Social epistemology of science as theoretical social science

Preprint, not peer reviewed.

Authors in alphabetical order:

Inkeri Koskinen, Faculty of Social Science, University of Helsinki

inkeri.koskinen@helsinki.fi ORCID: 0000-0002-9060-7011

Jaakko Kuorikoski, Faculty of Social Science, University of Helsinki

jaakko.kuorikoski@helsinki.fi ORCID: 0000-0001-5676-717X

Renne Pesonen, Faculty of Social Sciences, Tampere University

renne.pesonen@tuni.fi ORCID: 0000-0001-6425-5772

Samuli Reijula, Faculty of Humanities, University of Helsinki

samuli.reijula@helsinki.fi ORCID: 0000-0001-6968-5819

Kristina H. Rolin (corresponding author), Faculty of Social Sciences, Tampere University

kristina.rolin@tuni.fi ORCID: 0000-0003-3893-7184

Abstract: Naturalistic social epistemology has gone mainstream, but two of its central methods, case studies and modelling, are poorly understood. Even though case studies and modelling can be seen as two distinct approaches to social epistemology, both can be seen as mutually complementary methods for investigating social mechanisms relevant to the epistemic functioning of science. We argue, therefore, that naturalistic social epistemology of science should be seen as continuous with theoretical social science.

Keywords: Methods in Philosophy of Science, Methods in Social Epistemology, Case Study Methods, Modelling Methods, Naturalistic Approaches to Philosophy of Science

1. Introduction

Philip Kitcher has defined the social epistemology of science as the study of the properties of epistemically well-designed social systems (1993, p. 304). Scientific communities are epistemically well designed when their social practices and institutional arrangements promote the epistemic goals of science, including problem-solving (Kuhn, 1962), understanding, the discovery of significant truth (Kitcher, 1993) and empirical success (Solomon, 2001). During the last four decades, this understanding of social epistemology, which has grown out of a naturalized philosophy of science (Giere, 1985, 1989; Kitcher, 1992), has become an established research program in philosophy of science (Longino, 2022). Even in its moderate versions, a naturalistic approach to the social epistemology of science makes substantial use of methods and resources not usually seen in the pages of philosophy journals, such as empirical studies of science, empirical methods, and computational modelling, to understand whether and how social practices and features of social systems in science promote the epistemic goals of science (or lead to epistemic failures). This is in sharp contrast to more traditional analytic social epistemology concerning the transference of justification in testimony, the rational response to peer disagreement, as well as the status of shared beliefs or social groups as epistemic agents (see e.g., Goldman & Whitcomb, 2011). Such, more conceptual questions are rightfully addressed in terms of conceptual analysis, historical case examples, thought experiments, and reflective equilibrium (Goodman, 1965).

Given that a naturalistic orientation to the social epistemology of science has gone mainstream, there is surprisingly little systematic discussion of the methods that can be used to examine whether and how social practices and features of social systems in science promote epistemic goals. While case studies and computational models are widely used in the social epistemology

literature, it is not always clear what purposes these methods are meant to serve, what their limitations are, and whether and how they can be combined into a mixed-methods approach. Also, as the use of computational models and detailed socio-historical examinations of cases seem to be methodologically closer to social sciences than traditional philosophy understood as conceptual analysis (and the application of other a priori methods), social epistemologists would benefit from a discussion of how these methods have been theorized in the philosophy of social sciences.

Our aim in this paper is to begin a more systematic discussion of the use of case studies and modelling methods in the social epistemology of science. If social epistemologists adopt a naturalistic approach to the social epistemology of science as the study of social systems, the choice of methods can and should be justified by arguing that they help understand which features of social systems of science promote the epistemic goals of science. These epistemic properties are descriptively and normatively discontinuous from the epistemic statuses of individuals (Mayo-Wilson et al., 2011), and importantly, they are not analytical implications of epistemological concepts and normative principles. We argue that empirical methods are used in the social epistemology of science in the same way as they are used in theoretical social scientific inquiry understood as a study of social mechanisms, that is, causal schemes theorized as "abstract sketches of causal configurations that can be adapted and combined to serve as parts of causal scenarios" (Ylikoski 2019, 16). In theoretical social science, both case studies and modelling methods can be used to examine social mechanisms. The social mechanisms studied here are ones that mediate between particular features of social practices in science and their epistemic effects. We argue that as both case studies and models can be used to study social mechanisms, they can also be combined into a mixed-methods approach.

While method debates are not novel in philosophy of science (Hangel & ChoGlueck, 2023; Veigl & Currie, 2025; Wagenknecht, Nersessian & Andersen, 2015), a systematic focus on the use of empirical methods in the social epistemology of science is timely for two reasons. First, much of the recent work in the field is characterized by the use of methods not usually seen in philosophy or even in other areas of philosophy of science, e.g., agent-based computational models. As social epistemology is concerned with the properties of a social system, what is under analysis are the systematic consequences of interactions within some specific social and institutional organization. Often this involves tracking not only intended but also unintended consequences of interdependent epistemic actions under specific systems of incentives, resource allocations, and possibilities of communication. Such claims cannot rely merely on philosophers' intuitions. One reason for this is that reasoning about the consequences of social interaction exhibiting feedback and subject to multiple simultaneous constraints is next to impossible without external inferential aids, i.e., models (Reijula & Kuorikoski, 2019). Another reason is that even when such theorizing is highly general and abstract, the hypotheses concern causal dispositions of complex systems (not conceptual truths), and these hypotheses are ultimately empirical in nature. Social epistemology of science, thus understood, therefore requires an empirical component: the study of the causal properties of social systems - including the unintended consequences arising from the interrelated functioning of various social mechanisms.

Second, while case studies and case examples have been widely used in philosophy of science, in the social epistemology of science they are used in ways that differ from their role in more traditional philosophy of science – or this is what we will argue in this paper. In contrast to the recent emergence of agent-based computational models in social epistemology, the use and the methodological reflection concerning historical case studies in philosophy of science has a long

pedigree. Although Reichenbach (1938) tried to categorically demarcate genuine epistemological questions about the necessary relations between empirical evidence and hypotheses from the historical and psychological contingencies related to the discovery of hypotheses, the latter half of the 20th century witnessed an institutionalized growth of attempts at using historical case-based evidence for the justification of normative claims about the method of science. The hope was, very roughly put, that the *right rules of ampliative inference could be read off from a careful look at the successful historical episodes of scientific progress*. Or at least that such episodes should be taken into the reflective equilibrium ultimately *justifying* these rules. Such justificatory ambitions meant that the historical episodes were heavily rationally reconstructed - much to the irritation of sociologists and historians of science.

A classic example is Popper's (1974) analysis of the discovery of Neptune, where astronomers, instead of rejecting Newtonian mechanics when Uranus's orbit deviated from predictions, hypothesized the existence of an unknown planet, which was later discovered. Popper viewed this as a prime example of a critical test. Lakatos (1970) later challenged the idea of critical tests by referring to the same example. He argued that Newtonian mechanics would not have been abandoned even if Neptune had not been found.

Another example of a rational reconstruction is Hempel’s (1966) examination of how Semmelweis successfully used ampliative reasoning on the cause of childbed fever by testing and rejecting several hypotheses. Since then, the same example has been referred to numerous times, especially in discussions about inference to the best explanation (e.g. Lipton, 2004; Bird, 2010). The example has become a kind of fixed point in reflective equilibrium: all participants in the discussion accept it as a good example of what their philosophical theories are about, and

those involved in the discussion defend their own stance by demonstrating how they can highlight something important in the familiar example.

Now, contrast this traditional aim of discovering the right rules of ampliative reasoning to the one often held in contemporary social epistemology of science. Social epistemologists aim to assess whether some features of social systems in science are epistemically rational while some others are not, where epistemic rationality is understood as instrumental rationality, i.e., using the means that are not only believed to be, but are in fact conducive to achieving some predefined desired goals (Giere, 1989, p. 380). Given this understanding of social epistemology, epistemic rationality is a matter of employing the means that are conducive to achieving desired epistemic goals, such as significant truth, empirically adequate theories or understanding. Whereas the traditional use of case studies aims to discover the right rules of ampliative *reasoning*, contemporary social epistemology of science is concerned with a *potentially broad set of social practices and institutional arrangements* that promote the epistemic goals of science. Case studies and case examples are seen as epistemically interesting insofar as they shed light on causal processes that lead to epistemically desirable outcomes.

In sum, a naturalistic approach to the social epistemology of science consults empirical studies of science or uses empirical methods to understand which properties of social systems promote the epistemic goals of science. It recognizes that not all questions in philosophy of science are answerable with the same methods and that the proper understanding of mixed methods in social epistemology requires understanding what kind of questions these methods are expected to answer. By analysing the use of empirical methods, we argue that a naturalistic social epistemology of science is continuous with theoretical social science.

In what follows, we analyse how the naturalistic program has been implemented in the social epistemology of science by using case studies (Section 2) and agent-based computational models (Section 3). We show that the purposes for which case studies and computational models are used are not as different as they are often assumed to be. We argue that case study methods and modelling methods can be combined into a mixed-methods approach to study social mechanisms that are of interest to the social epistemology of science (Section 4). Our analysis will also shed light on some recent controversies over the epistemic functions that these methods have in philosophy of science.

2. Case studies in the social epistemology of science

In this section, we examine the role that case studies play in the kind of social epistemology that investigates the epistemic effects of social practices or systems of science. To understand this role, it is necessary to have a working definition of a case study at hand. In what follows, we adopt a broad conception of case studies, ranging from case studies as they are understood in social sciences to case examples (or vignettes). At one end of this conception, one can find a s*ocial scientific case study* as it is theorized by Morgan (2012). According to Morgan (2012), a social scientific case study is characterized by five features: it investigates a bounded whole object of analysis, examines the object in its context without a clearly defined boundary between the object and its context, involves deep engagement and rich evidential material, applies mixed methods, and presents results in a narrative that reflects the complexity of the topic. Ylikoski and Zahle (2019) further highlight that a social-scientific case study can be a conceptualisation of a naturally occurring item or process as a case of something, an example of a more general phenomenon under study. At the other end of our broad conception of case study, one can find

what we call *case examples*: stylised descriptions of historical or contemporary cases based on existing literature, not on philosophers' original empirical research. Social scientists, especially economists, would call these vignettes. Descriptions of cases in philosophy of science today often fall somewhere on the continuum between these two extremes. In the literature on the use of case studies in philosophy of science, all these different ways of producing descriptions of cases are typically called case studies (Grasswick & McHugh, 2021).

In the rest of the section, we briefly review the debate on the so-called dilemma of case studies in the philosophy of science and discuss how philosophers of social science have analysed the epistemic functions of case studies in the social sciences. Then we analyse their use in recent social epistemology of science. We argue that it is the philosophical literature on the epistemic functions of case studies in social sciences that properly addresses important issues when cases are used in the social epistemology of science. Case studies can be useful for the purpose of theorising social mechanisms that are of epistemic interest. As such, case studies should be viewed as instances of *process tracing*. Nevertheless, we also emphasize that other kinds of evidence are needed for generalizable causal claims, and that case studies alone are a poor evidential basis for practical recommendations.

2.1. The dilemma of case studies in philosophy of science

The use of case studies – broadly understood – has been debated especially by those philosophers of science who have been inspired by the historical turn initiated by Kuhn (1962). At the centre of these debates has been the question of whether historical case studies (or, in our terms, often case examples) of science can be used to provide inductive support for philosophical accounts of scientific knowledge. This question is often framed as the dilemma of case studies: "On the one hand, if the case is selected because it exemplifies the philosophical point, then it is

not clear that the historical data hasn't been manipulated to fit the point. On the other hand, if one starts with a case study, it is not clear where to go from there—for it is unreasonable to generalize from one case or even two or three." (Pitt 2001, p. 373). One horn of the dilemma is cherry-picking of data for empirical generalizations, whereas the other one is weak inductive reasoning.

In response to this dilemma, some philosophers of science argue that case studies should not be used as evidence for inductive reasoning, but that they serve other purposes like rhetorical or pedagogical ones. Currie (2015, p. 544) refers to this limitation as the "curse" of the case study: if case studies are interpreted as evidence, they are a poor basis to build inductive generalizations about science. Moreover, scientific practices are heterogeneous, so it is not clear that philosophers can justify the use of small sample sizes (Currie, 2015, p. 557). Instead, Currie proposes that case studies are "illustrative of conceptual tools" (p. 560) that serve to anchor philosophical analysis. Philosophers can use case studies to illustrate how one can use conceptual tools, and this amounts to non-inductive epistemic justification for a case study approach (p. 565).

Another response to the dilemma of case studies is proposed by Chang (2012). While he acknowledges that philosophers cannot leave behind inductive reasoning entirely, he suggests that the role of case study in philosophy of science is recast as a relationship between the concrete and the abstract, rather than between the particular and the general. Instead of treating a case study as a single data point for inductive reasoning, a case study can be seen as the concrete that can have a two-way relationship with the abstract (2012, p. 110). Philosophers can use abstract ideas to understand a concrete episode in the history of science or contemporary science, and they can learn from examining concrete cases and use this knowledge to revise their abstract

ideas. As Chang puts it: “If we extract abstract insights from the account of a specific concrete episode that we have produced ourselves, that is not so much a process of generalization, as an articulation of what was already put into it” (2012, p. 110).

We do not deny that case studies are valuable as the illustration of conceptual tools (as Currie claims), or that they have both important heuristic and analytic roles in history and philosophy of science (as Chang claims). However, these accounts do not capture all of the ways in which cases are used in the social epistemology of science today. In a naturalized social epistemology of science, the challenge is to understand under what conditions some scientific practices and social arrangements are conducive to desired epistemic goals. To meet this challenge, the use of case studies in the social epistemology of science has more in common with the use of case studies in social sciences.

In the remaining part of this section, we argue that social epistemologists do in fact use case studies in similar ways to social scientists: to understand the effects of scientific practices and social arrangements in science. More specifically, we argue for a legitimate use of cases in the social epistemology of science: case studies (in the broad sense) can help philosophers theorize about social mechanisms that mediate between scientific practices and social arrangements, on the one hand, and desired epistemic goals, on the other. This understanding of the use of cases will also suggest new solutions to the dilemma of case studies.

2.2. Lessons from philosophy of the social sciences

While social scientists are interested in explaining why something happened or why some state of affairs prevails, social epistemologists aim to answer the question of why some features of social practices or systems in science are epistemically rational while some others are not. As we argued in Section 1, epistemic rationality is understood as instrumental rationality which consists

in employing means that are not only believed to be but are in fact conducive to achieving desired epistemic goals (Giere, 1989, p. 380). In other words, epistemic rationality of a social system is a matter of it employing means that are conducive to achieving desired epistemic goals, such as significant truth, empirically adequate theories or understanding. In this section, we discuss the question of whether case studies can be used to answer the question of why some practices and social arrangements in science are epistemically rational and how they are conducive to desired epistemic goals. To do so, we turn to philosophy of the social sciences and review debates on the use of case studies in the social sciences.

Philosophical debates on the use of cases in social sciences have focused on two questions: whether case studies can provide evidence for causal claims, and whether and under what conditions they can support empirical generalizations. The former is a necessary condition for investigating the epistemic effects of social practices and the latter for applying the learned lessons for policy advice.

Like philosophers concerned with the dilemma of the case study, philosophers of social sciences are critical of the view that case studies can serve as evidence for empirical generalizations. Crasnow (2021) argues that in social sciences, case studies rarely aim to provide inductive support for empirical generalizations or theories. Even if case studies provide such support, it is weak, at best. Moreover, case studies are rarely used for the purpose of testing theories. While in principle a case study can provide counterevidence to a theory, it does not necessarily falsify the theory. In some cases, it may only restrict its domain. If case studies are used to test theories, the test is weak, at best. This is because a case study can tell that something is wrong with the theory, but it does not tell us what should be revised (Crasnow, 2021, p. 76).

Along similar lines, Ylikoski (2019) distinguishes between theory-mediated and empirical generalizations. Whereas empirical generalizations proceed from one or few cases to many cases, and hence, are problematic, theory-mediated generalizations proceed from one or few cases to social mechanisms that are of interest in other contexts too. Knowledge about social mechanisms is valuable because it can be used to construct explanations, i.e. to provide answers to how-questions that underlie causal why-questions. By contributing to the understanding of social mechanisms, case studies can reveal how causes actually produce their effects by describing the process by which this happens (Ylikoski, 2019, p. 16).

Ylikoski's (2019) account brings us to the other debate on the use of case studies in social sciences: the ability of case studies to provide evidence for causal claims. While not all case studies aim to provide evidence for causal claims (e.g., they may just aim to provide a comprehensive and rich description of their object), some case studies are used to support such claims. The literature on *process tracing* argues that case studies can provide evidence for causal claims but not in the same way as quantitative methods do (Crasnow, 2017). There are a number of competing accounts of process tracing, but most argue that while quantitative methods can produce evidence for general causation, case studies can generate evidence for singular causation. However, the evidence provided by case studies is highly dependent on theoretical background assumptions (Crasnow, 2012). Given a theoretical framework, it is possible for social scientists to make inferences about what would happen under some interventions. As Crasnow puts it, "Process tracing aims at giving a singular causal account of the event but can only do so against the backdrop of a general causal account" (2012, p. 664).

To summarize, philosophers of social sciences argue that while case studies can serve many purposes, they serve some purposes better than others (see also Ylikoski & Zahle, 2019, p. 2).

Case studies are apt for understanding new phenomena in depth rather than for justifying empirical generalizations. They are more suitable for identifying singular causation than for generating evidence for general causation. Most importantly, *case studies can be used to identify social mechanisms or processes that are relevant to a general account of social phenomena even though they are insufficient for empirical generalizations*. This means that there is an alternative to the dilemma of the case study. The use of case studies does not need to aim for empirical generalizations (and does not necessarily involve cherry-picking of data or weak inductive reasoning), if case studies are seen as evidence for theory-mediated generalizations aiming to illustrate social mechanisms.

2.3. A new way of using cases in philosophy of science: an example

Let us now turn to the use of cases in contemporary social epistemology of science. While philosophers of science still use cases (and case examples), in the manner prescribed by Currie and Chang – to illustrate conceptual tools and as concrete sounding boards for abstract ideas – cases are also used to study social mechanisms that mediate between the social systems of science and their epistemic effects. Rather than serving as agreed-upon examples for backing up or concretising normative claims about scientific methods, cases are used to trace and illustrate epistemically relevant empirical causal properties of social systems.

Good examples of this abound. Miriam Solomon, for instance, uses multiple cases in order to test whether the three conditions for normatively appropriate dissent in science that she has identified actually work as she argues. As she is committed to a naturalized approach to social epistemology and takes scientific rationality to be socially emergent (Solomon, 2001, p. 12), she argues that such testing is necessary. The argument for her account "depends on examination of case studies, which show that satisfaction of the three conditions is linked to scientific success"

(Solomon, 2006, p. 27). Another good example is Bennett Holman's and Sally Geislar's (2018) use of a case in which a pharmaceutical company managed to manipulate a participatory process intended to incorporate patient perspectives into the evaluation of a drug developed by the company. They use the case to argue that the institutional design of participatory programmes should be more resistant against attempts to manipulate the results. In brief, many philosophers of science today draw upon cases when examining the structures of social and institutional processes and organizations, and the systematic consequences of interactions within them, in much the same way as case studies are used in social sciences. As a result, both the benefits and the challenges related to the use of cases in this kind of social epistemology increasingly resemble those social scientists face.

We will now illustrate this empirical function with a more detailed example: Justin Biddle's (2007) use of the case of the painkiller Vioxx. He argues that its acceptance and later withdrawal were epistemically inadequate due to flawed institutional arrangements in privatized pharmaceutical research. He finds existing social epistemological proposals by Kitcher and Longino to be inadequate when trying to address the core problems of the case and emphasizes the need for the social epistemology of science to focus more on institutional structures and power dynamics. Finally, he suggests a way in which pharmaceutical research could be reorganised to avoid similar failures in the future.

The Vioxx case garnered considerable public attention. Biddle (2007) describes how the pharmaceutical company Merck was aware that its drug Vioxx could cause heart problems but did not investigate the matter because the company's marketing officials did not consider the research necessary. It was only after a study on a different question was halted due to emerging

heart symptoms that the drug was withdrawn from the market. It is estimated that Vioxx may have caused the deaths of up to 55,000 users.

Biddle argues that Longino's and Kitcher's theories that focus on the ideal organisation of scientific communities overlook the institutional contexts that allowed such a failure. He claims that the institutional arrangements of commercialised science, i.e. the fact that a pharmaceutical company's marketing officials could influence decisions on research, were one of the root causes of the catastrophe. Moreover, he argues that the numerous ties between experts at the US Food and Drug Administration (FDA) and pharmaceutical companies exacerbated the problem by leading to regulatory practices that favoured the companies. To prevent similar cases in the future, Biddle proposes that the FDA could adopt an adversarial system of research, where industry-sponsored scientists would have to argue on behalf of their companies. Their adversaries would be scientists with no funding from pharmaceutical companies, who in turn would argue on behalf of the public – and the FDA juries would have no ties to any pharmaceutical companies.

Biddle takes the case to be an example of a larger phenomenon, and he uses it to identify a complex social/institutional mechanism that led to an epistemic failure and significant harm. He also claims that the case can be used as evidence of a general institutional problem in pharmaceutical research: "The Vioxx case study, along with numerous case studies developed elsewhere, provides strong reasons for believing that the privatization of the biomedical sciences is epistemically worrisome, and it suggests that the primary response to this situation should be a social, or organizational, one" (Biddle, 2007, p. 21). In other words, he claims that the Vioxx case is an example of a systemic problem in current commercialised pharmaceutical research, and that the mechanism or mechanism type he identified is common. This is an empirical

generalization. While there is no reason to doubt it, a single case is not yet sufficient to support it. Biddle himself recognises this problem by referring to other examples and, among other things, information about the ties the FDA experts have to pharmaceutical companies.

It is not our purpose to engage in a debate on the Vioxx case. In what follows, we analyse how the case is used in the social epistemology of science.

2.4. Identifying social mechanisms

This kind of use of cases clearly differs from the ways in which cases are used in more traditional philosophy of science. While Biddle's case description shares some features with the familiar philosophical case examples (for instance, it is based on existing literature), it is not primarily an example of good or bad scientific reasoning. Nor is it used as a concrete sounding board for abstract ideas, or to illustrate conceptual tools. Instead, the Vioxx case is used to examine a phenomenon that Biddle claims to be common in actual science, and to identify a putative causal mechanism. It is used to explore the systematic, epistemically beneficial or harmful consequences of patterns of interaction within structures of social and institutional organisations. In other words, the use of the case resembles ways in which case studies are used in the social sciences.

Biddle makes causal claims about the role of non-scientific business staff in research and focuses on the institutional structures of industry and regulatory bodies which made the oversight of serious risks possible. This exemplifies well the functions of cases as they are understood in the methodological literature on case studies in the social sciences. Case studies are apt for understanding new phenomena in depth and for identifying singular causal processes (Crasnow, 2021). Moreover, they can be used to identify social mechanisms or processes that are relevant to general accounts of social phenomena (Ylikoski, 2019; Ylikoski & Zahle, 2019).

Philosophers of science who use case studies in this way suggest that cases can serve as evidence for more general claims about the phenomenon of interest. It is important to note, however, that this generalizability can only be achieved through the identification of a mechanism, and by relying on causal background knowledge supporting the generalizability of this mechanism (Ylikoski, 2019). The evidential value of case studies does not reduce to simple inductive support from a sample size of one. Nevertheless, it is also important to note that the value of the practical recommendations – for instance, policy recommendations – one can make on the basis of case studies depends on two things. First, it depends on whether the identified mechanism is actually common. Only in that case can the recommendation even in principle be conducive to achieving the desired goal. Second, even if the identified mechanism is common, trying to implement such a recommendation in the real, complex, social and institutional settings of science could be too costly or it could create unwanted side effects. Biddle is well aware of these limitations, and emphasises that his policy proposal needs further development, and that its ultimate test would be "whether it can be articulated in a clear fashion, whether it can be implemented, and whether it improves the quality of pharmaceutical research" (Biddle, 2007, p. 35).

In sum, we have argued that the use of case studies and case examples in the social epistemology of science comes methodologically closer to their use in social sciences than their use in traditional philosophy of science. In the social epistemology of science, cases are not merely used to illustrate the use of conceptual tools or to facilitate a back-and-forth movement between the abstract and the concrete. They are used to identify and study social mechanisms that mediate between the social systems of science and their epistemic effects. Next, we turn to the role of

agent-based models in the social epistemology of science. We argue that their function is surprisingly similar.

3. Modelling and simulation methods in the social epistemology of science

Although philosophy of science has always utilized mathematical modelling in a broad sense, such as Bayesian networks in formal epistemology (Bovens & Hartmann, 2003), current modelling in the social epistemology of science is characterized by the use of simulation models and agent-based models in particular. The literature is populated with simple simulated agents exploring their environment, carrying out experiments and communicating their results to other agents. As agent-based simulations are by far the most popular modelling technique in social epistemology, we limit our discussion to them.

Typically, the goal of these computational exercises is not to normatively evaluate the rationality of the behavioural rules of single agents, but to assess which forms of social interaction and institutional conditions best facilitate the attainment of some predetermined epistemic outcomes, such as maximizing "epistemic value" of research or convergence to "the best" theory among available options. In the abstract, these models are demonstrations of the *independence thesis* (Mayo-Wilson, Zollman & Danks, 2011), according to which collective epistemic rationality, understood as instrumental rationality, is not a simple aggregate of individual rationality. In particular, the models investigate causal hypotheses about the effects of alternative forms of social organization of cognitive labour.

In this section, we first introduce an example of an agent-based model and then argue that the most promising way of interpreting the epistemic role of such models is to view them as tools for

exploring social mechanisms understood as causal schemes. When models are understood as tools for theoretical exploration and argumentation, their epistemic role is not very different from the role that case studies are thought to have in social sciences and in the social epistemology of science.

3.1. Epistemic networks: An example of agent-based modelling of scientific inquiry

To give a quick example, one of the best-known and widely applied classes of models is the epistemic network model introduced by Zollman (2007) and generalized in Zollman (2010). Epistemic networks consist of a collection of agents, often formalized as Bayesian learners who are sensitive to evidence gathered from experimental outcomes. In each simulation round, the agents choose between two options, A or B, and observe either a successful or unsuccessful outcome based on the predetermined payoff probability of each option. Simulation rounds are conceptualized as episodes of gathering evidence for theory A or B. The agents aim to reap the rewards of adopting the better theory; however, they do not know which one is superior. At the start of the simulation, each agent is randomly assigned a preferred theory, and they switch if the evidence accumulated during the successive simulation rounds supports the alternative option. The agents monitor the experimental outcomes gathered by their network neighbours, and they base their theory choice both on their own information as well as information they obtain from their neighbours.

The question that interested Zollman (2007) was how the network structure of agents affects the epistemic performance of the community. The worse theory may outcompete the better one, if misleading early evidence supports the worse theory and the whole community locks onto it. Zollman found that densely connected networks are susceptible to such lock-in, and that the problem may be mitigated by restricting the information flow between agents. This result is

known as *the Zollman effect*. Moreover, Zollman (2010) suggested that the effect is not solely an abstract theoretical possibility. He argued that his model may actually explain why the research on peptic ulcer disease was locked on a wrong theory for decades, even though the correct theory about the bacterial etiology of the disease was well-known but not pursued during that period. Specifically, the misleading evidence gathered in the early 1950s, suggesting that no bacteria live in the human stomach, was distributed too quickly and widely, leading virtually the entire research community to abandon the correct theory.

Many of the assumptions in Zollman's model can be altered or relaxed. For example, Alexander (2013) investigated reinforcement learning instead of the Bayesian learning rule and a dynamic rather than a fixed network structure. In some models, agents explore alternative theoretical options in addition to their preferred one (Kummerfeld & Zollman, 2016). Some studies also investigate other mechanisms that affect theory selection of agents besides experimental knowledge, such as critical interaction and rational inertia, which refers to the tendency of scientists to retain their current hypotheses despite the mounting counterevidence (Frey & Šešelja, 2018).

Follow-up studies have found that the Zollman effect is not robust across parameter values but it depends on specific modelling assumptions. Furthermore, several subsequent modifications to Zollman's original model (e.g., Kummerfeld & Zollman, 2016; Frey & Šešelja, 2018) suggest that under various modelling conditions, more communication can be epistemically beneficial, contrary to Zollman's original proposal. However, the general point argued by Zollman (2010) is widely accepted that research benefits from *some* mechanisms that secure the diversity of opinions until sufficient evidence is gathered about all the theoretical options available.

In addition to epistemic network models, there are many other formal frameworks that are used for simulating scientific inquiry, and many other questions have been studied besides the convergence to the best theory. For example, some models investigate how diversity may, or may not, trump ability in collective problem solving (Hong & Page, 2004; Reijula & Kuorikoski, 2021) or how the tendency of some scientists to prefer novel approaches over already proven ones benefits the research community (Weisberg & Muldoon, 2009; Pöyhönen, 2017). In the next section, we analyse the epistemic uses of models in the social epistemology of science.

3.2. Uses and functions of agent-based simulations in social epistemology

We highlight two important features that are common to agent-based models in the social epistemology of science. First, their use resembles more theoretical social science than traditional philosophical analysis in that modelling aims to uncover complex, unintuitive and unintended causal consequences of social interactions (Epstein, 2008; Page, 2018). These models do not aim at analysing what the epistemic goal of inquiry is, but instead, which social factors facilitate its attainment. The resemblance between modelling in social epistemology and in social sciences is not surprising. For instance, Zollman's (2007) model was inspired by economics (Bala & Goyal, 1998), models similar to Weisberg and Muldoon's (2009) epistemic landscape have been widely utilized in social sciences (Gerrits & Marks, 2014), and Hong and Page's (2004) diversity model was developed for the purposes of organizational research.

Second, what is common to agent-based simulations of scientific inquiry is that they are highly idealizing. The actions and decision-rules of individual scientists are extremely simple, and the focus of simulation studies lies in the emergent macro-scale effects produced through agent interactions. While the representational distance between the models and their intended real-world targets remains vast, the general aim of these models is precisely to isolate, simplify, and

control selected aspects of hypothetical phenomena of interest. While this does not directly inform us about any actual real-world system, the computational investigations into the behaviour of model systems can provide causal hypotheses on real-world processes.

A broadly shared starting point in the literature on modelling in science is that models are surrogate systems that enable an indirect research strategy, where an object that is too big or inconvenient to study directly is investigated indirectly by interacting with the model system (Weisberg, 2013). In the first stage of modelling, a model system is constructed. Then, the researcher studies the model system in order to generate findings. Finally, these findings are interpreted as informing us about the target system.

Much of the philosophical debate on models concerns the following questions: How is it possible to learn something about the target by studying something else, the model? What kind of knowledge can models offer in general, and what is the epistemic function of such quasi-empirical models in philosophy in particular? Does modelling teach us something about the putative target systems (e.g., a particular scientific community or research episode), or should we instead see models as informing us about something more general, such as the epistemic properties of social mechanisms operating in scientific communities?

To answer these questions, we examine how models can offer general knowledge about social mechanisms that mediate between features of social systems and their epistemic effects. In the next section, we introduce conceptual tools from the methodological discussion in social sciences, especially the distinction between causal scenarios and causal schemes (Ylikoski & Aydinonat, 2014; Ylikoski, 2019; see also Frey & Šešelja, 2018).

3.3. The epistemic role of models: from causal scenarios to causal schemes

A typical answer in the literature to the question of how models produce knowledge, resorts to the representation relation between a model and its target. If a model is an accurate or otherwise good representation of its target, that representation relation makes it possible that experimentation with a model can inform us about the properties of the target system. Models, or more specifically, simulation runs, can be seen as simulated causal scenarios. Multiple runs can be aggregated and key variables manipulated, and the resulting set of outputs can be analysed quantitatively like actual data. This approach to seeing models as experiments on surrogate systems has been defended by several philosophers (e.g., Mäki, 2009; Morgan, 2005). It is also at least an implicit assumption in many modelling studies in social epistemology, and sometimes defended explicitly (e.g., Fazelpour & Steel, 2022).

However, a common argument against the usefulness of models in social epistemology concerns their lack of realism (Thicke, 2019; Bedessem, 2019; Martini & Pinto, 2016). The models used, such as network epistemology models (Zollman, 2007) and epistemic landscape models (Weisberg & Muldoon, 2009), depict very simple scenarios, and no empirical data are used in model construction or validation. In response to this criticism, we argue that the epistemic utility of models does not reside merely in the accurate representation of causal scenarios, which they do not even attempt to achieve. Even if these models cannot provide realistic simulations of real-world events, this does not make them epistemologically useless. In our view, the epistemic utility of models depends on their ability to provide understanding on how hypothetical social mechanisms may function in general. For that purpose, highly idealized theoretical models can be more useful than detailed simulations.

This understanding of the epistemic role of models can be found in analyses of theoretical models in the philosophy of social sciences. According to Ylikoski and Aydinonat (2014), the

main function of theoretical models in social sciences is to provide *causal mechanism schemes*. These schemes do not depict any specific causal scenarios, and the models that investigate them need not be compatible with evidence about any particular causal history. The purpose of theoretical models is to generate how-possibly explanations by demonstrating that some particular effect could in principle be produced by the mechanism investigated. If one aims to trace how an actual case or scenario unfolds, possible causal mechanisms need to be identified. Furthermore, Ylikoski and Aydinonat argue that the search for causal explanations in social sciences often has the form of eliminative induction, that is, it aims to eliminate as many of the competing scenarios as possible (2014, p. 26). In their view, the list of possible explanations also itemizes a set of mechanisms that potentially should be ruled out when one seeks to explain a particular event.

The value of computational modelling is most easily discernible in cases where unexpected and complex system-level consequences emerge from plausible assumptions concerning micro-level mechanisms. For example, Zollman's (2007) modelling results reveal that, contrary to our intuitions, the flow of information in scientific communities may not always be epistemically beneficial. Therefore, models are particularly useful for us boundedly rational agents, because they help us derive implications of assumptions that we could not reliably uncover merely by verbal theorizing. In this sense modelling could even be seen as a philosophical tool par excellence (see Mayo-Wilson & Zollman, 2021).

In sum, we have argued that by helping to us identify causal schemes and by offering insight into the functioning of hypothetical causal mechanisms, models can be epistemically valuable even when they do not track any specific real-world scenarios. However, even then, the model may be

epistemically useful as it may have isolated an important mechanism that is in play in many contexts, although rarely observed in its pure form in the real world (Mäki, 2020).

3.4. Theoretical exploration and argumentation with models

Thus far we have argued that the epistemic value of highly abstract theoretical models arises from their capacity to illuminate causal schemes rather than to account for detailed causal scenarios. Furthermore, the way models shed light on causal schemes is by functioning as tools for thinking and argument-making (Aydinonant et al., 2021; Pesonen, 2024). Even if it is not possible to get genuinely new empirical information about targets by modelling them, this does not mean that models are epistemically useless, far from it.

The way we understand the epistemic set-up is as follows. In constructing a model, the modeler makes and formalizes a set of assumptions about some target phenomena. In models in the social epistemology of science, such assumptions typically concern the social and cognitive mechanisms in the epistemic community, as well as some initial conditions. For example, in Zollman's (2007) epistemic networks, the key assumptions were that scientists are rationally sensitive to evidence, which was implemented by formalizing the agents as Bayesian learners, and that scientists gather evidence through social transmission as well as experimenting. By running the model and analysing it, the modeler uncovers further claims that are entailed by the assumptions. In social sciences and social epistemology, such entailments typically pertain to macro-level dynamics in the modelled population. For example, Zollman's (2007) epistemic networks make it possible to explore a hypothetical causal mechanism scheme between different patterns of information flow and the epistemic performance of the community. These features of modelling connect with our discussion on the uses of case studies above: *Modelling is a tool for*

the theoretical development of the mechanistic hypotheses putatively exemplified in empirical case studies.

Models boost the reasoning of individual modelers, but every inference from a theoretical model to any real-world phenomenon is inductive, as the model may misrepresent or omit relevant variables and mechanisms. Instead of being tools for demonstrative inference, models serve the argumentative functions of making our reasoning transparent and convincing others (Epstein, 2008). This aspect of modelling helps see why starting from a dyadic representation relationship between a model and its targets provides a poor starting point for model appraisal. From such a viewpoint, every abstraction and idealization appears as a potential weakness. However, when the inferential or argumentative *goal* of a model is clearly in view, such appraisal is more feasible. Insofar as models are adequate, they are adequate for a particular purpose (Parker, 2020).

Because highly idealized theoretical models do not directly explain any particular real-world event, some philosophers have proposed that, in addition to providing possible explanations, they serve heuristic functions, such as discovering conjectures and novel perspectives on historical case studies (Šešelja, 2022), and they act as tools for theoretical exploration (Šešelja, 2021) and argumentation (Aydinonat et al., 2021). These uses are surprisingly similar to the ones offered as potential uses of case studies in social sciences and social epistemology. As we have argued in Section 2, cases can also be seen as tools for theoretical exploration that aims to understand social mechanisms (Crasnow, 2012; Morgan, 2012).

The upshot is that it appears to be feasible to bring case studies and models together in the pursuit of joint goals in the social epistemology of science. Although case examples or vignettes are frequently used in modelling studies to explain the motivation and rationale behind the

model, simulations are almost never embedded as part of the mixed-methods suite for full case studies. However, there is no reason why this should remain so. Indeed, in Section 2, following Morgan (2012), we have characterized case studies not as a particular method but a multi-method approach that can accommodate modelling in interpreting and explaining specific causal episodes. In the next section, we discuss how and why to combine modelling with case studies and examine an example found in the literature.

4. Towards a mixed-methods approach

Let us summarize two philosophical lessons that emerge from the discussion in the previous two sections. First, both case studies and modelling methods are used in the social epistemology of science in a similar way to how they are used in social sciences, that is, for devising causal hypotheses instead of conceptual analyses. In studies using these methods, the research question generally is which social factors facilitate the attainment of a particular goal – in the context of social epistemology, an epistemic goal. Therefore, methodological reflections about the proper use of these methods should be informed about the methodological discussion in the philosophy of social sciences.

Second, while agent-based modelling and case studies in the social epistemology of science rarely intersect, they appear to serve surprisingly similar purposes and legitimate uses. They also share similar problems. Both case studies and modelling develop hypotheses and theoretical understanding about social mechanisms, but do not establish how general or robust these mechanisms and their presumed effects are. Consequently, while both methods in the social epistemology of science have frequently been used to formulate concrete policy advice for changing the social organization of science, their reliability in isolation is questionable when it

comes to predicting actual outcomes, unintended consequences, and the robust effectiveness of any proposed policy actions.

For example, on the basis of his modelling study, Zollman claims that an epistemically productive division of cognitive labour is best maintained by limiting the information available to scientists (2007, p. 586). But it is difficult to see how this recommendation could be implemented in a social system of science that favours open data-sharing and publication, and effective communication. Nor is it clear whether constraining the flow of information will actually lead to the desired epistemic goal, or whether the intervention triggers other unintended processes leading to undesired consequences (e.g., scientists wasting time and effort on a project that others have already shown to be a dead end). This suggests that other kinds of evidence is needed to decide whether the causal schemes examined in modelling studies are relevant when the aim is to intervene in actual practices of science.

Similar worries can be raised about Biddle's (2007) Vioxx case study. On the basis of a single case, Biddle makes "a preliminary proposal for reorganizing certain aspects of pharmaceutical research, namely by instituting an adversarial system of research within the FDA" (2007, p. 36). As we have argued, it is not clear whether Biddle's recommendation is feasible given that it requires not only a significant amount of expert labour but also experts who are well equipped to play the adversarial roles. As Zollman's advice, Biddle's advice may also lead to unintended, undesirable consequences (e.g., allegedly neutral experts are increasingly influenced by the companies).

In this section, we argue that these limitations of both modelling and case studies are reasons to explore further a mixed-methods approach to the social epistemology of science. A mixed-methods approach can make the discovery of social mechanisms more effective than using

merely one type of approach. For example, a case study can be used to discover causal schemes, and a model can be developed to examine further the functioning of the said causal mechanisms under a set of counterfactual conditions The results from simulation can suggest further cases to study, which, again can lead to new models. We argue that this kind of mixed-methods approach would be especially useful when social epistemologists aim not merely to explore social mechanisms but also to strengthen the relevance of their study to actual practices of science which are always embedded in a historical and social context. As our criticisms of Zollman (2007) and Biddle (2007) suggest, modelling studies and case studies alone are likely to provide an inadequate basis for socially relevant or engaged social epistemology.

To see how modelling methods and case studies can complement each other in the social epistemology of science, it is useful to briefly highlight their differences. As we have argued, agent-based models of scientific inquiry do not aim to trace actual causal scenarios and rarely utilize empirical data. They investigate hypothetical causal mechanisms which we have called causal schemes, and which are idealized and isolated from real-world details. Case studies, in contrast, are based on actual events and aim to describe them in rich detail. On the face of it, it seems plausible to suggest that case-based knowledge about an event and its context can help social epistemologists decide which causal schemes are actually relevant in a particular context, and model-based understanding of causal schemes can help decide whether lessons learned in one case can carry over to other contexts.

No matter how rich the empirical material supporting a case study is, causality and causal schemes cannot simply be read from even compelling empirical narratives. Mechanistic story telling is an ever-present danger in process tracing and the methodological literature in the social sciences provides many attempts at formulating principles for and taxonomies of testing causal

claims within purely observational case studies (e.g. Collier, 2011). Theoretical models can be of help in suggesting plausible empirical fingerprints of different causal mechanisms, and potential tests for distinguishing between them. Theoretical models can aid in the understanding of a particular causal scenario by embedding it into a broader causal system, and by investigating its interactions with other relevant mechanisms. Theoretical models can also provide hypotheses about the sensitivity of the causal schemes to changes in background assumptions and thus help in bridging the inductive gap in extrapolating the causal conclusion to similar, yet different, contexts. Such knowledge is a necessary step towards social epistemology that aims to be relevant to actual practices of science.

To show how modelling studies and case studies can work together to improve the exploration of social mechanisms and thereby strengthen the relevance of social epistemology, we now return to the example we have discussed above: Zollman's (2007, 2010) epistemic network model and his recommendation to limit the flow of information in scientific communities to ensure an epistemically productive division of cognitive labour.

Frey and Šešelja (2018) challenged Zollman's conclusions by re-examining the historical case of research on peptic ulcer disease and by revising the epistemic network model accordingly. Research on the disease started in the late 19th century, and there were two competing hypotheses: that the disease has a bacterial origin, or it is caused by excessive acidity in the stomach. Both theories were pursued until the 1950s when a widespread consensus on the acidity theory was reached. However, in the 1980s, the disease was demonstrated to have bacterial origins, and the new consensus on the bacterial hypothesis was settled in the 1990s. The decline of the bacterial theory is often attributed to Palmer's study published in 1954, which appeared to show that there were no live bacteria in the human stomach. However, Palmer's evidence was

misleading as it was based on a limited methodology. Based on his network model, Zollman (2010) claimed that the demise of the bacterial hypothesis can be attributed to Palmer's results disseminating too widely too fast.

Šešelja and Straßer (2014) and Radomski et al. (2021) produced textual and other evidence showing that the bacterial hypothesis was virtually abandoned prior to Palmer's landmark publication, despite promising research avenues for the bacterial theory and known problems with Palmer's method at the time. Moreover, Frey and Šešelja (2018) ran the simulation of their revised network model several times with different network structures and analysed only those simulation runs where the false consensus among agents was reached and it lasted less than a third of the total length of each simulation run. This choice was intended to reflect the fact that the research community spent three decades in the wrong consensus during the 90-year research period. More than 90% of the runs that satisfied this criterion were those where the information flow among scientists was restricted, suggesting that the wrong consensus was not caused by too effective information flow. That is, in only 10% of the simulation runs that fit the timeline of peptic ulcer disease research could the false consensus be attributed to the Zollman effect, rendering it an unlikely explanation for the actual case.

In sum, by bringing more empirical information from the case to bear on their model, Frey and Šešelja’s (2018) challenged Zollman's (2010) conclusion that the dynamics of peptic ulcer research resulted from excessive information sharing. Moreover, Frey and Šešelja (2018) explored two additional social mechanisms with their network model: rational inertia, which delayed scientists from premature theory switching, and critical interaction, which guided them toward the correct theory. Their simulation results suggest that the culprit in peptic ulcer disease

research may have been the absence of these two social mechanisms, and particularly incautious theory assessment because of the lack of rational inertia.

Frey and Šešelja (2018) admit that their modelling study is only a step toward better empirical calibration of models of scientific inquiry. Without such calibration, they contend, models remain merely tools for theoretical exploration, lacking evidential value. However, in our view, theoretical exploration and argumentation are precisely the core epistemic functions of models in social epistemology. Case studies help to recognize empirically relevant social mechanisms, and model building is essentially a form of theory construction. Good models help us theorize about actual and counterfactual causal scenarios by providing causal schemes with clearly specified variables and parameters that can be added, removed, and modified. Models need not be accurate descriptions of actual causal scenarios to be epistemically useful.

Note that the way Frey and Šešelja (2018) use modelling is compatible with our suggestion, despite our different understandings of the ultimate goals of modelling in social epistemology. Their aim was explicitly to use modelling along with other methods for investigating a particular historical case. Their study takes a stand on the causal mechanism that may have hampered the advancement of science in this instance. Moreover, it provided new insights into the causal story by calibrating the model to fit some relevant empirical facts of the actual case within the larger space of possibilities.

The lesson to be learned is that a mixed-methods approach to the social epistemology of science gives a boost to the exploration of social mechanisms and helps philosophers strengthen the relevance of their studies to actual practice of science embedded in particular historical and social contexts. It may not, however, provide a sufficient basis for policy advice or concrete recommendations to improve the social organization of science, and hence, social

epistemologists should proceed cautiously if they wish to move from relevant studies to such recommendations.

5. Conclusions

We have argued that a naturalistic approach to social epistemology of science is continuous with theoretical social science. By analysing the use of methods in social epistemology, we have argued that both case studies and modelling methods are used to examine social mechanisms which help philosophers understand how particular features of social systems can lead to epistemically beneficial outcomes (or sometimes epistemically neutral or harmful outcomes). We have also suggested that to understand the epistemic value of cases and modelling, social mechanisms are best interpreted as causal schemes, that is, abstract sketches of causal configurations that could mediate between particular features of social systems and their epistemic effects.

But is social epistemology understood as theoretical social science still philosophy, then? While some social epistemologists are likely to believe that the question is beside the point, it does highlight significant practical questions of disciplinary division of labour. We have argued that to the extent that social epistemologists ask social-scientific questions, they should also make use of social-scientific methods and evidence needed to address those questions. Whether social epistemologists should ask such questions remains an open question, but our argument suggests that if they do, they should strive to live up to the methodological and evidential standards adequate for such knowledge claims.

References

Alexander, J. MK. (2013). Preferential attachment and the search for successful theories. *Philosophy of Science*, 80(5), 769–782. https://doi.org/10.1086/674080

Aydinonat, N. E., Reijula, S. & Ylikoski, P. (2021). Argumentative landscapes: the function of models in social epistemology. *Synthese*, 199, 369–395. https://doi.org/10.1007/s11229-020-02661-9

Bala, V. & Goyal, S. (1998). Learning from neighbors. *Review of Economic Studies, 65*(3), 595–621. https://doi.org/10.1111/1467-937X.00059

Bedessem, B. (2019). The division of cognitive labor: two missing dimensions of the debate. *European Journal for Philosophy of Science, 9*(1), 1–16. https://doi.org/10.1007/s13194-018-0230-8

Biddle, J. (2007). Lessons from the Vioxx debacle: What the privatization of science can teach us about social epistemology. *Social Epistemology,* 21(1), 21–39. : http://dx.doi.org/10.1080/02691720601125472

Bird, A. (2010). Eliminative abduction: examples from medicine. *Studies in History and Philosophy of Science Part A*, 41 (4), 345–352. https://doi.org/10.1016/j.shpsa.2010.10.009

Bovens, L. & Hartmann, S. (2003). *Bayesian epistemology*. Oxford University Press.

Chang, H. (2012). Beyond case-studies: history as philosophy. In S. Mauskopf & T. Schmaltz (Eds.), *Integrating History and Philosophy of Science* (pp. 109-124), Boston Studies in the Philosophy of Science 263. https://link.springer.com/chapter/10.1007/978-94-007-1745-9_8

Collier, D. (2011). Understanding process tracing. *PS: Political Science & Politics*, *44*(4), 823–830. https://doi.org/10.1017/S1049096511001429

Crasnow, S. (2012). The role of case study research in political science: Evidence for causal claims. *Philosophy of Science*, *79*(5), 655–666. https://doi.org/10.1086/667869

Crasnow, S. (2017). Process tracing in political science: What's the story?, S*tudies in History and Philosophy of Science Part A*, 62, 6-13. https://doi.org/10.1016/j.shpsa.2017.03.002

Crasnow, S. (2021). Feminist science studies: Reasoning from cases. In H. Grasswick & N. A. McHugh (Eds.), *Making the case: Feminist and critical race philosophers engage case studies* (pp. 73-98). Albany: SUNY Press.

Currie, A. (2015). Philosophy of science and the curse of the case study. In Daly, C. (Ed.), *The Palgrave Handbook of Philosophical Methods* (pp. 553-572). London: Palgrave Macmillan. https://doi.org/10.1057/9781137344557_22

Epstein, J. M. (2008). Why model? *Journal of Artificial Societies and Social Simulation*, 11(4), 1-12.

Fazelpour, S., & Steel, D. (2022). Diversity, trust, and conformity: A simulation study. *Philosophy of Science*, 89(2), 209-231. https://doi.org/10.1017/psa.2021.25

Frey, D., & Šešelja, D. (2018). What is the epistemic function of highly idealized agent-based models of scientific inquiry? *Philosophy of the Social Sciences*, *48*(4), 1–27. https://doi.org/10.1177/0048393118767085

Gerrits, L. & Marks, P. (2015). The evolution of Wright's (1932) adaptive field to contemporary interpretations and uses of fitness landscapes in the social sciences. *Biology & Philosophy*, 30(4), 459–479. https://doi.org/10.1007/s10539-014-9450-2

Giere, R. N. (1985). Philosophy of science naturalized. *Philosophy of Science*, 52 (3), 331-356. https://doi.org/10.1086/289255

Giere, R. N. (1989). Scientific rationality as instrumental rationality. *Studies in History and Philosophy of Science Part A*, 20 (3), 377-384. https://doi.org/10.1016/0039-3681(89)90013-7

Goldman, A., & Whitcomb, D. (Eds.) (2011). *Social epistemology: Essential readings*. Oxford University Press.

Goodman, N. (1965). *Fact, Fiction, and Forecast*, 2nd edition. Indianapolis, IN: Bobbs-Merrill.

Grasswick, H. & McHugh, N. A. (2021). Introduction. In Grasswick, H. & McHugh, N. A. (eds.), *Making the case: Feminist and critical race philosophers engage case studies* (pp. 1-19). Albany: SUNY Press.

Hangel, N. & ChoGlueck, C. (2023). On the pursuit of qualitative methods in empirical philosophy of science. *Studies in History and Philosophy of Science*, 98, 29-39. https://doi.org/10.1016/j.shpsa.2022.12.009

Hempel CG. (1966). *Philosophy of Natural Scienc*e. Englewood Cliffs, NJ: Prentice-Hall.

Holman, B., & Geislar, S. (2018). Sex drugs and corporate ventriloquism: How to evaluate science policies intended to manage industry-funded bias. *Philosophy of Science*, 85(5), 869–881. https://doi.org/10.1086/699713

Hong, L., & Page, S. E. (2004). Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *PNAS, 101*(46), 16385–16389. https://doi.org/10.1073/pnas.0403723101

Kitcher, P. (1992). The naturalists return. *The Philosophical Review*, 101(1), 53-114. https://doi.org/10.2307/2185044

Kitcher, P. (1993). *The advancement of science: Science without legend, objectivity without illusions*. Oxford University Press.

Kuhn, T. (1962). *The Structure of scientific revolutions*. University of Chicago Press.

Kummerfeld, E., & Zollman, K. J. S. (2016). Conservatism and the scientific state of nature. *British Journal for the Philosophy of Science*, *67*(4), 1057–1076. http://doi.org/10.1093/bjps/axv013

Kuorikoski, J. & Ylikoski, P. (2015). External representations and scientific understanding. *Synthese*, *192*(12), 3817–3837. https://doi.org/10.1007/s11229-014-0591-2

Lakatos, I. (1970). Falsification and the methodology of scientific research programmes. In Lakatos, I. & Musgrave, A. (Eds.), *Criticism and the growth of knowledg*e. Cambridge University Press.

Lipton, P. (2004). *Inference to the best explanation*. 2nd ed. London: Routledge.

Longino, H. (1990). *Science as social knowledge: Values and objectivity in scientific inquiry*. Princeton University Press.

Longino, H. (2022). What's social about social epistemology? *Journal of Philosophy*, 119(4), 169-195. https://doi.org/10.5840/jphil2022119413

Martini, C., & Pinto, M. F. (2016). Modeling the social organization of science: Chasing complexity through simulations. *European Journal for Philosophy of Science, 7*, 221–238. https://doi.org/10.1007/s13194-016-0153-1

Mayo-Wilson, C. & Zollman, K. J. S. (2021). The computational philosophy: Simulation as a core philosophical method. *Synthese*, 199, 3647–3673. https://doi.org/10.1007/s11229-020-02950-3

Mayo-Wilson, C., Zollman, K. J. S. & Danks, D. (2011). The independence thesis: When individual and social epistemology diverge. *Philosophy of Science*, 78(4), 653-677. https://doi.org/10.1086/661777

Morgan, M. S. (2005). Experiments versus models: New phenomena, inference and surprise. *Journal of Economic Methodology*, 12, 317–329. https://doi.org/10.1080/13501780500086313

Morgan, M. S. (2012). Case studies: One observation or many? Justification or discovery? *Philosophy of Science*, *79*(5), 667–677. https://doi.org/10.1086/667848

Mäki, U. (2009). MISSing the world. Models as isolations and credible surrogate systems. *Erkenntnis*, 70(1), 29-43. https://doi.org/10.1007/s10670-008-9135-9

Maki, U. (2020). Puzzled by idealizations and understanding their functions. *Philosophy of the Social Sciences*, 50(3), 215-237. https://doi.org/10.1177/0048393120917637

Page, S. E. (2018). *The model thinker: What you need to know to make data work for you*. Basic Books.

Parker, W. (2020). Model evaluation: An adequacy-for-purpose view. *Philosophy of Science*, 87 (3), 457-77. https://doi.org/10.1086/708691

Pesonen, R. (2024). Highly idealized models of scientific inquiry as conceptual systems. *European Journal for Philosophy of Science, 14:*39. https://doi.org/10.1007/s13194-024-00601-9

Pitt, J. C. (2001). The Dilemma of case studies: Toward a Heraclitian philosophy of science. *Perspectives on Science*, 9 (4), 373-382. https://doi.org/10.1162/106361401760375785

Popper, K. (1974). Replies to my critics. In Schilpp, P. A. (Ed.), *The Philosophy of Karl Popper*, vol. 2 (pp. 961–1197). La Salle, IL: Open Court Press.

Pöyhönen, S. (2017). Value of cognitive diversity in science. *Synthese* 194, 4519–4540. https://doi.org/10.1007/s11229-016-1147-4

Radomski, B. M., Šešelja, D., & Naumann, K. (2021). Rethinking the history of peptic ulcer disease and its relevance for network epistemology. *History and Philosophy of the Life Sciences, 43*(4): 113. https://doi.org/10.1007/s40656-021-00466-8

Reichenbach, H. (1938). *Experience and prediction: An analysis of the foundations and the structure of knowledge*. Chicago: University of Chicago Press.

Reijula, S., & Kuorikoski, J. (2019). Modeling epistemic communities. In M. Fricker, P. J. Graham, D. Henderson & N. J. L. L. Pedersen (Eds.), *The Routledge Handbook of Social Epistemology* (pp. 240-249). New York and Oxon: Routledge.

Reijula, S., & Kuorikoski, J. (2020). The diversity-ability trade-off in scientific problem solving. *Philosophy of Science*, *88*(5), 894–905. https://doi.org/10.1086/714938

Šešelja, D. (2021). Exploring scientific inquiry via agent-based modelling. *Perspectives on Science*, 29(4), 537–557. https://doi.org/10.1162/posc_a_00382

Šešelja, D. (2022). Agent-based models of scientific interaction. *Philosophy Compass*, 17(7), e12855. https://doi.org/10.1111/phc3.12855

Šešelja, D., & Straßer, C. (2014). Heuristic reevaluation of the bacterial hypothesis of peptic ulcer disease in the 1950s. *Acta Biotheoretica, 62*(4), 429–454. https://doi.org/10.1007/s10441-014-9219-7

Solomon, M. (2001). *Social empiricism*. MIT Press.

Solomon, M. (2006). Norms of epistemic diversity. *Episteme* 3:1-2, 23–36. https://doi.org/10.3366/epi.2006.3.1-2.23

Thicke, M. (2019). Evaluating formal models of science. *Journal for General Philosophy of Science, 51*(3), 315–335. https://doi.org/10.1007/s10838-018-9440-1

Veigl, S., & Currie, A. (Eds.) (2025). *Method in the philosophy of science: A user's guide*. MIT Press.

Wagenknecht, S. Nersessian, N. & Andersen, H. (Eds.) (2015). *Empirical philosophy of science: Introducing qualitative methods into philosophy of science*. Springer.

Weisberg, M. (2013). *Simulation and similarity: Using models to understand the world*. Oxford University Press.

Weisberg, M., & Muldoon, R. (2009). Epistemic landscapes and the division of cognitive labor. *Philosophy of Science, 76*(2), 225–252. https://doi.org/10.1086/644786

Ylikoski, P. (2019). Mechanism-based theorizing and generalization from case studies. *Studies in History and Philosophy of Science Part A*, 78, 14–22. https://doi.org/10.1016/j.shpsa.2018.11.009

Ylikoski, P., & Aydinonat, N. E. (2014). Understanding with theoretical models. *Journal of Economic Methodology, 21*(1), 19–36. https://doi.org/10.1080/1350178X.2014.886470

Ylikoski, P. & Zahle, J. (2019). Case study research in the social sciences. *Studies in History and Philosophy of Science Part A*, 78, 1–4. https://doi.org/10.1016/j.shpsa.2019.10.003

Zollman, K. J. S. (2007). The communication structure of epistemic communities. *Philosophy of Science, 74*(5). 574–587. https://doi.org/10.1086/525605

Zollman, K. J. S. (2010). The epistemic benefit of transient diversity. *Erkenntnis, 72*(1), 17–35. https://doi.org/10.1007/s10670-009-9194-6