

DOI: 10.1017/psa.2022.8

This is a manuscript accepted for publication in *Philosophy of Science*.

This version may be subject to change during the production process.

## **Rethinking Ethnography for Philosophy of Science**

**Nancy J. Nersessian**

Harvard University

nancynersessian@fas.harvard.edu

**Miles MacLeod**

University of Twente

m.a.j.macleod@utwente.nl

## **Abstract**

We lay groundwork for applying ethnographic methods in philosophy of science. We frame our analysis in terms of two tasks, 1) to identify the benefits of an ethnographic approach in philosophy of science, and 2) to structure an ethnographic approach for philosophical investigation best adapted to provide information relevant to philosophical interests and epistemic values. To this end, we advocate for a purpose-guided form of cognitive ethnography which mediates between the explanatory and normative interests of philosophy of science, while maintaining openness and independence when framing such an investigation in order to achieve robust unbiased results.

**Keywords:** empirical methods; ethnography; cognitive ethnography; scientific practice

## 1. Introduction

With the exception of historical case study methods, philosophy of science has not traditionally relied on empirical methods to gather data to support its claims. Of course, strongly normative visions of philosophy of science would tend to exclude their value for philosophical analysis. As with historically oriented philosophers, we contend that even normative analyses of science need to be based on authentic practices. Indeed, since the naturalistic turn, there is an expectation that philosophical claims at least be aligned with, if not explicitly grounded in scientific practice. Usually accounts of practice have relied on available historical data, published and archival. A major problem with exclusive use of historical data is that we are at the mercy of “what’s left behind” – an even greater problem with 21<sup>st</sup> century science, which maintains scarce archival materials. Over recent years it has become increasingly common for philosophers to use a range of social science methods to generate their own empirical data to advance claims about scientific practice. This trend is supported, for instance, by the agenda of the Society for Philosophy of Science in Practice. Quite recently, philosophers have initiated discussions on the potential role and value of empirical methods, such as the LEAHPS (Learning from Empirical Approaches in HPS) workshops and the edited volume, *Empirical Philosophy of Science: Introducing Qualitative Methods into Philosophy of Science* (Wagenknecht et al., 2015).

A number of philosophy of science studies have been published that employ experimental methods and survey-based approaches, as well as bibliometric studies, to investigate for instance the meanings of certain concepts and the types and role of values in scientific practice. Such approaches are called “quantitative” in the social science literature because the research uses numerical and formal methods of analysis. These empirical approaches have generated some reflection on and critical discussion of their potential value within philosophy of science and empirical strengths (Waters 2004; Machery 2016; Mizrahi 2013; Robinson et al. 2019). Our concern in this paper, however, is with systematic ethnographic approaches, which comprise field observations, artifact collection, and interviews and interpretive forms of data analysis such as textual coding and thematic analysis. Such approaches are called “qualitative” in the social science literature because they use natural language methods of interpretive analysis. There are now a number of philosophical studies which do make systematic use of ethnographic methods (see, e.g., Nersessian 2012a; Nersessian 2019; Nersessian 2022; MacLeod and Nersessian 2013b;

Andersen and Wagenknecht 2013; Hardesty 2018; Hangel and Schickore 2018; Mansnerus and Wagenknecht 2014). A number of qualitative empirical approaches currently in use, however, do not appear to rely on any explicit systematic evidence gathering and analysis procedures. By-and-large, philosophers who use qualitative approaches have not reflected on these methods when they advance or make use of empirical claims.<sup>1</sup> Our intention here is to initiate that methodological discussion.

In this paper we explore what roles ethnography might play in an empirical philosophy of science and how it could be adapted to suite philosophical interests, goals, and values. In section 2 we position ethnography within philosophy of science by illustrating specific philosophical goals ethnography might serve. In section 3 we consider which adaptations of this social science method might be needed to achieve these goals. We stress the value of ethnography both as a means to study contextual relations, but also as a means to collect and analyze data systematically on real-world epistemic practices, and, thus, improving the reliability of philosophical claims about scientific practice. In these sections, we argue for the possibility of a more purpose-guided style of cognitive ethnography than is customary in social sciences and STS approaches, and provide guidance on how to approach a philosophical adaptation. Finally, in section 4 we give examples from our research that exemplify this more purpose-guided style for philosophical ends.

Since transfer of methods across domains is never straightforward, we intend our analysis not as definitive, but as a means to provoke more extensive methodological discussion, and provide an initial guide for those who seek to engage with ethnographic approaches.

---

<sup>1</sup> We note, but do not address here, that the same can be said for the selection, gathering, and analysis of historical data used in case studies, where method traditionally is learned implicitly in apprenticeship (Osbeck and Nersessian 2015). The major philosophical concern about using historical case studies that has been addressed has been the question of whether they are generalizable, rather than how they have been conducted.

## 2. Positioning Ethnographic Research

Our strategy in this section is to sketch various philosophical issues and goals to which ethnography can contribute. In this section we first provide a basic characterization of ethnography and distinguish it, briefly, from other empirical methods and then articulate some of the interests and goals of philosophy of science to which an ethnography can contribute.

### 2.1. Distinguishing ethnographic approaches

As mentioned, philosophers of science are now engaging with a variety of empirical methods, that can, roughly, be partitioned into quantitative and qualitative research. Experimental methods, natural language programming and bibliometric methods rely on either mathematical or statistical methods of data analysis to draw claims from data. Experimental methods, for instance, survey-based methods, have been used to design studies that investigate the meanings and uses scientists ascribe to concepts under controlled conditions which mimic practical contexts or to investigate the values motivating scientists' decisions (Stotz et al. 2004; Linqvist et al. 2011; Murdock et al. 2017; Robinson et al. 2019). Natural language programming methods are being used to analyze scientific texts in order to determine the structure of conceptual frameworks through mapping conceptual relationships (Mennes 2018). These can be combined with machine-learning driven topic-modeling methods to survey large collections texts and draw conclusions on the pluralities of conceptual meanings for central scientific concepts, such as the concept of "life" in biological science (Malaterre and Chartier 2021). Bibliometric and corpus analysis methods are used to help gauge the relevant importance of various concepts, methods, disciplines, and so forth within the scientific landscape, which in turn aids philosophers in forming better generalizations on the scope and relevance of particular scientific activities, practices and beliefs (Mizrahi 2013, 2016).

Given that these empirical methods have been carving out their own space and prominence within philosophy of science, it is important to consider what specific contributions a systematic ethnographic approach can make both as an alternative and complement to them. In order to do this, we need to begin with a basic account of how an ethnographic approach can be distinguished. In broad terms, we propose ethnography is:

*A naturalistic inquiry which aims to collect **in-situ** data and extract information on practices and their relations to their contexts through an intensive and detailed description and systematic analysis of those practices and their contextual relations.*

This characterization is sufficiently broad so as to capture the basic motivations of ethnographic researchers from a wide variety of fields. Guba (1978, Lincoln and Guba 1985) introduced the notion of a “naturalistic inquiry,” specifically, to contrast with empirical inquiry by means of controlled experimental design. A naturalistic inquiry is not only ecologically valid in that there is at most minimal manipulation of a real-world setting, but also there are no strict constraints, such as predetermined categories of interpretation, placed on analysis. Such studies are principally inductive rather than hypothetico-deductive. The scope of an inquiry is of course framed by research questions which serve to guide the focus, but an ethnographic approach is geared towards open exploration of practices rather than testing hypotheses. It is not a goal to confirm or apply a particular theoretical interpretation. These features support the second aspect of an ethnographic study raised in our characterization, which is that ethnography aims at intensive and detailed descriptions and analysis of in-situ practices under study and their relations to the environments in which they occur, which are otherwise inaccessible.

Naturalistic inquiry requires systematic data gathering and analysis to establish the reliability, to the extent possible, of its claims. In general, ethnographies are built around a family of evidence gathering tools, primarily field observation and interviewing, and interpretive data analysis techniques, such as grounded coding (Corbin and Strauss 2008) and thematic analysis (Braun and Clark 2006). Although context or situation can be understood broadly, philosophers of science are likely to focus on features of *problem-solving* contexts, which include reasoning, understanding, conceptual innovation and change, imagination, and collaboration. These contexts can provide detailed information on many issues of interest to philosophers, including the nature and structure of scientific problems, what affordances and constraints shape specific practices, how scientists develop and use methods and concepts, how they create and evaluate claims and explanations, and how they communicate results.

An ethnographic approach, broadly construed in these ways, differs from experimental and other empirical approaches used in philosophy of science. For example, in experimental approaches, researchers need to design a controlled experimental set-up or context. Such design requires prior

assumptions as to what is relevant to the phenomena under study. This method enables controlled testing of specific hypothesis regarding specific relations, which is argued to confer generality on the findings. However, extending findings about practices that derive from artificial contexts to real-world settings is problematic, as has been argued extensively by “environmental perspectives” in cognitive science (Nersessian 2005). Ethnographic studies, in contrast, have high ecological validity with respect to the ascertaining the intricate complexity of specific practices and their environmental dependencies.

Ethnographic studies and bibliometric and text-analysis methods differ, among other ways, in that they operate at different scales of investigation. Bibliometric and published text-analysis realize reliable results by surveying over a wide domain of scientific activity and considering the output of many scientific labs and disciplines, sometimes exhaustively. These methods, too, derive generalized inferences about scientific practices or scientific concepts. In contrast, an ethnographic approach works with a small sample of scientific practices, including use of scientific concepts and works towards a more general account by comparing findings across multiple cases.

As we see it, all these empirical methods have positive contributions to make to investigating scientific practice, both for philosophers’ descriptive and explanatory aims, as well as their normative ones. What is important is that philosophers understand what questions and goals each can be used to address and the affordances and constraints of each in their systematic application. In the next section we consider specific potential contributions of ethnographic methods.

## 2.2 Benefits of an ethnographic approach for philosophy of science

Among empirical methods, an ethnographic approach has distinct benefits to contribute to philosophy of science. As with any academic field, of course, philosophy of science is diverse in its interests, goals, and practices. We concentrate on those of practice-oriented philosophy of science, which has genuine *descriptive and explanatory* goals that, among other aims, are directed towards understanding or explaining the epistemological principles at work in scientific practice. A major contribution ethnography can make is to improve philosophical understanding of how and to what extent specific epistemic practices of science are shaped by situational features, including cognitive, material, and sociocultural ones. These features are largely erased in scientific publications, on which many philosophers rely.

As with other empirical approaches, an ethnographic approach operates through systematic procedures by which to gather sufficient and varied data and to conduct analyses. Such procedures serve to minimize the biases on what we might expect is relevant, based on prior philosophical theories, to unearth surprising or novel findings, and to increase the reliability of our claims, at least as they pertain to the particular labs or research groups studied. Through such procedures ethnographic research also serves to improve the robustness of philosophers' claims about scientific practice, and align them with actual practices.

We acknowledge that many philosophers of science do consider the field principally a *normative* one, in which case empirical methods should contribute to normative aims rather than just characterize or explain aspects of scientific practice. Ethnographic research does not preclude having normative aims to evaluate the epistemological or decision-making principles in play once the ethnographic analysis is complete, but in this case the evaluation operates on authentic practices. Such a normative analysis of practices can be relevant and practicable to scientists, as well as contribute to philosophical analysis. In this way, ethnographic research can make a unique contribution to evaluating normative claims, for instance, on what might or not be a rational procedure, against a background of knowledge of the actual practices scientists employ and the constraints under which they work.

Further an ethnographic investigation can help provide means to assess the generality and applicability of theories and concepts developed in philosophy, such as “mechanistic explanation” and “natural kind” (Mansnerus and Wagneknecht 2015). Ethnographic research can help us determine whether and the extent to which these apply in practice, or suggest how they might be modified, or assist our development of novel philosophical theories and concepts. Along these lines, we can assess the extent to which the rational norms or constraints posited to hold by a philosophical theory or account are truly explanatory for a particular set of scientific decisions or judgements.

Those applying other empirical methods have worked with a similar rationale in mind (see, e.g., Machery 2016). For instance, experimental methods have been used to test philosophical conceptions or reconstructions of the gene concept against scientists' own use in order to help assess the relevance of philosophical conceptions (Stotz et al. 2004). These methods do so by isolating the particular concepts and precise uses to be studied, and test a broad sample set to

enhance the reliability and generality of assessments. Such specificity and controlled testing are contrary to an ethnography approach. On the other hand, ethnography offers the possibility to examine how the concepts are used in situ and over extended periods. Thus, ethnography and other empirical methods can complement each other in terms of the kinds of information they can acquire, the level and breadth of detail, and the strength and generality of the inferences they afford. Indeed, as cognitive science has demonstrated with the “in vivo/in vitro” method (Dunbar 1995), ethnographic (in vivo), experimental (in vitro), and other empirical methods, can be used in combination to provide richer and more accurate accounts of scientific practices.

To sum up, what we aimed to show here is that, in principle, an ethnographic approach has important contributions to make with respect to various goals of philosophy of science.

Ethnography is used by a wide range of fields today, and they all set their own constraints on ethnographic investigation to align it with their interests, goals, and values. In the next section, we undertake a similar task and consider how to adapt ethnography for use in philosophical research.

### 3. Adapting Ethnographic Methods

What understanding philosophers of science have of ethnography derives, primarily, from how it has been portrayed and used in STS, which remains close to its anthropological roots. The main objection philosophers have to STS accounts is their epistemic relativism. When STS researchers introduced ethnographic methods as a means for studying scientific practices, for instance in “laboratory life” (Latour and Woolgar 1979), they programmatically constrained relevant explanatory factors to social and cultural, including personal motivations and interests. This move downplayed or denied entirely the relevance of epistemological and cognitive factors traditionally emphasized in philosophical analysis. Epistemic relativism, however, is a specific theoretical stance associated with what is known as the program of “the social construction of knowledge” (e.g., Bloor 1991) and is not necessary to an ethnographic investigation. STS has aligned ethnography with the interests, goals, and values of that perspective. There is nothing in principle about ethnographic methods that would prohibit focusing on other dimensions of practices, including the rational and cognitive dimensions that align with philosophical interests, goals, and values. We believe philosophers should not cede this important empirical method to STS.

Cognitive science is one area that has made productive use of ethnography to investigate a range of problem-solving practices in work and learning contexts. What the cognitive anthropologist, Edwin Hutchins, has dubbed “cognitive ethnography” (Hutchins 1995, see also, Lave 1988) has emerged as a methodological choice of “environmental perspectives” (Nersessian 2005) in cognitive science. Researchers who adopt this perspective aim to understand cognition in situ, as a process in which conceptual, social, and material resources do not simply scaffold, but are integral to – and integrated in – various cognitive practices. Cognitive ethnography focuses largely on problem-solving practices of individuals and groups, which include reasoning, understanding, conceptual change, learning, collaborating, and so forth, as situated in contexts, with their attendant resources, and as distributed across people and salient artifacts. Many practice-oriented philosophers also have an objective to surmount the perceived rational/social or cognitive/cultural divide (Longino 2001; Nersessian 2005) that permeates both STS and traditional normative philosophical approaches. Cognitive ethnography is a useful method for investigating how these dimensions are integrated in the epistemic practices of scientists.

The most widely influential research in cognitive ethnography, and the best known by philosophers, is Hutchins’ (1995) research on technologically rich and well-defined problem-solving environments, specifically piloting planes and ships, where “natural” is extended to comprise work contexts. “Naturally occurring culturally constituted human activity” (1995, xiii) of any kind can, in principle, be investigated with ethnographic methods. Although science research labs differ in important ways from the environments studied by Hutchins, they are such natural, or “in the wild,” environments for scientists, and it is possible to use an ethnographic approach in line with cognitive ethnography to examine rational dimensions of knowledge-making practices (see Nersessian 2022, forthcoming).

Another way cognitive ethnography differs importantly from STS-type ethnography and aligns better with philosophical objectives is with respect to the latter’s aim to build theory. STS aims to determine the particularities of a case, without abstracting findings to more general, theoretical conclusions. As with cognitive science accounts generally, data analysis in cognitive ethnography aims to move beyond just thick description (Geertz 1973) of richly nuanced details of the case to advance a theoretical account of cognitive processes. As Hutchins has framed this objective:

“There are powerful regularities to be described at the level of analysis that transcends the details of the specific domain. It is not possible to discover these regularities without understanding the details of the domain, but the regularities are not about the domain specific details, they are about the nature of cognition in human activity” (quoted in Woods 1997, 7).

Philosophical accounts, too, need to transcend the details of the case to meet both explanatory and normative goals. Data analysis strategies that facilitate abstraction are important in any philosophical adaptation of ethnography. Although rooted in the particular, philosophers can avail themselves of social science methods for qualitative data analysis, such as grounded coding and thematic analysis, to abstract candidate regularities from the particularities of a case, and evaluate the extent to which these transfer robustly across the epistemic practices in a scientific field (see, e.g., Corbin and Strauss 2008; Patton 2006; Braun and Clarke 2006). Such analyses aim to move from the specificity of the case to build a broader interpretive account by using systematic procedures to abstract and coalesce interpretive categories, and, where appropriate, formulate candidate hypotheses to transfer and assess across cases, using multiple cases to work back and forth between data and theory to attain a warranted degree of generality.

### 3.1 Distinguishing a philosophical approach to ethnography

We aim to place philosophical ethnography within the family of cognitive ethnographies because it aligns well with the interests, goals, and values of practice-based philosophy of science. However, some thought is required to distinguish a form more tailored to philosophy of science, with its descriptive, explanatory and normative goals. We follow the approach of Linden Ball and Thomas Ormerod (2000) to framing a field-specific cognitive ethnography, in their case, to study engineering design practices. To begin, they distill a set of principles central to a “prototypical” ethnography from a wide range of introductory guides; situatedness, richness, participant autonomy, openness, personalization, reflexivity, self-reflection, intensity independence, historicity. They proceed to discuss the limitations of traditional, anthropological interpretations of these with respect their field and consider adaptations to fit the objectives of cognitive ethnography in design by weighting the importance of these principles to their goals and refining their interpretation. We broadly agree that the principles they have distilled are prototypical, as applied to either data collection or analysis, and use them to discuss how to adapt cognitive ethnography, in our case to interests, goals, and values of philosophy of science.

The goals we outlined for ethnography in philosophy of science argue for the importance of collecting rich data on in situ problem-solving in order to build detail accounts of those practices and their relations to their contexts. Four of the principles Ball and Ormerod distill – situatedness, richness, intensity, and participant autonomy – concern how to collect data to support the production of detailed accounts and ensure their reliability, as philosophers of science would demand. Situatedness - meaning that data are collected by a participant observer who is located within the everyday context of interest - and participant autonomy - meaning that the study participants are not required to comply in any fixed, predetermined study arrangements - ensure that the study is naturalistic. Richness of data collection is related to intensity, where the presumption is that an ethnographic study needs to collect enough data from a variety of sources so as to be able to build warrant for an account through establishing consistency of findings across methods and sources of data collection (“triangulation”). For instance, an interview only study would be considered a legitimate qualitative investigation, but not an ethnographic one. Intensity requires the ethnographer, for instances, write field notes that document frequent observations of sufficient duration and undertake multiple interviews to ensure the collection of a significant sample of ongoing practices in the environment, as well as to rule out that a sample has been biased in a particular direction. The period, variety, and quantity of data collection take into consideration the community to be investigated and the purposes of the study. In sum, these principles establish that data collection in an ethnographic study will be of sufficient scope, variety, and depth to yield reliable information.

The remaining six principles Ball and Ormerod outline concern both data collection and data analysis, and might initially seem at odds with philosophical objectives. Four of these pertain to what might be called the ethnographer’s stance towards the investigation; namely, that an investigation requires personalization, reflexivity, self-reflection and historicism. Many STS investigations advance interpretations of these principles in line with their assumption of epistemological relativism. In particular, they interpret the requirement for reflexivity in this manner, but it need not be interpreted as such, and, indeed, is a much contested notion in

qualitative research.<sup>2</sup> Broadly construed, it requires that the ethnographer endeavor to take account of the participant's point of view as an individual ("subject"), rather than simply as an object of investigation, while also taking account of their own effects on the behavior of the participants.

When given a broad interpretation all of these principles can be consistent with philosophical goals and values. Personalization requires the researcher examine their feelings and attitudes about the situations they encounter in collecting or analyzing data. With historicism the researcher generates understanding principally through making connections to local historical and cultural contingencies. Self-reflection requires the researcher to scrutinize the framing of the study and the procedures for data collection and analysis for potential influence of prior theoretical interests. The main feature of these principles is that they all serve to mitigate researcher bias, which is an issue in all empirical research. There are many pitfalls in ethnographic research with respect to potential biases in the formulation of research questions and in data collection and analysis techniques, and so principled, critical reflection is required. Such self-scrutiny, for example, would control for asking leading questions in conducting an interview or for importing favored notions into data analysis. In general, it is important to keep in mind that all ethnographic research is interpretive. As such, the researcher is the instrument of data collection and analysis, and so the researcher's interests, values, and motivations are always present, and it is a necessary part of good research to confront and be explicit about these (Osbeck and Nersessian 2015). Even with a broader interpretation, these principles play an important role in ensuring evidential and interpretive reliability.

Finally, in STS fields these principles have sometimes been interpreted to require that research questions, data collection, and data analysis should not even be motivated by prior theoretical interests or values. This interpretation is, however, overly restrictive. Even traditional anthropologists were motivated by theoretical interests, such as kinship relations. Complete neutrality would make research impossible. There is an important difference between an inquiry

---

<sup>2</sup> See, for example, the 32 articles in the two special issues of the journal *Forum: Qualitative Research*, "Subjectivity and Reflexivity in Qualitative Research" (issue3(3), 2002 and 4(2), 2003).

being constrained by theory and being motivated by theory, such as in the formulation of its research questions. Self-scrutiny keeps the researcher aware of such motivations.

The final two principles Ball and Ormerod extracted from the literature – openness and independence – require special consideration when adapting ethnography to philosophy of science. Openness to novel findings that emerge from data analysis and independence of analysis from theoretical presuppositions are central strengths of ethnography in comparison to other empirical methods. However, within the spectrum of cognitive ethnographies, philosophy of science is likely to be more purpose-guided in its goals; in particular, there is an expectation that empirical findings have a role to play in evaluating philosophical theory and concepts (see, e.g., Mansnerus and Wagneknecht 2015). Ethnography conducted with evaluative purposes in mind will be framed around concepts, theoretical issues, and normative claims of interest, and as such requires data collection and data analysis to be more focused. Thus, research questions that guide data collection and analysis will need to be formulated with respect to issues pertaining to these interests. Philosophical ethnographies with evaluative aims will, of course, lack some of the explorative and open-ended dimensions of a traditional ethnography. What is important is that, although purpose-guided, data collection and analysis not “fit the data to the theory.” A major strength of an ethnographic approach is that it enables philosophers to go beyond “pre-analytic intuitions” and see the unexpected. In our experience, the most important concern for philosophers of science when setting up an ethnographic project and managing it, is that of all ethnographic research: to frame research questions and to collect data in a manner so as to facilitate, in particular, the potential discovery and emergence of unanticipated results in a robust manner.

In this quest philosophers can be guided by the long history of critical scrutiny of qualitative research methods with respect to issues of warrant (rigor, reliability, trustworthiness) in psychology and sociology, where it has been important to establish credibility and value for these with respect to dominant quantitative methods.<sup>3</sup> Patton (2002) provides a useful overview.

---

<sup>3</sup> Recently, the extent and variety of current fields adapting ethnography and qualitative methods to their interests, goals, and values, led the Journal of Qualitative Research to set up a taskforce to develop guidelines sufficient for evaluating methodological integrity in data collection and

Qualitative methods applications are not formulaic or recipe-like, so they can be tailored to research goals and questions, while adhering to accepted canons of what constitutes "trustworthy" (Lincoln and Guba, 1985) or "validated" (Eisner, 2003) data collection and analysis procedures.<sup>4</sup> In particular, Eisner's (2003) three principles serve as helpful guides: structural corroboration, referential adequacy, and consensual validation. Structural corroboration requires that a sufficient number of data points converge on a conclusion to support an interpretation. Referential adequacy addresses the richness and clarity of the description and interpretation and how it aligns with member understanding. Consensual validation refers to the level of inter-rater agreement that can be reached for the coding schemes and interpretation. Taken together these principles serve to ensure the investigation has collected systematic and sufficient data from multiple sources so as to corroborate and determine the trustworthiness of interpretations.

Finally, the question remains to what extent a philosophical ethnographic approach might legitimately achieve generality of sufficient scope to be useful for more purpose-guided aims. Philosophical theory often relies on generalizations which are expected to apply at the level of fields or disciplines, not only to individual labs or research groups, where the issue is whether these are, indeed, representative. Although we do not have the space to delve into details here, philosophers can draw on numerous analysis that have been advanced to think about such issues as they pertain to historical case studies. Insights from philosophy of science on how to use case-study material to build theory help to illuminate the difference. Historically oriented philosophy of science, too, has been concerned with how use empirical insights from data on scientific practices to develop or examine theoretical notions, while avoiding unwarranted generality. Early critiques, of a simple inductivist perspective on case data advocated that the way to

---

analysis across fields when reviewing papers and grant proposals. Levitt et al. (2017) provides a useful conceptual analysis of these.

<sup>4</sup> Although "validate" is reserved in philosophy for logical argumentation, it is often used to signify credible qualitative research. We prefer to use Guba's (1978, Lincoln and Guba 1985) more neutral term of "trustworthy" when considering issues of warrant.

understand the relation between specific cases and theory is as a bootstrapping method customarily used in the sciences (see, e.g., Nersessian 1991). Roughly, in such bootstrapping processes, “hypotheses are made within a background of beliefs and problems... They are refined, made more specific, modified, or rejected in light of more constraining data (a detailed case study). Surviving hypotheses are then tested against other data and other hypotheses to determine the extent of their validity” (Nersessian1991, 683). Bootstrapping entails a working back and forth between data and theory, until a satisfactory accommodation is achieved. It is an iterative and incremental, open-ended process.

Recently, in thinking about the use of qualitative data on scientific practices, such as from interview studies and ethnographies, Erika Mansneurus and Susann Wagenknecht (Mansneures and Wagenknecht, 2015) follow up on a recommendation of Hasok Chang (2012, p.111) to construe the relation between historical case studies and philosophical theorizing in terms of the “concrete” and the “abstract” instead of the customary inductive categories of the “particular” and the “general”. They use this suggestion to further articulate the bootstrapping account and argue that the way philosophers can arrive at “limited generalizations,” while they “avoid unwarranted generality,” is to “create a *dialogue between the abstract and the concrete* [italics in original]” (40). That is, to work back and forth between data and theory, which they, too, note is a bootstrapping procedure. Further, they contend, such “productive interplay” (40) makes it possible to examine and further develop philosophical concepts and theories with qualitative case study data, while avoiding the pitfalls philosophy has often succumbed to of fitting the data to the theory. The mechanism case we discuss in section 4 provides a good example of this. Such dialog is in line with traditional practices in ethnographic analysis. As the anthropologist Clifford Geertz has emphasized, “one does not start intellectually empty-handed. Theoretical ideas are not created wholly anew in each study... they are adopted and refined from other, related studies, and, refined in the process, applied to new interpretive problems. If they cease being useful... they stop being used... If they continue to be useful, throwing up new understandings, they are further elaborated and go on being used” (1983, 57).

### 3.2 Putting ethnography into practice

Given the desiderata discussed in the previous section, we illustrate briefly how we put these into practice in our own investigations. Here we outline and motivate the investigative strategies we used in our ethnographies of four research labs in the bioengineering sciences in the bioengineering sciences, two in biomedical engineering (BME) and two in integrative (computational) systems biology (ISB). We found these emerging interdisciplinary fields to be particularly good subjects for an ethnography of epistemic practices because methods, norms, and epistemologies were not yet stabilized. Our investigations in each lab began with a basic set of questions relevant to our philosophical interests, broadly summarized as<sup>5</sup>:

What are the representational and reasoning practices (“problem-solving”) in this community?

How is epistemic warrant developed for novel practices?

What are the different epistemic assumptions, values, and norms at play in the practices of these interdisciplinary communities?

What and how are concepts, methods, and theories being used from engineering? From biology? What is the nature or results of their interaction?

In what ways might cognitive, social, cultural, and material “factors” be mutually implicated in epistemic practices?

These questions are framed as open questions without asserting specific hypotheses. As such they guided but did not totally constrain data collection and analysis, and were broad enough to facilitate “productive interplay”. Although we chose to investigate interdisciplinary fields at the intersection of biology and engineering, the selection of labs was serendipitous and we had no prior understanding of the fields. We began investigation in each lab with a “pilot” study, which consisted of interviewing the director about the lab’s history and current research, touring the lab with descriptions of artifacts and activities by the director and lab manager, and meeting with

---

<sup>5</sup> We had additional questions related to their learning practices as part of our research funded by the NSF.

researchers who would participate in our investigation. This preliminary research enabled us to better-focus our research questions, and to identify relevant data sources. In particular, they focused our attention on the various model-building practices through which bioengineering scientists attempt to manage the complexity of biological systems by utilizing engineering methods, concepts, materials, and values (e.g., Nersessian and Patton 2009; Nersessian 2012a; Nersessian 2022; MacLeod and Nersessian, 2013a; MacLeod and Nersessian, 2018).

Our data collection in all labs comprised: audio-taped open and semi-structured interviews; field observation with note-taking; arranged demonstrations of experimental procedures and technologies; video and audio recorded lab meetings; note-taking during journal club discussions; photographs of white board text and diagrams; diagrams of the spatial layout of each lab and photographs of lab space evolution; and artifact collection, to the extent possible: grant proposals, paper drafts, presentations, dissertation proposals, emails, private lab wiki correspondence, diagrams/sketches, and so forth. Our data were analyzed using a variety of complementary qualitative methods, specifically, qualitative data coding, case study analysis, thematic analysis, and cognitive-historical analysis (see, (Osbeck, Nersessian, Malone, and Newstetter, 2011) for a thorough account of our qualitative data analysis procedures).

We were attentive to the concern that the data collected in an ethnographic study might not be representative. One novel mitigating strategy of our research was to practice "team ethnography." More than one ethnographer had responsibility for observations and interviews in a given lab, and our more senior members worked across the labs. Another strategy we used was to collect longitudinal data intensively for at least a year in each lab. The data comprised persistent observations, multiple interviews of each participant for multiple participants, and the kinds of archival data previously mentioned.

In general, our weekly research group meetings provided the venue to scrutinize and evaluate the ethnographic work together as it unfolded and to reach consensus on coding, theme development, and other forms of data interpretation. Such collaborative reflection greatly enhanced our individual abilities to make explicit and keep in check our interests, values, and motivations.

A look at a few of our systematic coding analysis procedures serves to illustrate ways in which inferences can be drawn from rich data sets through systematic means to ensure to the extent

possible robust, unbiased findings, while directed towards subjects of interest to philosophers. Once our coding progressed sufficiently, we related findings to philosophical theoretical frameworks, formulated hypotheses in interplay with issues within these, and evaluated, revised, and refined hypotheses in comparison to our emerging empirical insights.

Coding is an interpretive procedure for partitioning the data by attaching descriptive categories to units of interview texts and field observation notes. Our approach to coding can be broadly characterized as “grounded,” as characterized by Corbin and Strauss (2008); meaning that although guided by our research questions, we remained open to seeing what codes might emerge from the data.<sup>6</sup> Coding development took place in several phases. We began with “open coding” directed towards identifying, categorizing, and describing what the text of the interview is about. During this process, coding pairs worked collaboratively on each transcript. We analyzed a subset of interviews progressively, line-by-line with the aim to provide an initial description for as many textual passages as seemed appropriate. In our research meetings, the entire group discussed the clarity, fit, and logic of the codes assigned. In early coding, we presented interpretations to the research lab members to check for their alignment with our understanding. We used feedback from all sources to make adjustments. Mid-way through our investigations we hired an external coding auditor to provide a check on our procedures, and establish interrater reliability.

Consistent with the goals of analytic induction (codes emerging from data and leading to hypotheses) and constant comparison (codes compared against possible alternative interpretations) (Lincoln and Guba 1985, Corbin and Strauss 2008), we continued to code additional interviews and to revisit previous coding and assess descriptions for adequacy and fit.

---

<sup>6</sup> We describe coding procedures here consistent grounded theory. Alternatively, it can be efficacious for philosophers of science to apply other methods, such as thematic analysis or a content analysis (Braun and Clarke 2006, Cho and Lee 2014). Both facilitate coding with pre-given themes or targets in mind, and aim at giving a less comprehensive account of relations in the data than was our goal.

After about 20% (the standard) of interviews were coded intensively, we coded the rest more selectively, focusing on categories of most relevance to our research questions to build out those categories, such as “model-based reasoning.” We further grouped and arranged codes into superordinate categories and subcategories. In the final phase, we related codes and categories directly to our research questions as a start towards building “theory.” In this process, the researcher formulates “a set of well-developed categories (themes, concepts) that are systematically interrelated through statements of relationship to form a theoretical framework that explains some phenomenon” (Corbin and Strauss 2008, 55). “Theory” development, in effect, takes the form of developing increasingly refined conceptual models. We used the codes, categories, and themes to create multiple case analyses, which provide finely detailed descriptions that follow epistemic practices of a specific researcher, or small group, as they worked to solve a complex problem.

Beyond description, though, code development is an abstractive process and codes provide the basis for cross-case comparison when detached from case-specific details. Codes enable the researcher to formulate candidate hypotheses for transfer across cases. We assessed transfer of selected major categories and themes across the labs, especially to ascertain what commonalities there might be in the nature of the model-building practices in related fields of bioengineering sciences and how they achieve their epistemic goals.

#### 4. Sample fruits of an ethnographic approach

To show the fruitfulness of an ethnographic approach for descriptive, explanatory, and normative, goals of philosophy of science we provide sample findings from our investigations. The first case provides an example of how an ethnographic approach can uncover novel science practices and open new philosophical issues. The second case shows how ethnographic findings can be used to assess philosophical notions, in this case “mechanistic explanation.” Within the confines of this paper, it is not possible to detail the data from which these findings derive. These data can be found in the numerous publications we cite.

##### 4.1 Developing epistemic warrant in in vitro modeling practices: the role of hybrid devices

Our first ethnographic study began with our discovery of a modeling practice widespread within the field of biomedical engineering, but previously unknown in literature on models in philosophy of science. That practice is to investigate specific processes occurring in complex biological systems through “building” (which we take to comprise processes of designing, constructing, evaluating, experimenting, and redesigning) highly abstracted in vitro model-systems designed to “mimic” or “simulate” these processes in controlled experimentation. The model-systems (locally called “devices”) comprise living biological cells and tissues and engineered materials, and are thus hybrid entities both ontologically and epistemically. They are built in iterative and incremental processes, much of which we were able to capture for specific projects. Significantly, we were able to follow how norms, standards, and methods around model-building become justified.

The tissue engineering lab’s research focused on the effects of mechanical forces as blood flows through the arteries. The primary model-system comprises a “flow loop” device that replicates the shear forces of blood as it flows through the lumen and either endothelial cells on slides or tissue-engineered blood vessel wall models (“constructs”), which are tubular collagen scaffolds with embedded endothelial cells or smooth muscle cells. The neuroengineering lab’s research focused on learning in living neural networks. The primary model-system in that lab was “the dish,” which is a culture of approximately forty thousand cortical neurons plated on a specially designed multielectrode array from which to record the electrical activity of the neurons, as well as stimulate them with externally generated signals. The dish can be connected with various “embodiments,” such as robots (“hybrots”) or computational creatures (“animats”). A major goal of the research was to get the dish to learn to control the behavior of the embodiments.

We focused on the nature of the inferences afforded by in vitro model-systems (Nersessian, 2002, 2004, 2006, 2009; Nersessian and MacLeod, 2017; Nersessian and Patton, 2009); the processes through which epistemic warrant is built as they are designed and constructed (Nersessian 2022; Chandrasekharan and Nersessian, 2015; Chandrasekharan and Nersessian, 2017); their roles in concept transfer, formation, and change (Nersessian 2012a, b; Nersessian and Chandrasekharan 2009); and how they function as epistemic tools (all our references cited). One novel finding of the BME research is that contrary to the usual characterization of the initial process of analogical reasoning as retrieving a source analogy, there are no ready-to-hand

analogical sources. Instead, in vitro models are constructed towards serving, eventually, as analogical sources from which BME researchers can transfer warranted inferences as candidate hypotheses to understand and control the target phenomena. As a researcher explained about her model-system: “*We typically use models to predict what is going to happen in a system [in vivo]. Like people use mathematical models to predict... what’s going to happen in a mechanical system? Well, this is an experimental model that predicts what would happen – or you hope that it would predict – what would happen in real life.*” Such prediction is a form of analogical transfer from a model built to serve as a source. During our investigations the main focus in each lab was to gain sufficient understanding of the in vitro devices and control their behavior under experimental conditions. Only after this achievement could researchers seek to transfer hypotheses derived from the in vitro source to the in vivo target system.

In vitro models are the primary means through which researchers in numerous fields of BME gain epistemic access to complex biological phenomena. The researchers build epistemic warrant for a model through the *principled* decisions and rationalizations they make in the processes of building it. The warrant for using these kinds of models as epistemic tools (Boon and Knuutila 2009) is connected to how the models function as dynamic representations; that is, how they are built to instantiate and simulate in vitro features. As we followed out the researchers’ reasoning and justifications for selecting features to instantiate in a version of the model, as well as those that might be left out, we came to understand that there is an important connection between analogy and exemplification. The epistemic warrant for devices as sources for learning about in vivo phenomena derives from the ways in which they are designed to exemplify – instantiate and, though this instantiation, refer to – features of the target germane to the problem situation (Elgin 2018).

We found the building process is subject to many situational dependencies, including the current state of engineering methods, technologies, and materials, and is dominated by the engineering values of abstraction, isolation, and control. For instance, in the first line of research in the tissue engineering lab, researchers extracted and isolated endothelial cells from their supporting biological structures in the artery in order to produce measures of cell morphology and proliferation under controlled shear forces produced by the flow loop device. The researchers rationalized this move by arguing that the endothelial cells are the layer of an artery closest to the

blood and so are directly in contact with the blood force; therefore, examining them in isolation should yield reliable information about specific aspects of their behavior. However, although they recognized endothelial cell cultures on slides do not instantiate features of the blood vessel wall that are important to the functionality of the cells, for the first twenty years of the research it was not possible to engineer living three-dimensional tissue. When methods and materials became available, they were able to produce three-dimensional blood vessel wall models (“constructs”), which served as an analogical source with which they could simulate additional behaviors and also have a basis of comparison for their cell culture studies. Further, the ability to build this analogy opened a novel application goal, that is, the possibility to create a blood vessel implant to repair damage arteries. Most of the research we investigated centered on building a range of model-systems (different analogical sources) with the construct to understand more about the functionality of endothelial and smooth muscle cells, as well as associated blood vessel tissues, in response to various mechanical forces. They also now directed research towards the new application goal, which required researchers to determine how to strengthen constructs to withstand in vivo blood forces and as well as to find a way to produce endothelial cells, the most immune-sensitive cells in the body, that would not be rejected by the host. The latter requirement opened a line of in vitro research on the effects of mechanical forces on stem cell differentiation and on maturation of progenitor endothelial cells that circulate in the bloodstream. The latter research was conducted by building the most complex model the lab had undertaken: an animal model-system to provide an analogy to the target human cardiovascular system.

Through an intensive investigation of different in vitro model-building practices in two contexts, we were able to see how warrant for a model – as well as for the novel practice of in vitro modeling – is developed through many iterations. In particular, we were able to determine how researchers rationalize choices, commonly by arguments pertaining to control, biological plausibility, and current engineering feasibility, to identify features of the biological system to exemplify in the built analogy, with respect to the goals of the problem context. On the basis of our analysis, we advanced an understanding of models in relation to analogy largely missing from the philosophical literature (exceptions include Nersessian 2008, Knuuttila and Loettgers, 2014).

## 4.2 Assessing “mechanistic explanation” in integrative systems biology

Systems biology is a diverse field. In the labs we studied the researchers mainly built models of cellular metabolic networks, such as the glycolysis pathway in yeast or lignin production pathways in poplar tree cells. The networks or pathways represent successions of chemical reactions, which transform chemical elements into other elements. Such networks are complex even when pared-back or artificially isolated for the purposes of modeling. For instance, there are often feedback relations, which produce nonlinearities. ISB researchers in our labs model these networks with ordinary differential equations, which capture the concentration of an element at any time, and thus track how concentrations evolve from a given starting point, often to reach a steady state.

This description of the model-building strongly suggests that a relevant philosophical notion for understanding what modelers are trying to achieve is “mechanistic explanation” of various biological functions. This philosophical notion has been developed and applied to describe and explain the structure and purposes of models across a range of biological contexts (Machamer et al., 2000; Bechtel, 2011). At first glance, ISB models appear to fit the structure of a mechanistic explanation, being composed of clearly identified parts and mathematical descriptions of their interactions. In the context of our analysis, we found the ISB labs provided opportunities to examine the applicability of this notion for characterizing and interpreting computational model-building in this area. We tracked model development and probed the reasoning behind modelers’ decisions (MacLeod and Nersessian, 2013a, 2013b; MacLeod and Nersessian 2018; Chandrasekharan and Nersessian 2019).

Our data suggested that mechanistic explanation was not an effective characterization of their practices. One thing to emerge was the extent to which simplifications figure into identifying a problem sufficiently tractable to model, as well as in developing a model. Although simplifications are often announced in published work, rationalizations that go beyond the technical ones, as well as intermediate problem reformulations, though common, are often not announced at all. An ethnographic approach provides insight into the reasons modelers advance for these modifications, while also by helping to reveal deviations from normative standards such as those associated with mechanistic explanation. In practice, these modelers start with the goal to replicate systems sufficiently to generate robust predictions with respect to specific variables.

However, they were almost always forced by aspects of the context to scale-back and limit their representation goals. By tracking these contextual constraints and the responses to them, we became aware of significant deviations from standards of mechanistic explanation and how these responses were rationalized relative to the existing context.

Two of the main contextual elements we noted were the frequent lack of data or data of the right kind (time series) and the limits on available computation. The lack of data was partially a consequence of difficult collaborative relationships with experimentalists, whose work constraints and motivations for participating modeling, often do not align with those of modelers (MacLeod and Nersessian 2013c; Nersessian 2022). Longitudinal observations and interviews allowed us to track the modeling process and to reveal general patterns. Modelers responded to the constraints by paring down both their overall goals and their representations of the systems through a lengthy iterative and incremental process. The purposes of such processes are to understand how to align the various constraints with their ability to produce a model that can capture at least some central dynamic relationships predictively, as well as to use the affordances of multiple degrees of freedom in the representation to increase the opportunities of finding a parameter fit accurate for a limited set of variables. Limiting and refining representational goals allow modelers to identify variables and parameters of less importance for instance, based principally on information from the models themselves, which can then be simplified or removed.

However, since many of these processes are driven by the goal of finding well-fit models for a select set of variables within a system, given computational and data constraints – and not strictly according to biological plausibility – they have the effect of obscuring whether or not the model adequately represents the underlying mechanisms of their targets. Indeed, modelers themselves often claim at the end of the processes to have developed a “system-level” understanding of their systems, which, given their model-building practices, we interpret as the model providing insight into the numerical or mathematical relations between specific variables, but not an overall mechanistic explanation of the behavior.

While it is conceivable that these deviations from mechanistic explanation might be identified from published work, modelers will not necessarily report all important simplifications. In particular, the initial pathways usually are not, leaving out an aspect of the modeling process

which provides insight into how modelers trade-away more mechanistic representations for a good fit for a limited set of variables. Our ethnographic approach, by virtue of rich data collection, was better placed to recognize the limits of applying mechanistic explanation as a governing concept by which to interpret and evaluate these models. Nor would it be easy working with publications alone to understand why modelers engage in producing these limited representations with limited predictive capacity, particularly given that they often publicly advocate the need for large scale and much less simplified models as the route to effective prediction. The rich data we collected on contextual features of computational model-building was critical to our ability to map the constraints modelers operating under and their precise motivations and goals and to discover the iterative trade-off strategy underlying final published models. Although the published models might seem limited or insufficient from the point of view of “mechanistic explanation”, it arguably more reasonable, as in this case, to evaluate the models in terms of a rich understanding of the strategies modelers use, the goals they are trying to achieve, and the contextual constraints they face.

## 5. Conclusion

We anticipate the use of ethnographic and other empirical methods will increase in philosophy of science, given their current momentum. An important factor in this increase is the fact that 21<sup>st</sup> century science by-and-large lacks the archival records of previous centuries. A Michael Faraday, who kept extensive records, with commentary, of day-to-day in notebooks, diaries, and letters, was a rarity in the past, and is likely non-existent in today’s science. When we started our research, we asked the tissue engineering lab director for access to the lab’s notebooks, and his response was “what notebooks?” They only kept records of experimental data in computer files. Computational labs do not even have these. Although we had access to the neuroengineering lab’s discussion wiki, it was not archived when the lab closed. Thus, by-and-large, philosophers are left trying to discern practices from publications and dissertations, where conventions rule out detailing the investigative pathway to the findings and discoveries, especially where researchers run into obstacles or impasses. In this paper, we have opened a discussion about some of the potential contributions and advantages of an ethnographic approach, and have proposed a way to frame and structure it to meet philosophical objectives within the network of principles governing a family of cognitive ethnographic approaches. We have advocated for a

purpose-guided form of ethnographic investigation, which can mediate between the explanatory and descriptive and the normative goals of philosophy of science. We hope our outline can provide a guide for those wishing to pursue an ethnographic project, or who want to weigh up the relative advantages and disadvantages of various empirical approaches, and make informed decisions about which method might best suit their purposes. Ultimately, the various approaches should be understood as complementary, each with an important part to play in developing a more empirical underpinning to support philosophical claims about scientific practice.

Acknowledgments:

We appreciate the support of the US National Science Foundation in conducting this research: (DRL0106773; DRL0411825; DRL0909971). We also appreciate the comments from Lisa Osbeck and Michael Stuart on an early draft, as well as from participants at both the LEAHPS I (Learning from Empirical Approaches in HPS) workshop at the University of Pittsburgh Center for Philosophy of Science, Pittsburgh 2018, and LEAHPS II workshop at Leibniz University, Hannover 2019, for comments on our ethnographic approach. Finally we would like to thank three anonymous reviewers and the editor, whose comments improved the paper significantly.

## References:

Andersen, Hanne, and Susann Wagenknecht. 2013. "Epistemic dependence in interdisciplinary groups." *Synthese* 190 (11): 1881-1898.

Ball, Linden J., and Thomas C. Ormerod. 2000. "Putting ethnography to work: the case for a cognitive ethnography of design." *International Journal of Human-Computer Studies* 53 (1): 147-168.

Bechtel, William. 2011. "Mechanism and biological explanation." *Philosophy of science* 78 (4): 533-557.

Bloor, David. 1991. *Knowledge and Social Imagery* (2nd ed.). Chicago: University of Chicago

Boon, Mieke, and Tarja Knuuttila. "Models as epistemic tools in engineering sciences: A pragmatic approach." In: *Handbook of the Philosophy of Technological Sciences*, ed. Anthonie Meijers, 687-719. Amsterdam: Elsevier Science.

Braun, Virginia, and Victoria Clarke. 2002. "Using thematic analysis in psychology." *Qualitative research in psychology* 3 (2): 77-101.

Chandrasekharan, Sanjay, and Nancy J. Nersessian. 2011. "Building cognition: the construction of external representations for discovery." In *Proceedings of the Annual Meeting of the Cognitive Science Society*, 33 (33).

Chandrasekharan, Sanjay, and Nancy J. Nersessian. 2017. "Rethinking correspondence: how the process of constructing models leads to discoveries and transfer in the bioengineering sciences." *Synthese* 48 (6): 1-30.

Chang, Hasok. 2012. *Is water H<sub>2</sub>O? Evidence, realism and pluralism*. Springer.

Cho, Ji Young, and Eun-Hee Lee. 2014. "Reducing confusion about grounded theory and qualitative content analysis: Similarities and differences." *Qualitative Report* 19 (32): 1-20.

Corbin, Juliet, and Anselm Strauss. 2014. *Basics of qualitative research: Techniques and procedures for developing grounded theory*. Los Angeles: Sage Publications.

Dunbar, Kevin. 1995. How scientists really reason: Scientific reasoning in real-world laboratories. In *The Nature of Insight*, ed. Robert J. Sternberg and Janet E. Davidson, 365–395. Cambridge, MA: MIT Press.

Eisner, Elliot W. 2003. On the art and science of qualitative research in psychology. In *Qualitative Research in Psychology: Expanding Perspectives on Methodology and Design*, Paul M. Camic, Jean E. Rhodes, and Lucy Ed Yardley, ed., 17-29. Washington, DC: American Psychological Association.

Elgin, Catherine Z. 2018. *True Enough*. Cambridge, MA: MIT Press.

Geertz, Clifford. 1973. *Interpretation of Cultures: Selected Essays*. New York: Basic Books

Guba, Egon. G. 1978. *Toward a methodology of naturalistic inquiry in educational evaluation*. Los Angeles: UCLA Center for the Study of Evaluation.

Hardesty, Rebecca A. 2018. "Much ado about mice: Standard-setting in model organism research." *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 68: 15-24.

Hangal, Nora, and Jutta Schickore. 2017. "Scientists' conceptions of good research practice." *Perspectives on Science* 25 (6): 766-791.

Hutchins, Edwin. 1980. *Culture and inference: A Trobriand case study*. Cambridge, MA: Harvard University Press.

Hutchins, Edwin. 1995. *Cognition in the Wild*. Cambridge: MIT press.

Knuuttila, Tarja, and Andrea Loettgers. 2014. "Varieties of noise: Analogical reasoning in synthetic biology." *Studies in History and Philosophy of Science Part A* 48: 76-88.

Latour, Bruno, and Steve Woolgar. 1979. *Laboratory Life*. London: Sage Publications

Lave, Jean. 1988. *Cognition in practice: Mind, mathematics and culture in everyday life*. Cambridge, UK: Cambridge University Press.

Levitt, Heidi M., Sue L. Motulsky, Fredrick J. Wertz, Susan L. Morrow, and Joseph G. Ponterotto. 2017. "Recommendations for designing and reviewing qualitative research in psychology: Promoting methodological integrity." *Qualitative psychology* 4(1), 2-22.

Lincoln, Yvonne, and Egon Guba. 1985. *Naturalistic Inquiry London*. Beverly Hills, CA: Sage Publications.

Linquist, Stefan, Edouard Machery, Paul E. Griffiths, and Karola Stotz. (2011). "Exploring the folkbiological conception of human nature." *Philosophical Transactions of the Royal Society B* 366: 444–454.

Longino, Helen. 2001. *The Fate of Knowledge*. Princeton, NJ: Princeton University Press

Machamer, Peter, Lindley Darden, and, Carl Craver. 2000. "Thinking about mechanisms." *Philosophy of science* 67(1): 1-25.

MacLeod, Miles, and Nancy J. Nersessian. 2013a. "Building Simulations from the Ground-Up: Modeling and Theory in Systems Biology." *Philosophy of Science* 80 (4): 533–556.

MacLeod, Miles, and Nancy J. Nersessian. 2013b. "Coupling Simulation and Experiment: The Bimodal Strategy in Integrative Systems Biology." *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 44: 572-584.

MacLeod, Miles, and Nancy J. Nersessian. 2013c. "The creative industry of integrative systems biology". *Mind and Society* 12: 35-48.

MacLeod, Miles, and Nancy J. Nersessian. 2018. "Modeling complexity: cognitive constraints and computational model-building in integrative systems biology." *History and philosophy of the life sciences* 40(1): 1-28.

Machery, Edouard. 2016. "Experimental philosophy of science." In *A companion to experimental philosophy*, ed. Justin Sytsma and Wesley Buckwalter, 475-490. John Wiley and Sons

Malaterre, Christophe, and Jean-François Chartier. 2021. "Beyond categorical definitions of life: a data-driven approach to assessing liveness." *Synthese* 198 (5): 4543-4572.

Mansneures, Erika, and Susann Wagenknecht. 2015. "Feeling with the organism: A blueprint for empirical philosophy of science." In *Empirical Philosophy of Science: Introducing Qualitative Methods into Philosophy of Science*, ed. Susann Wagenknecht, Nancy J. Nersessian, and Hanne Andersen, 37-64. Cham: CH: Springer.

Mennes, Julie. 2018. "Sense Disclosure: A new procedure for dealing with problematically ambiguous terms in cross-disciplinary communication." *Language Sciences* 69: 57-67.

Mizrahi, Moti. 2013. "The pessimistic induction: a bad argument gone too far." *Synthese* 190 (15): 3209-3226.

Mizrahi, Moti. 2016. "The History of Science as a Graveyard of Theories: A Philosophers' Myth?." *International Studies in the Philosophy of Science* 30 (3): 263-278.

Murdock, Jaimie, Colin Allen, and Simon DeDeo. 2017. "Exploration and exploitation of Victorian science in Darwin's reading notebooks." *Cognition* 159: 117-126.

Nersessian, Nancy J. 1991. "Discussion: The method to 'meaning': A reply to Leplin." *Philosophy of Science* 58: 678-687.

Nersessian, Nancy J. 2002. "The cognitive basis of model-based reasoning in science." In *The Cognitive Basis of Science*, eds., Peter Carruthers, Stephen Stich, and Michael Siega, 133-153. Cambridge: Cambridge University Press.

Nersessian, Nancy J. 2004. "Interpreting scientific and engineering practices: Integrating the cognitive, social, and cultural dimensions." In *Science and technical thinking*, eds., Michael E. Gorman, Ryan D. Tweney, David C. Gooding, and Alexandra P. Kincannon, 17-56. Mahwah, NJ: Lawrence Erlbaum Associates.

Nersessian, Nancy J. 2006. Model-Based Reasoning in Distributed Cognitive Systems. *Philosophy of Science* 73 (5): 699-709.

Nersessian, Nancy J. 2008. *Creating Scientific Concepts*. Cambridge, MA: MIT Press.

Nersessian, Nancy J. 2009. "How do engineering scientists think? Model-based simulation in biomedical engineering research laboratories." *Topics in Cognitive Science* 1, 730-757.

Nersessian, Nancy J. 2012a. "Modeling Practices in Conceptual Innovation: An Ethnographic Study of a Neural Engineering Lab." In *Scientific Concepts and Investigative Practice*, ed. Uljana Feest and Friedrich Steinle, 245-269. Berlin: DeGruyter.

Nersessian, Nancy J. 2012b. "Engineering concepts: The interplay between concept formation and modeling practices in bioengineering sciences." *Mind, Culture, and Activity* 19 (3): 222-239.

Nersessian, Nancy. J. 2022 (forthcoming). *Interdisciplinarity in the Making: Methods and Models in Frontier Science*. Cambridge, MA: MIT.

Nersessian, Nancy J. and Sanjay Chandrasekharan. 2009. "Hybrid analogies in conceptual innovation in science." *Cognitive Systems Research* 10(3): 178-188.

Nersessian, Nancy J., and Chris Patton. 2009. Model-based reasoning in interdisciplinary engineering: Two case studies from biomedical engineering research laboratories. In *Philosophy of Technology and Engineering Sciences*, ed. Dov M. Gabbay, Paul Thagard, John Woods, and Antonie W. M. Meijers, 678-718. Amsterdam: Elsevier Science.

Nersessian, Nancy J., and Miles MacLeod. 2017. "Models and simulations." In *Springer Handbook of Model-Based Science*, Lorenzo Magnani and Tommaso Bertolotti, 119-132. Springer.

Osbeck, Lisa. M. and Nancy J. Nersessian. 2015. Prolegomena to an empirical philosophy of science. In *Empirical Philosophy of Science*. Lisa M. Osbeck and Nancy J. Nersessian. 13-35. Cham: CH: Springer.

Osbeck, Lisa M., Nancy J. Nersessian, Karen R. Malone, and Wendy C. Newstetter. 2011. *Science as Psychology: Sense-making and Identity in Science Practice*. Cambridge, UK: Cambridge University Press.

Patton, Michael Quinn. 2002. *Qualitative Research and Evaluation Methods*, Beverly Hills, CA: Sage.

Robinson, Brian, Chad Gonnerman, and Michael O'Rourke. 2019. "Experimental Philosophy of Science and Philosophical Differences across the Sciences." *Philosophy of Science* 86 (3): 551-576.

Stotz, Karola, Paul E. Griffiths, and Rob Knight. 2004. "How biologists conceptualize genes: an empirical study." *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 35 (4): 647-673.

Wagenknecht, Susann, Nancy J. Nersessian, and Hanne Andersen. 2015. *Empirical philosophy of science: Introducing qualitative methods into philosophy of science*. Springer

Waters, C. Kenneth. 2004. What concept analysis in philosophy of science should be (and why competing philosophical analyses of gene concepts cannot be tested by polling scientists). *History and philosophy of the life sciences* 26: 29-58.

Woods, David D. 1997. Towards a theoretical base for representation design in the computer medium: Ecological perception and aiding human cognition. In *The Ecology of Human - Machine Systems*, ed. John M. Flach, Peter A. Hancock, Jeff Caird, and Kim J. Vicente, 157-188. Hillsdale, NJ, Lawrence Erlbaum.