

Forthcoming in *Philosophy of Science*. This is a slightly truncated version — please cite the published version when available.

Experimental philosophy of science: Beyond categories

Abstract. Traditional philosophy of science tends to investigate scientific concepts by studying their *meanings* or *extensions*, aiming to illuminate the *properties* or *kinds* those concepts refer to. Call this the *categorical* approach: it studies scientific concepts by asking which things fall under an associated category, and in virtue of what they so fall. While this is often informative, it overlooks a significant part of what concepts contribute to science. I argue that although experimental philosophy of science has inherited the categorical approach, it is actually well-positioned to *critique* that approach and develop alternatives.

1. INTRODUCTION

Experimental philosophy aims to correct a pervasive tendency in philosophy: the reliance on *individual intuitions* about the meaning of concepts. Those intuitions may be too fragile, too context-sensitive, or too idiosyncratic to do the work that philosophers require (Ludwig, 2018; Weinberg, 2016). Experimental philosophers instead defer to a broader community's understanding of concepts, appropriating empirical methods to study that broader community.

Philosophers of science have gladly embraced this approach. We are, after all, used to setting aside our intuitions and deferring to the way a community (science) uses concepts. And we're even versed in some empirical methods, if you count case studies (Laudan et al., 1988). So it's no surprise that some of the most interesting work in experimental philosophy (X-phi) has been in experimental philosophy of science (X-phi-sci), with a slew of results concerning concepts like 'gene' (Stotz & Griffiths, 2004), 'innateness' (Griffiths & Stotz, 2008), and 'representation' (Favela & Machery, 2023).¹

I'm going to suggest that there is a second pervasive tendency for X-phi-sci to correct. This second tendency does not concern the *sources* of philosophical insight (intuitions), but rather the *kind* of insight philosophers seek, namely: insight into the *meaning* and *extension* of scientific concepts. When philosophers take up a scientific concept, they are usually interested in how that concept is defined, what it takes for that concept to apply, or to which things it is (correctly) applied. This tendency is present in work on concepts like 'gene' (Stotz & Griffiths, 2004), 'innateness' (Griffiths & Stotz, 2008), 'life' (Cleland, 2012; Cleland & Chyba, 2007), 'stem cell' (Suárez, 2023), 'mental disorder' (Ballesteros & Batalla, 2024), 'computation' (Piccinini, 2015; Shagrir, 2022), 'function' (Garson, 2019), 'representation' (Cummins, 1991; Neander, 2017; Ramsey, 2007; Shea, 2018), and virtually every other scientific concept philosophers have turned their analytic eye on. As one example, our questions about neural representation tend to overwhelmingly concern what neuroscience thinks representation *is*, how we should *define* 'representation,' which things *are* representations, or what *makes something count* as a representation (cf. Richmond, 2023, 2025a).

¹ Some of the most interesting *critiques* of X-phi also come from philosophy of science (Waters, 2004).

I'll call this the *categorical* approach: it studies scientific concepts by asking which things fall into an associated category, and in virtue of what they so fall.² In Section 2 I'll argue that the categorical approach overlooks a significant part of what concepts contribute to science, but is nonetheless pervasive in both philosophy of science and X-phi-sci. In Section 3 I'll argue that X-phi-sci is well-positioned to critique and correct this tendency, if only it can return to its original self-conception and deepen its connections with cognitive psychology.

2. X-PHI-SCI AND THE CATEGORICAL APPROACH

This should be an uncontroversial claim: what concepts do, and what they contribute to science, goes well beyond their meanings, extensions, or any categories they pick out. Consider the concept 'virus' in the study of disinformation (Kucharski, 2016; Panchal & Jack, 2022). Disinformation is not in the extension of 'virus.' It does not satisfy the typical meaning or definition of that concept (a string of RNA or DNA with various biological properties). It does not have the features something must have to *count as* a virus.

But thinking of disinformation as a virus has nonetheless shaped disinformation science. The concept 'virus' has introduced a range of resources from virology, especially models and modeling strategies, that help to capture and explain important features of disinformation. We could understand this as a metaphorical use of the concept 'virus.' Or we could understand it as a case of *domain transfer*, where strategies from one domain are applied in another (e.g., see Mollo, 2020, 2021; Richmond, 2025b). What matters is that the concept contributes *resources* to disinformation science, helping scientists perform scientific tasks. It does not contribute a *category* into which misinformation is supposed to fall.³

Disinformation science is not unique in this respect. Domain transfer is a ubiquitous feature of science (Dunbar, 2002; Nersessian, 2002), and there is no reason to expect that other cases — like the fluid dynamical models applied to traffic jams (Sun et al., 2011) — are any better understood by asking whether traffic *counts as* a fluid, or what makes something a fluid in the first place. Categorization is simply not all that scientific concepts do, and we miss a great deal if we ignore the other ways they shape scientific work (e.g., see Egan, 2014, 2021; Richmond, 2025a, 2025b). I won't explore these in detail, but aside from introducing explanatory strategies and modeling resources, some of the contributions concepts can make are:

- *Directing attention* by exploiting a concept's associations or the things it makes salient (Reuter, 2024). E.g., when we think of the brain as a dynamical system, we highlight certain of its features (like the way it evolves through a state space) as we build and appraise models of it (cf. Favela, 2021).

² Compare the 'semantic mentalist' approach in X-phi (Alexander et al., 2010).

³ Nor is the associated category very informative about the resources in question. Learning that viruses must be strings of RNA or DNA does little to illuminate the modeling resources that are associated with the concept 'virus' and taken up in disinformation science.

- *Framing questions* to make tools, hypotheses, or methods applicable. E.g., when we think of cognition as computation, we make it possible to frame hypotheses about cognitive capacities in terms of the algorithms underlying them (Richmond, 2025a).
- *Suggesting questions* in the first place. E.g., Bursten et al. (2016) point out the way that different conceptions of multimetallic nanoparticles prompt different sets of questions in materials science (Bursten et al., 2016, p. 1918).
- *Facilitating communication* between labs, programs, or fields. Lazebnik (2002) stresses the way that shared terms and concepts allow different biology labs to identify the same patterns and to understand each other's work, and Egan (2014) and Cao (2022) make similar points about concepts in cognitive science.
- *Providing frameworks* within which to understand a target system. E.g., consider John Craig's 1699 theory of probability, which borrowed the concepts of Newtonian physics: probabilities were "velocities of suspicion" propelled by "motive forces" through a "space" whose axes represented "degrees of assent" (quoted in Gigerenzer et al., 1989, p. 27). This provided a formal and a conceptual framework to capture the mathematical structure of probability.

In light of all this, it would be a problem if philosophers focused too closely on the categories that concepts pick out, neglecting their other functions. It would be a *big* problem if philosophers reached sweeping conclusions about the role of a scientific concept based only on the categories it names, e.g., concluding that it has no useful role in science simply because of its failure to delineate useful categories. In the rest of this section I will suggest that both philosophy of science and X-phi-sci currently have these problems.

To make that point, a more thorough literature review would be desirable. But in lieu of such a review I will describe, in more detail, one representative example from each of *traditional* and *experimental* philosophy of science and offer citations to further work in each field that suggest the categorical approach is prevalent beyond the examples I discuss.

To begin with traditional philosophy of science, consider the way we have approached the concept 'representation' in cognitive science. This century's most influential books on that topic are likely Ramsey (2007) and Shea (2018). Both aim to understand *what role the concept 'representation' plays in science* (Ramsey, 2007, p. 5; Shea, 2018, p. 10). Each, however, pursues that goal by offering an account of what features *make something* a representation (cf. Richmond, 2025a). Ramsey helpfully characterizes the broader literature on representation as attempting to state "the sort of physical conditions and relations that have been assumed to bestow upon an internal state the status of representation" (2007, p. 189), and offers his own condition: to count as a representation something must satisfy a particular "job description" (2007, p. 24). Shea offers a different account, to the effect that to be a representation (and a representation with particular content) is to have a certain set of historical or etiological properties (2018, p. 5).

It might be controversial to describe these projects as focused on *categories*. They may just aim to “naturalize,” “ground,” or even just “explain” representation, as another influential book puts it (Neander, 2017, pp. 3, 9). But regardless of how we label their project, the focus is on *what counts as a representation*, and *in virtue of what features* it counts. And that is precisely the project I outlined in the introduction. This work does not ask about the modeling resources the concept ‘representation’ introduces, how it helps scientists pose or frame questions, how it directs attention or guides thought and action, etc. It asks which things are representations, and what features make them representations. Discussions of representation are just one example, but they are a paradigmatic example of an area dominated by the categorical approach, and reflect a broader trend in the special sciences, whether the concept under discussion is ‘representation,’ ‘computation’ (Piccinini, 2015; Shagrir, 2022), ‘function’ (Garson, 2019), or something else.

What about *experimental* philosophy of science? Has X-phi, in its defiance of tradition, managed to buck the categorical approach? Unfortunately not. When experimental philosophers study scientific concepts, they do so almost exclusively by asking which things fall under those concepts, or what features it takes to fall under them. I’ll discuss just one example: a recent experimental study on (again) the concept ‘representation’ in cognitive science. Favela & Machery (2023) use a straightforward elicitation design, where participants read descriptions of a fictional brain imaging experiment and report whether they agree or disagree with statements like “that brain activity *represents* faces,” “that neuron *responds to* faces,” and some others. Their main finding was that, for statements like “the neuron *responds to* faces,” participants’ answers clustered closer to “strongly agree” and “strongly disagree” than they did for statements like “the neuron’s spiking *represents* faces,” which were closer to “neither agree nor disagree” (p. 7). Favela & Machery draw a number of conclusions from this, most importantly:

[P]sychologists’ and, to an even greater extent, neuroscientists’ concept of representation is imprecise: Psychologists and neuroscientists selected responses indicating uncertainty about what properties a brain pattern must have to count as a representation and what follows from calling a brain pattern a representation. (Favela & Machery, 2023, p. 11)⁴

The authors conclude on this basis that the concept ‘representation’ must be eliminated from the brain sciences, or at least reformed, since it is not currently playing a useful role. It would be hard to be clearer about one’s commitment to the categorical approach. The only evidence offered that the concept ‘representation’ does not play a useful role in science is that scientists are unsure *what counts as a representation* and *what features something must have* to count as a representation. To dramatize the point, imagine we studied disinformation scientists by showing them vignettes describing cases of disinformation, and asking whether they agree with statements like “that bit of disinformation is a virus.” I don’t know what we would find, but we

⁴ They make it clear that “what follows from calling a brain pattern a representation” does not refer to non-categorical functions of the concept ‘representation,’ but merely what properties something must have if we know it is a representation (p. 9).

would clearly not be asking the right questions. We would not be getting at the role of the concept ‘virus’ in disinformation science (cf. Richmond, 2023).

As with the previous example of traditional philosophy of science, Favela & Machery (2023) is a popular, influential, and paradigmatic example of X-phi-sci, representative of a major trend in the field, also exemplified by the papers I listed at the bottom of the first page. Importantly, this trend is also present in *foundational* work in X-phi-sci. E.g., Stotz & Griffiths’ (2004) influential studies of the concept ‘gene’ almost exclusively asked which things count as (which) genes, and why.

My point has been that the categorical approach is pervasive in both traditional and experimental philosophy of science. It is not *universal*. Some authors explicitly consider the non-categorical functions of concepts (Baker et al., n.d.; Cao, 2022; Egan, 2014, 2021; Richmond, 2023, 2025a, 2025b). But the majority approach in these literatures is so invested in the categorical approach that it can seem quite reasonable to infer from “this concept isn’t helping scientists sort things in and out of a category” to “this concept isn’t serving a useful scientific role.” As I argued at the beginning of this section, that constitutes a serious methodological problem.

3. BEYOND THE CATEGORICAL APPROACH

I’m going to argue that even though X-phi, and X-phi-sci especially, typically succumb to the categorical approach, they are actually in a perfect position to critique and improve on it. This is because X-phi’s *self-conception* offers a number of important resources for a non-categorical approach, even if the field doesn’t always take advantage of them.

So, what exactly *is* the self-conception of X-phi? Two aspects are worth discussing. First, an influential account of X-phi considers it to be a sort of “conceptual ecology” (Griffiths & Stotz, 2008), studying the role of concepts in their “epistemic niches” (Stotz, 2009, p. 235), and understanding concepts as *tools* for solving scientific problems (Stotz, 2009). On that conception, the categorical approach would make sense only if a concept’s niche consisted of categorization tasks. But we know that science and scientists do not just categorize: they model, explain, simplify, communicate, reason, predict, etc. If *that* is a concept’s niche, then conceptual ecology cannot limit itself to understanding the ways a concept categorizes its targets, any more than real ecologists can content themselves to just say which things count as foxes and which as rabbits. A genuine conceptual ecology would have to look beyond the meaning and extension of scientific concepts, or the categories associated with them.

Another, perhaps more influential understanding of X-phi casts it as a branch of cognitive science: an attempt to replace conceptual analysis with methods taken from psychology, and to investigate the “underlying cognitive process[es]” involved in concept use (Knobe, 2016).⁵ But

⁵ Alexander et al. make this point only a little more directly than most: “Experimental philosophy’s positive programs would do well, we think, to become more like experimental psychology” (2010, p. 311). And Cohnitz & Häggqvist (2018, p. 419) point out that “there are several philosophical projects in which the targeted subject matter is explicitly conceived as grounded in psychological competences ... [and] psychological processes are what these projects are after.”

as it happens, cognitive psychology is just the place to find a viable *alternative* to the categorical approach. Consider two examples.

First: the psychology of medical reasoning — a topic that is closely related, for obvious reasons, to scientific reasoning. The psychology of medical reasoning aims to understand the various reasoning tasks doctors perform and how different concepts (among other resources) contribute to the achievement of those tasks (Goldszmidt et al., 2013). There is no effort in this literature to privilege categorical tasks — *managing a patient encounter* sits alongside the task of *diagnosis*. And it is not assumed that any task must be served by concepts categorizing patients, health conditions, or anything else (e.g., Goldszmidt et al., 2012). Instead, the goal in this work is to understand the broader range of effects that concepts have on medical reasoning.

Second: the work of Tania Lombrozo and collaborators, who study *teleological explanations* and the concept ‘function.’ This work is diverse. Some asks which groups of people prefer explanations that appeal to an object’s function as opposed to its mechanistic features (Lombrozo et al., 2007). Some asks which kinds of categorizations (Lombrozo, 2009) and generalizations (Lombrozo & Gwynne, 2014) people make when they understand a system in terms of its functions as opposed to its mechanistic features (Lombrozo & Carey, 2006).

But some of the most ambitious work asks about the cognitive goals that using the concept ‘function’ helps people achieve (Lombrozo & Carey, 2006). For example, Lombrozo (2010) prompted two groups of laypeople with *mechanistic* vs *functional* characterizations of a toy system, and asked how the two groups went on to *assign causal responsibility* within that system. Participants in the mechanistic condition appeared to use a transference notion of causation, and participants in the functional condition appeared to use a counterfactual notion. That’s an important finding about “the impact of [teleological vs mechanistic] explanatory modes on reasoning” (p. 304), having little to do with which things count as (having) functions or mechanisms.⁶

Should X-phi-sci follow suit? The case here is stronger here than it was with the ecological self-conception, which is just an evocative metaphor. The psychological conception is literal: science plausibly *is* a set of cognitive tasks. In fact, this idea — the *cognitive theory of science* (Giere, 1987; Nersessian, 2002) — has been a preoccupation of philosophers arguing in support of X-phi-sci (e.g., Machery, 2016). And if science is a set of cognitive tasks, it should be studied as cognitive tasks are, which is to say, *not* exclusively or primarily on the categorical approach.

So, what would it mean to set aside the categorical approach? Most basically, it would mean taking concepts seriously as more than categorization devices, studying their roles in the variety of tasks that science undertakes. It would therefore offer theories of how (specific) concepts support (specific) tasks. And, among other possibilities, it could *test* those theories following cognitive psychology’s lead: creating conditions in which to observe scientists performing scientific tasks (or more experimentally tractable versions of them), perturbing their performance

⁶ It is worth noting that, despite studying laypeople, Lombrozo and colleagues chose a task and pair of concepts that are also paradigmatic *scientific* tasks/concepts. Their work is not just a model for X-phi-sci, but a direct contribution to the field.

by changing the concepts that scientists are using (e.g., prompting one group to think of a system in functional terms and one in mechanistic terms), and measuring changes in task performance.⁷

This is closely in line with Fisher’s suggestion for a “pragmatic experimental philosophy” (2015), which asks “when and how subjects regularly benefit from employing a concept” (p. 414) rather than under what conditions they apply the concept to real or hypothetical cases. But Fisher and other *functionalists* in conceptual engineering still conceive of concepts, insofar as they help produce some benefit for their users, as doing so in virtue of the classifications or categorizations they support — not other cognitive functions like directing attention and so on (Fisher, 2015, pp. 423-424). And the point of the non-categorical approach is that we have, from cognitive psychology, both *reasons to think* that concepts do a lot more than name and delineate categories, and *methods for studying* the way they do so — and especially how, in doing so, they support scientific research.

This suggestion also allows us to revisit a concern *within* X-phi-sci about methodological diversity. As Appiah put it: “If anything shadows the prospects of experimental philosophy, it is, I think, that our notion of experiment, our notion of the empirical, may be too conservative, too narrow” (Appiah, 2008, p. 19). Appiah’s hope was for a radical expansion of X-phi’s methods to include natural experiments and historical work, but less radical expansions are also discussed. Hangel & ChoGlueck (2023), e.g., argue that qualitative methods should have a more central role in X-phi-sci. And my argument so far suggests that they are right. While those authors restrict the role of qualitative methods to interrogating the “social norms of science” (p. 30), we can make a more general case. Qualitative methods are a central part of both the examples I gave of cognitive psychological research, above. They are used to elaborate the nature of medical and lay reasoning, and to probe the roles concepts play in them (Goldszmidt et al., 2013, Lombrozo & Carey, 2006). This is no surprise. Cognition is messy. Control is difficult. Variables multiply. In those conditions, qualitative and mixed methods allow for comparisons and discoveries that would otherwise be impossible.

Qualitative methods are also useful when you don’t yet know enough about your target system. By letting a phenomenon speak (to some degree) on its own terms, you find patterns you wouldn’t have expected and relationships you didn’t know to look for. This is particularly important for a new field, or a field in transition. Henne (2023) distinguishes between *framing inquiry* — the kind of research we do to better conceptualize a system, discover the right questions to pose, and find methods appropriate to that system and those questions — and *framed inquiry* — in which we apply those methods and concepts to devise and test hypotheses, answering the questions we’ve posed. Framing inquiry typically involves gathering qualitative data to help decide what questions are relevant and how they might be answered.

To come to the point, X-phi-sci, as “dangerous” as it might aim to be (Weinberg, 2016), has very often conducted *framed* inquiry on questions inherited from traditional philosophy (Fisher,

⁷ Many other experimental paradigms are possible. E.g., one ongoing project uses “consensus glossary” methods to bring researchers from different sub-fields into a type of conversation that can surface disagreements over central concepts, in order to reveal the role those concepts play in different research programs.

2015). This is by no means universal in X-phi. E.g., many experimental philosophers studying moral reasoning take up new questions about the psychological processes underlying moral judgment, not traditional questions about which moral judgments are correct or how moral terms should be defined. But X-phi-*sci* has generally not made this shift, as I argued in Section 2. Even if it suggests new methods to pursue its questions, those questions still tend to be about the definition of scientific concepts or the nature of the kinds or categories they name. What I'm suggesting now is that X-phi-*sci*'s deep connections to cognitive psychology suggest a broader set of *qualitative*, not just quantitative, methods. And those methods themselves open up the exciting (and maybe dangerous!) prospect for X-phi-*sci* to conduct its own framing inquiry, discovering the questions that it and its new methods can and ought to answer — not just applying new methods to the categorical questions that have been passed down to it.

As one more example of the non-categorical approach, I want to end this section with a look at the work of Daniela Bailer-Jones (2002, 2009), who has been relatively absent from X-phi-*sci*'s canon. Bailer-Jones studied how scientists conceive of *models* in their research, primarily by conducting and analyzing long-form interviews with scientists.⁸ This was undeniably X-phi-*sci*, with qualitative rather than quantitative methods.⁹ And despite focusing on how scientists understood and used a particular scientific concept, 'model,' Bailer-Jones was not asking categorical questions. She didn't try to define 'model,' tell us which things counted as models and which things didn't, or tell us what features made something count as a model. She used her interviews to probe the way that the concept 'model' *organized scientific thinking*.

She did ask scientists how they thought about the concept 'model,' and even asked them to define it. But here are some of the definitions she reports: "a model is the way you set up a problem for investigation" (Bailer-Jones, 2002, p. 282); "the real world does what it does and the model is some simplified representation of part of it in one form or another" (Bailer-Jones, 2002, p. 283). As definitions, as the meaning of the concept 'model,' or as a way of capturing its extension, these answers are implausible. But Bailer-Jones doesn't press the scientists on the definitions (despite pressing them at many other points in the interviews) or try to improve on their definitions. She isn't concerned with the definitions *qua* definitions. Instead, she uses these answers, and especially the rest of the interviews, to analyze how scientists think of models, how they relate models to the rest of their work, and how their understanding of models shapes their reasoning and helps organize their research projects. Her interviews and analysis are designed to answer those questions without detouring through puzzles about the meaning or extension of the concept 'model.'

Bailer-Jones reports a number of interesting findings about models in science: how they serve pragmatic (not just epistemic) roles (2002, p. 286); how they support and are supported by idealizations (2002, p. 285); and how can they provide *insight* into a system, not just descriptions of it (2002, pp. 286–291). I leave it as an exercise to the reader to review that work and imagine

⁸ In hindsight, that work may have benefited from more rigorous methods for analyzing qualitative data (cf. Braun & Clarke, 2006; Richards & Hemphill, 2018). But that won't affect my point here.

⁹ Methodologically, Bailer-Jones' work shares more with the work of Arne Naess (1938) on truth, seventy years before experimental philosophy emerged in its current form, than it does with, e.g., Favela & Machery (2023).

the lessons that would be lost if Bailer-Jones had instead worried about what ‘model’ means, or which things count as models and why. In short, Bailer-Jones’ work represents an important, pioneering, but too-little appreciated attempt to forge a non-categorical path in X-phi-sci, and one that reveals a role for qualitative methods.

4. CONCLUSION

The traditional philosophical approach to scientific concepts, which focuses on their *meaning* and *extension*, is incomplete. While that approach can be informative, it overlooks a significant part of what concepts contribute to science. Experimental philosophy is well-positioned to highlight and critique this oversight, and to rectify it by providing alternative ways to understand and investigate scientific concepts. While this calls for a major change to mainstream X-phi-sci, in light of that field’s self-conception this looks less like a proposal for reform and more like a renewal of vows.

The ways forward are not easy. Many of the methods required are more complex than elicitation studies. And a non-categorical approach will almost certainly be more expensive. A main reason elicitation studies have proliferated in X-phi seems to be that they can be done on the frayed shoestring that most philosophers must work with. But a more sophisticated approach should be our ideal. When we know we are working, out of necessity, with limited or inappropriate methods, that should be noted, so that we do not draw extravagant conclusions from limited findings. And if we can make a more sophisticated approach work, the difficulties will be worth it. They will allow us to broaden the canon of X-phi and X-phi-sci, learning more from philosophers like Bailer-Jones who are not always appreciated as pioneers. They will increase the methodological diversity of X-phi-sci — arguably a good in itself. And, most importantly, they will give X-phi-sci the methods and conceptual framework it needs to really understand scientific concepts and its role in studying them.

REFERENCES

- Alexander, J., Mallon, R., & Weinberg, J. M. (2010). Accentuate the Negative. *Review of Philosophy and Psychology*, 1(2), 297–314. <https://doi.org/10.1007/s13164-009-0015-2>
- Appiah, K. A. (2008). Experimental Philosophy. *Proceedings and Addresses of the American Philosophical Association*, 82(2), 7–22. <https://about.jstor.org/terms>
- Baker, B., Lange, R. D., Richmond, A., Kriegeskorte, N., Cao, R., Pitkow, X., Schwartz, O., & Achille, A. (n.d.). Use and usability: Concepts of representation in philosophy, neuroscience, cognitive science, and computer science. *Neurons, Behavior, Data, and Theory (Forthcoming)*.
- Bailer-Jones, D. M. (2002). Scientists’ thoughts on scientific models. *Perspectives on Science*, 10(3), 275–301. <https://doi.org/10.1162/106361402321899069>
- Bailer-Jones, D. M. (2009). *Scientific Models in Philosophy of Science*. University of Pittsburgh Press.
- Ballesteros, V., & Batalla, A. L. (2024). Valuable Harmful Dysfunctions. *Crítica*, 56(167), 45–69. <https://doi.org/10.22201/iifs.18704905e.2024.1516>

- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
- Bursten, J. R., Hartmann, M. J., & Millstone, J. E. (2016). Conceptual Analysis for Nanoscience. *Journal of Physical Chemistry Letters*, 7(10), 1917–1918. <https://doi.org/10.1021/acs.jpcclett.6b00694>
- Cao, R. (2022). Putting representations to use. *Synthese*, 200(151).
- Cleland, C. E. (2012). Life without definition. *Synthese*, 185(1), 125–144. <https://doi.org/10.1007/s1>
- Cleland, C. E., & Chyba, C. F. (2007). Does “life” have a definition? In W. T. Sullivan & J. A. Baross (Eds.), *Planets and Life: The Emerging Science of Astrobiology* (pp. 119–131). Cambridge University Press.
- Cohnitz, D., & Häggqvist, S. (2018). Thought Experiments in Current Metaphilosophical Debates. In M. T. Stuart, Y. Fehige, & J. R. Brown (Eds.), *The Routledge Companion to Thought Experiments*. Routledge.
- Cummins, R. (1991). *Meaning and Mental Representation*. MIT Press.
- Dunbar, K. N. (2002). Understanding the role of cognition in science: the Science as Category framework. In P. Carruthers, S. Stich, & M. Siegal (Eds.), *The Cognitive Basis of Science* (pp. 154–171). Cambridge University Press.
- Egan, F. (2014). How to think about mental content. *Philosophical Studies*, 170(1), 115–135.
- Egan, F. (2021). A Deflationary Account of Mental Representation. In J. Smortchkova, K. Dolega, & T. Schlicht (Eds.), *What are Mental Representations?* Oxford University Press.
- Fardin, M. A. (2014). On the rheology of cats. *Rheology Bulletin*, 83(2), 16–30. <https://doi.org/10.1002/actp.1984.010350801>
- Favela, L. H. (2021). The dynamical renaissance in neuroscience. *Synthese*, 199(1–2), 2103–2127. <https://doi.org/10.1007/s11229-020-02874-y>
- Favela, L. H., & Machery, E. (2023). Investigating the concept of representation in the neural and psychological sciences. *Frontiers in Psychology*, 14(1165622), 1–13. <https://doi.org/10.3389/fpsyg.2023.1165622>
- Garson, J. (2019). *What Biological Functions are and Why They Matter*. Cambridge University Press.
- Giere, R. N. (1987). The Cognitive Study of Science. In Nancy. J. Nersessian (Ed.), *The Process of Science: Contemporary Philosophical Approaches to Understanding Scientific Practice* (pp. 139–160). Martinus Nijhoff Publishers.
- Gigerenzer, G., Swijtink, Z., Porter, T., Daston, L., Beatty, J., & Krüger, L. (1989). *The Empire of Chance*. Cambridge University Press.
- Goldszmidt, M., Minda, J. P., & Bordage, G. (2013). Developing a unified list of physicians’ reasoning tasks during clinical encounters. *Academic Medicine*, 88(3), 390–397. <https://doi.org/10.1097/ACM.0b013e31827fc58d>
- Goldszmidt, M., Minda, J. P., Devantier, S. L., Skye, A. L., & Woods, N. N. (2012). Expanding the basic science debate: The role of physics knowledge in interpreting clinical findings. *Advances in Health Sciences Education*, 17(4), 547–555. <https://doi.org/10.1007/s10459-011-9331-2>
- Griffiths, P. E., & Stotz, K. (2008). Experimental philosophy of science. *Philosophy Compass*, 3(3), 507–521.

- Hangel, N., & ChoGlueck, C. (2023). On the pursuitworthiness 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>
- Henne, C. (2023). Framed and framing inquiry: a pragmatist proposal. *Synthese*, 201(2), 1–25. <https://doi.org/10.1007/s11229-023-04059-9>
- Knobe, J. (2016). Experimental Philosophy is Cognitive Science. In J. Sytsma & W. Buckwalter (Eds.), *A Companion to Experimental Philosophy* (pp. 37–52). Wiley Blackwell.
- Kucharski, A. (2016). Post-truth: Study epidemiology of fake news. *Nature*, 540(7634), 525. <https://doi.org/10.1038/540525a>
- Laudan, R., Laudan, L., & Donovan, A. (1988). Testing Theories of Scientific Change. In A. Donovan, L. Laudan, & R. Laudan (Eds.), *Scrutinizing Science: Empirical Studies of Scientific Change* (pp. 3–46). Kluwer.
- Lazebnik, Y. (2002). Can a biologist fix a radio? Or, what I learned while studying apoptosis. *Cancer Cell*, 2(3), 179–182. <https://doi.org/10.1007/s10541-005-0088-1>
- Lombrozo, T. (2009). Explanation and categorization: How “why?” informs “what?” *Cognition*, 110(2), 248–253. <https://doi.org/10.1016/j.cognition.2008.10.007>
- Lombrozo, T. (2010). Causal–explanatory pluralism: How intentions, functions, and mechanisms influence causal ascriptions. *Cognitive Psychology*, 61(4), 303–332. <https://doi.org/10.1016/j.cogpsych.2010.05.002>
- Lombrozo, T., & Carey, S. (2006). Functional explanation and the function of explanation. *Cognition*, 99(2), 167–204. <https://doi.org/10.1016/j.cognition.2004.12.009>
- Lombrozo, T., & Gwynne, N. Z. (2014). Explanation and inference: Mechanistic and functional explanations guide property generalization. *Frontiers in Human Neuroscience*, 8(September), 1–12. <https://doi.org/10.3389/fnhum.2014.00700>
- Lombrozo, T., Kelemen, D., & Zaitchik, D. (2007). *Inferring Design: Evidence of a Preference for Teleological Explanations in Patients With Alzheimer’s Disease*. 18(11), 999–1006.
- Ludwig, K. (2018). Thought experiments and experimental philosophy. In M. T. Stuart, Y. Fehige, & J. R. Brown (Eds.), *The Routledge Companion to Thought Experiments* (pp. 385–405). Routledge.
- Machery, E. (2016). Experimental Philosophy of Science. In J. Sytsma & W. Buckwalter (Eds.), *A Companion to Experimental Philosophy* (pp. 475–490). Wiley-Blackwell.
- Mollo, D. C. (2020). Content Pragmatism Defended. *Topoi*, 39(1), 103–113. <https://doi.org/10.1007/s11245-017-9504-6>
- Mollo, D. C. (2021). Deflationary realism: Representation and idealisation in cognitive science. *Mind and Language*, July 2020, 1–19. <https://doi.org/10.1111/mila.12364>
- Naess, A. (1938). “Truth” as Conceived by Those Who Are Not Professional Philosophers. I Kommisjon Hos Jacob Dybward.
- Neander, K. (2017). *A Mark of the Mental*. MIT Press.
- Nersessian, Nancy. J. (2002). The cognitive basis of model-based reasoning in science. In P. Carruthers, S. Stich, & M. Siegal (Eds.), *The Cognitive Basis of Science* (pp. 133–153). Cambridge University Press.
- Panchal, R., & Jack, A. (2022). The contagiousness of memes: containing the spread of COVID-19 conspiracy theories in a forensic psychiatric hospital. *BJPsych Bulletin*, 46(1), 36–42. <https://doi.org/10.1192/bjb.2020.120>
- Piccinini, G. (2015). *Physical Computation: A Mechanistic Account*. Oxford University Press.
- Ramsey, W. (2007). *Representation Reconsidered*. Cambridge University Press.

- Reuter, K. (2024). Salient semantics. *Synthese*, 204(2). <https://doi.org/10.1007/s11229-024-04669-x>
- Richards, K. A. R., & Hemphill, M. A. (2018). A practical guide to collaborative qualitative data analysis. *Journal of Teaching in Physical Education*, 37(2), 225–231. <https://doi.org/10.1123/jtpe.2017-0084>
- Richmond, A. (2023a). Commentary: Investigating the concept of representation in the neural and psychological sciences. *Frontiers in Psychology*, 14. <https://doi.org/10.3389/fpsyg.2023.1259808>
- Richmond, A. (2025a). What is a theory of neural representation for? *Synthese*, 205(14).
- Richmond, A. (2025b). How computation explains. *Mind & Language*, 40(1), 2–20. <https://doi.org/10.1111/mila.12521>
- Shagrir, O. (2022). *The Nature of Physical Computation*. Oxford University Press.
- Shea, N. (2018). *Representation in Cognitive Science*. Oxford University Press.
- Stotz, K. (2009). Experimental philosophy of biology: notes from the field. *Studies in History and Philosophy of Science Part A*, 40(2), 233–237. <https://doi.org/10.1016/j.shpsa.2009.03.008>
- Stotz, K., & Griffiths, P. (2004). Genes: Philosophical Analyses Put to the Test. *History and Philosophy of the Life Sciences*, 26(1), 5–28.
- Stotz, K., & Griffiths, P. (2004). Genes: Philosophical Analyses Put to the Test. *History and Philosophy of the Life Sciences*, 26(1), 5–28.
- Suárez, J. (2023). What is the nature of stem cells? A unified dispositional framework. *Biology and Philosophy*, 38(5). <https://doi.org/10.1007/s10539-023-09930-0>
- Sun, D., Lv, J., & Waller, S. T. (2011). In-depth analysis of traffic congestion using computational fluid dynamics (CFD) modeling method. *Journal of Modern Transportation*, 19(1), 58–67. <https://doi.org/10.1007/bf03325741>
- Sytsma, J., & Buckwalter, W. (2016). *A Companion to Experimental Philosophy* (J. Sytsma & W. Buckwalter, Eds.). Blackwell. <https://doi.org/10.5840/ancientphil200929120>
- The Gene Ontology Consortium. (2000). Gene Ontology: tool for the unification of biology. *Nature America*.
- Waters, C. K. (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(1), 29–58.
- Waters, C. K. (2019). An Epistemology of Scientific Practice. *Philosophy of Science*, 86(4), 585–611. <https://doi.org/10.1086/704973>
- Weinberg, J. M. (2016). Going Positive by Going Negative: On Keeping X-Phi Relevant and Dangerous. In J. Sytsma & W. Buckwalter (Eds.), *A Companion to Experimental Philosophy* (pp. 71–86). Wiley-Blackwell.