is in terms of the structure of the community, in terms of who is connected to whom and how information spreads throughout the network.² The sort of thing has been suggested before as a possible explanation of gender citation gaps (e.g., Sen, 2018; Brown & Samuels, 2018; Esarey & Bryant, 2018), but usually in an informal way, e.g. by noting women may be less 'well-networked'. This paper will give some more substance to this possible explanation, showing exactly how network structure can affect citation rates in different social identity groups.

It is generally very difficult to test the potential causes of something like citation gaps. While studies may show, for example, the existence of biases, in any real community there are too many factors at play to pin down one particular cause or set of causes. In order to explore causal pathways which are hard to study in the real world, Sect. 3 will present mathematical models. These models abstract away from the messy details of interactions in real academic communities to investigate causal relationships of interest, e.g. the relationship between the structure of the community and the size of citation gaps. Additionally, while many empirical studies on citation gaps only consider gender as a possible difference maker, the models here potentially apply to any aspect of social identity, so long as that aspect importantly shapes interactions. Therefore, much of what will be discussed in the rest of the paper could illuminate causes of citation gaps for underrepresented and minority groups in general. I will discuss the results of the models as applying to social identity in general, except when drawing on empirical evidence specifically related to gender.

² In calling this a 'structural cause', I am using the phrase to indicate that the cause pertains to how the group is organized, or that the relations between the parts within the whole are important. The use of 'structural' in this phrase should not be taken to encompass all types of causes that are referred to as 'structural' (e.g. structural oppression or systemic discrimination), nor to preclude there being being a myriad of such structural causes present in the community. Thanks to a reviewer for pointing out this possible confusion.

These models will demonstrate how citation gaps can arise from the social structure of academic communities. This discussion fits with recent work in philosophy of science [sparked by the dissertation work of Justin Bruner, later published as Bruner (2019)], showing how large-scale inequities for minority groups can arise from each individual merely performing a rational strategy. These inequities arise, for example, in everyday bargaining (Bruner, 2019; O'Connor et al., 2019; O'Connor, 2019), academic collaborations (Bruner & O'Connor, 2015; Rubin & O'Connor, 2018), and priority disputes over scientific discoveries (Rubin & Schneider, 2020). In addition to providing a new sort of phenomena where large-scale inequities are expected to arise for minority groups (citation gaps), the paper will demonstrate new mechanisms by which these inequities can emerge (detailed in Sects. 3, 4).³ This demonstration will then put me in the position to show how these structural causes can interact with other countervailing factors (internet searches), as well as empirical observations about publication rates and cognitive biases.

Before presenting these models, I should say something about a common reaction to empirical work on citation gaps. People often respond with something like the following claim: “I cite whatever paper is best/most relevant, I don’t look at demographics.” On the one hand, this sort of ‘merit defense’ does likely apply to many people who are well-intentioned in their citation practices, and asking people to spend time thinking about who they are citing can seem burdensome or unfair. On the other hand, if citation gaps are due, at least in part, to these structural causes, this sort of response misses the mark. It could be that everyone is using the unbiased and reasonable strategy of citing the most relevant and best papers they are aware of, and yet this still leads to inequities. This is because they are citing the best, most relevant papers *of the papers they know about*. However, due to the way information spreads through their academic network, they may not know about relevant papers from members of certain social identity groups.

3 Modeling structural causes of citation gaps

This section will describe how to mathematically represent the way academic communities are structured, then use this representation to show how citation gaps can emerge even in a community of people who do not pay attention to demographics when deciding who to cite.⁴

The structure of a community can be usefully captured by a network, which summarizes who is connected to whom. In these networks, there are nodes which represent different academics and edges or links between them which we can think of as representing regular communication channels, e.g. regular collaborations, people who read each other’s papers and give comments, etc. Within the network, there are people of different social identity ‘types’ and these different types can make up a larger or smaller proportion of the total community. This means that we

³ Building on Rubin and Schneider (2020), this paper can be thought of as demonstrating particular mechanisms regarding credit attribution within the ‘credit economy’ by which these inequities can emerge.

⁴ A copy of the code for this paper has been made available on the Open Science Framework at https://osf.io/9suty/?view_only=70022ee795654ee9a22bc311a97a763e.

can meaningfully talk about minorities and majorities within the population. The communities considered here are simplified in that there are only two different social identity types—for instance men and women, with women being underrepresented and thus a minority in a particular academic discipline.

One important factor that influences the structure of these communities is homophily, or the tendency of people to interact within their own social identity group. There is a preponderance of evidence that networks describing both academic collaborations and personal friendships are homophilic, that people of the same social identity often cluster in subdisciplines within the larger discipline, etc. (del Carmen & Bing, 2000; Currarini et al., 2009; West et al., 2013; Botts et al., 2014; Wang et al., 2019). There are a variety of reasons why people may, consciously or subconsciously, form links more often within their social identity group rather than outside it. For instance, an unfair distribution of labor in collaborative projects may lead minority group members to break ties with the majority group (Rubin & O'Connor, 2018). There can also be positive reasons to form within-group links, such as receiving support or relevant information from people within your own social identity group (see, e.g. Yang et al., 2019 and citations therein). In any case, there is good reason to include homophily in the structure of the communities we are trying to represent.

One way to capture homophily in a network is to use multi-type random networks (Golub & Jackson, 2012). These networks are generated in a fairly straightforward fashion: for every pair of nodes in the network, there is some probability a link is formed between them, which depends on whether the nodes are of the same type. When there is homophily, there is a higher probability of a link forming if the nodes are in the same social identity group, $p(in)$, and a lower probability if they are from two different social identity groups, $p(out)$.

3.1 The basic model

Let us imagine that we have two very similar papers making the same claim, which are published at exactly the same time, in similar journals, etc. We can ask: all else being equal, does the social identity of the authors make one paper more likely to be cited than the other? In this basic model, the only way that social identity influences citation chances is by influencing the place an author occupies on the network.

Since we are interested in what happens when these authors are from two different social identity groups, each time the model is run there is one randomly chosen majority member and one randomly chosen minority member who each publish a paper making the same claim at the same time. We then track how often the majority group's work (paper 1) gets cited versus the minority group's work (paper 2) as new papers come out. We draw 200 people at random (with replacement) from a network of 100 academics to publish a paper, and, when they do, there is some chance they cite paper 1 or paper 2 based on how likely it is that they know about the paper. Of course, these new papers coming out also have some chance of being cited, but to answer the current question of interest we will only track how many citations accumulate to the original two papers—does author 1 get more citations or author 2?

As new papers come out, authors of these new papers have some chance of citing each of the two original papers. If the new paper is by one of the original two authors, they will cite themselves, the assumption being that they know about the paper they themselves wrote. If the new paper is not by one of the two original authors, there is some chance they will hear about each of the papers through their network. This chance is determined by the shortest path length between the author of the new paper and the authors of the original papers—that is, how many links on the network it takes to get from one person to the other. With probability p an academic hears about and cites their neighbor's paper, with probability p^2 they cite their neighbors' neighbor, and so on.⁵ So, there is some chance a person hears about each of the two papers, and if they know about a paper they will cite it. It is possible to know about and cite both papers.

We look at how both homophily and representation in a field can affect the size of citation gaps. Recall that there is homophily in a community if a person has a higher probability of forming a link within their own social identity group than with someone outside that social identity group, or $p(in) > p(out)$. The greater the difference between these two linking probabilities, the greater the homophily of the community. Results are presented for a range of cases, starting with $p(in) = p(out)$, i.e. no homophily. Every time $p(in)$ was increased by .01, $p(out)$ was decreased by the same amount, creating a range of levels of homophily. For simplicity, only $p(in)$ values are shown in Fig. 1.

Figure 1a shows results for different sizes of the majority population and different levels of homophily.⁶ We track the size of the citation gap, i.e. the proportion of total citations that go to a majority group member minus the proportion that go to a minority group member. If the citation gap is $\frac{1}{3}$, for example, this means the majority group is cited two times for every time the minority group is cited (i.e. they get $\frac{2}{3}$ of the citations compared to the minority group's $\frac{1}{3}$). When there is no homophily, there is no citation gap—both groups are cited equally often. But, as we increase homophily, the relative sizes of the groups starts to matter more and more. In the extreme case of high homophily, and only 10% of the field comprised of the minority group, the paper by a majority group member is cited more than twice as often as the one by a minority group member, for a citation gap of roughly .34. This matches with empirical evidence that women are less likely to be cited in sub-disciplines which are more male dominated (Dion et al., 2018).

Figure 2a, b show that the citation gaps in this model can be explained by appeal to citation ratio gaps, e.g. men citing men more often and women citing women more often. Figure 2a shows the majority citation of majority ratio (i.e., the ratio of majority member's papers to minority member's papers cited in new papers coming out by majority group members). This ratio is not strongly affected by the size of the

⁵ The results presented here are for $p = .3$, but similar results can be obtained with higher or lower probabilities.

⁶ To get a reliable estimate of the expected citation gap, 100 different networks were formed for each combination of $p(in)$ and majority group size, and 100 simulations of the citation process were run on each (with authors of the original papers chosen at random each time).

Fig. 1 Homophily and group size affect citation gaps

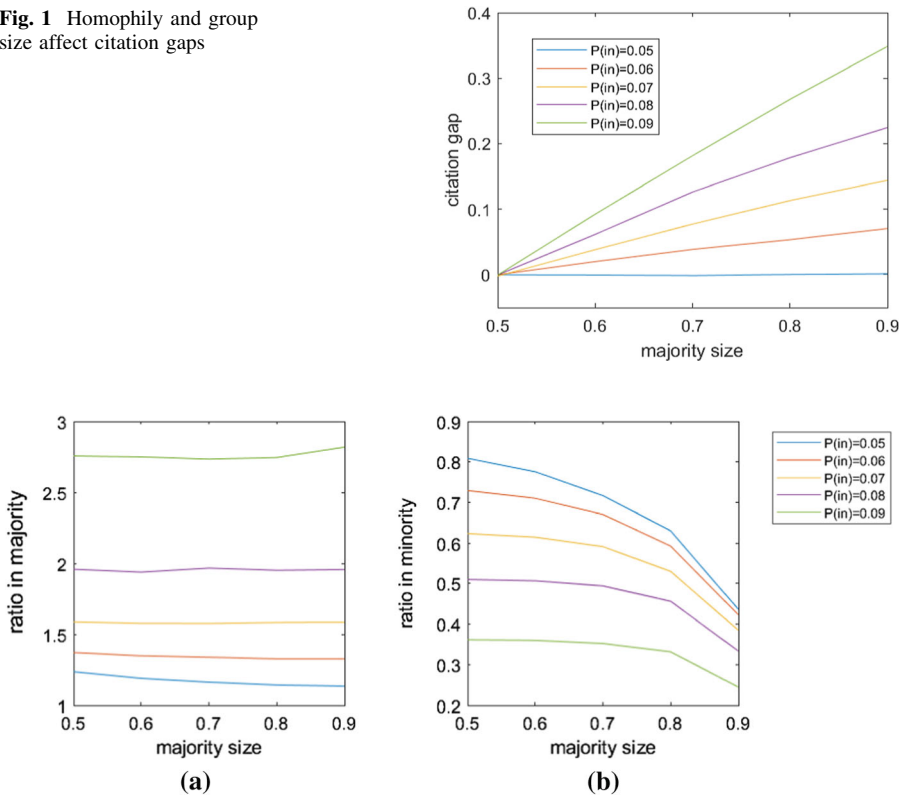


Fig. 2 Homophily and group size affect the citation gap by affecting how often **a** majority group members and **b** minority group members cite members of the majority group

majority group, but it is affected by homophily. The more people cluster into subgroups based on social identity, the more likely it is that authors will only hear about papers written by those that share their social identity. Note that even with no homophily, the ratio is slightly above 1 because of self citation; a majority group member citing themselves is an instance of a majority member citing a majority member, which occurs with probability 1.

Figure 2b shows the corresponding graph for the ratio of majority to minority member's papers cited in new papers coming out by minority members. Just as majority members are more likely to cite majority members with greater homophily, minority members are less likely to cite majority members (and more likely to cite minority members) as homophily increases. Their citation ratio of majority to minority decreases as we increase homophily. In this case, the ratio also decreases as the minority group gets smaller—this is because self citation becomes more important as there are smaller numbers in a group.⁷ The results presented here are

⁷ In a model which excludes the two original authors from the accumulation of citations, results show that this decrease does not occur.

for a network of 100 people, so, for example, if the minority group is 10% of the community, any new paper by a minority group member has a 1 in 10 chance of being by the author of one of the original papers.

The important relationship here is that, the more homophilic the network is, the larger the citation ratio gap. This ratio gap then gives rise to the overall citation gap. If majority group members are more often citing majority group members, then as we increase the proportion of the majority in the population, they will overall be cited more often than the minority group members.

3.2 Searches

One thing that could potentially mitigate this effect is the importance of search engines like Google Scholar. To investigate this possibility, we consider a second model, which is identical to the first except that authors have some chance of finding a paper by looking through their network and some chance of finding a paper to cite via internet search. For the results presented below, there is a 30% chance to cite someone you hear about through your network versus a 70% chance to go looking online. When an academic uses a search engine, the chance they find a paper is influenced by how many citations it already has, leading to a ‘rich get richer’ or ‘Matthew effect’ (West & Bergstrom, 2021). In general, a paper with more citations shows up on an earlier page in the search, and the earlier a paper appears in the search the more likely it is that someone will cite it.⁸

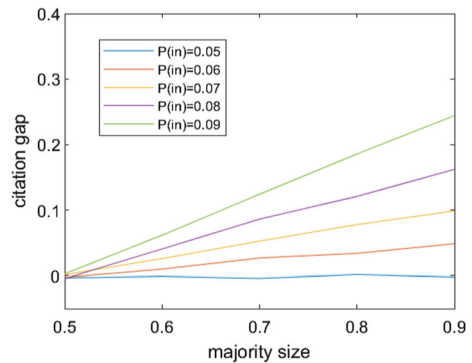
The possibility of internet searches dampens the effect both homophily and group size have on citation gaps, but not as much as one might expect. For contrast, the green line representing high homophily ranges from 0 to about .24 in Fig. 3 versus in Fig. 1, where it ranged from 0 to about .34. This means it is not the case that 30% of the time the majority members are cited more often due to network effects, whereas 70% of the time people are citing each social identity group with equal probability based on internet searches. If that were true, we would expect a 70% reduction in the majority’s advantage, from .34 to about .1, or 70% of their excess citations gone.⁹

Part of the reason the majority group maintains much of its advantage regarding citation chances is that the search algorithm unintentionally encodes some of its structural advantage. That is, because the chance you find a paper is influenced by how many citations it already has, and the majority group are receiving more citations because news of their paper spreads to more people in the network, the

⁸ For the results presented here, the likelihood to cite a paper based on a search, p is determined by the page, g , such that $p = .95^{10g}$ and $g = 10 - \frac{10c}{10+c}$, where c is the number of citation a paper has accumulated. Nothing depends on these particular equations, they merely capture the observation that more citations lead a paper to be on an earlier page, consequently making it more likely to be cited.

⁹ One reason for this is that people are more likely to find a paper to cite through their network than by searching through pages of internet searches, so slightly more than 30% of citations come from looking through the network. For example, again looking at the extreme of high homophily and low representation, data from these simulations shows that around 40% of citations come from looking through the network. However, this cannot fully explain the results in Fig. 3. If it did, we would expect about 60% reduction in the majority’s advantage (from .34 to .14).

Fig. 3 Homophily and group size affect citation gaps, with searches



chance that you find a paper by a majority group member is greater than the chance you find a paper by a minority group member when doing an internet search.

3.3 Different publication and citation rates

Thus far, we have considered models where the only difference social identity makes to one's chances of being cited is via place in the academic network. This section considers how these network effects can work in combination with some empirical observations regarding other ways social identity matters.

First, we incorporate the fact that men tend to publish more often than women (see Hesli & Lee, 2011 and references therein). How much more they publish varies field to field, but here we present results for when men publish 1.5 times as often as women. The rest of the model is the same as in Sect. 3.2, with academics both looking through their networks and using internet searches. We still track citations accumulating to one paper by a man and one paper by a woman, except that the potential papers coming out that could cite these two papers are more often authored by men.

When we include this factor, as one might expect, citation gaps increase in size. As Fig. 4a shows, even when there is no homophily and there are equal numbers of men and women, there is still a small citation gap. Again, as we increase homophily, the relative size of the two social identity groups makes more and more difference. The size of citation gaps starts higher than in the previous two models, and reaches levels somewhere between the first and second models for high homophily and smaller relative size of the minority group.

The results up until now show that citation gaps can arise even without any biases, e.g. thinking that a minority group member's work is less worthwhile. Now we consider how bias can interact with these structural causes. We can include bias in the model by simply saying that when a majority group member finds a paper by a minority group member, there is a chance they will not cite it.¹⁰ Figure 4b shows how just a 5% chance of failing to cite a minority group member affects citation

¹⁰ Of course, minority group members may also have some chance of not citing other minority group members. This will affect overall citation rates in a similar way.

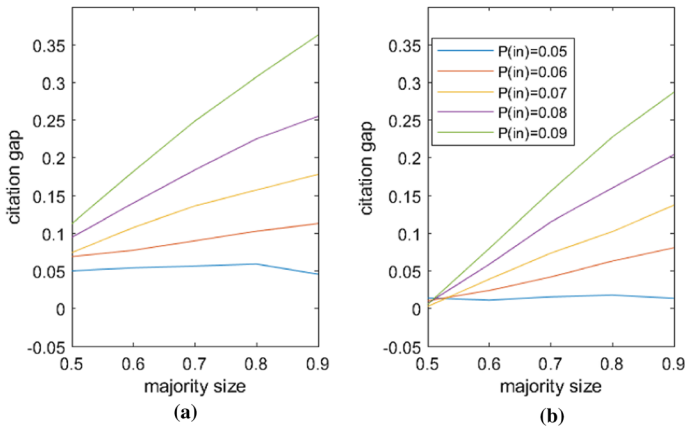


Fig. 4 Homophily and group size affect citation gaps, **a** with searches and different publication rates, and **b** with searches and bias

gaps. When we compare Fig. 4a, b, we can see that bias affects citation gaps in a qualitatively similar way to differences in publication rates when the minority group is small and network is homophilic, but has less effect in what we might think of as the ‘baseline’ cases where there is no homophily or where groups have equal representation.¹¹

3.4 Other potential factors

Of course, in real academic communities, these effects are not cleanly separable, which is, as I argued, one of the reasons mathematical models are so useful. However, it is important to remember that the factors discussed in Sects. 3.1–3.3 stand in complicated causal relationships with each other and to other factors that have not been considered. While the structural causes discussed above are important, they should not be taken to explain every facet of citation patterns. For instance, one observation that cannot be explained by appealing just to the network structure of the community is that, even though citation gaps decrease as there are more women in a subfield, citation gaps still exist in some subfields which are female dominated, and in fact, in many fields citation gaps been increasing as representation of women increases (Huang et al., 2020; Dworkin et al., 2020).

Some of these phenomena could be due to bias, or a publication gap (things discussed in the previous sections). Or, there could be any number of other factors at play. For instance, we might consider the fact that as more women enter fields, they are not necessarily given sought-after research or tenure track positions, but may instead often do more of the teaching or lab management labor (Dietrich & Tambasco, 2007; Eagly, 2020). We might also consider prestige bias—being considered one of the ‘big names’ in your field or being employed at a prestigious

¹¹ The magnitude of the effect of bias or publication rates on the citation gap may be different, depending on the values those parameters take.

institution can affect your citation chances, and many other aspects of your career (Morgan et al., 2018). We might also think that access to mentors, which has been found in a lot of cases to be a barrier for many social identity groups (Milkman et al., 2015; Martinez-Cola, 2020), is relevant as it would help improve quality of work, as well as likelihood of publication and citation. Citations may also be playing a sort of signaling role (to signal to the referees your competence in the field) and that referees are going to be looking for already well-cited people or papers, increasing citations to those who have already been well-cited and discouraging citation of others.¹² There are still other factors like the dropout gap, differences in funding, gatekeeping of particular subject areas, issues of who gets asked to speak at conferences, who publishes in higher prestige journals, and so on.

While I have argued that the structure of the community is likely partially responsible for citation gaps, it should hopefully be clear that this does not mean that these other factors are not relevant. That said, these structural causes are an important factor to consider, which may be hard to see from data. Additionally, as I will argue next in Sect. 4, the existence of these structural causes also has other implications for research communities.

4 The importance of structure

There are several reasons why illuminating these structural causes is important. First, and perhaps most obviously, identifying the cause(s) of an outcome is important if we want to consider possible interventions for changing that outcome. Second, that this type of structural cause may be at play has implications for those putting forth the ‘merit defense’ described at the end of Sect. 2—you should be concerned that you are not citing the best papers if it is likely the best papers are not all reaching you. Section 5 will return to both of these issues. The purpose of this section is to discuss a perhaps somewhat less obvious reason we should care that structural causes are playing a role, which is that these structural causes can generate feedback loops that increase inequity over time. I will support this claim using the particular example of citation gaps, though the general argument is more wide-reaching.

In studying behaviors of complex systems, such as academic communities, it is important to be aware of potential feedback loops, where initial inequities may feed back into greater and greater asymmetries over time. The general importance of feedback loops in studying complex systems has been noted by many, e.g. Mitchell (2009). This is important to keep in mind when using models to study the phenomena of interest, as you might run into a problem of *back-reaction*, where the features assumed to be in the background of a process under study (and either

¹² Thanks to Chris Weaver and Riet Van Bork for discussions on this topic. See also Rubin and Schneider (2020) for a discussion of the role signaling can play in the context of assigning priority for scientific discoveries.

described at low-fidelity, or assumed to be fixed) interact with the foreground features the model is designed to investigate.¹³ For example, in the context of the models provided in Sect. 3, we might say the foreground includes citation dynamics, while the network structure is in the background as a feature assumed to be fixed. However, if we want to know about, for instance, the longer term behavior of the system, it might be troubling that network structure in the model is not also influenced by the accumulation of citations (which can vary dramatically across the members of the community), but is instead fixed in the background.

More generally, back-reaction is a concern for these sorts of simple models, which hold fixed otherwise important factors in the background. These background factors may themselves change as a result of some proposed intervention, meaning we must then reanalyze the model in light of the new background. This is even more important when there is potential for feedback loops, as the changes in the background compound over time and more quickly depart from what was originally assumed in the model.

The damages of failing to account for back-reaction have been famously observed in the case of artificial intelligence algorithms approving credit or predicting recidivism, where it has been shown that by ignoring racial inequalities, these algorithms created a feedback loop where initial inequalities lead to future greater inequalities, not only entrenching but amplifying current injustices (O’Neil, 2016). Similarly, overlooking marginalization according to social identity may have major consequences when evaluating how academic communities (ought to) function. While we should be sensitive to any (potential) biases in these cases, paying attention to structural causes takes on increased importance because of the way they can compound over time. Section 4.2 provides an example of how this can occur.

This is, of course, not to say that that all models of academic communities must include considerations of social identity. Modelers always have to simplify, and we learn a lot from models which ignore important aspects of social identity. The point is that we ought to be careful in forming opinions on, arguing for, or enacting policies/reforms without doing due diligence in investigating how they impact already marginalized groups, especially when relying on models where we have reason to suspect a problem of back-reaction.

4.1 Citation gaps and peer review

There are many well-known problems with peer review, and suggestions for how to rectify these problems are everywhere. Many (e.g., Kriegeskorte, 2012; Nosek & Bar-Anan, 2012; Teixeira da Silva & Dobránszki, 2015; Vale, 2015; Heesen & Bright, 2019) advocate abolishing pre-publication peer review and replacing it with a system where academics publish papers by posting them to an archive, where they are reviewed post-publication (similar to what already happens in parts of mathematics and physics). While details of these proposals differ, under this sort

¹³ Thanks to Mike Schneider for discussions on this point.

of reform, feedback is no longer given by an assigned reviewer, but is performed by members of the community as they see fit, post-publication. Those proposing these sorts of reforms often ignore issues of marginalization and under-representation. Heesen and Bright (2019) are an exception; they engage with these issues in a significant way. As such, this section will focus on their argument for abolishing pre-publication peer review.

Heesen and Bright (2019) argue that abolishing pre-publication peer review is a sort of ‘Pareto improvement’ or ‘weakly dominant strategy,’ in the sense that for any standard for evaluating their proposal compared to the current system, abolishing pre-publication peer review is at least as good or better. This makes for a nice argument, because then one does not have to worry about how to weigh one factor against another. They argue there are a number of ways abolishing pre-publication peer review would be beneficial, including faster sharing of results, more efficient time allocation for scientists, decrease in gender skew of publications, and so on. There are also a number of aspects of science for which they argue the evidence is that this reform will basically make no difference, including ‘epistemic sorting’ (e.g., determining which articles are high quality), malpractice or fraud detection, effects of credit incentives, etc. And finally, they consider potential difficulties for their proposal. One—a guarantee for outsiders, e.g. journalists—they dismiss, arguing that peer review is not as much of a guarantee of quality as we think it is, and the other—a runaway Matthew effect—they take more seriously but think is highly speculative.

This is a very short overview, but the point of this section is not to argue for or against abolishing pre-publication peer review. Instead, it will focus on two issues related to citation gaps, which ought to be taken into account for any weighing of potential effects of this sort of reform to the peer review process.

First, current evidence does not support the claim that decreasing the gender productivity gap would be a benefit of abolishing pre-publication peer review. At best, this reform would make no difference for this factor. The observation that men publish more is often explained in terms of women anticipating bias in the peer review process. Whether or not there is bias in peer review, women certainly expect there to be (Lee, 2016). So, women spend more time on each of their papers to ensure its quality is above whatever threshold for publication, meaning they write fewer papers overall.¹⁴ So, Heesen and Bright (2019) reason that their proposed reform will ease the productivity gap as women will be primarily concerned to publish when their paper meets their own, rather than the community’s, standards.

Yet, this claim seems implausible as women (ought to) anticipate biased evaluation of their work post-publication as well; there are various stages post-publication at which gender bias can occur, including in the accumulation of citations (and therefore in the ‘impact’ or ‘visibility’ of work, according to most measures), as well as in hiring and promotions (Steinpreis et al., 1999; Sarsons, 2017), notoriety of researchers (Adams et al., 2019), and the uptake of ideas (Hofstra et al., 2020). Women are likely at least as aware of these things as they are

¹⁴ See Bright (2017) for a decision theoretic model supporting this argument.

of any potential bias in peer review, and so they will still be concerned to meet the community's (perceived to be biased) standards if we abolish pre-publication peer review.

Second, citation gaps are relevant to thinking about the runaway Matthew effect, which is a feedback loop whereby initial inequities lead to greater inequities over time. As Heesen and Bright (2019) explain:

The scientific community allocates the resources necessary for future work on the basis of its recognition of past performance. So if there is excess reward for some and unfair passing over of others at the present stage of inquiry, this will ramify through to future rounds of inquiry, misallocating resources to people whose accomplishments do not fully justify their renown. (p. 23)

Heesen and Bright consider this runaway Matthew effect carefully, but ultimately conclude that: "Our present thought is that this is a very speculative objection, and there is no empirical evidence to back up the claim that eliminating pre-publication peer review will have dire consequences in this regard." (p. 24) That is, they argue there is no evidence eliminating pre-publication peer review will make the runaway Matthew effect any worse than it is under our current peer review system.

While I agree with Heesen and Bright on the need for further empirical evidence, and that this runaway effect might possibly be counteracted by sufficient additional reforms to protect those most likely to be impacted, I will make a push for this being not so very speculative and an issue that deserves greater attention. In doing so, I will extend the reasoning put forth by Heesen and Bright (2019) and make the worry they describe more concrete. So, in one sense, I agree with Heesen and Bright, and am elaborating on their work. However, I will also provide an argument against their ultimate conclusion that abolishing pre-publication peer review is weakly dominant: the runaway Matthew effect is likely to be worse under their proposed reform than under the current peer review system.¹⁵

Here is a quick argument for this conclusion in the context of citation gaps. Unlike in the peer review system, in order for a paper to receive post-publication peer review, it must first be seen by others. As discussed throughout this paper, we have evidence that work by women and minority group members is less likely to be seen by others in an academic community. Furthermore, social positioning ultimately depends, in part, on the previous impact of a researcher's work - those who have accumulated citations or other recognition for their work are more likely to be central in a network (Yan & Ding, 2009; Hoffmann et al., 2016). Therefore, researchers or groups of researchers who are initially on the peripheries may accumulate less prestige, creating a runaway effect where they are pushed more and more to the peripheries over time. By contrast, under the peer review system, there is a mechanism by which an academic can gain reputation that is much less dependent on their centrality in the network, i.e. publishing a paper in a reputable journal, which would cut against the runaway Matthew effect.

¹⁵ Thanks to a reviewer for pushing me to clarify the points of agreement and disagreement.

Section 4.2 will provide more detail and support for this case, developing a model to demonstrate how the sort of system Heesen and Bright (2019) propose encourages a runaway Matthew effect. I will, because of the focus of this paper, be making this argument in the context of citation gaps according to social identity, but it should be easy to see how a similar argument could be given to raise concerns more generally about any situation where there are some members of a research community starting out on the peripheries.

To be clear, I do not take myself to be providing a definitive argument against abolishing pre-publication peer review. Instead, I take myself to be identifying a mechanism by which a runaway Matthew effect could plausibly occur—a feedback loop involving social positioning and visibility of work—which we have reason to believe is less influential in our current system of peer review compared to under the proposed reform. Therefore, given current evidence, we should conclude abolishing pre-publication peer review is not weakly dominant; there is at least one standard of evaluation under which the reform would likely be worse than the current system. I take this as both as a note of caution for those attempting to replace our current peer review system, and as an impetus for further work, e.g. comparing the magnitude of the runaway affect under different peer review systems, identifying factors that could impede it or prevent it entirely, and determining how to suitably modify a post-publication peer review system to protect those affected by the runaway Matthew effect.

4.2 A runaway Matthew effect

In order to demonstrate the existence and nature of the feedback loop behind this runaway Matthew effect, I present a model capturing the likelihood that researchers engage with each others' work. Here is the basic set-up: In order for a paper to receive post-publication peer review, it must first be seen by others. This can be achieved in a variety of ways: word of mouth, being promoted on an archive, the archive posting being re-tweeted, and so on (Vale, 2015). The likelihood that someone will read and/or share a paper by another researcher depends on many factors, e.g. whether they know that person personally, whether they are familiar with their work, whether they are employed by the same institution, and the reputation of that person. This likelihood of one researcher sharing another's work can be represented by a weighted directed edge in a network. Directed edges point from one researcher to another, while the weight determines the chance that the first researcher will share or engage with a paper by the second researcher.

This formalism is well suited to identify the feedback loop described above because we can track how these weights change over time. For instance, more people may become familiar with a person's work as it is shared more often or someone can gain reputation from publishing papers with high impact (i.e. that are cited often). These sort of factors will increase the weight of edges pointing toward that person, meaning that others are more likely to share and review their work, leading to a further increase in weights, and so on.

The model was implemented as follows. First, weighted directed networks were formed in such a way as to be close to the networks in the previous models, taking

those networks to represent the starting point of a model investigating the possible effects of abolishing pre-publication peer review. So, we start with a multi-type random network and replace each undirected edge with two directed edges of weight 1. Then, since, presumably, there is some chance you will engage with someone's paper even if you are not in regular communication with its author, everywhere there previously was no link, two directed links with a small weight of .1 were added.¹⁶

These weights will change as papers are shared. For instance, more people may become familiar with a person's work as it is shared more often or someone can gain reputation from publishing papers with high impact (i.e. that are cited often). These sort of factors will increase the weight of edges pointing toward that person, meaning that others are more likely to share and review their work. In the model, each researcher starts with one paper posted to the archive. Each round, researchers engage with five papers and, in addition, have a small chance to publish an additional paper.¹⁷ Papers are engaged with based on the weights the person engaging puts on each of the other researchers, i.e. the more weight they put on a paper's author, the more likely they are to engage with that paper.¹⁸ Each time a paper is engaged with, its author gains a small amount of reputation, i.e. the edges pointing from each other researcher in the community to that author are increased by a small amount.¹⁹

To quantify and be able to measure asymmetries due to social identity, we can measure *centrality*, which captures how central a person is to the network, i.e. how well-connected they are. In particular, we will measure a *centrality gap*, or how much more well-connected majority group members are, by looking at their average centrality compared the average centrality of minority group members. The results in Fig. 5 use a simple measure of centrality based on the weights of incoming arrows:

$$C_i = \sum_j w_{ji}$$

The centrality of a node i , C_i , is found by summing up the weight each other person in the network puts on i (i.e., by summing up the weights of all the arrows pointing to node i , which measure how likely those people are to engage with i 's work).²⁰

Figure 5 shows the centrality gap at the beginning of simulations versus the centrality gap after 50 rounds of 20 researchers publishing and engaging with

¹⁶ The method of forming these networks should not make a difference to the results, as long as we form a homophilic weighted directed network.

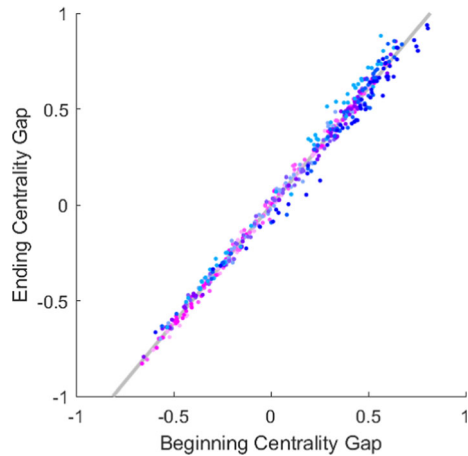
¹⁷ Turn order each round is determined randomly, and the chance to publish an additional paper was set to 10% for the results below.

¹⁸ The five papers were chosen by a weighted random sampling procedure. It is possible for a researcher to engage with a paper in multiple ways, e.g. by commenting on it and by sharing it with others.

¹⁹ For the results presented here, this increase is .005, but the exact amount does not change the qualitative results. Additionally, weights were normalized at the start of the simulation and each time they evolve, so that the sum of each person's outgoing arrows is one.

²⁰ Results are similar for other measures of centrality, e.g. closeness centrality, which is based on shortest path lengths between nodes.

Fig. 5 Beginning and ending centrality gap, for various levels of homophily and sizes of the majority group



others' work.²¹ A positive centrality gap means the majority group members are, on average, more central to the network, meaning their work is more likely to be engaged with. A negative centrality gap then corresponds to the situation where the minority group's work is more likely to be engaged with. There are two important features to note. First, and most importantly for showing the runaway Matthew effect, is that the slope of the regression is greater than one. This means that if you start with a positive number, you expect it to increase by the end, e.g. .5 would increase to .6 with the slope of roughly 1.2 in Fig. 5. Similarly, if you begin with a negative centrality gap (where the minority group is favored), you expect to get a more negative number by the end (i.e. the initially favored minority group is more favored at the end). That is, starting with an advantage tends to create a feedback loop whereby the advantaged group will become more advantaged as time goes on.

The second thing to note has to do with the size of the minority group. The data points in Fig. 5 are color coded according to majority group size, and Fig. 6 separates out these data points to make comparison easier. We can see from Fig. 6 how minority groups can be disadvantaged in terms of centrality, meaning the feedback loops will generally serve to increase their disadvantage over time. When both groups are evenly represented—Fig. 6a—the data points are evenly spread out. That is, although there can be differences in centrality between social identity groups, which then feeds back into greater differences over time, this is not expected to disadvantage one group more than the other. To look at the case where the minority group is very small—10%, as in Fig. 6c—the data points are more clustered in the upper right portion, where the minority group starts off disadvantaged and ends up more disadvantaged at the end. Figure 6b, where the majority is 70% of the total group, is intermediate between these cases. The

²¹ To get an estimate of how the process is expected to go, for each combination of parameters, 50 networks were formed randomly and five simulations were run on each of these networks. Data points in Fig. 5 represent each network that was formed, averaging over the five simulations. Results are very similar if instead each simulation is considered a data point.

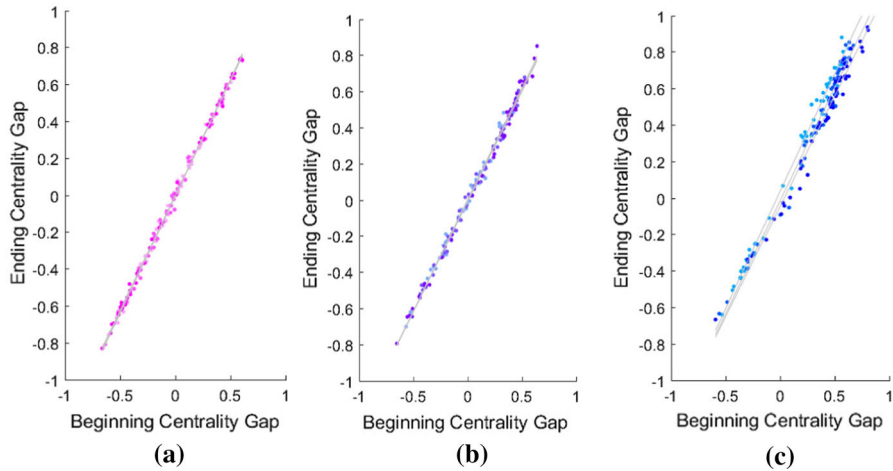


Fig. 6 Beginning and ending centrality gap, for various levels of homophily and **a** 50% majority, **b** 70% majority, and **c** 90% majority. Data points are coded according to different levels of $p(in)$, ranging from .05 (lightest) to .09 (darkest)

runaway Matthew effect tends to disfavor minority groups, and the smaller the minority group is, the more likely it is to be disadvantaged.

5 Discussion

We have seen that citation gaps can arise due to the structure of academic communities. In homophilic networks, people tend to cite papers written by authors sharing their own social identity, which disadvantages those in minority or underrepresented groups. We have also seen that internet searches may unintentionally encode some of the structural advantage to the majority group, and that phenomena like different publication rates and bias against minority group members may interact with the structure of the community to increase the size of citation gaps.

The existence of these structural causes has implications not only for how we evaluate proposals for things like altering the peer review process, but also for proposals to address citation gaps themselves. One common proposal is that we ought to require or encourage scholars to cite women and minorities, for example by increasing maximum word counts to allow people to cite women or minorities. This may help to combat citation gaps due to bias, by encouraging people to rethink why they have not cited particular papers, but will not necessarily be effective if the lack of citation is due to not having heard of work by members of certain social identity groups. Authors could spend more time conducting searches, but there is bound to be resistance if those authors feel they have already done due diligence in finding

appropriate citations and, in general, proposing fixes for structural problems in terms of individual actions can be fighting an uphill battle.²²

In fact, these sorts of individual level fixes are often met with some comment of the form: “We should not have quotas for our bibliographies”. This type of response likely rests on the attitude behind the merit defense, where people argue that they already cite the best/most relevant papers. In most of the models in Sect. 3 the merit defense is true by design, yet citation gaps emerge. However, there would be reasons to want to eliminate citation gaps even if everyone was, individually, unbiased in their citation practices. In addition to being unfair, the existence of citation gaps is epistemically undesirable as good work goes unnoticed.²³ To address structural causes of citation gaps, supporting projects that make the work of marginalized groups more visible by creating public lists or otherwise promoting their work may be more effective, and can perhaps prevent defensive reactions from those who feel as though individuals are being unfairly blamed or burdened.

The results presented here also indicate that overlooking how structural features affect certain social identity groups may have major consequences when evaluating how research communities (ought to) function. Section 4.2 provided reason to expect a runaway Matthew effect under a post-publication peer review system, which has been argued by many to be beneficial in terms of efficiency and knowledge production. Switching to this system, which increases the importance of social positioning for recognition of work, is likely to encourage a runaway Matthew effect. While Heesen and Bright (2019) are right that we ought to substantiate these claims with empirical evidence, in the meantime, there is reason to believe that under at least one standard of evaluation abolishing pre-publication peer review would be worse than our current system.

Of course, there are many relevant factors not discussed in the models here. We might also be interested how these structural features interact not just with psychological biases and differing publication rates, but with other phenomena such as prestige bias. Additionally, this paper did not discuss further potential feedback effects, for example, how being pushed further to the peripheries might affect scholars’ quality of work or productivity. If work by women and minorities is less likely to be widely seen, there is reason to think it would also receive less critical engagement when researchers are allowed to pick which papers to engage with in the review process. Therefore, one might argue that the quality of their work may suffer, be less impactful, and therefore their reputation would suffer, leading to their work being less likely to be engaged with and receive critical engagement in the future. There is further work to be done investigating these possibilities.

²² This is not to disparage those who make a concerted effort to be cognizant of citation gaps when compiling bibliographies, or journals which have made efforts to encourage authors to be cognizant. As emphasized, this paper does not deny the existence of individual level causes, such as implicit bias, which these efforts may counterbalance effectively.

²³ See Schneider et al. (2020) for an argument that exchange of ideas between social identity groups is epistemically important.

Acknowledgements Many thanks to Mike Schneider for ideas, discussions, criticisms, and so on. The initial models here are similar to those developed in a paper co-authored with Mike (Priority and Privilege in Scientific Discovery). Thanks to Justin Bruner, the Spring 2021 Fellows at the Pitt Center for Philosophy of Science, and two reviewers for comments on the paper. Also, thanks to audiences at University of Groningen, the Women in Academia conference at UC Irvine, and the Pitt Center for Philosophy of Science lunchtime talk for feedback. This material is based upon work supported by the National Science Foundation under Grant No. 2045007.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

References

- Adams, J., Brückner, H., & Naslund, C. (2019). Who counts as a notable sociologist on wikipedia? Gender, race, and the “professor test”. *Socius*. <https://doi.org/10.1177/2378023118823946>.
- Aksnes, D. W., Rorstad, K., Piro, F., & Sivertsen, G. (2011). Are female researchers less cited? A large-scale study of Norwegian scientists. *Journal of the American Society for Information Science and Technology*, 62(4), 628–636.
- Beaudry, C., & Larivière, V. (2016). Which gender gap? Factors affecting researchers’ scientific impact in science and medicine. *Research Policy*, 45(9), 1790–1817.
- Botts, T. F., Bright, L. K., Cherry, M., Mallarangeng, G., & Spencer, Q. (2014). What is the state of blacks in philosophy? *Critical Philosophy of Race*, 2(2), 224–242.
- Bright, L. K. (2017). Decision theoretic model of the productivity gap. *Erkenntnis*, 82(2), 421–442.
- Brown, N. E., & Samuels, D. (2018). Beyond the gender citation gap: Comments on dion, sumner, and mitchell. *Political Analysis*, 26(3), 328–330.
- Bruner, J. P. (2019). Minority (dis) advantage in population games. *Synthese*, 196(1), 413–427.
- Bruner, J. P., & O’Connor, C. (2015). Power, bargaining, and collaboration. In T. Boyer, C. Mayo-Wilson, & M. Weisberg (Eds.), *Scientific collaboration and collective knowledge*. Oxford: Oxford University Press.
- Cameron, E. Z., White, A. M., & Gray, M. E. (2016). Solving the productivity and impact puzzle: Do men outperform women, or are metrics biased? *BioScience*, 66(3), 245–252.
- Corley, E. A., & Sabharwal, M. (2010). Scholarly collaboration and productivity patterns in public administration: Analysing recent trends. *Public Administration*, 88(3), 627–648.
- Curran, S., Jackson, M. O., & Pin, P. (2009). An economic model of friendship: Homophily, minorities, and segregation. *Econometrica*, 77(4), 1003–1045.
- del Carmen, A., & Bing, R. L. (2000). Academic productivity of African Americans in criminology and criminal justice. *Journal of Criminal Justice Education*, 11(2), 237–249.
- Di Vaio, G., Waldenström, D., & Weisdorf, J. (2012). Citation success: Evidence from economic history journal publications. *Explorations in Economic History*, 49(1), 92–104.
- Dietrich, M. R., & Tambasco, B. H. (2007). Beyond the boss and the boys: Women and the division of labor in drosophila genetics in the united states, 1934–1970. *Journal of the History of Biology*, 40(3), 509–528.
- Dion, M., & Mitchell, S. M. (2012). *Gender, participation, and citations: Comparing peace science, political methodology, state politics, and IPE conferences*. Poster presented at the annual conference of the Peace Science Society, Savannah, GA.
- Dion, M. L., Sumner, J. L., & Mitchell, S. M. (2018). Gendered citation patterns across political science and social science methodology fields. *Political Analysis*, 26(3), 312–327.
- Dworkin, J. D., Linn, K. A., Teich, E. G., Zurn, P., Shinohara, R. T., & Bassett, D. S. (2020). The extent and drivers of gender imbalance in neuroscience reference lists. *Nature Neuroscience*, 23(8), 918–926.
- Eagly, A. H. (2020). Do the social roles that women and men occupy in science allow equal access to publication? *Proceedings of the National Academy of Sciences*, 117(11), 5553–5555.
- Esarey, J., & Bryant, K. (2018). Are papers written by women authors cited less frequently? *Political Analysis*, 26(3), 331–334.

- Ferber, M. A. (1986). Citations: Are they an objective measure of scholarly merit? *Signs: Journal of Women in Culture and Society*, 11(2), 381–389.
- Ferber, M. A. (1988). Citations and networking. *Gender & Society*, 2(1), 82–89.
- Ferber, M. A., & Brün, M. (2011). The gender gap in citations: Does it persist? *Feminist Economics*, 17(1), 151–158.
- Golub, B., & Jackson, M. O. (2012). How homophily affects the speed of learning and best-response dynamics. *The Quarterly Journal of Economics*, 127(3), 1287–1338.
- Håkanson, M. (2005). The impact of gender on citations: An analysis of college & research libraries, journal of academic librarianship, and library quarterly. *College & Research Libraries*, 66(4), 312–323.
- Heesen, R., & Bright, L. K. (2019). Is peer review a good idea? *The British Journal for the Philosophy of Science*, 72(3), 635–663. <https://doi.org/10.1093/bjps/axz029>
- Hesli, V. L., & Lee, J. M. (2011). Faculty research productivity: Why do some of our colleagues publish more than others? *PS: Political Science & Politics*, 44(2), 393–408.
- Hoffmann, C. P., Lutz, C., & Meckel, M. (2016). A relational altmetric? Network centrality on research Gate as an indicator of scientific impact. *Journal of the Association for Information Science and Technology*, 67(4), 765–775.
- Hofstra, B., Kulkarni, V. V., Galvez, S.M.-N., He, B., Jurafsky, D., & McFarland, D. A. (2020). The diversity-innovation paradox in science. *Proceedings of the National Academy of Sciences*, 117(17), 9284–9291.
- Huang, J., Gates, A. J., Sinatra, R., & Barabási, A.-L. (2020). Historical comparison of gender inequality in scientific careers across countries and disciplines. *Proceedings of the National Academy of Sciences*, 117(9), 4609–4616.
- Knobloch-Westerwick, S., Glynn, C. J., & Huge, M. (2013). The matilda effect in science communication: An experiment on gender bias in publication quality perceptions and collaboration interest. *Science Communication*, 35(5), 603–625.
- Kriegeskorte, N. (2012). Open evaluation: A vision for entirely transparent post-publication peer review and rating for science. *Frontiers in Computational Neuroscience*, 6, 79.
- Larivière, V., Ni, C., Gingras, Y., Cronin, B., & Sugimoto, C. R. (2013). Bibliometrics: Global gender disparities in science. *Nature News*, 504(7479), 211.
- Leahey, E., Crockett, J. L., & Hunter, L. A. (2008). Gendered academic careers: Specializing for success? *Social forces*, 86(3), 1273–1309.
- Lee, C. J. (2016). Revisiting current causes of women's underrepresentation in science. In J. Saul & M. Brownstein (Eds.), *Implicit bias and philosophy, Volume 1: Metaphysics and epistemology* (pp. 265–282). Oxford, UK: Oxford University Press.
- Lee, C. J., Sugimoto, C. R., Zhang, G., & Cronin, B. (2013). Bias in peer review. *Journal of the American Society for Information Science and Technology*, 64(1), 2–17.
- Long, J. S. (1992). Measures of sex differences in scientific productivity. *Social Forces*, 71(1), 159–178.
- Maliniak, D., Powers, R., & Walter, B. F. (2013). The gender citation gap in international relations. *International Organization*, 67(4), 889–922.
- Martinez-Cola, M. (2020). Collectors, nightlights, and allies, oh my! White mentors in the academy. *Understanding and Dismantling Privilege*, 10(1), 25–57.
- Merritt, D. J. (2000). Scholarly influence in a diverse legal academy: Race, sex, and citation counts. *The Journal of Legal Studies*, 29(S1), 345–368.
- Milkman, K. L., Akinola, M., & Chugh, D. (2015). What happens before? A field experiment exploring how pay and representation differentially shape bias on the pathway into organizations. *Journal of Applied Psychology*, 100(6), 1678.
- Mitchell, S. D. (2009). *Unsimple truths: Science, complexity, and policy*. Chicago: University of Chicago Press.
- Mitchell, S. M., Lange, S., & Brus, H. (2013). Gendered citation patterns in international relations journals. *International Studies Perspectives*, 14(4), 485–492.
- Morgan, A. C., Economou, D. J., Way, S. F., & Clauset, A. (2018). Prestige drives epistemic inequality in the diffusion of scientific ideas. *EPJ Data Science*, 7(1), 40.
- Nosek, B. A., & Bar-Anan, Y. (2012). Scientific utopia: I. Opening scientific communication. *Psychological Inquiry*, 23(3), 217–243.
- Nosek, B. A., Graham, J., Lindner, N. M., Kesebir, S., Hawkins, C. B., Hahn, C., Schmidt, K., Motyl, M., Joy-Gaba, J., Frazier, R., et al. (2010). Cumulative and career-stage citation impact of social-

- personality psychology programs and their members. *Personality and Social Psychology Bulletin*, 36(10), 1283–1300.
- O'Connor, C. (2019). *The origins of unfairness: Social categories and cultural evolution*. Oxford: Oxford University Press.
- O'Connor, C., Bright, L. K., & Bruner, J. P. (2019). The emergence of intersectional disadvantage. *Social Epistemology*, 33(1), 23–41.
- O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. New York: Broadway Books.
- Østby, G., Strand, H., Nordås, R., & Gleditsch, N. P. (2013). Gender gap or gender bias in peace research? Publication patterns and citation rates for journal of peace research, 1983–2008. *International Studies Perspectives*, 14(4), 493–506.
- Powell, A., Hassan, T. M., Dainty, A. R., & Carter, C. (2009). Note: Exploring gender differences in construction research: A European perspective. *Construction Management and Economics*, 27(9), 803–807.
- Rubin, H., & O'Connor, C. (2018). Discrimination and collaboration in science. *Philosophy of Science*, 85(3), 380–402.
- Rubin, H., & Schneider, M. D. (2020). Priority and privilege in scientific discovery. *Studies in History and Philosophy of Science Part A*, 89, 202–211.
- Sarsons, H. (2017). Recognition for group work: Gender differences in academia. *American Economic Review*, 107(5), 141–45.
- Schneider, M. D., Rubin, H., & O'Connor, C. (2020). Promoting diverse collaborations. In G. Ramsey & A. De Block (Eds.), *The dynamics of science: Computational frontiers in history and philosophy of science*. Pittsburgh: Pittsburgh University Press.
- Sen, M. (2018). Response to dion, sumner, and mitchell. *Political Analysis*, 26(3), 335–337.
- Slyder, J. B., Stein, B. R., Sams, B. S., Walker, D. M., Jacob Beale, B., Feldhaus, J. . J. ., & Copenheaver, C. A. (2011). Citation pattern and lifespan: A comparison of discipline, institution, and individual. *Scientometrics*, 89(3), 955–966.
- Stack, S. (2002). Gender and scholarly productivity: The case of criminal justice. *Journal of Criminal Justice*, 30(3), 175–182.
- Steinpreis, R. E., Anders, K. A., & Ritzke, D. (1999). The impact of gender on the review of the curricula vitae of job applicants and tenure candidates: A national empirical study. *Sex Roles*, 41(7–8), 509–528.
- Teixeira da Silva, J. A., & Dobránszki, J. (2015). Problems with traditional science publishing and finding a wider niche for post-publication peer review. *Accountability in Research*, 22(1), 22–40.
- Vale, R. D. (2015). Accelerating scientific publication in biology. *Proceedings of the National Academy of Sciences*, 112(44), 13439–13446.
- Wang, Y. S., Lee, C. J., West, J. D., Bergstrom, C. T. & Erosheva, E. A. (2019). Gender-based homophily in collaborations across a heterogeneous scholarly landscape. 1–33.
- Ward, K. B., Gast, J., & Grant, L. (1992). Visibility and dissemination of women's and men's sociological scholarship. *Social Problems*, 39(3), 291–298.
- West, J. D., & Bergstrom, C. T. (2021). Misinformation in and about science. *Proceedings of the National Academy of Sciences*, 118(15), 1–8.
- West, J. D., Jacquet, J., King, M. M., Correll, S. J., & Bergstrom, C. T. (2013). The role of gender in scholarly authorship. *PloS One*, 8(7), e66212.
- Yan, E., & Ding, Y. (2009). Applying centrality measures to impact analysis: A coauthorship network analysis. *Journal of the American Society for Information Science and Technology*, 60(10), 2107–2118.
- Yang, Y., Chawla, N. V., & Uzzi, B. (2019). A network's gender composition and communication pattern predict women' leadership success. *Proceedings of the National Academy of Sciences*, 116(6), 2033–2038.