UNIVERSITY OF CALIFORNIA,
IRVINE

Those Who Know Better

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Philosophy

by

Jingyi Wu

Dissertation Committee:
Professor Cailin O'Connor, Co-Chair
Professor James Owen Weatherall, Co-Chair
Associate Professor Liam Kofi Bright
Associate Professor Davin L. Phoenix
Distinguished Professor Brian Skyrms

2023

DEDICATION

To

My dear friend S and my grandmother Wenxian Guo.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

Thank you to my co-chair Jim Weatherall, without whom I would not be in graduate school, let alone progressing through it. Jim's marvelously clear and rigorous lecture on spacetime structures at the Munich Center for Mathematical Philosophy in Summer 2016 first inspired me to pursue philosophy as a career. In the past seven years, Jim supported me unfailingly, providing detailed and perspicuous comments on almost everything I have written, opening doors to professional opportunities, and lending an ear in times of personal and professional crises. I learned from Jim to stay inquisitive about anything that interests me, to navigate delicately through thorny administrative matters, and to always be prepared to engage, learn, and grow.

Thank you to my co-chair Cailin O'Connor, without whom this dissertation would not have existed. It was Cailin's seminar on Network Epistemology in Spring 2020 when the ideas in the first chapter were first conceived. Cailin has provided thorough, clear, and insightful comments on everything in this dissertation. Cailin has the talent for communicating the most complex philosophical ideas in the most accessible way, and her love for the art of living is extremely inspiring. I learned from Cailin that philosophy is at its best when it is intimately related to its social ecology—when it is effectively communicated to the public, when it draws inspiration from lived experiences, and when it helps us make sense of messy social realities and our places in them.

Thank you to my committee member Liam Kofi Bright. Liam's work on the formal social epistemology of race and gender provides existence proof that what is in this dissertation has a place in philosophy. Since Liam and I met (virtually) in Fall 2020, he has been extraordinarily generous with his time. Liam has provided prompt, observant, and constructive feedback on almost everything I have written since, and is always available when I seek his advice. Liam provides mentorship of this quality to so many (minoritized) students in philosophy. I do not know how he does it all. I am fortunate to have Liam as a future colleague.

Thank you to my other committee members Davin Phoenix and Brian Skyrms. Davin's writing workshop on race and justice in Spring 2021 provided a much-needed sanctuary for me. Davin, through his enthusiasm and encouragement, assured me that my work has relevance outside of academic philosophy. I learned from Davin to write *for* the community that nourishes and inspires me. Brian's depth of knowledge in anything about social dynamics is awe-inspiring. I learned from Brian to be humble yet confident. The HPS Fund that Brian administers has provided financial support for many of my research travels that make this dissertation possible.

This dissertation is also made possible through generous support from the Social Science Merit Fellowship, the Provost PhD Fellowship, the Summer Inclusive Excellence Fellowship, the Associate Dean's Fellowship, the Graduate Dean's Recruitment Fellowship, and National Science Foundation Grant No. 1922424. Materials from this dissertation were presented at the PSA 2021 in Baltimore, the PSA 2022 in Pittsburgh, the Seminar on Formal and Social

Epistemology (of Science) in Paris, the Workshop on Agent-based Modeling of Epistemic Communities in Bochum, the Pittsburgh Formal Epistemology Workshop, and to department members at York, Caltech, Penn, BU, and LSE. I am grateful for the comments and questions from the audience members.

My progression through this Ph.D. would not be possible without the research, writing, and support groups I am part of over the years. At UC Irvine, I am especially grateful to members of the Philosophy of Physics Research/Reading Group (PPRG), the Network for Modelers research group, the Philosylvania writing group run by Helen Meskhidze, and the daily check-in group with Jessica Gonzalez, Helen Meskhidze, Chris Mitsch, and Will Stafford during the stay-at-home period of the pandemic. Outside of UC Irvine, I benefit from the (almost daily) dissertation writing time with Sena Bozdag in Winter and Spring 2023, the writing accountability group with Bixin Guo, Chris Mitsch, and Siddharth Muthu Krishnan in 2022-23, and the Rosalind peer support Discord server started by Chrisy Xiyu Du. I am also fortunate to have excellent outside mentors who have supported me in various forms along the way, especially Kareem Khalifa, Carole Lee, Hannah Rubin, Laura Ruetsche, Alison Wylie, and Kevin Zollman. Thanks to Dr. Travis LaCroix for sharing this LaTeXtemplate.

Though it may not be explicit in the main text, one core methodology I adopt in this dissertation is auto-ethnography. Many ideas became crystalized after regular and extensive discussions (a.k.a. gossip sessions) with a community of co-conspirators about our lived experiences in the messy social structures around us. These lovely people include: Nidhi Banavar, Liam Kofi Bright, Nathan Gabriel, Itzel Garcia, Bixin Guo, Em Hernandez, Kareem Khalifa, Hannah Kim, Milana Kostic, Shiqi Lin, Chris Mitsch, Stella Moon, Jiangxue Ning, Allie Richards, Ellen Shi, Angela Sun, Jocelyn Wang, Jim Weatherall, Kino Zhao, and members of the Asian American Philosophy working group at UC Irvine in Winter and Spring 2022.

Several friends on this list deserve extra thanks. Thanks to Nidhi Banavar, for our conversations about grief, family, guilt, social injustice, and how to stay socially relevant as a scientist. Thanks to Em Hernandez, for our conversations about how academia (dys)functions and about (un)healthy human relationships, and for our peaceful time co-working and playing games. Thanks to Milana Kostic, for discussing our shared struggles over the years about balancing the urge to do hands-on activism at home and the lure of intellectual life in exile. Thanks to Shiqi Lin, for processing our shared vulnerabilities in this geopolitical reality, and for being a great sounding board throughout, especially during this year on the market. Thanks to Ellen Shi, for homemade food and for our conversations about family, childhood, educational trajectories, and how there are *so few of us* in academia. Thanks to Kino Zhao, for all the solidarity in times of need, and for being a truly dependable co-organizer.

Many friends who have lived in the area over the years have provided encouragement, support, soulful conversations, and camaraderie. Some were already mentioned above; others include: Clara Bradley, Yunlong Cao, Cindy Chen, Adam Chin, Margaret Farrell, Ben Genta, Rebecca Korf, Charles Leitz, JB Manchak, Helen Meskhidze, Anubhav Nanavaty, Sarita Rosenstock, Will Stafford, Gary Wu, Xinyue Yuan, and Lichen Zhen.

Many thanks to the friends I made during my year in Munich, from 2017-18, especially Sena Bozdag, Alfredo García Cid, Nakul Heroor, Conny Knieling, Milana Kostic, Alex Markham, Alex Niderklapfer, Maksim Nikiforovski, Allie Richards, and Nata Yang. This dissertation is about the benefits of diversity in our epistemic lives. But having a group of truly diverse friends in Munich allows me to experience first-hand how *fun* a cosmopolitan life can be. Thank you for making it such a remarkable year.

Many thanks to some of my "old" friends, especially Liting Fan, Danna Fang, Chujia Ji, Fan Wu, and Fan Yang. We don't talk often, but when we do, I am always so moved by how well we have aged during our times apart, and how we invariably come to similar realizations about many aspects of life.

Even though we are no longer together, many thanks to David Yang, for traveling around the world with me, for providing the stability during the three years he moved to California for me, for reading through many of my draft papers, for discussing early ideas of this dissertation with me, and above all, for supporting my dream in the philosophy academia. I am fortunate to have shared part of my life with him. Thanks to Hua Su and Fan Yang, for their hospitality during many of our visits to Canada.

Many thanks to my parents, Yuping Chen and Jiangfei Wu. Things between us have not always been smooth, but I understand that they have raised me in the best way they thought possible, even though that means my current life is, in many ways, incomprehensible to them. Thanks for the risk they took.

This dissertation is dedicated to S. I often mourn the unrealized possibility in which we would go through graduate school together, and I often wonder what she would think of my philosophical ideas. I hope one day she will read this dissertation, and I will tell her all about the intervening years. I am sure the dissertation she would have written/will write is much better than this, and this dissertation would have been better had she played a role in shaping it. I miss her.

This dissertation is also dedicated to my grandmother, Wenxian Guo. She is so smart and never lost the thirst to learn. I am sure she would have been the first person in our family to get a Ph.D. had she been born into better circumstances. She always tells me how unwilling she was for her formal education to end abruptly after elementary school due to poverty. She supported her granddaughters to move and study abroad, to pursue educational opportunities she never had herself, even though that meant she would see less of us. Now I will be the first person in our family to get a Ph.D. This achievement is hers as much as it is mine.

Chapter 1 is published Open Access under license CC BY 4.0 in *Philosophy and Phenomenological Research*: 'Epistemic Advantage on the Margin: A Network Standpoint Epistemology' by Jingyi Wu, 2022 (Early View). The published version is accessible at https://doi.org/10.1111/phpr.12895. No changes were made.

VITA

Jingyi Wu

Ph.D in Philosophy **2023**
University of California, Irvine *Irvine, California*

M.A. in Social Science **2021**
University of California, Irvine *Irvine, California*

B.A. in Philosophy, Mathematics **2017**
Middlebury College *Middlebury, Vermont*

PUBLICATIONS

The Cultural Evolution of Science (with Cailin O'Connor and Paul E. Smaldino) **Forthcoming**
The Oxford Handbook of Cultural Evolution

The Next Generation Event Horizon Telescope Collaboration: History, Philosophy, and Culture (with 24 interdisciplinary co-authors) **2023**
Galaxies

How Should We Promote Transient Diversity in Science? (with Cailin O'Connor) **2023**
Synthese

Epistemic Advantage on the Margin: A Network Standpoint Epistemology **2022**
Philosophy and Phenomenological Research

Explaining Universality: Infinite Limit Systems in the Renormalization Group Method **2021**
Synthese

A Bohr Mollerup Theorem for Interpolating the Triangular Numbers (with Stephen Abbott) **2018**
Journal of Convex Analysis

ABSTRACT OF THE DISSERTATION

Those Who Know Better

By

Jingyi Wu

Doctor of Philosophy in Philosophy

University of California, Irvine, 2023

Professor Cailin O'Connor, Co-Chair
Professor James Owen Weatherall, Co-Chair

How do social identities and social injustice impact knowledge production in a group? Does diversity improve group learning? In what ways? How do we promote such diversity? These are the kinds of questions I tackle in my dissertation, which is situated at the intersection of social epistemology, network modeling, and the philosophy of race and gender.

A theme running through my findings is that having a diversity of approaches within a group can facilitate the production of better knowledge. I propose three novel mechanisms that lead to epistemically-beneficial diversity and use agent-based models to investigate their often surprising downstream consequences. I show that (1) marginalized social groups sometimes develop better beliefs because their testimony is devalued by dominant groups; (2) industrial scientists can gain epistemic benefits by failing to share their research; and (3) a group can ultimately learn better when its members explore many, possibly sub-optimal, solutions to a problem instead of always choosing the best available solution.

Introduction

How do social identities and social injustice impact knowledge production in a group? Does diversity improve group learning? In what ways? How do we promote such diversity? These are the kinds of questions I tackle in my dissertation, which is situated at the intersection of social epistemology, network modeling, and the philosophy of race and gender.

A theme running through my findings is that having a diversity of approaches within a group can facilitate the production of better knowledge. I propose three novel mechanisms that lead to epistemically-beneficial diversity and use agent-based models to investigate their often surprising downstream consequences. I show that (1) marginalized social groups sometimes develop better beliefs because their testimony is devalued by dominant groups; (2) industrial scientists can gain epistemic benefits by failing to share their research; and (3) a group can ultimately learn better when its members explore many, possibly sub-optimal, solutions to a problem instead of always choosing the best available solution.

In the first chapter, "Epistemic Advantage on the Margin," I consider situations where a dominant group ignores or devalues evidence from a marginalized group. I find that the marginalized group develops better beliefs more often and faster than the dominant group, and in some cases even outperforms a community without epistemic injustice. When the dominant group ignores data, they take longer to explore the options. The marginalized group, however, learns from their own experience and the exploration of the dominant group, and ultimately benefits from this epistemic diversity. These results support a contested standpoint epistemology thesis that marginalized groups can know better, by connecting it to a more widely discussed phenomenon that marginalized groups' testimony is often devalued.

In the second chapter, "Withholding Knowledge," I argue that there are epistemic incentives for subgroups of scientists to unilaterally withhold evidence from the larger community. In so doing, they develop better beliefs more frequently and faster than the rest of the community. The withholding scientists gain epistemic advantage for similar reasons as the marginalized agents in chapter one. I build two new models from different modeling paradigms to provide robustness tests for these findings. I further analyze the scientific sharing dynamics from a game theoretic perspective and show that groups face an epistemic prisoner's dilemma: they learn worst when others withhold evidence, best when they unilaterally withhold, and in between when everyone shares.

In the third chapter, "Better than Best," I show that a group can learn better when its members do not always employ the best available solution to a problem. If group members randomly select a better solution than their own, they preserve a diversity of approaches that ultimately makes them more successful. In a slogan, "better" beats "best."

My findings generate a range of potential behavioral and policy interventions for empirical testing. For instance, dominant groups may epistemically benefit from listening to previously marginalized voices; the public may epistemically benefit from policies requiring industry to share research; and institutions may epistemically benefit from policies that preserve non-optimal but promising diverse perspectives.

On top of the thematic unification, the three chapters are linked methodologically. This dissertation utilizes two fundamentally different modeling paradigms: the bandit problem and the NK landscape problem, both combined with network modeling techniques. The bandit problem is featured in Chapters 1 and 2, and the NK landscape problem is featured in Chapters 2 and 3. These two models represent different types of scientific and social inquiry. In the bandit model agents figure out which of the two probabilistic epistemic options is better, representing, e.g., clinical doctors finding out the efficacy of two different drugs by conducting trials. In the NK landscape model, agents search in a vast epistemic landscape with multiple "peaks," representing, e.g., researchers adopting different approaches to solve a problem. Together these models represent a wide array of possible problems that different epistemic communities face.

Despite the differences, at the core of both modeling paradigms, there is a trade-off between exploration and exploitation: Do agents exploit the solutions that they currently have, or do they explore other possibilities in the hope of finding better solutions? This trade-off explains why having diverse approaches in a community is beneficial to social learning. If a community, for whatever reasons, consists of members who are testing a diverse range of options, then agents can have access to data from options they have not explored, while continuing testing their own options. This dissertation, then, proposes some of the reasons that can lead to epistemically-beneficial diversity of practice, and test the robustness of these mechanisms in two different modeling paradigms. This kind of trade-off between exploration and exploitation is likely present in many of our endeavors, since we humans are cognitively and computationally limited. Insofar as that is right, the results I present will likely hold more broadly, even though I only test the mechanisms in two modeling paradigms here.[1]

Although not officially included, several projects I contribute to over the years serve as natural companion pieces to this dissertation. In Wu and O'Connor (2023), we evaluate different mechanisms for generating epistemically-beneficial diversity in scientific communities, keeping in mind practical and ethical considerations. In Wu et al. (2023), we offer a brief overview of formal models that treat science as a cultural evolutionary system, and consider

[1] I am, after all, cognitively and computationally limited too!

how distinct cultural patterns emerge in scientific communities by appealing to mechanisms of cultural transmission and social learning. Both articles include extensive discussions of models and mechanisms proposed in this dissertation. In an ongoing project with Liam Kofi Bright and another with Sina Fazelpour, Juan Paris, Hannah Rubin, and Courtney Sharpe, we explore other kinds of injustice in our society that can be represented by asymmetries in network relations, and explore their consequences. These projects serve as natural extensions to Chapter 1. Readers may also be interested in a recent blog post I wrote introducing the results from the first two chapters of this dissertation (Wu, 2023b).

Chapter 1

Epistemic Advantage on the Margin

> *What people of color quickly come to see—in a sense, the primary epistemic principle of the racialized social epistemology of which they are the object—is that they are not seen at all.*
>
> *Mills (2007, 17)*

1.1 Introduction

Social epistemologists working on race or gender have written extensively on dominant social groups' widespread practice of ignoring or devaluing testimony arising from marginalized groups. For example, Dotson (2011) uses *epistemic quieting* to describe situations in which an audience, often from a dominant social background, fails to identify a speaker, often from a marginalized background, as a knower. Other forms of this practice include *epistemic smothering* (Dotson, 2011), *testimonial injustice* (Fricker, 2007), and a form of *white ignorance* (Mills, 2007). Central to all these cases is a failure of testimonial reciprocity between a speaker and an audience.[1] Moreover, this failure of reciprocity is often unidirectional

[1]For more on testimonial reciprocity, see Hornsby (1995).

because members of marginalized groups cannot afford to engage in the devaluation and ignoration of testimony from the dominant group, due to sociopolitical power imbalance (c.f. Mills, 2007, Section 2). I call the situation in which the dominant group ignores testimony from the marginalized group *one-sided testimonial ignoration*, and the situation in which the dominant group devalues testimony from the marginalized group *one-sided testimonial devaluation*. Together, they constitute a unidirectional failure of testimonial reciprocity. That such situations occur is widely claimed in social epistemology, and some authors (e.g. Wylie (2003); Mills (2007); Saint-Croix (2020)) regard the unidirectional failure of testimonial reciprocity as a key claim of standpoint epistemology, which is a strand of social epistemology that takes as epistemically salient the social positions (or standpoints) that knowers are situated in.[2]

Besides the unidirectional failure of testimonial reciprocity, another key claim of standpoint epistemology, most notably advocated by Hartsock (1983), contends that "certain [socially marginalized] locations *themselves* foster more accurate beliefs, not only concerning one's own social position, but also the social and natural world more broadly" (Saint-Croix, 2020, 493, emphasis in the original). This claim is often called the *inversion thesis*, after the inverse relation between knowers' sociopolitical power and epistemic privilege. The interpretations of and justifications for the inversion thesis are often highly contested (see, e.g., Wylie (2003); Intemann (2010); Toole (2020)). In what follows, I propose a possible mechanism that gives rise to the inversion thesis, by connecting it to the other key claim mentioned above.

[2]There are two other distinct forms of epistemic marginalization discussed in the standpoint epistemology literature that are worth mentioning here. First, one might think that sometimes marginalized groups are not even included in the epistemic community in such a way that they can provide testimony/evidence. Second, one might think that sometimes marginalized groups do not have access to dominant groups' testimony/evidence at all (Narayan, 1988). Interestingly, the base model on one-sided testimonial ignoration that I will present in the paper can be reinterpreted to model this first alternative form of epistemic marginalization as well. This is because the asymmetry in evidence updating dynamics in the model can be interpreted both as marginalized agents providing testimony that is subsequently ignored by the dominant group, *and* as marginalized agents not providing testimony to dominant agents at all (either due to unwillingness or due to epistemic exclusion). These are two very different forms of epistemic marginalization, which interestingly share the same structural form. The models I present in this paper unfortunately do not apply to the second alternative form of epistemic marginalization. I address this limitation in more detail in §1.5 and leave the additional modeling work for future research. Thanks to an anonymous reviewer for raising these points.

Specifically, I ask, is it the case that simply by virtue of their testimony being ignored or devalued, members of the marginalized group gain epistemic advantages that foster more accurate beliefs?

I use computer simulations to investigate this question. In my models, members of the marginalized group end up with a number of epistemic advantages, by virtue of their testimony being ignored or devalued by the dominant group. Failure of testimonial reciprocity *can* hence render the inversion thesis true. Though the models I use are highly idealized and abstract, I argue that my simulations provide one *possible* explanation for the inversion thesis, by casting it as a consequence of the unidirectional failure of testimonial reciprocity.[3]

I construct three models to support my argument. My models are adapted from the network epistemology framework developed by Bala and Goyal (1998) and introduced to philosophy of science by Zollman (2007).[4] In all previous implementations of the model in philosophy, the network connections are reciprocal, meaning that if agent Y updates on agent Z's evidence, then Z updates on Y's evidence in the same fashion. In contrast, the network connections in my models are not reciprocal when agents interact with outgroup members.

I start with a base model of one-sided testimonial ignoration. Here, dominant agents ignore testimony shared by marginalized agents, but marginalized agents update on all testimony shared with them.[5] I find that marginalized agents arrive at the true belief more frequently and faster, and select epistemically better actions during the learning process, as compared to dominant agents. Moreover, marginalized agents arrive at the true belief more frequently

[3]See Bokulich (2014) for discussions on how-possibly explanations and how-actually explanations. Moreover, though my results may inform real world processes, I do not claim that unidirectional failure of testimonial reciprocity necessarily underlies *all* real world scenarios where the inversion thesis holds. I leave open the possibility that other phenomena may also lead to the inversion thesis. That is, I provide a sufficient condition for the inversion thesis under reasonable assumptions, but not a necessary condition.

[4]Variations of this model have seen fruitful applications in the philosophy of science and social epistemology, e.g. Zollman (2007), Zollman (2010), Mayo-Wilson et al. (2011), and O'Connor and Weatherall (2018).

[5]I use "marginalized agents" to denote "members of the marginalized group," and "dominant agents" to denote "members of the dominant group."

even compared to agents in a model with perfect testimonial reciprocity, and dominant agents do so less frequently. The entire community in this model learns the true belief less frequently and more slowly than the community with perfect testimonial reciprocity. These results show that the one-sided testimonial ignoration practiced by the dominant group is epistemically detrimental to its members and to the entire community, but is epistemically advantageous to the marginalized group.

I then construct two variations of the base model—one for one-sided testimonial ignoration with the homophilic network structure and one for one-sided testimonial devaluation. I use homophilic networks—where agents prefer to connect with ingroup members—because many real world networks are homophilic (McPherson et al., 2001) and because these networks allow me to vary agents' information access based on group membership. I find that, regardless of their information access, members of the marginalized group arrive at the true belief more frequently than the dominant group in homophilic networks. The degree to which marginalized agents gain other epistemic advantages, such as their speed of learning, depends on their information access. Finally, I build a model of one-sided testimonial devaluation, where dominant agents discount, rather than ignore, testimony from the marginalized group. Here, marginalized agents arrive at the true belief faster and select epistemically better actions during the learning process, as compared to dominant agents.[6]

The paper will be organized as follows. §1.2 introduces and motivates the base model, as well as presents the key results. §1.3 discusses the first variation with homophilic network structures. §1.4 presents the second variation on testimonial devaluation. §1.5 discusses how my results relate to standpoint epistemology. In closing, I will briefly note how my models, by virtue of introducing nonreciprocity to network connections, complicate the understanding and applications of previous network results.

[6]Due to model design, the marginalized and dominant groups necessarily learn the true belief with the same frequency.

1.2 The Base Model: one-sided testimonial ignoration

1.2.1 The Model

The base model consists of a network of agents who are presented with the same learning problem. Agents are tasked with learning which of the two available options is better, by updating on evidence from their neighbors and themselves. The network has two subgroups—the marginalized group and the dominant group—with their members distinguished only by the updating dynamics. Dominant agents only update on evidence shared by ingroup members, whereas marginalized agents update on evidence shared by everyone.

The Bandit Problem

To motivate my model, let us consider a toy example. The scenario is not meant to be realistic, but rather to illustrate the cases to which my models are applicable in a high-level way. Suppose that an organization hires for a position, and eventually offers the position to a candidate from a particular social group X.[7] Suppose further that this is not the first time that a candidate from X is hired, and Hana, a consultant who has access to some details of the case, is tasked to investigate why the candidate was hired. There is an available individualistic meritocratic explanation, A, according to which a candidate from X is hired because they are the best at doing the job out of all candidates. Explanation A is well understood, but is only known to succeed about half the time when applied to similar cases. For instance, there might have been multiple candidates who were equally good at the job but only one was hired. Explanation A is all right, but is inadequate as a catch-all explanation. Recently, a new explanation called structural bias explanation, B, is also hypothesized to account for this kind of social phenomena. Explanation B says that a

[7]To make it feel more concrete , readers can substitute “X” with “White,” “Male,” “Able-bodied,” etc.

candidate from social group X is hired because during the hiring process there was structural bias against candidates who are not members of X. Explanation B is not well understood, and the community is unsure whether it is better or worse than A. Suppose that Hana happens to have the initial belief that B would be better than A in this case, so she decides to solicit evidence to test B by, for instance, investigating whether the company's job description contains biased language.[8] She then learns from the evidence she discovers, forms posterior beliefs on the explanations, and continues to test B if she thinks that it is better. If she keeps getting good evidence for B, eventually she should believe that B is better than A with overwhelming confidence. Or, if the evidence Hana gets for B is unsatisfactory, she would give up on testing B because she thinks it is inferior to the available alternative.

The above scenario can be modeled by what is called a two-armed bandit problem. The name "bandit problem" comes from applying it to a gambling scenario, where a gambler, facing a many-armed "bandit," aims at maximizing their profit and choosing the best-performing arm by interacting with the machine. Here is how the problem is set up for one agent. For every time step (or "round"), the agent selects between two actions: A and B, and gets payoffs based on their choice. Each of the actions is associated with a fixed probability of success. The success rate for A is well-known to the agent, and is set to .5. The success rate for B, however, is uncertain to the agent: the agent knows that action B is either slightly better than A, with a success rate of $.5 + \epsilon$, or it is slightly worse than A, at $.5 - \epsilon$. When an action generates success, the agent receives a payoff of 1, and they receive no payoffs otherwise. In the models I implement, unbeknownst to the agents, I set the success rate of B to $.5 + \epsilon$. The goal for the agent is to determine which action has a higher success rate by learning from their own actions and payoffs.[9]

[8] Here we stipulate that, depending on what hypotheses to test, Hana performs different actions, which then provide Hana with evidence for the chosen hypothesis. Furthermore, Hana has limited resources to test explanations, so she is incentivized to test the better explanation each round. See §1.5 for potential limitations of this stipulation.

[9] Agents learn by applying Bayes' rule. For more detail, see §1.2.1.

Thus constructed, the two-armed bandit problem is suitable to model learning situations where there are two competing choices. This model is often applied in epistemology to scenarios with two competing theories, explanations, or hypotheses available for a given phenomenon. Besides Hana's quest to figure out which of the two explanations best accounts for the company's hiring decision, another classic application of the model is the clinical trials of drugs. Here, action A represents a drug that is well-understood, and the doctor's goal is to figure out whether a new drug, B, is better or worse than A.

I use the bandit problem because we can set the true state of the world one way and observe how well the agent learns the true belief, which naturally models epistemic advantage. Moreover, when we cast the bandit problem in a social setting, the evidence sharing dynamic becomes a suitable place to implement the unidirectional failure of testimonial reciprocity, as I will discuss shortly.

Going Social

In many real cases, learning is not an isolated, individualistic activity. Doctors are often not alone when they test new drugs; they may be in contact with other doctors in the same clinical trials. Hana may also have a team of consultants on the same case.

When multiple individuals figure out the same problem together, they can share evidence and incorporate others' evidence in their own learning. To model this, I introduce a network of agents who face the same two-armed bandit problem. Each agent is connected to some or all of the other agents and I call the agents they are connected to their "neighbors."[10] In each round, each agent selects their action based on their belief in the proposition "B is better than A," obtains evidence from their action, shares their evidence with their neighbors, and

[10] A *network structure* describes how agents are connected to each other.

updates on the evidence they receive. Further, because the evidence is passed on between agents, a piece of evidence is a form of testimony—it has a speaker and an audience.

Zollman (2007) builds a model just like this. He simulates the model on different network structures and finds, perhaps counterintuitively, that a more sparsely connected community has epistemic advantages over a more connected one, in the sense that the former learns the true belief more frequently. However, a more sparsely connected community also learns the truth more slowly. Zollman's findings are collectively dubbed "the Zollman effect." They uncover a trade-off between the speed and accuracy in social learning. The effect will become relevant later in two ways. First, I will use the reasoning behind the effect to explain some of my modeling results. Second, I will argue that my results complicate the interpretation and application of the Zollman effect.

De-idealizing Testimonial Relationships

In Zollman (2007)'s model, if two agents are connected, then they share their evidence with each other and fully update on the evidence they receive. That every agent treats all testimony they receive equally is, of course, an idealization, and perhaps an unwarranted one. Social epistemologists working on race and gender, notably standpoint epistemologists, have written extensively on dominant social groups' widespread practice of silencing testimony from marginalized perspectives.

Recall the epigraph of this paper. Mills (2007) argues that the primary epistemic principle of a racialized social epistemology is that people of color are not seen as knowers. In the same chapter, Mills gives the example of Kant's dismissal of a Black carpenter's epistemic credibility: "and it might be, that there were something in this which perhaps deserved to be considered; but in short, this fellow was *quite black* from head to foot, a clear proof that what he said was stupid" (Kant, 1960, 113, emphasis in the original, qtd. Mills 2007, 32). Moreover, several authors point out that people of color's testimony is often not taken seriously unless they have white authenticators (Mills, 2007; Fatima, 2017; Bright, 2023), and women's testimony is often ignored until repeated by men, a phenomenon dubbed "hepeating" (Deo, 2019). It is hypothesized that Ida B. Wells-Barnett (2012), when arguing against lynching in 1895, only includes evidence from white sources because her intended white audience would trust these sources rather than the testimony of Black people (Bright, 2023).[11] As a contemporary example, Deo (2019) conducts an empirical study on how race and gender influence legal academia,[12] and finds that "most women in the [study] sample, regardless of racial/ethnic background, have endured silencing, harrassment, mansplaining, hepeating, and gender bias" (Deo, 2019, 47). For instance, one study participant has "counted over ten times on [her] faculty where [she has] said something and [nobody has responded; then] a male faculty has repeated it and another male colleague has said, 'Good idea!"' (Deo, 2019, 45). Writing about her experience as a woman of color in the predominantly white and male field of professional philosophy (an experience that corroborates Deo's findings), Fatima (2017, 151) claims that "if the only way that a woman of color's testimony is given any uptake is if dominant members of academia verify it, then we have already discounted the epistemic credibility of the speaker." The existence and prevalence of silencing testimony from marginalized perspectives is widely recognized.

Importantly, this dismissal of testimony is one-sided. Mills (2007, 17) argues that "often for their very survival, blacks have been forced to become lay anthropologists, studying the strange culture, customs, and mind-set of the 'white tribe' that has such frightening power over them, that in certain time periods can even determine their life or death on a whim." Mills quotes Baldwin's brutally honest line, "I have spent most of my life, after all,

[11] Though Bright (2023) eventually favors an alternative, statistically-based explanation for Wells-Barnett's decision, the original hypothesis is still plausible as testimonial ignoration was undoubtedly salient at the time.

[12] In this comprehensive study, Deo (2019) presents quantitative and qualitative results from a core sample comprising almost 10% of all US women of color law professors, together with a comparison sample of white or men of color law professors.

watching white people and outwitting them, so that I might survive" (Baldwin, 1993, 217, qtd. Mills 2007, 18). In Deo (2019)'s study, a woman of color participant admits that she always acquiesces to requests from university administration, sometimes even unreasonable ones, because she fears professional repercussions if she declines. In order to survive in a hegemony dominated by other people, marginalized folks *cannot afford to* ignore testimony and demands from the dominant group.

Given this, it is appropriate to de-idealize the testimonial relationships in the model to incorporate one-sided testimonial ignoration.[13] I implement testimonial ignoration by dividing the population into two groups: the marginalized group and the dominant group. Marginalized agents update on evidence from all their neighbors, but dominant agents only update on evidence shared by ingroup neighbors.[14] Here, we have a failure of testimonial reciprocity—marginalized agents take testimony from the dominant group as they are meant to be taken, but dominant agents fail to do so reciprocally. To bring us back to the hiring scenario, there might be a few consultants who do not trust Hana's evidence, as (supposedly) Hana is not a member of the social group X and, for them, Hana might have bias against X.[15]

[13]I will not address *how* we identify situations with one-sided testimonial ignoration. Dotson (2011) has gracefully tackled this question.

[14]One might worry that the evidence I presented in the previous paragraph only supports the claim that *insofar as* marginalized agents receive evidence from dominant agents, they cannot afford to ignore the evidence, but not so much that marginalized agents receive all the evidence from their dominant neighbors. This is a very fair concern, especially given that sometimes marginalized knowers are excluded from participating in epistemic communities (more in §1.5). The base model I present in this section makes the idealizing assumption that marginalized agents receive evidence from all their neighbors. However, one can reinterpret the first variation presented in §1.3 with homophilic networks, where p_{outgroup} is small, as modeling some version of epistemic exclusion. Here, marginalized agents only have very sparse evidential access to the dominant group. We can think of this as the poignant situation where a small number of marginalized knowers are invited to participate in dominant epistemic spaces, but their testimony is still ignored. Thanks to an anonymous reviewer for raising this concern.

[15]For a more realistic example, consider Blanche, a Black fill-in maid for a white family in the novel *Blanche on the Lam* (Neely, 1993). As a lower class Black woman, Blanche's epistemic credibility is not fully recognized by other members of the family that employs her; but Blanche continues to listen to and in on the family members during her work. This eventually allows Blanche to gather enough evidence and solve a series of murder mysteries in the family. This example of Blanche's standpoint is discussed at length in Wylie (2003). Though the example of Blanche is fictional, the phenomenon that Black domestic helpers have their epistemic credibility suppressed by their employers, but still gain an outsider-within status in white middle-class families is discussed in detail in Collins (2002).

Further Technical Details

A few technical details are in order before I present my results.

Initiation

I set the success rate for A to be .5, and the success rate for B to be $.5+\epsilon$. At the start of the simulation, every agent is assigned a credence randomly selected from a uniform distribution between 0 and 1 (exclusive). The credence reflects their belief in the proposition "B is better than A."

A Typical Round

At the start of every round, each agent selects one of the two actions—if their credence is $> .5$, they choose action B; otherwise they choose A. The agent then performs the chosen action a number of times, n, and receives payoffs.

Then, each agent uses Bayes' rule to update their credence based on both their own experience and the experiences of their neighbors. For example, suppose that $\epsilon = .1$ (i.e. the success rate of B is .6). If an agent has prior credence of .7 that B is better than A, and pulls action B one time, which succeeds in generating a payoff of 1, then their posterior credence after updating on their own experience is

$$P(H|E) = \frac{P(E|H)P(H)}{P(E|H)P(H)+P(E|\neg H)P(\neg H)} = \frac{.6\times.7}{.6\times.7+.4\times.3} = .78.$$

Here, H (hypothesis) stands for "B is better than A," and E (evidence) is "taking action B once yields 1 payoff." It is worth noting that, whether or not the action succeeds in generating a payoff, performing action A will not change the posterior credence. To see that, we observe that $P(E'|H) = P(E'|\neg H) = P(E''|H) = P(E''|\neg H) = .5$, where E' is "taking action A once yields 1 payoff," and E'' is "taking action A once yields 0 payoffs." Consequently, $P(H|E') = P(H|E'') = P(H)$. Agents similarly update on neighbors' evidence by applying Bayes' rule. Note that dominant agents only update on evidence from ingroup neighbors and marginalized agents update on evidence from all neighbors.

After each agent finishes updating, we increase the time step by 1 and repeat the procedure for a typical round.

End of Learning

There are three stable end states for this model:

- Community convergence to the true belief: every agent has a credence of $> .99$ that B is better than A. In this state, it is increasingly unlikely that agents would switch from B to A. Everyone succeeds in learning.
- Community convergence to the false belief: every agent has a credence of $\leq .5$ that B is better than A. In this state, nobody would be actively testing B. Everyone fails in learning.
- Polarization: every marginalized agent has a credence of $> .99$ that B is better than A, and every dominant agent has a credence of $\leq .5$. In this state, no dominant agent would be actively testing B. Every marginalized agent succeeds in learning, every dominant agent fails in learning, and the entire community fails in learning.

Due to one-sided testimonial ignoration, polarization is a new end state for my model compared to Zollman (2007). In this state, even though marginalized agents are still testing action B, their testimony is ignored by dominant agents. We have a stable situation where the agents' beliefs are split along group membership. It is worth noting that a polarization with the opposite distribution of credence cannot be stable, since in this state, marginalized agents would still update on evidence from dominant agents, and the model would evolve. If the network reaches one of the stable end states, the community has finished learning.

1.2.2 Results

I simulate this model using the following values for the key parameters, for $10,000$ runs each:

- Total number of agents ("size") of the network (k) : $3, 6, 12, 18$.
- Proportion of the marginalized group in the population(d): $\frac{1}{6}$, $\frac{1}{3}$, $\frac{1}{2}$, $\frac{2}{3}$.
- Number of pulls per round (n): $1, 5, 10, 20$.
- Probability of B (P_B): .51, .55, .6, .7, .8.
- Network structure: complete.[16]

For all parameter values, I run a comparison model with perfect testimonial reciprocity; this is equivalent to the complete model from Zollman (2007). Moreover, all the results in this paper are robust across all parameter values, unless otherwise noted.

I employ three ways to measure how well subgroups learn: the (1) frequency and (2) speed at which subgroups learn the true belief, and the (3) frequency at which subgroup members select the epistemically better action during the learning process. The marginalized group holds epistemic advantages compared to the dominant group according to all three measures.

I measure how frequently a subgroup learns the true belief by calculating the proportion of simulation runs (out of $10,000$ runs) where the subgroup succeeds in learning. This measure captures how often an average agent of a given subgroup eventually learns the true belief. Here, the marginalized group learns the true belief more frequently (Figure 1.1) because polarization counts as success for marginalized agents and failure for dominant agents. As long as there are simulation runs that end in polarization, the marginalized group would learn better in this respect. In fact, for all parameter values, the marginalized group learns the true belief more frequently.

[16] A network is complete when everyone is connected to everyone. The network structure here is complete prior to adding one-sided testimonial ignoration.

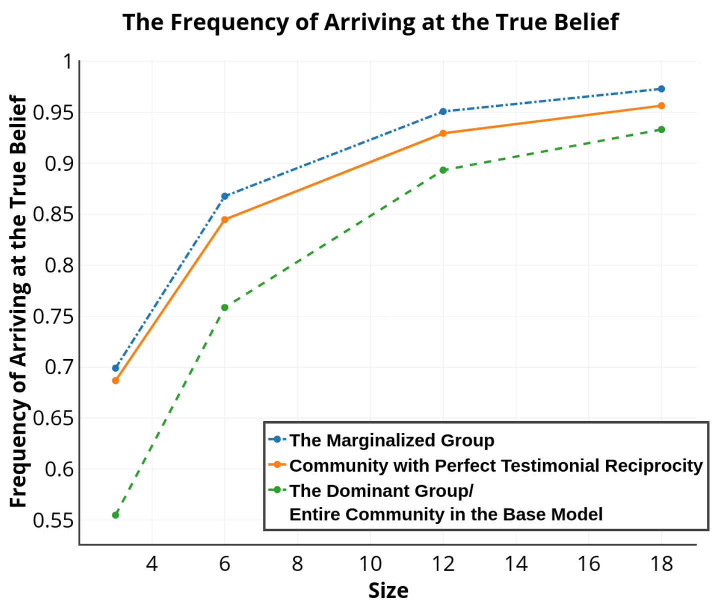

Figure 1.1: Base Model, $P_B = .6, n = 1, d = \frac{1}{3}$, 10,000 simulation runs.

A subgroup's speed of learning the true belief is measured as follows. For each agent and each of their successful runs, I document the earliest round after which the agent maintains a credence of $> .99$. Call this an individual agent's rounds to successful learning. Then, for each subgroup, I compute the subgroup's average rounds to successful learning by taking the average of all individual agents' rounds to successful learning, out of all successful simulation runs and all agents in the subgroup. Unlike some previous measures of speed of learning, here I only consider cases where the subgroup learns the truth.[17] This measure represents how long an average agent from a given subgroup takes to learn the true belief. Marginalized agents learn the true belief faster (Figure 1.2) because they have access to more information per round.

[17]My measure differs from previous ones in the literature. Different from Zollman (2007)'s "average time to success," I measure a subgroup's speed of learning not by observing when the entire community reaches the true belief, but by taking the average of individual agent's rounds to successful learning. "Rounds to consensus" in O'Connor and Weatherall (2018) measures the entire community's rounds to consensus, regardless of truth or falsity. In contrast, I only consider cases where the subgroup learns the truth. This measure allows me to capture possible differences in the speed of successful learning between the two subgroups. For instance, there could be possible variations in subgroups' rounds to successful learning even when the community converges to the true belief. Later, I will introduce another measure: the *entire community*'s average rounds to successful learning, which is the same as Zollman (2007)'s.

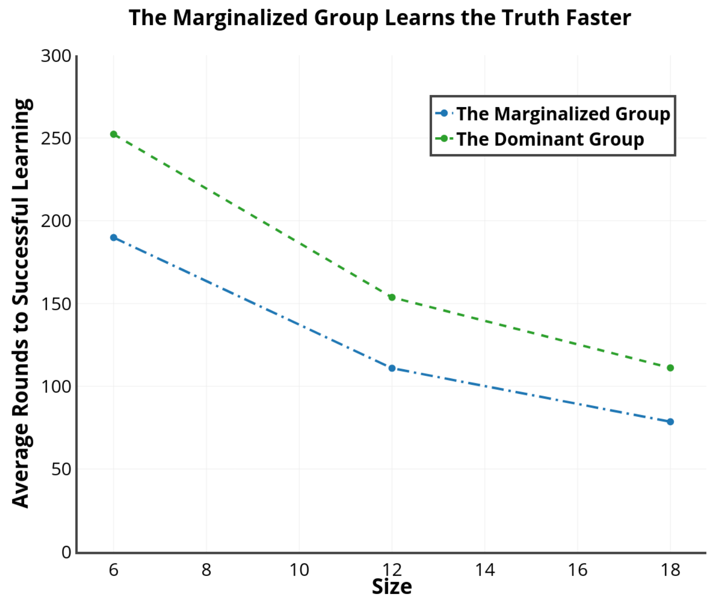

Figure 1.2: Base Model, $P_B = .51, n = 10, d = \frac{1}{2}$, $10,000$ simulation runs.

I measure how frequently subgroup members select the epistemically better action during the learning process by calculating the average proportion of rounds that an agent of this subgroup selects action B, out of all simulation runs and all agents in the subgroup, without considering every agent's selection at round 0.[18] This measure captures how frequently an average agent from a given subgroup chooses the epistemically better action during the learning process. It is not surprising that marginalized agents select the better action more frequently, since they also eventually learn the truth more frequently and faster.

So far, my results adhere to a basic empiricist intuition, that access to more information provides epistemic advantages. However, the following results, together with results from the first variation (§1.3), suggest that marginalized agents' epistemic advantages cannot solely be explained by having access to more information.

Compared with the community with perfect testimonial reciprocity, where all evidence is fully updated by the receiver, the marginalized group in my model learns the true belief

[18] An agent's action at round 0 only depends on their initial credence as randomly selected by a uniform distribution. Including this round would add noise to the data, especially when agents learn very fast (when P_B and n are large).

more frequently, and the entire community as well as the dominant group learns the true belief less frequently (Figure 1.1). A marginalized agent holds this epistemic advantage even though they update on as many pieces of evidence as an agent in the model with perfect testimonial reciprocity. This is because in simulation runs that start with initially unpromising results for action B,[19] the community with perfect testimonial reciprocity would quickly settle on action A, resulting in a convergence to the false belief. In my model, though marginalized agents would change their actions quickly, the dominant agents would not due to their limited information access. Consequently, the epistemically better action remains active in the network for longer, making it more likely that marginalized agents would turn around to correct their action. This is an instance of the Zollman effect, where more sparsely connected network structures produce epistemic benefits.[20]

The entire community in my model learns the true belief less frequently than the community with perfect testimonial reciprocity because a necessary condition for the former community to learn the true belief is that the dominant group learns it. But the dominant group, because of the one-sided testimonial ignoration, acts as an isolated model of size $k \cdot (1 - d)$ with perfect testimonial reciprocity.[21] For models with perfect testimonial reciprocity, the smaller the size of the network, the less frequently the community learns the true belief (Zollman, 2007). The dominant group, and hence the entire community, learns the true belief less frequently than the community of size k with perfect testimonial reciprocity.

Moreover, the entire community in my model learns the true belief more slowly than the community with perfect testimonial reciprocity (Figure 1.3). The speed of successful learn-

[19]Because whether each action succeeds is probabilistic, these scenarios occur in my simulations.

[20]This result is robust with parameters such that the average rounds to successful learning for the entire community is > 3 (i.e. excluding "easy" learning situations with large P_B and n). One reason for the non-robustness in the edge cases is related to the trade-off between the learning speed and learning accuracy. For cases where the learning is really fast, the speed of successful learning is often very close between the marginalized group and the community with perfect testimonial reciprocity, so their learning accuracy is also comparable. Moreover, when agents finish their learning quickly, their learning accuracy is highly dependent on their initial beliefs, which are randomly assigned and highly variable.

[21]Recall that k is the network size and d is the proportion of the marginalized group.

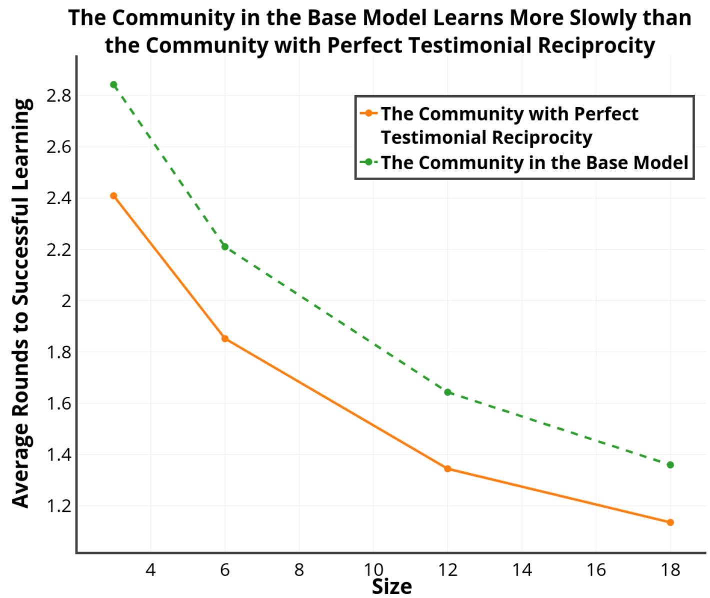

Figure 1.3: Base Model, $P_B = .7, n = 5, d = \frac{1}{3}, 10,000$ simulation runs.

ing for the *entire community* is measured differently from the subgroups, to facilitate a direct comparison with Zollman (2007)'s model. I first record, for each simulation run with community success, the round at which the community finishes learning. I then define the average rounds to successful learning for the entire community as the average of these rounds out of all successful simulation runs. This measure captures how long an entire community takes to learn the true belief. The entire community in my model learns the true belief more slowly for a similar reason—the dominant group, as an isolated group with perfect testimonial reciprocity of size $k \cdot (1 - d)$, learns the truth more slowly than a community with perfect testimonial reciprocity of size k (Zollman, 2007).

Finally, the proportion of the marginalized group (d) impacts the degree of epistemic (dis-)advantage. As d increases, the marginalized group learns the true belief more often,[22] and the dominant group less often. This is because the size of the dominant group decreases as d increases. In the face of unpromising initial results for B, dominant agents are even more less likely to quickly give up on B due to their further limited information access, resulting

[22]This is robust with parameters such that the average rounds to success for the entire community is > 3.

in more gain in learning accuracy for the marginalized group. Furthermore, as d increases, the entire community learns the truth more slowly for similar reasons.

In this model, I do not restrict the proportion of the marginalized group in the population. What defines a subgroup's marginalized status is the one-sided testimonial ignoration, rather than its size. This is a merit since in many real cases, the marginalized group can be in the majority, such as during the Apartheid in South Africa.

1.3 Variation 1: Homophilic Networks

1.3.1 The Model

As mentioned, some results from the based model can be explained by marginalized agents having access to more information, but other results cannot solely be explained by information access. Simulating the model with homophilic networks—where agents prefer to connect with ingroup members—allows me to further investigate the extent to which information access influences epistemic advantages. Moreover, homophily is a natural choice, as many real human networks are homophilic based on race, gender, class, etc. (McPherson et al., 2001). My results show that marginalized agents still hold a number of epistemic advantages, even when they have equal or fewer expected connections compared to dominant agents.

This variation differs from the base model in network structure only. I use two-type random graphs to generate the homophilic networks.[23] First, every agent is connected to themselves. Then, I divide the agents into marginalized and dominant groups. Each agent has some probability of connecting with ingroup members, P_{ingroup}, and some other probability of connecting with outgroup members, P_{outgroup}. Finally, I require that, prior to adding one-

[23]See Golub and Jackson (2012) and Rubin and O'Connor (2018) for previous implementations of homophilic networks using this method.

sided testimonial ignoration, the network structure is undirected. This means that Y is connected to Z if and only if Z is connected to Y. The network is homophilic when $P_{\text{ingroup}} > P_{\text{outgroup}}$.

When a subgroup is in the minority, its members have fewer expected connections than members of the other group, prior to adding one-sided testimonial ignoration. To see this, first observe that the expected number of connections for an agent in this subgroup is

$$P_{\text{ingroup}} \cdot (k \cdot d' - 1) + P_{\text{outgroup}} \cdot k \cdot (1 - d') + 1,$$

where k is the size of the network, and d' is the proportion of this subgroup in the population. For an agent in the other group, their expected number of connections is

$$P_{\text{ingroup}} \cdot (k \cdot (1 - d') - 1) + P_{\text{outgroup}} \cdot k \cdot d' + 1.$$

When $d' < \frac{1}{2}$ and $P_{\text{ingroup}} > P_{\text{outgroup}}$, an agent in this subgroup has fewer expected connections than an agent in the other group.[24]

When one-sided testimonial ignoration is added and when the marginalized group is in the minority, we can specify the values of P_{ingroup}, P_{outgroup}, and d such that a marginalized agent would have fewer, equal, or more expected connections as compared to a dominant agent. For example, a marginalized agent would have the same number of connections as a dominant agent when $P_{\text{ingroup}} = .8$, $P_{\text{outgroup}} = .4$, and $d = \frac{1}{3}$, fewer expected connections when $P_{\text{ingroup}} = .8$, $P_{\text{outgroup}} = .3$, and $d = \frac{1}{3}$, and more expected connections when $P_{\text{ingroup}} = .8$, $P_{\text{outgroup}} = .5$, and $d = \frac{1}{3}$. Simulating with homophilic networks, then, can reveal the extent to which information access influences marginalized agents' epistemic advantages.

[24]The result holds probabilistically and is not necessarily true for individual simulation runs.

1.3.2 Results

I simulate this variation using the following values for the key parameters, for $10,000$ runs each:[25]

- Size of the network (k) :18.
- Number of pulls per round (n): 1, 5, 10, 20.
- Probability of B (P_B): .51, .55, .6, .7, .8.
- Proportion of the marginalized group (d): $\frac{1}{6}$.
  - $P_{\text{ingroup}} = .8, .9, 1$.
  - $P_{\text{outgroup}} = .6, .7, .8$.
- Proportion of the marginalized group (d): $\frac{1}{3}$.
  - $P_{\text{ingroup}} = .7, .8, .9$.
  - $P_{\text{outgroup}} = .3, .35, .4, .45, .5$.

I choose different values for P_{ingroup} and P_{outgroup} based on d because depending on the value of d, the values of P_{ingroup} and P_{outgroup} needed for a marginalized and a dominant agent to have the same number of connections are different.

Furthermore, I only simulate connected* networks in order to reduce noise in the data.[26] Because of the one-sided testimonial ignoration, I define connectedness* as follows. A network is connected* when (1) there exists a path from any marginalized agent to any arbitrary agent in the network, and (2) there exists a path from any dominant agent to any arbitrary dominant agent. Moreover, there exists a path from agent Y to agent Z iff there are agents $A_0, A_1, ..., A_i$ with $i \geq 1$ in the network such that (1) $Y = A_0$ and $Z = A_i$, and (2) A_k updates on evidence shared by A_{k+1}, with $0 \leq k \leq i-1$. I only test the network size of 18 for two practical reasons. First, the total number of simulations is already large due

[25] I randomly generate a homophilic network for every simulation run.

[26] If the network is not connected*, then there would necessarily be two or more isolated communities without any evidence sharing in between. This network would produce less than typical learning speed and learning accuracy, compared to connected* counterparts with same parameter values.

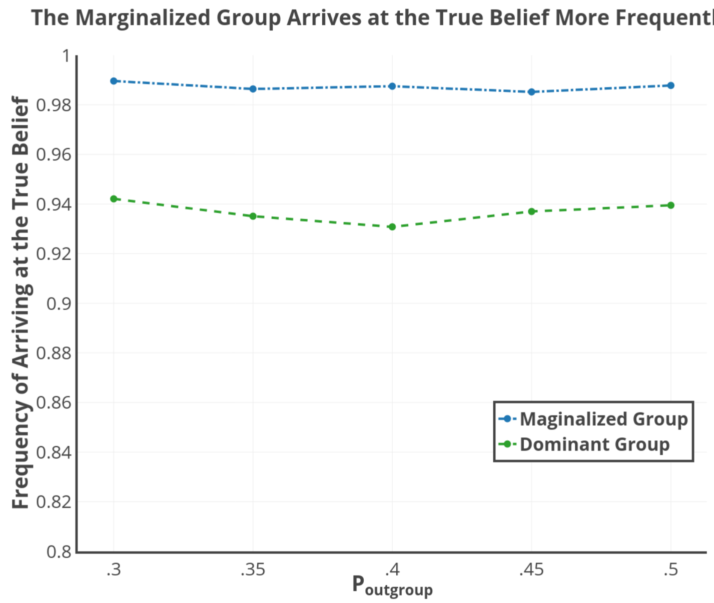

Figure 1.4: Variation 1, $P_B = .55, n = 5, d = \frac{1}{3}, P_{\text{ingroup}} = .8, k = 18, 10,000$ simulation runs.

to variations in P_{ingroup} and P_{outgroup}. Moreover, networks with large populations are more likely to be connected* after the two-type random graph generation process.

My results show that the marginalized group learns the true belief more frequently than the dominant group, regardless of their expected numbers of connections. Similar to the base model, this is due to polarization. When polarization occurs, the marginalized group succeeds in learning but the dominant group fails, creating a disparity in learning accuracy. Moreover, the frequency of learning the true belief for the two subgroups does not change drastically as P_{outgroup} changes (Figure 1.4). This is because P_{outgroup} does not influence the epistemic behavior of the dominant group as an isolated community. As a result, the epistemic benefits the marginalized group gains remain the same.

How the two subgroups compare regarding their learning speed and the frequency of choosing the epistemically better action depends on their members' number of connections. When a marginalized agent has the same expected number of connections as a dominant agent, the former selects the epistemically better action (i.e. action B) more frequently than the latter

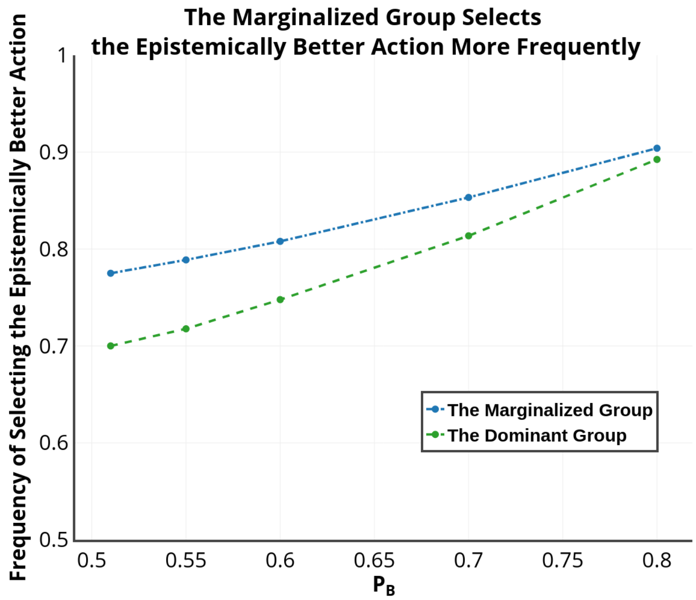

Figure 1.5: Variation 1, $n = 1, d = \frac{1}{3}, P_{\text{ingroup}} = .9, P_{\text{outgroup}} = .45, k = 18, 10,000$ simulation runs.

during the learning process (Figure 1.5).[27] However, unlike the base model, marginalized agents in general learns the true belief more slowly than dominant agents.[28]

The reason is that marginalized agents learn much more slowly when simulations end in polarization than when simulations end in community success, but polarization still counts as success for them.[29] If we discount polarization from the marginalized group's average rounds to successful learning, then the marginalized group also learns the truth faster than the dominant group in many but not all cases.

When a marginalized agent has more expected connections than a dominant agent, the marginalized group's epistemic advantage in speed of successful learning becomes more robust, while its members continue to hold the other advantages. As the difference in the expected numbers of connections grows, eventually the average rounds to successful learn-

[27]The behaviors of the two subgroups get closer as P_B increases. When P_B is large, the two possible states of the world become easier to distinguish. Therefore, the agents finish quickly, at less than 2 rounds. The frequency of selecting the better action conveys less information as P_B gets large.

[28]Measured in the first way introduced in §1.2.2.

[29]The marginalized group takes around two to five times more rounds to learn the true belief when simulations end in polarization than those that end in community success.

ing for the marginalized group, including polarization, would be lower than that for the dominant group.[30]

When a marginalized agent has fewer expected connections than a dominant agent, the marginalized group loses its epistemic advantage in the speed of successful learning. However, in many cases where the difference in numbers of expected connections is small (≤ 2), the marginalized group retains its epistemic advantage in the frequency of selecting the epistemically better action during learning. Hence, the epistemic advantages brought to the marginalized group by one-sided testimonial ignoration is sometimes strong enough to offset the potential disadvantages from the loss of information access.

1.4 Variation 2: one-sided testimonial devaluation

1.4.1 The Model

I now simulate testimonial devaluation, where dominant agents devalue evidence from the marginalized group, rather than ignore it. The model differs from the base model only in updating rules. Here, I introduce Jeffrey conditionalization, which allows agents to update on evidence according to how much they trust the accuracy of it. The formula for Jeffrey conditionalization is the following:

$$P_f(H) = P_i(H|E) \cdot P_f(E) + P_i(H|\neg E) \cdot P_f(\neg E).$$

The agent's final credence for the hypothesis H ($P_f(H)$) is defined as the agent's initial credence for H after Bayesian conditioning on the evidence E being true ($P_i(H|E)$) times the agent's credence that E is accurate ($P_f(E)$), plus the the agent's initial credence for H

[30]The base model is a special case of a two-type random graph, where $P_{\text{ingroup}} = P_{\text{outgroup}} = 1$. The results in §1.2.2 fits with those here.

after Bayesian conditioning on E being false ($P_i(H|\neg E)$) times the agent's credence that E is inaccurate ($P_f(\neg E)$). When $P_f(E) = 1$, the agent fully trusts the accuracy of E, and Jeffrey conditionalization reduces to Bayes' rule. When $P_f(E) = P_i(E)$, i.e. the agent's final credence for E equals their initial credence for E,[31] then $P_f(H) = P_i(H)$, i.e. the agent keeps their original credence for H and ignores the evidence altogether. When $P_f(E)$ is between $P_i(E)$ and 1, the agent positively updates on the evidence, though not fully.

O'Connor and Weatherall (2018) use Jeffrey conditionalization to model situations where agents do not fully trust the evidence gathering practices of other agents. In their models, agents' final credence for the evidence $P_f(E)$ is based on how similar the sharer's belief is to the updater's. Agents find the evidence shared by someone with similar beliefs more trustworthy. In my model, however, the distrust is not based on the relative similarity of beliefs, but rather on group membership.

As before, the entire population is divided into the marginalized group and the dominant group. Marginalized agents fully update on evidence shared by all neighbors; dominant agents, in contrast, fully update on evidence shared by ingroup neighbors, but devalue evidence shared by outgroup members by applying Jeffrey conditionalization, with $P_f(E)$ calculated using:

$$P_f(E) = 1 - m \cdot (1 - P_i(E)).^{32}$$

Here, m is a parameter between 0 and 1, capturing how much dominant agents devalue testimony from the marginalized group. When $m = 0$, $P_f(E) = 1$ and dominant agents fully update on evidence from marginalized agents—the model is equipped with perfect

[31] $P_i(E)$ can be calculated from $P_i(H)$ in the following way: $P_i(E) = P_i(E|H)P_i(H) + P_i(E|\neg H)P_i(\neg H)$.

[32] There are several formulae for $P_f(E)$ that would satisfy the desiderata below equally well. However, my results would remain largely the same had I chosen the alternatives. Furthermore, because Jeffrey conditionalization is non-commutative (c.f. Lange, 2000), I require that a dominant agent randomly selects the order according to which they update. The order of updating, to my knowledge, does not influence the qualitative results.

testimonial reciprocity. When $m = 1$, $P_f(E) = P_i(E)$ and dominant agents fully ignore evidence from marginalized agents—the model becomes my base model. For this variation, I simulate cases with $m \in (0, 1)$, i.e. I consider cases where dominant agents devalue but do not completely ignore evidence from the marginalized group. The higher the value of m, the more dominant agents devalue.

Because dominant agents still positively update on evidence from marginalized agents, polarization is no longer a stable end state. The two remaining stable end states are community convergence to the true belief and community convergence to the false belief. Thus, for every simulation run, marginalized and dominant groups end with the same belief state. They learn the truth with the same frequency.

1.4.2 Results

I simulate this model using the following values for the key parameters, for $10,000$ runs each:

- Size of the network (k) : 3, 6, 12, 18.
- Proportion of the marginalized group (d): $\frac{1}{6}$, $\frac{1}{3}$, $\frac{1}{2}$, $\frac{2}{3}$.
- Number of pulls per round (n): 1, 5, 10, 20.
- Probability of B (P_B): .51, .55, .6, .7, .8.
- Degree of devaluation (m): .2, .5, .8.
- Network structure: complete.

I run the model with perfect testimonial reciprocity for all parameter values except for m for comparison.

I find that the marginalized group arrives at the true belief faster than the dominant group (Figure 1.6).[33] The marginalized group's advantage in learning speed depends on both m and d. As m increases, the difference in the average rounds to successful learning between

[33]Measured in the first way introduced in §1.2.2.

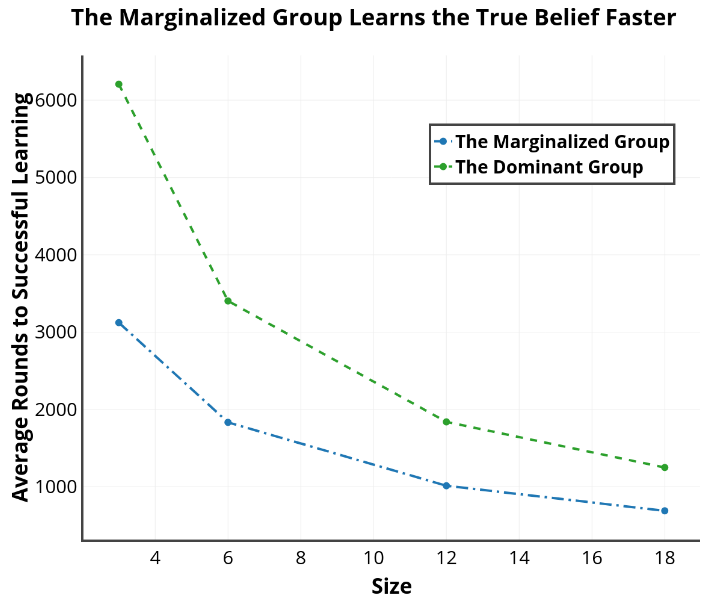

Figure 1.6: Variation 2, $P_B = .51, n = 1, d = \frac{2}{3}, m = .8., 10,000$ simulation runs.

the two subgroups widens. As d increases, the difference in learning speed also widens. Moreover, marginalized agents select the epistemically better action more frequently than dominant agents during the learning process.[34] As m increases, the difference in frequency between the two subgroups widens.[35]

Compared with the model with perfect testimonial reciprocity, the entire community in this variation learns the true belief more slowly.[36] Moreover, as m increases, the entire community in this variation learns more slowly; as d increases, it also learns more slowly. This shows that one-sided testimonial devaluation is detrimental to the entire community's learning speed, and the adverse effect becomes more salient the more dominant agents devalue marginalized agents' testimony.

[34]This result is robust for parameters such that the average rounds to successful learning is > 2 for the entire community.

[35]This result is robust for parameters such that the average rounds to successful learning is > 3 for the entire community. The degree to which marginalized agents obtain this epistemic advantage similarly depends on d, but the result is not as robust, especially when m is small.

[36]Measured in the second way introduced in §1.2.2.

Finally, comparing the community learning accuracy between this variation and the community with perfect testimonial reciprocity does not bring robust results. The frequency at which the entire community arrives at the true belief fluctuates around that of the model with perfect testimonial reciprocity. In general, as m and d grows, the entire community is more likely to arrive at the true belief more frequently than the community with perfect testimonial reciprocity. What is robust, however, is that the community in this variation always learns the true belief less frequently than the marginalized group in my base model.

1.5 A Network Standpoint Epistemology

In the above three models, marginalized agents end up with several epistemic advantages, by virtue of their testimony being ignored or devalued by the dominant group. Here, the testimonial ignoration and devaluation practiced by the dominant group is largely epistemically detrimental to its members and the entire community, but is epistemically advantageous to the marginalized group. The inversion thesis, which states that marginalized social groups hold epistemic advantages, is a key claim of standpoint epistemology, though its interpretations and justifications are often contested (see Wylie, 2003; Intemann, 2010; Toole, 2020). My modeling results contribute to standpoint epistemology in two ways. First, I provide a clear interpretation of the inversion thesis by making epistemic advantages precise using several measures. Second, I provide one possible way in which the inversion thesis can arise by showing that it follows from another key claim of standpoint epistemology, namely, the unidirectional failure of testimonial reciprocity.

Standpoint epistemology started as an application of Marx's analysis of the proletarian standpoint to the effect of the sexual division of labor in knowledge production (Hartsock, 1983). It was later extended to cover other unequal power relations' influence on knowledge production. For Hartsock (1983, 298), women's material lived experiences, such as their "relationally defined existence, bodily experience of boundary challenges, and activity of transforming both physical objects and human beings," foster more accurate beliefs for all human activities.

Hartsock (1983)'s argument faces a number of interpretive and justificatory questions. For one, it is unclear exactly what she means by "more accurate beliefs," or more broadly, epistemic advantages. Moreover, it is unclear whether the lived experiences she cites are descriptively true for all women, and it is unclear how they further lead to epistemic advantages. Lurking in the background is also a question about intersectionality—as individual human beings are subjected to different dimensions of oppression, how do we identify the subgroups that hold the most superior knowledge (see Longino, 1993)? Do we look to the ones who are the "most" oppressed within the oppressed groups for the "best" knowledge? These interpretive and justificatory questions have led to heated debates.[37] For instance, Harsock's argument is charged by Hekman (1997) with essentializing women, though Hartsock vehemently denies the charge (Hartsock, 1997). Partly due to intense debates and despite its fruitful applications, standpoint epistemology has been marginalized in contemporary philosophy (Toole, 2020).

This paper is part of a recent effort (e.g. Toole, 2020; Saint-Croix, 2020) at addressing issues facing standpoint epistemology by articulating interpretations of the theory that are neither obviously false nor trivially true,[38] and offering explanations for its claims using novel philosophical methods. First, I precisely interpret a subgroup's epistemic advantages using three measures in my models—the frequency at which a subgroup eventually learns the true belief, the speed at which a subgroup learns the true belief, and the frequency at which a subgroup selects the epistemically better action during the learning process. Second, my modeling results show one possible way in which the inversion thesis can be true; namely,

[37] For instance, Alison Wylie (2003) calls standpoint theory "one of the most controversial theories to have been proposed and debated in the 25-30 year history of second wave feminist thinking about knowledge and science."

[38] Contra Intemann (2010)'s comments.

when there is unidirectional failure of testimonial reciprocity. Thus, I provide a sufficient condition for the inversion thesis under reasonable assumptions, but not a necessary one. One might regard the claim that marginalized groups' testimony is ignored or devalued as far less controversial than the inversion thesis. Insofar as this is right, my models also have the virtue of explaining a controversial thesis by showing that it follows from something more widely accepted.[39] Given the widespread nature of one-sided testimonial ignoration and devaluation, my results may also shed light on real world cases where the inversion thesis holds.

Note that the mechanism I identified—the unidirectional failure of testimonial reciprocity—is not necessarily the ones that standpoint epistemologists such as Hartsock had in mind. Perhaps it will turn out that women's lived experiences, as Hartsock understands it, lead to their testimony being ignored or devalued, or perhaps Hartsock's reasoning would independently lead to the inversion thesis. In these cases her original justification would still be vindicated. Investigating this, however, is beyond the scope of this paper.

My models admittedly have a few limitations. To start, all agents in my models face the same learning problem, thus they share the same "reality." However, many works in philosophy of race[40] concern the fundamentally different realities faced by marginalized and dominant groups, for instance, the race disparity in policing in the US. Some might suggest that the marginalized group has more accurate beliefs because they are "fluent" in both worlds (see Mills, 2007). My models do not incorporate this notion of "dual realities." As such, though I show that the marginalized group has epistemic advantages when learning the shared "reality," I do not rule out the possibility that the marginalized group might have other kinds of (dis-)advantages due to "dual realities." I plan to explore models without a shared "reality" in my future work.

[39]Thanks to an anonymous reviewer for raising this point.
[40]E.g., the notion of "double consciousness" in Du Bois (2008).

Moreover, my models only have two groups, and I do not consider more groups or agents with multiple group membership. One would expect models with more groups to follow similar epistemic trends, but with slightly altered dynamics. For instance, a testimonial ignoration model with a third "bridge" subgroup might not have polarization as an end state.

In addition, one might worry that some of the assumptions in bandit models may be too idealized. For instance, in some real world situations, investigating one hypothesis may bring insights into other hypotheses as well. One would expect that in this situation, marginalized agents would still gain epistemic advantage in learning speed and the frequency of selecting the better action, though the entire community would eventually be able to reach the true belief reliably.[41] Thinking about how some of the assumptions can be relaxed is a worthwhile direction of future research. Moreover, as Wu and O'Connor (2023) recently notes, some of the network effects in the bandit model paradigm, such as the Zollman effect, are independently discovered in other modeling paradigms like the NK landscape model (Lazer and Friedman, 2007; Fang et al., 2010). It would be worthwhile to test if the marginalized group would end up with epistemic advantages as we apply one-sided testimonial ignoration and devaluation to the NK landscape model. If the results replicate, then this could indicate that the ways in which these modeling paradigms differ are in some sense irrelevant to the phenomena that we aim to explain (see Batterman and Rice (2014)).

Furthermore, in my models, the marginalized agents are participating members of the epistemic community in the sense that they still have access to others' evidence. However, in real epistemic situations, sometimes the very manifestation of marginalization is the *exclusion* of certain agents from epistemic communities. This concern would rightly constrain the applicability of my models. But I would like to suggest that this reflects a merit of my approach as well. Recall that Hartsock's argument faces the following problem of intersectionality: if marginalized groups have epistemic advantages for all human activities, then

[41]This is because all agents always get information about both actions in some form.

should we go to the most marginalized group in the world to seek the best knowledge? This problem becomes more puzzling given Narayan's observation that "oppression is often partly constituted by the oppressed being denied access to education and hence to the means of theory production" (Narayan, 1988, 36). My models resist this slippery slope by focusing on situations where marginalized agents are participating members of the epistemic communities, in that they have access to dominant agents' evidence. In fact, when dominant agents refuse to share their evidence with the marginalized group, the two subgroups effectively function as isolated communities. In this case, the marginalized group, often also in the minority, may learn worse as a result.[42] Moreover, the variation I presented in §1.3 with homophilic networks, where p_{outgroup} is small, can be reinterpreted to model a version of epistemic exclusion. Here, marginalized agents have very sparse access to dominant groups' evidence, but their informational access to the dominant group is not completely cut off. This is akin to the all-to-familiar situation where very few marginalized knowers are invited to participate in dominant epistemic spaces, but their testimony is still ignored (see, e.g., Settles et al. (2020)). My results from §1.3 suggest that marginalized agents in this situation still learn the true belief more frequently.

Finally, I will preempt a tempting but misguided response to my modeling results. One might suggest that, since the marginalized group ends up with epistemic advantages, we should now start to ignore or devalue testimony arising from some members of our community, as long as we eventually listen to what they say. This response is misguided for two reasons. First, one-sided testimonial ignoration and devaluation is a textbook case of testimonial injustice according to Fricker (2007). It is unjust because the audience is not giving enough credit to the speaker as they justly deserve. Ignoring or devaluing testimony is committing injustice.

[42] To be sure, non-participating marginalized agents could have epistemic advantages in other aspects than the shared "reality" for all agents; Narayan (1988) offers a few examples of these. Moreover, when dominant agents refuse to share their evidence with marginalized agents, but still updates on evidence shared by the marginalized group—an epistemic exploitation scenario that may underlie some real cases—, my base model can be reinterpreted to account for this situation too, with the dominant and marginalized groups, and thus their epistemic (dis-)advantages, exchanged.

Second, one-sided testimonial ignoration is epistemically detrimental to the entire community as it learns the true belief less frequently and more slowly. Moreover, in the case of one-sided testimonial devaluation, where the community *might* learn better than the community with perfect testimonial reciprocity, my results show that the marginalized group can always learn better by refusing to share their evidence with the dominant group.[43] When marginalized agents are in a situation where their evidence is constantly devalued, they would have little incentive to continue sharing their epistemically better informed evidence with dominant agents. Consequently, the entire community would be epistemically worse off, since the learning situation reverts to the base model. Rather than prompting individuals to ignore or devalue certain community members' evidence, I hope my modeling results would motivate individuals and communities to identify cases of preexisting failure of testimonial reciprocity, to give epistemic credit where it is overdue, and to recognize the epistemic advantages that marginalized agents may already hold.

Besides offering one possible way in which the inversion thesis could arise by casting it as a consequence of the unidirectional failure of testimonial reciprocity, my modeling results complicate the understanding and applications of certain network effects. Before closing the paper, I will briefly discuss my contribution to network epistemology and note directions of future work.

As previously mentioned, the Zollman effect is usually understood as a claim that "a sparser network structure can benefit an epistemic community" (Rosenstock et al., 2017). Zollman (2007, 2010) finds that the more connected the networks, the less frequently but faster the community learns the true belief. However, Rosenstock et al. (2017) tests Zollman (2007, 2010)'s models using an expanded range of parameter values, and finds that the Zollman effect is not robust for a considerable portion of the parameter space. It is worth noting that Rosenstock et al. (2017) does not find a reversal of the Zollman effect, i.e.

[43]My base model applies to both when dominant agents do not update and when marginalized agents refuse to share.

more sparsely connected communities never learn the truth less frequently than the more connected counterparts.

The results of my base model further complicate our understanding of the Zollman effect. The community in my model learns the true belief less frequently *and* more slowly than the community with perfect testimonial reciprocity. This shows that for the benefits of the Zollman effect to obtain, the sparse structure cannot manifest in a cutoff of information channels for a subgroup. Otherwise, the Zollman effect may be reversed. For the marginalized group in my base model, on the surface level its members gain epistemic benefits because they do not lose connections, so it seems to counter the spirit of the Zollman effect. However, the marginalized group gains epistemic benefits precisely because they get information from another group that is disconnected. In a sense, the reasoning behind the Zollman effect explains the situation here: the marginalized group benefits from the disconnectedness of the dominant group, and as a result, its members learn the truth more frequently.

This also constrains how the Zollman effect could be applied to real epistemic communities. If, say, a group of scientists decides to interpret the Zollman effect as suggesting that they should stop reading papers from others in the scientific community, but they continue to publish and post to the arXiv, then my modeling results show that as long as the authors of the papers that they ignore continue to read their papers, members of this group may learn worse according to all measures behaving that way.

Earlier I mentioned another interpretation of my base model, which treats the loss of testimonial reciprocity not as the audience refusing to update, but as the speaker refusing to share. Under this interpretation, the audience in general no longer commits testimonial injustice. As it turns out, this alternative interpretation has fruitful applications in social epistemology. For one, it provides another instance of the Independence Thesis, which roughly states that the prescriptions for individual and group decision-making can come apart (Mayo-Wilson et al., 2011). Indeed, a subgroup may learn the truth more frequently by refusing to share their evidence with outgroup members; but the entire community suffers as a result. Industry scientists who have proprietary knowledge but still have access to information from academia can be modeled this way. This interpretation also applies to situations where the dominant group makes their evidence inaccessible to marginalized groups, but exploits evidence from marginalized perspectives. In both cases, the industry scientists and the dominant group would have epistemic advantages. Moreover, applied to scientific communities, this interpretation sheds light on the recent debate on whether the communist norm, which prescribes that scientists share their findings as widely as possible, is an additional contract that scientists should sign (Strevens, 2017; Heesen, 2017). My results would suggest that by following the communist norm, scientists may not learn the true belief as frequently as theoretically possible, but they avoid the epistemic pitfall when no one shares. I explore this interpretation in a follow-up paper (Wu, 2022b).

Chapter 2

Withholding Knowledge

2.1 Introduction

The communist norm in science mandates that scientists should share their work as widely as possible (Merton, 1973). It is a core norm in science, but not a given. While academic scientists largely adhere to the communist norm, industrial scientists do not typically share proprietary findings.[1] There seem to be community-level benefits for sharing one's work. For instance, scientific discoveries could be made faster, so the public as well as the rest of the scientific community can reap the benefits earlier (Strevens, 2017; Heesen, 2017). But do *individual* scientists, or *subgroups* of scientists, have an incentive to share?

Previous work in the philosophy of science answers this question by appealing to credit incentives by which scientists may be instrumentally motivated. Strevens (2017) appeals to the priority rule, which stipulates that credit—a proxy for recognition for one's scientific work—will be allocated only to the scientist who first makes a discovery. He argues that because of the priority rule, individual scientists would have credit incentives to withhold

[1] In fact, publicly-funded scientists in the US *must* share their research by 2025 (Brainard and Kaiser, 2022). Many European funders have similar requirements.

their evidence from others during the process of inquiry. In so doing, they increase their own chances of making a credit-worthy finding first.[2] However, Heesen (2017) uses a multi-stage game theoretic model to argue that in many cases, it is still rational for credit-maximizing scientists to share, because scientists can publish and claim credit for intermediate results.

Credit is only one of many incentives facing scientists. They are also, or at least nominally, motivated to find the truth (see, e.g., Bright (2017); Zollman (2018)). Curiously, this epistemic incentive underlying scientific sharing has been relatively under-explored. The received position seems to be that a scientist's decision to share or not during the learning process does not impact the likelihood of them eventually making a discovery. Not only do Bright and Heesen (2023) recently explicitly claim that the communist norm is non-epistemic just in this sense, but both Strevens (2017) and Heesen (2017)'s models also implicitly assume that the probability that a scientist will eventually make a discovery remains unchanged regardless of whether they share during the learning process. It is unclear what grounds this position, especially given the growing literature from network epistemology that shows how our social connections and evidence-sharing dynamics can significantly shape knowledge production (Zollman, 2010; O'Connor and Weatherall, 2018; Fazelpour and Steel, 2022; Wu, 2022a).

In this paper, I focus on the epistemic incentives underlying scientific sharing. Specifically, I ask, if scientists are purely motivated by the truth or epistemic significance of their own findings, are they incentivized to share evidence? I investigate this question by simulating models of an epistemic community with two subgroups, one capable of withholding evidence from out-group members, and one adhering to the communist norm by sharing evidence.[3] I find that the subgroup that withholds ends up with several epistemic advantages, both compared to the subgroup that shares and to an entirely "communist" community where

[2] Partha and David (1994) make a similar argument more informally.

[3] Note that my models have two *subgroups*, whereas previous models (e.g. Strevens (2017); Heesen (2017)) typically have two *agents*. But this is not a major departure—results presented here still hold when there are two agents, one withholding and one sharing.

everyone shares regardless of subgroup membership. This suggests that a truth-seeking scientist may be incentivized to withhold their evidence.[4]

I construct two models from different paradigms for this purpose—a generalized multi-armed bandit model with a network structure based on Zollman (2010) and an NK landscape model with a network structure based on Lazer and Friedman (2007). These two models represent two different types of scientific inquiry. In one model agents figure out which of the two probabilistic epistemic options are better, representing, e.g., clinical doctors finding out the efficacy of two different drugs by conducting trials. In the other, agents search in a vast epistemic landscape with multiple "peaks," representing, e.g., researchers adopting different approaches to solve a problem. Together these models represent a wide array of possible scientific problems.

Results from the bandit model show that members of the "withhold" group reach the true belief more frequently and faster, and select the epistemically better action more frequently during the learning process, as compared to the "share" group. In the NK landscape model, members of the share group end up with worse solutions than both the withhold group and a generally communist community. Moreover, the withhold group gains epistemic advantage even compared to the communist community in terms of arriving at the true belief more frequently (in the bandit model) and ending up with epistemically better solutions (in the NK landscape model in most cases).[5]

These results are troubling, especially given that proprietary industrial scientists (see, e.g. DeAngelis (2003); Michaels (2008); McGarity and Wagner (2010)) and scientists working on classified research (see, e.g. Galison (2004)) routinely withhold their evidence from others. Academic scientists, on the other hand, largely conform to the communist norm (Louis et al.,

[4]Note that in my models the scientists are motivated by *themselves* finding the truth. So the epistemic benefits may translate into practical benefits for the scientists too—such as recognition of or financial gains from their work.

[5]These results extend and provide robustness checks for recent results from a simpler multi-armed bandit model, built for another context (Wu, 2022a).

2002; Macfarlane and Cheng, 2008). My results suggest that even if these scientists are not as instrumentally motivated by credit as academic scientists, they may still have epistemic incentives to withhold evidence. Moreover, while not explicitly modeled in this paper, industrial scientists likely have further financial incentives to withhold evidence. Consequently, academic scientists suffer epistemically.

I will then observe that, based on simulation results, this share-withhold dynamics gives rise to what one might call an Epistemic Weak Prisoner's Dilemma. Each subgroup receives the highest epistemic payoff when they withhold while the other subgroup shares. Their payoffs are the worst when the other group withholds, and their payoff is intermediate when both subgroups share. I will discuss features of this game and explore strategies that may shift communities into mutual sharing.

This paper is organized as follows. In §2.2, I introduce the generalized bandit model and present my simulation results. In §2.3, I introduce the NK landscape model and discuss my results. I also note common features of the two models that lead to robust qualitative findings and explain why certain results are not as robust. In §2.4, I show that a weak prisoner's dilemma may represent the epistemic dynamics of scientific sharing. §2.5 concludes. This paper is supplemented by two technical appendices. Appendix A provides definitions of the end states in the generalized bandit model. Appendix B describes how the solution space of the NK landscape model is generated.

2.2 The Generalized Bandit Model

In this section, I construct a model where a group of agents is tasked with a learning problem. There are two subgroups: the withhold group, whose members only share evidence with in-group neighbors and withhold from out-group neighbors, and the share group, whose members share evidence with every neighbor. The goal of this model is to explore the epistemic consequences of this asymmetric evidence sharing dynamics across subgroups.

Let us motivate the model by considering a toy example. Suppose that a group of clinical scientists is tasked to figure out which of the two drugs, A or B, are more effective at treating a disease. Each scientist tests the drugs on patients, receives evidence from their tests, and updates their beliefs from the evidence they and others receive. The scientists have limited resources, so they each only test one drug per round. Moreover, they want to minimize patient suffering, so they always assign the drug that they currently think is better. As the scientists keep getting evidence for the drugs, eventually they should reach stable beliefs about which drug is better.

This example is well modeled by what is called a two-armed bandit problem with a network structure. The name "bandit problem" comes from applying the model to a gambling situation, where a gambler aims at maximizing their profits when playing with a multi-armed "bandit" (or slot) machine. Let us first consider the problem for one agent, before thinking about it in a group setting. Every round, the agent selects between two options, A and B, tests their choice a fixed number of times n, and gets evidence about how many times their tests succeed. Each option has a fixed probability of success. Unbeknownst to the agent, I set the success rate of A, P_A, to be .5, and the success rate of B, P_B, to be lower than .5. This means that A is objectively the better choice. The agent, however, is uncertain about the success rate of either option, and their credence for each is represented by a beta distribution with two parameters, α and β.[6] Details about the beta distribution do not matter for our purpose. What matters is that in the context of Bayesian learning, we can

[6] A beta distribution is a function of the following sort:

$$f(x) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{B(\alpha, \beta)}$$

where $B(\alpha, \beta) = \int_0^1 u^{\alpha-1}(1-u)^{\beta-1}du$ and $\alpha, \beta > 0$.

interpret α as the agent's estimate of the "successes" of the arm, and β as their estimate of the "failures." For instance, suppose the agent starts with $\alpha_A = 1$ and $\beta_A = 3$ for option A, tests this option 10 times, and receives 8 successes and 2 failures, then, after Bayesian updating, their posterior will still be a beta distribution, with parameters $\alpha_A = 1 + 8 = 9$ and $\beta_A = 3 + 2 = 5$. If this arm continues to be tested, over time, $\frac{\alpha_A}{\alpha_A+\beta_A}$ will approach the true success rate of A. Every round, the agent selects the option that they currently think has a better chance at succeeding,[7] tests it n times, and updates on their evidence in the (myopically) Bayesian way described above. After sufficient rounds, the agent's belief about which option is better stabilizes.

Now suppose a group of agents, connected via a network structure, solves the same bandit problem together. Each agent starts with four randomly assigned parameters $\alpha_A, \beta_A, \alpha_B, \beta_B$,[8] representing their own credence for the two options. The network structure determines who they are connected to, or who their "neighbors" are. Typically, as in the case of Zollman (2010), every round, after collecting their evidence, agents share the evidence with all their neighbors and update on their neighbors' evidence in the Bayesian way.

In my model, however, not everyone shares evidence with all their neighbors. I divide the community into two subgroups: the withhold group, whose members share their evidence with in-group neighbors but withhold evidence from out-group neighbors, and the share group, whose members share their evidence with all their neighbors. All agents then update their beliefs based on all the evidence they receive (including their own evidence) in the Bayesian way. This modification creates an asymmetry along group membership in the evidence-sharing dynamics. To go back to the previous example of clinical trials of drugs, my model may represent a situation where a community of scientists is trying to figure out the efficacy of two different drugs, but within this community, there is a subgroup of

[7] If the mean of their beta distribution for A is higher than the mean of their beta distribution for B, i.e. $\frac{\alpha_A}{\alpha_A+\beta_A} > \frac{\alpha_B}{\alpha_B+\beta_B}$, then they choose A; they choose B otherwise.

[8] These are real numbers between 0 and 4 (exclusive).

industry-affiliated scientists who withhold their evidence from the outside, even though the rest of the community continues to share their evidence widely.[9]

I run each simulation for $10,000$ rounds. At the end of the simulation, there are typically three stable end states: (1) community convergence to the true belief, where everyone thinks that A is better; (2) community convergence to the false belief, where everyone thinks that B is better; and (3) polarization, where the withhold group thinks that A is better, but the share group thinks that B is better. An overwhelming majority of simulations end and stay in these states.[10] Detailed definitions of the end states are available in Appendix A.

Polarization is only possible in this model because one group withholds evidence from another. In this state, the withhold group succeeds in learning, but the share group fails in learning. This state is stable because here, the share group has already settled on B and they do not receive evidence about A anymore, even though the withhold group continues to test A. It is important to note that a polarized state in the other direction, i.e. the share group succeeds but the withhold group fails, is not stable. This is because the withhold group would continue to receive evidence for action A from the share group, and since A is in fact better, the evidence they receive would prompt them to eventually switch.

It is important to note that this model is a generalization of a previous model (Wu, 2022a) involving a simpler bandit problem. There are two subgroups in Wu (2022a)'s model, one ignoring evidence from out-group neighbors, and one updating on evidence shared by all their neighbors. The asymmetric evidence *updating* dynamics in Wu (2022a)'s model is structurally equivalent to the asymmetric evidence *sharing* dynamics in this model, with the ignored group in Wu (2022a) and the withhold group in this model occupying structurally the same position. This model uses a more complex bandit problem where agents start

[9]Note that industrial scientists mutually share with each other. I briefly explore what would happen if there are multiple mutually withholding "industry" groups in §2.4.

[10]Simulations with learning problems that are "hard" (e.g. where P_A and P_B are close or when n is small) are more likely to not finish in these end states than problems that are "easy," though for each set of parameter values, more than 97% of simulations reach one of these end states.

with not knowing the success rates of either epistemic option and there are infinitely many possible success rates for both options, whereas in Wu (2022a)'s model, agents start with knowing the success rate of one option and there are only two possible success rates for the other. As we will see, my results on the epistemic advantage of the withhold group provide a robustness check in a more generalized setting for the results in Wu (2022a).[11]

2.2.1 Results and Discussions

For each set of parameter values, I run the model for $10,000$ simulations. The parameters tested include:

- Total number of agents (k): $3, 6, 12, 18$;
- Proportion of the withhold group in the population (d): $\frac{1}{6}, \frac{1}{3}, \frac{1}{2}, \frac{2}{3}$;
- Number of tests per round (n): $1, 10, 100, 1000$;
- Success rate of B (P_B): $.4, .45, .49$;
- Network Structure: complete.[12]

In addition, I run a "communist" model, where everyone shares with everyone, for all the parameter values. This is equivalent to models from Zollman (2010), and it provides a comparison to my results. Unless otherwise noted, results reported here are robust across all parameter values.

I find first that the withhold group succeeds in learning more frequently than the share group (Figure 2.1). I measure the frequency of successful learning for a subgroup by calculating the proportion of simulations that the subgroup succeeds in learning out of all simulations where the community ends in one of the three end states. In this model, the withhold group

[11] Wu (2022a) uses her results on the epistemic advantage of the ignored group to justify a standpoint epistemology thesis that marginalized group can sometimes have better knowledge.

[12] A network structure is complete when every agent is connected to every other agent. I also consider the directed Erdős-Rényi random networks of size 18 for robustness checks.

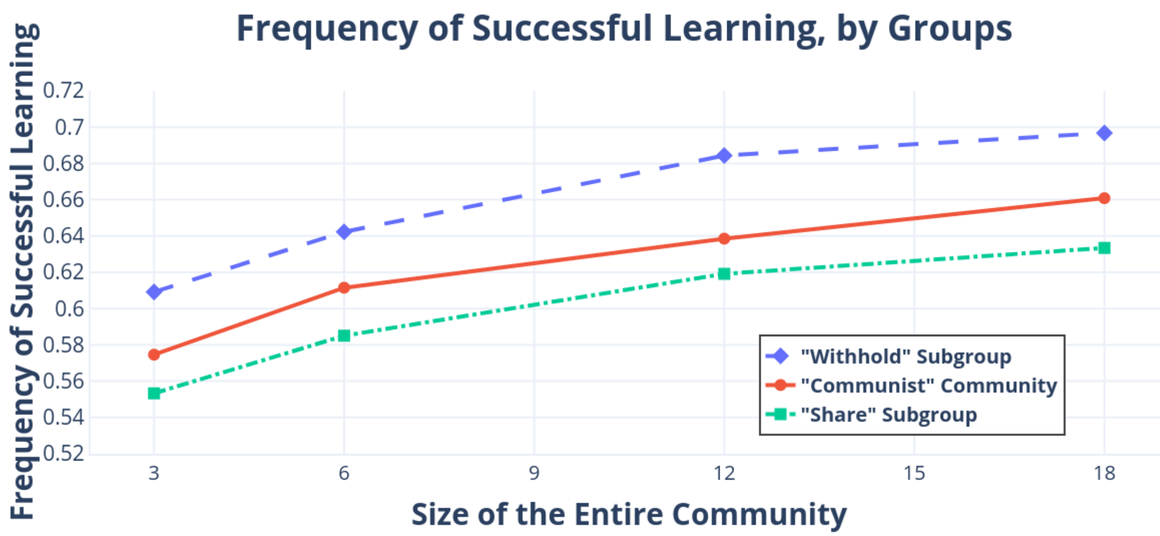

Figure 2.1: $P_B = .49$, $n = 10$, $d = \frac{1}{3}$, complete network structure, $10,000$ simulations.

succeeds in learning more frequently because polarization is an end state. In this state, the withhold group succeeds in learning but the share group fails.

Moreover, and perhaps surprisingly, the withhold group succeeds in learning even more frequently than a communist community where everyone shares with everyone (Figure 2.1).[13] The reason is that in simulations that start with initially promising evidence for B and unpromising evidence for A,[14] members of the communist community may settle on B as the better option. But in my model, since the share group is effectively an isolated communist community of a smaller size, it takes them longer to reach a stable state. This in turn means that evidence from both actions circulates in the network for longer, and the community, including the withhold group, may revert to choosing A. This explanation is related to what is often called the Zollman effect (Zollman, 2007), which states that a more sparsely-connected network more frequently succeeds in learning. Here, the sparsity of the network slows down learning, so the community spends more time exploring different options before settling down. This period of time where members of the community test different

[13]This result holds for more than 92.2% of parameter combinations. It is more likely to fail when the learning is "easy," i.e. when P_B is low and n is high. When $P_B = .49$, this result holds 100%. This pattern of robustness levels is consistent with Rosenstock et al. (2017)'s findings on the Zollman effect.

[14]These simulation runs are possible because of the probabilistic nature of epistemic options in the model.

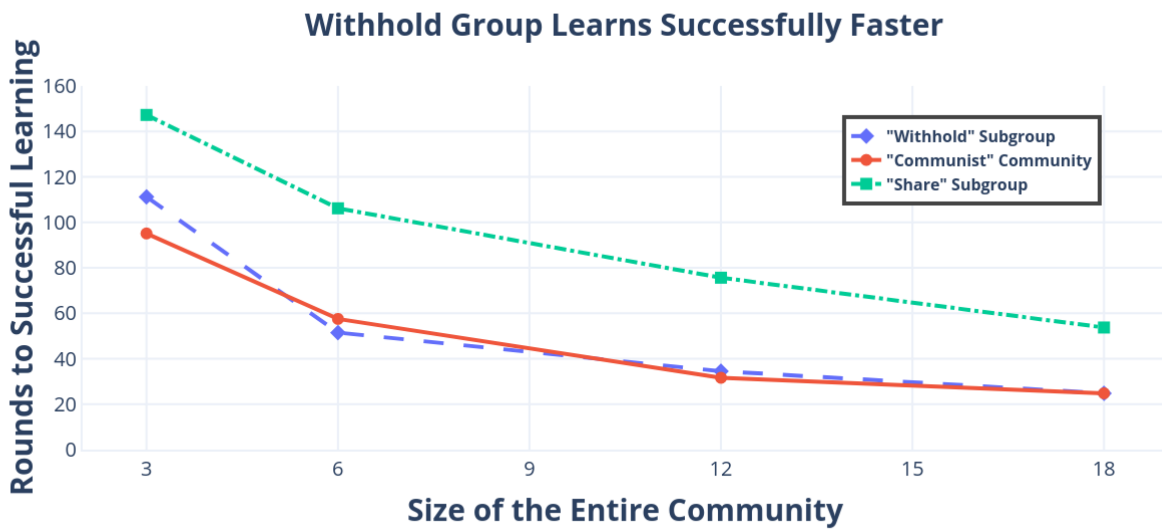

Figure 2.2: $P_B = .45$, $n = 1$, $d = \frac{2}{3}$, complete network structure, $10,000$ simulations.

options is called a period of transient diversity (Zollman, 2010; Wu and O'Connor, 2023). When transient diversity lasts longer, it brings epistemic benefits to some or all community members.

It may be instructive to think about the withhold group's epistemic advantage through a trade-off between exploration and exploitation inherent in the bandit model. Due to the probabilistic nature of both arms of the bandit, agents typically have to test each arm sufficiently many times to form a reliable estimate of its success rate. This gives rise to a dilemma for individual agents in the model—do they keep exploiting the option that they currently think is best, or do they explore the other option that seems inferior? By *not* sharing their evidence with out-group members, the withhold group in a sense takes the best of both worlds. They continue to exploit the option they currently think is better, while benefiting from the exploration of the share group.[15] This is related to the "free rider problem" (Kummerfeld and Zollman, 2015), which describes situations where it is rational for individual agents to leave the exploring to others in the community. In my model, the

[15]Note that the share group explores more not because of differences in their behavioral rules, but because they may be testing different options from the withhold group since they do not have access to the withhold group's evidence.

free rider problem is even more insidious, since here the free riders, i.e. the withhold group, reap even more epistemic benefits than the rest of the community.

Furthermore, the withhold group takes a shorter time to succeed in learning than the share group in simulations that end in community convergence to the true belief.[16] The withhold group's speed to successful learning is comparable to the communist community's (Figure 2.2). The time a subgroup takes to succeed in learning is measured as follows. For every agent in the subgroup, I record the first round after which the agent's credence satisfies the success conditions (see Appendix A for details). I then take the average out of all agents in the subgroup and all simulations that end in community convergence to the true belief. Here, the withhold group succeeds in learning faster because they update on more pieces of evidence than the share group. However, in simulations that end in polarization, the withhold group takes longer to learn the truth. This is because in these cases, the share group converges to believing that B is superior early on, but the withhold group slowly tests A until they have sufficient evidence for its superiority. This means that the withhold group consistently has epistemic advantage over the share group in all situations—when they both succeed in learning, the withhold group succeeds more quickly; when the share group fails in learning, the withhold group still may succeed, albeit slowly.[17]

The withhold group on average selects the epistemically better action more frequently during the learning process than the share group (Figure 2.3). For this, I count the number of rounds that an agent in the said subgroup selects action A, then divide it by all agents in the subgroup and all simulations.[18] As the withhold group gets larger, it selects the epistemically better action more frequently, and the share group selects it less frequently.

[16]This result holds for more than 92.8% of parameter combinations. It is more likely to fail when the learning is "hard," i.e. when P_B is closer to P_A. In this situation, it is difficult to distinguish P_A and P_B, so it is more likely to have a small number of simulations that take a long time to finish, thus skewing the average rounds to successful learning.

[17]Contrast this with the Zollman (2007, 2010)'s finding that a less connected community succeeds in learning more frequently but less quickly. The withhold group in my model truly takes the best of both worlds.

[18]Including simulations that do not end in the three end states.

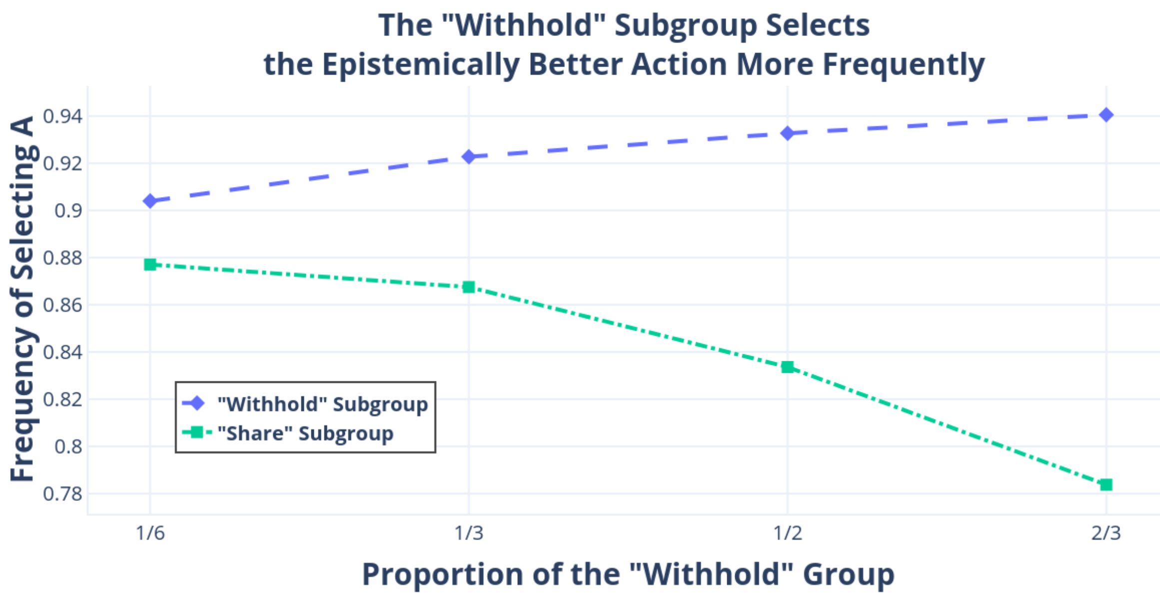

Figure 2.3: $P_B = .4$, $n = 1$, $k = 6$, complete network structure, $10,000$ simulations.

The other epistemic advantages for the withhold group similarly depend on its proportion in the community. As the proportion of the withhold group grows, the withhold group's epistemic advantage tends to increase in terms of the frequency and speed of successful learning.

To further test for robustness, I simulate this model and Zollman (2010) with directed Erdős-Rényi random networks and find qualitatively similar results.[19]

[19]Such a network is generated at the start of each simulation in the following way. First, every agent is linked to themselves. Then, for every agent X and every other agent Y, the probability that X is connected to Y is a fixed number b ($0 < b < 1$). I ran the simulations for communities of size 18 for all other parameters listed, with probability of connection $b = .6, .7, .8, .9$, for $10,000$ simulations each. Note that this network is directed because X could be connected to Y without Y be connected to X. I expect a model with undirected network structures to produce similar results. Furthermore, I only consider connected* networks in this case, defined in the following way. A network is connected* if (1) there exists a path from any sharing agent to any arbitrary agent in the network; and (2) there exists a path from any withholding agent to any arbitrary withholding agent. Moreover, there exists a path from agent Y to agent Z if there are agents $A_0, ..., A_n$ such that $A_0 = Y$, $A_n = Z$, and A_i shares evidence with A_{i+1}, for $0 \leq i < n$. This definition is analogous to the definition of connected* in Wu (2022a).

2.3 The NK Landscape Model

Next, I consider an influential epistemic landscape model—the NK landscape model.[20] My model again consists of two subgroups—a withhold group and a share group—searching in the same epistemic landscape for solutions to a problem. In what follows, I first introduce the general idea behind an epistemic landscape model. I then introduce the solution space of the NK landscape model, the network structure, and agents' behavioral rules in my model. After that, I compare and contrast this model with the generalized bandit model in §2.2, before discussing the results.

We can think of an epistemic landscape as containing a large number of research approaches to a particular topic of inquiry. Each research approach is a point on the landscape and has a score associated with it, representing its "epistemic significance." Following Alexander et al. (2015, 426), I interpret research approaches in a broad way—they have a number of components, including research questions, methods, skills, instruments, background assumptions and theories, etc. We can then introduce a group of agents searching the landscape, i.e. choosing research approaches and solving problems. This models important aspects of scientific problem solving—scientists constantly communicate with others in their community and decide whether they should stick with the research approach they currently have or try new ones, either by exploring on their own or adopting an approach from the community.

The NK landscape is a sophisticated multi-dimensional landscape with multiple "peaks."[21] The solution space is N-dimensional, with binary strings (0s and 1s) of length N as its points. At the start of each simulation, an algorithm with parameter K ($0 < K < N - 1$) is used to randomly assign scores between 0 and 1 to each string. A full description of the algorithm

[20] While the NK landscape model is influential in other fields, it is under-explored in philosophy (with the exception of Alexander et al. (2015)).

[21] C.f. lower-dimensional epistemic landscape models, e.g. Weisberg and Muldoon (2009); Hong and Page (2004). The NK landscape model was originally developed in biology to model "synergies" among genes, see Kauffman and Levin (1987); Kauffman and Weinberger (1989).

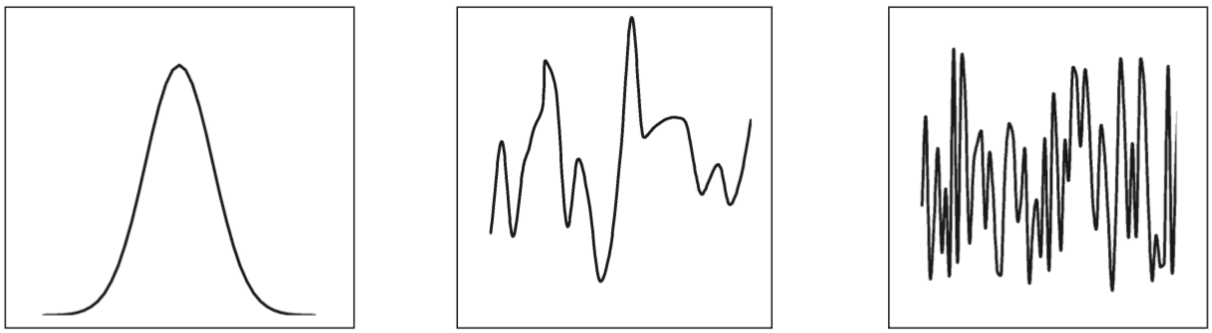

A: A simple problem space, similar to K = 0.

B: A complex rugged space with local maxima and minima, similar to 0 < K < N – 1.

C: A chaotic space in which the value of every point is independent of adjacent points, similar to K = N – 1.

Figure 2.4: Stylized Representation of the Solution Space (Lazer and Friedman, 2007).

is available in Appendix B. Roughly speaking, the parameter K determines how "rugged" the solution space is and how correlated "nearby" scores are. When $K = 0$, the solution space is smooth with one single peak. When $K = N - 1$, the solution space becomes totally chaotic, where the score of every point is totally independent of adjacent points. When $0 < K < N - 1$, the landscape is rugged, with multiple local optima, and with some correlation between adjacent solutions. In this case, some peaks in the landscape may only be reached via some solutions but not others. As K grows, the landscape becomes increasingly harder for agents to search. Though it is impossible to sketch a higher dimensional space, Figure 2.4 provides a stylized representation of the solution space as we vary K. [22] In this paper, we focus on the $0 < K < N - 1$ regime. In the context of scientific problem solving, we can think of these binary strings as different research approaches to a topic of scientific inquiry, and the N digits of each string as possible components of these approaches, e.g. research questions, lab instruments, standards of induction, etc. The epistemic significance of an approach, then, depends on how well its components work together.

[22]Note that though the figure is two dimensional, it does not represent the problem where $N = 2$. Rather, it provides a stylized representation of the complexity of the problem as K increases, for an unspecified $N > 2$.

I now turn to the network structure and behavioral rules of my model. The model I construct is a variation of Lazer and Friedman (2007)'s, with the addition of asymmetric evidence sharing.

At the start of each simulation, we have 100 agents, connected to each other via a directed Erdős-Rényi random network with a probability of connection b, as generated by the algorithm in Footnote 19.[23] This means that each agent on average has $100 \cdot b$ neighbors. I divide the agents into two subgroups, the withhold group and the share group. Every agent is randomly assigned a starting solution, in the form of a binary string of length N.

Every V rounds, the epistemic community goes through social learning.[24] When this happens, each member of the withhold group shares their solution and its score with their in-group neighbors, and each member of the share group shares their solution and its score with all their neighbors. Then, every agent in the community looks at all solutions shared with them, and chooses the solution with the highest score to copy, if it is higher than their own.[25] If their own score is higher than all those shared with them, then the agent goes through a local search. This means that they randomly choose a bit in their binary string to alter (1 to 0 or 0 to 1), and, if the altered string has a higher score, they switch to that solution. Otherwise, they maintain their current location. In other rounds (rounds indivisible by V), each agent conducts a local search to try to improve their score. In scientific problem solving, local searches represent situations where each scientist or lab tries to improve their own research approach by making changes to one component of their approach. Note that when $V = 1$, agents undergo social learning every round. I simulate the model for 200

[23]I use directed networks to allow for situations where an agent is aware of a solution of another agent, but not vice versa, even if they are in the same group. The results I report does not depend on the directedness of the network structure.

[24]I start the simulation at round 1, and require that the community go through social learning when the current round is divisible by V. This ensures that agents go through a few rounds of local search before starting social learning, unless $V = 1$.

[25]If there are multiple solutions with the highest score, they randomly select one to copy.

rounds. As we will see, in most simulations, the community reaches stable behaviors much more quickly.

In the NK landscape model, we have a community of agents searching in a vast epistemic landscape with multiple "peaks," so depending on how agents search, they may fail to ever discover the global optimum. There are two characteristics that distinguish the NK landscape model from the bandit model. First, instead of the two options considered in the bandit model, in the NK landscape model we have a myriad of solutions.[26] Second, in the NK landscape model, each solution's epistemic significance is readily known to an agent who chooses it, whereas in the bandit model both options are probabilistic, making their success rates harder to estimate. The exploration and exploitation trade-off for agents in this model is thus the following: do they exploit the epistemic significance of their current solution, or do they keep exploring the landscape in the hope of finding a better solution? Many authors (e.g. March (1991); Lazer and Friedman (2007); Fang et al. (2010)) find results similar to the Zollman effect in the NK landscape model, and these results are empirically confirmed to some extent (Mason and Watts, 2012; Derex et al., 2018). That is, they find that in the NK landscape model too, a less connected community may end up with better solutions. This is because in more connected networks, agents may quickly settle onto a local optimum, thus failing to explore better alternatives elsewhere. In less connected networks, agents retain a diverse range of solutions for longer, and thus may ultimately discover better solutions.

Simulating the asymmetric evidence-sharing dynamics in the NK landscape model, therefore, serves as an important robustness check. Moreover, we may think that many of the real scientific problems involve *a wide range of probabilistic options* (Wu and O'Connor, 2023), not just a wide range of certain options, as in the case of the NK landscape model, or a limited range of probabilistic options, as in the case of a two-armed bandit model. A replication of my qualitative results here may thus increase our confidence that similar results would hold

[26]In the $N = 20$ case considered in this paper, there are $1,048,576$ solutions.

in a model that combines features of the bandit model and NK landscape model.[27] As I show shortly, the qualitative results exhibit similar patterns across the two modeling paradigms, though some results are more robust than others.

Before presenting the results, it is important to note another salience difference between the NK landscape model and the bandit model. In the NK landscape model, some high-performing solutions are only accessible if agents explore from certain other solutions via local search, so agents concentrating on one patch of the landscape may not ever discover promising solutions in other regions. In the bandit model, however, different options are completely independent from each other. This fact is relevant in explaining some results below.

2.3.1 Results and Discussions

I run the model $1,000$ times for each parameter combination:

- N: 20;[28]
- K: $5, 10, 15$;[29]
- V: $1, 3, 5$;
- Proportion of the withhold group (d): $.2, .4, .6, .8$;
- Probability of connection in directed Erdős-Rényi random networks (b): $.4, .6, .8, 1$.[30]

[27] Such models may take the form of a bandit model with sufficiently many arms, or an epistemic landscape model where the score of each solution is probabilistic. I leave the detailed modeling work to further research.

[28] N here is different from n in §2.2.1. n is the number of tests per round in the bandit model.

[29] K here is different from k in §2.2.1. k is the number of agents in the bandit model.

[30] When $p = 1$, the network structure is complete.

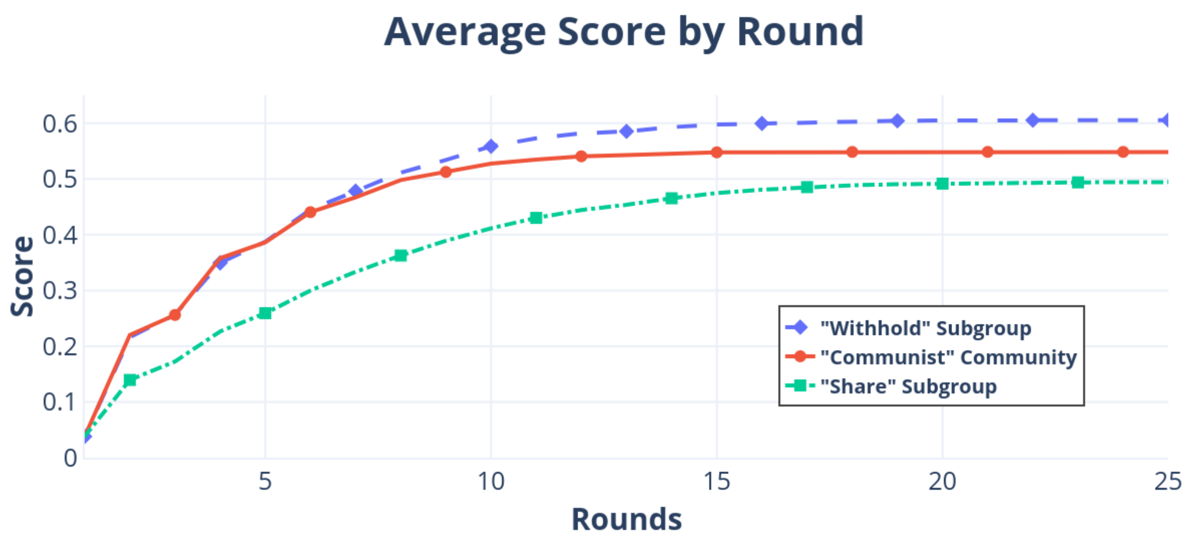

Figure 2.5: $N = 20$, $K = 5$, $V = 1$, $d = .8$, $b = .6$. 1,000 simulation runs. All three groups reach stable scores after 25 rounds.

In addition, for each parameter combination, I run the model of a communist community where everyone shares with all their neighbors.[31] As noted before, each community has 100 agents. The results I present are robust across parameter values unless otherwise noted.

The first result is that the withhold group ends up having better scores on average than the share group. This result always holds because the withhold group has the option of switching to the better solutions from the share group, but not vice versa. Because of this, the withhold group always ends up with at least as good a score as the share group at the end of each simulation.

For almost all ($> 97\%$) parameter combinations, the share group ends up having worse scores on average than the communist community. This is because the share group is smaller in size than the communist community and thus suffers from two epistemically detrimental consequences. First, the initial solutions present in the share group are not as diverse as in the communist community, so high scoring solutions may be "farther away" from agents in

[31] This is equivalent to Lazer and Friedman (2007). For ease of comparison, the two communities always search in the same solution space, and have exactly the same random network structure and initial solutions, all of them randomly generated for each simulation.

the share group. Second, agents in the share group on average have fewer neighbors than agents in the communist community. So when an agent in the share group and an agent in the communist community both have the best solution among their neighbors, it takes longer for the neighbors of the former agent to conduct a thorough local search for better solutions. In the face of unpromising results in early local search, then, these agents may give up and switch to another solution from elsewhere in the community. Figure 2.5 shows the average scores per round for different groups.

For most sets of parameter values ($> 56\%$) the withhold group ends up with better scores on average than the communist community. When this happens, the withhold group benefits from the extra exploration done by the share group, while the communist community converges to a suboptimal solution too quickly. This is a similar mechanism for epistemic advantage as that identified in the bandit model, but the withhold group's advantage here is not as robust.

To see why, let us consider a simplified case. Suppose that we have four solutions, A, B, C, and D. Suppose further that $A < B < C < D$ in epistemic significance, but D is only available when we explore from B, and C is only available when we explore from A. Now suppose that A is the current best solution among the share group, and B is the current best solution among the withhold group. Then in cases where the share group discovers C before the withhold group discovers D, the withhold group would switch to C prematurely, without ever discovering D. Curiously, when a communist community encounters this case, they would not make the same mistake, since agents would choose B over A first, and the community would, after sufficient local searches, settle on D. Qualitatively similar cases can arise in the NK landscape model, and in these situations, the withhold group learns worse than the communist community. The same situation does not apply to the two-armed bandit problems.

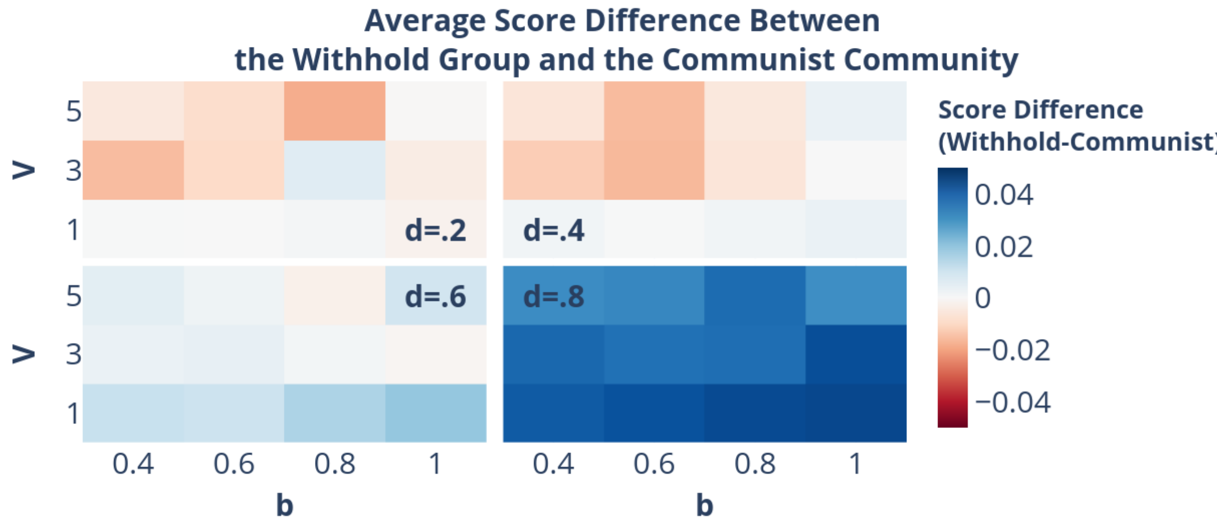

Figure 2.6: $N = 20$, $K = 10$, 1,000 simulation runs. Gray-scale version available upon request.

For this reason, when d is small, i.e. when the withhold group is in the minority, the withhold group is more likely to lose its epistemic advantage compared to the communist community. The withhold group does not have enough agents to do as thorough a local search around the most promising solution as the communist community, given the vast number of possible local variations on a solution. They may instead give up on their local search if the more populous share group reaches a seemingly promising solution in the meantime. As d increases, the withhold group becomes more effective at local search, and it gains epistemic advantage.[32]

Moreover, as V increases, i.e. when social learning becomes less frequent, the withhold group tends to have less of an epistemic advantage. To understand this, we observe that when $V = 1$, because of the quick social learning, the communist community is more likely to prematurely settle down to a local optimum, without ever exploring other regions of the landscape. In contrast, in the withhold-share community, diverse solutions are still present even when social learning is fast, because the withhold group does not share. In cases where solutions among the share group are initially unpromising, its members can keep exploring them, and the withhold group benefits if the share group's exploration proves

[32]The withhold group does better than the communist community 100% when $d = .8$.

fruitful. However, when V increases, both communities know more about the landscape terrain before commencing social learning. In this case, the communist community simply has more agents conducting local searches around promising solutions after social learning, making it more likely that it would find something.

Finally, as K increases, the withhold group tends to lose its epistemic advantage over the communist community as well. The reason for this is slightly different from before. When K grows, the landscape becomes increasingly chaotic, and the epistemic communities do not typically learn very well at all.[33] In these "hostile" landscapes, it is again better to have more agents conduct targeted local exploration around one solution, as in the case of the communist community, than dividing up labor such that neither subgroup has enough agents to do thorough local searches, as in the case of the withhold-share community. One might find this result surprising because previously Rosenstock et al. (2017) finds that the Zollman effect is the most robust in hard bandit problems, but here it seems that the opposite holds in the NK landscape problem. This conclusion may be too quick, because here what is holding the withhold-share community back is the insufficient *local* exploration caused by the size of the subgroups, not a lack of global diverse solutions. This nonetheless points to a significant difference between the bandit model and the NK landscape problem.

Figure 2.6 shows the degree of the withhold group's epistemic (dis-)advantage over the communist community as we vary several parameters.

2.4 Epistemic Prisoner's Dilemma

Results from §2.2-§2.3 suggest that a subgroup that withholds evidence can have epistemic advantages over the rest of the epistemic community that shares, and over a communist community in many cases. If scientists are purely motivated by the truth of their own

[33] When $K = 15$, all communities and subgroups' average final scores do not exceed .4.

findings, the withhold-share dynamics gives rise to what one might call an Epistemic Weak Prisoner's Dilemma (EWPD).[34]

The game is set up as follows. We have two groups who are the primary players of the game. Group members internally share their evidence with each other. On top of that, each group has two strategies, sharing evidence with out-group members (Share), and withholding evidence from out-group members (Withhold). Each group's payoff is epistemic, corresponding to the probability that the group converges to the truth in that situation. Specifically, the epistemic payoffs in Tables 2.1 and 2.2 are calculated from the frequencies of successful learning measure from the bandit model in §2.2.[35] By assuming that members of the community are already sorted into different groups or coalitions, this game is an instance of a cooperative game (see §9 of Ross (2021)). Because of this assumption, the game may be interpreted to represent certain groups (e.g. industrial scientists) better than others (e.g. academic scientists), as I will discuss shortly.

Symmetric EWPD		Group 2 (1/2)	
		Share	Withhold
Group 1 (1/2)	Share	.58, .58	.56, .63
	Withhold	.63, .56	.56, .56

Table 2.1: $k = 12$, $n = 1$, $P_B = .49$. $10,000$ simulation runs. Groups 1 and 2 are both 1/2 the size of the community.

Asymmetric EWPD		Group 2 (1/3)	
		Share	Withhold
Group 1 (2/3)	Share	.58, .58	.57, .61
	Withhold	.67, .55	.57, .55

Table 2.2: $k = 12$, $n = 1$, $P_B = .49$. $10,000$ simulation runs. Group 1 is 2/3 the size of the community, and Group 2 is 1/3.

In the EWPD, each group receives the highest payoff when they withhold evidence from the other group while the other group shares with them. Their payoffs are the worst when the

[34]In so doing, I provide new examples of how game theory can be used as a tool for the social epistemologists, see Zollman (2021).

[35]For Share-Share, I calculate the frequency of successful learning for the communist community. For Withhold-Withhold, I calculate the same frequency for an isolated communist community of a smaller size.

other group withholds evidence from them (no matter what they themselves choose to do), and their payoff is intermediate when both groups share.

There are two features of this game worth noting. First, the game is symmetric only when each group is exactly half of the entire community in size (Table 2.1); otherwise, the game is asymmetric (Table 2.2). In the latter case, the group that is the majority in size would gain more when they withhold evidence while the other group shares, and lose less when both subgroups withhold. This may seem counter-intuitive, given that in a public goods game, where players decide whether to contribute to a public pot which will then grow and be distributed to everyone, players have the incentive to be in the very few that withhold. However, in the EWPD, members of each group still share with each other, and this mechanism is absent in a public goods game. Furthermore, social knowledge is unlike public goods. It is not the case that the more people share, the better the group is at learning—the Zollman effect tells us otherwise.

The second feature is that this prisoner's dilemma is *weak*, since the situation where both groups withhold is a weak Nash equilibrium, not a strict one. The situations where one group switches their strategy (top right and bottom left) are two additional weak Nash equilibria of this game, and they are both Pareto improvements from the situation where both groups withhold, since one group would gain and no groups lose.[36] Given this feature, we might imagine a new game where both groups are slightly altruistic, and their payoff is calculated by .9 of their own original epistemic payoff, plus .1 of the other group's original epistemic payoff in the same situation. Then, both situations where one group withholds and the other shares are (strict) Nash equilibria in this game, and Withhold-Withhold is no longer Nash (see Table 2.3). One might call this new game an Epistemic Hawk-Dove. Furthermore, we might imagine another game where only one group is slightly altruistic. The result is an

[36]The EWPD can also be seen as a weak Hawk-Dove, with Withhold-Withhold as an additional weak Nash equilibrium.

Asymmetric Epistemic Hawk-Dove (Table 2.4), and the situation where the altruistic group shares and the other self-interested group withholds is the only strict Nash equilibrium.[37]

Epistemic Hawk-Dove		Group 2 (1/2)	
		Share	Withhold
Group 1	Share	.58, .58	.567, .623
(1/2)	Withhold	.623, .567	.56, .56

Table 2.3: New game constructed from Table 2.1. Both groups are slightly altruistic.

Asymmetric Epistemic Hawk-Dove		Group 2 (1/2)	
		Share	Withhold
Group 1	Share	.58, .58	.567, .63
(1/2)	Withhold	.623, .56	.56, .56

Table 2.4: New game constructed from Table 2.1. Group 1 is slightly altruistic.

One might take this Asymmetric Epistemic Hawk-Dove to explain our current situation where slightly altruistic academic scientists generally share their evidence while self-interested industrial scientists withhold.[38] However, the strength of this explanation depends on whether we can successfully interpret academic scientists and industrial scientists as group agents able to play the two strategies. While the industrial scientists have centralized decision-making power, it is difficult to coordinate decentralized academic scientists to play the Withhold strategy, especially given that they have additional credit incentive to publish.[39] The assumption that academic scientists can be thought of as a group agent in a cooperative game may not be realistic. Instead, our current situation may be closer to one where academic scientists are stuck with playing the Share strategy, and the industrial sci-

[37]The other Nash equilibrium, where the altruistic group withholds and the self interested group shares, is weak.

[38]For those who are skeptical that academic scientists are altruistic, we can instead introduce academic scientists' credit incentive to the EWPD by adding a small payoff to one group in situations that they share, representing the recognition they receive when they publish. The result would be an asymmetric Hawk-Dove as well. Furthermore, we can introduce industrial scientists' financial incentive to the game by adding a small payoff to the second group in situations that they withhold. The two additions together would produce a game where the situation where the first group shares and the second group withholds is the only Nash equilibrium, strict or weak. Thanks to Hannah Rubin for thinking through this point with me.

[39]Thanks to Cailin O'Connor for raising this point. The case for classified scientists is a bit complicated. We might think that they are the opposite of academic scientists, since they have centralized decision-making power, but are stuck with the Withhold strategy. Analyzing how this complicates sharing dynamics is beyond the scope of this paper.

entists, having centralized decision-making power, choose to Withhold to maximize their epistemic payoffs.

Given that industrial scientists can be interpreted as group agents, are there ways in which industrial scientists may come to share from a game theoretic perspective? There is a large literature on how cooperative behaviors could evolve in a prisoner's dilemma (See, e.g., Chapter 6 of O'Connor (2020) for an overview). The key is that when players' actions become correlated—when it is likely that sharers only interact with sharers and withholders with withholders—cooperative behaviors would be expected to evolve over time. One way to achieve correlated interaction is when players keep track of the actions of other players, and reciprocate their actions accordingly (See, e.g., Trivers (1971); Skyrms (2001); Binmore (2005)). For instance, industrial scientists may start by sharing their evidence, then on successive rounds choose the action that their partner performed before (a strategy called Tit-For-Tat). Other ways to achieve correlated interaction include secret handshakes (sending a signal to identify cooperators before playing, see Robson (1990)), network reciprocity (players disproportionately interact with certain "neighbors" and adopt a reciprocal strategy, see Alexander (2007)), etc. Under these mechanisms, cooperative behaviors are expected to evolve. But note that this analysis only outlines how mutual sharing can be evolutionarily advantageous between an industry group and other similarly agential groups (e.g. other industry groups). When interacting with a decentralized group that is stuck with the Share strategy (e.g. academic scientists), the best action for an industry group is still Withhold.

Another possibility comes from the observation that in my models, the withholding agents are mutually sharing, but in reality, industrial groups may not share with each other. One might think that if there are multiple mutually withholding industrial groups, then even though they each learn better than the sharing academic scientists, they could reach even better beliefs if the whole community comes to share. This reasoning is (surprisingly) not supported by simulation results. Suppose that we have an epistemic community with three equal-sized subgroups, X that shares with everyone, and Y and Z that only share with in-group members and no one else. My simulation results show that in many cases Y and Z reach the true belief even more frequently than in the situation where X, Y, and Z are all mutually sharing.[40] In other words, the epistemic advantage that Y and Z gain from free-riding off of X's exploration is strong enough that they can even outperform a communist community of a larger size. Furthermore, even if Y and Z *would* perform better if the whole community is communist, they would perform even better if they mutually share but continue to withhold from X, especially when X is stuck with the Share strategy.

Finally, throughout the section, I relied on simulation results from the bandit model in §2.2 to construct the payoff tables, but we know very little about the problem space of our scientific inquiries (Wu and O'Connor, 2023). As §2.3 shows, for a considerable portion of the parameter space, especially when the withhold group is in the minority, it does not end up scoring better than the communist community in the NK landscape problem. If industrial scientists constitute a minority of the scientific community, then since they may not have enough powers to conduct thorough local searches, perhaps it would be rational for them to share. Then, the base game may instead be more like a stag hunt where mutual sharing offers the best payoffs. However, industrial science is rather sizable. For instance, in the US, industry conducts 75% and funds 72% of research and experimental development in 2019, according to the National Science Foundation (Boroush and Guci, 2022). In this case, industrial scientists would still gain an epistemic advantage if they withhold. Moreover, there may be other advantages associated with withholding than the ones modeled here, such as reaching conclusions faster so they can reap financial benefits sooner. These additional advantages may shift the game back to a prisoner's dilemma, even if the problem space is more like an NK landscape.

[40] I tested this situation in the bandit model with 18 total agents, where $P_B = .45$ and $.49$, and $n = 1$.

2.5 Conclusion

In this paper, I use network models to show that, scientists who withhold evidence can obtain *epistemic* advantages. My results suggest that industrial scientists and classified researchers may have additional epistemic incentives to withhold their work, even if the academic credit economy does not apply to them in the same way. I further use my modeling results to construct epistemic games that illustrate the underlying sharing dynamics. Before closing, I will briefly discuss some implications for other models of scientific sharing, connect my models to topics in social epistemology, note limitations of my analysis, and suggest directions for future work.

To start, my modeling results may complicate previous credit-based models of scientific sharing. Both Strevens (2017) and Heesen (2017)'s models assume that the probability that a scientist would make a discovery remains fixed regardless of whether they share during the process of discovery. But my results show that the probability of a scientist making the *right* discovery (as in §2.2) or an *epistemically more significant* discovery (as in §2.3) does depend on whether they share. In the context of Heesen (2017)'s multi-stage model, if agents consistently share throughout the process of discovery, then they may risk prematurely settling on a worse theory, without ever discovering better ones. Furthermore, in the situation where one agent consistently withholds while the rest of the community shares, then even though the sharing agents can publish and claim credit for intermediate results as Heesen (2017) argues, the withholding agent may reach the final stage much faster, likely with epistemically more significant findings, which could bring them even more credit.

Moreover, my models offer new support for the independence thesis in social epistemology, which states that individual and group rationality can come apart (Mayo-Wilson et al., 2011).[41] Indeed, in my models, a subgroup may be epistemically better off if they withhold

[41] See Bradley (2022) for an analysis of different versions of the independence thesis.

evidence, but the epistemically optimal situation for the entire community is when everyone shares. The action that maximizes a subgroup's epistemic payoff and the action that maximizes the community's epistemic payoff come apart.

My analysis admittedly has a few limitations. First, agents in both models all choose the best available epistemic option when they undergo social learning. However, recent work (e.g. Kummerfeld and Zollman (2015); Wu (2023a)) shows that alternative behavioral rules can significantly change epistemic outcomes. Combining these alternative behavioral rules with the asymmetric evidence-sharing dynamics is a worthwhile direction for future research. Moreover, I have only tested a small range of cases where there are multiple withhold groups. It is worthwhile to conduct a fuller analysis with more than two groups, since my limited testing shows that the dynamics may be more complicated than expected.

Furthermore, while it is clear from my analysis that epistemic consideration alone does not incentivize a group to share, I have said very little about what to do as a result. My analysis may support policies that require industrial scientists to share for the benefit of the community, or mechanisms that make academic scientists more agential as a group, or recommendations that ensure a small amount of exploration in the community to reduce the epistemic drawbacks of mutual sharing. We need models that combine instrumental and epistemic incentives of scientific sharing before fuller recommendations are made.

Finally, one may find the differential robustness levels between the bandit models and the NK landscape model independently interesting from a philosophy of modeling perspective. It is worthwhile to explore whether there is a similar difference in robustness for other mechanisms that lead to transient diversity (e.g. those reviewed in Wu and O'Connor (2023)). This can offer insight into how the exploration and exploitation trade-offs differ in these models, and how results obtained from either model can be used to represent and explain reality.

Chapter 3

Better than Best

3.1 Introduction

When solving a complex problem in a group, should group members always choose the best available solution that they are aware of? This question arises when there is a group of people coming together to figure something out. They may be solving a scientific problem, or generating social or cultural innovation. If they care about their epistemic success, should they always choose the epistemically most successful option at the moment?

In this paper, I build simulation models to show that, perhaps surprisingly, a group of agents who individually randomly follow a *better* available solution than their own can end up outperforming a group of agents who individually always follow the *best* available solution. The reason for this result relates to the concepts of transient diversity (Zollman, 2010; Wu and O'Connor, 2023; Smaldino et al., 2022) and cognitive division of labor (Kitcher, 1990; Weisberg and Muldoon, 2009; Thoma, 2015) in epistemic communities. The "better" strategy preserves a diversity of practice in the community for some time, so the community can survey a range of solutions before settling down.[1] The "best" strategy, by contrast, may lock the group in a suboptimal position that prevents further exploration. In a slogan, "better" beats "best."

My models are adapted from Lazer and Friedman (2007)'s model where a network of agents is tasked with solving a sophisticated epistemic landscape problem called the NK landscape problem. Here, agents search in a space with multiple "peaks." They only know the solutions of their neighbors on the network, and their own results of (limited) local exploration, so they may fail to ever discover better solutions globally. Agents in the model face an exploration and exploitation trade-off: do they exploit the solution they currently have and the local peak nearby, or do they explore other regions of the landscape for possibly better solutions? In my models, the "better" strategy allows for a high degree of exploration within the community, even though at every single time step, every agent's expected payoff is strictly no greater than what they would have gained were they follow the "best" strategy.

This result is significant because first, it reveals a tension between individual and group decision-making. Here, groups learn better in the long run when their members do not always choose the best for themselves in the short run. This tension itself is not new in social epistemology. For instance, Mayo-Wilson et al. (2011) proposes the Independence Thesis, which states that individual and group rationality may come apart. My results demonstrate this thesis in another modeling paradigm (c.f. the bandit problem in Mayo-Wilson et al. (2011)).

Second, many feminist philosophers of science (Longino, 1990; Fehr, 2011) suggest that different social groups tend to adopt different approaches to problem-solving, which can be represented by different starting points on an epistemic landscape. The "better" strategy explored in this paper, then, would to be a good way to preserve these diverse approaches.

[1]Though, the "better" strategy preserves a diversity of practice *only* when social learning is not too frequent (see §3.3).

Some of the solutions brought by marginalized groups may not seem promising, perhaps due to a historical lack of resources, but they may nevertheless become epistemically significant after explorations.

The model described here also makes technical contributions to the modeling science literature. First, the modeling paradigm I use, the NK landscape model, is very under-explored in the philosophy of science.[2] But I think this model represents a different yet (I will argue) important type of scientific inquiry. Second, the "better" strategy I introduce here is a new mechanism for transient diversity (see, again, Wu and O'Connor (2023); Smaldino et al. (2022)). Moreover, this strategy, unlike other mechanisms, never makes an individual worse off from one round to another, and yet the community still performs relatively well. This may be a more practical strategy for generating epistemically-beneficial diversity.

This paper is organized as follows. In §3.2, I first provide a general interpretation of epistemic landscape models in the context of scientific problem solving. I then introduce the details of my model, including the two behavioral rules: "better" and "best." In §3.3, I present the main simulation results of this paper. In §3.4, I consider a variation of the model: a mixed community where some agents adopt the "better" strategy, while others adopt the "best." In §3.5, I draw implications of the results in the social and cognitive diversity literature.

[2]With the exception of Wu (2022b); Alexander et al. (2015). Though, the NK landscape model, with or without a network structure, is more widely used in theoretical biology (Kauffman and Levin, 1987; Kauffman and Weinberger, 1989), cultural innovation (Lazer and Friedman, 2007; Gomez and Lazer, 2019; Barkoczi and Galesic, 2016), and organizational design (Ethiraj and Levinthal, 2004; Siggelkow and Levinthal, 2003; Marengo et al., 2000).

3.2 The Model

3.2.1 Interpreting the Epistemic Landscape Model

A landscape model contains a large number of points with varying heights. In the context of scientific problem-solving, we can use a landscape model to represent a group of scientists coming together to solve a problem by trying different research approaches. Each approach is a point on the landscape and has a score associated with it, representing its "epistemic significance," e.g. how truth-conducive or fruitful it is. I interpret research approaches on such an epistemic landscape broadly. They can have many components, including research questions, methods, skills, instruments, etc. (Thoma, 2015).

Epistemic landscape models can represent important aspects of scientific problem solving—scientists constantly communicate with others in their community and decide whether they should stick with the research approach they currently have or try new ones, either by exploring on their own or adopting an approach from the community. Epistemic landscape models, especially lower-dimensional ones with one or two peaks, have been used in the philosophy of science literature to model scientific problem-solving (Hong and Page, 2004; Weisberg and Muldoon, 2009; Thoma, 2015).

3.2.2 The NK Landscape Model

The NK landscape model is a sophisticated multi-dimensional landscape with multiple peaks. This model was originally developed in theoretical biology to study how different variants of a gene work together to produce fitness (Kauffman and Levin, 1987; Kauffman and Weinberger, 1989). The solution space is N-dimensional, with binary strings (consisting of 0s and 1s) of length N as its points. For instance, if $N = 3$, then 001 is a point on the landscape, so is

101, 010, etc. Then, an algorithm with parameter $K (0 < K < N - 1)$ is used to assign a score between 0 and 1 to each point.[3] In the context of scientific problem-solving, we can think of each of these dimensions as a component of a research approach (research questions, tools, skills, etc.).[4] The score of a research approach then depends on how well different components work together in synergy.

Roughly speaking, the parameter K in the NK landscape model determines how rugged the landscape is and how correlated "nearby" scores are. As K increases, the landscape becomes increasingly difficult to search. Though it is impossible to sketch a higher dimensional space, Figure 3.1 provides a stylized representation of the solution space as we vary K.[5] When $K = 0$, the landscape is smooth with one single peak, and when $K = N - 1$, the landscape is totally chaotic, with the value of every point totally uncorrelated with nearby points. The most interesting space is when $0 < K < N - 1$. The landscape becomes rugged with multiple peaks, with nearby points somewhat correlated with each other. We will focus on this parameter space for the rest of the paper, since we can use it to represent complex scientific problems for which similar research approaches are somewhat correlated in epistemic significance. In this regime, it is often the case that a high-scoring solution is only accessible when exploring from a limited patch of the landscape.

Since the NK landscape model contains multiple peaks and a large number of solutions,[6] it contains the following exploration and exploitation trade-off: do agents exploit the epistemic significance of their current solution (and the local optimum nearby), or do they keep exploring the landscape in hope of finding better solutions? If exploration is not enough, then agents may be stuck in local optima, without discovering more promising solutions. If

[3] See Appendix B for details of the algorithm.

[4] That each dimension only has binary options is an idealization. Investigating whether the results hold when this idealization is relaxed is an interesting follow-up project.

[5] This figure represents a landscape where $N > 2$, though it looks like a one-dimensional landscape.

[6] In the $N = 20$ case considered here, there are over a million solutions.

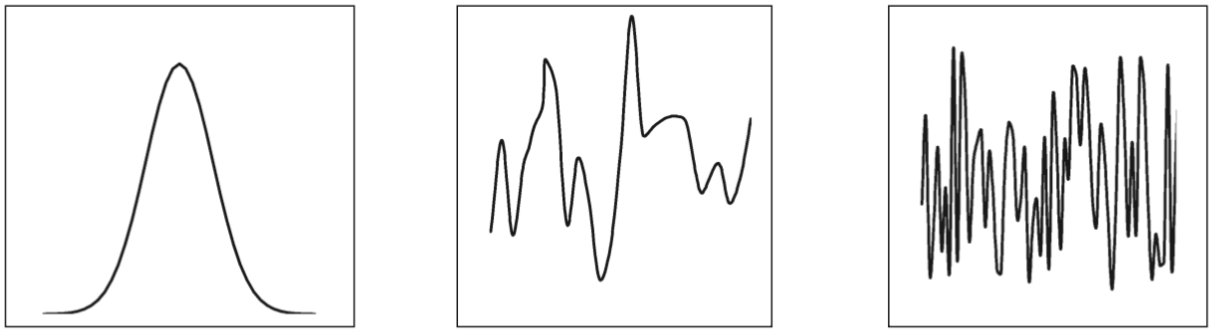

A: A simple problem space, similar to K = 0.

B: A complex rugged space with local maxima and minima, similar to 0 < K < N – 1.

C: A chaotic space in which the value of every point is independent of adjacent points, similar to K = N – 1.

Figure 3.1: Stylized Representation of the Solution Space (Lazer and Friedman, 2007).

exploration is too much, then agents waste time wandering around, trying out potentially low-scoring options.

3.2.3 Network Structures and Initialization

Having specified the solution space of the NK landscape problem, we now turn to the network structure of the model. At the start of each simulation, I have two separate communities of 100 agents, connected to each other via a directed Erdős-Rényi random network. This means that for every agent, the probability that it would form a link with another agent is a fixed number p in $[0, 1]$, and the links formed are not necessarily bidirectional. Every agent is assigned a starting solution, in the form of a binary string of length N. To facilitate direct comparisons, the two communities have identical initial conditions, including the solution space, starting solutions for individual agents, and network structure.

3.2.4 Behavioral Rules

Next, I introduce the two behavioral rules, "best" and "better." All agents in one community follow the "best" behavioral rule, and all agents in the other community follow "better." The only difference between the two behavioral rules is that, during social learning, agents with the "best" rule choose the best solution they are aware of, and agents with the "better" rule randomly choose a better solution than their own.

According to the "best" behavioral rule, in every V round ($1 \leq V \leq 5$), every agent chooses the best-performing solution among all their neighbors' solutions to copy (if there are multiple with the same highest score, they randomly select one of the highest to follow).[7] If they themselves have the best score among their neighbors, then they conduct a *local search* to try to improve their score. This means that they randomly choose a bit in their solution to change (1 to 0, and 0 to 1), and if the change brings a higher score, they switch to that solution.[8] Otherwise, they maintain their current solution. In other rounds (rounds indivisible by V), they do a local search to try to improve their score. We call V the frequency of social learning. This behavioral rule is the same as the one explored in Lazer and Friedman (2007).

According to the "better" behavioral rule, every V round, every agent randomly chooses a better-performing neighbor and copy their solution.[9] If they themselves have the best score among their neighbors, then they do a local search to try to improve their score. In other rounds, they do a local search.

[7] Each simulation starts at round 1, so agents do not go through social learning immediately unless $V = 1$.

[8] As a follow-up, one may be interested in models where agents alter more than one bit during local search. This may represent particular creative or resourceful individuals..

[9] Since an agent chooses a better performing *neighbor* to follow, if there are multiple neighbors with the same solution, the probability that this solution will be selected is proportionally increased. This is different from randomly choosing a better performing *solution* to follow.

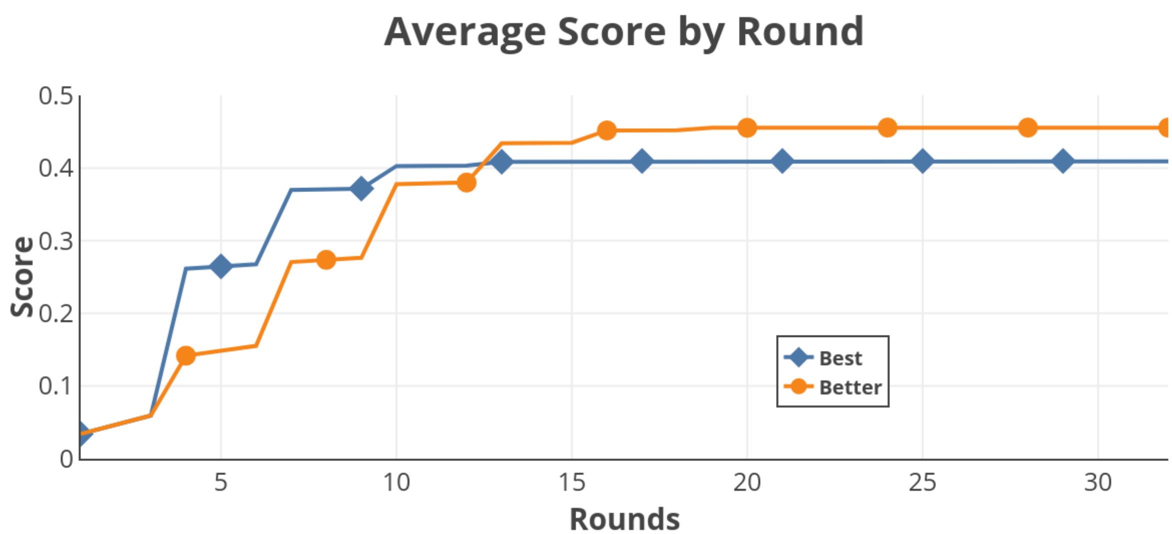

Figure 3.2: $N = 20$, $K = 10$, $V = 3$, $p = .5$. Both communities are stable after 30 rounds.

In the context of scientific problem-solving, we can think of social learning as when scientists improve their own approaches by learning from their friends and collaborators. The "best" behavioral rule requires scientists to always adopt the best solution during social learning, and the "better" behavioral rule requires scientists to randomly choose a better one to mimic. Local search, on the other hand, is when scientists try to improve their own research approach by making small adjustments to it. Because it is typically the case that a high-scoring solution is only accessible from a limited patch of the landscape, agents can only discover their local peak when conducting local search, and social learning is the main mechanism for agents to move to other regions of the landscape.

3.3 Results

I run the model for long enough so that communities stabilize in their solutions. I set $N = 20$ and vary K, V, and p. I run $1,000$ simulations for each parameter combination and present the average results.

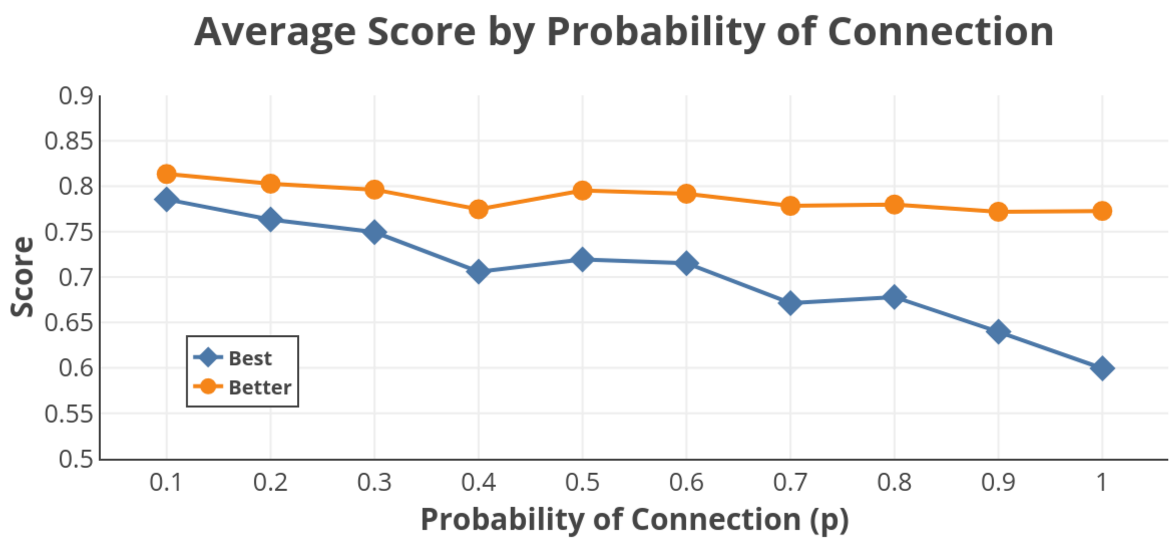

Figure 3.3: $N = 20$, $K = 5$, $V = 5$. Scores shown are average final scores of communities.

The main result is that the community with the "better" behavioral rule ends up having a higher score than the community with the "best" behavioral rule (Figures 3.2 and 3.3). However, the "best" community performs better at the beginning of a simulation, and in general is faster at converging to a consensus solution.[10] This is because a community with the "best" behavioral rule would quickly converge to the vicinity of the most promising solutions that they are aware of, while the "better" community would explore a variety of decent options and their close-by local peaks before building consensus, generating a diversity of practice. It takes the "better" community longer to survey the landscape, but its members are less likely to be stuck in low-scoring peaks.

This trade-off between speed and accuracy in social learning has previously been explored in models about how network connection impacts learning (Lazer and Friedman, 2007; Zollman, 2007, 2010). These models (from two different paradigms) show that a more sparsely connected community learns more accurately but more slowly, precisely because the com-

[10] Here I only present agents' eventual epistemic significance, their average performance per round, and their speed of convergence. For space reasons, I do not investigate other measures, e.g., an agent's average epistemic significance across rounds (Pöyhönen, 2017).

munity experiences a transient period during which a diversity of options is being tested. A more connected community is more likely to settle for an inferior option too early.

The present model introduces more texture to these previous results, since in a sense, the "better" behavioral rule seems to be an effective strategy to counter-balance the dangers of too much connection. As Figure 3.3 shows, the learning accuracy of the "best" community drops significantly as the network becomes more connected. But the "better" community maintains a high level of learning accuracy even as the community becomes more connected. Since limiting connectivity has been criticized as an impracticable way of improving epistemically-beneficial diversity (Rosenstock et al., 2017), encouraging individual community members to adopt something that mimics the "better" behavioral rule may be a plausible alternative.

While the result of "better" beating "best" is fairly robust overall, it is less robust when social learning happens every round.[11] To see why, let us consider a simplified case. Suppose that both the "better" and the "best" communities have A, B, and C as their starting solutions. Further, suppose that solution D is a local peak accessible from local search from B, and E is a local peak accessible from local search from C. Finally, suppose that $A < B < C < D < E$ in epistemic significance. In this scenario, the "better" community would be split between B and C first, and in cases where the agents with B solution discover D before agents with C solution discover E, the community would converge to D, without ever discovering E. In the "best" community, however, agents would all quickly converge to C first, and with sufficient local exploration, would discover E. When social learning slows down, this scenario happens less often, because agents in the "better" community would have enough "time" to explore the vicinity around both B and C sufficiently, so it is more likely that both D and E are discovered before the next social learning. This is a simplified case, but qualitatively similar situations happen with non-negligible probability in the full complex model. This suggests that in order for a diversity of practice to be beneficial to social learning, it has

[11] The result holds more than 74% across all parameter combinations, and 100% when $V > 1$.

to be sustained in the community for some time to allow for sufficient local exploration; infrequent social learning makes this possible.

3.4 Variation: Mixed Community

I now introduce a variation of the model: a mixed community where some agents adopt the "better" strategy, while others adopt the "best." In this community, even though all agents eventually converge to the same solution, agents who adopt the "best" solution reap more epistemic benefits at the beginning, as compared to agents who adopt "better" (Figure 3.4).

Comparing this community with the two communities studied in §3.3, we see that the mixed community ends up outperforming the "best" community, and does not do as well as the "better" community (Figure 3.4). The mixed community outperforms the "best" community precisely because of the agents who adopt the "better" strategy, and that creates a transient diversity of approaches in the community. The agents that adopt the "best" strategy here are essentially free-riding on the epistemic benefits that the "better" strategists provide. In so doing, they get the best of both worlds—they do (relatively) well eventually, while adopting high-scoring solutions in early game.

This creates a dilemma: the "better" strategists are useful to have in the community, but those agents may not be epistemically incentivized to keep their strategy. How should we encourage this epistemically exploratory behavior that benefits the community? I think promising solutions involve structuring the community in such a way that some agents find the "better" strategy attractive for other (intrinsic or extrinsic) reasons. For instance, Nguyen (2022) argues that *intellectual playfulness*—a disposition to try out new ideas for fun—functions as an intellectual "insurance policy" against what he calls *epistemic traps*. This, in a sense, is in line with my results—if some scientists are intrinsically motivated to try

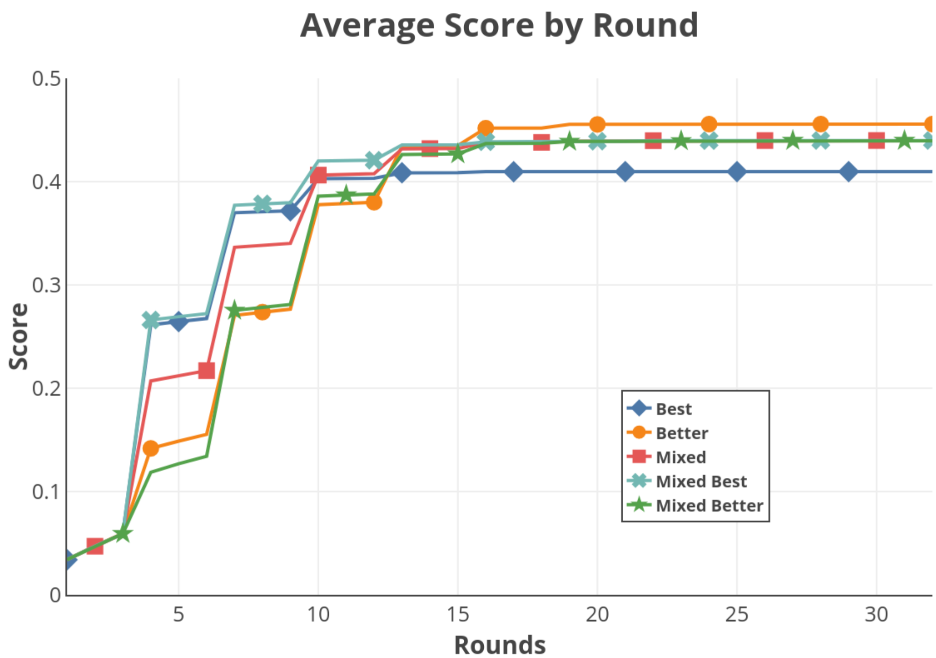

Figure 3.4: $N = 20$, $K = 10$, $V = 3$, proportion of the "better" group=.4. All communities are stable after 30 rounds. "Mixed" denotes the average score of the mixed community; "mixed best" denotes the average score of the "best" strategists in the mixed community; and "mixed better" denotes the average score of the "better" strategists in the mixed community.

out exploratory solutions for the fun of it, then having them in an epistemic community and learning from them can be epistemically beneficial to the community.[12] Another idea is to offer extrinsic incentives for exploratory ideas, such as coordinating funding agencies such that some amount of exploratory work is always funded and promoted. Moreover, various authors argue that even without dedicated funding, credit considerations alone may incentivize scientists to pursue exploratory research, since fewer individuals work on these topics (Kitcher, 1990; Strevens, 2003).

[12]While they share an exploratory attitude, an intellectually playful person completely disregards the truth value of the belief system they try out, but a "better" strategist still chooses a solution that is higher in epistemic significance. However, if some agents in the community choose to "randomly walk" in the landscape without any concern for epistemic significance, they will generate epistemically-beneficial diversity too.

3.5 Coda: Social and Cognitive Diversity

In this paper, I present an epistemic landscape model in which a group of agents who randomly choose a better solution than their own can outperform a group of agents who always choose the best available solution. I argue that this result has a natural interpretation in the context of scientific problem-solving. A group of scientists who entertain a diverse range of reasonable research approaches for some time can outperform a group of scientists who always choose the best available research approaches. Before I close the paper, I will draw some implications of these results in the social and cognitive diversity literature.

First of all, note that the epistemically-beneficial diversity of practice in the main model is *not* the same as cognitive diversity. Cognitive diversity is usually understood as the presence of agents with different cognitive styles—different ways of gathering, processing, or acting on data (Hong and Page, 2004; Pöyhönen, 2017). The most epistemically successful community considered in this paper is a *homogeneous* group of agents who all follow the "better" rule.[13] Moreover, even though a mixed community also performs well, it does well by virtue of having "better" strategists in the community, not by virtue of the *diversity* in cognitive styles.

Second, as briefly discussed in §3.3, in order for the diversity of practice to be epistemically beneficial in this model, the diverse range of solutions needs to be *sustained* in the community for some time. This allows for sufficient exploration of the local region around individual solutions, so local peaks are more likely to be discovered in between social learning. This means that if social learning is too frequent, it brings the same kind of epistemic harm as too much network connection (c.f. Lazer and Friedman (2007)). Infrequent social learning and sparse network structures can both ensure that a local region is sufficiently explored before an agent moves elsewhere.

[13] This is in line with Weisberg and Muldoon (2009)'s (controversial) result that a homogeneous group of mavericks perform well.

Finally, feminist philosophers of science have written extensively about how diverse social groups tend to have diverse background beliefs and approaches to problem solving (Longino, 1990; Fehr, 2011), which may then be represented by diverse initial patches on an NK landscape model. If this is right, then the "better" behavioral rule is one way to help preserve those initially plausible but not outstanding solutions from diverse (social) locations. Perhaps a solution from some marginalized social group does not stand out initially due to a historical lack of resources, but a stellar solution may be reachable if we explore in its vicinity. Moreover, I have also shown that having a diverse range of solutions simpliciter (more crudely—having diverse bodies in the room) is not enough, exploration in its vicinity needs to be supported in a sustaining fashion, so that agents do not prematurely switch to a mainstream approach without realizing the potential of their local perspectives.

Appendix A

End States of the Generalized Bandit Model

After $10,000$ rounds, the community typically settles into one of three states, defined below:

- **Community convergence to the true belief** is reached when, for each agent, the following three conditions are satisfied: (S1) $\frac{\alpha_A}{\alpha_A+\beta_A} > \frac{\alpha_B}{\alpha_B+\beta_B}$; (S2) $\frac{\alpha_A}{\alpha_A+\beta_A} > P_B$; (S3) $\frac{\alpha_B}{\alpha_B+\beta_B} < P_A$. The first condition ensures that all individual agents think that option A performs better than option B. The third condition ensures that agents' actions are, for the most part, stable, because agents continue to receive evidence about A.[1] The second condition is included for consistency reasons—we will need it as a stability check in simulations that lead to polarization.[2]

[1]Because of the probabilistic nature of this model, this condition does not ensure that individual actions are always stable when it is satisfied.

[2]As a technical aside, the "success conditions" listed here are not the analogue of the success condition discussed in Wu (2022a), in the following sense. In Wu (2022a), each agent's credence is given by a single number between 0 and 1, representing their credence in the statement "B is better than A." Then, an agent reaches the success condition if this number is $> .99$ (in that model, B is in fact better than A). For the success condition in our current model to be analogous to the one before, we need to calculate the probability that one beta distribution performs better than the other beta distribution and ask whether this probability is $> .99$. This amounts to calculating the following formula

- **Community convergence to the false belief** is reached when, for each agent, the following two conditions are satisfied: (F1) $\frac{\alpha_A}{\alpha_A+\beta_A} < \frac{\alpha_B}{\alpha_B+\beta_B}$; (F2) $\frac{\alpha_A}{\alpha_A+\beta_A} < P_B$. Condition (F1) says that all agents think that the worse action, B, is better. (F2) is a stability condition similar to (S3).

- **Polarization** is reached when, for every member of the withhold group, the success conditions are satisfied, and for every member of the share group, the failure conditions are satisfied. Condition (S2) is needed because the withhold group continues to receive evidence about B from the share group.

$$Pr(p_B > p_A) = \int_0^1 \int_{P_A}^1 \frac{{p_A}^{\alpha_A-1}(1-p_A)^{\beta_A-1}}{B(\alpha_A,\beta_A)} \frac{p_B^{\alpha_B-1}(1-p_B)^{\beta_B-1}}{B(\alpha_B,\beta_B)} dp_B dp_A.$$

It turns out that even after simplification, this formula is still too computationally taxing. So I adopted the current success conditions. See Miller (2015) for details on Bayesian A/B testing.

Appendix B

The Solution Space of the NK Landscape Model

I now describe how scores are assigned to solutions in the NK landscape model. I run this algorithm at the start of each simulation. So even for the same value of N and K, solution spaces are different in each simulation.

I start with randomly selecting a K-tuple, $J = (j_1, ..., j_K)$, where $j_1, ..., j_K$ are integers between 1 and N inclusive, without repeating. I then generate a valuation function f that takes binary strings of length $K+1$ to $(0,1)$ by assigning each possible binary combination a random number from $(0,1)$. For instance, if $K = 2$, then $f(0,0,1) = .549$, $f(0,1,0) = .235$, $f(1,1,0) = .652$, etc. would be a valuation function. I randomly generate J and f at the start of every simulation, and they remain fixed throughout the simulation.

Remember that each solution S is a binary string of length N, i.e. $S = (s_1, s_2, ..., s_N)$, where $s_i \in \{0,1\}$ for $1 \leq i \leq N$. Each S can now be associated with a score, via the following function

$$F(S) = \frac{1}{N}(\sum_{i=1}^{N} f(s_i, s_{i+j_1}, ..., s_{i+j_K})).$$

Here, for cyclicity, we set $s_l = s_{l-N}$ for $l > N$. What this function does is to calculate the score of a point as a sum of the "value" of each bit of the string, while taking into consideration its interaction with K other bits.[1] This allows for correlations among nearby solutions to be reflected in the final score. The score is then normalized (multiplied by $\frac{1}{N}$).

To give an example of this, suppose that $N = 3$, $K = 1$, and $J = (1)$. Suppose further that $f(0,0) = .235$, $f(0,1) = .367$, $f(1,0) = .785$, and $f(1,1) = .954$. Then the score of the binary string $S = (0,1,1)$ is

$$F(S) = \frac{1}{3}(\sum_{i=1}^{N} f(s_i, s_{i+1})) = \frac{1}{3}(f(0,1) + f(1,1) + f(1,0)) = \frac{1}{3}(.367 + .954 + .785) = .702.$$

In general, every landscape generated is going to have a different maximum score. For ease of comparison across simulations, I further divide the score of every point by the maximum score in the landscape to normalize it, i.e. $\tilde{F}(S) = F(S)/\text{MaxScore}$. This way, the highest score in every landscape generated is 1.

Finally, because of how the NK landscape model is generated, the distribution of scores is similar to that of a normal distribution, where the majority solutions are moderately good at solving a problem. This can be unrealistic since usually haphazard solutions to a problem do not perform well. I thus follow Lazer and Friedman (2007) and apply the following transformation $\hat{F}(S) = (\tilde{F}(S))\hat{\ }8$ to every score, such that most solutions receive

[1] In management and organizational science, there is a fairly vibrant literature on alternative functions called "interaction matrices" that allow for coupling between only selected bits. See, e.g., Siggelkow and Levinthal (2003); Ethiraj and Levinthal (2004).

a low score, and high scores are distinguished. Note that none of the qualitative results I present depend on this transformation, since the transformation preserves the order among the scores. $\hat{F}(S)$ is what I mean by score throughout the paper.

Bibliography

Alexander, Jason McKenzie (2007). *The structural evolution of morality.* Cambridge University Press.

Alexander, Jason McKenzie, Johannes Himmelreich, and Christopher Thompson (2015). Epistemic landscapes, optimal search, and the division of cognitive labor. *Philosophy of Science*, 82(3): 424–453.

Bala, Venkatesh and Sanjeev Goyal (1998). Learning from neighbours. *The review of economic studies*, 65(3): 595–621.

Baldwin, James (1993). *Nobody Knows My Name: More Notes of a Native Son.* Vintage.

Barkoczi, Daniel and Mirta Galesic (2016). Social learning strategies modify the effect of network structure on group performance. *Nature communications*, 7(1): 13109.

Batterman, Robert W and Collin C Rice (2014). Minimal model explanations. *Philosophy of Science*, 81(3): 349–376.

Binmore, Ken (2005). *Natural justice.* Oxford university press.

Bokulich, Alisa (2014). How the tiger bush got its stripes: 'how possibly' vs. 'how actually' model explanations. *The Monist*, 97(3): 321–338.

Boroush, Mark and Ledia Guci (2022). Research and development: Us trends and international comparisons. *National Science Foundation.*

Bradley, Clara (2022). The many independence theses.

Brainard, Jeffrey and Jocelyn Kaiser (2022). Us to require free access to papers on all research it funds. *Science*, 377(6610): 1026–1027.

Bright, Liam Kofi (2017). On fraud. *Philosophical Studies*, 174(2): 291–310.

Bright, Liam Kofi (2023). Ida b. wells-barnett's the red record. In Schliesser, Eric, editor, *Neglected Classics of Philosophy, II.* Oxford University Press.

Bright, Liam Kofi and Remco Heesen (2023). To be scientific is to be communist. *Social Epistemology.*

Collins, Patricia Hill (2002). *Black Feminist Thought: Knowledge, Consciousness, and the Politics of Empowerment*. routledge.

DeAngelis, Tori (2003). Does industry funding deserve a bad rap? *Monitor on Psychology*, 34(7).

Deo, Meera E (2019). *Unequal Profession: Race and Gender in Legal Academia*. Stanford University Press.

Derex, Maxime, Charles Perreault, and Robert Boyd (2018). Divide and conquer: intermediate levels of population fragmentation maximize cultural accumulation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1743): 20170062.

Dotson, Kristie (2011). Tracking epistemic violence, tracking practices of silencing. *Hypatia*, 26(2): 236–257.

Du Bois, William Edward Burghardt (2008). *The Souls of Black Folk*. Oxford University Press.

Ethiraj, Sendil K and Daniel Levinthal (2004). Modularity and innovation in complex systems. *Management science*, 50(2): 159–173.

Fang, Christina, Jeho Lee, and Melissa A Schilling (2010). Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning. *Organization Science*, 21(3): 625–642.

Fatima, Saba (2017). On the edge of knowing: Microaggression and epistemic uncertainty as a woman of color. In Cole, Kirsti and Holly Hassel, editors, *Surviving Sexism in Academia: Feminist Strategies for Leadership*, pages 147–157. Routledge.

Fazelpour, Sina and Daniel Steel (2022). Diversity, trust, and conformity: A simulation study. *Philosophy of Science*, 89(2): 209–231.

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

Fricker, Miranda (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press.

Galison, Peter (2004). Removing knowledge. *Critical Inquiry*, 31(1): 229–243.

Golub, Benjamin and Matthew O Jackson (2012). How homophily affects the speed of learning and best-response dynamics. *The Quarterly Journal of Economics*, 127(3): 1287–1338.

Gomez, Charles and David Lazer (2019). Clustering knowledge and dispersing abilities enhances collective problem solving in a network. *Nature communications*, 10(1): 1–11.

Hartsock, Nancy CM (1983). The feminist standpoint: Developing the ground for a specifically feminist historical materialism. In *Discovering reality*, pages 283–310. Springer.

Hartsock, Nancy CM (1997). Comment on hekman's "truth and method: Feminist standpoint theory revisited": Truth or justice? *Signs*, pages 367–374.

Heesen, Remco (2017). Communism and the incentive to share in science. *Philosophy of Science*, 84(4): 698–716.

Hekman, Susan (1997). Truth and method: Feminist standpoint theory revisited. *Signs: Journal of Women in Culture and Society*, 22(2): 341–365.

Hong, Lu and Scott E Page (2004). Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences*, 101(46): 16385–16389.

Hornsby, Jennifer (1995). Disempowered speech. *Philosophical topics*, 23(2): 127–147.

Intemann, Kristen (2010). 25 years of feminist empiricism and standpoint theory: Where are we now? *Hypatia*, 25(4): 778–796.

Kant, Immanuel (1960). *Observations on the Feeling of the Beautiful and Sublime*, volume 97. University of California Press.

Kauffman, Stuart and Simon Levin (1987). Towards a general theory of adaptive walks on rugged landscapes. *Journal of theoretical Biology*, 128(1): 11–45.

Kauffman, Stuart and Edward Weinberger (1989). The nk model of rugged fitness landscapes and its application to maturation of the immune response. *Journal of theoretical biology*, 141(2): 211–245.

Kitcher, Philip (1990). The division of cognitive labor. *The journal of philosophy*, 87(1): 5–22.

Kummerfeld, Erich and Kevin JS Zollman (2015). Conservatism and the scientific state of nature. *The British Journal for the Philosophy of Science*, 67(4): 1057–1076.

Lange, Marc (2000). Is jeffrey conditionalization defective by virtue of being non-commutative? remarks on the sameness of sensory experiences. *Synthese*, pages 393–403.

Lazer, David and Allan Friedman (2007). The network structure of exploration and exploitation. *Administrative science quarterly*, 52(4): 667–694.

Longino, Helen (1993). Subjects, power, and knowledge: Description and prescription in feminist philosophies of science. *Feminist Epistemologies*, pages 101–120.

Longino, Helen E (1990). *Science as social knowledge*. Princeton university press.

Louis, Karen Seashore, Lisa M Jones, and Eric G Campbell (2002). Macroscope: Sharing in science. *American Scientist*, 90(4): 304–307.

Macfarlane, Bruce and Ming Cheng (2008). Communism, universalism and disinterestedness: Re-examining contemporary support among academics for merton's scientific norms. *Journal of Academic Ethics*, 6(1): 67–78.

March, James G (1991). Exploration and exploitation in organizational learning. *Organization science*, 2(1): 71–87.

Marengo, Luigi, Giovanni Dosi, Paolo Legrenzi, and Corrado Pasquali (2000). The structure of problem-solving knowledge and the structure of organizations. *Industrial and Corporate Change*, 9(4): 757–788.

Mason, Winter and Duncan J Watts (2012). Collaborative learning in networks. *Proceedings of the National Academy of Sciences*, 109(3): 764–769.

Mayo-Wilson, Conor, Kevin JS Zollman, and David Danks (2011). The independence thesis: When individual and social epistemology diverge. *Philosophy of Science*, 78(4): 653–677.

McGarity, Thomas O and Wendy E Wagner (2010). *Bending science: How special interests corrupt public health research*. Harvard University Press.

McPherson, Miller, Lynn Smith-Lovin, and James M Cook (2001). Birds of a feather: Homophily in social networks. *Annual review of sociology*, 27(1): 415–444.

Merton, Robert K (1973). *The Sociology of Science: Theoretical and Empirical Investigations*. University of Chicago Press.

Michaels, David (2008). *Doubt is their product: how industry's assault on science threatens your health*. Oxford University Press.

Miller, Evan (2015). Formulas for bayesian a/b testing. `https://www.evanmiller.org/bayesian-ab-testing.html`. Accessed: 2021-10-08.

Mills, Charles (2007). White ignorance. In Sullivan, Shannon and Nancy Tuana, editors, *Race and Epistemologies of Ignorance*, pages 11–38. Albany, NY: State Univ of New York Pr.

Narayan, Uma (1988). Working together across difference: Some considerations on emotions and political practice. *Hypatia*, 3(2): 31–48.

Neely, Barbara (1993). *Blanche on the Lam*. Penguin Group USA.

Nguyen, C Thi (2022). Playfulness versus epistemic traps. In *Social Virtue Epistemology*, pages 269–290. Routledge.

O'Connor, Cailin (2020). *Games in the Philosophy of Biology*. Cambridge University Press.

O'Connor, Cailin and James Owen Weatherall (2018). Scientific polarization. *European Journal for Philosophy of Science*, 8(3): 855–875.

Partha, Dasgupta and Paul A David (1994). Toward a new economics of science. *Research policy*, 23(5): 487–521.

Pöyhönen, Samuli (2017). Value of cognitive diversity in science. *Synthese*, 194(11): 4519–4540.

Robson, Arthur J (1990). Efficiency in evolutionary games: Darwin, nash and the secret handshake. *Journal of theoretical Biology*, 144(3): 379–396.

Rosenstock, Sarita, Justin Bruner, and Cailin O'Connor (2017). In epistemic networks, is less really more? *Philosophy of Science*, 84(2): 234–252.

Ross, Don (2021). Game Theory. In Zalta, Edward N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2021 edition.

Rubin, Hannah and Cailin O'Connor (2018). Discrimination and collaboration in science. *Philosophy of Science*, 85(3): 380–402.

Saint-Croix, Catharine (2020). Privilege and position. *Res Philosophica*, 97(4): 489–524.

Settles, Isis H, Martinque K Jones, NiCole T Buchanan, and Kristie Dotson (2020). Epistemic exclusion: Scholar(ly) devaluation that marginalizes faculty of color. *Journal of Diversity in Higher Education*.

Siggelkow, Nicolaj and Daniel Levinthal (2003). Temporarily divide to conquer: Centralized, decentralized, and reintegrated organizational approaches to exploration and adaptation. *Organization Science*, 14(6): 650–669.

Skyrms, Brian (2001). The stag hunt. In *Proceedings and Addresses of the American Philosophical Association*, volume 75, pages 31–41.

Smaldino, Paul E, Cody Moser, Alejandro Pérez Velilla, and Mikkel Werling (2022). Maintaining transient diversity is a general principle for improving collective problem solving.

Strevens, Michael (2003). The role of the priority rule in science. *The Journal of philosophy*, 100(2): 55–79.

Strevens, Michael (2017). Scientific sharing: Communism and the social contract. *Scientific collaboration and collective knowledge*, pages 1–50.

Thoma, Johanna (2015). The epistemic division of labor revisited. *Philosophy of Science*, 82(3): 454–472.

Toole, Briana (2020). Demarginalizing standpoint epistemology. *Episteme*, pages 1–19.

Trivers, Robert L (1971). The evolution of reciprocal altruism. *The Quarterly review of biology*, 46(1): 35–57.

Weisberg, Michael and Ryan Muldoon (2009). Epistemic landscapes and the division of cognitive labor. *Philosophy of science*, 76(2): 225–252.

Wells-Barnett, Ida B (2012). *The Red Record Tabulated Statistics and Alleged Causes of Lynching in the United States*. Tredition.

Wu, Jingyi (2022a). Epistemic advantage on the margin: A network standpoint epistemology. *Philosophy and Phenomenological Research*.

Wu, Jingyi (2022b). Withholding knowledge.

Wu, Jingyi (2023a). Better than best: Epistemic landscapes and diversity of practice in science.

Wu, Jingyi (2023b). Modeling injustice in epistemic networks. https://philosophyofbrains.com/2023/02/14/cognitive-science-of-philosophy-symposium-network-modeling.aspx. Accessed: 2023-04-13.

Wu, Jingyi, Cailin O'Connor, and Paul E Smaldino (2023). The cultural evolution of science. In *Oxford Handbook of Cultural Evolution*. Oxford University Press.

Wu, Jingyi and Cailin O'Connor (2023). How should we promote transient diversity in science? *Synthese*, 201(2): 37.

Wylie, Alison (2003). Why standpoint matters. In Figueroa, Robert and Sandra G. Harding, editors, *Science and Other Cultures: Issues in Philosophies of Science and Technology*, pages 26–48. Routledge.

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

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

Zollman, Kevin JS (2018). The credit economy and the economic rationality of science. *The Journal of Philosophy*, 115(1): 5–33.

Zollman, Kevin JS (2021). The theory of games as a tool for the social epistemologist. *Philosophical Studies*, 178(4): 1381–1401.