The Mismeasure of Woman: The Epistemic and Social Impacts of Gendered Citation Practices

Darcy McCusker

A dissertation

submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

University of Washington

2021

Reading Committee:

Carole Lee, Chair

Conor Mayo-Wilson

William Talbott

Program Authorized to Offer Degree:

Philosophy

University of Washington

Abstract

The Mismeasure of Woman: The Epistemic and Social Impacts of Gendered Citation Practices

Darcy McCusker

Chair of the Supervisory Committee:
Dr. Carole Lee
Philosophy Department

This dissertation will explore the philosophical problems that arise from the finding that women are cited less frequently than men in a number of sciences, what I'll call *gendered citation practices*. Gendered citation practices have received little philosophical attention despite the importance of citations to the practice of science. Citations are used by scientists in every publication to connect their work to the community that they are a part of. Citations are public (to the extent that publications are public--minimally they are visible to members of the scientific community with institutional or other access to published articles), which makes them easier to track and quantify than many other types of barriers women face in science. Citations are also an increasingly important metric that are either currently used or raised as a potential metric for informing decisions about funding, hiring, and promotion. Disparities in citations have the potential for substantial impacts on the career trajectory of women in science, especially when

compounded with other gender disparities. The three papers of this dissertation discuss some of the philosophical issues that arise from gendered citation practices. The first paper argues that GPCs are a form of epistemic injustice, where women are denied access to uptake, a term borrowed from Longino. The second paper provides a sympathetic extension of Longino's views and argues that citations can be used to measure uptake and ensure that all members of the community are treated equitably. The third paper considers what would happen if citations were given even more weight in key decisions in scientists' careers and shows that women may be subject to even more epistemic injustice when grants are given to the most highly cited researchers.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

There are any number of incredible people who have helped me get to this point, and I feel honored to have had their help and support in getting to this point. If you've been a part of my graduate school experience, thank you.

I have so many people to thank in the philosophy department. Thank you to Melanie Tate, Sam Sumpter, and Michelle Pham for being there to vent to about the process of writing and everything else; having each of you to turn to throughout my graduate career has been invaluable and I know I will always have your friendship.

I am grateful to have been a part of the Neuroethics group; thank you to Tim Brown, Paul Tubig, Erika Versalovic, Marion Boulicault, Ishan Dasgupta, Kate MacDuffie, Andreas Schönau, and Eran Klein. I want to specifically thank Sara Goering for her role in bringing me into the Neuroethics group, but also for her mentorship, kindness, and willingness to host a great party; Sara has served in so many roles administratively, academically, and personally for me throughout my time at UW, and I am so appreciative of all the kinds of support I've gotten from her.

Justin Lawson, Olin Robus, and Jon Rosenberg taught logic with me my first quarter; since then, I've benefited enormously from their perspectives on my work and their friendships. Thank you to Ramona Marquez, Anna Bates, and Kayla Mehl-Hutchinson for indulging my need to frequently yell about things. Thank you to Janice Moskalik and Karen Emmerman for sharing their experiences with parenting in grad school and for showing me how to be an awesome parent and an awesome scholar. Thank you to the other members of the philosophy department who I have taken classes with, taught with or for, and learned with over the past seven years.

Thank you to Leah Ceccarelli, my GSR as well as my mentor for the STSS certificate for providing a perspective on my work from outside of the department. And thank you to Bill Talbott for serving on committee and helping me shape the epistemic injustice paper that was published in *Philosophy of Science*, and substantially modified for this dissertation.

Conor Mayo-Wilson has been a source of endless support throughout this process; I don't know how many hours of meetings we might have had over the past seven years, but it is surely more than I have had with any other faculty member. This is both because Conor is always willing to meet with me, even if I don't have much to say or any progress to report, and also because he enjoys telling long, elaborate stories to illustrate a small point. I will always be grateful for Conor's insistence that I was capable of succeeding, especially when I did not see it myself.

Thank you to Carole Lee, my chair, for guiding me through this process. I am indebted to Carole for her help finding sources, talking through conceptual confusions, and providing copies amounts of feedback on my writing. I am constantly impressed by Carole and everything she gets done in a day. I appreciate that this past year and a half has been especially challenging for anyone with kids, but Carole has always been willing to find time to talk when I need it. I hope Carole will be able to take an incredibly deserved break sometime soon.

Finally, I'd like to thank my family for their never ending support. Thank you to my parents and sister for your love and support. Most important to my success has been my family at home; Mike and Chris have both been there for me in every way I needed throughout this process, and Kieran helped more than I ever could've imagined. It turns out that it's much easier to solve your chronic procrastination problem when you really only have nap time to get things done. Kieran is my greatest accomplishment, but this dissertation is my second greatest.

Chapter 1. INTRODUCTION

1.1 GENDERED CITATION PRACTICES

This dissertation will explore the philosophical problems that arise from the finding that women are cited less frequently than men in a number of sciences. I will first summarize and explain the empirical evidence for and significance of these gender disparities, which I'll call *gendered citation practices*.

The careers of women scientists differ from those of scientists who are men in a number of ways. Women tend to publish less on average (Lariviere et al, 2013). Women leave science at higher rates than men at many points along their career paths and are underrepresented at the highest levels of academic positions (Ceci et al, 2014). Many of these differences are well documented and much discussed. One difference that has gotten comparatively less attention is that in many scientific fields, there is a gap between the average citation rates for men and women. This pattern is what I've termed *gendered citation practices* (GCPs).[1]

Fields where gendered citation practices occur include: astronomy (Caplar, Tacchella, and Birrer 2016), archaeology (Hutson 2002), communication science (Knobloch-Westerwick and Glynn 2013), ecology and evolutionary biology (Kelly and Jennions 2006), epidemiology (Schisterman et al. 2017), gynecological oncology (Hill et al. 2015), international relations (Maliniak, Powers, and Walter 2013), neuroscience (González-Álvarez and Cervera-Crespo 2017), neurosurgery (Khan et al. 2009), pediatric neurosurgery (Klimo et al. 2014), and psychology (Geraci, Balsis, and Busch 2015). This represents an incredibly diverse section of science; it includes social sciences, medical fields, biological sciences, and physical sciences. A

[1] Hutson (2002) uses the same term.

difference in citation rates or another citation metric is found in every one of these fields. In some fields, the effect is seen for the field as a whole, and in others, it is only seen within academic ranks, such as only for more junior faculty. However, many fields are left with citation disparities by gender that cannot be accounted for by controlling for any known characteristic of the authors or the work.

Table 1.1. contains a summary of the empirical work on citation rates by gender. The empirical work presented here forms the groundwork for each of the following papers. Each paper takes it as a starting place that gendered citation practices occur in the sciences discussed in the fields here, and I will simply refer to 'science' for convenience, when I am referring to the fields where gendered citation practices occur.

Gendered citation practices have received little philosophical attention despite the importance of citations to the practice of science. Citations are used by scientists in every publication to connect their work to the community that they are a part of. Citations are public (to the extent that publications are public--minimally they are visible to members of the scientific community with institutional or other access to published articles), which makes them easier to track and quantify than many other types of barriers women face in science. Citations are also an increasingly important metric that are either currently used or raised as a potential metric for informing decisions about funding, hiring, and promotion. Disparities in citations have the potential for substantial impacts on the career trajectory of women in science, especially when compounded with other gender disparities. The three papers of this dissertation discuss some of the philosophical issues that arise from gendered citation practices.

Table 1.1. Summary of Empirical Research on GCPs

Field	Methods	Effect size	Significance test performed and p values
Psychology	Randomly sampled 70 male and 70 female psychologists with tenure working at NRC Top 100 colleges and universities and calculated h indexes for each	Sampled women have h-index score of 17.17, compared to 21.55 for men	T test ($p < .05$)
Pediatric neurosurgery	Calculated h-index for all 312 practicing pediatric neurosurgeons in US	Women (n=52) had mean h-index of 8, men (n=260) had mean h index of 14	Test unspecified, ($p < .001$)
Neurosurgery	Calculated h-indexes for 1225 neurosurgery faculty at 99 (out of 101) US programs; only non-neurosurgery faculty were excluded	The median h-index was 11 for males ($n = 1144$) and 8 for females ($n = 81$). No significant effect was found when controlling for academic rank, however	2-tailed Mann-Whitney, ($p < .001$)
Gynecologic oncology	Calculated h-index for all 507 gynecologic oncology faculty at US medical schools	Men had a median h-index of 16, women had a median h-index of 8	Chi-squared test ($p < .001$)
Astronomy	ML-algorithm trained on papers by men used to calculate expected citations for papers written by women in five top astronomy journals from 1950-2015	Algorithm predicted women would receive 4% more citations than men but instead men receive 6% more citations; average citations per paper is 50, so this is a difference of 3 citations per paper	N/A
Epidemiology	Citation counts for 4,194 papers published between 2008 and 2012 in six top journals were calculated	Female first authors received four fewer citations on average than those with male first authors, and articles with female last	None performed

		authors received two fewer citations than those with male last authors	
Archaeology	Calculated percentage of women cited in articles published in four archaeology journals from 1978-1998	Male and female authors cite female authors at a lower rate than expected, where the expected value is determined by the percentage of articles authored by women in that journal	None performed
Communication Studies	Citation counts were calculated for all 1,020 papers from 1991 to 2006 in two top communications journals	Women were cited on average 12.77 times; men were cited on average 17.73. Women were cited less than expected when controlling for percentage of publications (36%)	ANOVA ($p = .011$)
Ecology and evolutionary biology	H-indexes were calculated for 187 editorial board members of seven journals	Women had lower average h-indexes when controlling for scientific age (time since first publication)	None performed
Neuroscience	Citation counts were calculated for 53,351 papers that were published from 2009 to 2010 in 30 neuroscience journals with the highest impact factor	Men were cited an average of 51.36 times, while women were cited an average of 48.78 times	ANOVA ($p < .001$)
International relations	3,000 articles published in twelve of the most prestigious journals in international relations between 1980 and 2006	An article written by a woman has about 20% fewer citations than a similar article written by a man when controlling for a number of variables such as topic, publication venue, time, and career level	Logistic regression (“statistically significant”, no p-value given)

The first paper argues that GPCs are a form of epistemic injustice, where women are denied access to uptake, a term borrowed from Longino. The second paper provides a sympathetic extension of Longino's views and argues that citations can be used to measure uptake and ensure that all members of the community are treated equitably. The third paper considers what would happen if citations were given even more weight in key decisions in scientists' careers and shows that women may be subject to even more epistemic injustice when grants are given to the most highly cited researchers.

1.2 Paper 1: Epistemic Injustice and Gendered Citation Practices

In the first paper of this dissertation, I show that the failure of uptake of women's work in the form of fewer citations constitutes an epistemic injustice. Miranda Fricker coined the term "epistemic injustice" in her 2007 book, although earlier philosophers discussed the same idea under other names (Code 1991, Collins 1991). An epistemic injustice is "a wrong [is] done to someone specifically in their capacity as a knower" (2007, 1). I argue that gendered citation practices constitute an epistemic injustice. In particular, these citation practices are a form of what Christopher Hookway calls a *participatory epistemic injustice* because they prevent women scientists from fully engaging in the epistemic practices of science.

Next, I draw on Helen Longino's notion of "uptake" to argue that the epistemic practice that women are excluded from is uptake of their work. Hookway's account clarifies why preventing women from participating is harmful, but participatory epistemic injustice still is focused on interactions between small groups of individuals. The failure of uptake of women's ideas happens at the scale of the scientific community, and the harm done to women is

cumulative in a way that points towards a broader scale of injustice that changes the trajectories of women's careers.

1.3 Paper 2: Revising Longino in Light of Gendered Citation Practices

In the second paper of this dissertation, I extend Longino's influential framework and her concept of uptake to show that citations and bibliometrics can be used to track whether a scientific community is adequately open to criticism. In doing so, measurements of uptake become a way to determine if a community consistently excludes or marginalizes particular sub-groups and can be used to track whether what Longino refers to as tempered equality occurs. I argue that in scientific communities with gendered citation practices, certain groups are not receiving uptake and, thereby, these communities cannot be characterized as meeting Longino's criterion of tempered equality.

To do so, I first describe Longino's view of science, that scientific knowledge is both rationally and socially justified. Scientific knowledge must be supported by reasons and evidence, but it also must be criticized by a wide set of perspectives. Longino develops a set of criteria that a community must meet to ensure that adequate criticism occurs and is considered by members of the community: public venues, uptake, shared standards, and tempered equality. When a community is characterized by all four criteria, that community can be said to successfully transform information into scientific knowledge.

I argue that citations are a way to track uptake in a scientific community. If an idea is taken up by the community, it should be clearly visible in the citations of the community. Quantifying uptake in terms of citations allows us to evaluate whether uptake is distributed across scientists as we would expect it should be. Citations are an effective tool to measure

uptake because they are used to connect authors and scientific work, visible to the community, and already understood to be an important part of the credit economy of science. I will argue that GCPs illustrate the power of using citations to measure uptake to ensure transformative criticism's proper functioning. I argue that using citations to measure uptake provides a concrete way to measure whether a scientific community treats all members of the community as having equal intellectual authority. This allows us to make the comparative claim that a community where GCPs occur is one where there is less tempered equality than one with more equitable citation practices.

1.4 Paper 3: Modeling Gendered Citation Practices

In the third paper of this dissertation, I present an agent based model (ABM) that simulates the effects of GCPs on a cohort of scientists as they progress through their careers by acquiring grants, publishing, and accumulating citations. I also investigate how different granting schemes may impact the community and interact with GCPs

I explore whether proposals focused on giving grants to the most highly cited researchers have disparate impacts on women. The evidence for or against the proposal is not definitive (nor could it be given the constraints of an ABM), but it does suggest that such a proposal would be ill advised in some circumstances. In particular, it seems that the impacts for women are worse in fields where average publication counts are low, especially early in a career and when the observed citation bias against women is high.

To do so, I first will argue that an ABM can be used to isolate the impact gendered citation practices have on publications, grants, and citations over the long term. An ABM can

provide another line of evidence that supports the causal story that women are disadvantaged in science due to gendered citation practices.

I also argue that an ABM is an effective way to test proposals to change grant funding that weigh citations much more heavily than the current system. While (so far) no grant agencies have proposed to make all decisions about grants based solely on citation counts, citations are seen as a potential proxy for measuring impact of past work and may provide information about the impact of future work. If grant agencies wanted to maximize the impact of the work they fund, heavily weighting citations might be a way to do so. The model developed here can test such a proposal and consider if such a proposal should be shelved because it causes disproportionately worse outcomes for women.

The model shows that there is a negative correlation between the variable for level of bias against women and the average grant rate for women compared to men; the stronger the bias, the more negative the outcomes; they receive fewer citations and are less likely to be in the top 25% of most cited researchers. The effect sizes are generally small; the average proportion of women's average citations compared to men's is reduced by about 1%, even when bias is high. I also investigate the potential impact of bias in astronomy using empirical publication and citation rates for the discipline. Women are also less likely to be in the top 25% of most cited researchers and have a lower average proportion of citations. As a result, I conclude that grants agencies should not give grants solely on the basis of citation counts, given the possibility for gender bias and the difficulty in implementing such a proposal in a thoughtful way to mitigate bias.

1.5 CONCLUSION

All three of these papers provide compelling evidence for why the scientific community must take GCPs seriously. Women are negatively impacted by gendered citation practices, and scientists should look for ways to both mitigate the impacts of GCPs and to prevent the occurrence of GCPs in the future. Some ways to mitigate the impacts of GCPs include focusing on how citation metrics are used, encouraging collaboration, and increasing women's opportunities to publish. Some ways to prevent GCPs might include improved training for researchers about how to find sources, implementing effective implicit bias training, ensuring women are included on journal editorial boards and regularly serve as peer reviewers, and introducing requirements to check publications for equitable citation practices. Several of these interventions focus specifically on what journals could do; given that citations are only generated when an article is published, it makes sense that journals should take a more involved role in preventing GCPs and other kinds of citation bias. For example, increasing the number of women who serve as peer reviewers and journal editors would increase the likelihood that at least one person who reviews the paper during the publication process is a woman, who may be more likely to recognize when some kinds of papers are missing from the reference list of a manuscript. Journals can also encourage high quality citation practices by explicitly checking reference lists for their gender (or racial) distribution; just saying that they are committed to improving equity in citation practices may be enough for researchers who submit manuscripts to make sure they have sought out relevant work from a wide variety of scientists. If journals are committed to citation equity, researchers will be more motivated to use resources such as Cite Black Authors ("Cite Black Author Database") and databases of underrepresented scholars such as the UP Directory ("UPDirectory") to seek out work by underrepresented authors. The

effectiveness of these interventions in reducing GCPs is plausible, but researchers should investigate if and to what extent they succeed, and if there may be unintended consequences.

These papers explore some of the philosophical problems of GCPs, but there are many avenues to pursue further. I have focused on comparisons between men and women, but those two categories do not cover the gender spectrum; it would be valuable to think about who, other than women, might be subject to GCPs and how that may change my analysis. There are also many other marginalized groups in science who also may be subject to inequitable citation practices, and the kinds of problems racial differences in citations rate pose for science may be importantly different from those discussed here with respect to GCPs. Citations are an important part of scientific practice, and I hope other philosophers will continue to explore their effects in science.

Chapter 2. EPISTEMIC INJUSTICE AND GENDERED CITATION PRACTICES

2.1 INTRODUCTION

Starting in the early 20th century, the field of Paleoindian studies--sometimes referred to as Early Man studies--was dominated by male archaeologists focusing on the lives of the first groups of people who came to the Americas from Asia roughly 15,000 years ago (Gero 2000, 307). These male archaeologists predominantly concentrated on the activity of the men in Paleoindian groups, specifically the hunting of megafauna such as mammoths and elk. Paleoindian societies were identified by archaeologists by the types of tools that men used for hunting; the kinds of rocks shaped into tools by Paleoindians had a particular fluted shape, and "[w]ithout the fluted point, there is no Paleoindian" (2000, 310). Paleoindian culture was identified with the hunting activity of men. The mostly male archaeologists did not consider what happened after the hunt, how the meat was prepared, or other alternatives for how Paleoindians ate.

During the same period, female archaeologists were concentrated in fields that focused on tools other than hunting tools, such as nutting stones and flake tools that were used for more 'domestic' purposes such as preparing food and cleaning animal skins, which were presumably done by the women in Paleoindian groups (2000, 313). The vast majority of books and articles on this topic were written by men, or pairs of men; 89% of articles in nine of the most important collected volumes from 1965 to 1990 had at least one male author, while 10% had only female authors. Both male and female archaeologists were doing lithic analysis of artifacts, but few women were cited by the male archaeologists or included in the volumes on lithic analysis edited by men.

Furthermore, women were cited less frequently than their male counterparts in archaeology at large (2000, 314). Gero points to the case of Eileen Johnson, a woman who was working in the male-dominated areas of Paleoindian studies. She had two single-authored publications in 1978 and 1980 that were cited a total of three times; in contrast, two male authors writing about findings from the same archaeological sites were cited eight and nine times for their respective articles. When Johnson co-authored pieces with men, she was cited a similar number of times as single male authors, but much more frequently than when she was a solo female author; her two co-authored publications received seven and eleven citations.[2]

Many scientific fields, such as archaeology, have increased the representation of women in the last few decades; given the changing demographics of science, it may seem that gender bias in science is gone, or at least rapidly disappearing. However, women continue to face challenges in science. In particular, women are cited less frequently than men in a number of disciplines, even after controlling for their relative frequency of publication and other potential confounding factors. This occurs in a wide variety of disciplines, including some in which women make up a larger proportion of the field than men.

The fields in which women are cited less frequently than men include: astronomy (Caplar, Tacchella, and Birrer 2017), archaeology (Hutson 2002), communication science (Knobloch-Westerwick and Glynn 2013), ecology and evolutionary biology (Kelly and Jennions 2006), epidemiology (Schisterman et al. 2017), gynecological oncology (Hill et al. 2015), international relations (Maliniak, Powers, and Walter 2013), neuroscience (González-Álvarez and Cervera-Crespo 2017), neurosurgery (Khan et al. 2009), pediatric neurosurgery (Klimo et al. 2014), and psychology (Geraci, Balsis, and Busch 2015). The methods vary across studies which

[2] Hutson (2002) looks at citations to women in the four archaeology journals from 1980s and 1990s and corroborates this pattern more broadly.

document these disparities and focus primarily on the most prestigious journals in their fields; but almost all find a statistically significant difference in citation rates by gender. For example, consider the evidence of GCPs in astronomy, epidemiology, and psychology. Caplar, Tacchella, and Birrer (2017) examine the papers published in five top astronomy journals since the 1950's and find that women received four percent fewer citations than men. Women receive an average of four fewer citations per paper than men for articles published in the top six journals in epidemiology (Schisterman et al. 2017). There are also differences in other kinds of citation metrics such as the h-index. The h-index of a researcher reports the number of papers of their papers, *h*, that have at least *h* citations. Women working in the top 100 departments in psychology have an average h-index of 17.17, while men working in similar departments have an average h-index of 21.64, a statistically significant difference (Geraci, Balsis, and Busch 2015, 2027). Combined, these studies give strong evidence that gender has an impact on citation rates.[3]

In this paper, I will show that this failure of uptake in the form of fewer citations constitutes an *epistemic injustice*. Miranda Fricker coined the term "epistemic injustice" in her 2007 book, although earlier philosophers discussed the same idea under other names (Code 1991, Collins 1991). An epistemic injustice occurs when "a wrong [is] done to someone specifically in their capacity as a knower" (Fricker 2007, 1). I will also argue that GCPs are a *first-order epistemic oppression*, a term introduced by Dotson (2012), which is a much broader category than that discussed by Fricker. I will next consider what Christopher Hookway calls a *participatory epistemic injustice*, another kind of first-order oppression; I argue that GCPs are a

[3] This list is likely not an exhaustive list of all the communities in science where gender bias may exist. More analyses of citation data are being done by members of scientific communities, especially as more attention is paid to citation rates for important decisions in a scientist's career. While it is possible that additional studies will shift the balance of the evidence from supporting gender bias to disconfirming it, the effect of gender bias in citation rates accords with both the subjective experience of women and a large body of literature on the biases women face in a number of different arenas, both academic and not (Valian 1999).

participatory epistemic injustice because they prevent female scientists from fully engaging in the epistemic practices of science.

Hookway's conception of participatory epistemic injustice is useful to help define the problem with gendered citation practices, but it does not fully capture the critical features of the epistemic harms: namely, the harm done to the epistemic status of the entire community's knowledge as well as the cumulative nature of these harms to individual and cohorts of women. I will use Helen Longino's notions of "uptake" and "objectivity" to demonstrate that gendered citation practices constitute an epistemic harm, not only to women in science, but also to the community of knowers as a whole. Additionally, I will explicate the cumulative nature of these harms for individuals in science and for the community as a whole.[4]

2.2 The Function(s) of Citations

Citations are a complex epistemic practice. They serve (at least) two primary purposes. First, they track the connections between and development of ideas across subsequent publications. Second, they function as a form of academic currency; citation counts are regarded as a proxy for impact and are used as one piece of evidence in tenure portfolios and grant applications. Merton (1979) delineates the two purposes of citations such that one is cognitive, and one is 'moral': citations are "designed to provide the historical lineage of knowledge and to guide readers of new work to sources they may want to check or draw upon for themselves. In their moral aspect, they are designed to repay intellectual debts in the only form in which this can be done: through open acknowledgment of them" (viii). I disagree that the distinction is as stark

[4] In my analysis, I will focus on women, but the same conclusion likely holds for any sub-group of science that has lower than expected citation rates and has experienced a historical background of exclusion from science.

as he paints it; both uses of citations have both an epistemic and ethical component to them, but he has correctly characterized which aspect is dominant for each purpose.

Given these two distinct purposes, it only makes sense to consider what kinds of citation practices are just or unjust and how they are unjust by focusing on the context. In science, some exclusions are appropriate. When using functions in the first sense, to connect publications, there are many perfectly just reasons to fail to cite something. Not every article deserves to be cited; some articles are redundant, poorly written, boring, or methodologically weak. Clearly, the author of a low quality article is not wronged when they fail to accumulate any citations to their paper. But collectively, the evidence for GCPs shows us that women are less likely to be cited in the aggregate. It's reasonable to assume that women and men have roughly the same distribution of article quality, and that women are no more likely to write poor articles than men.[5] But when GCPs occur, it seems that something is going wrong; it is unlikely that the differences are due to quality, yet women are, on average, worse off than men when it comes to their citation counts. Given the epistemic function of citations to connect new scientific work to what has come before it, GCPs likely constitute an epistemic injustice of some kind.

Further, there also seems to be a different harm happening when considering citations as a form of academic currency. Citations do seem like an effective proxy for research impact; a citation does indicate some amount of engagement with a previous work, so if a paper is highly cited, it is likely that it contains something important, novel, or at least controversial. But there is no reason to think that men and women will do high impact work at different rates. If citations are a measure of research impact and women receive fewer citations, their work will seem less

[5] There is even some evidence that women submit higher quality publications to offset potential gender bias in the peer review process. See Lee (2016) for a summary of the relevant research and Bright (2017) for a decision theoretic model that provides evidence for this claim.

important. If citations are used in decisions like grant or tenure applications, women may be less likely to succeed in having their grants funded or receiving a promotion. Again, the harm seems to be epistemic in nature, but in contrast to the previous case, less directly so. Citations as a form of prestige can enable (or stymie) future scientific research; they have an instrumental role in furthering the epistemic goals of scientists.

Once these two functions of citations are delineated, it is clear that there is some kind of epistemic injustice occurring, but the type may differ based on what function citations are fulfilling. Next, I will discuss some of the major views on epistemic injustice and oppression.

2.3 Epistemic Injustice and Epistemic Oppression

2.3.1 *Fricker's Testimonial and Hermeneutical Injustice*

In this section, I will discuss Fricker's conception of epistemic injustice and her definition of a negative identity prejudice, a necessary component of establishing that an epistemic injustice has occurred. Miranda Fricker's book, *Epistemic Injustice,* defines an epistemic injustice, broadly, as "a wrong done to someone specifically in their capacity as a knower" (2007, 1). She defines and contrasts two types of epistemic injustice: testimonial injustice and hermeneutical injustice. Her concepts of epistemic injustice establish the importance of a negative identity prejudice to defining an epistemic injustice, which I take to be a crucial component of the injustice of GCPs. However, her two types of epistemic injustice do not accurately capture the nature of GCPs.

In a testimonial injustice, a speaker is given less credibility for their testimony due to a prejudicial stereotype about members of the speaker's group (Fricker 2007, 21). One such example that Fricker discusses is the in-court testimony of Tom Robinson in *To Kill a*

Mockingbird. The all-white jury has lowered credibility in Tom's testimony, given the racial prejudices they have about black people, and as a result, do not believe him when he says he did not harm Mayella Ewell (2007, 23-24).

Hermeneutical injustice occurs when members of marginalized groups are prevented from developing the epistemic resources to interpret their own experiences due to prejudices about their identities on the part of the dominant group (2007, 155). Fricker's primary example of this is sexual harassment. Before the term 'sexual harassment' was coined by women in the 1970's, women were kept from the appropriate epistemic resources to make sense of this experience (2007, 151).

Both testimonial and hermeneutical injustice require negative identity prejudices in whoever is committing the injustice: the hearer in the case of testimonial injustice and the dominant group in the case of hermeneutical injustice. A negative identity prejudice is a negatively-valanced judgment made about a person on the basis of their membership in a social group (2007, 27-28). Negative identity prejudices are closely related to prejudicial stereotypes; when a member of a marginalized group provides testimony, the negative stereotypes about that group can undermine the hearer's ability to evaluate the sincerity or competency of the speaker. The hearer may be less likely to give an appropriate amount of credibility to the speaker's testimony as a result of the negative stereotype (2007, 30-32).

These prejudices are importantly resistant to counter-evidence; if someone has a mistaken belief about a social group that is corrected when they encounter contradictory evidence, it will fail to be an identity prejudice (2007, 35). Maintaining a prejudice in light of counter-evidence requires an affective investment on the part of the prejudiced, and any negative identity prejudice

that is maintained in light of such evidence necessarily involves "an ethically bad affective investment" (2007, 35).

The addition of 'epistemic injustice' to the philosophical vocabulary has generated a great deal of discussion among epistemologists and ethicists about other types of wrongs that are distinctly epistemic. While Fricker's concepts of testimonial and hermeneutical injustice cannot fully capture the wrong done by gendered citation practices, her definition of 'negative identity prejudice' will be necessary in showing why gendered citation practices constitute an epistemic injustice.

If gendered citation practices are an epistemic injustice in the form that Fricker provides, it must be due to an identity prejudice on the basis of gender. While there might be explicit sexists in academia who refuse to cite women on principle, these bad actors are almost certainly in the minority due to the liberal values held by most scholars across academia (Prentice 2012). These explicitly sexist scientists are therefore unlikely to substantially contribute to the robust gendered citation practices seen in some sciences. While explicit sexism might explain some small part of gendered citation practices, it seems much more likely that implicit biases are what causes less uptake for papers authored by women. Implicit biases are "typically thought to involve associations between social groups and concepts or roles" (Saul and Brownstein 2016, 1-2). Implicit association tests (IAT) have shown that most people have stronger associations between men and science than they do between women and science (Nosek et al. 2009). In a sample of over 500,000 IAT results from around the world, about 70% showed at least somewhat stronger association between men and science than women and science (2009, 10594). If scientists are implicitly biased towards perceiving men as belonging in science, then they may be

more likely to treat the scientific work done by men favorably and are therefore more likely to cite in it their own work, whether favorably or negatively.

Both women and men exhibit the same patterns of bias towards men and science on IAT measures. Farrell, Cochrane, and McHugh (2015) replicate the findings of Nosek et al (2009) on the IAT, but they also use another measure of measure of implicit association, the implicit relational assessment procedure (IRAP). Using the IRAP, they find that women have a moderate bias towards men and science and men did not have a statistically significant bias towards men in science in their sample of 12 male respondents (2015, 126). Carli et al (2016) investigate the connections between stereotypes about men, women, and scientists and find that the stereotypes of research participants about scientists are more similar to the stereotypes they hold about men than those that they hold about women.

The implicit bias against women in science is resistant to counter-evidence. The number of women participating in science has consistently increased over the past 50 years, creating a great deal of counter-stereotypical evidence that may have reduced implicit biases against women in science. Yet, most people still hold these biases. There is a robust literature on the ways in which our thoughts are not totally transparent to us (Talbott 1995, Kahneman et al 1982), and these implicit biases seem fairly entrenched.

However, Fricker requires that the bias is not responsive to counter evidence specifically because of an affective judgment on the part of the hearer. This affective judgement is not clearly present in the case of implicit associations against women in science; many scientists explicitly support women in science, but still maintain stronger implicit associations between men and science. Saul (2016) and Holryod and Puddifoot (2020) have discussed the connections between implicit biases and epistemic injustice and seem to assume implicit biases can lead to epistemic

injustices with or without an affective component, since implicit biases generally go against our stated commitments. Instead, we might think of the resilience of these biases as something closer to inertial resistance to counter evidence than an affective resistance.[6] So implicit biases against women in science constitute a negative identity bias against women in science.

However, GCPs don't seem to fit into Fricker's categories of epistemic injustice. Gendered citation practices do not prevent women from making sense of their experience in science; women often explicitly identify the ways in which they receive lowered credibility in science through gender bias in citations. Thus, gendered citation practices cannot be a hermeneutical injustice. Gendered citation practices don't seem to be a form of testimonial injustice, either. In a testimonial injustice, the hearer gives a lower credibility judgment to a speaker on the basis of a negative prejudice the hearer holds about people like the speaker. When considering citations in their first function, to connect scientific work, it's not clear that failing to cite is due to a reader giving a credibility deficit to the writer of a particular paper. Citations are often to a paper as a whole, not specific pieces of testimony; citations are used to indicate agreement but can also be used in the service of criticism. Failing to cite someone is not the same kind of judgment as not believing them; a failure to cite is closer to failing to recognize that someone has contributed to a conversation at all. When we consider citations in their second role, as a form of status in scientific communities, GCPs are clearly not a testimonial injustice, since the function of citations in that instance has almost no similarities to testimony. Gendered citation practices are not testimonial or hermeneutical injustices, but the harm is epistemic in nature.

[6] Thanks to Bill Talbott for the term 'inertial resistance'.

2.3.2 *Dotson's Epistemic Oppression*

In the previous section, I argued that Fricker's epistemic injustice framework shows that a negative identity prejudice is necessary for an epistemic injustice, but that the problem of GCPs does not clearly map onto either of her categories. In this section, I will discuss Kristie Dotson's concept of epistemic oppression. I take epistemic oppression to be a wider phenomenon than epistemic injustice, where any epistemic injustice is an epistemic oppression. Dotson specifically defines a new kind of epistemic oppression, a contributory injustice, that goes beyond Fricker's two kinds of injustices, and which she describes as third-order epistemic oppression. She also argues that epistemic oppressions such as testimonial and hermeneutical injustice are reducible to social and political oppression, while her new category of third-order epistemic oppressions is irreducible to such oppression. I will then consider whether this new type of epistemic oppression can help understand the problem of GCPs.

Dotson (2014, 2012) takes a different approach to epistemic oppression than Fricker does to epistemic injustice. She categorizes epistemic oppression by first-, second-, and third-order epistemic oppressions, where each type is more complex and requires more complicated remedies than the order below it. She distinguishes between epistemic exclusion, "an unwarranted infringement on the epistemic agency of knowers" (2014, 115), and epistemic oppression, "persistent epistemic exclusion that hinders one's contribution to knowledge production" (2014, 115). The shift from exclusion to oppression occurs when thinking about how to correct epistemic exclusions when the exclusion is related to a particular social identity.[7]

[7] In (2012), Dotson describes Fricker's concept of negative identity prejudice and provides her own examples of beliefs about Black people that are maintained by affective resistance to counter evidence. In (2014), however, Dotson places epistemic oppression in the context of already existing social and political oppression. She does continue to provide examples of "negative prejudices" against speakers, so she has not totally given up Fricker's language. While it is somewhat unclear whether the negative stereotypes against women in science are due to affective resistance to change, it is much harder to deny that women have been a historically oppressed group, both

Dotson describes first-order epistemic exclusions as an inefficiency of the epistemic resources of a community (2014, 126). Her first-order epistemic oppressions include testimonial injustices (2012, 26). When some speakers are given less credibility, it is not that the concept of credibility doesn't exist in the community or that something about the epistemic resources of the group needs to change. Instead, all members of the community need to extend credibility equally to every member, although Dotson recognizes what a challenging task this is (2014, 126).

While I have already shown that GCPs are not a testimonial injustice, Dotson's first-order epistemic exclusions seem like a better fit. In the case of citations as connections, there is no obligation to cite every paper, or even every paper one reads or knows of when preparing a manuscript. So, failing to cite a paper is, most of the time, an epistemic exclusion, but not a particularly problematic one. There are all sorts of perfectly legitimate reasons for choosing not to cite a publication. We may worry that if we ask researchers why they choose to include or exclude a particular source, their post hoc explanation will not mention bias, even if it played a factor in the actual choice. But it would be very difficult to prove or show that bias is the cause. So GCPs in the context of connecting literature aren't necessarily first-order epistemic oppression; it is not that the epistemic resource of connecting with other papers should be shared more widely.

In the case of citations as credit, it seems plausible to say that GCPs represent a failure to distribute the epistemic resource of credit fairly. If citation counts are taken as a reliable source of information about research impact but are subject to gender bias, it will result in unfair, reduced assessments of the value of research done by women. These assessments may feed back

politically and socially. If the background conditions of social and political oppression are sufficient for epistemic oppression, then the affective component of the negative identity prejudice may not be needed. In that case, the implicit biases against women in science are adequate to show that epistemic oppression occurs.

into the choices researchers make about who to cite but can also have impacts on the ability of women researchers to compete for grants and tenure. When considering the role of citations as a form of academic currency, it does seem that first-order epistemic oppression occurs.

A second-order epistemic exclusion occurs when the shared epistemic resources of the group are not adequate for a specific epistemic task (126-127, 2014). A second-order epistemic oppression then corresponds to a hermeneutical injustice for Fricker (2012, 29). The community (or some of its members) recognizes that the epistemic resources needed for a task, such as making sense of the experiences of the community, are lacking (2014, 127). So, to fix second-order epistemic oppression, the community needs to augment or change its available epistemic resources (2014, 128-129).

Importantly, both first- and second-order oppression are reducible to political or social oppression for Dotson. She does not deny that there is an epistemic aspect of first- and second-order oppressions but argues that social oppression is what maintains them and the changes necessary to correct both types require undoing the social oppression, not providing a specifically epistemic response (2014, 116). In first-order epistemic oppression, providing equality to the currently marginalized group is the way to rectify the oppression--where equality means that the stereotypes used against members of the marginalized group are genuinely gone.

Second-order epistemic oppression contains hermeneutical injustices, and I have already shown that GCPs do not constitute a hermeneutical injustice. When considering citations in their function as markers of prestige, GCPs do not seem to constitute second-order epistemic oppression, since GCPs do not point to the inadequacy of some epistemic resource of the scientific community.[8] But when considering GCPs in their two functions, it does seem that

[8] GCPs do contribute to a critique of the use of citations as a proxy for research impact; if some groups have lower citation rates for reasons that are not related to the quality of their work, citations are not as valuable a source of

GCPs when focusing on citations as connections across papers can, in some instances, constitute second-order epistemic oppression. In the case of archaeology in the 1980s described in the introduction, women may suffer from a second-order epistemic oppression, because their papers that are not being cited show that the questions and methods of the discipline need to expand. Such cases point to a need to expand the epistemic resources available to the community, not just to distribute them more equitably.

A third-order epistemic exclusion is "a compromise to epistemic agency caused by inadequate dominant, shared epistemic resources. To say that such resources are inadequate is to throw into question the relevance of a given community's overall dominant resources in light of the knowledge production activities in question" (2014, 129). A type of third-order exclusion is a contributory injustice (2012, 31). In a contributory injustice, one member of an epistemic exchange, the perceiver, fails to recognize the relevance of an alternative epistemic framework offered by the knower. Dotson departs from Fricker because Fricker assumes there is only one set of epistemic resources available to all epistemic agents and Dotson recognizes that there are often alternative frameworks developed and used by marginalized epistemic groups. In a hermeneutical injustice, marginalized members of a community are prevented from developing new epistemic resources to make sense of their own experiences; in a contributory injustice, the epistemic resources already developed by marginalized groups are not recognized as relevant by the dominant members of the epistemic community (2012, 32).

Third-order epistemic oppressions are irreducible to social oppression because it is a feature of epistemological systems, not social ones, that drive third-order oppressions (2014, 132). Epistemological systems, which "include[...] operative, instituted social imaginaries, habits

information about research impact as they are presented as. The impact of using citations in grant determinations will be discussed in depth in Chapter 3.

of cognition, attitudes towards knowers and/or any relevant sensibilities that encourage or hinder the production of knowledge" (2014, 121), are resilient to change. Even if the social oppression that has generated the conditions for the third-order epistemic oppression to occur disappeared, the epistemological system will not change without substantial effort and the third-order oppression would persist. While second-order oppression requires a change to the epistemic resources, third-order oppression requires changes to the broader epistemological system that can only be brought about by taking seriously the claims of people that the epistemological system deems untrustworthy or unreliable.

It is reasonable to next ask if GCPs might be a form of third-order epistemic oppression. The idea of contributory injustice sounds promising; when considering citations as either a form of connections across papers or as a marker of prestige, GCPs undermine women's ability to contribute to science. But it is not clear that GCPs are a form of exclusion that points to a problem with the epistemological system of science; women are not providing evidence that the entire enterprise of science should be rejected in favor of an alternative system that can better accommodate their observations.[9]

But Dotson's account is an important extension of Fricker's and does provide a helpful framing for how to think about GCPs. Testimonial and hermeneutical injustices are just types of first- and second-order epistemic oppressions, respectively, which opens up the possibility of

[9] This determination may also depend on what we take to be the epistemic resources of a system vs the epistemological system itself. If we think the differences in methods offered by women archaeologists in the 80s (or another example where women's increasing representation in science brought about specific changes in the methods or questions of a community) would require a change to the epistemic system of science, then GCPs may constitute third-order epistemic oppression. I would treat the scientific method as an example of an epistemological system, though, and I don't think that there are scientists who are writing papers that would undermine our belief in the scientific method. Dotson seems to have a somewhat narrower conception of an epistemological system; she discusses an example of a contributory injustice where an academic fails to recognize the contributions of an activist, because the academic does not see the relevance or importance of the activist's contribution due to differing epistemological systems (2012).

many other kinds of epistemic oppressions existing within each order. Further, she emphasizes the kinds of changes necessary for correcting each type of oppression.

2.3.3 *Hookway's Participatory Epistemic Injustice*

Both the types of epistemic injustices laid out by Fricker and Dotson's third-order oppressions do not seem able to make sense of GCPs; the problem is not that GCPs undermine testimony or point to a lack of conceptual understanding. It's not that GCPs point to a fundamental problem in our epistemic system. Instead, GCPs seem to be a form of first-order epistemic oppression, although not a testimonial injustice. Christopher Hookway provides a name and description for a category very similar to first-order epistemic oppressions: participatory epistemic injustices.[10]

Participatory epistemic injustices include testimonial injustices, but the uniting feature is that these injustices prevent a knower from fully participating in an epistemic *practice*, which can include more than just testimony (Hookway 2010, 155). Hookway argues that Fricker views epistemic injustice from an informational perspective--in which only exchanges of information are important. Instead, starting from a participatory perspective--in which the ability of a knower to engage with epistemic practices is at the forefront--can capture a wider variety of injustices that are specifically epistemic (2010, 157). Testimony is one epistemic practice, but we also participate in others such as discussion and questioning. A participatory epistemic injustice occurs whenever a participant in any epistemic practice is prevented from fully engaging with the practice due to a prejudicial stereotype held by the other member in the exchange.

[10] Hookway's paper was published before Dotson's and does not use the term 'epistemic oppression' and thus cannot argue for himself that participatory epistemic injustices are (or are not) first order epistemic oppressions.

Hookway provides a number of examples that are not testimonial injustices, but in which an epistemic injustice still takes place. In the first, a teacher does not recognize her student as a potential partner in inquiry due to a prejudicial stereotype she holds about the student (2010, 156). The teacher answers the student's questions when he asks for clarification or additional information; when the student asks a non-informational question meant to engage with a deeper issue, the teacher ignores the question or assumes the question is irrelevant to discussion. She does not see the student as being able to contribute due to her stereotype about members of his group. This example fails to be a testimonial injustice, because the student is not testifying to any particular fact that the teacher disbelieves; instead, he is participating in a different epistemic practice: questioning.

Another example Hookway gives focuses on participants in philosophical discussions in which hypotheticals and counterexamples are offered and refuted; the content of the discussion contains assertions, but not exclusively (2010, 158). If a participant in the discussion is perceived as being a poor philosopher because their contributions are judged in light of some group membership, a participatory epistemic injustice is committed--although, once again, a testimonial injustice does not necessarily occur.

Hookway is most interested in epistemic injustices in the classroom; all of his cases revolve around educational epistemic practices, such as discussion and question asking. In a classroom situation, even the most equitable one, there are real differences between the level of expertise and understanding between the students and teacher. In contrast, GCPs represent exchanges between intellectual equals.

A participatory epistemic injustice seems to fall under the umbrella of first-order epistemic oppressions. Hookway identifies a number of other epistemic practices, in line with

Dotson's epistemic resources, where members of an epistemic exchange are not treated equally. It is not that there is a lack of understanding of how questioning works when a participatory injustice occurs; it is that a member of the exchange is not treated as an intellectual equal.

Hookway's notion of participatory epistemic injustice is useful, because it substantially broadens the domain of epistemic injustice: it can include all varieties of epistemic practices, including those that occur in science and are relevant to citation. Gendered citation practices prevent women from fully engaging with the epistemic practices of science. In particular, failing to cite women in science prevents them from receiving uptake for their work.

2.4 Uptake and Gendered Citation Practices

Hookway's concept of participatory epistemic injustice is helpful in broadening the epistemic practices from which knowers can be excluded; however, it does not capture how GCPs are connected to the epistemic status of the knowledge community as a whole. Helen Longino's account of science can elucidate the importance of citations to science and to help show why gendered citation practices can constitute an epistemic harm, not only to women in science, but also to the entire scientific community. By Longino's account, only through the uptake of its content can original work be conferred the status of knowledge by the community. If women are suffering participatory epistemic injustice in the form of gendered citation practices, the knowledge produced by the community is less objective, or even fails to become knowledge.[11]

[11] Longino uses the criteria she develops in (1990) and (2002) as an ideal for scientific communities to strive for to ensure objectivity. I use her discussion of objectivity to emphasize the stakes for a community: GCPs can directly threaten the success of the entire endeavor of science, if Longino is correct. Even if the connection between objectivity and uptake is not what Longino describes, we should still be concerned about the effectiveness of a

For Longino, criticism is what ensures the objectivity of the knowledge created by a community; any background beliefs held by the members of the community that influence their work will be appropriately scrutinized and can be rejected by the community. The four criteria she proposes to ensure adequate criticism are: public venues, shared standards, uptake, and tempered intellectual authority (Longino 2002, 129-131). The criterion of most interest here is uptake. Uptake requires that members of the community hear the criticisms made by others and respond appropriately; not every criticism will need a response, but the community members must show that they are willing to modify their background beliefs in light of criticisms.

While Longino's account applies specifically to criticism, uptake is necessary for original work as well. If some members of the community who contribute original work to a community are not being subject to criticism – or if some members of the community who provide criticism are not being taken up – then the knowledge produced by the community is not objective. Citations represent one such form of uptake; they show that an original work has been read by other scientists, and the use of a citation requires some kind of comment on, or integration with, the content of the work.

For gendered citation practices to constitute a participatory epistemic injustice, a participant must be denied from full participation in an epistemic practice due to an identity prejudice on the part of their partner in the practice. Longino's concept of uptake is helpful to understand from what epistemic practice women are excluded. Uptake is an epistemic practice; it is a necessary step in the process of transforming hypotheses and data into objective scientific knowledge, and citations are an important part of that codification into knowledge (Longino 1990, 69). If particular groups are being excluded from uptake, either by not having their original

scientific community that does not treat all of its members as intellectual equals. The connections between uptake, intellectual equality, and objectivity will be explored much more fully in Chapter 3.

work be appropriately criticized, or by failing to have their criticisms of others' work heard by other members of the community, then a participatory epistemic injustice is done.

Thus, given the evidence about gendered citations practices and the implicit biases against women in science, we should conclude that women are being excluded from the epistemic practice of uptake, and are subject to an epistemic participatory injustice.

2.5 Cumulative Harms

Citations and other measures of publication impact are becoming increasingly important in institutional decisions affecting scientists' careers, such as employment, tenure and grant applications (Abbott et al. 2010). Journal impact factor, a frequently used indicator of research quality, has fallen out of favor in light of criticisms that high journal impact factor does not always correspond to high impact for a paper published in a given journal (Casadevall and Fang 2015; Seglen 1997; Lariviere and Sugimoto 2018; "San Francisco Declaration of Research Assessment"). There has been a shift toward using article- and individual-level citation counts and related measures as the important metrics of impact instead, because such measures focus on the particular work, as opposed to the venue of the work. Given the importance of citations to the individual careers of scientists, it is clear that a participatory epistemic injustice in the form of gendered citation practices will have a substantial impact on women in science.

However, thinking of GCPs as a participatory injustice is not enough; women are consistently excluded from uptake, and just focusing on one instance fail to capture the harm done by repeated participatory epistemic injustices, both diachronically for an individual and synchronically for a marginalized group in a community. This is because the harm of gendered citation practices is not fully captured by summing up each instance of injustice. First, I will

consider the larger harms done to an individual across her career. I will then discuss the harms done to women as a group in scientific communities.

Suppose that two scientists, one male and one female, are both assistant professors working at similar types of institutions on projects of similar quality and importance in the same field, with the same number of publications of similar quality. If the female scientist suffers from one participatory epistemic injustice in the form of her work not being cited in a publication by a colleague due to the colleague's identity prejudice about women, she has been harmed by being denied access to full participation in science. The harm, however, may not have a substantial effect on her career path. She may still be able to secure grants for her future projects, continue to publish, and receive tenure. However, if she is subjected to the repeated harms of gendered citation practices, this may impact her career.

For example, if the female academic from the previous paragraph is working in epidemiology, each of her papers will receive roughly four fewer citations than her similarly situated male colleague (Schisterman et al 2016, 163). If she and her male colleague both have five papers as assistant professors, this means she is already behind her colleague by 20 citations. If citation counts are used as a measure of quality in grant applications, as has been suggested by Ioannidis and Khoury (2014), she may fail to secure at least one grant she applies for due to her lower citation count. She not only loses the prestige attached to earning the grant but may be less able to fund her research and ultimately publish fewer papers, which further decreases her odds of obtaining tenure.

Now, consider the effects if such a scenario repeatedly plays out for all the women in a discipline. Many fields in science are still dominated by men at the top of the field. For example, in epidemiology, women have gone from representing a third of the field in 1984 to over half by

2006 (Schisterman et al 2016, 159) while much of the leadership of epidemiology has remained male; 59% of full professors in the discipline are men, compared to 30% of the associate professors and 27% of the graduate students (2016, 162). While there are more women entering science than ever before, small differences in the way that men and women are evaluated can lead to fewer women at the top of a field, even when starting from equal pools of men and women (Martell, Lane, and Emrich 1996).

For example, consider a group of 100 assistant professors, 50 male and 50 female, at comparable universities doing work of similar quality. Let's assume that this group starts with an average of five publications each, and the women of the group receive four fewer citations on average. If they all apply for the same set of grants, and citation counts are used to decide who is awarded a grant, a disproportionate number of women will fall below the cutoff. The effects of these differences will be further amplified through subsequent application cycles as the entire cohort progresses through their careers.[12]

2.6 Conclusion

In this paper, I have shown that the gendered citation practices in science constitute a participatory epistemic injustice against women, and that the harms they cause in aggregate are greater than the sum of their individual impacts. The cumulative harms of multiple participatory epistemic injustices can damage women's careers, keep them out of leadership positions in their disciplines, and force them out of the field altogether. These epistemic injustices cause serious harm to the state of scientific knowledge, reducing the objectivity of that knowledge and potentially preventing some scientific works from being recognized as knowledge. While

[12] This proposal will be investigated in more detail in Chapter 4.

gendered citation practices are not the only problem women face in science, such poor citation practices are on their own a substantial problem for science. Understanding why gendered citation practices constitute an epistemic injustice, and attending to the cumulative harms they bring, can help shape a response to the problem.

Chapter 3. REVISING LONGINO IN LIGHT OF GENDERED CITATION PRACTICES

3.1 INTRODUCTION

Many scientists and philosophers have argued that information must be critically evaluated before it can become scientific knowledge; scientists must confirm that what constitutes scientific knowledge is appropriately justified (Kitcher, 2002; Solomon, 2001). Scientific communities also must be open to criticism and alternate perspectives; science progresses by changing in light of new data, methods, and theories (Kuhn, 1962). While an openness to criticism is widely recognized as an important aspect of scientific practice, not enough attention has been paid to quantify that openness nor how to determine if all members of a community are equally likely to be listened to.

In this paper, I will extend Longino's influential framework and her concept of uptake to show that citations and bibliometrics can be used to track whether a scientific community is adequately open to criticism. In doing so, measurements of uptake become a way to determine if a community consistently excludes or marginalizes particular sub-groups and can be used to track whether tempered equality occurs. I argue that in scientific communities with inequitable citation practices, certain groups are not receiving uptake and, thereby, these communities cannot be characterized as meeting Longino's criterion of tempered equality.

In Section 3.2, I describe Longino's view of science, that scientific knowledge is both rationally and socially justified. Before a piece of information becomes scientific knowledge, it must have reasons and evidence to support it, but it also must be socially justified in the sense that it be adequately scrutinized from multiple perspectives to ensure unacceptable background

assumptions are rooted out. Longino develops a set of criteria that a community must meet to ensure that criticism occurs and is considered by the community: public venues, uptake, shared standards, and tempered equality. When a community is characterized by all four criteria, that community can be said to successfully transform information into scientific knowledge.

In Section 3.3, I argue that citations are a way to track uptake in a scientific community. If an idea is taken up by the community, it should be clearly visible in the citations of the community. As information science and bibliometrics improve, it will be easier and easier to track the spread of ideas through a community. While there are other ways that uptake might occur, looking at citation patterns has several advantages. Citations are mechanisms for attributing previous ideas to their authors, visible to the community, and already understood to be an important part of the credit economy of science. Quantifying uptake in terms of citations allows us to evaluate whether uptake is distributed across scientists as we would expect it should be. I will argue that *gendered citation practices* (GCP), the persistent differences in citation rates for men and women, illustrate the power of using citations to measure uptake to ensure transformative criticism's proper functioning.

In Section 3.4, I argue that using citations to measure uptake has the added advantage of providing a concrete way to measure whether a scientific community treats all members of the community as having equal intellectual authority. To close, I will discuss how this work can be extended in future empirical work and theoretical work as the sophistication of bibliometric research continues to develop.

3.2 LONGINO'S VIEW

In *Science as Social Knowledge* (1990) and *The Fate of Knowledge* (2002), Helen Longino argues for a view that she calls *critical contextual empiricism* (CCE). According to CCE, knowledge must be rationally justified, through careful observation and reason, but it also must be socially justified, through a process that she calls *transformative criticism* in which background assumptions undergo a rigorous process of critical scrutiny from many members of the scientific community. To ensure that transformative criticism can occur, Longino argues that particular norms of community engagement must be in place.

These norms are motivated by concerns about how the underdetermination of hypotheses by their evidence can lead to idiosyncratic science. Scientific theories are not logically entailed by observations. This means that scientists must always rely on background assumptions about methods, instrumentation, and values in conjunction with observations to provide adequate justification for theories.[13] Background assumptions are necessary, but unlike observations, they cannot always be tested empirically. When background assumptions that license inferences from data to hypothesis remain uncontested due to the lack of diversity in perspectives, critique, and uptake within the scientific community, this damages the normative standing of its knowledge claims.

[13] Longino distinguishes between two types of background assumptions; she calls epistemic background assumptions *constitutive* and non-epistemic ones *contextual* (1990, 4). She cites simplicity and predictability as constitutive assumptions and beliefs about intrinsic differences between people of different races and sexes as a contextual assumption. Longino does not think that contextual values are illegitimate in science. Instead, she argues that they are a necessary part of science, and instead of trying to remove all contextual values in background assumptions, we should simply ensure that the background assumptions are ones that the community can accept (1990, 81-82). Much has been said about the possibility of distinguishing epistemic and non-epistemic values since Longino first wrote (see, for example, Douglas, 2009; Steel and Whyte, 2012; Elliott, 2011). In general, though, the possibility of distinction does not make much of a difference to the process of transformative criticism; both kinds of background assumptions are equally subject to revision under CCE.

As one of her key case studies, Longino critiques the research on the biological basis of differences in mathematical ability between the sexes (1990, 117).[14] Researchers argued that prenatal testosterone exposure best accounts for the differences in mathematical ability between the sexes found on standardized tests such as the SAT. Longino observes that this research program depends on several assumptions: that there is one stable attribute of mathematical ability which the SAT accurately tests for, and that gendered question content will not affect scores for girls and women differently from those of boys and men (1990, 123-126). These assumptions, among many others, are necessary to license the inference from observations about differences in hormone levels and SAT scores across the sexes to a theory that says hormone differences cause differences in mathematical ability.

To launch her critique, Longino need not say these background assumptions are obviously false, wrong, or unwarranted; instead, she simply observes that they are taken for granted when moving from observation to theory - the researchers accept the background assumptions as obvious or not even as assumptions at all (1990, 184-85). Because the knowledge produced by the community that depends on those assumptions is contextualized to that community and time, the theory is only acceptable to the extent that the assumptions that support it have withstood relevant criticism.

The process of transformative criticism is designed to provide the best opportunity for the community to both unearth and scrutinize such hidden assumptions. For example, the assumptions she highlights about the reliability of the SAT in testing for some single, stable

[14] The research Longino discusses has advanced substantially in the 30 years since she wrote *Science as Social Knowledge*. The hypothesis Longino discusses, that hormones are responsible for differences in mathematical ability, has largely been rejected in favor of a more sociocultural explanation (see Ceci et al 2014, 84-93, for a summary of the research on these differences). However, differences in hormone levels are still provided as a possible explanation for a wide variety of phenomena, often relying on the same assumption Longino highlights.

attribute of mathematical ability are largely now rejected: most education researchers do not think that the SAT can capture 'mathematical ability' because they now have a more inclusive view of what mathematical ability is.[15] The assumption made about the SAT's ability to measure mathematical ability is not an empirically tested observation; it is not possible to go out into the world and find what constitutes the limits of mathematical ability. It is instead a background assumption that motivates methodological choice, how we interpret evidence, and, thereby, must itself be subjected to a process of transformative criticism.

3.2.1 *Community Norms*

What community-wide norms enable a community to identify and scrutinize its background assumptions? Longino's criteria of CCE are venues for criticism, uptake, public standards, and tempered equality of intellectual authority.[16] I will explicate each in what follows.

For background assumptions to be made available for community-wide scrutiny and to facilitate open critique, there must be *public venues* for original research and criticism (2002, 129). For Longino, "[t]he avenues for the presentation of criticism include such standard and public forums as journals, conferences, and so forth" (1990, 76). For Longino, the community should treat criticism as an equally valuable activity as original research.

[15] See, for example, Kattou et al (2013) which defines mathematical ability as "a multidimensional construct, including quantitative ability (number sense and pre-algebraic reasoning), causal ability (examination of cause-effect relations), spatial ability (paper folding, perspective and spatial rotation abilities), qualitative ability (processing of similarity and difference relations) and inductive/deductive ability" (167).

[16] I will use the names for the criteria from *The Fate of Knowledge*. In *Science as Social Knowledge*, Longino uses slightly different names: recognized venues for criticism, shared standards, community response, and equality of intellectual authority (1990, 76-81). *The Fate of Knowledge* provides a more refined version of the criteria, where Longino has taken up (!) many of the criticisms of her view, clarifying and sharpening her position.

Second, criticism should not just be heard, but be taken up by the community. For uptake, "[t]he community must not merely tolerate dissent, but its beliefs and theories must change over time in response to the critical discourse taking place within it" (2002, 129-130). That "change may comprise the acceptance of new beliefs, the modification of beliefs, [and] the development of new data, reasons, and arguments" (2002, 129-30). Longino suggests that "this responsiveness is measured by such public phenomena as the content of textbooks, the distribution of grants and awards, [and] the flexibility of dominant world views [sic]" (1990, 78). Even when critiques do not fully persuade individual community members, they must be ready to "defend their work against criticism" (1990, 78) and, more generally, "pay attention to and participate in the critical discussion taking place" so that "the assumptions that govern their group activities remain logically sensitive to it" (2002, 130). Indeed, those who fail to ever uptake criticisms of their own views may lose their intellectual authority (133). Uptake has a central role in ensuring transformative criticism can occur: "[u]ptake is what makes criticism part of a constructive and justificatory practice" (130).

Third, there must be public standards for how theories are evaluated and what determines whether a criticism is worthy of engagement (2002, 130). These public standards allow community members to "share some referring terms, some principles of inference, and some values or aims" which is necessary "for the identification of points of agreement, points of disagreement, and what would count as resolving the former or destabilizing the former" (2002, 130-1). Public standards exclude certain kinds of criticisms such as creationists who prioritize biblical evidence over empirical evidence. The public standards themselves are also subject to revision over time through the same processes of critical discourse (2002, 131).

Fourth, a community must have tempered intellectual equality. While scientists will differ in both their innate abilities for certain kinds of tasks and their education in science, the community must treat those with relevant knowledge as worthy of consideration. Longino claims, “The exclusion of women and members of certain racial minorities from scientific education and the scientific professions constitutes not only a social injustice but a cognitive failing” (132); science is worse off by having fewer kinds of dissenting voices.[17] No one should be excluded from the conversation because of their social position alone. However, equality is tempered in that public standards do rule out some kinds of criticism as not relevant since differences in expertise, education, and experience can and should still affect the weight given to what researchers say. If a community agrees, it cannot be because some voices have been left out; it must be because all voices were integrated into the deliberation about what should be accepted. Tempered equality is important because it ensures that criticism is offered from as broad a range of perspectives as possible.

To sum up, Longino defines her notion of *epistemic acceptability* in this way:

[17] Standpoint theorists endorse a similar claim, although for different reasons, and rely on many of the same examples as feminist empiricists, such as Longino, do. For example, Haraway, a prominent standpoint theorist, argues in *Primate Visions* that the changes to the dominant views in primatology on the sexual behavior of female primates came about as a direct result of the increasing representation of women in primatology (1989). Haraway argues that women were in an epistemically superior position to men to be able to identify the androcentric assumptions that were operating in primatology. Wylie, another foundational standpoint epistemologist, argues for what she calls the situated knowledge thesis: "social location systematically shapes and limits what we know, including tacit, experiential knowledge as well as explicit understanding, what we take knowledge to be as well as specific epistemic content" (2003, 4). She further develops what a standpoint is, which goes beyond situated knowledge; a standpoint requires one to both recognize their own situated knowledge and to understand how the epistemic practices of their communities enable the exclusions of certain types of knowledge. Wylie highlights the contributions of women in archaeology in the late 80s and early 90s that drew attention to the sexist assumptions most male archaeologists had relied on in generating their research questions as a particular instance of the positive impact of a feminist standpoint on a scientific discipline. Longino directly cites Haraway and uses other similar cases, where the increasing representation of women lead to more opportunities for criticism from new perspectives, and science was improved as a result. Longino does not go so far as to endorse the view that women were especially suited to provide those criticisms. Intemann (2010) argues that feminist standpoint epistemology and feminist empiricism (especially Longino's CCE) have become more similar and should ultimately merge into what she calls "feminist standpoint empiricism". Longino and feminist standpoint epistemologists agree that having more women scientists is beneficial for science, even if they differ on some of the details.

> Some content A is epistemically acceptable in community C at time t if A is or is supported by data d evident to C at t in light of reasoning and background assumptions which have survived critical scrutiny from as many perspectives as are available to C at t, and C is characterized by venues for criticism, uptake of criticism, public standards, and tempered equality of intellectual authority. (2002, 135)

This definition exemplifies each aspect of CCE: the content must go through appropriate critical examination; the content is accepted by a particular community at a particular time; the content must be supported by empirical evidence in addition to the necessary reasoning and background assumptions. Additionally, a community must have all four criteria of transformative criticism in place to judge some content as epistemically acceptable. Although this definition of epistemic acceptability may appear to be binary--either the information has survived adequate scrutiny or it has not; either the data is evident to the community or not; either the community is characterized by the four criteria or not-- she makes it clear that it is not intended to be so: "Because the norms represent ideals that can be partially satisfied, this qualification is a matter of degree. Communities may be more or less dependable and effective as producers of knowledge." (2002, 134).[18] Because communities can achieve this ideal to varying degrees, we can make comparative claims: a community that is better characterized by the criteria is more effective at producing knowledge than one less characterized by the criteria.

3.2.2 *Changes to Other Norms*

Longino's work has been hugely influential, and I am not the first person to suggest that the criteria are too vague and merit refinement. Other philosophers have also discussed ways in which the practice of science shows us that the criteria are not sufficient for preventing

[18] See also her discussion of this point in Science as Social Knowledge (1990, 76-81).

undesirable content of background assumptions from sneaking in. Most philosophers working in this vein have articulated more specific, concrete ways of interpreting the criteria to improve transformative criticism's functioning.

Most philosophers have focused on public venues, shared standards, and tempered equality. For example, Borgerson (2011) and Solomon and Richardson (2005) both discuss ways in which the existence of public venues is not sufficient to ensure adequate critical scrutiny is achieved. Borgerson highlights that some medical researchers fail to use the public venues, by either failing to publish negative results or purposefully publishing in obscure journals. Solomon and Richardson (2005) are similarly concerned that the existence of public venues is not enough to compel researchers to disclose all relevant research results.

Other philosophers have discussed how Longino's description of shared standards is inadequate. Lee and Schunn (2011) show several ways that public standards can be applied in biased ways against researchers based on their social identities, which impacts which ideas are made available in public venues, which are taken up, and the attribution of intellectual authority. Leuschner (2015) discusses how shared standards can be inconsistently applied to women and people of color, and she provides suggestions such as standardized authorship order to reduce the likelihood of bias in the application of the standards.

Tempered intellectual equality is also vague. Borgerson and Fehr (2013) both argue that tempered equality requires stronger cultivation of dissent than Longino indicates, although they disagree on what kind of dissent is most productive to promote. Borgerson advocates for any dissent that generates new theories, while Fehr favors demographic diversity because it will generate additional epistemic diversity.

Uptake has not received as much attention as the other three of her criteria even though criticism is incredibly important to the functioning of science.[19] Criticism is a major driver of change in science. The four criteria all work together to ensure that transformative criticism occurs. But uptake is clearly central to the success of transformative criticism.

In the next section, I will discuss how uptake is threatened by gendered citation practices. I will provide a refinement of uptake, that like many of the above discussions of Longino, is meant to ensure that the criteria she delineates can better achieve her aims.

3.3 Uptake and Gendered Citation Practices

In the previous section, I discussed some of the concrete ways Longino seems to think uptake should be seen: presentations, papers, textbooks, the pattern of funded grants, and the like. Although Longino identifies some ways we can measure change in a community, she misses a powerful and informative metric. In this section, I will argue that citations should be used to measure uptake at the community-wide level. By foregrounding the pattern of citation-level uptake, we can measure whether uptake is distributed as we would expect across the scientific communities. I will argue that quantifying uptake in terms of citations provides a way to measure GCPs, which damage the community's capacity for appropriate uptake and transformative criticism. Updating uptake so that it should be measured in terms of citations, therefore, enhances Longino's account so that it can better track how different ideas, proposed by

[19] Borgerson (2011) says that she offers amendments to all four of the criteria. The original version of uptake (which she calls responsiveness to criticism) that she attributes to Longino says "the community as a whole must be responsive to such criticism" (2011, 438). The revised version she offers says "the community must be responsive to criticism" (2011, 443). The change in scope isn't remarked upon in the text, and it's not clear why she says she has modified all four criteria since the changes to uptake are quite minor and not explicitly argued for.

members belonging to different social groups, are woven into the community's store of knowledge.

3.3.1 *Why Citations Should Be Used to Measure Uptake*

Longino recognizes that citations are artifacts that trace the integration of ideas and critiques into the broader body of scientific claims that stand as the community's store of knowledge. She says, "Publication in a journal does not make an idea or result a brick in the edifice of knowledge. Its absorption is a much more complex process, involving such things as subsequent citation, use and modification by others, et cetera. Experimental data and hypotheses are transformed through the conflict and integration of a variety of points of view into what is ultimately accepted as scientific knowledge" (1990, 69). However, when characterizing uptake, she does not recognize patterns of citation as a valuable tool for investigating patterns of uptake. This is a missed opportunity. In contrast to many discursive spaces where uptake might take place—e g., a chat with a colleague or a question at a talk—citations are visible, traceable, and archived in the public output of science, including conference papers, talks, and publications. Other aspects of the publication process, such as peer review, may appear promising as sites of uptake, something Longino also recognizes: "Peer review is often pointed to as the standard avenue for such criticism, and indeed it is effective in preventing highly idiosyncratic values from shaping knowledge" (1990, 76). But reviewer reports are not generally available for study. Citation data and bibliometrics, on the other hand, are readily accessible. Databases such as Web of Science and Google Scholar collate the citations of nearly all papers in nearly all disciplines. Citations are an accessible way to find the connections between research programs and will show us which researchers are discussing each other's ideas.

How can citations reveal patterns of uptake? Criticism of previous work is common even in largely original work. Bertin et al show that approximately 17% of citations within a large corpus of papers have a negative valence, using phrases such as 'disagrees with' or 'does not show', when discussing a previous work (2016, 1427). Not all, or even most, of these negative citations are likely true criticisms; it may be that the authors of the original study are merely noting that their results did not align with the results found in another paper, without saying it is due to a difference in methods or background assumptions that is flawed in some way. But presumably, at least some of these instances *do* correspond to genuine criticisms. Additionally, many of the positive citations may refer to the choice to use a method, observation, theory, or background assumptions over another; the authors may use citations to show that they agree with particular criticism instead of citing in a negatively valanced way what it is they disagree with. So, by looking at the citations of a community, we can have a reasonable proxy for uptake: if people are not being cited, their ideas are not being engaged with and their criticisms are unlikely to be taken up.[20]

How can citations be used to see a community-level change in response to criticism? Imagine Author A writes paper X. Author B reads paper X and disagrees with *i (*where *i* is a theory, observation, methodological choice, or background assumption*)* that A endorses, so B writes a paper Y that directly criticizes *i.* Author A and some members of community C read paper Y and consider the criticisms offered. In this case, A finds the criticism in paper Y to be compelling enough to revise their beliefs, and they write paper Z that acknowledges and cites the

[20] There may also be relevant differences in how much uptake is or could be visible in a discipline's citation patterns. One might also see uptake at a conference: some researchers are directly responding to someone else in their presentation, or there is an author meets critics session. Some disciplines are so fast-paced that conferences are likely the best place to look for uptake. Even in a field where most of the criticisms occur in conference papers, not publications, those debates likely will eventually filter into publications and therefore citations.

criticism in paper Y, and ultimately, A rejects *i*. Some members of the community C are similarly convinced by paper Y and reference it (and perhaps paper Z as well) in future works when discussing the dispute. One can see that paper Y has more citations than paper X, because the community accepts the arguments offered in paper Y. [21]

Uptake does not require that every possible criticism be capitulated to, but it should be considered. Even in Paper Y represents a criticism of Paper X that isn't strong enough to convince A or other members of C to reject background assumption *i,* Paper Y still should be cited as a foil to those who continue to accept *i*; it should be recognized in the form of citations that the criticisms in paper Y were at least considered.

Citations also have an already recognized function in the credit economy of science. Citations are considered a better proxy for research impact than measurements such as journal impact factor because they more directly focus on the impact of individual papers (Seglen, 1997; DORA; see also Hicks et al, 2015 and Wilsdon et al, 2015 for a more cautious endorsement of some citation metrics).

When the community is genuinely responsive to criticism, we should be able to trace it through the citation practices of the community. Once we can measure uptake, it is possible to use uptake to measure the extent to which a community is characterized by tempered equality. In the next section, I will show how we can use bibliometric methods to ensure that particular sub-communities are having their criticisms taken up when we have a quantifiable way to look at

[21] Mere counts may sometimes misrepresent the relationship between papers X, Y, and Z. Even when paper Y provides a compelling criticism of paper X, it will not always be the case that paper Y has more citations. Paper X may be older and has had a chance to accumulate more citations. Other members of the community may consistently cite both paper X and Y to discuss the dispute between them, even if they ultimately side with paper Y. Paper X may be cited in other disciplines, even though most people in A's discipline have rejected *i.* But on the whole, scientists will likely cite the paper that changed their minds more than the one they rejected. The empirical patterns of how to see uptake in citations will only be easier to determine as more sophisticated techniques develop.

uptake. If a community is not providing uptake to a particular subset of the community, that community is inadequately characterized by the norms.

3.3.2 *Gendered Citation Practices*

Longino recognizes the value of bibliometric research and alludes to its value in measuring uptake. In the context of the contributions of women to the philosophy of science, she says this: “We should ask ourselves whether and how questions introduced by female colleagues fare in the profession’s conversations. Here, bibliographic and citation research that tracks when questions were first introduced to the profession, when they started to gain traction, and whose work related to them was taken up as worthy of serious discussion could be revealing” (Longino, 2013). In her 2013 piece, Longino is concerned that the contributions of women and feminist philosophers of science have not been adequately taken up by the wider community of philosophers of science. But the problem she highlights is much broader than just philosophy of science: women are not receiving adequate uptake in a variety of academic fields.

The literature on gendered citation practices demonstrates that entire sub-groups of researchers with similar intellectual authority are not being taken up at the same rates as other sub-groups. Gendered citation practices occur in a variety of scientific disciplines. These fields include astronomy (Caplar, Tacchella, and Birrer 2016), archaeology (Hutson 2002), communication science (Knobloch-Westerwick and Glynn 2013), ecology and evolutionary biology (Kelly and Jennions 2006), epidemiology (Schisterman et al. 2017), gynecological oncology (Hill et al. 2015), international relations (Maliniak, Powers, and Walter 2013), neuroscience (González-Álvarez and Cervera-Crespo 2017), neurosurgery (Khan et al. 2009), pediatric neurosurgery (Klimo et al. 2014), and psychology (Geraci, Balsis, and Busch 2015).

This represents an incredibly diverse section of science; it includes social sciences, medical fields, biological sciences, and physical sciences. A difference in citation rates or another citation metric is found in every one of these fields. In some fields, the effect is seen for the field as a whole, and in others, it is only seen within academic ranks, such as only for more junior faculty. However, many fields are left with citation disparities by gender that cannot be accounted for by controlling for any known characteristic of the authors or the work.

3.3.3 *Why Gendered Citation Practices Threaten Uptake*

Uptake is crucial for the norms of CCE to function as they should. All four of the criteria work together to ensure that scientific information is scrutinized from a wide variety of perspectives, but without any uptake whatsoever, science would be unrecognizable. It is difficult to even conceptualize how scientific progress could be made without scientists listening to each other's critiques and building on each other's new ideas. Clearly, some amount of uptake is a standard part of scientific practice. But getting uptake right is crucial for Longino's theory to succeed: if the process of transformative criticism is what elevates scientific information into scientific knowledge, then criticisms must be taken up in line with their quality and relevance, not the social identity of the person offering them.

In a field with GCPs, on the surface, it may look like uptake in Longino's sense occurs. The community does change background assumptions, methods, and research questions in response to criticism as measured by content in textbooks, funded grants, papers, and conference presentations. The fact that women are cited less frequently is not evidence that *some* uptake does not occur. Consider, for example, astronomy: astronomy has GCPs, but also is characterized by uptake and produces knowledge. Astronomers have changed their minds about

methods, observations, and background assumptions, and we can see these changes in citation patterns as a result.

However, the occurrence of gendered citation practices means either women are not being read in the first place, or other authors are reading them but failing to find or recognize their relevance to future research.[22] If women are consistently less cited, the criticisms they do offer have less chance of being taken up. Longino's criteria are meant to ensure that scientific knowledge is adequately scrutinized to deserve its label. If certain perspectives--namely the perspectives of women--are not being listened to and engaged with, then background assumptions may not be adequately criticized.

Community C should be citing paper Y if the criticisms offered in Paper Y are to be genuinely taken up by the community. If Paper Y fails to accumulate citations due to GCP, then the community is not providing the uptake for a criticism it should. Longino's theory of uptake needs to explicitly include citations to measure whether certain perspectives are systematically ignored or taken up.

3.4 Uptake and Tempered Equality

In this section, I argue that using citations to measure uptake provides a powerful way to determine if tempered equality exists within a scientific community. Many discussions of tempered equality have focused on what kind of diversity is important for science and how to best fulfill Longino's mandate that diverse perspectives be cultivated. Borgerson (2011) and Fehr

[22] I am not going to take a position on which option is more likely, but my analysis does not depend on what the causes of GCP are. It is enough to note they occur to discuss what might happen as a result without delving into the psychology of researchers about why they are less likely to cite women.

(2013) disagree about what kind of diversity is important for science, but both agree that more diversity of some kind is important. Hicks (2011) worries that Longino's theory is too open to diverse views, in that it may allow in perspectives of Nazi scientists. Steel et al (2021) argue that Longino does not recognize that diversity in science can hinder the epistemic activities of science.

But none of the discussions of tempered equality have provided a way to quantify the extent to which a community is characterized by tempered equality. Tempered equality, like the satisfaction of all four criteria, exists on a scale. It is sometimes clear that a community does not qualify as having adequate tempered equality, such as most historical scientific communities, where women and people of color were consistently excluded from the practice of science. Longino specifically highlights the positive epistemic effect of increasing the representation of women in science; women were able to recognize androcentric assumptions and spawned many new research areas (2002, 132). There are clear epistemic benefits of increasing representation, but increasing representation is not enough. The demographic of a scientific community is not the only way, nor the most effective way, to measure whether tempered equality exists.

Citations allow us to investigate if there are any differences in whose ideas are taken up. In fields with GCPs, women are cited less frequently, and their ideas receive less uptake. In a community with less tempered equality, this is exactly the kind of pattern we would expect to see: some sub-group--in this case women--is not being treated equally to another sub-group--in this case men per their intellectual authority. Tempered equality requires that "[w]here consensus exists, it must be the result not just of the exercise of political or economic power, or of the exclusion of dissenting perspectives, but a result of critical dialogue in which all relevant perspectives are represented" (Longino 2002, 131). If certain sub-groups of the community, such

as women, are not included or deemed not relevant, tempered equality is reduced. Using citations as uptake and looking for disparities across groups also gives a way to compare the level of tempered equality across communities. Consider astronomy, a field where GCPs occur. We could compare astronomy to astronomy*; in astronomy*, women and men are cited with equal frequency. It seems clear that astronomy* is doing a better job than astronomy in listening and responding to criticisms from the largest possible number of perspectives; the perspectives of women are less likely to be included in astronomy than in astronomy*. Not only does astronomy have less uptake than astronomy*, but it also has less tempered equality.

Astronomy may seem like a strange example, since the epistemic benefits of women's perspectives are not as obvious as when the subject matter has clearly gendered implications, such as in social sciences, medicine, and biology. I will not argue here that women qua women offer any particular new perspective in astronomy.[23] But the exclusion of any subgroup leaves open the possibility that an important perspective is excluded, even if women are no more likely to have such a perspective than any other subgroup of the discipline.

If novel ideas are evenly distributed across a community and one group is less likely to have their novel idea taken up, the community will miss out on the opportunity to engage with those novel ideas. Depending on how rare good ideas are and how much gender bias there is in citations, this could result in a substantial decrease in the epistemic output of a scientific community or a substantial slowing down of the scientific process. For example, assume a

[23] One regularly noted practice in science (and most academic disciplines) is that women and men work on different topics; the case studies feminist philosophers of science use depend on this fact. Caplar, Tacchella, and Birrer (2016) also note this difference in astronomy, even though the choice of research area is clearly not tied to the gendered aspects of the research, since the objects of study don't have genders. But the clustering of women in certain subdisciplines doesn't change my analysis; instead of comparing astronomy and astronomy*, we can instead think of comparing the subdisciplines of solar physics and solar physics* and find the same concerns.

community has 50 men and 50 women. In a particular year, every member publishes two papers. One out of every ten papers has a substantial criticism that would move the discipline forward, so there are 20 of these papers, equally distributed across men and women. If women are 10% less likely to be cited than men, then at least one of the papers by a woman author will be missed despite its valuable contribution. While the case of actual scientific practice is much more complex, the scientific community should ensure that all criticisms have an equal chance of receiving uptake for maximum epistemic benefit.

In a field where the research questions are more closely tied to gender, such as psychology, the need to fully include women is more obvious. Women are more likely to have different kinds of criticism when gender is the direct object of study, even if their criticisms are no more likely to be better in any way than those given by men. Improving citation practices will ensure that the criticisms of both men and women can be fully evaluated by the community and taken up when deemed reasonable.

A community where citation patterns are equitable for all members of the community does not guarantee that tempered equality occurs, however. It may be that the community has equitable citation practices but does not adequately cultivate diverse perspectives. So having equitable citation practice is not sufficient for tempered equality, but it is a necessary condition. Because equality of intellectual authority is tempered, and women are represented more heavily among junior scientists, someone might object that these citation patterns simply reflect differences in seniority.

3.5 CONCLUSION

Tracking uptake with citations has substantial benefits for Longino's theory. It provides both a way to understand when a community is taking up criticism appropriately and a way to monitor whether there is tempered equality in how those criticisms are taken up. This approach to quantifying Longino's account of uptake and tempered equality of intellectual authority opens up new avenues for research on the health of scientific communities. I have focused on the difference in citation rates by gender in this paper, but there are likely other kinds of disparities in citation patterns for other underrepresented groups such as scientists of color and scientists from the Global South; Hopkins et al (2012) and Lavriere et al (2013) discusses some of the disparities in publications for scientists from underrepresented racial groups and the Global South, respectively. While the sense of uptake I have provided will likely provide the same impetus to improve citation practices for any underrepresented group in science, there may be issues that are more specific to particular underrepresented groups. For example, whenever race is a direct subject of study in a particular discipline, we should expect that scientists who belong to racial minorities will have a different set of criticisms from those of white practitioners. Ensuring that all criticisms have an equal chance to be heard and taken up when appropriate has clear epistemic benefits. Any evidence of citation bias also strengthens the case that more needs to be done to fully include members of all underrepresented groups in science.

It would also be worthwhile to investigate what kinds of interventions might improve disparities in citations for a social group; encouraging more equitable citation practices will improve both the uptake and tempered equality of a community. The representation of women in sciences has been steadily increasing; it is likely having more women to cite would eliminate some differences, but it would be worthwhile to investigate how likely this is, as well as what

interventions are most effective at encouraging women to enter a field and retaining them once they get there.

Citation bias also likely interacts with many of the other challenges women face in sciences, so improving outcomes for women in other areas may also help. Citations are used to evaluate scientists for grants and tenure; decreasing citation disparities could increase the success rate for grants and increase the representation of women in more prestigious academic positions. Having one's work recognized as valuable and contributing to an ongoing conversation is a kind of validation women scientists do not always get; improving citation practices may also keep more women in science. Women publish less frequently, and similarly to increasing representation, finding ways to encourage women to publish more would provide more opportunities for their work to be cited.

It may also be possible to more concretely measure the extent to which a community has public venues and shared standards. Analyses like Bertin et al (2016) could be used to ensure there are adequate public venues by investigating whether criticisms are published in prominent journals in a particular discipline and comparing it to other fields. The inclusion of women as journal editorial board members, peer reviewers, and grant funding boards may give information about the extent to which the community has shared standards; if some voices are being left out of the conversations about what counts as quality work, the resulting standards may skew towards certain perspectives.

Citations are also not the only kind of attribution of credit that might be relevant to quantifying uptake. As discussed in Section 3.3, uptake can occur in other locations such as informal conversation or at a conference Q&A. These types of uptake may be difficult to study since they are largely unrecorded. There are other written artifacts of uptake; for example,

articles and books will also include acknowledgment sections and it would be worthwhile to investigate if there are gender differences in those attributions.

This analysis also assumes that citations will continue to function in roughly the same way they do in current scientific practice. Citations provide a convenient way to measure uptake, and it is necessary for CCE to be able to know if uptake is happening, but nothing here requires that citations specifically are how uptake is determined. Instead, my proposal for uptake is only specific to communities where citations are already a regular practice of the community (which, to my knowledge, would include pretty much every scientific discipline currently). But it is worthwhile to think about other ways to measure uptake that will ensure Longino's criteria function in the way she intends, even if there are no citations whatsoever.

Chapter 4. MODELING GENDERED CITATION PRACTICES

4.1 INTRODUCTION

In this paper, I will argue that an ABM can be used to isolate the impact gendered citation practices (GCPs) have on publications, grants, and citations over the long term. Many studies show that women face several disadvantages in science (Moss-Racusin et al, 2012; Ley and Hamilton 2008; Helmer et al, 2017; Lariviere et al, 2013; West et al, 2017). However, few of these studies investigate if gender citation practices are a cause of disadvantage.[24] An ABM can provide another line of evidence that supports the causal story that women are disadvantaged in science due to gendered citation practices.

I also argue that an ABM is an effective way to test proposals to change grant funding that weigh citations much more heavily than the current system. While (so far) no grant agencies have proposed to make all decisions about grants based solely on citation counts, citations are seen as a potential proxy for measuring impact of past work and may provide information about the impact of future work and there are proposals to encourage and increase the use of citation metrics (Ioannidis and Khoury, 2014). If grant agencies wanted to maximize the impact of the work they fund, heavily weighting citations might be a way to do so. The model developed here can test such a proposal and consider if such a proposal should be shelved because it causes disproportionately worse outcomes for women. I conclude that grants agencies should not give grants solely on the basis of citation counts, given the possibility for gender bias and the difficulty in implementing such a proposal in a thoughtful way to mitigate bias.

[24] Lariviere et al (2013) does discuss gender differences in citation and publication rates and concludes that increasing women's representation and full participation in science is necessary, but there are no simple ways to do so.

The model shows that there is a negative correlation between the variable for level of bias against women and the average grant rate for women compared to men; the stronger the bias, the more negative the outcomes; they receive fewer citations and are less likely to be in the top 25% of most cited researchers. The effect sizes are generally small; the average proportion of women's average citations compared to men's is reduced by about 1%, even when bias is high. Using empirical publication and citation rates for the discipline, I use the model to investigate the potential impact of bias in astronomy. My results suggest that women astronomers are also less likely to be in the top 25% of most cited researchers and have a lower average proportion of citations.

4.1.1 *Agent Based Models*

Agent based models are an increasingly common tool for philosophers. Models can be used to provide a possible explanation for observed macro level empirical phenomena by testing whether a set of agents will produce the expected phenomenon given a set of simple rules (Epstein 2006). In contrast to statistical models, ABMs allow the agents to interact with their simulated environment and the environment responds to the individual choices over the course of the model (Bruch and Atwell, 2015).

Philosophers of science have used models to investigate questions such as how information should be shared in a scientific community (Zollman, 2013), how the productivity gap may arise between men and women (Bright, 2017), and how discrimination can arise in epistemic communities (Rubin and O'Connor, 2018).

Agent based models have also been used to inform policy decisions. Urban planners have used detailed simulations of traffic conditions in a city and simple rules about how people in the

simulation move around to show what impact certain transit policies might have on the flow of traffic through the city or how it might change land use (Waddell, 2002; Waddell et al, 2003). Epidemiological models are used to show that a policy intervention such as mask mandates are effective in controlling the spread of an illness by comparing what happens in the model when agents wear masks to what happens when they do not (Auchincloss and Roux, 2008). Some philosophers have also used the results of their models to make suggestions about public policy, such as Heesen (2017), which considers some proposals about grant allotment and will be discussed in more detail later.

4.1.2 *Modeling Gendered Citation Practices*

This chapter's ABM has two main aims: to investigate whether gendered citation practices alone make a difference to the long term outcomes of agents of different genders, and to test how gendered citation practices would interact with a proposal to give grants to more highly cited researchers.

There is a great deal of literature about the problems women face in science: a chilly climate (Wylie et al, 2007; Chisholm et al, 1999), lowered publication output (Larivière et al, 2013), a leaky pipeline leading to lower levels of career attainment and an underrepresentation of women (Avolio, Chavez, and Vilchez-Roman 2020), less access to mentorship (Blood et al, 2012), implicit biases (Moss-Racusin et al 2012), and (sometimes) lower grant rates.[25] While

[25] Bornmann et al (2007) conducts a meta-analysis of data from a variety of international grant agencies and finds that women receive grants at lower rates than men; Van der Lee and Ellemer (2015) show that women receive fewer grants than men in the Netherland's national research council; Witterman et al (2019) shows that women receive fewer grants from Canada's national research council. But Pohlhaus et al (2010) finds no difference between men and women in NIH grants and Ceci and Williams (2010) conduct a meta-analysis of the literature on bias against women in grants and conclude that grant review processes are not biased in favor of either gender. Both Van der Lee and Ellmer (2015) and Witterman et al (2019) find that the gaps are due to lower evaluations of women principal

gendered citation practices are obviously connected to these other issues, some studies indicate that there is bias even after controlling for many of these factors, such as fewer papers published, differences in citations for subdisciplines or topics, and lower average academic rank (Caplar, Tacchella, and Birrer, 2017; Way, 2017).

The productivity of scientists is already a key factor in evaluating researchers for grants and tenure and promotion; the number of publications and prestige of the journal they are published in is seen as a good proxy for the quality of the research. Moher et al (2018) look at the tenure and promotion criteria at a variety of US academic biomedicine institutions and find that they are generally focused on productivity measures, including number of articles published, journal impact factor, and the number of citations.

Using productivity measures in grant-funding decisions raises worries because women are less likely to be cited than men in a variety of scientific fields. Larivière et al (2013) find that papers with women as the solo author or in prestigious authorship positions--generally the first or last, depending on the discipline--are cited less than papers with men in the same type of authorship role. Their dataset includes over 5 million papers from 2008 to 2012, from nearly every country in the world. Other fields where women are cited less than men include astronomy (Caplar, Tacchella, and Birrer 2017), archaeology (Hutson 2002), communication science (Knobloch-Westerwick and Glynn 2013), ecology and evolutionary biology (Kelly and Jennions 2006), epidemiology (Schisterman et al. 2017), gynecological oncology (Hill et al. 2015), international relations (Maliniak, Powers, and Walter 2013), neuroscience (González-Álvarez and Cervera-Crespo 2017), neurosurgery (Khan et al. 2009), pediatric neurosurgery (Klimo et al. 2014), and psychology (Geraci, Balsis, and Busch 2015).

investigators, not due to lower proposal quality, meaning that if women are missing out on grants, it may be because of the compounded effects of gender difference.

An ABM can attempt to isolate the impacts of gendered citation practices. An ABM can assume that women and men are, for example, equally likely to publish and equally likely to receive a grant but assume that there is some small difference between the citation rates of men and women and investigate the impact of that difference alone. The difference of men and women scientists' actual careers can be directly measured by empirical investigation, but it is incredibly difficult to determine what caused those differences. An ABM can provide evidence that citation bias alone is responsible for some of the differences in career outcomes between men and women.

4.1.3 *Modeling Citation-focused Grant Funding*

Citation counts are an increasingly popular way to measure impact of scientific work; many assume that the best papers with the greatest ability to shape future science are the ones that are mostly widely cited. Ioannidis and Khoury (2014) and Heesen (2017) discuss, with varying levels of endorsement, some policy proposals that give grants to researchers that are the most highly cited. In this paper, I will consider what might happen if these trends were taken to their logical conclusion: what if grants were given to researchers solely based on being among the most highly cited?

Ioannidis and Khoury (2014) provide a suggestion for what this policy might look like. They argue in favor of giving grants and tenure to researchers who do well in the areas they highlight as important: productivity, quality, replicability, shareability, and translatability (PQRST). By productivity, they specifically advocate for using a metric such as the top 10% of most highly cited articles to determine who is most productive. They stress the importance of using all of the measures, so would not approve of solely using citation metrics alone. But much

of the rest of their proposal is purposefully vague; they want grant agencies and tenure committees to decide for themselves how important each criterion is, and how to evaluate them. The most concrete suggestion they offer is to give grants to those who are most highly cited, and this paper will test this citation-focused implementation of their proposal.[26]

Heesen (2017) discusses the type of scientist who might find such a proposal compelling: a scientist who is convinced that the most highly cited papers are such because they are the most informative papers and deserve all their citations on merit alone. He calls this the competence-based view of academic superstars, where superstars are just those who are very highly cited (4511). Given that granting agencies want to fund work that will be successful and impactful, if the competence-based model is correct, grant agencies should award grants to those who are most highly cited. Heesen ultimately does not endorse the competence-based model because he shows that luck can explain the emergence of superstars as well as the competence-based model can.

The proposal I test in this paper is: give grants to the most highly cited scientists. I also consider what happens when the cutoff for who constitutes a top performer is higher than the grant rate, and what happens when additional grants are given to those outside of the top performers after each top performer has earned a grant.

There are several reasons to use a model to test this proposal. A model can investigate what impacts a policy proposal might have on the careers of scientists without having to

[26] Ioannidis and Khoury's (2014) proposal is specifically bounded to biomedical research. In this paper, I will consider it as a general suggestion for all sciences. Biomedical research does have different aims than many other scientific disciplines; the focus is generally on treatments and interventions in biomedical research, while other sciences have a much broader range of intentions. Other parts of the PQRST are more specific to biomedicine; replicability, shareability, and translatability are all focused on the ability of a particular intervention succeeding in other applications, including as treatments. But both the productivity and quality metrics apply broadly to all sciences. Since I am interested in focusing on productivity metrics, it is warranted to apply a form of their proposal to all sciences.

experiment with the careers of actual people. Trying to test the impact of changing how grants or tenure is decided is both impractical and unethical. It would require massive changes to how science currently operates, which is unlikely to happen given how many institutions would have to agree to the change.

Further, if the practical and ethical issues of changing the reward structure of science were overcome and the intervention could be implemented in a randomized way, there are still many confounds that would make it hard to interpret any changes caused by changing methods for allocating grant funding; changing the grant structure of science would have an impact on the careers of scientists, as do many other factors that couldn't be controlled for. While an ABM cannot provide conclusive evidence about the effects of a major policy change, neither could an empirical investigation. An ABM has the advantage that it can be implemented by one person and provide some evidence within a much shorter timeframe and with many fewer resources than a large experiment would. Using an agent based model in conjunction with empirical investigations would provide the strongest source of evidence, but an ABM can give evidence that even attempting the empirical work is necessary.

Finally, a model can investigate what, if any, disparate impacts there will be on women on the basis of GCPs alone. Given that gender bias exists in several scientific disciplines, focusing solely on productivity measures may further disadvantage women in science. The model can also test what the specific impacts of gendered citation practices alone will be even if women published at equal rates and were equally likely to receive grants. But if grants are given solely to top performers and women may be less likely to be among the top performers, women's careers may suffer. The model can investigate how large an impact such a proposal would have when gendered citation practices are the only kinds of bias. If women are less likely to get grants

when grants are given to top performers while at the same time, they are also less likely to be cited, then it is likely the impact will be even greater when other forms of bias are also in play.

4.2 METHODS

The model was coded in NetLogo 6.0.4, a programming language designed for agent based models. A description of the model follows, and the full code of the model is in Appendix A.

4.2.1 *Initializing the Model*

In the initial phase, the model generates scientists, based on a user input for the size of the community. Half are men and half are women; using similarly sized populations of men and women removes the possibility that differences in population sizes are driving any effects seen in the model.

The starting agents are meant to represent a cohort of researchers at roughly the same career stage, near the beginning of their careers. It is reasonable to think of these agents as scientists starting in postdoc positions, for example, where the papers generated in this initial round represent those published while in graduate school. The starting set of agents could also represent researchers starting a tenure track faculty position, where the initial publications are those published during both graduate school and previous postdocs. The model can be adjusted to represent an accurate size for this community based on graduation rates or available postdoc or tenure track positions for a particular discipline. Large community sizes require large amounts of memory and are only feasible with distributed computing or a supercomputer, so to keep the processing power small and accessible, I only focused on small communities. So, the model

could not capture the entire field of chemistry, for example, where 2,810 doctorates were awarded in 2018 ("Survey of Earned Doctorates") without using distributed computing. It can capture a smaller subset of chemistry PhDs, however, such as a cohort of PhD graduates in polymer chemistry, where the number of PhDs awarded in 2018 was 132[27]. The model was tested to represent communities of a maximum of 250, but may work for larger sizes, as well.

Once the appropriate number of scientists has been created, the next step is to generate initial publications for each person. The model uses a log normal distribution to generate the initial number of publications. A log normal distribution is described by the mean and standard deviation of a normal distribution where each value from the normal distribution is then raised to *e*. The user input value for the average number of publications is not the mean used to generate the initial values. Instead, the user input average is used to calculate the mean of the normal distribution according to the formula:

$$mean = exp(\mu + \frac{\sigma^2}{2}) \quad (4.1)$$

where the mean is the user input average and σ is the standard deviation, which is held constant at .6. Next, the number of initial publications for each agent is generated by drawing a number *x* from a normal distribution with mean μ and a standard deviation of .6. Then, *x* is used as the exponent for *e*, rounded up to the nearest integer, then 1 is subtracted from the value. This ensures that there are only integer values of publications, and that it is possible for an agent to have no publications in the initial phase.

[27] It also may be that the model can better represent a subdiscipline than a larger discipline. In chemistry, it may be rare for a polymer chemist to cite an organic chemist, for example, but highly likely that they will cite other polymer chemists. The model does not account for any kind of clustering across topics, so it better fits the assumptions of the model to focus on a community where a paper would plausibly cite any other paper, because the topics are close enough. West et al (2013), for example, uses citations to create a multilevel map of the disciplines of science, including fields, subfields, and specialties; the model would be best suited to modeling subfields or specialties.

The choice to use a log normal distribution is supported by some empirical results. There is limited data about what the distribution of publications early in a researcher's career looks like. However, I was able to secure access to two data sets, one from sociology and one from computer science that contain information about the number of publications assistant professors have at time of hire.[28] Both also included information about prior position status; the sociology data set noted whether an individual had been ABD, a postdoc, an assistant professor, or some other position prior to appointment, while the computer science data only includes whether each individual had a postdoc immediately prior to appointment. The sociology data covers from 2003 until 2020; the computer science data covers the 1970's to 2011. I tested both datasets, focusing on the most recent years of each--2007-2011 for computer science and 2016-2020 for sociology--to determine what distribution fit best. I estimated the parameters of the log normal distribution that most closely fit our data. I tested Poisson, gamma, log normal, negative binomial, and Weibull distributions, because they all have zero as a lower bound and infinity as an upper bound[29]. Most people have a low number of publications, but there are always academic superstars with many, many publications; each distribution considered has a long right tail. The log normal distribution provided the best fit for both data sets, including when looking at subsets by year or type of prior position; this was determined visually. The mean number of publications for each data set varied based on what subset was included, but the standard deviation of the

[28] The data set from sociology was discussed in Bauldry (2013), and he has continued to add and maintain the data set. The data set from computer science was reported on in Way, Larremore, and Clauset (2016). Both were acquired through personal communication with the authors.

[29] There is no hard upper limit to the number of papers someone could write; even the most productive author ever could potentially write one more paper. However, there are of course practical limits to the number of papers one author can produce. Ioannidis et al (2018) identify Akihisa Inoue as the researcher with the most publications; he had published 2,566 papers between 1976 and 2016. (Ioannidis fails to note that he is likely to surpass Akihisa Inoue as the most productive; he has over 2,000 publications on his google scholar profile, starting in 1994, meaning he has an average of 80 publications per year to Inoue's mere 64). While an impossible value for the number of publications pulled from a log normal distribution is always possible, it is a very unlikely event, given how numbers are generated from a distribution.

normal distribution that describes the logarithmic normal graph was generally between .5 and .75, so I selected .6 to be the standard deviation for the normal distribution used to generate the initial number of publications. This value could easily be changed but including it as a user input would likely reduce the clarity of the model.

Every publication has exactly one author. It is common in most scientific fields that multi-author collaborations are the norm, so this assumption of the model may seem inaccurate. However, the model only includes a small cohort of scientists entering the field at roughly the same time. There are many other members of the community than just this small set of agents, and those others could be acting as co-authors and collaborators for each of these papers. Given that each agent represents someone working as a postdoc or early faculty member in a lab, it is reasonable to interpret these agents as belonging to distinct labs and therefore being unlikely to collaborate with other agents in the cohort. This assumption may be less realistic as the model progresses; as each agent has been in the community longer, they are more likely to interact and collaborate with a larger network of researchers, including those who incidentally entered the community at roughly the same time. However, the current version of the model does not account for this. It would require additional empirical investigation to what extent such collaborations do take place. If they do, then a future direction for the model could be to integrate such interactions.

Once every agent has an appropriate number of publications generated, citations from each publication are generated. The model generates a directed network/graph in which edges between publications are citations. The number of citations per paper is governed by a user input for the average proportion of papers cited within the community. The average citations may vary based on the field considered, and the model is meant to be general enough to investigate a

variety of scientific fields. To find the average number of citations per publication this represents, the user input proportion is multiplied by the total number of publications. For each publication, a number is drawn from a log normal distribution with the natural log of the average number of citations as its mean and a standard deviation of 1. If the value drawn from the distribution is less than 1, or greater than the total number of publications that can be cited, a new number from the same distribution is generated.

Log normal is used again since citations are also bound at zero but with no upper limit and have a long right tail. There is no empirical evidence or even investigation, to my knowledge, of what the shape of citation counts for early career researchers are. Most empirical work on citation counts looks at a whole field at a time, with researchers of all career stages, and finds that citations obey a power law (Clauset, Shalizi, and Newman 2009). Using log normal to generate citations over time does lead to a distribution that appears to be a power law. Additionally, I was able to collect and analyze a small dataset about citations of early career researchers in astronomy, which provides some empirical support that log normal is a reasonable distribution.

Each publication has an equal chance to be cited in this set up phase. A publication cannot cite itself, but it can cite a publication written by the same author[30]. The number of citations generated by the log normal function for each person represents finding papers to cite and compiling a bibliography for their paper. The number of bibliographic entries in a paper in most scientific fields is quite large, but again, the model is meant to represent a small slice of the community, so it is reasonable to use input values for the average number of citations that are not representative of how many entries are in a bibliography.

[30] The model, if anything, likely undercounts the extent to which self-citation occurs; King et al (2017) find that approximately 10% of citations in a paper are to the author's own previous publications.

4.2.2 *Running the Model*

Once the community is generated, here are two modes that the user can turn on or off that govern how the model runs. First, one can control whether the model considers the gender of the author of a publication when deciding to make a new citation; if this mode is in effect, there is also an input that adjusts how strong the bias towards authors who are men is. Second, one can control whether the model gives grants to those agents who are highly cited. Both modes work independently of each other, so there are four conditions for the model to run under: bias off and grants for top performers off; bias on and grants for top performers off; bias off and grants for top performers on; and bias on and grants for top performers on. I will first discuss the most basic case, where both modes are off.

In each step of the model, three functions are called that represent the work of scientists. First, each agent has an opportunity to receive a grant. Next, each agent has an opportunity to publish. Finally, new publications generate new citations.

The first function determines which agents get a grant in this round of the model. The rate at which grants are given is a user input, expressed as a decimal. The user input determines how many grants are given, then distributes the grants randomly. This way of distributing grants is essentially a lottery; any agent can receive a grant, and previous grants do not increase the likelihood of future grant success.

Next, the agents who have received a grant in this round are able to publish. Every grant results immediately in one publication. The model also does not allow an agent to publish without a grant in that turn of the model. Both assumptions represent a simplification of the scientific process, since of course many grants do result in multiple publications which may not

be published immediately and since scientists with other resources at hand (e.g., startup funds, shared instrumentation, graduate students, etc.) are able to publish without grant funding.

Finally, new publications generate new citations. The number of average citations at any given step is equal to the current number of publications times the user provided value for the average proportions of papers cited. This means that each new paper will cite more and more of the publications in the community, meaning that the average number of citations per paper and per author will increase exponentially instead of linearly.[31] This makes sense, given that as the cohort of researchers age, they will be more likely to be familiar with and build on (or critique) each other's work.

In this phase, unlike the initial phase, each publication does not have an equal chance of being cited. First, a random publication is chosen. Then, one of that publication's edges is chosen. Then, one end of that link is chosen; it could either be the original publication or the existing paper cited by the new paper. Roughly 50% of the time, the original publication will be chosen, resulting in an equal chance for every paper to be selected. However, the other 50% of the time the model will tend to pick publications with more citations; a paper with more edges connecting it to other papers is more likely to 'win' this lottery. The choice of the publication connected to will be driven by the degree distribution of the network--the distribution of the number of edges each publication has. The distribution of new citations in the model will not be uniform as a result. The process can be thought of as the author of the new publication doing a literature search and finding a relevant paper, then choosing to cite either the original paper, or a paper in its bibliography.

[31] This accords with empirical citation patterns across science (Barabási and Albert 1999).

In the case where gender bias is introduced into the model, the grant and publish functions are identical to the case without bias; it is only in generating citations that the gender of the author matters. The evidence for whether the rate of success in receiving grants is affected by gender is mixed. One reason women may receive grants less frequently than men is that women publish less, on average than men. I am interested in investigating one possible mechanism, gendered citation practices, for citation disparities and attempting to show what would happen if only this mechanism were operating. This tells us how much impact this mechanism alone has on citation disparities.

In the biased method of generating citations, the function starts out in the same way as the unbiased case. It chooses a publication at random, then chooses one of its citations at random. Next, however, the function checks what gender the two authors at either end of the link are. If they are the same, it will function just like the unbiased version; it will choose one end at random. If they are different, then the amount of bias is considered. The amount of bias can vary from 1% to 10%, where women are that percentage less likely to receive a citation when the authors selected are of different genders. So, when there is bias in the model, authors of either gender will choose to cite the male author over the female one that percentage of the time.[32]

The other variation in the model changes how grants are allotted. When the option for grants for top performers is turned off, grants will be equally distributed across all agents. But when the user turns on grants for top performers, grants will instead mostly go to the individuals

[32] This would likely reduce the total number of citations for women by less than the amount of bias would indicate. The author only discriminates when the choice is to choose a man over a woman to cite. However, given that there are only author pairs with different genders approximately 33% of the time (at least in early stages of the model), this means a 10% chance to cite a man over a woman is only realized one third of the time, so the 'true' bias level is closer to 3.3%. This value may evolve over the course of the model, however, if the number of heterogeneous author pairs changes over time. But if the parameter for 10% bias does not accurately represent 10% fewer citations for women, this may make a substantial impact to how the results of the model are interpreted.

with the highest citation counts. How the model decides who gets grants under this condition varies based on user inputs. There are two relevant user inputs: the grant rate and the cutoff for who counts as a top performer. When the grant rate and cutoff are the same, all grants are given to the top performers. When the grant rate is lower than the cutoff--such as when the grant rate is 10% but the cutoff is 25%--the model will select 10% of the total number of authors to receive grants, but the only eligible authors are those in the top 25%. When the grant rate is higher than the top performer cutoff--for example, the grant rate is 50% and the cutoff is 30%--the top 30% of performers will all receive grants, and another 20% of the total number of authors will receive grants, although only those not in the top 30% will be eligible for the additional grants (no one can receive more than one grant in this iteration of the model).

4.3 RESULTS

4.3.1 *Parameter Sweep*

A parameter sweep is necessary to determine if results are robust across a number of parameter values. It is done by varying the inputs of the model across a set of plausible values.

4.3.1.1 Choosing Parameter Sweep Values

I ran the parameter sweeps in four subsets of conditions for the model: bias off and grants for top performers off; bias on and grants for top performers off; bias off and grants for top performers on; and bias on and grants for top performers on. This helps to avoid duplication of runs; for example, when both bias and grants for top performers are off, there would be 90 identical runs where the program iterates through the nine combinations for those two values. As

each run gets longer and more complex, this duplication of runs can add significant time to the parameter sweep and increases the likelihood of errors related to inadequate memory.

The values selected and summarized in Table 4.1 are meant to represent plausible values for the inputs. For community size, 50, 100 and 250 are meant to represent a small, medium, and large size community. The average number of publication values of 3, 7, and 10 may be low for most scientific disciplines, but larger values, especially for community size of 250, cause memory errors. The average proportion of citations within the community is difficult to estimate, but even for the smallest community of 50 with the lowest number of average publications of 3, a proportion of .01 means that each paper will have approximately two citations from within the community. For a community size of 250, lower proportions of average citations are used, since using the same proportion of .01 would represent substantially more citation per papers than it does in smaller community sizes, so the proportions are smaller to keep better in line with a similar absolute number of initial citations per paper. Even at the lowest value of three initial publications and with an average citation proportion of .001, each paper will have an average of one citation (and will quickly accumulate more, since the proportion is rescaled to the number of papers after each step of the model). The values for grant rate are given to represent a high, medium, and low value, although 50% may be an unrealistically high grant rate for many funding agencies. The values chosen for grants given to top performers ensures that each case where the grants to top performers is greater than, less than, or equal to the grant rate will occur. The bias level is chosen to represent a high, medium, and low value for bias.

Table 4.1. Initial Values for the Parameter Sweep

Community Size	50, 100, 250
Average publications	3, 7, 10, [25, 50 for community size 50]
Average proportion of citations	[.01, .005 for community sizes of 50 and 100] [.005, .001 for community size of 250]
Grant rate	10%, 25%, 50%
Bias level	1%, 5%, 10%
Percent of top performers to receive a grant	10%, 25%
Gender bias	True, false
Grants given to top performers	True, false

4.3.1.2 Linear Models of Proportional Averages of Citations, Grants, and Funding

I generated a series of linear models[33] to analyze the results of the simulation and determine whether women are negatively impacted by GCPs. This required making linear models for each of the relevant outcomes: grants, publications, and citations. Since the absolute number of grants, publications, and citations would be largely determined by the initial values for the ABM, it would be meaningless to compare the values across all the runs of the parameter sweep without standardizing them in some way, so I calculated the average proportional differences between men and women. For each outcome, the average value for women was divided by the sum of the means of men and women. This ensures that the differences are scaled in terms of the outcome variable, making it possible to compare the differences between grant rates for different average numbers of publications. If men and women have a difference of two

[33] I am forced to use two senses of model in this paper: I discuss both agent-based models and statistical, in this case linear, models. If the model in question is linear, it will always be noted. If there is no additional adjective, the model in question refers to the agent-based model. I have tried to clarify within a sentence or paragraph with the appropriate adjectives if there is any possibility of ambiguity about which type of model is discussed.

publications when everyone has three initial publications on average, it would not be equivalent to a difference of two publications when there is initially an average of 50 publications. Dividing by the sum of the means scales the differences so that they can be more reasonably compared.

Each linear model used the same set of predictors. Gender bias and bias level were treated as one variable; where when gender bias was false, it was treated as being 0. Grants given to top performers and the percent of top performers to receive a grant were also treated as the same variable, where when giving grants to top performers was false, it was treated as grants were given to the top 100% (i.e., anyone in the community was equally likely to receive a grant). I created a categorical variable representing the various combinations of community size and average proportion of citations used. This was necessary because different values for the proportion were used across the different community sizes; although .005 was used for all three, it was the high value for a community of 250 and the low value for communities of 100 and 50. Publication count and grant rate were also included as predictors. The interaction of each other predictor with gender bias were also included in each model. All predictors were treated as categorical variables.

For the first linear model, the difference in citations between men and women across a single run was the outcome variable. First, I investigated what, if any, parameters can be relaxed. If a parameter does not increase the predictive power of the linear model, it can be relaxed. One way to test this is to generate new reduced linear models to compare the predictive power. In the reduced linear model, the parameter value is dropped from the linear model. If the reduced linear model yields equal predictive power to the original linear model, then the parameter can be relaxed. Dropping community size and the grant rate given to top performers (including whether it was on, since the linear model treated grants for top performers off the same as having the

input set to 100%) did not reduce the predictive power of the linear model. So, these inputs, community size and grant rate for a top percentage of performers, did not add to the predictive power of the linear model. This is only true for communities between 50 and 250 and grants given to the top 10-100% of performers, since those parameter values were the ones tested.

Once the linear model parameters that do not add predictive power are removed, the linear model can be fit again, and the results interpreted. The summary table for the linear model on average proportion of citations for women is shown in Table 4.2. Bias level is a significant main effect. When there is bias of five or ten points, women on average have significantly fewer citations; women's proportion of citation was lowered by 1% when bias is set to five versus the case where there is no bias, and 3% lower when bias is set to ten.

There are also a number of significant interactions. As the average number of publications increases, the negative impact on women's average citations decreases; when bias is set to ten, but there are an average of 50 initial publications, women on average perform the same as men. This may be explained by the fact that there is no bias in the set up phase, regardless of what the bias level is set to. The agent-based model runs for 50 turns, so when the average number of publications at the start is large (such as 25 or 50), the citations generated at the start of the ABM without bias protect against any effect the bias level might otherwise have. The bias level also interacts with grant rate. Higher grant rates are worse for women, with the largest effect when bias is the highest, too; at a grant rate of 50% and a bias level of 10%, women have 2% fewer citations than when there is no bias.

Table 4.2. Summary of Linear Model of Differences in Citations

	Estimate	P-value	
(Intercept)	0.50022	< 2e-16	***
Average publications: 7	0.00152	0.500622	
Average publications: 10	-0.00086	0.704221	
Average publications: 25	0.002587	0.452918	
Average publications: 50	-0.00304	0.377031	
Bias level: 1	-0.00017	0.954138	
Bias level: 5	-0.01022	0.000568	***
Bias level: 10	-0.03027	< 2e-16	***
Grant rate: 0.25	-0.0022	0.288728	
Grant rate: 0.5	0.001184	0.568487	
Average publications: 7 x Bias level: 1	-0.00471	0.161047	
Average publications: 10 x Bias level: 1	-0.00066	0.844636	
Average publications: 25 x Bias level: 1	-0.00031	0.949765	
Average publications: 50 x Bias level: 1	0.009659	0.050167	.
Average publications: 7 x Bias level: 5	0.001658	0.622135	
Average publications: 10 x Bias level: 5	0.007036	0.036464	*
Average publications: 25 x Bias level: 5	0.008579	0.08196	.
Average publications: 50 x Bias level: 5	0.018292	0.000209	***
Average publications: 7 x Bias level: 10	0.013692	4.72E-05	***
Average publications: 10 x Bias level: 10	0.020002	2.83E-09	***
Average publications: 25 x Bias level: 10	0.023999	1.15E-06	***
Average publications: 50 x Bias level: 10	0.033971	6.03E-12	***
Bias level: 1 x Grant rate: 0.25	-0.0001	0.973574	
Bias level: 5 x Grant rate: 0.25	-0.00688	0.024898	*

Bias level: 10 x Grant rate: 0.25	-0.00927	0.00251	**
Bias level: 1 x Grant rate: 0.5	-0.00482	0.116201	
Bias level: 5 x Grant rate: 0.5	-0.01107	0.000305	***
Bias level: 10 x Grant rate: 0.5	-0.02083	1.14E-11	***

Significance codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Table 2

The next linear model predicts the proportion of publications for women. The linear model was selected for as in the previous example, and community size, average publications and grant rate did not add to the predictive power of the linear model. The summary table is shown in Table 4.3. Bias level of ten was the only significant main effect, and it lowered the proportion of papers authored by women by .9% compared to when there is no bias. The only significant interaction was between bias level 10 and grants to top performers off; this raised the proportion of papers by women by roughly .9% when compared to the case where publications went only to the top 10% of performers. However, there was not a significant effect for the difference between grants to the top 10% and 25% of performers at the highest bias level. This effect does indicate that women have fewer publications when grants are given to top performers and gender bias is high compared to when grants are given equally.

Table 4.3. Summary of Linear Model of Differences in Publications

	Estimate	P-value	
(Intercept)	0.50102	<2e-16	***
Grants for top performer proportion: 0.25	-0.000106	0.964961	
Grants for top performer proportion: 1	-0.001212	0.637173	
Bias Level: 1	6.43E-05	0.980037	

Bias Level: 5	-0.002726	0.288754	
Bias Level: 10	-0.009516	0.000214	***
Grants for top performer proportion: 0.25 x Bias Level: 1	-0.003896	0.283653	
Grants for top performer proportion: 1 x Bias Level: 1	0.000328	0.930159	
Grants for top performer proportion: 0.25 x Bias Level: 5	-0.002534	0.485654	
Grants for top performer proportion: 1 x Bias Level: 5	0.004061	0.277705	
Grants for top performer proportion: 0.25 x Bias Level: 10	-0.001262	0.728332	
Grants for top performer proportion: 1 x Bias Level: 10	0.009656	0.00987	**

Significance codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Finally, the last linear model focuses on the proportion of grants received by women. The results are very similar to those of the linear model for publications. Community size, average publications, and grant rate also do not add to the power of the linear model. The summary table is provided in Table 4.4. The highest level of bias is the only significant main effect, although the medium bias level of five bias points is nearly significant ($p \approx .07$). High bias reduces the average proportion of grants for women by 1.6%. The only significant interaction is between ten points of bias and grants given to top performers off, where women fare better than when grants are given to the top 10% of performers. The negative effect of bias on grant rates is nearly erased when grants are given equally.

Table 4.4. Summary of Linear Model of Differences in Grants

	Estimate	P-value	
(Intercept)	0.502283	<2e-16	***
Bias Level: 1	-0.000334	0.93477	
Bias Level: 5	-0.007435	0.0688	.
Bias Level: 10	-0.016291	6.72E-05	***

Grants for top performer proportion: 0.25 x Bias Level: 1	-0.001588	0.7834	
Grants for top performer proportion: 1 x Bias Level: 1	-0.000314	0.95786	
Grants for top performer proportion: 0.25 x Bias Level: 5	0.000677	0.90673	
Grants for top performer proportion: 1 x Bias Level: 5	0.007796	0.19	
Grants for top performer proportion: 0.25 x Bias Level: 10	0.0018	0.75532	
Grants for top performer proportion: 1 x Bias Level: 10	0.016714	0.00496	**

Significance codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

4.3.1.3 Linear Model of Percentage of Women in the Top 25%

So far, I have investigated the impact of GCPs on women's grants, publications, and citations. However, given that the proposal being tested is directed towards the top performers, it is important to consider whether giving grants to only top performers would reduce the occurrence of women in the set of top performers. Each run of the parameter value also reported how many men and how many women were in the top 25% of performers; the top 25% was always reported, even when grants were given randomly, and when grants were given to the top 25% and top 10%.

To normalize for differences in community size, the number of women and the number of men reported in the top 25% for each run was converted into a percentage of the top performers who were women. An ANOVA analysis was performed on the data, and the summary table is shown in Table 4.5. The analysis shows that when there is the highest level of bias, women are significantly less likely to be in the top 25% of performers; the percentage of women was reduced by two percentage points for the highest level of bias. Figure 1 shows the differences in the proportion of women across each level of bias; there is a clear downward shift from each box to the next.

Average publications were also a significant main effect. However, the effect size was very small; changing the number of average publications lowered the proportion of women by less than .1 percentage points. Additionally, there are significant interactions between the level of bias and grant rate; when bias and grant rates increase, the percentage of women in the top performers decreases. The effect is largest when bias is high and the grant rate is high; when bias is high and the grant rate is 50%, the proportion of women in the top 25% is reduced by 6 percentage points, relative to no bias and a 10% grant rate. Since the grant rate drives publication rates, this makes sense; when the grant rate is higher and bias is higher, there are more opportunities for new publications to fail to cite women.

Table 4.5. Summary of Linear Model of Women in the Top 25%

	Estimate	P-value	
(Intercept)	4.89E-01	<2e-16	***
Bias Level: 10	-2.28E-02	2.73E-05	***
Bias Level: 1	8.27E-04	0.87879	
Bias Level: 5	-7.29E-03	0.17896	
Grant Rate: 0.25	-2.96E-03	0.58528	
Grant Rate: 0.5	4.31E-03	0.42732	
Average publications	6.45E-04	2.98E-14	***
Average citation proportion: 0.005	-1.74E-03	0.61074	
Average citation proportion: 0.01	-6.04E-03	0.09125	.
Grants for top performer proportion: 0.1	4.85E-03	0.0738	.
Grants for top performer proportion: 0.25	1.60E-02	3.77E-09	***
Bias Level: 10 x Grant Rate: 0.25	-2.02E-02	0.00861	**
Bias Level: 1 x Grant Rate: 0.25	-6.66E-03	0.3851	

Bias Level: 5 x Grant Rate: 0.25	-1.47E-02	0.05478	.
Bias Level: 10 x Grant Rate: 0.5	-6.30E-02	2.39E-16	***
Bias Level: 1 x Grant Rate: 0.5	-1.45E-02	0.05796	.
Bias Level: 5 x Grant Rate: 0.5	-3.37E-02	1.13E-05	***

Significance codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

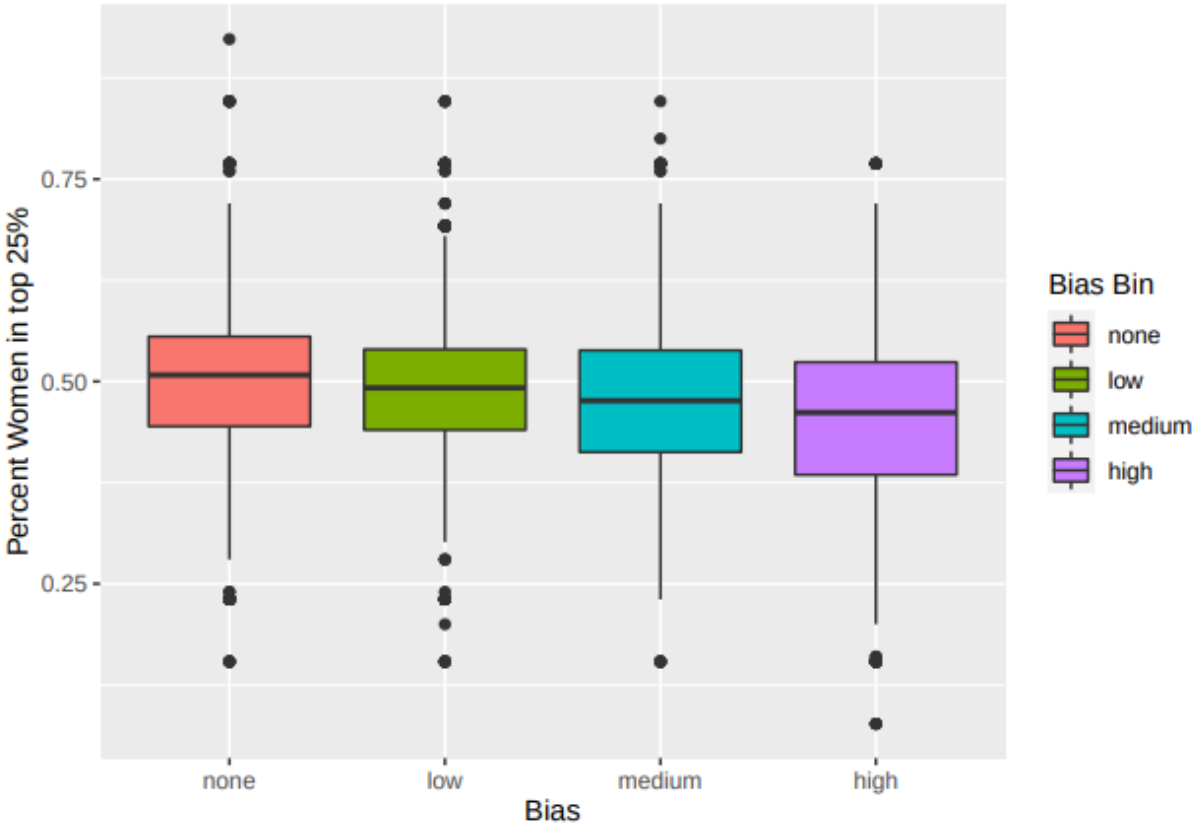

Figure 4.1 Women in the top 25% as bias increases

4.3.2 *Experiment: Astronomy*

The agent-based model can be used to test what might happen to a community if a policy to give grants to top performers was instituted. The agent-based model can use input values

about a specific scientific discipline and provide a prediction about what might occur if that discipline implemented the policy.

I selected astronomy for several reasons. First, there are roughly 150 graduates from PhD programs every year, and approximately 60-75 faculty positions advertised per year (Trump et al 2020). The small size makes it easier to collect information about what tenure track assistant professors look like at the time of hire, and astronomy has a jobs wiki that records both searches underway and, in some cases, the applicant who received the job. Additionally, the small size is well within the community size where the ABM runs efficiently, so a large number of replications can be done. Second, while there are several possible funding sources in astronomy, the two most prestigious for US-based researchers are NSF and NASA, both of which publicly report their grant success rates. As of 2014, NSF funded 17% of grant applications, while NASA funded approximately 20% (Astronomy and Astrophysics Advisory Committee, 2016). Finally, Caplar, Tacchella, and Birrer (2017) shows that papers authored by women in the top five astronomy journals receive 10% fewer citations than predicted, given other characteristics of their papers; astronomy is a field where gendered citation practices seem to occur. Therefore, it is likely good estimates are available for all the parameters needed for the model.

4.3.2.1 Choosing Model Values for the Experiment

First, I determined the appropriate community size for the model. There were approximately 70 tenure track positions listed on the astronomy jobs wiki in the 2019-2020 academic year at universities in the US and Canada.[34] Positions at universities outside of the US and Canada were excluded in the count, as were research institutes. The astronomy jobs wiki is

[34] The wiki also includes information about postdocs. Given the nature of astronomy, many of these postdocs are based internationally at various telescopes, but many candidates will be from the US and Canada, making it difficult to determine which institutions to include or exclude.

an unofficial listing of posted jobs, so it may contain inaccuracies, but the wiki is clearly used by many and frequently updated. There are notes about when first round interviews went out, updates on additional application materials requested, lists of campus visit candidates, and sometimes who the job was offered to and accepted by. Job applicants have a strong interest in having a centralized resource for job listings and information, so despite the informal nature of the list, it is likely the best list of jobs and hired candidates available.[35] Given that the model is meant to represent researchers when starting an academic position such as a postdoc or assistant professor, 70 is a reasonable value for the community size parameter.

The values for the grant rate can easily be taken from the average success rates for NSF and NASA grants, 17% and 20% respectively. Given the closeness of the values, I choose to only test the NASA grant rate of 20% to reduce the number of runs necessary, since the grant rate does not seem to have a large impact on the outcomes of the model. I choose to test what would happen if grants were given to the top 10%, 20%, and 30%; this ensures that there is a case where the grant rate is higher, lower, and (for the runs using the 20% grant rate from NASA), equal to the percentage of top performers who receive grants.

Caplar, Tacchella, and Birrer (2017) provides good estimates for the values of bias points to be tested. They found that in the top five astronomy journals, men are cited more than women. The authors used a random forest algorithm to model the number of citations based on "the seniority of the first author, the number of references, the total number of authors, the year of publication, the journal of publication, the field of study and the geographical region of the first

[35] Some of these searchers were ultimately cancelled due to Covid-19, so there were not 70 tenure track hires. However, it is unclear how many of the job searches were cancelled, and it is unclear what impact Covid-19 may have had on who was hired and when. While this was not a typical academic job cycle, it is not possible to collect data on any previous year's hires, at least with the expectation that the data collected will represent what their bibliometric profile looked like at time of hire.

author's institution" (2, 2017) and trained it on a set of papers authored by men. Their model predicted the numbers of citations expected based on the non-gendered characteristics of a second set of papers, which were written by women, which the authors then compared with the actual numbers of citations made to these papers. They found that women receive 10% fewer citations than expected. Their model predicts that, given the features of the papers written by women, women should receive 4% more citations than men, but instead men receive 6% more citations. Given that the bias level functions as the percentage by which women are less likely to be cited, this makes it reasonable to set the bias level at 6%, to represent the 6% more citations men receive, and with 10%, to represent the 10% more citations they predict women should receive. Additionally, I tested with just one bias point to see what would happen with the minimum amount of gender bias.

For the values for average number of publications, I collected a small dataset based on information from the astronomy jobs wiki. 24 of the 70 listings indicated who received the position, so I compiled a list of names and new institutional affiliations. Then, I looked for Google Scholar profiles for each of these researchers. Only 18 researchers had Google Scholar profiles. For each of these 18 profiles, I collected the number of total citations as calculated by Google Scholar, as well as information about how many peer reviewed papers they had written, how many citations each paper had, and when their first peer reviewed academic publication was published. Most Google Scholar profiles included a number of conference presentations as well as papers posted on arXiv, and these entries were not counted in the peer reviewed paper totals, even if they had citations. I also tried to exclude any papers that seemed to be obviously included in error; one profile had a few medical research papers from the 80s that could not be the work of the same researcher who wrote papers in astronomy, for example.

It is unclear how representative this sample may be of assistant professor hires for 2019-2020, given that is a convenience sample of researchers whose names were listed on the jobs wiki website as receiving a position (many of whom likely contributed their own names to the website) and who also had a Google Scholar profile; these profiles may systematically differ from those who were hired but not listed on the jobs wiki, or from those who were listed but did not have Google Scholar profiles. This data was collected in November 2020, so it includes publications from 2020 that were not listed on the CVs of applicants when they compiled their application materials in the fall of 2019. However, despite its limitations, this data should represent a snapshot of some assistant professors in astronomy early on in their professorial careers.

The average number of publications for this group was 35, with a standard deviation of 33.37; the minimum number of publications was 7 and the maximum was 139. Several distributions were fit to the data including gamma, Poisson, and log normal distributions, and the log normal gave the best fit[36]; the parameters for the log normal function are a mean of 3.31 and a standard deviation of .696. I used the formula to calculate the mean of the log normal function:

$$mean = exp(\mu + \frac{\sigma^2}{2}) = 34.3 \quad (4.2)$$

Given the closeness of the values for the log normal distribution and the actual data, I used 35 for the model parameter. The model is programmed to use .6 as its standard deviation and it is not an input that can be changed without changing the code, but I left it at .6, since the values are close,

[36]Log normal was the best fit, but Weibull, gamma, and negative binomial were also good fits; Poisson was a very poor fit. This dataset is very small, so it is easier to fit a curve to it than a large dataset but given that log normal was also the best fit for the data from computer science and sociology, log normal is most likely the best way to describe the shape here as well.

and if anything, there is likely less spread to the actual number of initial publications, given that .696 comes from a small sample with a large range.

The final parameter for the model is the average proportion of citations. I collected data about total citations for each individual in the astronomy data set as well as information about citations for each paper. The average number of total citations on peer reviewed papers per individual was 2,655 with a standard deviation of 3,412; the minimum was 141, the maximum was 12,940, and the median was 1468. The average citations per paper was 75.5 with a standard deviation of 195; the minimum was 0, the maximum was 2032, and the median was 29.[37]

The large spread of citations can be partially explained by the fact that researchers have had different length careers. I collected information about the first year each individual published a peer reviewed article in astronomy, since papers from five years ago will generally have more citations than papers from one year ago. Using this information, I was able to calculate an average number of citations per paper per year: 19.5, with a median of 12.9, a minimum of 5.5 and a maximum of 104.3.

However, this data does not immediately help when determining what is a reasonable value for the *proportion* of papers in the model (i.e., written by other early career researchers) that are cited by another researcher in the model. The Astrophysics Data Service (ADS) has 13 million database records including peer reviewed and non-peer reviewed publications and preprints. If every article has 1,000 citations, it would only cite a vanishingly small proportion of the total articles. However, those 13 million entries cover publications since essentially the start of the discipline. In a rapidly changing and developing field, it is likely that more recent articles get the bulk of citations. An analysis of citation practices in astronomy from 1981 shows that

[37] The data sets from sociology and computer science that were discussed in the methods section did not have any information about the number of citations per paper at the time of hire (or at any other time).

most articles get the bulk of their citations in the first five years after publication (Abt, 1981); that timeline may be even more compressed at this point. So, it is likely that early career authors are citing one another, given that they are publishing in the same time frame and citations tend towards more recent articles.[38]

But when experimenting with values for the proportion of articles, it became clear that values such as .01 were likely far too high. When exploring the possible parameters for astronomy, I found that using a .01 proportion of papers within the model cited lead to authors having an average of roughly 2,000 citations in total at the start of the model. While this is not so far off from the average number of citations per individual calculated above, it does not make sense given the assumptions of the model: the model only represents the proportion of papers cited from within the model, not the entire astronomical community. So, the average number of citations per person should be much lower than the observed number, given the assumption of the model[39]. As plausible estimates, I choose .005 and .001 since these result in average initial citations at more reasonable values.

Again, the experiments were run in four groups: bias off and grants for top performers off; bias on and grants for top performers off; bias off and grants for top performers on; and bias on and grants for top performers on. Each iteration of the model was run 1000 times, resulting in 32,000 runs.

[38] The model does not explicitly account for age in the calculations for likelihood to cite. If anything, it is more likely to cite older papers, since they will have a higher degree of connection.

[39] How much lower is, of course, an open question for which it would be very difficult to collect data on. I also want to avoid putting too much confidence in the value of average citations per person based on such a small sample that may be unrepresentative.

Table 4.6 Initial Values for the Astronomy Experiment

Test Parameters: Astronomy	
Community Size	70
Average publications	35
Average proportion of citations	.005, .001
Grant rate	20%
Bias points	1, 6, 10
Percent of top performers to receive a grant	10%, 20%, 30%
Gender bias	True, false
Grants given to top performers	True, false

Significance codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

4.3.2.2 Results of the Astronomy Experiment

Although there were supposed to be 32,0000 runs of the experiment, five are missing.[40] Thus, the analysis was performed on 31,995 runs. I used the same methods as described above: an ANOVA was performed using all the parameters that varied, then compared to a series of reduced linear models to determine, which, if any parameters could be safely dropped from the

[40] The five that are missing are all from the same combination of parameters and occurred sequentially. Either there was some issue internal to what happened in each run, or the issue was external to the simulation. There were no errors recorded to immediately explain what happened to these five runs; NetLogo does record errors when the problem is inadequate memory, for example, although it may not exhaustively report errors. When a simulation gets stuck in an infinite loop, NetLogo will not report it as an error, but in my experience, it also does not terminate the program at any point. This set of simulations ran to their (seemingly) successful endpoint without having to be terminated or otherwise interfered with. Given that the five runs occurred sequentially in a group of 16,200 runs, it seems unlikely that five sequential runs ran into the same error and that all were an error that Netlogo fails to report. I do not know what other kinds of errors external to the simulation likely culprits for these five runs failing might be, but I do not believe it's a problem internal to the simulation. While the five are all from the same combination of parameters, the difference between 1,000 results and 995 results to analyze in terms of the power of the model is negligible, so the runs were not repeated and added to the analysis.

linear model. The outcomes were again the average proportion of citations, publications, and grants for women and the percent of women in the top 25% of all scientists by total citation counts. The summary table for the ANOVA performed on the reduced linear model is reported for each outcome.

For the average proportion of grants and publications, there was no significant impact for any of the parameters, including bias. Even at the highest level of bias, it did not have a significant effect on the difference between men and women.

For the average proportion of citations, bias level did have a significant effect at the medium and high level. As the bias level increased, the average proportion of citations going to women decreased. However, the effect is quite small; even at the highest level of bias, women receive only .5% less of the proportion of citations. The summary table is shown in table 4.7.

Table 4.7 Summary of Linear Model of Differences in Citations for Astronomy

	Estimate	P-value	
	Estimate	P-value	
(Intercept)	0.49938	<2e-16	***
Bias Level: 1	-0.000329	0.601	
Bias Level: 6	-0.003582	1.26E-08	***
Bias Level: 10	-0.005395	<2e-16	***

Significance codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

For the proportion of women in the top 25% of researchers by total citations, bias level did have a significant effect at the highest two levels. The summary table for this linear model is provided in Table 4.8. Again, the effect size is fairly small; women still represent 49% of the top 25% of performers, even at the highest level of bias.

Table 4.8 Summary of Linear Model of Women in the top 25% for Astronomy

	Estimate	P-value	
(Intercept)	0.498938	<2e-16	***
Bias Level: 1	-0.001028	0.523	
Bias Level: 6	-0.00825	3.01E-07	***
Bias Level: 10	-0.011362	1.76E-12	***

Significance codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

The agent-based model so far has assumed that there are equal numbers of men and women. But in astronomy, more men than women graduate with PhDs each year. Women represented 40% of the PhD graduates in 2014 and 40% of recently hired faculty in astronomy in 2016 (Porter and Ivie, 2019). Women (appeared to) represent 10 out of the 24 researchers identified as tenure track assistant professors starting in 2020 on the astro job wiki. A lower percentage of women could amplify the effects of bias against women, as shown, for example in Rubin and O'Connor (2018), where biased strategies against minority groups are more likely to develop as the size of the minority group decreases.

I added a parameter that allowed changing the percent of women in the agent-based model, and ran 8,600 additional simulations, where women represented 40% of the researchers in the ABMl. Since the average proportion of citations did not affect the outcomes, it did not vary. Instead, only bias level and grants for top performers varied in the same way as described in the table above.

When adding these additional observations to the dataset of 31,995 runs where women represented 50% of the researchers, the results were almost identical. Bias level was still the only significant main effect for the average proportion of citations; publications and grant average

proportions were unaffected by any of the parameters. The summary chart for the linear model of average proportion of citations is shown in Table 4.9.

Table 4.9 Summary of Linear Model of Differences in Citations for Astronomy with fewer women in the model

	Estimate	P-value	
(Intercept)	0.499732	<2e-16	***
Bias Level: 1	-0.000348	0.616	
Bias Level: 6	-0.002858	3.96E-05	***
Bias Level: 10	-0.00597	<2e-16	***

Significance codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

The percentage of women was a predictor of the composition of the top 25% of most cited researchers, but essentially having 10% fewer women in the ABM simply reduced their percentage in the top 25% by 10%. The summary table for the linear model is provided in table 4.10. Lowering the percentage of women in the agent-based model did not seem to make a substantial difference to the results for astronomy, although the percentage of women in a community may begin to make a difference as the percentage of women decreases.

Table 4.10 Summary of Linear Model of Women in the Top 25% in Astronomy with fewer women

	Estimate	P-value	
(Intercept)	0.497484	<2e-16	***
Percent Women: 40	-0.096837	<2e-16	***
Bias Level: 1	-0.000557	0.75688	
Bias Level: 6	-0.005526	0.00212	**

Bias Level: 10	-0.010546	4.52E-09	***

Significance codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

4.4 DISCUSSION

The two main questions this agent-based model intended to answer were: does bias alone, independent of any differences in productivity or grant attainment by gender, negatively affect the careers of women scientists in the model? And will productivity based grant proposals have a disparate impact on women if implemented? Based on the model output, the answer to the first question is yes, and the answer to the second question is maybe, under some circumstances. The impact on women of only funding the most highly cited researcher will depend on the amount of citation bias in the community, and the amount and kinds of collaboration that exist in the field.

In both the parameter sweep and experiment, bias consistently lowered the average proportion of citations for women and lowered the percentage of women in the top 25% of performers, as shown in tables 4.2 and 4.5. This shows that women are affected negatively by gender bias in their discipline in terms of citations and being highly cited relative to male peers.

In the parameter sweep, bias also impacted the average proportion of both grants and publications, as shown in tables 4.3 and 4.4. Women were less likely to receive grants, and since grants lead directly to publications, bias also lowered the rates at which women published relative to men. Awarding grants to top performers further reduced the rates of grants and publications for women. The base parameterization of the model awarded grants to the top 10% of performers. When grants were given equally to all researchers, there was a positive effect on grant and publication proportions for women. The significant interactions between the highest level of bias and grants given equally to every researcher show that women were disadvantaged

when grants were given to only the top 10%. This interaction was only seen at the highest level of bias, but lower levels of bias did not differ enough to constitute a main effect, so it is not surprising that there was no significant interaction at lower levels of bias, when there was no significant main effect, either.

Additionally, the difference between giving grants to the top 10% and top 25% at the highest level of bias did not have a significant interaction, meaning that women are disadvantaged at either threshold.[41] So, in the parameter sweep, this does support that women may be disadvantaged by productivity based grant funding, and provides some evidence against implementing a new grant funding system.

However, when using values taken to represent astronomy, the negative effect of bias for women and any interaction with giving highly cited researchers grants in terms of grants and publications disappears, while the effect of bias on citations and proportion of top performers persists. So, the evidence against implementing productivity based grants may be weaker than the parameter sweep indicates.

There may be other explanations for why the results from astronomy look different than the results from the parameter sweep. The initial conditions of the experiment seem to be driving the results. In particular, the relatively high number of initial publications coupled with a relatively low grant rate may have a somewhat protective effect on the outcomes for women. There is no bias in how citations are allotted when the model is generated. The rate at which women and men occur in the top 25% at the start of the model will be roughly equal. When the

[41] Erosheva et al (2020) note that even small differences--differences of 0.3 points on a scale of 9--in the evaluation of candidates as measured by impact scores can affect grant success when grant rates are low, around 10% or 20% for the NIH grants they consider. While the NIH funding process is substantially more complex than the one implemented in the model, they show that small differences between candidates can result in disparate grant outcomes.

grant rate is only 20% and the model runs for 50 turns, each author will add an average of 10 publications over the course of the model, which is lower than what one would expect over an entire academic career.[42] So a large proportion of the total publications are generated in the initial stages of the model, and there are fewer publications generated later, when bias can occur. Many women who started the model in the top 25% will persist, even if there is a high level of bias, simply due to the size of their starting citations, and the relatively small effect of any additional biased citations.

As a result, the evidence from the model does not clearly indicate that women will be disadvantaged by citation-based grant funding, where grants go to the most highly cited researchers. But the difference between the results of the parameter sweep and astronomy experiment do show some potentially important implications for such a policy.

First, the amount of citation bias makes a difference. In fields where women and men are cited at similar rates, even if there are small but significant differences, these disparities may not result in statistically significant differences over longer time scales and may not have a negative impact on women's careers. Lower productivity and lower grant attainment may also impact citation values in the real world, but the impact of gendered citation practices alone is not significant at low levels. If women and men have similar publications and grant success rates and there are only small differences in citation rates in a field,[43] emphasizing citation metrics in grant decisions is less likely to disadvantage women substantially.

---

[42] This probably means that 50 rounds of the model are not adequate to represent what happens at say, retirement, but it may be enough to represent a career at time of tenure or promotion to full professor. Some of the possible ways to address this issue will be discussed in the next section.

[43] Exactly how small these differences might need to be is unclear. The model uses the level of bias to represent the extent to which women are less likely to be cited. The empirical studies that show citation bias report on their findings in a variety of ways, sometimes as a magnitude of difference or sometimes as a percentage. There is not a clear indication of what constitutes a low level of bias, beyond subjective judgments. For example, Caplar, Tacchella, and Birrer (2017) report a 10% gap between the expected and actual rates of citation of women. This sounds like a high level of bias and given the publication and citation patterns of astronomy discussed in this paper,

Additionally, having many early publications with no citation bias has a seemingly protective effect on future grant success when citations are used to determine grants. In astronomy, the bulk of publications were generated at the start of the model, and the function for citation bias was only used in subsequent rounds. The initial publications are meant to represent those written as postdocs or graduate students. In astronomy, publications with a large number of authors are the norm. The model treated every paper as single authored (assuming any co-authors would not be represented within the model). Depending on the authorship norms of astronomy, it may be that graduate students and postdocs are in a less prestigious authorship position, where their names are likely not generally used to determine whether to cite a paper or not.[44] So, it may be that in their early publications, any gender differences in citations are unlikely to appear. Presumably, as the researchers in the model age, they are more likely to move into prominent authorship positions, where gendered citation practices are more likely to be in effect. In fields with similar authorship norms of large, multi-author publications, this protective effect of many early publications somewhere in the authorship list may line up with assumptions of the model and lessen the impact of any later gender bias on grant success when citations are used.

Further, astronomers have a high degree of collaboration across domestic and international universities. Adams et al (2007) shows that almost 65% of astronomy papers published in major journals in 1999 included a research team with multiple affiliations within the

10% would lead to a large magnitude of difference. Using the median value for early career researchers reported above of 1468, a 10% difference would reduce a researcher's citations by 140, and again, this is just at the point where they are leaving grad school or a postdoc. On the other hand, women in epidemiology receive an average of four fewer citations per paper than men for articles published in the top six journals (Schisterman et al. 2017). This sounds like a much more modest difference, especially if average citations per paper is high in epidemiology. So, epidemiology may be below the threshold for small enough differences, and astronomy is above it but it would require a more exact understanding of how bias level in the model maps onto the real world.

[44] The model is agnostic as to why a researcher chooses to cite a paper written by a man instead of one written by a woman. The topic, publication venue, and relevance to the current paper all surely make a difference to the choice to cite a given paper in the 'real world'. But whether the cause is explicit or implicit bias, it seems most plausible that it is directed at the researchers in the more prominent authorship positions.

US and 24.5% of papers included at least one internationally affiliated author. If these large, multi-authored papers often contain multiple researchers from the same cohort, spread out across multiple domestic labs, that would violate the assumptions of the model. However, it is unclear if relaxing that assumption would substantially change the outcome and testing it would require a more complex model.

Despite the definitive answer from the model, I would still oppose implementing a proposal to give grants to the most highly cited researchers. The results of the parameter sweep show that women may be disadvantaged by such a proposal. It is unclear how generalizable the results from astronomy might be, given some of the features of astronomy highlighted above. It would be difficult for a granting agency to fund the most highly cited researchers without perpetuating gender bias, unless more is known about what can mitigate the effects of GCPs.

4.5 Conclusion

4.5.1 *Policy Implications*

In this paper, I have shown that GCPs alone can reduce the productivity of female researchers in the model, and that proposals to give grants to the most highly cited researchers may disproportionately impact women. As a result, I conclude that such proposals should not be implemented.

However, there are a number of important policies that the scientific community--including granting agencies, institutions, graduate schools, and scientists themselves--might implement as a result of the findings of the model. Women are negatively impacted by gendered citation practices, and scientists should look for ways to mitigate the impacts of GCPs and to prevent the occurrence of GCPs in the future. Some ways to mitigate the impacts of GCPs

include focusing on how citation metrics are used, encouraging collaboration, and increasing women's opportunities to publish. Some ways to prevent GCPs might include improved training for researchers about how to find sources, implementing effective implicit bias training, ensuring women are included on journal editorial boards and regularly serve as peer reviewers, and introducing requirements to check publications for equitable citation practices.

GCPs currently exist in the publication records of many scientific disciplines and even if we could ensure no bias in citation metrics going forward, it is still necessary to consider ways to reduce their effects now. This paper tested what might happen if grant agencies only used citation metrics and I concluded such a policy would likely harm women. But citations are already used in grant decisions as well as in hiring and promotion. Citations are generally assumed to be an effective proxy of research impact, but GCPs show that other factors intervene. If committees continue using citation metrics, they should know that there is a real possibility of bias and that citations may not be an effective way to compare the impact of two scientists' work. Increasing collaboration between men and women and between early career and more senior researchers also can help. More diverse collaborators may reduce the possibility that the name of a female researcher in the list of authors impacts the decision of who to cite. Encouraging women to publish more, especially at early career stages, gives more chances to increase their overall citation rates and specific measures such as h-index which also considers the number of publications.

To prevent GCPs, scientists should develop better research and citation practices and journals should do more to encourage diverse citations. Early career researchers need better training to find sources and evaluate what to include in a bibliography. Librarians are an incredibly valuable resource, readily available at every academic institution, but are

underutilized in teaching explicit research skills to do exhaustive literature searches (Delaney and Bates, 2018). Researchers can use resources such as Cite Black Authors ("Cite Black Author Database") and databases of underrepresented scholars such as the UP Directory ("UPDirectory") to seek out work by underrepresented authors. Journals also shoulder some of the responsibility of ensuring women are cited when appropriate. Women should be regularly used as peer reviewers, so that they can highlight relevant literature the author missed[45]. Similarly, improving the representation of women serving in editorial roles at journals can also provide another opportunity to inform the author of relevant research they may have missed. Journals should also consider developing explicit procedures for looking at citation lists and ensuring they adequately include the relevant literature. Trainings that are effective at reducing implicit biases can also reduce the likelihood that GCPs will occur[46].

While these interventions are intuitively plausible, their impact on the prevalence of gendered citation practice is unknown and researchers should investigate whether and the extent to which they succeed. But the impacts of GCPs show that the scientific community has a responsibility to act.

4.5.2 *Future Directions for the Model*

This model is simple, and there are a number of ways to extend and improve the model. The model assumes that all scientists and all papers are created equal. There are no differences between each scientist in the model beyond having been assigned a gender. However, scientists do vary in the quality of their work. Some papers are better and more informative than others and

[45] There are of course significant structural barriers to increasing the rates at which women serve as peer reviewers, such as the Covid-19 pandemic (see Squazzoni et al, 2020). One proposal is to pay peer reviewers for their work, a proposal I will not argue for here, but I am strongly in favor of.

[46] For examples of effective types of training, see Devine et al (2012) and Jackson, Hillard, and Schneider (2014).

deserve to be read and cited more frequently. Some grant proposals are stronger than others and deserve to be funded over weaker ones. While there is no reason to think women or men are systematically better at science than one another, there is good reason to think that talent, access to resources, and luck are not distributed equally to all members of the community. Implementing a metric for quality of scientists, papers, or grant proposals would be one way to attempt to capture these complexities.

Deciding how scientific quality or importance is distributed across scientists, papers, or grant applications is, of course, a challenging empirical proposition. However, building any model requires making assumptions about agents, and there are certainly a variety of assumptions about how scientists differ that can be defended or backed with at least some empirical support. For example, NIH has data about the scores for each grant submitted for a particular grant type. With this data, it would be possible to find the shape of the distribution of grant scores, and to build into the model the process of assigning grant scores to each scientist based on this distribution. Additionally, a future version of the model could use the relationship between past and future grants to give more grants to past recipients.

The model currently does not allow more than one publication resulting from a single grant. This likely is not an accurate assumption for most fields; grants generally lead to many publications, especially if they are multi-year grants. Scientists also may publish papers that are not specifically tied to any particular grant, which this version of the model also does not represent. There are many possibilities for how to address this concern but allowing scientists to publish more frequently for each grant or providing a function to publish without earning a grant in each cycle would increase the accuracy of the model.

The model also only represents a small slice of the community. There are no papers that are not tied to a scientist in the model. Papers without authors in the community being modeled could be added to the model in the initial stage and in subsequent steps of the model that could also cite and be cited by the researchers. This would also make it easier to estimate what the average proportion of citations should be, since it would be more closely tied to the average number of citations per paper for a field. Further, the only members are those who enter the community at the same time; there are no pre-existing members, no one is added later, and no one leaves the model. Adding more scientists to the model would increase its complexity quite substantially but may also provide more accurate results.

While the model could be extended in any number of these ways, it has answered the primary questions of this paper. Citation bias lowers the proportion of citations, grants, and publications women earn over the course of the model. The model does not definitively show that women would or would not be disproportionately affected by citation-focused grant funding but does provide some evidence about the conditions under which it is more likely to negatively affect women.

BIBLIOGRAPHY

Abbott, Alison, David Cyranoski, Nicola Jones, Brendan Maher, Quirin Schiermeier, and Richard Van Noorden. "Metrics: Do Metrics Matter?" *Nature* 465, no. 7300 (2010): 860-862.

Abt, Helmut A. "Some trends in American astronomical publications." *Publications of the Astronomical Society of the Pacific* 93, no. 553 (1981): 269.

Adams, James D., Grant C. Black, J. Roger Clemmons, and Paula E. Stephan. "Scientific teams and institutional collaborations: Evidence from US universities, 1981–1999." *Research policy* 34, no. 3 (2005): 259-285.

Astronomy and Astrophysics Advisory Committee. 2016. Competed Grant Success Rates in US Astronomy and Astrophysics.

Auchincloss, Amy H., and Ana V. Diez Roux. "A new tool for epidemiology: the usefulness of dynamic-agent models in understanding place effects on health." *American journal of epidemiology* 168, no. 1 (2008): 1-8.

Avolio, Beatrice, Jessica Chávez, and Carlos Vílchez-Román. "Factors that contribute to the underrepresentation of women in science careers worldwide: a literature review." *Social Psychology of Education* 23, no. 3 (2020): 773-794.

Barabási, Albert-László, and Réka Albert. "Emergence of scaling in random networks." *Science* 286, no. 5439 (1999): 509-512.

Bauldry, Shawn. "Trends in the research productivity of newly hired assistant professors at research departments from 2007 to 2012." *The American Sociologist* 44, no. 3 (2013): 282-291.

Bertin, Marc, Iana Atanassova, Cassidy R. Sugimoto, and Vincent Lariviere. "The linguistic patterns and rhetorical structure of citation context: an approach using n-grams." *Scientometrics* 109, no. 3 (2016): 1417-1434.

Blood, Emily A., Nicole J. Ullrich, Dina R. Hirshfeld-Becker, Ellen W. Seely, Maureen T. Connelly, Carol A. Warfield, and S. Jean Emans. "Academic women faculty: are they finding the mentoring they need?." *Journal of Women's Health* 21, no. 11 (2012): 1201-1208.

Borgerson, Kirstin. "Amending and defending critical contextual empiricism." *European Journal for Philosophy of Science* 1, no. 3 (2011): 435-449.

Bornmann, Lutz, Rüdiger Mutz, and Hans-Dieter Daniel. "Gender differences in grant peer review: A meta-analysis." *Journal of Informetrics* 1, no. 3 (2007): 226-238.

Bright, Liam Kofi. "Decision theoretic model of the productivity gap." *Erkenntnis* 82, no. 2 (2017): 421-442.

Bruch, Elizabeth, and Jon Atwell. "Agent-based models in empirical social research." *Sociological methods & research* 44, no. 2 (2015): 186-221.

Caplar, Neven, Sandro Tacchella, and Simon Birrer. "Quantitative Evaluation of Gender Bias in Astronomical Publications from Citation Counts." *Nature Astronomy* 1, no. 6 (2017): .

Casadevall, Arturo and Ferric C. Fang. 2015. "Impacted Science: Impact is Not Importance." *mBio* 6 (5): e01593.

Carli, Linda L., Laila Alawa, YoonAh Lee, Bei Zhao, and Elaine Kim. "Stereotypes about gender and science: Women≠ scientists." *Psychology of Women Quarterly* 40, no. 2 (2016): 244-260.

Ceci, Stephen J., Donna K. Ginther, Shulamit Kahn, and Wendy M. Williams. "Women in academic science: A changing landscape." *Psychological science in the public interest* 15, no. 3 (2014): 75-141.

Ceci, Stephen J., and Wendy M. Williams. "Sex differences in math-intensive fields." *Current directions in psychological science* 19, no. 5 (2010): 275-279.

"Cite Black Authors.", https://citeblackauthors.com/.

Clauset, Aaron, Cosma Rohilla Shalizi, and Mark EJ Newman. "Power-law distributions in empirical data." *SIAM review* 51, no. 4 (2009): 661-703.

Code, L. (1991) *What Can She Know? Feminist Theory and the Construction of Knowledge*, Ithaca: Cornell University Press.

Collins, P.H. (1991) *Black Feminist Thought: Knowledge, Consciousness, and the Politics of Empowerment*, New York, NY: Routledge.

Delaney, Geraldine, and Jessica Bates. "How can the university library better meet the information needs of research students? Experiences from Ulster University." *New Review of Academic Librarianship* 24, no. 1 (2018): 63-89.

Devine, Patricia G., Patrick S. Forscher, Anthony J. Austin, and William TL Cox. "Long-term reduction in implicit race bias: A prejudice habit-breaking intervention." *Journal of experimental social psychology* 48, no. 6 (2012): 1267-1278.

Dotson, Kristie. "A cautionary tale: On limiting epistemic oppression." *Frontiers: A Journal of Women Studies* 33, no. 1 (2012): 24-47.

Dotson, Kristie. "Conceptualizing epistemic oppression." *Social Epistemology* 28, no. 2 (2014): 115-138.

Douglas, Heather. *Science, policy, and the value-free ideal*. University of Pittsburgh Pre, 2009.

Elliott, Kevin C. "Direct and indirect roles for values in science." *Philosophy of Science* 78, no. 2 (2011): 303-324.

Epstein, Joshua M. 2006. *Generative Social Science : Studies in Agent-Based Computational Modeling*. Princeton: Princeton University Press.

Erosheva, Elena A., Sheridan Grant, Mei-Ching Chen, Mark D. Lindner, Richard K. Nakamura, and Carole J. Lee. "NIH peer review: Criterion scores completely account for racial disparities in overall impact scores." *Science Advances* 6, no. 23 (2020): eaaz4868.

Farrell, Lynn, Andy Cochrane, and Louise McHugh. "Exploring attitudes towards gender and science: The advantages of an IRAP approach versus the IAT." *Journal of Contextual Behavioral Science* 4, no. 2 (2015): 121-128.

Fehr, Carla. 2011. "What is in it for Me? the Benefits of Diversity in Scientific Communities." In *Feminist Epistemology and Philosophy of Science*, 133-155: Springer.

Fricker, Miranda. Epistemic Injustice: Power and the Ethics of Knowing. Oxford; New York: Oxford University Press, 2007.

Geraci, Lisa, Steve Balsis, and Alexander J. Busch Busch. 2015. "Gender and the H Index in Psychology." *Scientometrics* 105 (3): 2023-2034.

Gero, Joan. "The Social World of Prehistoric Facts: Gender and Power in Paleoindian Research." In *Interpretive Archaeology: A Reader*, 304-16. London: Leicester University Press, 2000.

González-Álvarez, Julio and Teresa Cervera-Crespo. 2017. "Research Production in High-Impact Journals of Contemporary Neuroscience: A Gender Analysis." *Journal of Informetrics* 11 (1): 232-243.

Haraway, D. J. "Primate visions: gender, race, and nature in the world of modern science." (1989).

Heesen, Remco. "Academic superstars: competent or lucky?." *Synthese* 194, no. 11 (2017): 4499-4518.

Helmer, Markus, Manuel Schottdorf, Andreas Neef, and Demian Battaglia. "Gender bias in scholarly peer review." Elife 6 (2017): e21718.

Hicks, Daniel. "Is Longino's conception of objectivity feminist?." *Hypatia* 26, no. 2 (2011): 333-351.

Hicks, Diana, Paul Wouters, Ludo Waltman, Sarah De Rijcke, and Ismael Rafols. "Bibliometrics: the Leiden Manifesto for research metrics." *Nature News* 520, no. 7548 (2015): 429.

Hill, Emily, Rachel Blake, Jenna Emerson, Peter Svider, Jean Eloy, Christina Raker, Katina Robison, and Ashley Stuckey. 2015. "Gender Differences in Scholarly Productivity within Academic Gynecologic Oncology Departments." *Obstetrics & Gynecology* 126 (6): 1279-1284.

Holroyd, J. D., and Katherine Puddifoot. "Implicit bias and prejudice." (2019): 313-326.

Hopkins, Allison L., James W. Jawitz, Christopher McCarty, Alex Goldman, and Nandita B. Basu. "Disparities in publication patterns by gender, race and ethnicity based on a survey of a random sample of authors." *Scientometrics* 96, no. 2 (2013): 515-534.

Hookway, Christopher. "Some Varieties of Epistemic Injustice: Reflections on Fricker." *Episteme-A Journal Of Individual And Social Epistemology* 7, no. 2 (2010): 151-63.

Hutson, Scott R. 2002. "Gendered Citation Practices in American Antiquity and Other Archaeology Journals." *American Antiquity* 67 (2): 331-342.

Intemann, Kristen. "25 years of feminist empiricism and standpoint theory: Where are we now?." *Hypatia* 25, no. 4 (2010): 778-796.

Ioannidis, John P. A, and Muin J Khoury. "Assessing Value in Biomedical Research: The PQRST of Appraisal and Reward." *JAMA* 312, no. 5 (2014): 483-84.

Ioannidis, John PA, Richard Klavans, and Kevin W. Boyack. "Thousands of scientists publish a paper every five days." (2018): 167-169.

Jackson, Sarah M., Amy L. Hillard, and Tamera R. Schneider. "Using implicit bias training to improve attitudes toward women in STEM." *Social Psychology of Education* 17, no. 3 (2014): 419-438.

Kahneman, Daniel, Stewart Paul Slovic, Paul Slovic, and Amos Tversky, eds. Judgment under uncertainty: Heuristics and biases. Cambridge university press, 1982.

Kattou, Maria, Katerina Kontoyianni, Demetra Pitta-Pantazi, and Constantinos Christou. "Connecting mathematical creativity to mathematical ability." *Zdm* 45, no. 2 (2013): 167-181.

Kelly, Clint D., and Michael D. Jennions. "The h index and career assessment by numbers." *Trends in Ecology & Evolution* 21, no. 4 (2006): 167-170.

Khan, Nickalus R, Clinton J Thompson, Douglas R Taylor, Garrett T Venable, R Matthew Wham, L Madison Michael, and Paul Klimo. "An Analysis of Publication Productivity for 1225 Academic Neurosurgeons and 99 Departments in the United States." *Journal of Neurosurgery* 120, no. 3 (2014): 746-55.

King, Molly M., Carl T. Bergstrom, Shelley J. Correll, Jennifer Jacquet, and Jevin D. West. "Men set their own cites high: Gender and self-citation across fields and over time." *Socius* 3 (2017): 2378023117738903.

Klimo, Paul, Garrett T. Venable, Nickalus R. Khan, Douglas R. Taylor, Brandon A. Shepherd, Clinton J. Thompson, and Nathan R. Selden. 2014. "Bibliometric Evaluation of Pediatric Neurosurgery in North America." *Journal of Neurosurgery: Pediatrics* 14 (6): 695-703. doi:10.3171/2014.8.PEDS1488.

Knobloch-Westerwick, Silvia and Carroll J. Glynn. 2013. "The Matilda Effect—Role Congruity Effects on Scholarly Communication." *Communication Research* 40 (1): 3-26.

Kuhn, Thomas S. "The Structure of Scientific Revolutions." (1962).

Larivière, Vincent, Chaoqun Ni, Yves Gingras, Blaise Cronin, and Cassidy R. Sugimoto. "Bibliometrics: Global gender disparities in science." *Nature News* 504, no. 7479 (2013): 211.

Lariviere, Vincent, and Cassidy R. Sugimoto. "The Journal Impact Factor: A Brief History, Critique, and Discussion of Adverse Effects." 2018.

Lee, Carole J. 2016. "Revisiting Current Causes of Women's Underrepresentation in Science." In *Implicit Bias and Philosophy, Volume 1*. Oxford: Oxford University Press.

Lee, Carole J., and Christian D. Schunn. "Social biases and solutions for procedural objectivity." *Hypatia* 26, no. 2 (2011): 352-373.

Leuschner, Anna. "Social exclusion in academia through biases in methodological quality evaluation: On the situation of women in science and philosophy." Studies in History and Philosophy of Science Part A 54 (2015): 56-63.

Ley, Timothy J., and Barton H. Hamilton. "The gender gap in NIH grant applications." *Science* 322, no. 5907 (2008): 1472-1474.

Longino, Helen E. *Science as Social Knowledge : Values and Objectivity in Scientific Inquiry*. Princeton, N.J.: Princeton University Press, 1990.

Longino, Helen E. 2002. *The Fate of Knowledge*. Princeton, N.J.: Princeton, N.J. : Princeton University Press.

Longino, Helen E. "Data, please.*" HOPOS: The Journal of the International Society for the History of Philosophy of Science* 3, no. 1 (2013): 144-146.

Maliniak, Daniel Ryan Powers, and Barbara F. Walter. 2013. "The Gender Citation Gap in International Relations." *International Organization* 67 (4): 889-922.

Martell, Richard F., Lane, David M., and Emrich, Cynthia. "Male-female Differences: A Computer Simulation." *The American Psychologist* 51, no. 2 (1996): 157-158.

Merton, Robert. Foreword to *Citation Indexing -- Its Theory and Application in Science, Technology, and Humanities*, by Garfield, Eugene, and Robert K. Merton. Vol. 8. New York: Wiley, 1979.

Moher, David, Florian Naudet, Ioana A. Cristea, Frank Miedema, John PA Ioannidis, and Steven N. Goodman. "Assessing scientists for hiring, promotion, and tenure." *PLoS biology* 16, no. 3 (2018): e2004089.

Moss-Racusin, Corinne A., John F. Dovidio, Victoria L. Brescoll, Mark J. Graham, and Jo Handelsman. "Science faculty's subtle gender biases favor male students." *Proceedings of the national academy of sciences* 109, no. 41 (2012): 16474-16479.

Nosek, Brian A., Frederick L. Smyth, N. Sriram, Nicole M. Lindner, Thierry Devos, Alfonso Ayala, Yoav Bar-Anan, et al. 2009. "National Differences in Gender-Science Stereotypes Predict National Sex Differences in Science and Math Achievement." *Proceedings of the National Academy of Sciences of the United States of America* 106 (26): 10593-10597.

Prentice, Deborah A. "Liberal Norms and Their Discontents.*" Perspectives on Psychological Science : A Journal of the Association for Psychological Science* 7, no. 5 (2012): 516-8.

Pohlhaus, Jennifer Reineke, Hong Jiang, Robin M. Wagner, Walter T. Schaffer, and Vivian W. Pinn. "Sex differences in application, success, and funding rates for NIH extramural programs." *Academic Medicine* 86, no. 6 (2011): 759.

Porter, Anne Marie, and Rachel Ivie. "Women in Physics and Astronomy, 2019. Report." *AIP Statistical Research Center* (2019).

Rubin, Hannah, and Cailin O'Connor. "Discrimination and collaboration in science." *Philosophy of Science* 85, no. 3 (2018): 380-402.

"San Francisco Declaration on Research Assessment." DORA. Accessed February 17, 2018. https://sfdora.org/.

Saul, Jennifer. "Implicit bias, stereotype threat, and epistemic injustice." In *The Routledge Handbook of Epistemic Injustice*, pp. 235-242. Routledge, 2017.

Saul, Jennifer and Michael Brownstein, eds. 2016. *Implicit Bias and Philosophy, Volume 1*. Vol. 1. Oxford: Oxford University Press.

Schisterman, Enrique, Chandra Swanson, Ya-Ling Lu, and Sunni Mumford. 2017. "The Changing Face of Epidemiology: Gender Disparities in Citations?" *Epidemiology* 28 (2): 159-168.

Seglen, P. O. 1997. "Why the Impact Factor of Journals should Not be used for Evaluating Research." *Bmj* 314 (7079): 497. doi:10.1136/bmj.314.7079.497.

Solomon, Miriam. "Social Empiricism." (2001).

Solomon, Miriam, and Alan Richardson. "A critical context for Longino's critical contextual empiricism| The fate of knowledge, Helen Longino, Princeton University Press, Princeton (2002), pp. 288, Price£ 12.95/$19.95 paperback, ISBN: 0-691-08876-4." Studies in History and Philosophy of Science Part A 36, no. 1 (2005): 211-222.

Squazzoni, Flaminio, Giangiacomo Bravo, Francisco Grimaldo, Daniel Garcıa-Costa, Mike Farjam, and Bahar Mehmani. "No tickets for women in the COVID-19 race? A study on manuscript submissions and reviews in 2347 Elsevier journals during the pandemic." *SSRN Electronic Journal* 10 (2020).

Steel, Daniel, Sina Fazelpour, Bianca Crewe, and Kinley Gillette. "Information elaboration and epistemic effects of diversity." *Synthese* 198, no. 2 (2021): 1287-1307.

Steel, Daniel, and Kyle Powys Whyte. "Environmental justice, values, and scientific expertise." *Kennedy Institute of Ethics Journal* 22, no. 2 (2012): 163-182.

"Survey of Earned Doctorates.", https://ncses.nsf.gov/pubs/nsf20301/data-tables.

Talbott, William J. "Intentional self-deception in a single coherent self." *Philosophy and Phenomenological Research* 55, no. 1 (1995): 27-74.

Trump, Jonathan R., Asantha Cooray, Shea Garrison-Kimmel, Boris Haeussler, and Anna Pancoast. Astro2020 State of the Profession White Paper Realistic Job Training for Astro PhDs. (2020).

"UPDirectory.", https://www.apaonline.org/mpage/theupdirectory.

Valian, Virginia. 1998. *Why so Slow? The Advancement of Women*. Cambridge, Mass.: Cambridge, Mass.: MIT Press.

Van der Lee, Romy, and Naomi Ellemers. "Gender contributes to personal research funding success in The Netherlands." *Proceedings of the National Academy of Sciences* 112, no. 40 (2015): 12349-12353.

Waddell, Paul. "UrbanSim: Modeling urban development for land use, transportation, and environmental planning." *Journal of the American planning association* 68, no. 3 (2002): 297-314.

Waddell, Paul, Alan Borning, Michael Noth, Nathan Freier, Michael Becke, and Gudmundur Ulfarsson. "Microsimulation of urban development and location choices: Design and implementation of UrbanSim." *Networks and spatial economics* 3, no. 1 (2003): 43-67.

Way, Samuel F., Daniel B. Larremore, and Aaron Clauset. "Gender, productivity, and prestige in computer science faculty hiring networks." In Proceedings of the 25th International Conference on World Wide Web, pp. 1169-1179. 2016.

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.

Wilsdon, James, Liz Allen, Eleonora Belfiore, Philip Campbell, Stephen Curry, Steven Hill, Richard Jones et al. "The metric tide." *Report of the independent review of the role of metrics in research assessment and management* (2015).

Witteman, Holly O., Michael Hendricks, Sharon Straus, and Cara Tannenbaum. "Are gender gaps due to evaluations of the applicant or the science? A natural experiment at a national funding agency." *The Lancet* 393, no. 10171 (2019): 531-540.

Wylie, Alison, Jakobsen, Janet R., and Fosado, Gisela. “Women, work and the academy: Strategies for responding to ‘post-civil rights era’ gender discrimination”. *New feminist solutions* (series). New York: Barnard Center for Research on Women. (2007).

Wylie, Alison, "Why standpoint matters." in *Science and other cultures: Issues in philosophies of science and technology*, edited by Robert Figueroa, and Sandra Harding. 26 (2003): 48.

Zollman, Kevin JS. "The epistemic benefit of transient diversity." *Erkenntnis* 72, no. 1 (2010): 17.

APPENDIX A

```
breed [people person]
breed [pubs pub]
directed-link-breed [citations citation]
undirected-link-breed [authored author]
people-own [pubcount totcitcount grantcount intpubs newgrant?]
pubs-own [citcount age]
globals [avgcitswomen avgcitsmen avgpubswomen avgpubsmen avggrantswomen avggrantsmen
sdmc sdwc sdmp sdwp sdmg sdwg avgcits intpublist]



to setup
  clear-all

  if grantrate <= 0 or grantrate >= 1
  [print "Please enter a decimal for the grant rate."
    stop]
   if avgcitsprop <= 0 or avgcitsprop >= 1
  [print "Please enter a decimal for the average proportion of papers cited."
    stop]
  set-default-shape people "person"
  set-default-shape pubs "book"
  set-default-shape citations "cits"
  createpeople
  createpubs
  setupauthors
  setupcits
  setcounts
  layout

  reset-ticks

end

to createpeople
  create-people comsize [   ;;makes the number of people in the community based on the input
comsize
   ifelse who >= (comsize / 2)  ;;makes half blue andf half pink; if the size is odd, there will be one
add'l blue
  [set color pink]
    [set color blue]
  ]
  layout-circle people 15   ;;lays them out in a circle
end
```

```
to createpubs
  let totintpubs 0  ;;keeps track of how many pubs to make in total
  set intpublist []
  ask people [
    let x (ln avgpubs - .18)
    let y random-normal x .6
    let w exp y
    let z ((ceiling w) - 1)
    set intpubs z  ;;sets that as the initial number of pubs for a person, so that it will be possible to
connect the person with the right number of pubs
    set totintpubs (totintpubs + z) ;;adds it to the counter of total pubs being made
    set intpublist lput z intpublist
  ]
  create-pubs totintpubs [set color green  ;;makes the right number of publications
  set age 0]  ;;keeps track of the age of pubs; it's not that important to do, but it's helpful for checking
things
 end

to setupauthors
ask people [
    while [(count author-neighbors) < intpubs]  ;; the loop won't end until each person has the right
number of pubs;
    [let pubwho one-of pubs      ;;finds a pub
      let i 0
      ask pubwho [ set i (count author-neighbors)]  ;;checks if the pub already has an author
      if  i < 1
      [create-author-with pubwho] ;;if not, makes the person the author. if yes, the loop repeats until
the person has enough pubs

  ]


  ]

end


to setupcits
  let x count pubs
  set avgcits ceiling( x * avgcitsprop)
   ask pubs
 [ let m ln avgcits
    let n random-normal m 1
    let p exp n
    let i ceiling p
```

```
    while [ i > ((count pubs) - 1) ]    ;;checks that the number of outgoing citations is not greater than the total number of citable articles
    [ set m ln avgcits
    set n random-normal m 1
    set p exp n
    set i ceiling p]
    while [(count my-out-citations) < i]  ;;repeats until this turtle has the right number of outcitations
    [let newwho [who] of one-of pubs  ;;finds a random pub
     if who != newwho
        [create-citation-to pub newwho] ;;as long as it's not itself, it makes a citation to that paper. it also won't create a link if there already is one, but the loop won't end until the right number of out citations is created
    ]

  ]

end

to layout         ;;makes it look nice
ask people
  [let x xcor
    let y ycor
    ask author-neighbors
    [setxy (x +  random-normal 0 3) (y +  random-normal 0 3)]
  ]
end


to go

  ifelse grantfortopx?
  [getgranttopx]
  [getgrant]
  publish
  ifelse genderbias?
  [setnewcitesbiased]
  [setnewcites]
  setcounts
  layoutagain
  tick

end

to go-once
   ifelse grantfortopx?
  [getgranttopx]
```

```
  [getgrant]
  publish
 ifelse genderbias?
  [setnewcitesbiased]
    [setnewcites]
  setcounts
  layoutagain
  tick
end

to getgrant
  ask pubs [set color green
  set age (age + 1)]      ;;turns previously new pubs back to green
  ask people [
    set newgrant? 0              ;;overwrites previous value for having gotten a grant
      let i random-float 1
    if i < grantrate
    [ set grantcount (grantcount + 1)     ;;if the random number is greater than the grant rate, the
person gets a grant
    set newgrant? 1
    ]

  ]
end


to getgranttopx  ;;this version gives grants to top performers
  ask pubs [set color green
  set age (age + 1)]
  ask people [set newgrant? 0]

 if grantrate = grantfortopx  ;;calls the right function, depending on the two inputs for grant rate
and what top percent to consider
   [grandgftxequal]
  if grantrate > grantfortopx
   [grgreaterthangftx]
   if grantrate < grantfortopx
   [grlessthangftx]


ask people [if newgrant? = 1   ;;once the new grants are distributed, this just tracks how many total
grants a person has
    [ set grantcount (grantcount + 1)]]

end
```

```
to grandgftxequal   ;;used when the two rates are equal; the effect is that only the top x% will get grants and no one else
  let a topzpeople (grantrate * 100)  ;;generates a list of top x%
  let i 0
      while [i < length a]   ;;goes through the list giving each member a grant
     [let b item i a
       let c item 1 b
      ask person c [set newgrant? 1]
    set i (i + 1)]
end

to grgreaterthangftx   ;;used when the grant rate is higher than the top x% value; the effect is that all of the top x% will receive a grant, plus grantrate - topx% additional people
  let a topzpeople (grantfortopx * 100)  ;;makes the list of top performers
  let i 0
      while [i < length a]  ;;goes through the list giving each member a grant
     [let b item i a
       let c item 1 b
      ask person c [set newgrant? 1
        ]
    set i (i + 1)]
      let grantsleft (grantrate - grantfortopx)  ;;sets the value for how many 'leftovers' there are
     ask people[ if newgrant? = 0      ;;makes sure only people who don't already have a grant can get one that's 'leftover'
         [let j random-float 1    ;;uses a random check like the version where everyone has an equal chance
          if j < grantsleft
            [set newgrant? 1]
          ]
    ]

end

to grlessthangftx  ;;used when the grant rate is lower than top x% percent input; the effect is that individuals from the top x% are choosen for the right number of total grants
  let m topzpeople (grantfortopx * 100)  ;;makes the list of top performers
    let n (grantrate * comsize)  ;;determines the number of grants to be given
   let p 0
      while [p < n]
      [let q length m
    let r random q  ;;picks a random element from the top performer list
    let s item r m
    let t item 1 s
        ask person t [set newgrant? 1]  ;;gives that person a grant
          set m remove s m  ;;removes them from the list
          set p (p + 1)
```

```
   ]
end

to-report countnewgrants    ;;checks if a person has a new grant to generate the correct number of new pubs
  let i 0
  ask people [set i (i + newgrant?)]
  report i

end


to publish
  let newpublist []
  create-pubs countnewgrants [set color red    ;;creates a new publication for each new grant and makes them red so they can be tracked more easily
    set age 0
     set newpublist lput who newpublist]
  ask people [
    if newgrant? = 1
    [let i first newpublist
    create-author-with pub i       ;;gives each person who received a new grant one of the newly created pubs
    set newpublist remove i newpublist]
    ]

end

to setnewcites
    let x count pubs
  set avgcits ceiling( x * avgcitsprop)
  let red-pubs pubs with [color = red]
  ask red-pubs [

       ;;only adds new citations to the new pubs
   let r ln avgcits
    let s random-normal r 1
    let t exp s
    let i ceiling t
   while[ i > ((count pubs) - 1) or i = 0]
      [set r ln avgcits
    set s random-normal r 1
    set t exp s
    set i ceiling t]
    while [(count my-out-citations) < i]  ;;repeats until this turtle has the right number of outcitations
    [let j one-of pubs    ;;picks a random pub
```

```
      let k 0
      ask j [set k one-of citations    ;;picks one if its citations
      ]
      let l [one-of both-ends] of k       ;;then picks one end. This does drive citations towards publications with more citations, but not as aggressively as picking a citation directly
    if l != self
      [create-citation-to l]

    ]


  ]

end

to setnewcitesbiased   ;;the biased version pays attention to what gender the author is
   let a count pubs
 set avgcits ceiling( a * avgcitsprop)
 let red-pubs pubs with [color = red]
 ask red-pubs [

  let r ln avgcits
   let s random-normal r 1
   let t exp s
   let i ceiling t
  while[ i > ((count pubs) - 1) or i = 0]
    [set r ln avgcits
   set s random-normal r 1
   set t exp s
   set i ceiling t]
  while [(count my-out-citations) < i]
  [let j one-of pubs ;;picks a random pub
      let k 0
      ask j [set k one-of citations]   ;;finds a citation of that pub
      let m [end1] of k
      let n [end2] of k
      let p author-gender m    ;;finds out the color of each end of the citation links
      let q author-gender n
      if p = q       ;;if they're the same, then the procedure is the same as the unbiased version
      [let x [one-of both-ends] of k
        if x != self
        [create-citation-to x]]
      if p < q      ;;if the values are different, there is a random check. if a random float between 0 and 1 returns above the threshold set with 'bias points', then the code will always pick the blue author.
```

```
      [ifelse random-float 1 < (1 - (biaspoints / 100))   ;;otherwise, the procedure is the same as the unbiased version
        [let x [one-of both-ends] of k
        if x != self
        [create-citation-to x]]
        [if m != self
          [create-citation-to m]]
      ]
      if p > q
      [ifelse random-float 1 < (1 - (biaspoints / 100))
        [let x [one-of both-ends] of k
        if x != self
        [create-citation-to x]]
        [if n != self
          [create-citation-to n]]
      ]
   ]

  ]
    end

to-report author-gender [x]  ;;gets the color of the author of a pub
  let i 0
  ask x [
    ask author-neighbors [set i color]
  ]
  report i
end


to layoutagain      ;;makes it look nice (again!) although it does occasionally give a point off the plot, so at some point, I should add offsets
  ask pubs [
    let x 0
    let y 0
    if color = red
    [ask author-neighbors
    [set x xcor
      set y ycor]
      setxy (x +  random-normal 0 3) (y +  random-normal 0 3)]
  ]
end

to-report topx%? [x]     ;;compiles info about the papers with the most citations
  let xlist []
 ask pubs[
```

```
    set xlist lput list citcount who xlist    ;;makes a list of lists for each pub, with its citcount and its who
  ]
set xlist sort-by [[i j] -> ( (item 0 i) > (item 0 j))] xlist  ;;thanks to substack for this code on sorting by the first item of the list
  let y length xlist
  let z (round((x / 100) * y))
  let topxlist []
  let i 0
  while [i < z]
  [let j item i xlist
    set topxlist lput j topxlist    ;;once sorted, a new list is made with just the top x % of turtles
  set i (i + 1)]
  report topx-gender topxlist x
  report topxlist

end

to-report topx-gender [x y]    ;;finds out the gender of the authors of the papers with the most citations
  let a length x
  let i 0
  let mencount 0
  let womencount 0
  while [i < a]
  [let j item i x
    let k item 1 j
    let l author-gender pub k
    if l = 105
    [set mencount (mencount + 1)]
    if l = 135
    [set womencount (womencount + 1)]
  set i (i + 1)]
  report (word "There are " womencount " papers written by women and " mencount " papers written by men in the top " y "% of citations")
end

to-report topzpeople [z]  ;; does the same as previous reporters, but for the total citation counts of individuals
  let peoplelist []
  ask people [
    set peoplelist lput list totcitcount who peoplelist
  ]
  set peoplelist sort-by [[i j] -> ( (item 0 i) > (item 0 j))] peoplelist
  let a length peoplelist
  let b (round((z / 100) * a))
```

```
  let c 0
  let topzlist []
  while [c < b]
  [let d item c peoplelist
    set topzlist lput d topzlist
    set c (c + 1)]
 ;report topz-gender topzlist z

  report topzlist
end


to-report topz-gender [x y]  ;;again, finds the gender of the most cited people
  let a length x
  let i 0
  let mencount 0
  let womencount 0
  while [i < a]
  [let j item i x
    let k item 1 j
    let l [color] of person k
    if l = 105
    [set mencount (mencount + 1)]
    if l = 135
    [set womencount (womencount + 1)]
  set i (i + 1)]
  report (word "There are " womencount " women and " mencount " men in the top " y "% of
citations")
end


to setcounts
  let mcitlist []
  let mgrantslist []
  let mpubslist []
  let wcitlist []
  let wpubslist []
  let wgrantslist []

  ask people [
    set pubcount (count author-neighbors)
    let j 0
    ask author-neighbors [
      let i (count in-citation-neighbors)
      set j (j + i)
    ]
```

```
    set totcitcount j
    ifelse color = blue
    [set mcitlist lput totcitcount mcitlist
     set mgrantslist lput grantcount mgrantslist
      set mpubslist lput pubcount mpubslist
    ]
    [set wcitlist lput totcitcount wcitlist
      set wgrantslist lput grantcount wgrantslist
   set wpubslist lput pubcount wpubslist
   ]
  ]
  set avgcitsmen mean mcitlist
  set avgcitswomen mean wcitlist
  set avgpubsmen mean mpubslist
  set avgpubswomen mean wpubslist
  set avggrantsmen mean mgrantslist
  set avggrantswomen mean wgrantslist
  set sdmc standard-deviation mcitlist
  set sdwc standard-deviation wcitlist
  set sdmp standard-deviation mpubslist
  set sdwp standard-deviation wpubslist
  set sdmg standard-deviation mgrantslist
  set sdwg standard-deviation wgrantslist

  ask pubs [set citcount (count in-citation-neighbors)]
end
```

VITA

[A short bio of the author is required for a Ph.D. dissertation at the University of Washington. The vita section does not go into the Table of Contents. The formatting style follows the text of the dissertation.]