Networks and Values

Rebecca Korf and Cailin O'Connor

Abstract

The values in science literature considers deep questions around value laden aspects of inquiry like: Where and when can values of different sorts legitimately influence science? What sorts of values, and related ethical practices, best promote scientific inquiry? How do individual beliefs and biases impact science? Earlier work in this area focused mostly on how values interact with individual practice. But increasingly there has been a move towards thinking about how the impacts of values can reverberate through scientific communities, and how the dynamics of social groups may lead to unexpected outcomes when values shape science. At the same time, social epistemologists and others have been using network models to think about how social network structures and group features shape knowledge generation. These models are often useful in thinking about the emergent, group-level effects of values on scientific progress. This paper will briefly survey this network literature and connect it with the literature on values in science.

1 Introduction

The values in science literature has considered deep questions around value laden aspects of inquiry like: Where and when can values of different sorts legitimately influence science? What sorts of values, and related ethical practices, best promote scientific inquiry? How do individual beliefs and biases impact science? Earlier work focused mostly on how values interact with individual practice. But increasingly there has been a move towards thinking about how the impacts of values can reverberate through scientific communities, and how the dynamics of social groups may lead to unexpected outcomes when values shape science. At the same time, social epistemologists and others have been using network models to think about how social network structures and group features shape knowledge generation. These models are often useful in thinking about the emergent, group-level effects of values on scientific progress.

This paper will briefly survey this network literature and connect it with the literature on values in science. As we will see, the causal links between values and epistemic outcomes are not always simple or predictable. Seemingly unproblematic types of value influence can sometimes have unexpected negative impacts on epistemic outcomes, and seemingly problematic sorts of value influence can sometimes improve outcomes. In other cases network models reveal that policy proposals to improve scientific functioning may not work as suggested. Still other models point to places where new work on values in science may be useful.

A central take-away is that the impacts of value influence are often dependent on emergent, group-level dynamics, which are hard to predict. Thus philosophers should take these sorts of phenomena into account, and generally take care, when thinking about value influence. This sort of caution is especially important in the domain of policy discussions where philosophical reasoning may impact real world decision-making.

The rest of the paper will proceed as follows. Section 2 briefly surveys relevant literature on values in science. The heart of the paper is section 3 which discusses results from network models and their relevance to the values in science literature. And section 4 briefly concludes.

2 Values in Science

The values and science literature is concerned with both epistemic impacts and ethical issues arising from value influence in science. While the literature has persuasively argued that science has responsibilities beyond simply providing as much knowledge as possible, it is also driven by the fact that an important goal of science is to provide helpful knowledge on a desirable time frame (and that there is an ethical responsibility to do so). A major question is how different instances of value influence might impact science's ability to fulfill this goal.

Early work in values and science generally focused on the ways that values or social beliefs held by an individual scientist might affect their decisions at various points of the scientific practice. The classic paper from Rudner (1953), for example, is about value judgements individual scientists make and much work that followed retained this individual lens. Recent prominent proposals for value management strategies like those from Douglas (2009), Elliott (2017), Steel (2017), and Brown (2020) also take the individual scientist as their target.

This individual lens has yielded important insights but, of course, scientists do not work alone. Scientists collaborate with colleagues, interact within communities of researchers, contend with social structures in their disciplines and institutions, work within societies that include political and industrial influences, and are part of wider publics which are interested in and affected by scientific results. For this reason some researchers have focused more on the group or community level when thinking about value influence in science. Notable early examples of this sort include Longino (1990) and Solomon (2001) who call for objectivity and knowledge to be reframed as social concepts.

Recently philosophers have worked on the impact of institutional structures on values in science. Douglas (2018, 2021), for example, explores which norms are useful for institutions, as well as individuals, in promoting ethical responsibility in science. She also focuses on how to hold scientists and institutions accountable when they fail to act ethically. Rolin (2015) argues that scientific collaboration illuminates how values can act differently on the individual and scientific level.

It is clear that including a social focus is important for the values in science literature. One persistent challenge for socially focused work is that predicting and understanding the dynamics of social groups is notoriously difficult. In the next section we turn to network models, which are specifically designed to help theorists get a handle on phenomena in epistemic communities, like groups of scientists. These models have proved useful to social epistemology—a sub-field with ongoing links to values in science. As we will show, they are useful to values in science as well.

3 Values and Network Models

Social group dynamics are complex and sometimes difficult to predict. In many cases individual decisions, preferences, or traits can lead to surprising emergent group effects. Simple models in the social sciences have been used to demonstrate and underpin this fact. Models are especially helpful here because they are easy to control and understand. Thus, unlike in most, complex, real-world communities, researchers using models can definitively show how particular features lead to emergent group level effects, at least in principle.

The values in science literature is interested in epistemic communities, and especially scientific ones. Work on this sort of community provides ample evidence for emergent, surprising, group level phenomena, as we will see (Mayo-Wilson et al., 2011). Some models demonstrate how value influence on the small scale can add up to emergent group effects. Others are relevant to policies intended to manage value influence. Still others point to under-explored topics where the values in science literature has an important role to play. Here we introduce this literature and discuss its relevance for values in science.

3.1 Industrial Selection

One important question for values in science is how and whether industrial values—those related to the desire to develop and sell new technologies or products—can legitimately influence science. Some argue we should use the roles values play in science, and the context in which their influence occurs, to determine the legitimacy of value influence, rather than the source or identity of such values (most prominently, Douglas (2009)). Under this sort of account, the "external" stages of science, including funding choices, can be legitimately influenced by a wide variety of values, as long as "internal" parts of research are not affected. So, for example, a pharmaceutical company could legitimately fund or even directly produce research on the efficacy of a drug, as long their interests did not directly influence the research process.

As Holman and Bruner (2017) and Korf (2024) argue, though, recent results in network modeling raise worries for this sort of account. Holman and Bruner (2017) develop a network model to demonstrate how industrial influence over research funding can lead to surprising negative outcomes. The type of model they use was developed in economics by Bala and Goyal (1998) and brought into philosophy by Zollman (2007). Agents learn which of two actions, A or B, is more successful by gathering data about their preferred action and sharing it with network neighbors. Over time, they update their beliefs about the actions in light of data. The agents can represent a community scientists testing two action-guiding theories and sharing their research over time—such as medical researchers learning whether bacteria or stomach acid is the primary cause of ulcers, whether smoking is likely to cause lung cancer or not, or which of two pills is more efficacious (Zollman, 2007, 2010; Weatherall et al., 2020). We do not go into details of the model for space reasons.

Inspired by historical cases Holman and Bruner (2017) develop a version where agents use different methods to explore the world. These methods are better or worse at tracking the actual underlying quality of the two actions. For instance, one medical researcher might use a highly skewed sample population, and thus develop a poor understanding of a drug, while another researcher using a representative sample might get better data. Thus agents testing A and B actually draw data from slightly different distributions.

The key feature of their model is that industry agents choose which researchers to fund, and, in doing so, shape their eventual influence on the community. Agents with more funding gather more data, which then influences the beliefs of network neighbors. The model also assumes that new agents join the network, and when they do they tend to adopt the methods of those who have been more productive (i.e., those with more funding).

They consider cases where industry has a stake in what theories scientists adopt. Suppose traditional treatment B is actually better than new drug A, but industry wants to sell new drug A. Industrial actors find and fund only those researchers who happen to use poor methods, i.e., ones that tend to show A is better than B. While the researchers maintain their academic integrity—they do not alter their methods in light of industry funding—this funding can nonetheless sway the community away from the correct belief.[1]

Holman and Bruner (2017) argue their model raises worries for proposals from Douglas et al. (2014) about mitigating industry harms. While the latter authors are mostly focused on how individuals can preserve their integrity in the face of industry influence, personal integrity cannot mitigate industry harms based in selection.

Korf (2024) points out that this case demonstrates a larger feature of science, which raises worries for accounts like that of Douglas (2009). (This account, remember, limits industrial influence to the external parts of science). When scientists update their beliefs in these models on the results of network neighbors the relative productivity of different agents impacts how their colleagues later evaluate hypotheses. By changing what a community of scientists believes about the world, industrial selection should change how these scientists understand and evaluate evidence later. Thus, external influence that systematically affects productivity—which grants are funded, who is given prestigious univer-

[1] See also Pinto and Pinto (2023) who confirm and extend this result.

sity positions, and which work is chosen to be published, along with industrial funding—ends up shaping internal scientific processes at later stages. Value influence can propagate and reverberate throughout a community in this way. Hence, accounts that limit certain kinds of value influence to external stages of science may fail to actually exclude such influence on the internal level in the longer term.

In thinking about real world communities, we might also expect external influence to become internal via other channels. Industrial selection shapes which theories scientists hold by shaping the distribution of research they see. Background beliefs deeply shape research choices at every stage of the scientific process (Longino, 1990; Okruhlik, 1994; Solomon, 2001; Anderson, 2020). Without further work we do not know what form these dynamics might take, but expect them to be complex and unpredictable. It will thus be hard to assess up front how external value influence should impact internal decision making later on.

The models described here, and their surprising, emergent social effects have several lessons for values in science. 1) Attempts to limit value influence to external stages of science may not prevent unwanted harms. 2) More generally, the demarcation between internal and external influence is not as simple as some have previously thought. And 3) attempts to mitigate industry harms on science cannot focus on individual researcher integrity alone.

3.2 Dissent

A topic that has received substantial attention in the values in science literature, as well as in metascience, is dissent (Solomon, 2001; de Melo-Martín and Intemann, 2018; Intemann and de Melo-Martín, 2014; Paetkau, 2024). Many worry about harmful dissent by industrial interests, or "epistemically detrimental dissent". But normal dissent is a crucial part of the scientific process. This leads to two key questions: Can we delineate inappropriate dissent from appropriate dissent? And if so, should we attempt to suppress inappropriate dissent?

Coates (2024) uses a network model to address the latter question, and especially the debate between those who do think it appropriate to suppress cynical dissent (Nash, 2018; Oreskes, 2017; Cook, 2017; Biddle and Leuschner, 2015; Biddle et al., 2017; Leuschner, 2018) and those who worry that this sort of policy risks suppressing useful dissent in science (de Melo-Martin and Intemann, 2014; de Melo-Martín and Intemann, 2018).

He starts with a network of agents engaging in a two-armed bandit problem. Following Zollman (2009), he assumes that instead of informally sharing results through their social network, scientists read a "journal", which chooses results to share. He compares the journal 1) randomly selecting results to share, and 2) "hiding dissent". In the latter, results are reviewed by community members. If all the reviewers believe the results support the incorrect theory (on their own lights), they are not shared. Thus less popular theories tend to be suppressed.

Coates (2024) finds that hiding dissent tends to harm group outcomes. Communities in this sort of model benefit from a period of diversity of practice, where multiple agents test multiple theories long enough to ensure they pick the best one. Hiding dissent tends to decrease diversity of practice. Thus the community more often preemptively settles on the worse theory. The exception, on his findings, is that after some period of exploration, where the community has already tested both theories extensively, suppressing dissent speeds consensus and does not harm outcomes.

Coates also considers versions of the model where one agent represents some cynical interest, such as an industry representative, following Holman and Bruner (2015). This agent persistently shares data that tends to support the worse theory. Even in this case, Coates finds that hiding dissent is typically harmful as it can quickly push the entire community to the incorrect (industry supported) consensus.

Notably, though, his model does not include any way for agents to detect epistemically detrimental dissent. Thus, the failure mode in hiding dissent—where the community is influenced towards the worse belief by industry, and starts to suppress the better one—might be avoidable in more realistic cases. Holman and Bruner (2015), in an earlier model, study the impacts of an industrial agent who shares misleading data. They consider situations where other agents can stop listening to the industrial agent by disconnecting from others whose data does not match their beliefs. The result is that the community often learns to ignore the biased agent, and performs better.

One take-away might be that even simple learners can successfully detect and avoid some sorts of misleading data. It might very well be possible, with more sophisticated cognitive abilities, to detect and avoid cynical dissenting evidence. In addition, a number of network models have demonstrated the extensive harms of industrial propaganda, lending weight to arguments that industry harms should be taken seriously (Holman and Bruner, 2015, 2017; Weatherall et al., 2020; Lewandowsky et al., 2019). If so, minor risks of suppressing appropriate dissent may be worth taking.

Here models from both Coates (2024) and Holman and Bruner (2015) directly inform debates in values in science over policy proposals by illuminating complex social processes. These models should not be taken to directly predict outcomes in real communities. But they can be used as aids to normal reasoning, or rigorous versions of thought experiments, which help inform the effects of policy choices in science.

3.3 Open Science

The open science movement aims to improve research integrity and access by promoting the open sharing of products of science.[2] To this point it has generally been assumed that increased sharing, including during intermediate stages of science, will improve research outcomes. Some extant results in network modeling, though, raise worries about unexpected downsides. These results do not

[2]Thanks to Kaetlin Taylor for suggesting the relevance of network models to the open science movement.

directly bear on previous work in values and science on sharing or transparency. Still we include them here given the general focus in values in science on science communication, and on the embodiment of ethical values (like transparency) in science policy.

Zollman (2007, 2010) looks at how network structure in a bandit model influences whether communities settle on correct or incorrect consensus. His surprising result is that across many variations of the model groups perform better the less they communicate. Subsequently it has become clear that a host of models, with different varieties of problems, structures, and assumptions, yield similar results (March, 1991; Lazer and Friedman, 2007; Fang et al., 2010; Grim et al., 2013; Frey and Šešelja, 2018, 2020; Rosenstock et al., 2017).

These results goes against the intuitive notion, underlying the open science movement, that more communication is generally better for science. But although initially surprising, once understood the result is intuitive. As noted, mechanisms that promote a period of diverse practice—agents testing multiple theories of some length of time—improve community output in these models. When agents communicate too much, they tend to settle on consensus too quickly, without a sufficient period of exploration. With respect to the open science movement, there is thus a risk in promoting sharing—that it might flatten the diversity of beliefs within a community and thus reduce diversity of practice.

To test this risk more explicitly, Torsell and Korf (2024) develop a different sort of network model. Agents explore an NK-landscape with "peaks" of different magnitudes corresponding to levels of epistemic success. Agents communicate with network neighbors about their current success levels, and tend to follow neighbors who are more successful than themselves, while also exploring locally. Torsell and Korf find that communities who share information more frequently can lead the entire group to herd onto local optima, but miss better solutions because they do not explore enough. Thus their results, again, suggest that too much sharing in science may inhibit progress by reducing diversity of practice.

It should be noted that there may be effective ways to both embody values related to openness and sharing in scientific communities, and simultaneously protect diversity of practice. Wu and O'Connor (2023) point to areas of science that purposefully isolate research teams for periods of time to improve outcomes. In this way open sharing might be preserved as a general value in science, while in specific instances it could be limited to improve diversity of practice.

3.4 Public Communication

Values in science has argued that scientists have ethical and epistemic responsibilities to publicly communicate their findings. Good communication by scientists may improve science literacy, reduce scientific misinformation, and counteract distrust in science (Kitcher, 2011; Desmond, 2023; Gerken, 2022, 2018; Intemann, 2024, 2023; John, 2018).

There are a few models relevant to this topic. The first center around industry attempts to control scientific narratives. Both Weatherall et al. (2020) and

Lewandowsky et al. (2019) consider industrial agents who select what to share from real, independently produced data, and, in doing so, manage to mislead the public. Neither of these models involve fraud or lying, but rather curation. Agents who wish to mislead curate which pieces of evidence the public sees and, in doing so, dramatically reshape public beliefs.

Another relevant model considers how this happens not with industry, but through channels of science journalism. Mohseni et al. (2022) consider a public receiving news items that are filtered through a journalistic process. Two communication strategies they consider are hyperbole—where journalists take real events and exaggerate them to sound more extreme or novel—and extremity—where journalist communicate accurately about events but select only extreme ones to share. They find that extremity—a form of curation—can be just as misleading as hyperbolizing, even though it involves sharing (just some of the) facts accurately.

There are a few take-aways for the values in science literature. First, ethical standards focusing primarily on truthfulness and accuracy of communication can be insufficient. Which facts are shared are often influenced by our values and this value-laden curation of accurate facts can, and very often does, mislead. Thus those worried about public belief should focus, at least to some degree, on how scientific sharing is curated rather than primarily on reducing falsehood.

Second, as noted, there has been a focus on individual responsibilities of scientists to share their work in useful ways. This responsibility is undeniable. But attention to the emergent effects of industry strategy in science communication points to places where even the best individual communication will be insufficient. These strategies are often subtle, successful, well-funded, and carried out by professionals. Thus it may be impractical for individual scientists to counter such efforts. A bigger picture should also address wider media and government policies that ensure successful public communication of science.

The main take-away for the values in science literature here is the suggestion that attention be given to a wider set of emergent, social phenomena in thinking about science communication. It is not that these phenomena have been entirely ignored, but under-discussed. As a result, more focus is needed on the responsibilities of policies makers, rather than individual scientists, to protect public belief.

3.5 Out-Group Mistrust

Two recent models present surprising results about how individual prejudices can influence outcomes in epistemic groups. Both show that in some cases ignoring or devaluing data shared by outsiders can improve outcomes for some, or all, community members. These results are deeply connected to claims from standpoint epistemology and epistemic injustice—two areas of thought with longstanding connections to values in science.

Wu (2023) starts with two communities in a network facing a two-armed bandit problem. She assumes that one of these—the dominant community—can entirely ignore or else partially devalue the data gathered by the other, marginalized, community.[3] Across many versions of this model Wu finds that the marginalized community tends to arrive at accurate beliefs more often, and also, typically, learns faster. In fact, they typically do better even than an entirely networked community. When a dominant group ignores data from a marginalized group in this model, they tend to learn more slowly, and explore for a longer period of time. The marginalized group then learns more quickly, but at the same time benefits from exploration happening in the dominant group. In other words, they are able to exploit the best current evidence, because they do not ignore data, without risking a premature, incorrect consensus. Wu (2023) takes this model to provide a causal pathway for what is sometimes called the "inversion thesis" in standpoint epistemology—that those who are disempowered or socially marginalized are sometimes, paradoxically, in a better position to learn about certain aspects of the world (Wylie, 2003).

Fazelpour and Steel (2022) consider a similar model, but where two groups both devalue data coming from their out-group. In an, again, somewhat surprising finding, they show that moderate out-group mistrust allows the whole group to perform better. This is, again, because mistrust slows learning and promotes exploration within the entire community. They also explore versions of the model where agents tend to conform with neighbors—a tendency that generally harms learners in these models (Weatherall and O'Connor, 2021; Mohseni and Williams, 2021)—with similar findings.

None of these authors advocate that scientists should cultivate mistrust towards identity based out-groups. Rather both take their models as further support of arguments—from feminist philosophy of science and values in science—in favor of identity diversity in scientific communities (Longino, 1990; Okruhlik, 1994). Wu points out that including marginalized groups within science may mean including insights that benefit from historical exclusion. And Fazelpour and Steel argue this sort of diversity may benefit learning groups directly by decreasing herding onto suboptimal theories. In this way, both of these models inform values in science.

3.6 Industrial Withholding

Wu (2022) uses a version of the model just described to think about a quite different topic—the communist norm in science, which says that scientists should share their findings quickly, widely, and freely (Merton, 1942; Bright and Heesen, 2023). In particular, she is interested in the asymmetry between academic science, which adheres to the communist norm, and industrial science, which does not. Typically, industry science is proprietary, meaning that results are kept private and used for a profit rather than freely shared.

As Wu points out, though, this means that industry and academic science stand in a relationship much like the marginalized and dominant communities in her previous model. One group receives data from both communities, and

[3] She directly relates this assumption to work on epistemic quieting (Dotson, 2011) and testimonial injustice (Fricker, 2007).

the other can only learn from one. She develops a different version of the model where actors face an NK-landscape problem. The details are not important here, but the result supports her previous findings—if an industrial community withholds their data while simultaneously learning from an academic one, they learn faster and better. The academics, on the other hand, spend more time exploring possibly sub-optimal theories.

This points to a deep unfairness in the structure of research communities. Academic science is often publicly funded, but industrial science can use academic insights in developing new products and findings, which they do not need to then share. They, in turn, profit off these products and findings (O'Connor, 2023). This asymmetry has not been widely discussed because it is not necessarily obvious a priori, but the model developed by Wu (2022) makes clear why it should be worrying.

These results have obvious implications for the acceptability of certain types of value influence to the structure of epistemic communities. Industry science is aimed at profit—a very different value from those driving academic progress. While profit motives are useful and important to discovery in many cases, there is, as noted, an unfairness here, especially to members of the public whose tax dollars fund academic science. Thus it might be appropriate to regulate the structure of industry science to ensure that industry does not unfairly profit from public funding.

4 Conclusion

As we have seen, interactions in learning groups—even simpler groups than those constituted by real humans—are often complex. Emergent phenomena can arise that are not easy to predict a priori. One important conclusion is that in thinking about normative standards for scientific communities, especially those crafted to promote scientific success, it is not always easy to assess which policies and norms will be best.

Here we have pointed to a number of network modelling results with interesting implications for the values in science literature—a literature deeply interested in the improvement of scientific functioning. We do not think that these models themselves should be generally taken to show what will happen in real groups of learners under analogous conditions. Rather they show some of the possible outcomes—including ones that sometimes support and sometimes press against central arguments and conclusions from the literature. These sorts of possibilities should be taken, at minimum, to complicate some theoretical arguments about science policy.

Some of these models, though, are further supported by examples, cases, or else empirical results that help validate their findings. Holman and Bruner (2017) argue that industrial selection actually occurred during the testing of anti-arrhythmic heart drugs, to unfortunate effect. The "Zollman effect"—where less connection can mean more success—has been detected in a human subjects experiment (Derex and Boyd, 2016). As predicted by Wu (2023) and

Fazelpour and Steel (2022) there is some evidence that in groups with heterogenous identities there is a reduction of herding, with positive results (Bear and Woolley, 2011; Phillips and Apfelbaum, 2012; Sommers, 2006). In each of these cases, there is more reason to think that the sorts of emergent phenomena identified in the models might be occurring in the real world, and to take them seriously in the values in science literature.

To summarize, given its topic matters, the values in science literature should generally be sensitive to the possibility of emergent, surprising social phenomena. Models are one, but not the only, way to explore where and when these phenomena occur.

Acknowledgments

Thanks to Kaetlin Taylor for conversations about the relevance of network models to the open science movement. Many thanks to Kevin Elliott and Ted Richards for their work editing this volume.

References

Anderson, Elizabeth (2020). "Feminist Epistemology and Philosophy of Science." *The Stanford Encyclopedia of Philosophy*. Ed. Edward N. Zalta. Spring 2020 edition. Metaphysics Research Lab, Stanford University.

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

Bear, Julia B and Anita Williams Woolley (2011). "The role of gender in team collaboration and performance." *Interdisciplinary science reviews*, *36*(2), 146–153.

Biddle, Justin B, Ian James Kidd, and Anna Leuschner (2017). "Epistemic corruption and manufactured doubt: The case of climate science." *Public Affairs Quarterly*, *31*(3), 165–187.

Biddle, Justin B and Anna Leuschner (2015). "Climate skepticism and the manufacture of doubt: Can dissent in science be epistemically detrimental?." *European Journal for Philosophy of Science*, *5*, 261–278.

Bright, Liam Kofi and Remco Heesen (2023). "To be scientific is to be communist." *Social Epistemology*, *37*(3), 249–258.

Brown, Matthew J (2020). *Science and moral imagination: A new ideal for values in science*. University of Pittsburgh Press.

Coates, Matthew (2024). "Does it Harm Science to Suppress Dissenting Evidence?." *Philosophy of Science*.

Cook, John (2017). "Response by Cook to "beyond counting climate consensus"." *Environmental Communication*, *11*(6), 733–735.

Melo-Martin, Inmaculadade and Kristen Intemann (2014). "Who's afraid of dissent? Addressing concerns about undermining scientific consensus in public policy developments." *Perspectives on Science*, *22*(4), 593–615.

Melo-Martín, Inmaculadade and Kristen Intemann (2018). *The fight against doubt: How to bridge the gap between scientists and the public.* Oxford University Press.

Derex, Maxime and Robert Boyd (2016). "Partial connectivity increases cultural accumulation within groups." *Proceedings of the National Academy of Sciences*, *113*(11), 2982–2987.

Desmond, Hugh (2023). "The ethics of expert communication." *Bioethics*, *38*(1), 33–43.

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

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

Douglas, Heather (2018). "From tapestry to loom: Broadening the perspective on values in science.".

Douglas, Heather (2021). "The rightful place of science: Science, values, and democracy.".

Douglas, Heather, Kevin Elliott, Andrew Maynard, Paul Thompson, and Kyle Whyte (2014). "Guidance on Funding from Industry." *SRPoiSE. org. http://srpoise. org/wp-content/uploads/2014/06/Guidance-on-Funding-from-Industry-Final. pdf.*

Elliott, Kevin Christopher (2017). *A tapestry of values: An introduction to values in science.* Oxford University Press.

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.

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

Frey, Daniel and Dunja Šešelja (2018). "What is the epistemic function of highly idealized agent-based models of scientific inquiry?." *Philosophy of the Social Sciences*, *48*(4), 407–433.

Frey, Daniel and Dunja Šešelja (2020). "Robustness and idealizations in agent-based models of scientific interaction." *The British Journal for the Philosophy of Science*.

Fricker, Miranda (2007). *Epistemic injustice: Power and the ethics of knowing*. OUP Oxford.

Gerken, Mikkel (2018). "Expert trespassing testimony and the ethics of science communication." *Journal for General Philosophy of Science*, *49*, 299–318.

Gerken, Mikkel (2022). *Scientific Testimony: Its roles in science and society*. Oxford University Press.

Grim, Patrick, Daniel J Singer, Steven Fisher, Aaron Bramson, William J Berger, Christopher Reade, Carissa Flocken, and Adam Sales (2013). "Scientific networks on data landscapes: Question difficulty, epistemic success, and convergence." *Episteme*, *10*(4), 441–464.

Holman, Bennett and Justin Bruner (2017). "Experimentation by industrial selection." *Philosophy of Science*, *84*(5), 1008–1019.

Holman, Bennett and Justin P Bruner (2015). "The problem of intransigently biased agents." *Philosophy of Science*, *82*(5), 956–968.

Intemann, Kristen (2023). "Science communication and public trust in science." *Interdisciplinary Science Reviews*, *48*(2), 350–365.

Intemann, Kristen (2024). "Value transparency and promoting warranted trust in science communication." *Synthese*, *203*(2), 42.

Intemann, Kristen and Inmaculada de Melo-Martín (2014). "Are there limits to scientists' obligations to seek and engage dissenters?." *Synthese*, *191*, 2751–2765.

John, Stephen (2018). "Epistemic trust and the ethics of science communication: Against transparency, openness, sincerity and honesty." *Social Epistemology*, *32*(2), 75–87.

Kitcher, Philip (2011). "Science in a democratic society." *Scientific realism and democratic society*. . Brill, 95–112.

Korf, Rebecca (2024). "Taking the Social Structure of Science Seriously when Debating Value Influence.".

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

Leuschner, Anna (2018). "Is it appropriate to 'target'inappropriate dissent? On the normative consequences of climate skepticism." *Synthese*, *195*, 1255–1271.

Lewandowsky, Stephan, Toby D Pilditch, Jens K Madsen, Naomi Oreskes, and James S Risbey (2019). “Influence and seepage: An evidence-resistant minority can affect public opinion and scientific belief formation.” *Cognition*, *188*, 124–139.

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

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

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.

Merton, Robert K (1942). “A note on science and democracy.” *J. Legal & Pol. Soc.*, *1*, 115.

Mohseni, Aydin, Cailin O’Connor, and James Owen Weatherall (2022). “The Best Paper You’ll Read Today.” *philosophical topics*, *50*(2), 127–153.

Mohseni, Aydin and Cole Randall Williams (2021). “Truth and conformity on networks.” *Erkenntnis*, *86*, 1509–1530.

Nash, Erin J (2018). “In Defense of “Targeting” Some Dissent about Science.” *Perspectives on Science*, *26*(3), 325–359.

O’Connor, Cailin (2023). “The Public Cost of Private Science.” *Nautil.us*.

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

Oreskes, Naomi (2017). “Response by Oreskes to “Beyond counting climate consensus”.” *Environmental Communication*, *11*(6), 731–732.

Paetkau, Tyler (2024). “Inductive risk and epistemically detrimental dissent in policy-relevant science.” *European Journal for Philosophy of Science*, *14*(1), 1.

Phillips, Katherine W and Evan P Apfelbaum (2012). “Delusions of homogeneity? Reinterpreting the effects of group diversity.” *Looking back, moving forward: A review of group and team-based research.* . Emerald Group Publishing Limited, 185–207.

Pinto, Manuela Fernández and Daniel Fernández Pinto (2023). “Epistemic diversity and industrial selection bias.” *Synthese*, *201*(5), 182.

Rolin, Kristina (2015). “Values in science: The case of scientific collaboration.” *Philosophy of Science*, *82*(2), 157–177.

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

Rudner, Richard (1953). "The scientist qua scientist makes value judgments." *Philosophy of science*, *20*(1), 1–6.

Solomon, Miriam (2001). *Social empiricism*. MIT press.

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

Steel, Daniel (2017). "Qualified epistemic priority: Comparing two approaches to values in science." *Current controversies in values and science.* . Routledge, 49–63.

Torsell, Christian and Rebecca Korf (2024). "Network Study on Open Sharing.".

Weatherall, James Owen and Cailin O'Connor (2021). "Conformity in scientific networks." *Synthese*, *198*(8), 7257–7278.

Weatherall, James Owen, Cailin O'Connor, and Justin P Bruner (2020). "How to beat science and influence people: Policymakers and propaganda in epistemic networks." *The British Journal for the Philosophy of Science.*

Wu, Jingyi (2022). "Withholding Knowledge." *Department of Logic and Philosophy of Science, University of California at Irvine: Irvine, CA, USA*.

Wu, Jingyi (2023). "Epistemic advantage on the margin: A network standpoint epistemology." *Philosophy and Phenomenological Research*, *106*(3), 755–777.

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." *Science and other cultures.* . Routledge, 26–48.

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

Zollman, Kevin JS (2009). "Optimal publishing strategies." *Episteme*, *6*(2), 185–199.

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